

Modelling both the continual accumulation and erosion of discolouration material in drinking water distribution systems

William Robert Furnass

B.Sc. (Computer Science, University of Leeds) 2003

M.Sc. (Environmental Management of Urban Land & Water, University
of Sheffield) 2010

A thesis presented for the degree of
Doctor of Philosophy



Department of Civil and Structural Engineering
The University of Sheffield
United Kingdom
October 2015

Modelling both the continual erosion and accumulation of discolouration material in drinking water distribution systems

Copyright © William Robert Furnass 2015.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License¹.



¹<http://creativecommons.org/licenses/by-nc-nd/4.0/>

Abstract

Discolouration material has previously been shown to accumulate continually and ubiquitously as cohesive layers in all drinking water distribution network pipes. Discolouration risk cannot therefore be eliminated by one-off strategies such as pipe renewal or invasive cleaning. However maintenance programmes of regular controlled flow increases have potential for long-term risk management.

To facilitate the design and optimisation of these programmes, a new discolouration model is proposed. This combines a novel sub-model of how the shear strength profile of material layers changes due to *accumulation* with a material *erosion* sub-model that has similar behaviour to the validated shear-stress-dependant PODDS erosion model. The accumulation sub-model reflects the observation that material with shear strengths exceeding the imposed shear stress accumulates at a rate invariant to this hydraulic force. These differing behaviours are facilitated by modelling how the amount of wall-bound material varies over time for distinct shear strength bands. This results in a model state that is more complex and powerful than PODDS but the new model has fewer parameters.

The validity of the model was tested by calibrating model instances of a pipe rig experiment, four distribution main flushes and three long-term trunk main monitoring programmes using a verified software implementation of the model. The empirical model can be automatically calibrated using swarm optimisation (as has been demonstrated) but in the one case where its predictive power was tested it was found to be limited. However, the model's present value is as a *framework* for furthering the understanding of erosion and regeneration and, as more accurate and complete datasets become available, refinements could result in a tool for medium to long-term network management, including the design of flow conditioning programmes. The model is now implemented in commercial network modelling software.

Acknowledgements

All the research undertaken during this PhD was part of the multidisciplinary project 'Pipe Dreams', which was funded by a Challenging Engineering grant from the U.K. Engineering and Physical Sciences Research Council (EP/G029946/1).

The author is grateful to the following for their contributions to this project:

- Prof. Joby Boxall, Dr. Steve Mounce and Dr. Richard Collins for their supervision, support and encouragement.
- Others directly or indirectly involved in the Pipe Dreams project for providing critical perspective (Dr. Katherine Fish, Dr. Rebecca Sharpe, Dr. Isabel Douterelo, Anthony Leeder) data that was used to support the modelling approach (Dr. Rebecca Sharpe) and the opportunity to collaborate on a three-month experimental study (Dr. Isabel Douterello).
- Dr. Paul Bentley for designing and building a robust, intuitive monitoring and control solution for the pipe rig system used in the aforementioned study.
- The PODDS project's Steering Group for giving the work context, offering suggestions and feedback on model concepts, formulations and the associated software and for offering field data used for model verification and validation. Special mention should go to Dr. Stewart Husband of the Pennine Water Group, Martin Jackson of Wessex Water, Hervé Vicente of Northumbrian Water Ltd, Dr. Dominic Cook of Yorkshire Water and to Anglian Water and Unities Utilities.
- Others within the University of Sheffield's Pennine Water Group have offered invaluable external perspective (Dr. Jonathan Edwards, Dr. John Machell, Dr. Fred Sonnenwald and Iftekhar Sunny), information and opinion relating to water chemistry and industrial practice (Dr. John Machell) and advice on technical matters (Dr. Fred Sonnenwald).
- Dr. John Machell and Dr. Steve Mounce provided the unpublished district metered area (DMA) temperature data used in fig. 2.15.
- Much of the software used to implement, explore and evaluate the discolouration model presented in this thesis (Python; IPython; NumPy; pandas; Cython; matplotlib; SciPy; git) along with software used in the writing of the thesis itself (L^AT_EX; Zotero; L^AT_EX; Inkscape; ImageMagick; Okular) are the products of open-source software projects, with the majority of the contributors to those projects being hard-working volunteers.
- The author would particularly like to thank Suzannah Evans and his parents for their unceasing, invaluable encouragement and support.

Related publications

Material from several prior publications written/co-written by the author have been incorporated within this thesis.

The conceptual model and formulations presented in §4 and some of the supporting evidence presented in §2 were previously presented as [Furnass et al. \(2012\)](#) then published as [Furnass et al. \(2014b\)](#). Note that the quantities α and dN/dt were erroneously associated with each other's units in those papers.

The experiment to assess the variability of the model parameters with shear strength, presented in §6, featured in [Furnass et al. \(2014a\)](#). The motivation, method and interpretation of results were expressed differently in that conference paper to avoid having to explain in detail the model developed in this thesis.

The TM-YR trunk main, modelled in §6.3.3, was previously modelled and featured in [Furnass et al. \(2014b\)](#). The modelling and analysis presented here is more extensive and corrects a mistake made in [Furnass et al. \(2014b\)](#) when estimating the model boundary conditions.

“I have come to believe that the whole world is an enigma, a harmless enigma that is made terrible by our own mad attempt to interpret it as though it had an underlying truth.”

Umberto Eco

Contents

1	Introduction	1
1.1	Discolouration in Water Distribution Systems	1
1.2	Reader's guide	2
2	Literature review	4
2.1	Discolouration processes in drinking water distribution system (DWDS)	4
2.1.1	The need for research into DWDS water quality	4
2.1.2	What is discoloured water?	4
2.1.3	Why is discolouration a concern?	5
2.1.4	Characterisation of discolouration material	8
2.1.5	Light scattering and measuring turbidity	8
2.1.6	Guidance and regulation	11
2.1.7	Discolouration material sources	11
2.1.7.1	Iron corrosion	11
2.1.7.2	Bulk water sources	12
2.1.7.3	Other sources	12
2.1.8	The relationship between turbidity and total suspended solids (TSS)	12
2.1.9	Transport of material from bulk water to wall	16
2.1.10	Mechanisms for binding to the pipe wall	17
2.1.11	Macro- and micro-scale heterogeneity of discolouration material	18
2.1.12	Shear-stress-driven discolouration material erosion	18
2.1.12.1	A velocity or shear-stress driven process?	18
2.1.12.2	The erosion of conditioned material by the 'excess shear stress'	21
2.1.12.3	Unsteady hydraulics and transient pressure waves	23
2.1.12.4	Effect of temperature and wall-bound discolouration material quantity on shear stress	24

2.1.13	The cycle of material accumulation and erosion	24
2.1.13.1	Introduction	24
2.1.13.2	Quantifying discolouration material accumulation; using the integral of material flux to quantify amounts of discolouration material	24
2.1.13.3	Accumulation: an ubiquitous, continual and repeatable process	27
2.1.13.4	Variation over time in the discolouration material strength versus quantity relationship	28
2.1.14	The possibility of an upper bound to the accumulation over a shear strength range	30
2.1.15	The forces required to 'clean' a pipe	34
2.1.16	Factors that could influence accumulation rates	36
2.1.16.1	Introduction	36
2.1.16.2	Temperature	37
2.1.16.3	Characteristics of raw water and water treatment works (WTW) finals	38
2.1.16.4	Asset characteristics	39
2.1.16.5	Pipe hydraulics	39
2.1.16.6	Network topologies	40
2.1.16.7	Water age	41
2.1.16.8	Combined effects	41
2.1.16.9	Season	41
2.1.17	Discolouration risk management	41
2.1.17.1	Quantifying the current discolouration risk	41
2.1.17.2	Interventions for risk reduction	42
2.1.17.3	The use of flow conditioning in trunk mains	45
2.1.17.4	Whole-life costing of flow conditioning	45
2.1.18	Discolouration processes: a summary	46
2.2	Modelling environmental systems	47
2.2.1	Introduction	47
2.2.2	Types of environmental models	47
2.2.3	Workflows for developing environmental models	49
2.2.4	Methods for fitting models to data: model fitting as an optimisation problem	50
2.2.4.1	Introduction	50

2.2.4.2	Evolutionary optimisation and Genetic Algorithms	51
2.2.4.3	Swarm optimisation and the particle swarm optimisation algorithm	52
2.2.4.4	Handling non-trivial constraints when using metaheuristics	55
2.2.4.5	Examples of the use of stochastic optimisation techniques in water engineering	56
2.2.5	Quantifying the dissimilarity between observations and predictions: correlation metrics	57
2.2.6	Methods for signal decomposition and filtering	57
2.2.7	Modelling water quality in DWDSs	58
2.2.7.1	Importance of accurately characterising DWDS hydraulics	58
2.2.7.2	Fundamentals of modelling water quality in DWDSs	59
2.2.7.3	Advection and mixing	59
2.2.7.4	Sources of error	60
2.2.7.5	Modelling the growth/decay of a single-specie	61
2.2.7.6	Modelling multiple reactants	61
2.2.8	Discolouration modelling: the need for a data-driven approach	61
2.2.8.1	Introduction	61
2.2.8.2	The Prediction of Discolouration in Distribution Systems (PODDS) model	62
2.2.8.3	PODDS concepts for material mobilisation	62
2.2.8.4	Calibrating and validating the PODDS erosion model	63
2.2.8.5	Limitations of the PODDS model	65
2.2.8.6	Other discolouration models	67
2.2.8.7	Data-driven models of material accumulation	69
2.3	Conclusions	69
3	Aims and objectives	70
4	Model development	71
4.1	Introduction	71
4.2	Requirements for a new discolouration model	71
4.2.1	Use cases	71
4.2.2	The ability of PODDS to satisfy the use cases in §4.2.1	74
4.3	Conceptual model	78

4.3.1	Introduction	78
4.3.2	Treating shear strength and material quantities as distributions and functions of distributions rather than simple scalars	78
4.3.3	Modelling material accumulation	79
4.3.4	Modelling erosion	82
4.3.4.1	The need for a new erosion model	82
4.3.4.2	Shear stress as an eroding force	83
4.3.4.3	Sequential or simultaneous or layer erosion	85
4.3.4.4	More complicated erosion models	88
4.3.4.5	Chosen erosion model	88
4.3.5	A complex system state	89
4.3.6	Modelling discolouration material in the bulk water: advection, mixing and reactions	90
4.3.6.1	Mobilised material does not settle	90
4.3.6.2	Turbidity is a pseudo-concentration; pseudo-masses of material	91
4.3.6.3	Dispersion and diffusion are negligible; lateral mixing is instantaneous	91
4.3.6.4	The upstream turbidity is known or can be estimated	92
4.3.6.5	The modelling scope is the extents of a single pipe length	92
4.3.7	Summary	93
4.4	Model formulations	93
4.4.1	Units for quantifying discolouration material: TPMU	93
4.4.2	Formulations for state change and material release from the pipe wall	94
4.4.3	Sub-models and formulations for material transport	95
4.4.4	Boundary conditions	97
4.4.5	Variable dependencies	98
4.4.6	A note on parsimony	98
4.5	Implementation as software	98
4.5.1	Requirements	98
4.5.2	Implementation options	100
4.5.2.1	Encoding state change, material release and material transport	100
4.5.2.2	Model scalability	101
4.5.2.3	Performance scalability	101
4.5.2.4	Data input, output and user interface	102

4.5.2.5	Dependencies	103
4.5.2.6	Chosen approach	103
4.5.3	Overview of implementation	103
4.6	Summary of model assumptions	105
4.7	Conclusions	107
5	Model verification, sensitivity analysis and fitting to data	109
5.1	Introduction	109
5.2	Demonstrations of model functionality	109
5.2.1	First example: model response to a simple, stepped shear stress profile	110
5.2.2	Second example: signal transformation due to advection and mixing; the peak turbidity occurring one turnover after a shear stress increase	112
5.2.3	Third example: partial erosion and compounded turbidity responses	112
5.2.4	Fourth example: a comparison of the Variable Condition Discolouration Model (VCDM) and PODDS models	114
5.2.5	Fifth example: Erosion being driven by the excess shear	114
5.3	Sensitivity of the model to the mesh resolution	119
5.3.1	Introduction	119
5.3.2	Scalar measures of model outputs	120
5.3.3	Resolving the depletion of a shear strength band within a timestep for a simple modelling case	121
5.3.4	Sensitivity w.r.t. Δt : more complex boundary conditions	122
5.3.5	Sensitivity w.r.t. Δt : dependence on β_e and τ_e	123
5.3.6	Sensitivity w.r.t. Δt : complex flow profiles	126
5.3.7	Relationship between the mesh discretisation steps	126
5.3.8	Smaller values of $\Delta \tau$ result in lengthier material release profiles	126
5.3.9	Sensitivity w.r.t. Δt and $\Delta \tau$: numerical exploration	128
5.3.9.1	General approach	128
5.3.9.2	Error in the total amount of material released from the pipe wall	129
5.3.9.3	Error in the maximum rate of material release from the pipe wall	130
5.3.9.4	Error in the total amount of eroded material as observed at the downstream end of the pipe	132
5.3.9.5	Error in the maximum turbidity experienced at the downstream end of the pipe	133

5.3.9.6	Repeating the numerical sensitivity analysis using refined mesh resolution ranges and constants	136
5.3.9.7	Discussion and Summary	138
5.4	Sensitivity of the model to its parameters	140
5.4.1	Sensitivity to α and β_r	140
5.4.2	Sensitivity to β_e	141
5.4.3	Implications for fitting the model to data	141
5.5	Sensitivity to the model boundary conditions	141
5.6	Model sensitivity to hydraulics and pipe attributes	145
5.6.1	Sensitivity to flow magnitude	145
5.6.2	Sensitivity to pipe length	147
5.6.3	Sensitivity to the pipe internal diameter	147
5.6.4	Sensitivity to the pipe roughness	147
5.6.5	Summary of sensitivity analysis and implications for model fitting	147
5.6.6	Assessing the need for an equivalent term to the PODDS excess shear power term	151
5.7	Model calibration: fitting the model parameters	151
5.7.1	Introduction	151
5.7.2	An idealised method for fitting the proposed model to data	154
5.7.3	Objective function	154
5.7.4	Fitting by optimisation	155
5.7.4.1	Approaches for finding a global optimum	155
5.7.4.2	Chosen optimisation algorithm: particle swarm optimisation (PSO)	156
5.7.4.3	Chosen PSO implementation	156
5.7.5	Optimising model parameters	156
5.7.5.1	Fitting model parameters simultaneously	158
5.7.5.2	Fitting model parameters in a nested fashion	158
5.7.5.3	Fitting a subset of model parameters after calculating others	158
5.7.5.4	Estimating/bounding the α parameter using known values of the PODDS k parameter	161
5.7.5.5	Validating the three parameter estimation methods	161
5.8	Fitting/estimating the wall sub-model boundary conditions	168
5.8.1	Introduction	168

5.8.2	Fit the initial amount of material at every modelled shear strength band using PSO	168
5.8.3	Reducing the complexity of the wall state boundary condition vector by approximation using cusps	169
5.8.4	Estimate the wall state boundary conditions using the shear stress history	170
5.8.5	Estimate the wall state boundary conditions using the shear stress history and an approximate regeneration rate	171
5.8.6	Estimate the wall state using VCDM conditioning	174
5.8.7	Estimating the wall state boundary condition using generalisations/ observations of the data	174
5.8.8	Discussion and summary	175
5.9	Determining the 'net turbidity response' within a pipe	175
5.10	Conclusions	176
6	Model validation	179
6.1	Introduction	179
6.2	Testing model assumptions	180
6.2.1	Introduction	180
6.2.2	Experimentally verifying the assumption that model parameters are invariant with shear strength ²	181
6.2.2.1	Theory	181
6.2.2.2	Experiment	183
6.2.2.3	Tests adapted for a recirculating system with only one turbidity meter	186
6.2.2.4	Data analysis results	186
6.2.2.5	Data analysis discussion	187
6.2.2.6	Modelling method	188
6.2.2.7	Modelling results	190
6.2.2.8	Conclusions and implications	191
6.3	Model validation using field datasets	193
6.3.1	Requirements and potential issues when calibrating models of real mains	193
6.3.2	Validating the VCDM erosion model by fitting to DMA flushing data .	194
6.3.2.1	Objectives	194
6.3.2.2	Characteristics of flushed reaches	195

²This section contains text that has been quoted and adapted from [Furnass et al. \(2014a\)](#).

6.3.2.3	Fitting method	196
6.3.2.4	PODDS-CI1 results (downstream-most reach)	197
6.3.2.5	PODDS-CI2 results	198
6.3.2.6	PODDS-PE1 results	199
6.3.2.7	PODDS-PE2 results	200
6.3.2.8	Conclusions	201
6.3.3	Trunk main case study TM-YR: fitting to flow trials ³	202
6.3.3.1	Motivation and system configuration	202
6.3.3.2	Data collation and pre-processing	203
6.3.3.3	Modelling TM-YR: Fitting attempt 1	204
6.3.3.4	Modelling TM-YR: Fitting attempt 2	208
6.3.3.5	Modelling TM-YR: Discussion and conclusions	210
6.3.4	Trunk main case study TM-NR: fitting to a longer, continual dataset	211
6.3.4.1	Motivation and system configuration	211
6.3.4.2	Shear stress event detection: identifying the times in lengthy datasets when increases turbidity are likely to occur	214
6.3.4.3	Estimating the net turbidity response from trunk main TM-NR	215
6.3.4.4	Model fitting	219
6.3.4.5	Discussion and summary	223
6.3.5	Trunk main case study TM-WB: fitting to a long dataset containing flow trials then testing predictive performance	224
6.3.5.1	Site description	224
6.3.5.2	Previous PODDS modelling	227
6.3.5.3	Objectives for modelling TM-WB using the VCDM	227
6.3.5.4	Data pre-processing	228
6.3.5.5	Model fitting method	229
6.3.5.6	Fitting results	231
6.3.5.7	Accuracy of predictions using the calibrated model	232
6.3.5.8	Discussion and conclusions	235
6.4	Conclusions	237

³This section contains text that has been quoted and adapted from [Furnass et al. \(2014b\)](#).

7 Discussion and recommendations for future work	240
7.1 Holistic appraisal of model performance	240
7.2 Value of the modelling framework including practical applications	244
7.2.1 Designing optimal flow conditioning programmes	244
7.2.2 Scenario testing and smart alarms	249
7.2.3 Revisiting the modelling use cases	249
7.3 Comparing the VCDM to other discolouration models / risk assessment methods	251
7.4 Ideal field test for validating and demonstrating the advantages of the new model	252
7.5 Revisiting the model assumptions and future research opportunities	255
7.6 Future work	258
8 Conclusions	261
References	264

List of Figures

2.1	Example drinking water distribution network	5
2.2	Solid, light-scattering material continually accumulates within drinking water distribution systems over time but can become suspended and advected through to consumers as a result of an anomalous increase in hydraulic conditions.	6
2.3	ATi A15/76 turbidity meter	10
2.4	Material accumulation in perspex pipe (after Vreeburg, 2007)	17
2.5	Internal surface of ductile iron trunk main	19
2.6	SEM image of biofilm on pipe wall.	19
2.7	\bar{u} vs τ_a and Q vs τ_a	21
2.8	Turbidity responses from increasing τ_a in steps in DMA pipes	22
2.9	Discolouration responses from UK trunk main sections with greatest τ_e , not greatest \bar{U}	23
2.10	Difference between turbidity and material flux	27
2.11	Turbidity data from the repeated flushing of DMA mains.	29
2.12	Material developed with range of shear strengths during laboratory rig experiment.	31
2.13	Material accumulation could increase linearly indefinitely or plateau.	31
2.14	Daily-averaged turbidity and flow records from a UK trunk main over 2006-2008	33
2.15	Temperature variation in UK DMA over one year.	38
2.16	The 'sinc' function has local optima	52
2.17	Demonstration of PSO (five particles; 2D objective function)	54
2.18	Lagrangian transport in pipes	60
2.19	Graphical output from PODDS software for calibrated model	64
2.20	h_L is indeterminate w.r.t. D and k_s if Q but not \bar{U} known	65
2.21	Sensitivity of PODDS to error in the n parameter	68
4.1	Use case: explore model sensitivity to regeneration rate.	73
4.2	Use case: explore model sensitivity to erosion rate.	73
4.3	Use case: calibrate model so output correlates with historic observations.	74

4.4	Use case: use model to design flow conditioning programmes.	75
4.5	Change in PODDS model state and output during and after an erosion event.	75
4.6	Changes in the PODDS model state and output due to the unvalidated material regeneration mechanism.	76
4.7	Discrepancy between conceptual understanding of material accumulation gained through field and laboratory experiments and PODDS model regeneration mechanism (different internal state but same turbidity response).	77
4.8	Discrepancy between conceptual understanding of material accumulation gained through field and laboratory experiments and PODDS model regeneration mechanism (different internal state <i>and</i> different turbidity response).	77
4.9	Interdependent factors influence accumulation rates	81
4.10	Regeneration rate model can be refined over time.	81
4.11	Conceptual model: regeneration only occurs for $\tau > \tau_a$	83
4.12	Conceptual model: discolouration material appears to bind to the wall in cohesive layers.	84
4.13	Conceptual model: erosion driven by excess shear stress, not absolute shear stress.	84
4.14	Options for conceptual model of erosion: shear strengths erode <i>sequentially</i> or shear strengths erode <i>simultaneously</i> at rate that is constant with t and τ	86
4.15	Flow profiles for potentially distinguishing erosion models that are sequential and simultaneous with respect to τ	87
4.16	Option for conceptual model of erosion: shear strengths erode simultaneously at rate that is constant with t for each τ but varies with τ	88
4.17	Option for conceptual model of erosion: shear strengths erode simultaneously at rate that is constant with τ but varies with t	89
4.18	Quantity of wall-bound material not always linear with shear strength.	90
4.19	T_{ds} is a function of material erosion, accumulation and T_{us}	92
4.20	Modelling pipes in series by chaining the inputs and outputs of multiple simulations.	94
4.21	The model permits some shear strengths of material to erode whilst others regenerate.	96
4.22	$\varphi(\tau_i, t = 0)$ is monotonically increasing w.r.t. τ	97
4.23	Variable/data dependency diagram for proposed model.	99
5.1	Model demonstration #1: response to simple, stepped τ_a profile.	111
5.2	Model demonstration #1: heatmap of $\varphi(\tau_i, t_j)$ for fig. 5.1	113
5.3	Relationship between dN/dt vs t and T_{ds} vs t ; peak turbidity occurs $1 \cdot t_t$ after τ_a increase.	113

5.4	Changes in the spatial variation of turbidity during a modelled erosion event due to advection and mixing.	113
5.5	Model demonstration #3: partial erosion and compound turbidity responses caused by a complex τ_a profile.	115
5.6	Model demonstration #3: heatmap of $\varphi(\tau_i, t_j)$ for fig. 5.5.	115
5.7	Time-series simulation results showing PODDS and VCDM simulation results can differ following a period of regeneration.	116
5.8	Heatmap of $\varphi(\tau_i, t_j)$ for VCDM simulation results shown in fig. 5.7.	116
5.9	C vs t and τ_c vs t for PODDS simulation results shown in fig. 5.7.	117
5.10	Model demonstration #5: turbidity response for series of τ_a increments of constant magnitude.	117
5.11	Turbidity spikes in fig. 5.10 all follow same curve up to the turnover time.	119
5.12	Model output less accurate if tracking very few τ bands.	120
5.13	Options for defining critical relative material quantity that want the model to be able to resolve the erosion of.	124
5.14	$(dN(t)/dt) / (max(dN(t)/dt))$ vs t/t_d collapses onto a single curve, regardless of the value of β_e	125
5.15	$(dN(t)/dt) / (max(dN(t)/dt))$ vs t/t_d collapses onto a single curve, regardless of the value of τ_e	125
5.16	τ_e depends on $\Delta\tau$	127
5.17	The time (in s) required for the complete depletion of the slowest eroding shear strength band given different shear strength discretisation step sizes	128
5.18	Relative error from discretely calculating (a) the total amount of material to have been released from the pipe wall and (b) the maximum release rate (plotted against non-dimensional time)	130
5.19	Relative error from discretely calculating (a) the total amount of material to have been released from the pipe wall and (b) the maximum release rate (plotted against non-dimensional shear stress)	131
5.20	Relative amount of material eroded per τ over dt at start of ideal erosion event.	131
5.21	Relative error from discretely calculating (a) the total amount of material to pass out of the pipe and (b) the maximum downstream turbidity (plotted against non-dimensional time)	134
5.22	Relative error from discretely calculating the total amount of material to pass out of the pipe (using a different expression to fig. 5.21a; plotted against non-dimensional time)	134
5.23	Relative error from discretely calculating (a) the total amount of material to pass out of the pipe and (b) the maximum downstream turbidity (plotted against non-dimensional shear stress)	135

5.24	As per fig. 5.18 but for narrower $\Delta t/t_t$ range and different fixed value of $\Delta\tau/\tau_e(\tau_{a_lo}, \tau_{a_hi})$.	136
5.25	As per fig. 5.21 but for narrower $\Delta t/t_t$ range and different fixed value of $\Delta\tau/\tau_e(\tau_{a_lo}, \tau_{a_hi})$.	137
5.26	As per fig. 5.19 but for narrower $\Delta\tau/\tau_e(\tau_{a_lo}, \tau_{a_hi})$ range and different fixed value of $\Delta t/t_t$.	137
5.27	As per fig. 5.23 but for narrower $\Delta\tau/\tau_e(\tau_{a_lo}, \tau_{a_hi})$ range and different fixed value of $\Delta t/t_t$.	137
5.28	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to α .	142
5.29	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to β_r .	142
5.30	Sensitivity of $\int_{t_0}^{t_{max}} Q(t) \cdot T_{ds}(t) dt$ to β_r .	143
5.31	$\tau_a(t)$ profile for distinguishing between α and β_r .	143
5.32	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to higher values of β_e .	143
5.33	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to lower values of β_e .	144
5.34	Sensitivity of $\varphi(\tau, t)$ to β_e .	144
5.35	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to the shape of $\varphi(\tau, t_0)$.	146
5.36	Q error impacts on τ_a and t_t .	146
5.37	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to Q .	148
5.38	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to L .	148
5.39	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to D .	149
5.40	Sensitivity of $dN(t)/dt$ and $T_{ds}(t)$ to k_s .	149
5.41	Sensitivity of PODDS to n parameter	151
5.42	Flow chart of FIT_SIMULT fitting method.	158
5.43	Flow chart of FIT_NESTED fitting method.	159
5.44	Flow chart of CALC_AND_FIT fitting method.	160
5.45	Attempts to fit VCDM to target prediction VCDM-2 (no Q noise) using FIT_SIMULT, FIT_NESTED and CALC_AND_FIT.	163
5.46	Attempts to fit VCDM to target prediction PODDS-2 (no Q noise) using FIT_SIMULT, FIT_NESTED and CALC_AND_FIT.	164
5.47	Attempts to fit VCDM to target prediction VCDM-1 (see table 5.5; with Q noise) using FIT_SIMULT, FIT_NESTED and CALC_AND_FIT.	164
5.48	Fitting artificial erosion events given known non-trivial boundary conditions: $\varphi(\tau_i, t_0)$ shapes of target predictions.	165
5.49	Target VCDM dataset VCDM-5, generated using noisy flow profile and 'two steps' $\varphi(\tau, t_0)$ shape.	166

5.50	Attempt to fit VCDM to target dataset VCDM-5 (generated using noisy flow profile and ‘half ramp’ $\varphi(\tau_i, t_0)$ profile) using the FIT_NESTED method.	168
5.51	Different methods for approximating $\varphi(\tau_i, t_0)$ using cusps.	169
5.52	Calibrate model by estimating $\varphi(\tau_i, t_0)$ using data from one time period then fit β_r, β_e and α over a successive period.	170
5.53	Maximum τ_a over a period can be used to polarise $\varphi(\tau_i, t)$ w.r.t. τ at the end of that period.	171
5.54	Estimating $\varphi(\tau_i, t_0)$ using both the τ_a profile over a period and an estimate of β_r (period $> \beta_r^{-1}$).	172
5.55	Estimating $\varphi(\tau_i, t_0)$ using both the τ_a profile over a period and an estimate of β_r (period $< \beta_r^{-1}$).	173
5.56	Decomposing T_{ds} into T_{us} and T_{net}	177
6.1	Model parameters could be constant, linear or non-linear with τ	181
6.2	Can study how erosion and regeneration processes vary with τ by increasing τ_a in equal increments after a controlled ‘growth phase’.	183
6.3	Schematic of each pipe system.	184
6.4	Daily Q , Re and τ_a profiles during growth phase of rig experiment.	185
6.5	Flow $Q(t)$ and measured turbidity $T_2(t)$ over the flushing phases of systems A (subplot a) and B (b). $T(t)$ is the turbidity per τ_a step with the background turbidity of the previous step subtracted.	187
6.6	Processed turbidity in the pipe rig for the last turnover at each τ_a level	188
6.7	188
6.8	The schematic of each pipe system as modelled. Roman numerals correspond to the locations shown in fig. 6.3.	189
6.9	VCDM fitted to the earlier parts of the flushing phase turbidity profiles.	190
6.10	VCDM fitted to the later parts of the flushing phase turbidity profiles.	191
6.11	VCDM fitted to the flushing phase turbidity profiles after discounting the turbidity responses from Steps 0 and 1.	191
6.12	Results of fitting the VCDM to the PODDS-CI1 dataset	198
6.13	Results of fitting the VCDM to the PODDS-CI2 dataset	199
6.14	Results of fitting the VCDM to the PODDS-PE1 dataset	200
6.15	Results of fitting the VCDM to the PODDS-PE2 dataset.	201
6.16	Schematic of trunk main TM-YR	203
6.17	Time-series observations in trunk main TM-YR	204
6.18	Time-series observations in trunk main TM-YR (Trial 1)	206

6.19	Trunk main TM-YR fit attempt 1 (Trial 1)	206
6.20	Trunk main TM-YR fit attempt 1 (Trial 2)	206
6.21	Trunk main TM-YR fit attempt 1 (Trial 3)	207
6.22	Trunk main TM-YR fit attempt 2 (Trial 1)	208
6.23	Trunk main TM-YR fit attempt 2 (Trial 2)	209
6.24	Trunk main TM-YR fit attempt 2 (Trial 3)	209
6.25	Trunk main TM-YR fit attempt 2 (material condition at the start of trial 1)	210
6.26	Network schematic showing trunk main TM-NR	212
6.27	Dataset for trunk main TM-NR used for VCDM calibration and validation	213
6.28	Demonstration of detecting likely turbidity events from just a shear stress time-series	216
6.29	Attempting to remove the background turbidity by applying a low-pass Butterworth filter to T_{wtw_obs} then subtracting the result from T_{ds_obs}	218
6.30	Attempting to remove the background turbidity by applying a centred rolling mean (middle subplot) and centred rolling 5th percentile to T_{ds_obs} then in each case subtracting the result from T_{ds_obs} .	220
6.31	Best model fit to TM-NR: fitting attempt 1, event 1	222
6.32	Best model fit to TM-NR: fitting attempt 1, event 3	222
6.33	Best model fit to TM-NR: fitting attempt 1, event 4	223
6.34	Best model fit to TM-NR: fitting attempt 1, event 6	223
6.35	Best model fit to TM-NR: fitting attempt 1, event 8	223
6.36	Best model fit to TM-NR: fitting attempt 1, event 11	223
6.37	Network schematic showing trunk main TM-WB	225
6.38	Flow and turbidity data from the first two flow trials in trunk main TM-WB	226
6.39	Internal condition of the TM-WB trunk main	227
6.40	Dataset for trunk main TM-WB used for VCDM calibration and validation	229
6.41	Trunk main TM-WB, VCDM fitting attempt, flow trial 1.	232
6.42	Trunk main TM-WB, VCDM fitting attempt, flow trial 2.	232
6.43	Trunk main TM-WB, VCDM fitting attempt, flow trial 3	233
6.44	Trunk main TM-WB, VCDM fitting attempt, flow trial 4.	233
6.45	Trunk main TM-WB, VCDM fitting attempt, flow trial 5.	233
6.46	Trunk main TM-WB, VCDM fitting attempt, flow trial 6.	234
6.47	Trunk main TM-WB: the relative material quantity versus shear strength relationship at the end of the VCDM fitting attempt.	234

6.48 Predicting turbidity in trunk main TM-WB using the model calibrated during the fitting attempt	236
7.1 Design for a flow conditioning programme where the background flow is diurnally varying	246
7.2 Design for a flow conditioning programme where the background flow is constant	248
7.3 Example operational management chart	249
7.4 Demonstration of the use of VCDM alarms	250

List of Tables

2.1	DWI water acceptability events in 2012	6
2.2	Some commonly used turbidity meters.	10
5.1	Comparing time-integral of turbidity and of material flux for turbidity spikes shown in fig. 5.10.	118
5.2	Summary of the ways in which the mesh resolution impacts on the model outputs.	139
5.3	Summary of model sensitivity to parameters, boundary conditions hydraulics and pipe attributes.	150
5.4	Generic PSO configuration used for model fitting.	157
5.5	VCDM target attribute sets for generating target VCDM predictions (used for testing the efficacy of several VCDM parameter fitting methods) given known simple boundary conditions and negligible regeneration.	162
5.6	PODDS target attribute sets for generating target PODDS predictions (used for testing the efficacy of several VCDM parameter fitting methods) given known simple boundary conditions and negligible regeneration.	163
5.7	VCDM target attribute sets used to generate VCDM predictions datasets, in turn used for testing efficacy of methods for fitting both erosion and regeneration parameters given non-trivial boundary conditions	165
6.1	Dimensions of pipe systems A and B	184
6.2	Q and corresponding τ_a steps imposed during flushing phase of rig experiment.	185
6.3	Several isolated flushing events for which the PODDS model has been successfully calibrated.	196
6.4	PSO Problem space boundaries used when validating the VCDM using the datasets listed in table 6.3.	197
6.5	Results of fitting the VCDM to the PODDS-CI1 dataset	197
6.6	Results of fitting the VCDM to the PODDS-CI2 dataset	198
6.7	Results of fitting the VCDM to the PODDS-PE1 dataset	199
6.8	Results of fitting the VCDM to the PODDS-PE2 dataset	200
6.9	Results of first attempt to fit the VCDM to the TM-YR trunk main dataset	205

6.10	Results of second attempt to fit the VCDM to the TM-YR trunk main dataset	210
6.11	Dataset for trunk main TM-NR used for VCDM calibration and validation . . .	213
6.12	Fitting the VCDM to data from trunk main TM-NR: parameters and correlation metrics resulting from the application of the FIT_SIMULT fitting method. . . .	221
6.13	Key events in the history of trunk main TM-WB regarding discolouration risk .	226
6.14	Stats for part of trunk main TM-WB dataset used for VCDM calibration and validation	230
6.15	Results of attempt to fit the VCDM to the TM-WB trunk main dataset	231
7.1	Details of the best model calibration fits for all DWDS pipe lengths modelled in §6. The first set of parameters for each pipe system is the result of trying to fit to the earlier part of the flushing phase and the second set from trying to fit to the later part.	240

Nomenclature

Acronyms

AC asbestos cement

API application programmer's interface

ARMA auto-regressive moving average

BPT break pressure tank

CI cast iron

CLSM confocal laser scanning microscope

DI ductile iron

DMA district metered area

DOMS Distribution Operational Maintenance Strategies

DPM Discolouration Propensity Model

DRM Discolouration Risk Model

DWI Drinking Water Inspectorate (England and Wales)

DWDS drinking water distribution system

EPR evolutionary polynomial regression

EPS extracellular polymeric substances

FCV flow control valve

GIS geographic information system

GUI graphical user interface

HDPE high-density polyethylene

HPC heterotrophic plate count

ISE integral of squared errors

LUT look-up table

MAC Maximum Admissible Concentration

NOM natural organic matter

NSEI Nash Sutcliffe (Model) Efficiency Index

NTU nephelometric turbidity unit

PaVoC particle volume concentration

PE polyethylene

PODDS Prediction of Discolouration in Distribution Systems

PSD particle size distribution

PSM Particle Sediment Model

PSO particle swarm optimisation

PVC polyvinyl chloride

RPM Resuspension Potential Method

SCADA Supervisory Control and Data Acquisition

SNR signal-to-noise ratio

SR service reservoir

TOC total organic carbon

TPMU turbidity pseudo-mass unit

TSS total suspended solids

UoS University of Sheffield

UTC Coordinated Universal Time

VCDM Variable Condition Discolouration Model

VSS volatile suspended solids

WPS water pumping station

WSZ Water Supply Zone

WTW water treatment works

WHO World Health Organisation

Symbols

A_w Total internal surface area of the pipe [m^2]

A_S Area of pipe wall swept [m^2]

b **PODDS** exponential coefficient of the discolouration potential [-]

b_p Used when quantifying scattering effects. Particle shape factor (1.5 for spherical particles)

c Mass of particles per unit volume of fluid (i.e. **TSS**) [$kg \cdot m^{-3}$]

C **PODDS** discolouration potential per m^2 of wall area [$NTU \cdot m$]

ΔC_e	PODDS change in discolouration potential per m^2 of wall area due to erosion [$NTU \cdot m$]
C_{max}	The maximum theoretical PODDS discolouration potential per m^2 of wall area [$NTU \cdot m$] (can only be approached asymptotically)
D	Pipe internal diameter, \varnothing [m]
D_H	Hydraulic diameter, equal to the diameter for a pipe with a circular cross-section [m]
D_p	Particle diameter [m]
f	Darcy Weisbach friction factor [–]
g	Acceleration due to gravity [$m \cdot s^{-2}$]
h_f	Flow-dependent head loss [m]
I_0	Used when quantifying scattering effects. Intensity of incident light [$J \cdot m^{-2} \cdot s^{-1}$]
I_t	Used when quantifying scattering effects. Intensity of transmitted light [$J \cdot m^{-2} \cdot s^{-1}$]
k	PODDS gradient coefficient, which relates the discolouration potential to the shear strength [$NTU \cdot m \cdot Pa^{-1}$]
k_s	Pipe roughness [m]
K	Used when quantifying scattering effects. Extinction coefficient accounting for the extent of reformation of the light beam after scattering behind the particle
K_0	Used when quantifying scattering effects. Particle-size-independent component of K
L	Pipe length [m]
L_L	Used when quantifying scattering effects. Optical path length [m]
m	Used when quantifying scattering effects. An (empirical) exponent of the wavelength which is dependent on the particle size and refractive index
n	PODDS exponential coefficient of eroding force [–]
n'	A function of D_p used when quantifying light scattering effects. Used to take account of reduced exposure of particles to the light source as concentration increases
N	The absolute amount of material in the VCDM that is bound to each m^2 of pipe wall [$TPMU \cdot m^{-2}$]
N_p	Number of particles per mL of fluid
\mathbb{N}	The set of natural numbers (non-negative integers)
P	PODDS linear coefficient of the eroding force [$NTU \cdot m \cdot Pa^{-1} \cdot s^{-1}$]

P'	PODDS linear coefficient of material accumulation [$NTU \cdot m \cdot Pa^{-1} \cdot s^{-1}$]
R	PODDS rate of supply from a m^2 of wall area [$NTU \cdot m \cdot s^{-1}$]
\mathbb{R}	The set of real numbers
Re	Reynolds number $[-]$, the ratio of inertia to viscous forces
$Q(t)$	Flow in pipe at time t [$m^3 \cdot s^{-1}$]
Q	Bulk flow [$m^3 \cdot s^{-1}$]
$T_{ds}(t)$	Abbreviation of $T_{ds_obs}(t)$ or $T_{ds_pred}(t)$ (depends on context as to which) [NTU]
$T_{ds_obs}(t)$	Turbidity observed at the downstream end of a pipe [NTU]
$T_{ds_pred}(t)$	Turbidity predicted at the downstream end of a pipe [NTU]
$T_{net}(t)$	Component of the turbidity observed at the downstream end of a pipe that is solely due to material erosion within that pipe [NTU]
$T_{us}(t)$	Turbidity observed at the upstream end of a pipe [NTU]
t	Time since a datum [s]
t_0	The amount of material at each tracked shear strength at time t_0 be estimated using the shear stress history over the period t_L to t_0 [s]
t_d	VCDM characteristic time of material depletion [s]
t_L	See t_0 .
t_t	Pipe turnover time (length divided by bulk velocity at a given time) [s]
\bar{u}	Bulk velocity [$m \cdot s^{-1}$]
V_p	Particle volume [m^3]
α	VCDM scaling parameter for translating from relative to absolute quantities of discolouration material [$TMPU \cdot Pa^{-1} \cdot m^{-2}$]
β_e	VCDM relative material erosion rate factor [$Pa^{-1} \cdot s^{-1}$]
β_r	VCDM relative material regeneration rate [s^{-1}]
γ	To resolve a change in φ of 0.1 for a given shear strength, Δt must be $\leq \gamma \cdot t_d$, where $\gamma \in [0, 1]$ $[-]$
$\varphi(\tau, t)$	The relative amount of material bound to each m^2 of pipe wall with strength τ at time t in the VCDM $[-]$
μ	Dynamic viscosity [$Pa \cdot s$]
ν	Kinematic viscosity [$m^2 s^{-1}$]
ρ	Fluid density [$1000 kg \cdot m^{-3}$]

ρ_p	Particle density [$1000 \text{ kg} \cdot \text{m}^{-3}$]
τ_a	Shear stress at the pipe wall [Pa]
τ_c	PODDS (scalar) material shear strength [Pa]
τ_e	Excess shear stress of some discolouration material bound to the pipe wall i.e. $\min(\tau_a - \tau, 0)$ [Pa]. τ is τ_c if using the PODDS model
τ_i	One of a range of shear strengths at which the quantity of wall-bound material is tracked by the VCDM [Pa]
τ_{max}	The maximum of the range of shear strengths tracked by the VCDM [Pa]
\in	Set membership e.g. if x belongs to the set of real numbers between a and b then $x \in [a, b]$
\forall	Universal quantification in predicate logic. Can be read as 'for all' e.g. $x/2 = x \times 0.5 \forall x$
$[a, b]$	The infinite set of real numbers from a to b inclusive
$[a, b)$	The infinite set of real numbers from and including a to but not including b

Chapter 1

Introduction

1.1 Discolouration in Water Distribution Systems

In recent years, water service providers and water industry regulators in the developed world have come to recognise the importance of holistic approaches to maintaining the quality of water provided to consumers. Both the World Health Organisation ([WHO](#)) and regulatory bodies such as the Drinking Water Inspectorate (England and Wales) ([DWI](#)) promote proactive, risk-orientated strategies for managing water quality that consider the entire journey from catchment to consumers' taps ([WHO, 2005](#); [DWI, 2005](#)).

Traditionally water providers have focused on the impact of catchment management and treatment processes on water quality, but the ways in which quality changes as water passes through distribution networks are now receiving much more attention. The hydraulics and water quality of distribution systems are difficult to characterise: distribution networks are topographically complex bio-reactors with very large surface areas but are predominantly buried infrastructure with limited monitoring opportunities.

Networks containing old unlined cast iron pipework may be associated with the supply of orange, black or brown water, of chlorine depletion and of bacterial regrowth ([Kirmeyer, 2000](#)). In the past, old mains associated with water quality problems have been replaced or (re)lined. However, studies have shown that the supply of discoloured water may be associated with newer plastic pipes as well as older corrodible pipes ([Vreeburg et al., 2008](#)). Also, recent initiatives ¹have stressed the importance of striking a better balance between capital and operational expenditure through consideration of whole life costs and so more water providers are appraising and developing management strategies to help prolong the life of assets and ensure satisfactory water quality.

The detection of orange, black or brown discoloured water at customers' taps has been shown to be the effect of a change in hydraulic conditions within the distribution network upstream. Particulate matter continually accumulates on pipe walls, then is mobilised when the typical peak flow per pipe is exceeded. Proactive, operational strategies have already been shown to be a cost-effective means for reducing the likelihood of discolouration events in mains historically associated with a high discolouration risk ([Husband et al., 2011](#)). This is achieved by

¹E.g. Distribution Operational Maintenance Strategies ([DOMS](#)) in England and Wales ([DWI, 2002](#)).

conditioning the particulate matter bonded to the pipe wall using controlled flow increases. In addition, discolouration modelling can be used to predict the turbidity response from exceeding the prevailing flows in distribution network. Such a management strategy can also be deployed for non-metallic mains for which there may be a discolouration risk (Vreeburg and Beverloo, 2011); in this case asset replacement or refurbishment have little impact on discolouration risk as the particulate matter originates from the bulk water, not directly from corrosion processes (Vreeburg et al., 2008). Flow conditioning of mains results in systems that are more resilient to increases in flow (Husband et al., 2011), such as those that result from bursts, valve operations and anomalous customer demands.

A key issue relating to the use of flow conditioning for managing discolouration risk is that the rates at which particulate matter re-accumulates as cohesive layers on pipe walls following layer erosion are not well characterised. Models of material *mobilisation* processes exist and have been verified in the field (Husband and Boxall, 2010a) but current models of material *accumulation* do not reflect the processes knowledge gained through studying accumulation in the field and laboratory .

An improved model of material erosion and accumulation plus the factors that influence the latter would equip water providers with the means to better schedule and design flow conditioning maintenance programmes. This would reduce the risk of customers being supplied with discoloured water and would also benefit the whole life costing of those programmes.

1.2 Reader's guide

Chapter 2: Literature review

The first part of the chapter focusses on discolouration as topic of concern for water companies, what is presently known about discolouration material, how it accumulates on and erodes from pipe walls and fourthly on the various strategies currently used for managing discolouration risk. The second part discusses the development and usage of environmental and drinking water network-specific models, progressing to an appraisal of the state of the art in discolouration modelling as described in the academic literature.

Chapter 3: Aims and objectives

Chapter 4: Model development

This chapter assesses the requirements for discolouration modelling and proposes a conceptual model for a new approach that has the potential to better fulfil those requirements than current discolouration models. This conceptual model is encoded as a set of empirical numerical formulations and algorithms for modelling discolouration processes at the pipe wall and in the bulk water. The associated assumptions are explicitly stated.

Chapter 5: Model verification, sensitivity analysis and fitting to data

The chapter begins with several demonstrations of the functionality of the model developed in the previous chapter. After verifying at a high level that the model satisfies the conceptual specification, the sensitivity of the model to its inputs, the mesh resolution and parameter values are investigated. Finally, a method for fitting the model to data is developed and validated using artificial datasets.

Chapter 6: Model validation

First, the model assumptions that need to be and can be validated are identified. The assumption that the rates of discolouration material accumulation and erosion are invariant with shear strength is then investigated using a realistic-scale laboratory pipe rig. In the second part of the chapter the validity of the model as a whole is tested using four DMA flushing datasets and three long-term trunk main datasets.

Chapter 7: Discussion and recommendations for future work

This chapter begins with an appraisal of the model's ability to represent and predict discolouration processes in real DWDS. Next, several practical applications for the model are explored. The modelling requirements and assumptions are then revisited. The chapter concludes with ideas for future related research.

Chapter 8: Conclusions

Chapter 2

Literature review

2.1 Discolouration processes in DWDS

2.1.1 The need for research into DWDS water quality

The development of regulations and guidelines and the application of stringent practices have ensured that potable water in the developed world is of high quality. In 2012, 99.96% of regulatory water samples in England were compliant with the standards for England and Wales (DWI, 2013a). However, there is a need to ensure that water quality remains high as distribution system infrastructure deteriorates; ageing unlined cast iron pipework, common in the UK, is often associated with water quality problems (Archuleta and Manwaring, 2002) and the proportion of water quality incidents attributable to events and activities downstream of WTW is on the rise (Gray, 2008).

Confidence in the quality of water entering distribution systems is gained through high-resolution monitoring of treatment works finals. However, the processes and reactions that occur downstream within pressurised distribution systems are difficult to study. Observing the physical, chemical and biological reactions occurring within the bulk water and at the pipe wall is made difficult by the subterranean nature, complex topologies and geographical extents of these systems (see Fig. 2.1 and Costello et al. 2007).

2.1.2 What is discoloured water?

In the water industry the *true colour* of water is defined as the measurable effect of dissolved substances on the absorption spectrum of water. Untreated water may have a noticeable colour if it is from an upland source due to the presence of humic and fulvic acids. These substances are removed during treatment (typically by coagulation and flocculation) (Ratnayaka et al., 2009).

Water *discolouration* as experienced by consumers is typically due not to colour but to suspended solids that slowly settle in fully quiescent conditions (Vreeburg and Boxall, 2007) or to air bubbles coming out of solution as water with a high dissolved oxygen content that had previously been at high pressures emerges from customers' taps (DWI, 2010; WHO, 2011, p. 229). So-called

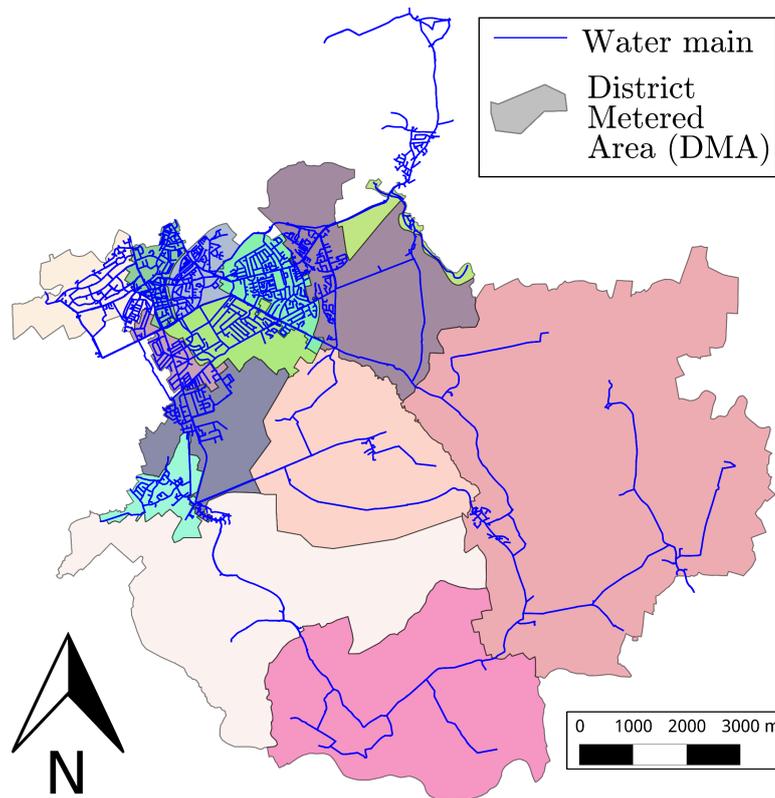


Figure 2.1: Drinking water distribution network infrastructure is mostly buried and is both topologically complex and spatially extensive, as can be seen in this plan of the mains and management areas that comprise part of a DWDS (after Furnass et al. (2013)).

'white water', which is milky/cloudy due to aeration or the presence of fine chalk particles, is not of great concern as there is no health risk associated with consuming white water.

Of greater interest and concern is discolouration due to suspended solids that have an orange, black or brown hue due to their metal content (Seth et al., 2004; DWI, 2010). It is widely accepted that such solids continually accumulate within distribution systems over time and can be mobilised by a hydraulic disturbance such as an anomalous increase in flow then conveyed through to consumers (see Fig. 2.2 and Vreeburg and Boxall, 2007).

Discolouration is typically quantified as *turbidity* in nephelometric turbidity unit (NTU). Turbidity is a measure of how incident light is scattered by particulate matter entrained in the bulk flow (Ratnayaka et al., 2009, p. 201, 218, 280).

2.1.3 Why is discolouration a concern?

In England and Wales utility companies have a duty to report events that have affected consumers to the Drinking Water Inspectorate (England and Wales) (DWI). 9.36% and 22.84% of the Not Significant/Minor and Significant events (respectively) that required investigation by an inspector were due to the supply of discoloured water (table 2.1, compiled using data from DWI, 2013b). Although discolouration incidents may have been the most common type of Significant event, it must be said that a portion of these have causes that lie upstream of the treated water transmission and distribution infrastructure.

The emergence of discoloured water from customers' taps is of concern for several reasons:

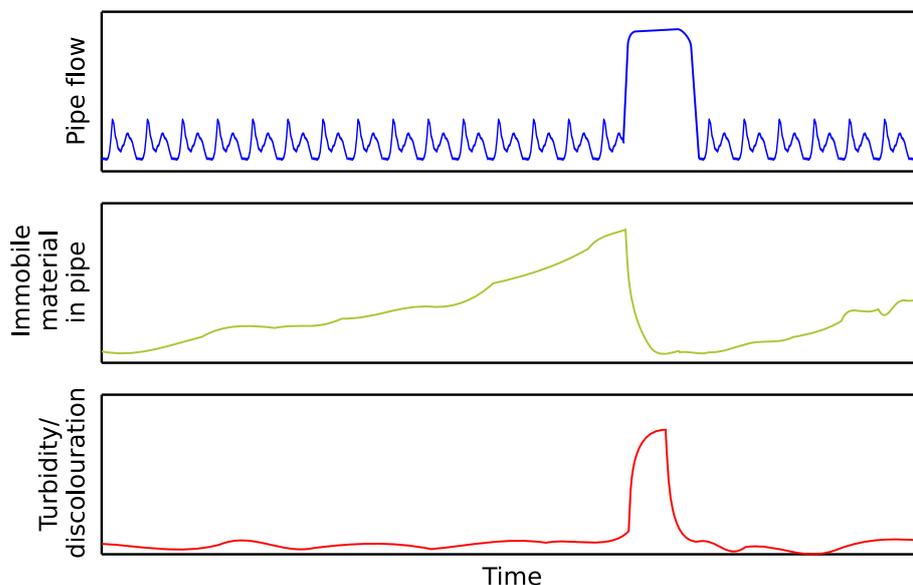


Figure 2.2: An illustrative figure demonstrating the current understanding that solid, light-scattering material continually accumulates within drinking water distribution systems over time but can become suspended and advected through to consumers as a result of an anomalous increase in hydraulic conditions.

Table 2.1: Water acceptability events reported to the Drinking Water Inspectorate (England and Wales) (DWI) (England and Wales) in 2012 where an inspector was required to investigate.

	Not Significant/Minor	Significant	Serious/Major
Air in water	2 (0.85%)	0	0
Chemical	16 (6.81%)	9 (4.57%)	0
Discoloured water	22 (9.36%)	45 (22.84%)	0
Inadequate treatment	1 (0.43%)	27 (13.71%)	1 (20.00%)
Loss of supplies/pressure	49 (20.85%)	15 (7.61%)	0
Microbiological	34 (14.47%)	41 (20.81%)	0
Taste/odour	27 (11.49%)	10 (5.08%)	0
Health concern	3 (1.28%)	24 (12.18%)	0
Public concern	73 (31.06%)	17 (8.63%)	2 (40.00%)
Other	8 (3.40%)	9 (4.57%)	2 (40.00%)
Total	235	197	5

Unsatisfactory aesthetics Turbidities above 4.0 *NTU* can be detected by the naked eye (WHO, 2011, p. 228-229); discolouration can therefore be detected by customers, unlike certain chemical and microbiological water quality issues, albeit in a subjective manner. Discoloured water can also stain laundry (DWI, 2013b).

In England and Wales the number of customer contacts regarding discolouration has fallen in recent years from 70,643 in 2008 to 50,456 in 2012 (DWI, 2013b), which is in part due to capital and operational works such as mains replacement and cleaning (DWI, 2013b). However, the supply of discoloured water is still a cause for concern as 50,456 contacts regarding discolouration is a significant proportion (38.1%) of the 132,470 contacts made in 2012 regarding water being unacceptable to consumers (DWI, 2013b).

The Maximum Admissible Concentrations (MACs) for regulated inorganic substances may be exceeded Unacceptable water aesthetics can 'mask' and distract from changes in water quality that may impact on human health (DWI, 2009). The World Health Organisation state that the presence of inorganic suspended solids in distribution systems may not always be a health risk e.g. if the material is from "groundwater minerals or from post-precipitation of calcium carbonate" (WHO, 2011, p. 228-229). However, when Cook (2007, p. 92) applied regression to turbidity and iron concentrations recorded during hydraulic disturbances he found that the UK regulatory limit for iron could be exceeded at relatively low turbidities (< 0.6 *NTU*).

Potential microbial risk The turbidity of supplied water has been linked to health risks in several studies (Mann et al., 2007). An increase in turbidity has been shown to negatively correlate with residual disinfectant concentrations under certain conditions (LeChevallier et al., 1981) through increasing chlorine demand (Gauthier et al., 1999a) and stimulating microbial growth (WHO, 2011).

The majority of micro-organisms in chlorinated DWDSs are sessile bacteria that are attached to the pipe wall or particle surfaces within *biofilms* (Batté et al., 2003; Srinivasan et al., 2008). A biofilm is a heterogeneous matrix of microorganisms and extracellular polymeric substances (EPS) that develops on and adheres to wetted surfaces and can protect organisms from disinfectant (Gauthier et al., 1999b). They can contain "viruses and parasites" (LeChevallier, 2005; Douterelo et al., 2013) and provide a mechanism for harbouring species that would not be able to grow in a planktonic state. Increases in flow can result in the sloughing of microorganisms and pieces of biofilm from the pipe wall (USEPA, 2006, p. 12), which could then pass through DWDSs to consumers. However, the risk posed by the development and erosion of biofilms is limited as the vast majority of microorganisms in DWDS biofilms are not pathogenic.

In many countries microbial risk is in part controlled using a disinfectant residual but, as high concentrations of suspended material provides a chlorine demand, all chlorine residual may be undesirably consumed by suspended particulate matter following a discolouration event, which could increase the microbial risk downstream within the DWDS (Gauthier et al., 1999b).

2.1.4 Characterisation of discolouration material

The particle size distribution (PSD) determined by Boxall et al. (2001) for a variety of UK distribution networks had a modal diameter of $8 - 16\mu\text{m}$ ($n=30$). 70% of the material (by counts) was less than $50\mu\text{m}$ in diameter. In the Netherlands, particle diameters have been found to mostly lie within the range $3 - 12\mu\text{m}$ (van Thienen et al., 2011b). The modal particle diameter for samples from pipes within two Australian DWDSs was found to be $11\mu\text{m}$ (Ryan et al., 2008, p. 26).

In the UK the **specific gravity** of discolouration material was found to be 1 to 1.3 (Boxall et al., 2001). The mean specific gravity for twenty sites (pipes and tanks) in a study in Melbourne, Australia, was found to be 1.64 (Ryan et al., 2008, p. 27).

Turbidity has been shown to positively correlate with **iron** and **manganese** concentrations (Boxall et al., 2003b) in UK distribution systems, although the ratio of iron to manganese varies considerably. Field studies have shown that increases in turbidity can correspond to increases in lead, aluminium, zinc, copper and calcium concentrations (Seth et al., 2004; Barbeau et al., 2005).

Discolouration material is also partly comprised of organic material (Gauthier et al., 1999a) including microorganisms that have become entrained in the bulk flow through the partial erosion of highly heterogeneous biofilms from the pipe wall (see §2.1.3 and Douterelo et al., 2013). Numerous studies have reported on the relationship between turbidity and microbial concentrations, including the following:

- Whilst monitoring hydraulic disturbances in DWDSs, Barbeau et al. (2005) found that bacterial concentrations correlated significantly ($P < 0.01$) with TSS ($r = 0.78$) and the log of turbidity ($r = 0.69$);
- Lehtola et al. (2006) found that turbidity, heterotrophic plate count (HPC) and enumerated bacterial counts increased within a bench-top rig following sudden, large changes in pressure.

The relationships between the microbial, abiotic organic and inorganic components of heterogeneous discolouration material can be complex: biofilms can corrode the walls of unlined iron pipes (Lee et al., 1980; McNeill and Edwards, 2001; Wang et al., 2012) and can mediate the deposition of iron and manganese from the bulk water (Sly et al., 1990; Douterelo et al., 2014).

To circumvent having to characterise the composition of discolouration material and the aforementioned relationships, the material mobilised by a hydraulic disturbance is typically considered to be homogeneous and principally defined by its light-scattering properties i.e. quantified as turbidity.

2.1.5 Light scattering and measuring turbidity

Light in a medium can be scattered by suspended particles. For a given volume of medium, the amount of scattering can be quantified in terms of the *intensity* of the incident light (I_0) and

intensity of the light that passes out of the volume (the transmitted light, I_t). The turbidity, T , of the volume can then be expressed using a modified version of the *Beer-Lambert Law*:

$$T = \frac{1}{L_L} \ln \left(\frac{I_o}{I_t} \right) \quad (2.1)$$

where L_L is the light path length through the volume (Clifford et al., 1995). The amount and direction of light scatter is a function of various factors including the ratio of the representative diameter (volume over surface area) of particles suspended in the medium to the wavelength of the incident light (WRC plc, 1994).

For particles that are small relative to the wavelength, light is dispersed by symmetric Rayleigh scattering. However, for relatively large particles (where the diameter exceeds the wavelength) light scattering is asymmetric: some light is symmetrically scattered due to reflection and refraction but there is also much 'forward scatter' (away from the direction of the light emitter) due to diffraction (Hulst, 1981). Particles of discolouration material should be considered to be large with regards to this distinction as typical diameters (§2.1.4) are larger than the incident wavelengths specified by the most common methods of measuring turbidity in DWDS (860 ± 30 nm for the ISO (2008) 7027 method and 400 – 600 nm for USEPA (1993) 180.1). Each of the standards for measuring turbidity specifies the angle between emitter and detector (including a tolerance), the wavelength of light, the bandwidth of that light, the units of turbidity, the maximum optical path length and the range of turbidities over which the method is valid (Ziegler, 2002). In Europe, the most common method is ISO (2008) 7027, whereas the USEPA (1993) 180.1 method is used in the US. Both methods require that light emitter and detector be positioned perpendicularly (see fig. 2.3); a *nephelometer* is a turbidity meter that quantifies scatter through 90° in *NTU*. With both methods, measurements are related to a *formazin* standard where, in general terms, 1 formazin turbidity unit (FTU) = 1 NTU (Ratnayaka et al., 2009, p. 218). For both ISO 7027 and USEPA 180.1, turbidity is linear with the amount of light scattering over the range 0.012 *NTU* to 40 *NTU* at a site-specific level (Sadar, 2003).

Turbidity should be measured on-site (WHO, 2011, p. 75) as turbidity is sensitive to particle size and the particulate matter in a sample may flocculate over time. Care should be taken that gases do not come out of solution due to depressurisation whilst monitoring turbidity as the resulting bubbles can significantly affect readings (Husband and Boxall, 2008, 2011).

Some turbidity instruments are used for discrete offline sampling (e.g. HACH 2100Q; see table 2.2). Others contain flow cells and can either be temporarily attached to hydrants/wash-outs (e.g. Siemens Hydraclam, ATi A15/76) or permanently installed at strategic locations (e.g. ABB 4690 or ATi A15/76 installed at a WTWs or service reservoirs (SRs)) for continuous sampling. UK water providers have recently begun installing WTW-style permanent turbidity monitoring in key trunk mains to allow for prompt, accurate feedback via Supervisory Control and Data Acquisition (SCADA) whilst adjusting flows.

In recent years instruments have been developed for monitoring multiple water quality species simultaneously at distribution system hydrants (Aisopou et al., 2010a). At present these devices are insufficiently robust and accurate over multi-week or multi-month periods to be attractive to water providers (Aisopou et al., 2011). However, as the technology improves these devices

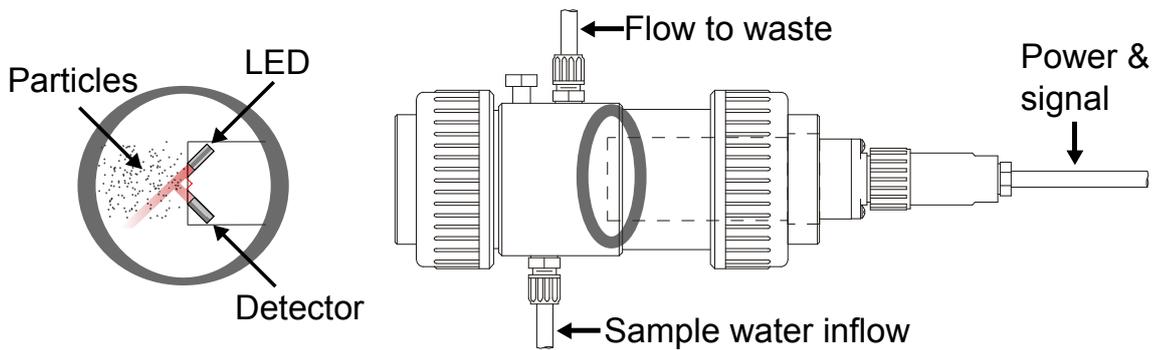


Figure 2.3: Turbidity is a measure of the proportion of light that is scattered through an angle (typically 90°) by suspended particles. Here turbidity is being measured using an ATi A15/76 turbidity meter (ATi, 2011), which is encapsulated within a flow cell.

Table 2.2: Some commonly used turbidity meters.

Instrument	Type	Range	Accuracy	Repeatability	Resolution
HACH 2100Q (HACH, 2010)	Portable discrete offline sampling	0–1000 NTU	±2% of reading plus stray light (≤ 0.02)	±1% or 0.01 NTU (whichever is greater)	0.01 NTU on lowest range
Siemens Hydra-clam (Siemens, 2013)	Portable continuous sampling (fits onto hydrants)	0.1-10 NTU	± 5% of reading ± 0.1 NTU	-	0.05 NTU
ATi A15/76 (ATi, 2011)	Portable or permanent continuous sampling (flow cell or immersed head)	4/400 NTU (0-4.000, 0-40.00, 0-400.0) 40/4000 NTU (0-40.00, 0-400.0, 0-4000)	40/400 scale: ± 5% or ± 0.02 NTU, whichever is greater 400/4000 scale: ± 10% or ± 2 NTU, whichever is greater	-	-
ABB 4690 (ABB, 2012)	Permanent continuous sampling (flow cell)	0-40 or 0-400 NTU	Low range: ± 2 %; High range: ± 5 % or 0.3 NTU	0 to 200 NTU: <1 % 200 to 400 NTU: 2 %	Low range: 0.003 NTU High range: 0.3 NTU

could be deployed to produce detailed datasets showing how parameters such as turbidity, temperature, conductivity, Cl concentration and dissolved oxygen concentration vary temporally and spatially over a distribution network (Gaffney and Boulton, 2011; Leeder et al., 2012; Mounce et al., 2014a).

Verberk et al. (2009) argued that turbidity data can provide a better understanding of cyclic variation in particle loads for turbidities greater than 0.1 *NTU* but particle counts can highlight changes in water quality at low loads. Particle counter instruments can be used to give an understanding of particle size distributions (Vreeburg, 2007) but the analysis of such data is complicated by each timestep being associated with a discrete distribution rather than a scalar.

2.1.6 Guidance and regulation

The WHO recommends that turbidity in distribution networks should not exceed 1 *NTU* so as not to adversely affect the efficacy of the disinfection residual (WHO, 2011, p. 228-229): the suspended solids that result in turbidity typically have an oxidant demand. The European Commission's Drinking Water Directive restates this guideline (for surface waters) (EC, 1998, p. L 330/47). The WHO believe "large, well-run municipal supplies should be able to achieve less than 0.5 *NTU* before disinfection at all times and should be able to average 0.2 *NTU* or less" (WHO, 2011, p. 228-229). In the United States national regulations specify that turbidity cannot exceed 1 *NTU* for systems that use conventional or direct filtration and 5 *NTU* for systems that use other filtration methods (USEPA, 2009). The DWI specify that turbidity should not exceed 4 *NTU* at customers' taps. In regulations regarding indicator parameters the DWI states that the maximum turbidity at the outlets of treatment works should be 1 *NTU* (prior to disinfection if pre-treatment is used) (DWI, 2010); Dutch regulations are similar (Verberk et al., 2009).

In England and Wales water quality regulations stipulate that turbidity should be sampled at treatment works, service reservoirs and at customers' taps (DWI, 2010). Continuous sampling at treatment works using permanent instrumentation is sufficiently precise and accurate to capture spikes in turbidity but step increases in turbidity within the distribution system are very unlikely to be detected via regulatory sampling programmes due to the low spatio-temporal resolution of the latter. Water providers are therefore typically notified of discolouration by consumers (Mounce et al., 2012a).

2.1.7 Discolouration material sources

Developing strategies for quantifying changes in discolouration risk over time in a DWDS requires an understanding of the possible origins of the particulate matter that accumulates within its pipework.

2.1.7.1 Iron corrosion

Through the oxidation of cast iron pipework and fittings soluble ferrous iron is released, scale develops and corrosion by-products can be further oxidised to form ferric oxide precipitates

(McNeill, 2000). Pipe flushing studies have shown that cast iron pipework exhibits a far greater turbidity response to flow increases than plastic pipes (Husband and Boxall, 2011) due to the mobilisation of those precipitates.

2.1.7.2 Bulk water sources

Notable discolouration responses have been detected in networks that do not feature any iron mains (Vreeburg et al., 2008). The material source in these systems is likely to be the bulk water itself. Highly-elevated turbidities are very occasionally due to sudden changes in raw water quality or to major treatment works failures (DWI, 2011b, p. 64) but more frequent turbidity spikes have also been correlated with the backwashing of treatment works filters (Vreeburg et al., 2008). In addition, Fe or Al-bearing coagulants used at treatment works may form discolouration materials (Husband and Boxall, 2011).

Discolouration material may develop within the distribution system itself: Vreeburg et al. (2008) found that installing 0.1 μm ultra-filtration immediately downstream of the treatment works significantly reduced but did not eliminate the discolouration risk. Intra-network formation of particulate matter may be the result of coagulation plus flocculation or bacterial growth. Verberk et al. (2009) showed that particle volume concentration (PaVoC), the total dry-weight TSS and inorganics in particulate form can increase with distance downstream of a WTW, indicating that particulates form from the flocculation of oxidised solutes.

2.1.7.3 Other sources

Rapid changes in velocity such as can result from a pump trip or valve closure can generate pressure waves within DWDS pipework (Brunone et al., 2000). Both very high and low transient pressures can cause structural damage but low pressures can also cause water to be sucked into the DWDS via cracks and leaks (LeChevallier et al., 2003; Besner et al., 2007a; Collins and Boxall, 2013). This ingress has the potential to contaminate the system with hazardous chemicals, viruses, bacteria (including fecal coliforms) (LeChevallier et al., 2003) and presumably also inorganic particulate material. DWDS contamination with undesirable solutes and particulates may also result from cross-connections and back-flow (USEPA, 2001) and also from maintenance operations (e.g. DWI, 2011a, p. 64) and the associated cleaning flushes (Besner et al., 2007b).

These sources of material are likely to be less significant than WTW breakthroughs or iron corrosion due to their localised, transient nature.

2.1.8 The relationship between turbidity and TSS

Discolouration is primarily quantified as turbidity because changes in the aesthetic appearance of drinking water result in complaints from customers and therefore drive action from water companies. However, understanding how the light-scattering of suspended particulates relates to the concentration of that suspension (total suspended solids (TSS); typically measured in mg L^{-1} rather than as a relative volume) allows for the study of the processes that drive

discolouration. If **TSS** can be measured then changes in particle loadings due to hydraulic disturbances can be assessed; furthermore, if the size and density distributions of that material are also known then the volume of material mobilised by a hydraulic disturbance can be estimated, as can the thickness of that material prior to its detachment from the pipe wall.

Turbidity is not a precise measure of the volume nor the mass of suspended solids in a sample. The “*shape, size and refractive indices of the particles in suspension all affect their light-scattering properties*” (Ratnayaka et al., 2009, p. 218), with the peak intensity of scattered light being for diameters of $0.5\mu m$ and intensity decreasing rapidly either side of that peak (WRC plc, 1994). Over a very wide range of turbidities the relationship with particle loading is far from linear, in part due to shadowing effects (WRC plc, 1994).

The nature and generality of the relationship between turbidity and **TSS** determines how easily the former can be used as a proxy for the latter. If turbidity increases approximately linearly with particle *mass* (and therefore *volume*) then turbidity will more linear with *TSS* than if turbidity increases linearly with particle *diameter*, as for a spherical particle its volume V_p is linear with its mass but is a non-linear function of its diameter D_p ($V_p = (\pi D_p^3)/6$). The literature offers various means for relating turbidity to particle size using the Beer-Lambert Law but no consensus regarding whether this relationship is more linear if size is quantified by particle volume or particle diameter. According to Kissa (1999):

$$T \propto \frac{\pi N \int_0^{\infty} D_p^3 \cdot f(D_p) dD_p}{6\overline{D_p}} \quad (2.2)$$

where N_p is the number of particles per mL and $f(D)$ is a normalised **PSD** function. If all particles are the same diameter then:

$$\begin{aligned} T &\propto \frac{\pi N \overline{D_p}^3}{6\overline{D_p}} \\ &\propto \frac{\pi N \overline{D_p}^2}{6} \end{aligned} \quad (2.3)$$

i.e. turbidity is proportional to $\overline{D_p}^2$. Secondly, according to Clifford et al. (1995), the transmitted intensity can be expressed as:

$$I = I_0 \exp\left(\frac{-b_p L_p}{\rho_p} \frac{K}{D_p} c^{n'}\right) \quad (2.4)$$

where c is the mass of particles per unit volume of fluid (i.e. **TSS**), ρ_p is the density of individual particles, b_p is a shape factor (1.5 for spherical particles) and K is the *extinction coefficient* accounting for the extent of reformation of the light beam after scattering behind the particle. For particles $> 4\mu m$, K is very close to 1.0. The exponent n' is a function of D_p , and is included to take account of reduced exposure of particles to the light source as the concentration increases. Combining eqs. 2.1 and 2.4 gives:

$$T = \frac{b_p K}{\rho D_p} c^{n'} \quad (2.5)$$

so turbidity is likely to be non-linear with D_p . Thirdly, according to [Reddy and Fogler \(1981\)](#):

$$T = \frac{1}{L_p} \cdot \ln\left(\frac{I_0}{I_t}\right) = K\pi\left(\frac{D_p}{2}\right)^2 N \quad (2.6)$$

N is the concentration of particles and K is the scattering coefficient, where:

$$K = K_0 \left(\frac{D_p}{2\lambda}\right)^m \quad (2.7)$$

Here, K_0 is the size-independent component of K , m is an (empirical) exponent of the wavelength which is dependent on the particle size and refractive index. Therefore:

$$T = K_0 \left(\frac{D_p}{2\lambda}\right)^m \pi \left(\frac{D_p}{2}\right)^2 N \quad (2.8)$$

so turbidity is proportional to $D_p^m \cdot D_p^2$ and will only be linear with D_p if $m = -1$.

The experimental evidence for turbidity being linear with [TSS](#) (and therefore particle volume) was stronger than the theoretical evidence but still mixed. Several European and North American studies explored the linearity of this relationship over the range of turbidities typically seen during [DWDS](#) flushing and for different source waters and pipe materials:

- [Boxall et al. \(2003a\)](#) found turbidity could serve as a useful proxy for [TSS](#) as the particle size distributions associated with a number of discolouration studies were similar (UK, French, Canadian and US [DWDSs](#)) and presented an approximately linear relationship for converting turbidity to [TSS](#) ($R^2 = 0.78$; zero intercept forced; turbidities up to $\cong 1000$ *NTU*; sample size and pipe materials not stated).
- When flushing two sub-regions of a Dutch [DWDSs](#) containing only polyvinyl chloride ([PVC](#)) and asbestos cement ([AC](#)) pipes, one of which was supplied with ultra-filtrated water and both having six flushing locations and ~ 500 supply points, [Vreeburg et al. \(2008\)](#) found different approximately linear relationships between turbidity and [TSS](#) for each ($R^2 = 0.65$, $n = 12$; $R^2 = 0.87$, $n = 12$). This discrepancy was attributed to the suspended material in the water that had been ultra-filtrated having different scattering properties, and indicates that the turbidity to [TSS](#) relationship is source-water dependent.
- The light scattering properties of mobilised discolouration matter appear to vary with pipe material and water source. [Schaap and Blokker \(2012\)](#) analysed turbidity and [TSS](#) data from 300 Dutch distribution mains flushes. They found that the relationship between the two variables was most linear when the data were grouped by water source and the proportion of unlined cast iron ([CI](#)) within the distribution area (for all data $R^2 = 0.80$; for the aforementioned groups $R^2 \in [0.85, 0.93]$). [Besner et al. \(2012\)](#) detail a study in Waterloo, Canada, a city supplied from a variety of sources. Each of a number of distribution pipes were each flushed once. A poor non-linear relationship was found when

all data were collated ($\ln(TSS) = 0.5968 \cdot \ln(NTU) + 1.0232$, $R^2 = 0.53$, $n = 180$ (45 flushed pipe lengths, 4 samples per flush)) but more linear relationships were seen when that data (plus data from flushes where TSS was only measured once) were grouped by pipe material (CI: $n = 16$, $R^2 = 0.78$; ductile iron (DI): $n = 15$, $R^2 = 0.68$, PVC: $n = 54$, $R^2 = 0.75$).

- There is limited evidence of linearity in larger diameter trunk mains: Vreeburg and Beverloo (2011) flushed two trunk Dutch mains, one AC and one PVC, and found $TSS = 0.3793 NTU$ in the former ($R^2 = 0.95$, $n = 12$, $\varnothing = 315 mm$) but the relationship was far from linear in the latter ($TSS = 0.2571 NTU$, $R^2 = 0.46$, $n = 10$, $\varnothing = 500 mm$).
- In contrast, Ryan et al. (2008) monitored a DWDS in each of six locations around Australia and found no associations between particle size distribution, TSS and turbidity. The modal particle size for 'background' flows was found to decrease with distance from the WTW and within 100mm distribution mains towards the extremities of the network most particles lay within the range $5 - 20 \mu m$, although second local PSD mode of $37 \mu m$ was found at all sites. When flows were elevated, the mobilised discolouration material was found to be similar but possibly slightly smaller in diameter. The lack of a correlation between TSS and turbidity may or may not be due to a difference in discolouration material composition (and therefore light-scattering properties) between Australian and European DWDSs: the material in some but not all Australian systems is dominated by clay (Prince et al., 2000; Ryan et al., 2008).
- Chanson et al. (2008) summarised correlations between turbidity and TSS for several studies of open waters. Eight of the quoted thirteen expressions were linear but the degree of linearity was not specified.

Few have explored how the light scattering properties at a sampling point within a DWDS vary over time and between hydraulic disturbances of different magnitudes. The aforementioned study by Vreeburg et al. (2008) in part explored variation in suspended particle size in two sub-networks over a year using particle counters but the collected data has been aggregated in time in a way that does not allow for an exploration of the correlation between particle loading, size distribution, turbidity and other factors of interest such as flow and metals concentrations. Verberk et al. (2009) monitored size distribution (particle counter) data and turbidity in various Dutch DWDSs over a week and found that trends in turbidity were not obvious or present in total particle count time-series.

Verberk et al. (2009) highlight that, although there is merit in studying variation in light-scattering properties when investigating discolouration processes, interpreting particle counter data for this purpose can be difficult: each sample comprises a value for each of several size bins (31 in the case of some of the sites monitored by Verberk et al. (2009)) rather than just a scalar as is the case when measuring turbidity.

Particle count data can be aggregated to yield a univariate time series by converting to particle volume concentration (PaVoC) [m^3/m^3] (Vreeburg et al., 2008; Verberk et al., 2009) but a failure to account for asymmetry in particle geometries could result in large errors. Alternatively

material fluxes past a sampling point can be quantified in units of mass (TSS) rather than volume (PaVoC) by measuring dry weight (Vreeburg et al., 2008) but the temporal resolution of this method is limited by needing to capture water for several hours to have enough mass for analysis and by the drying process being offline and time-consuming.

In conclusion, turbidity may be sensitive to particle size and shape but taking and interpreting measurements is simpler and quicker than for PaVoC and dried weight TSS analysis. Also, it is not TSS that drives customer complaints regarding aesthetics but turbidity.

After conducting flushing operations, Cook and Boxall (2011) in the UK and Vreeburg and Beverloo (2011) in the Netherlands made efforts to quantify the thickness of the cohesive layers that had been eroded by the increases in flow. In both studies a linear relationship between turbidity and TSS was assumed (Cook and Boxall, 2011) or calculated (Vreeburg and Beverloo, 2011) as was a particle density. The total volume of mobilised material could be calculated then normalised by pipe internal surface area to give a thickness. Cook and Boxall (2011) found that the flushing of previously undisturbed mains could generate notable turbidities from layers estimated to be $< 0.57\text{mm}$ thick. Therefore, discolouration in DWDSs is most likely due to hydraulic disturbances causing the erosion of these thin layers rather than from the erosion of large turbercles.

2.1.9 Transport of material from bulk water to wall

Discolouration material has traditionally been viewed as sediment that settles under gravity along the pipe invert. However, Boxall et al. (2001) found that the size and density distributions of discolouration material prohibits particles from settling under all but quiescent conditions. Such conditions are rarely found in distribution systems so Boxall et al. (2001) concluded that particles must be transported to the wall by other means and also that they will remain entrained in the bulk flow as a *wash load* once mobilised. van Thienen et al. (2011b) concurred that the migration of material from the bulk flow to the pipe wall cannot solely be due to sedimentation under typical hydraulic DWDS conditions.

Vreeburg (2007, p. 66) dosed a laboratory rig containing a perspex pipe ($\varnothing = 0.1\text{m}$) with iron flocs and found that the flocs collected over time around the entire circumference at bulk velocities of 0.14ms^{-1} but only around the lower part of the pipe at 0.06ms^{-1} (fig. 2.4). This indicates that radial transport processes may be significant at the higher velocity. It was not stated whether the flocs used in the experiment were representative of discolouration material with regards to diameter and density.

A correlation has been found between the quantity of wall-bound discolouration material, measured using turbidity, and sessile biofilm, measured using microscopy techniques (Sharpe, 2013). Sharpe (2013) and Fish (2014) monitored biofilm growth around the circumference of a high-density polyethylene (HDPE) pipe ($\varnothing = 79.3\text{mm}$) and found no statistical difference in the microbial community or biofilm physical structure between the soffit, middle and invert of the pipe.

The processes that drive solutes to the pipe wall are dependent on hydraulic regime and so it is possible that the transport of fine suspended particles to the wall may also be dependent on the Reynolds number (Re). There are several validated approaches to modelling *solute transport*

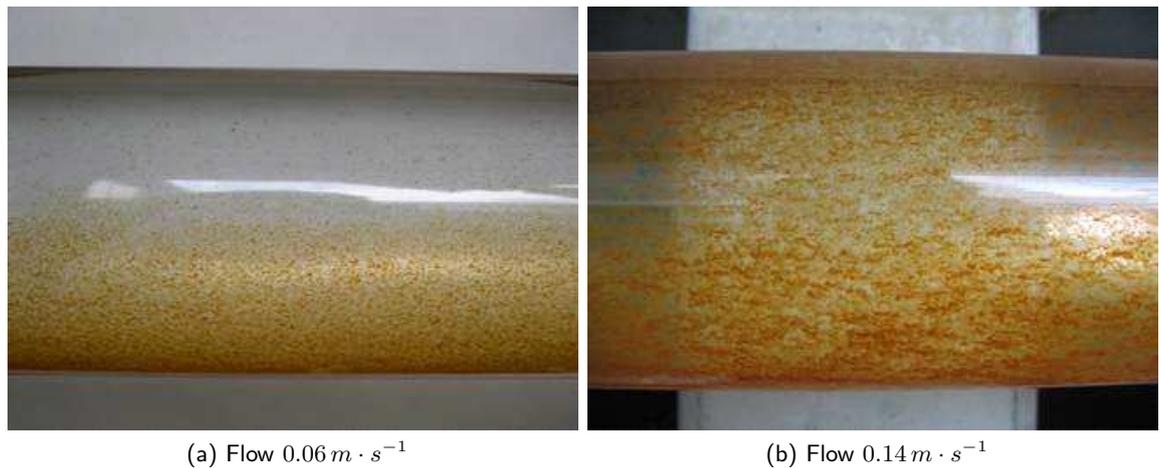


Figure 2.4: Photographs of accumulated material within a $\varnothing = 100\text{ mm}$ Perspex pipe loop after four days of re-circulation with high concentrations ($10\text{ mg} \cdot \text{L}^{-1}$) of ferric chloride (Vreeburg, 2007, p. 62).

to the pipe wall; these are dependent on Re and molecular diffusivity (Biswas et al., 1993; AWWA Research Foundation, 1996; Rossman, 2000). The radial transport of discolouration *particulates* may be driven by similar processes due to particles being light enough to exhibit Brownian motion.

van Thienen et al. (2011a) investigated whether radial transport processes such as *turbophoresis*, the transport of particles down a turbulence gradient due to particle inertia, could contribute to the migration of material from the bulk flow to the pipe wall. *Turbulent diffusion* was found to be the radial transport processes that dominates radial transport as particle diameters and bulk flow velocities are typically too small/low for inertia-driven mechanisms such as *turbophoresis* to be significant.

2.1.10 Mechanisms for binding to the pipe wall

If discolouration material were held at the wall by radial transport processes alone then it would all be evacuated upon the draining of the system. Vreeburg (2007, p. 67) found that some material remained bound to the pipe wall of a rig after draining, indicating that the material had cohesive properties. When Mn reacts with oxidants such as oxygen or chlorine it precipitates and then accumulates on pipe and tank internal surfaces as a black 'slime' (Kohl and Medlar, 2006; Ratnayaka et al., 2009). Fe-bearing corrosion scales can line the walls of Fe pipes (McNeill, 2000).

Micro-organisms can become attached to pipe walls when they get sufficiently close that attractive forces come into play; Dunne (2002) termed this primary adhesion mechanism 'docking'. These sessile organisms then release EPS that allow them to more strongly adhere to the wall ('locking'). The resulting biofilm, which is predominantly comprised of extracellular carbohydrates, proteins, DNA and live and dead organisms (Denkhaus et al., 2007; Jakubovics et al., 2013) then provides a sticky surface along which the co-deposition of organisms (Batté et al., 2003) and inorganic material (Zacheus et al., 2001) can occur. Husband et al. (2008) found that metal uptake rates in a laboratory pipe rig were greater after four weeks into an experiment than one week, suggesting that biofilm growth may be a precursor to metal oxides

attaching to pipe walls. The two other stages of the biofilm life cycle are maturation (growth) of the biofilm and detachment (Dunne, 2002), which could be the active detachment of single organisms or could be due to sloughing driven by an increase in hydraulic forces (Stoodley et al., 2001b). The cycle can then repeat through detached organisms then ‘docking’ at pipe/tank walls again further downstream (Vaerewijck et al., 2005).

2.1.11 Macro- and micro-scale heterogeneity of discolouration material

Discolouration material does not form on pipe walls in entirely smooth layers. Spatial heterogeneity has been seen with the naked eye through the wall of a perspex pipe in a laboratory pipe rig (fig. 2.4) and also through conducting internal inspections of drained-down trunk mains (fig. 2.5).

Spatial heterogeneity has also been observed at the macro scale. Biofilms, a component of discolouration material in DWDS (§2.1.3; §2.1.4), have been shown to vary somewhat in the ratio of volume to coverage area (a proxy for thickness) between fields of view when imaged in three dimensions using a Confocal Scanning Laser Microscope (Fish, 2014, §5.2.1.4). Also, complex three-dimensional structures can be seen in biofilms that have developed over twenty-eight days under conditions not unlike those in DWDS (fig. 2.6).

2.1.12 Shear-stress-driven discolouration material erosion

2.1.12.1 A velocity or shear-stress driven process?

It is widely acknowledged that particulate material at the pipe wall becomes entrained in the bulk flow upon an increase in hydraulic forces. There are however several views on which factors drive material erosion.

One theory of erosion is that discolouration material is comprised of loose sediments that are mobilised when a threshold velocity is exceeded. However, this conflicts with the evidence in §2.1.9 that the material is not sediment as particles are sufficiently small and light for them not to typically settle under gravity in all but fully quiescent conditions.

Boxall et al. (2001) took an alternative view on discolouration material erosion and accumulation: the particle size and density analysis described in §2.1.9 and “the observation of Mehta and Lee (1994) that particles of this size range ($< 20\mu m$) normally exhibit cohesive properties” led them to theorise that discolouration material binds to the pipe wall in *cohesive layers* that are conditioned by the hydraulic shear stress (Boxall and Saul, 2005).

Vreeburg (2007) concurred with this argument after considering the gravitational settling of non-cohesive sediment by Stokes’ law and its resuspension according to a Shields-based empirical formula. Vreeburg (2007) concluded that the velocities required for the resuspension of discolouration material (particle diameter $\in [1, 25] \mu m$ and density $\in [1050, 2600] kg \cdot m^{-3}$) are sufficiently frequently exceeded in real DWDS for it to only be possible for material to accumulate if it adheres to the pipe wall in cohesive layers. Subsequently, Pothof and Blokker (2012) compared the Soulsby (1997) critical shear stress for mobilisation to shear stresses experienced in real DWDSs. They also concluded that typical discolouration material



Figure 2.5: The internal surface of a $\varnothing=350\text{mm}$ unlined *DI* trunk main, taken after the main had been drained down. Image supplied by Wessex Water.

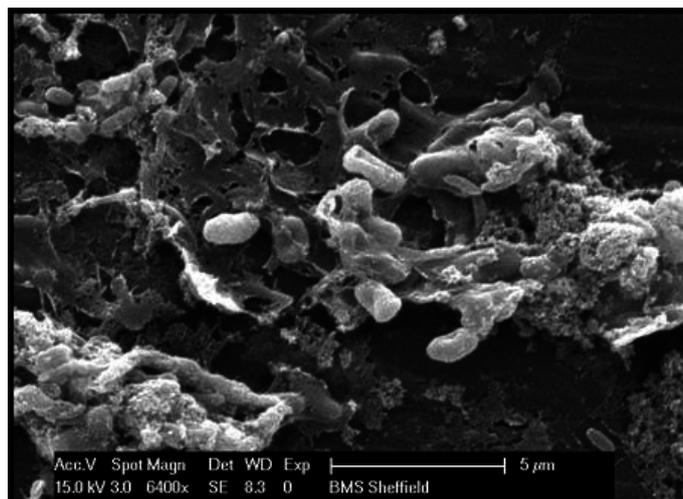


Figure 2.6: A Scanning Electron Microscope (SEM) image of a biofilm on the wall of a $\varnothing=79.3\text{mm}$ *HDPE* main (from Fish, 2014). The main was subjected to a low-varying diurnal flow pattern at 16°C for twenty eight days, during which time a biofilm developed on the pipe wall. This image is of a biofilm sample from after the main was subsequently flushed at a much higher flow rate.

($D_s < 25 \mu m$; $\rho_s < 1.3 kg \cdot m^{-3}$) can not accumulate within DWDS pipework if it is non-cohesive unless the daily peak hydraulic forces are always very low, suggesting that material most likely develops in cohesive layers and classical non-cohesive sediment erosion models such as those of Ackers and White (1973) and Soulsby (1997) are not representative of discolouration material mobilisation processes.

Pothof and Blokker (2012) express the daily peak hydraulic force required for the mobilisation of non-cohesive discolouration material of typical densities and diameters as a bulk velocity ($0.06 m \cdot s^{-1}$) yet the erosion of cohesive materials is typically considered to be driven by the shear stress τ_a at the material/fluid boundary, which is a function of not just velocity but also hydraulic radius and boundary roughness. Under steady state conditions the shear stress at the pipe wall is as per eq. 2.9.

$$\tau_a = \rho g D_H \frac{h_f}{L} \quad (2.9)$$

where

τ_a	applied shear stress [$N m^{-2}$]
ρ	fluid density [$kg m^{-3}$]
D_H	hydraulic diameter, equal to the diameter for a pipe with a circular cross-section [m]
L	pipe length [m]
h_f	flow-dependent head loss over the length of the pipe [m], which is defined as:

$$h_f = f \cdot \frac{\bar{u}^2}{2g} \cdot \frac{L}{D_H} \quad (2.10)$$

known as the Darcy Weisbach formula (Chadwick, 2004), where

\bar{u}	bulk velocity [$m s^{-1}$]
ρ	fluid density [$kg m^{-3}$]
g	acceleration due to gravity ($9.81 m s^{-2}$)
f	Darcy Weisbach friction factor [-], which is equal to:

$$f = \begin{cases} \frac{64}{Re} & Re < 2000 \\ \frac{0.25}{\left(\log_{10}\left(\frac{k_s}{3.7D} + \frac{5.74}{Re^{0.9}}\right)\right)^2} & 4000 \leq Re \end{cases} \quad (2.11)$$

assuming the Hagen Poiseulle formula for laminar flow (Bhave, 1991) and the Swamee Jain approximation to the Colebrook White formula for turbulent flow (Bhave, 1991). For transitional flow ($2000 \leq Re < 4000$), Re can be found by cubic interpolation within the Moody diagram (Dunlop, 1991). Here:

k_s Nikuradse equivalent sand roughness [m]

Re Reynolds number [–], the ratio of inertia to viscous forces, which is equal to:

$$Re = \frac{\rho U D_H}{\mu} = \frac{U D_H}{\nu} \quad (2.12)$$

where

μ dynamic viscosity [Pa · s]

ν kinematic viscosity [$m^2 s^{-1}$]

For reference, the influence of the bulk velocity and bulk flow on the shear stress at the pipe wall is shown in fig. 2.7.

2.1.12.2 The erosion of conditioned material by the ‘excess shear stress’

Wall-bound discolouration material does not have a single shear strength but has a strength versus quantity profile, where the lower bound of this strength range is typically the prevailing shear stress i.e. the discolouration material is *conditioned* by the usual pipe hydraulics. The strength profile of wall-bound material can be studied by increasing the shear stress in steps above and beyond the ‘background’ shear stress, as illustrated in fig. 2.8 (after Husband and Boxall, 2010a). In these examples note that more discolouration material is mobilised by each flow step; were it non-cohesive sediment rather than cohesive layers then it would all be mobilised by the first flow increase.

Intuitively one might assume that a trebling rather than doubling of the absolute shear stress at the pipe wall would cause greater and quicker material erosion, but studies have strongly indicated that the rate of erosion is dictated by the amount by which the shear stress at a given time *exceeds the current shear strength* of the wall-bound material (henceforth this quantity will be referred to as the *excess shear stress*), with this shear strength being shaped by the prevailing shear stress.

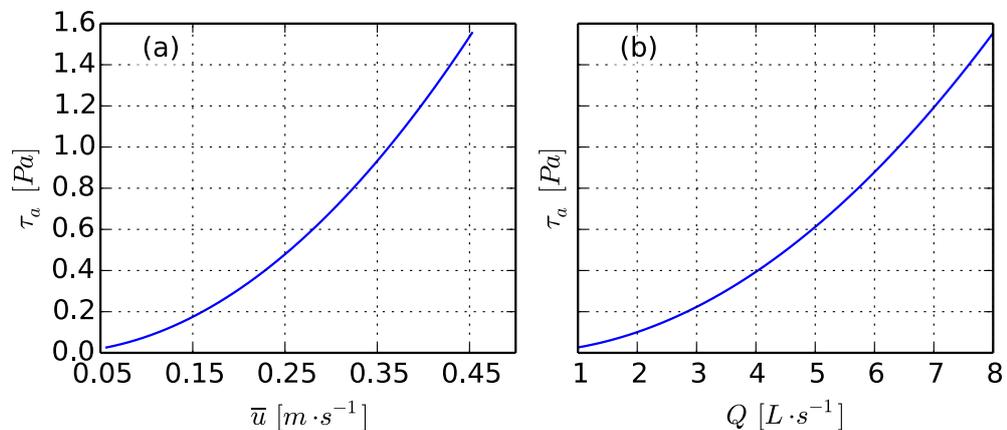


Figure 2.7: The impact of (a) bulk velocity \bar{u} and (b) bulk flow Q on the shear stress τ_a at the pipe wall (diameter $\varnothing = 150$ mm; pipe roughness $k_s = 5$ mm)

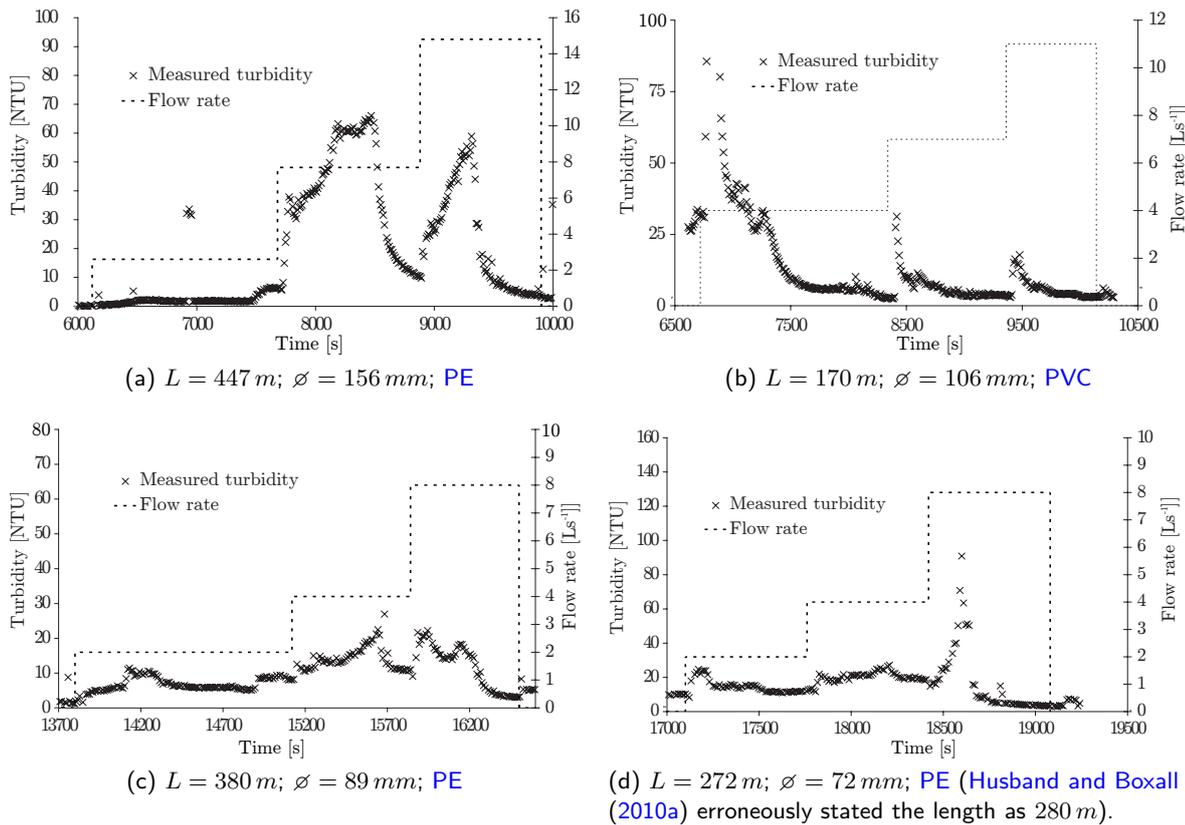


Figure 2.8: The discolouration material bound in cohesive layers to pipe walls has a strength versus quantity profile that can be explored by repeatedly increasing the shear stress in steps above prevailing levels (after Husband and Boxall, 2010a).

The importance of the excess shear stress as an eroding force is concisely illustrated in the first of two case studies presented in Husband and Boxall (2012): (see fig. 2.9). In the studied trunk main system, enabling works required that flows be increased in three distinct pipe lengths (A , B and C). The relative increase in absolute shear stress was greatest in pipe C (it quadrupled) but large turbidity responses were only seen for pipes A and B . Husband and Boxall (2012) attributed this to A and B having experienced greater excess shear stresses as a result of the increase in flow: in those two mains the relative difference in shear stress pre- and post-flow increase was greater than for C (fig. 2.9), which would equate to a larger excess shear stress (if the prevailing shear stress prior to the flow increase had *conditioned* the wall bound material so it and the material shear strength were similar).

Studies of cohesive layers of particulate material in various research areas have found the layer strengths to be conditioned by prevailing hydraulic forces and the rate and amount of material erosion to be driven by the excess shear stress: such behaviour has been witnessed for discolouration material (as mentioned above and in Husband and Boxall, 2010a), for biofilms (through sloughing; Telgmann et al., 2004) and, in the field of marine geology, for cohesive estuarine muds (Sanford and Maa, 2001).

Learning of the importance of the excess shear as an eroding force can reshape one's understanding of which pipes in a DWDS are at greatest risk of causing discolouration. For example, pipes such as dead ends that usually experience low, variable shear stresses can be associated with the greatest discolouration risk.

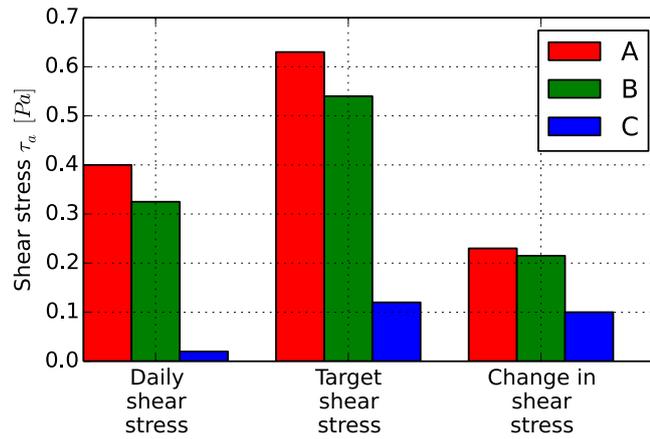


Figure 2.9: Flows in three pipe lengths (A, B and C) within a UK trunk main system were increased above prevailing levels. The relative increase in absolute shear stress was greatest in pipe C but large turbidity responses were only seen in A and B, demonstrating that the excess shear stress (the target shear stress minus the typical daily shear stress) is the primary driver of the erosion of cohesive layers of discolouration material from pipe walls (after Husband and Boxall, 2012).

2.1.12.3 Unsteady hydraulics and transient pressure waves

Discolouration events are due to sudden increases in hydraulic forces yet eq. 2.9, an expression for the Darcy-Weisbach shear stress, assumes simple steady-state hydraulic conditions (the fluid has no inertia, has a constant density and is conveyed through inflexible conduits). The total shear stress at the wall under unsteady conditions can far exceed that predicted by eq. 2.9 (Brunone et al., 2000). Naser et al. (2006) and Aisopou et al. (2010b) argue that discolouration material erosion could be significantly affected by the unsteady shear stress as it can be three to ten times the steady shear stress (Aisopou et al., 2010b).

Characterising and modelling unsteady shear stresses is difficult due to the complexity of the transverse velocity profile under transient conditions. The sudden closure of a valve results in the generation of transient pressure waves and in pressure gradients that decelerate the (compressible) fluid; Brunone et al. (2000) showed experimentally that this can result in momentary flow reversals near the wall due to the velocity being lower there than at the centreline prior to the closure. In addition, under turbulent, unsteady conditions velocity profiles can be temporarily asymmetric. As a result of the complexity and rapid variation in velocity profile unsteady shear stresses can be high even though the bulk discharge for a pipe is relatively low (Brunone et al., 2000).

Few have experimentally studied the impact of unsteady hydraulics on discolouration material erosion. Mustonen et al. (2008) generated 'pressure shocks' within very narrow diameter lab rig pipework after first allowing material to develop on its walls. These transients were found to mobilise biofilms and 'loose deposits', however no attempt was made to quantify and model the relationship between the dynamic pressure changes and material erosion.

There is also an unsteady shear stress at times where the bulk velocity is not changing rapidly: this is due to turbulence. Its contribution to material erosion will most likely be dwarfed by the viscous shear stress (eq. 2.9) as it is negligible within the laminar sub-layer (Graebel, 2001) and the velocity fluctuations due to turbulence are minor compared to the time-averaged velocity (Tennekes, 1972).

2.1.12.4 Effect of temperature and wall-bound discolouration material quantity on shear stress

The shear stress at the pipe wall is a function of Reynolds number (the ratio of inertia to viscous forces) and so of the dynamic viscosity and density (eq. 2.12), both of which are temperature dependent. However, the significance of this temperature dependency has not been explicitly explored in the literature, most likely because of too many other discolouration-related parameters such as biofilm growth and corrosion also being temperature dependent but not being easy to control for.

Shear stress is also a function of wall roughness (eq. 2.11) but the impact of time-variance in the amount of discolouration material bound to the pipe wall on roughness and shear stress has also not been explored to date. Changes in roughness could potentially be detected through measuring headloss over a pipe length whilst upping the flow in step increments.

2.1.13 The cycle of material accumulation and erosion

2.1.13.1 Introduction

Discolouration risk would potentially be simple to manage were it just a case of ensuring a pipe had been cleaned (abrasively or hydraulically) once before an unavoidable increase in shear stress. However, as is discussed in this section, discolouration material has been shown to continually, repeatably and ubiquitously accumulate on pipe walls following a period of erosion. Therefore, although discolouration risk can be reduced by one-off capital and operational measures, these commonly need to be coupled to an ongoing operational programme (at least in networks that cannot be entirely reconfigured/replaced to 'design out' the risk).

2.1.13.2 Quantifying discolouration material accumulation; using the integral of material flux to quantify amounts of discolouration material

Water providers are keen to characterise material accumulation processes¹ to inform the design of operational strategies for managing discolouration risk. This is made difficult by the complexity of subterranean DWDS and the physical and biochemical reactions that occur within them. Erosion processes can be qualified and quantified through the direct measurement of turbidity but the study of accumulation processes in the field and lab typically requires more indirect methods given the processes' largely latent nature. Most methods involve comparing the turbidity response from two controlled hydraulic events (typically unidirectional pipe flushing) that bookend a period over which accumulation is believed to have occurred (Blokker et al., 2011; Husband and Boxall, 2011; Cook and Boxall, 2011; Vreeburg et al., 2008). The approach taken by Husband and Boxall (2011) when studying accumulation in DMA's was as follows:

1. Select the pipe(s) to be flushed. Ideally, these should not have been affected by hydraulic anomalies during the previous 1-4 years (duration depends on pipe material).

¹Also referred to as material *regeneration*, although a more specific definition of regeneration is given in §4.

2. Ensure that the pipes upstream are either cleaned first or are of a much larger diameter so that a significant increase in flow in the 'target' pipe(s) will only have a negligible impact on shear stresses upstream.
3. Open a hydrant at the downstream end of the target pipe(s) and flush whilst continually monitoring turbidity and flow at/through the hydrant. In 'through mains' (not dead ends) the flushing may need to be undertaken at night to reduce the error in flow measurements should flow be monitored via a standpipe attached to the hydrant.
4. Turbidity, flow and pressure should be monitored until the turbidity of the water emerging from the hydrant reaches the background levels seen prior to flushing. Turbidity should ideally be monitored at a high temporal resolution ($\lesssim 10s$)
5. The amount of material mobilised per flush is then approximately quantified using several different metrics such as (a) the peak turbidity during the flush, (b) the average (presumably the arithmetic mean) turbidity or (c) turbidity integrated with respect to time (with units of $NTU \cdot s$), shown below:

$$\int_{flush_start}^{flush_end} T_{ds}(t) dt \quad [NTU \cdot s] \quad (2.13)$$

where

$T_{ds}(t)$ Observed turbidity at downstream end of the pipe (assuming that the water at the upstream end has negligible turbidity) [NTU]
 t time [s]

6. A second flush is then conducted X months later during which flushing technicians aim to achieve the same flow through the target pipe(s). The quantity of mobilised material from this flush is calculated using the same metric as before.
7. The difference between the measures of the amount of mobilised material and time difference between the flushes provides either an absolute or relative measure of accumulation or a relative measure (where the use of a relative measure requires the assumption that the discolouration potential immediately prior to the first flush was maximal given the prevailing hydraulic conditions and layers were not regenerating significantly). Hydraulic disturbances between the flushes would of course affect such measures (Cook, 2007, p. 81).
8. Subtracting the amount of material mobilised by the second flush from the amount mobilised by the first and dividing by the time difference gives an absolute accumulation rate. Alternatively, division can be used in place of subtraction to give a relative accumulation rate; this requires the assumption that the discolouration potential immediately prior to the first flush was maximal given the prevailing hydraulic conditions and layers were not regenerating significantly. Hydraulic disturbances between the flushes would of course affect the accuracy of these accumulation rate estimates (Cook, 2007, p. 81).

Methods based on this approach have been applied within [DMAs](#) to study if and how accumulation processes vary with time/season, source water, asset characteristics and network topologies ([Cook and Boxall, 2011](#); [Husband and Boxall, 2011](#)).

In the aforementioned method, the amount of material mobilised per flush is best quantified as turbidity integrated over the flush duration that metrics insensitivity to noise. The metric can be thought of as being proportional to the mass of mobilised material so long as flow is constant and turbidity is assumed to be proportional to [TSS](#) so can therefore be treated as a concentration. This second assumption has been shown to be valid in many [DWDS](#) (§2.1.8).

If the flow is *not constant*, as is often the case during field studies of material shear strength versus quantity profiles (see fig. 2.8), then this metric will produce erroneous results due to it not including flow as a factor. Fig. 2.10 demonstrates the importance of quantifying the total amount of material mobilised by a shear stress increase using both turbidity *and* flow to calculate the *material flux*. Here the turbidity at the downstream end of the pipe, T_{ds} , has been discretely sampled every Δt since time t_0 . Fluctuations in T_{ds} have been driven by fluctuations in pipe flow, Q , the latter having been sampled with the same frequency. Each turbidity measurement at time t is assumed to be representative of the parcel of water that passed the observation point over say $t \pm (\Delta t/2)$. If turbidity is considered to be a pseudo-concentration then the total amount of material per parcel is $T_{ds}(t) \cdot Q(t) \cdot \Delta t$ (shown as green bars in fig. 2.10). Integrating the material flux $T_{ds}(t) \cdot Q(t)$ that passes out of the pipe over the period t_0 and t therefore gives the total amount of material to have passed that observation point between t_0 and t (green line in fig. 2.10, hence the importance of considering both turbidity *and* flow (eq. 2.14) when estimating the amount of material that has been mobilised from the pipe wall using observed time-series data.

$$A_w \cdot N = \int_{flush_start}^{flush_end} Q(t) T_{ds}(t) dt \quad (2.14)$$

where

$A_w \cdot N$ the amount of material mobilised from each m^2 of the pipe's internal surface ($N, [NTU \cdot m]$) scaled by total wall area $A_w [m^2]$ to give the total amount of mobilised material $[NTU \cdot m^3]$;

$T_{ds}(t)$ Turbidity at the downstream end of the pipe (here assuming that the water passing into the upstream end has negligible turbidity) $[NTU]$;

$Q(t)$ Pipe flow $[m^3 \cdot s^{-1}]$;

t Time $[s]$.

Note that in circumstances where the turbidity at the upstream end is non-zero the material flux and amount of mobilised material should be calculated using the *net turbidity response* from the pipe, $T_{net}(t)$, rather than just the turbidity measured at the downstream end, $T_{ds}(t)$.

Returning to methods for quantifying material accumulation, others have used different approaches to that presented in [Husband and Boxall \(2011\)](#):

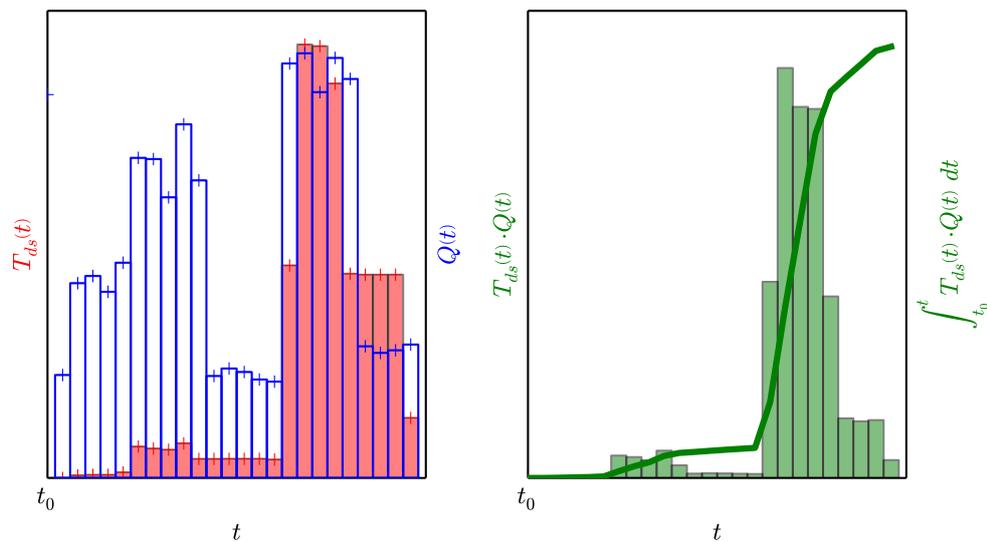


Figure 2.10: An illustration of the importance of measuring the amount of material mobilised from a pipe wall using the material flux passing the downstream end. The material flux (green) is the product of the discretely sampled flow (shown in blue; Q) and downstream turbidity (red; T_{ds}). Here, the cumulative material flux is also shown (green line), this being the total amount of material to have passed the observation point since time t_0 .

- [Schaap and Blokker \(2013\)](#) presented a variant of this method where the total amount of material mobilised per flush is calculated in the same manner but is then divided by pipe length to give a measure of *Locally Accumulated Material* (LAM). Previous studies referenced by [Schaap and Blokker \(2013\)](#) indicated that most material is mobilised during the first turnover per flush so [Schaap and Blokker \(2013\)](#) assumed mobilisation was near-instantaneous and so were able to quantify LAM per metre of flushed pipework.
- Discolouration potential and thus accumulation can be quantified using the Resuspension Potential Method (RPM) ([Vreeburg, 2007](#)). A 0.35 m.s^{-1} increase in bulk velocity is imposed within a pipe for fifteen minutes, during which time turbidity is monitored. Five aspects of the turbidity trace can then quickly be scored to give a RPM score between 0 (no discolouration risk) and 15.
- [Gaffney and Boulton \(2011\)](#) explored whether mass fluxes could be used to quantify erosion from and/or accumulation within a pipe segment but sampled turbidity at such a low rate (15-30 mins) that it is possible turbidity spikes would not have been detected.

2.1.13.3 Accumulation: an ubiquitous, continual and repeatable process

Intuition suggests that the development of enough discolouration material on pipe walls to cause detectable turbidity during flushing would be limited to unlined corrodible iron pipes and/or pipes conveying water with a high particle loading. However, notable quantities of discolouration material can accumulate in most if not all DWDS systems (groundwater and surface water-fed) within pipes of a wide range of materials (including but not limited to unlined CI, PE, PVC, AC and cement-lined CI) and diameters ([Husband and Boxall, 2011](#); [Blokker et al., 2011](#)). Material accumulation is therefore understood to be an **ubiquitous** process.

When [Blokker et al. \(2011\)](#) flushed a variety of pipes within a Dutch distribution area on three occasions (returning 17 then 24 months after the initial flushes) they found that decreasing *amounts* of material were mobilised per flush yet their measure of the accumulation *rate* was fairly constant, suggesting that material accumulation is **continual, repeatable** and potentially linear. [Lehtola et al. \(2004b\)](#) showed diurnal variation in turbidity and other water quality parameters for a Fe pipe with which there was a known discolouration risk; this suggests that cohesive material layers are continually eroding and regenerating at the pipe wall. Studies in the UK have also shown accumulation to be continuous ([Husband et al., 2008](#), figs. 4-7, table 1), potentially linear ([Husband and Boxall, 2008](#); [Cook and Boxall, 2011](#); [Husband et al., 2008](#)) over periods as short as a month and repeatable ([Husband and Boxall, 2008](#)).

Intuition suggests that accumulation may not be linear as seasonal fluctuations in parameters such as temperature will accelerate/decelerate the biochemical aspects of material accumulation processes. Such variation has not been detected in a number of studies as repeated flushing has been conducted on an annual basis (e.g. [Husband and Boxall, 2008](#)) and/or the number of flushes per flushing point has been very few ([Blokker et al., 2011](#); [Cook and Boxall, 2011](#)). All laboratory investigations into accumulation using pipe systems with hydraulics and water quality comparable to real DWDSs have been too short in duration (up to one month) to detect the effects of seasonal variation. However, [Schaap and Blokker \(2013\)](#) did identify a positive relationship between temperature and material accumulation when a distribution main was flushed thirty times in eighteen months. Also, single-month experiments in temperature-controlled facilities have shown temperature to have a notable effect on material build-up ([Sharpe, 2013](#); see also §2.1.16.2 and §2.1.16.9).

The discolouration cycle repeats over different time frames: both between controlled or unexpected anomalous flow increases but also over much shorter periods due to frequent variation in demand, pumping and reservoir flows ([Gaffney and Boulton, 2011](#)). Diurnal turbidity patterns in distribution pipes typically show an increase overnight, attributable to corrosion/stagnation ([Vreeburg, 2007](#); [Cook, 2007](#)) but [Cook \(2007\)](#) presented examples of the turbidity reaching a maximum at the same time as flow in areas dominated by domestic demand, indicating that the patterns also reflect how material that has accumulated whilst flow was relatively low is then eroded by the daily peak turbidity. Suboptimal WTW processes may also give rise to frequent fluctuations in turbidity ([Vreeburg et al., 2008](#)).

2.1.13.4 Variation over time in the discolouration material strength versus quantity relationship

For any sudden increase in flow the risk of exceeding regulatory turbidity limits is a function of a) how much material is present at the pipe wall, b) the strength profile of that material and c) the increase in shear stress. The relationship between the quantity and strength of wall-bound material can be explored at any given moment by imposing a stepped shear stress profile and measuring how much material is mobilised by each excess shear. This highlights one of the principal issues associated with the study of discolouration material: it is difficult if not impossible to characterise the system well without greatly disturbing it. However, once the system has been disturbed and its state is known then this information could be used along

with an understanding of how the relationship between strength and quantity changes over time due to accumulation and erosion to estimate the state at a future time.

One way of exploring how material accumulates over different strength ranges is to impose a stepped shear stress profile whilst measuring turbidity then repeat this exercise after a period of accumulation. Husband and Boxall (2011) flushed 67 pipes in England then returned to each site to flush one or more times using the same flow (and therefore shear stress) profiles, with each visit being separated by a number of months. As can be seen from fig. 2.11, the initial flush of each main generated more turbidity than subsequent flushes, implying that the cohesive material layers had not accumulated to the point where they had reached a steady state over the intervening period(s). However, a turbidity response can be seen during later flushes for each increase in shear stress, indicating that all shear strengths of cohesive layers are regenerating simultaneously.

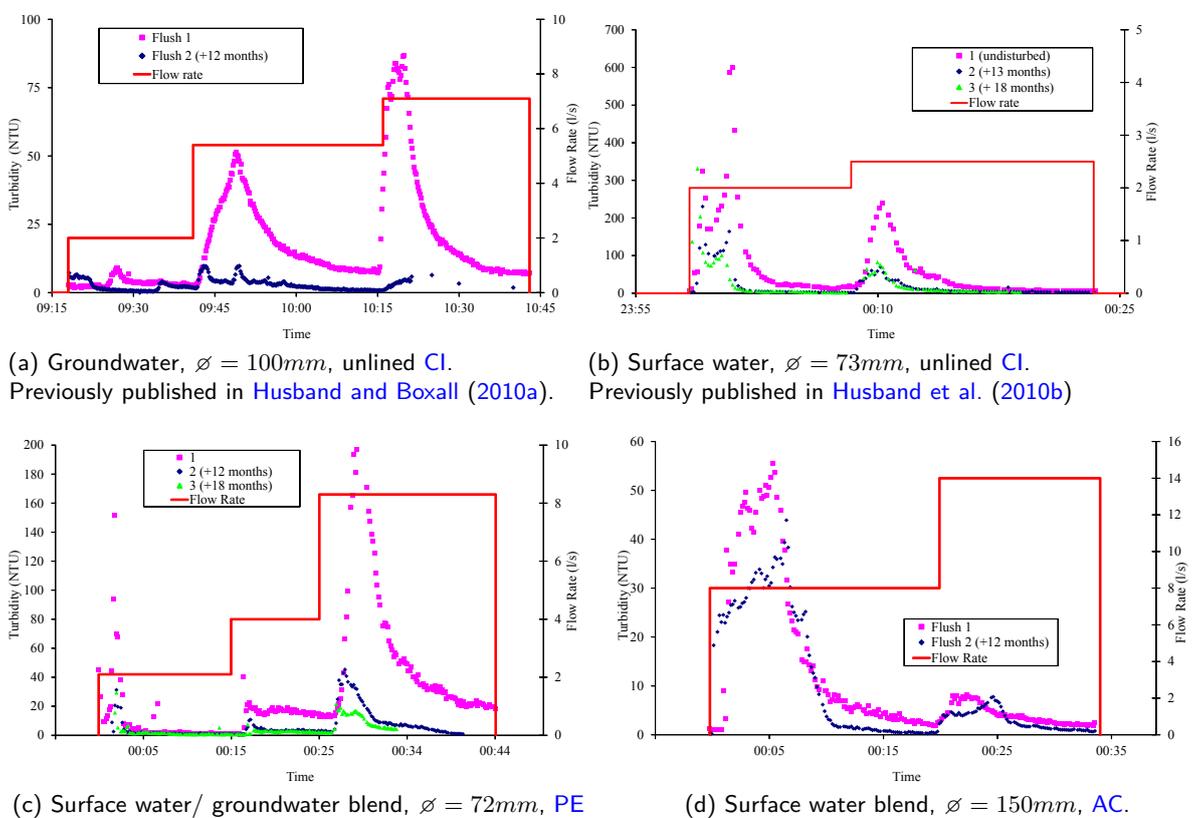


Figure 2.11: Turbidity data from the repeated flushing of DMA mains.

The findings of these field studies concur with those of laboratory experiments. Sharpe (2013) connected three parallel pipe loops (HDPE; 203m long; $\varnothing = 79.3\text{mm}$) to a tank with a 24h system residence time. The tank was fed with treated moorland run-off via an unlined CI trunk main. After being cleaned using high flushing flows and super-chlorination, the pipes were subjected to a 28-day material accumulation phase during which a different steady shear stress was imposed at the wall of each pipe. During a subsequent material mobilisation phase, the flow was increased in each pipe in turn, draining down the tank between the flushing of each pipe. The flow and therefore shear stress were increased in steps, each of which was sustained for three turnovers. Turbidity was monitored at the downstream ends of each pipe loop and was normalised by the background turbidity prior to the mobilisation phase. The entire experiment was conducted at first 8°C then 16°C . The recorded turbidity data (fig. 2.12) showed that

discolouration material with a variety of shear strengths had accumulated simultaneously over the 28 days, rather than the material at the wall being dominated by stronger-bound material.²

Material quantity has been shown to increase across the strength profile over time, rather than only very weak or quite strong material accumulating over periods of one to twelve months. Whether this accumulation is invariant with shear strength at values greater than above the shear stress has not yet been conclusively proven, although Cook (2007, p. 76) presented data from the flushing of two DMAs where the proportions of material that accumulated over three different strength ranges appeared to be constant between a number of pairs of flushes. It should be noted that this conclusion in Cook (2007) was the result of the visual interpretation of graphs rather than statistical analysis.

The possibilities of there being upper limits to the amount of material that can accumulate within a given strength range (§2.1.14) and the strength with which material can adhere to the wall (§2.1.15) are considered in the following sections.

2.1.14 The possibility of an upper bound to the accumulation over a shear strength range

If it is known that material accumulates at all shear strengths over a range simultaneously then one of the questions this gives rise to is whether the accumulation process can continue within this range near-indefinitely or whether it ceases or significantly slows in time long before the pipe becomes hydraulically impassible (fig. 2.13).

The question is difficult to answer conclusively in the field using techniques such as paired mains flushes (§2.1.13.2), principally due to it being difficult to conduct paired flushing operations under semi-controlled conditions that are separated by a sufficient duration to observe material accumulation reaching the levels prior to the first flush. The majority of paired flushes conducted to date in DMAs have been up to and including a year apart (Blokker et al., 2011; Cook and Boxall, 2011; Husband and Boxall, 2011) and 'full' accumulation has not been observed over such intervals; estimations of the duration required for 'full' accumulation are from linear extrapolation (1.52 to 4.17 years for DMA pipework in England and Wales (Husband and Boxall, 2011)). Repeatedly flushing mains where the individual flushes are separated by intervals of multiple years is made difficult by having to ensure that the shear strength range over which one is measuring accumulation is *only* eroded by the scheduled flushes themselves and not by flow increases in the several years before the first flush or between the flushes. In addition each site would need to be revisited multiple times to ascertain whether the accumulation rate at a given site only appears to decrease beyond a particular duration of accumulation period.

Laboratory rigs offer finer measurement and greater control of water quality and quantity during such studies but to date the longest flushing studies explored material accumulation over just twenty eight days (Sharpe, 2013; Fish, 2014), even shorter than the aforementioned field studies so less suited to detecting a reduction in accumulation rates with time.

As an alternative to paired flushing, continuous monitoring of turbidity and flow also has potential for studying temporal changes in discolouration processes. Husband et al. (2010b)

²This and the previous paragraph contain text that has been quoted from Furnass et al. (2014b).

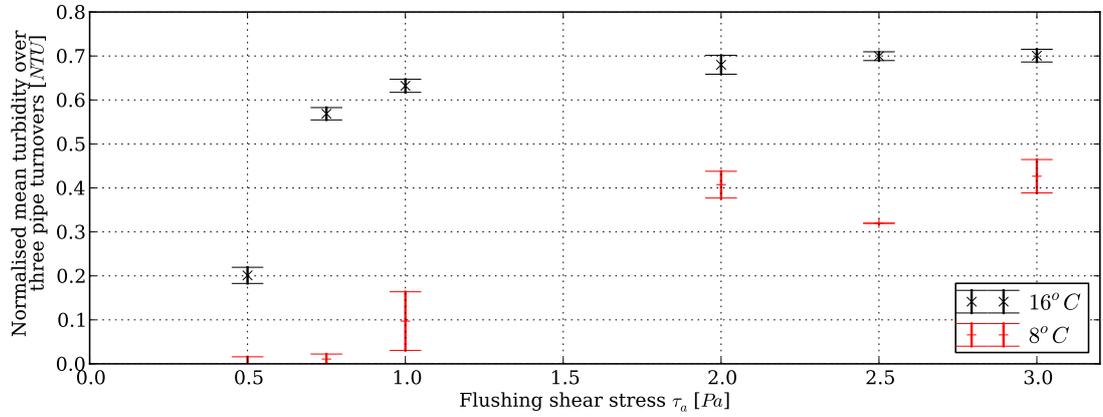


Figure 2.12: Results of laboratory studies into discolouration material accumulation (after [Sharpe, 2013](#)). Material layers were developed over 28 days under a steady shear stress of 0.1 Pa ($0.2 \text{ L} \cdot \text{s}^{-1}$) at 8°C and 16°C . These layers were then eroded through imposing a series of discrete increases in shear stress. The way in which turbidity data was normalised is described in the body of the text.

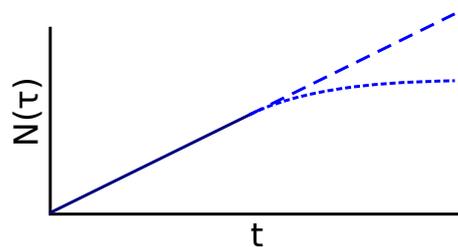


Figure 2.13: The amount of material N that accumulates on each unit area of pipe wall with strength τ could either increase linearly with time or could trend towards a maximal level, as illustrated in this diagram.

reported the influence of interventions on discolouration processes in a pair of parallel trunk mains ($\varnothing = 1.8m, 1.5m; L = 20km$) with a common inlet that had previously been operated at particular flows for several years. The first intervention involved increasing the flow in the $\varnothing = 1.8m$ main to deliberately erode material in a controlled fashion to allow the $\varnothing = 1.5m$ main to be taken out of service for jet washing. The wider was later also jet washed. Continuous monitoring at the common inlet and the downstream end of the narrower main showed that prior to any interventions the turbidity at the downstream end was approximately five times greater than the upstream end. This discrepancy was far less during the months following the brief, controlled flow increase and was less again following the jet washing (fig. 2.14). One possible explanation for these changes is that the material in the main had reached a maximal state prior to the interventions but then fresh material could accumulate following the flow conditioning and jet washing. Before the interventions, the particulate material forming within the main through processes such as precipitation, coagulation and flocculation may have been passing out at the downstream end due to it not being possible for more material to accumulate; such material was then able to bind to the wall after the interventions. Alternatively it could have been the case that the calibration slope of either the upstream or (more likely) the downstream turbidity instrument was incorrect and the interventions caused a reduction in the proportion of particulate material that passed out of the main having entered it at the upstream end. Unfortunately the downstream turbidity monitoring data is limited to the period shown in fig. 2.14 so it is not known whether the turbidity at the downstream end returned to its original level with time, which, if seen, could indicate that the material quantity at the pipe wall has reached a maximal level again.

Others have assessed the impact of flushing on regular turbidity fluctuations within DMA mains (Cook, 2007; Gaffney and Boulton, 2011; Mounce et al., 2014a). Gaffney and Boulton (2011) monitored pressure and turbidity at two locations within one DMA and four locations within another at fifteen-minute resolution for five days before and fourteen days after flushing. Through plotting the magnitude of turbidity events $> 0.5 NTU$ against their duration Gaffney and Boulton (2011) found that the resulting curves were similar for the first DMA but that the magnitude and percentage of total monitoring time of events detected in the other DMA actually increased after flushing. This contradicts the hypothesis stated previously that the turbidity baseline and fluctuations would decrease following mains cleaning interventions due to material being able to accumulate once again through flushing bringing the amount of material at the wall below maximal levels. However, when Mounce et al. (2014a) undertook a similar study but applied the CANARY event detection system (USEPA, 2010) to six months of historical turbidity and pressure data from four monitoring locations within a DMA they found flushing notably decreased the frequency of turbidity events. This could indicate that accumulation had reached a maximal level prior to flushing.

Mounce et al. (2014a) also studied how for the four monitoring locations turbidity correlated with (a) flow and (b) pressure at different temporal scales using a wavelet-based technique known as Semblance Analysis (Cooper and Cowan, 2008). The hypothesis was that for fouled mains turbidity would be positively correlated with flow and negatively correlated with pressure at the daily scale. This was partly demonstrated in the results, where the correlation between turbidity and the hydraulic signals 'reversed' following pipe flushing at two of the four sites, B and D. The other two sites, C and E, were thought to be subject to atypical conditions so would not reflect

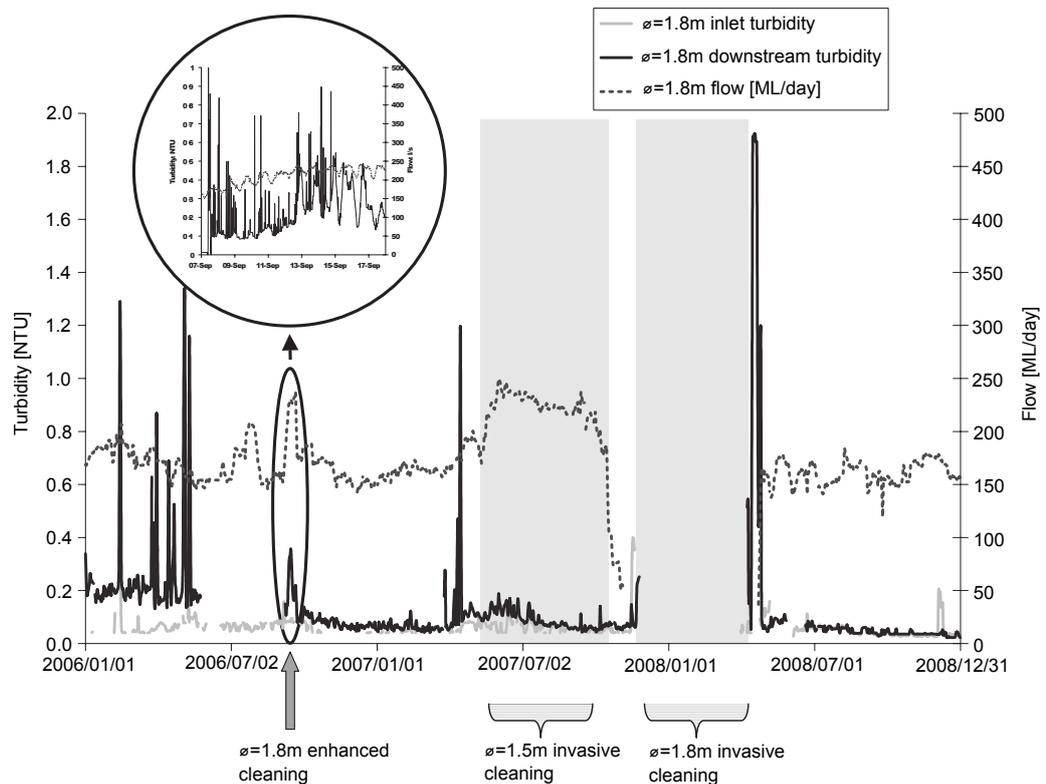


Figure 2.14: Daily-averaged turbidity and flow records from a UK trunk main over 2006-2008, after Husband et al. (2010b). Inset: 15-min flow and turbidity data during a period where flow was increased to forcibly remove discoloration material (the absence of downstream data from May-Sept 2006 was due to instrument recalibration).

this hypothesis. The change in correlation at B and D following flushing could be indicative of the material bound the pipe walls then no longer being at maximum capacity regarding discoloration material, allowing for material accumulation again. Mounce et al. (2014a) also assessed whether any change in the daily turbidity cycle following flushing interventions could be quantified using the daily standard deviation of turbidity at each monitoring location, possibly averaged over each week. The calculated values were typically higher for the weeks before flushing for two out of four sites.

There are also theoretical arguments for material accumulation within a strength range being or not being finite. The thickness of biofilms may be limited by the diffusion of nutrients through the EPS matrix. However, scale does not appear to be subject to such limitations: a significant portion of the pipe diameter can become occupied by scale/tubercles over time (although it is not known whether tubercles are erodible and can therefore be classed as discoloration material), and in Australia much capacity in raw water mains has been lost due to microbially-mediated iron deposition (Forward, 2004). Note that the conditions in raw water mains are not necessarily comparable with those in DWDS and no drinking water mains are known to have become as fouled due to discoloration material accumulation.

Another argument for the material that can accumulate with a given shear strength having an upper bound is that the turbidity response from flushing previously undisturbed mains is not infinite. However, it could be argued that accumulation in mains may be indefinite yet slow enough so that the turbidity responses seen during flushing following a decade or more of accumulation are not many orders of magnitude larger than what is seen after material has

been accumulating over a year.

The evidence and theoretical suggestions for and against material accumulation over a given strength range being bounded are mixed. It will be difficult if not impossible to demonstrate bounded or indefinite accumulation given the need for long term monitoring and control and the fact that the quantity of wall-bound discolouration material can presently only be easily measured by disturbing it. In reality, it may be that the accumulation process neither stops suddenly when the material at the wall reaches a certain quantity nor continues linearly until the pipe is impassible but instead slows after an accumulation period longer than any studied to date (see fig. 2.13). Also, the question may be irrelevant if shear stresses in DWDSs are sufficiently variable that material at the corresponding shear strength ranges is rarely permitted to accumulate for long enough for bounded and indefinite accumulation to be distinguishable.

2.1.15 The forces required to 'clean' a pipe

Knowing the minimum hydraulic force required to reduce discolouration risk to acceptable levels is important for the operational management of DWDSs. One method of quantifying this force has been to impose increasingly large velocities or shear stresses within pipes of different materials and/or diameters to see if there is a limit beyond which those pipes emit no further turbidity response. In the Netherlands such studies have resulted in the recommendation that fouled distribution mains be pro-actively flushed by opening hydrants and raising the velocity in the main to 1.5 m s^{-1} for three turnovers (Vreeburg and Boxall, 2007). In the UK, Miller (1994), cited by Vreeburg and Boxall (2007), proposed target flushing velocities that are a function of pipe diameter and specific gravity, in keeping with classical sediment transport theory.

In the Netherlands a change in the requirements for fire-fighting flows has permitted the development of 'self-cleaning' DWDSs (Vreeburg and Boxall, 2007). These are dendritic structures in which pipes get narrower towards the extremities of the network. Regular high velocities and uni-directional flows ensure sediments do not accumulate in distribution systems but are passed through to points of supply at acceptable concentrations (Blokker et al., 2010) i.e. accumulation can be limited by capitalising on the diurnal discolouration cycle (§2.1.13). The design maximum daily velocity is 0.4 m s^{-1} (Vreeburg and Boxall, 2007) but through field trials Blokker et al. (2010) found that peak velocities of $0.2 - 0.25 \text{ m s}^{-1}$ every two days may be sufficient to prevent fouling in AC and PVC pipes. Blokker et al. (2011) show how existing networks can be made self-cleaning through adjusting valve positions to increase daily peak velocities. However, for many DWDSs pipe downsizing, extensive valve reconfigurations and material replacement are not practical options for reducing discolouration risk and may impact on the ability to meet fire-fighting demands.

Increasing the bulk velocity to the Dutch recommended self-cleaning level on just one occasion is not necessarily sufficient to remove all material from the walls of plastic pipes (Husband et al., 2008, experiment series 2 (HDPE pipe)).

It may not be appropriate to quantify the cleaning force required to remove all material from the pipe wall as a velocity; reference to critical velocities implies that erosion is driven by classical sediment transport processes yet discolouration material has been shown to develop

around the pipe circumference in cohesive layers and be too light once entrained to easily settle as sediment (§2.1.9). Also, any maximum shear strength of discolouration material cannot be deduced from the aforementioned target cleaning velocities as shear stress is a function of the Darcy Weisbach friction factor (eq. 2.9) and therefore of pipe roughness and diameter (eq. 2.11) as well as bulk velocity).

Whether there is a maximum shear strength of wall-bound discolouration matter appears to differ between pipe materials and the value of this threshold appears to vary for pipe materials where a cleaning shear stress has been identified:

- During the flushing of two English DMAs Cook and Boxall (2011) found the maximum shear strength of plastic-bound material to be $\leq 0.7 Pa$, yet no cleaning shear stress was found for the AC and cement-lined Fe pipes in the vicinity: the shear stress was increased to $8 Pa$ in steps and more material was mobilised at each level.
- In a HDPE laboratory pipe system in which the hydraulics and water quality were largely representative of a real DWDS, pipe turbidity responses were observed during one set of experiments for increases in shear stress up to and including $5.7 Pa$ (following a material growth period) (Husband et al., 2008). In another set that used similar equipment, no cleaning shear stress³ was detected ($> 1.97 Pa$) in certain experiments but in others little material was removed when the shear stress was increased above $0.43 Pa$ (Sharpe, 2013). There is presently little understanding of how the cleaning shear stress in smooth pipes such as those made of HDPE can be predicted from asset characteristics, bulk water characteristics and system history.
- At the maximum shear stress³ imposed by Sharpe (2013) ($2.45 Pa$), biofilms, known to correlate with discolouration potential (Sharpe, 2013), were only partially eroded as the majority of all proteins but a much smaller proportion of the cells and carbohydrates in the biofilm had been removed. There may therefore be a (site-specific) upper limit to the quantity of discolouration material in smooth-walled pipes but parts of the biofilm are very difficult to erode and may still provide an adhesive coating to the internal surfaces of DWDSs even after aggressive flushing.
- In rough-walled, unlined English CI pipes there does not appear to be a maximum shear strength: during DMA flushing operations, Husband and Boxall (2010a) did not find a shear stress above which no more material was mobilised from unlined CI. Husband and Boxall (2010a) suggested that this could indicate that the corrosion process is not shear-stress-limited but the other material accumulation processes that affect ferrous and non-ferrous pipes are. However, this appears not to be true as no evidence of maximum shear strengths has been found in non-corrodible AC and lined Fe pipes.
- Also, some English unlined CI pipes were found not to be self-cleaning even though the daily maximum bulk velocity was approximately equal to the Dutch recommended self-cleaning velocity of $0.4 ms^{-1}$ (Husband and Boxall, 2011). However this recommended self-cleaning velocity may not be applicable for unlined CI pipes due to it having been

³These shear stresses are not as quoted by Sharpe (2013) and Fish (2014); they were recalculated using eq. 2.9, the quoted diameter ($\varnothing = 79.3mm$), roughness ($k_s = 0.075mm$), flow and the Swamee-Jain approximation to the Colebrook-White formula.

developed in a country where corrodible mains, with their higher roughnesses and therefore shear stresses for the same diameter, make up a very small proportion of DWDS pipework (Vreeburg, 2007).

In the few studies of discolouration processes in large diameter ($\varnothing \geq 200mm$) trunk (transmission) mains (Vreeburg, 2007; Vreeburg and Beverloo, 2011; Seth et al., 2009; Husband et al., 2013, 2011, 2010a,b; Saldarriaga et al., 2010) no evidence has been found for or against there being a cleaning shear stress.

An understanding of the stresses required to remove biofilm from pipe walls is important as, following the erosion of much discolouration material, the remaining sticky EPS could expedite material accumulation. Abe et al. (2012) imposed mechanical shear stresses on laboratory-grown biofilms using atomic force microscopy and found that a shear stress of $2.5 \times 10^5 Pa$ was required for the removal of almost all biofilm, a stress that is far higher than those required to reduce discolouration risk to negligible levels in plastic pipes and is very unlikely to ever be imposed within DWDSs.

There is no published evidence of discolouration material having a minimum strength.

In summary, the current discolouration risk can be minimised in plastic pipes by imposing a shear stress greater than a threshold; however, there is disagreement as to the value of this threshold, which may be due to it being partly site-specific. However, an understanding of the range of values that this threshold could fall within for a given pipe can help bound discolouration risk. There is mixed evidence for there being similar thresholds in AC and concrete-lined Fe pipes and none for unlined CI mains, possibly due to material accumulation being partly due to corrosion in unlined CI pipes. Further studies where much higher shear stresses are imposed may help identify discolouration risk shear strength thresholds for a greater variety of pipes. However, biofilm studies have shown that flushing can minimise discolouration risk at a given moment in time that some cohesive biofilm may remain strongly adhered; this material has the potential to promptly restart the discolouration cycle.

It should be noted that in some systems the system configuration (pump, reservoirs and pipes) may not allow the flow and therefore the shear stress to be raised to the levels needed to clean plastic pipes.

2.1.16 Factors that could influence accumulation rates

2.1.16.1 Introduction

System hydraulics are the primary driver of the *maximum* discolouration potential (the quantity of material adhered to the wall at any given time) whereas water quality and asset attributes such as pipe material appear to control accumulation *rates* (Cook and Boxall, 2011; Husband and Boxall, 2011, 2007; Lehtola et al., 2006); for this reason it is sensible to discuss accumulation in relative rather than absolute terms when investigating accumulation rates and the factors that influence them.

Husband and Boxall (2011) presented expectations and measures of variance for the rates of accumulation in various environments (e.g. Fe vs Al vs no coagulation; surface vs ground

vs blended water source; corrodible vs non-corrodible mains). Basic bi-variate categorical breakdowns of discolouration rates are presented by Husband and Boxall (2007, §4.2). However, given the complex, interrelated nature of the physical, chemical and biological reactions that are considered to contribute to discolouration material accumulation, it is thought that predictive models of accumulation rates require through the application of *multi-variate*, regressive methods to sufficient volumes of representative data. Work has already begun in this area: Mounce et al. (2014c) used DMA flushing data from Husband and Boxall (2011) to demonstrate how multivariate data analysis techniques can be used to identify and explore the relationships between accumulation rate predictor variables (using a non-linear dimensionality reduction technique known as Kohonen Self-Organising Maps) and also identify formulae of varying complexity for accumulation rates using subsets of potential predictor variables (using evolutionary polynomial regression (EPR)). Opher and Ostfeld (2011) also used multi-variate data analysis techniques to investigate a similar problem: *model trees* (Quinlan, 1992), optimised using genetic algorithms (Holland, 1992), were used to learn pipeline 'biofouling rates' from a large number of predictor variables.

The following factors have the potential to influence discolouration material accumulation rates.

2.1.16.2 Temperature

The temperature of the bulk water and pipe wall influences all biochemical reactions within a distribution system. Research indicates that microbial growth increases significantly when temperatures exceed 16°C (Dukan et al., 1996). The maxim that reaction rates double for every 10°C increase in temperature (Connors, 1990) appears valid for biofilm growth in distribution systems between 7°C and 20°C (Hallam et al., 2001). Increases in temperature also have the capacity to expedite the corrosion of pipes, fittings and fixtures through influencing the solubility of oxygen, viscosity, ferrous iron oxidation rates and microbial-mitigated corrosion processes (McNeill, 2000). Disinfection residuals serve as biocides but can also expedite corrosion; the bulk decay constant of chlorine has been found to vary significantly with temperature (Powell et al., 2000); this could lead to deficient residuals at network extremities.

Temporal variation in water temperature is significant in distribution systems supplied with surface water ($3.9^{\circ}\text{C} - 32.2^{\circ}\text{C}$ over a year for a US DWDS supplied from a river (Potter, 2000); $3.2^{\circ}\text{C} - 20.2^{\circ}\text{C}$ over a year for a UK DWDS supplied with surface water (see Fig. 2.15) but much less so for groundwater-fed systems.

The temperature of advected water can also vary as it moves from source to tap due to heat transfer through pipe walls to/from backfill and the surrounding soil. Blokker and Pieterse-Quirijns (2013) developed and validated a model of heat transfer from the subsurface to advected water in DWDS and found that residence time had a greater effect on water temperature than atmospheric temperature or the temperature of the source.

Several studies have explored how temperature influences discolouration material accumulation rates. A strong positive relationship between bulk water temperature and regeneration rates was identified by Schaap and Blokker (2013) when a Dutch distribution main was flushed thirty times in eighteen months. Sharpe (2013) studied the impact of temperature (comparing 8°C and 16°C) and prevailing shear stress on accumulation rates in a realistic-scale HDPE

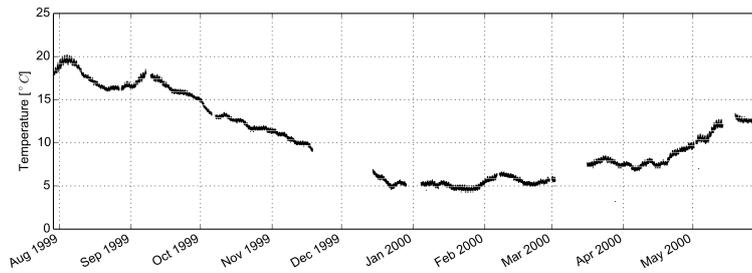


Figure 2.15: Temperature variation in a surface-water-supplied DMA in the UK over the course of one year (previously unpublished).

pipe rig over 28-days and found that accumulation, as measured by turbidity and Fe and Mn concentrations during flushing, was most greatly influenced by temperature. She also noted that temperature affected the strength profile of the wall-bound material: turbidity data recorded during stepped flushing showed that the majority of material grown at 16°C was eroded during the first few step increases in shear stress but that material grown at 8°C was distributed over a wider range of strengths (fig. 2.12). It should be noted that it is not known if/how these results would differ for non-plastic pipes, for different diameters and for different source waters.

2.1.16.3 Characteristics of raw water and water treatment works (WTW) finals

Water from upland sources is typically much richer in natural organic matter (NOM) than groundwater and the concentration of organic substances is seasonal as it reflects the degree of run-off. Full layer regeneration (relative accumulation of 100%) in the UK DMAs is typically 1.5 years for mains supplied with surface water versus 3 years for ground waters (Husband and Boxall, 2011). To remove the NOM that gives water its natural colour it is common to dose at treatment works with Fe or Al-bearing coagulants (Sharp et al., 2006). Coagulant demand is a function of NOM concentration, composition and character (Sharp et al., 2006); failure to appropriately tune the coagulation, flocculation and filtration processes at treatment stage could lead to NOM or coagulant material passing into the distribution system (Cook, 2007, p. 124). In addition, hydrophilic NOM fractions are much less easily removed by coagulation than hydrophobic substances (Sharp et al., 2006). Husband and Boxall (2011) found that more discolouration material collected in surface water-fed systems than ground-water supplied networks and accumulation rates were greatest for systems in which Fe coagulants were used, less for Al-coagulated networks and less again for non-coagulated systems.

Other aspects of the chemical composition of WTW finals could influence accumulation rates: a decrease in pH could expedite corrosion; metals could precipitate as slimes on pipe walls; the nutrients nitrate (difficult to remove at WTW) and phosphate (often added WTW to limit cuprosolvency and plumbosolvency) could affect biofilm growth rates, as could dissolved CO₂ and oxygen levels. Hard waters could provide calcite linings that have a positive impact on corrosion rates (McNeill and Edwards, 2001) but provide the key nutrient Ca.

The concentration of effective disinfectant residual could also impact on accumulation rates. Both oxygen- and chlorine-bearing compounds are capable of oxidising dissolved Fe and Mn to more soluble forms (Sly et al., 1990).

The particle loading downstream of the [WTW](#) should also be considered. [Vreeburg et al. \(2008\)](#) studied material accumulation in a [DWDS](#) containing only non-corrodible pipes ([AC](#) and [PVC](#)) and found that $0.1\mu m$ ultra-filtration reduced material regeneration times from 1-1.5 years to 10-15 years. Interestingly, the particulate matter that collected downstream of the ultra-filtration unit was primarily organic whereas the material in mains not supplied by the unit was mostly inorganic. The nature of [WTW](#) filter backwashing could also influence accumulation rates ([Vreeburg et al., 2008](#)).

2.1.16.4 Asset characteristics

Material layers accumulate at on average 52% per annum in corrodible, unlined Fe [DMA](#) pipes (calculated using the method presented in §2.1.13.2) but only 28% per annum in other [DMA](#) mains ([Husband and Boxall, 2011](#)). This is most likely due to plastic pipes only being supplied with discolouration material from the bulk water whereas Fe pipes also collect particulate matter from corrosion processes ([Husband and Boxall, 2010a](#), p. 92). No [DWDS](#) pipe materials are immune to the accumulation of discolouration materials ([Husband and Boxall, 2011](#)). Other material-related properties that have the potential to influence accumulation rates are: whether Fe pipes are cast, ductile or galvanised; whether any protective scale in Fe pipes has recently been removed through activities such as pigging ([Vreeburg et al., 2004](#)); the type and condition of any pipe lining; the age of the main, which correlates with hydraulic roughness and the likelihood of the pipe containing 'legacy' materials; the type of plastic as new ([Lehtola et al., 2004a](#)) and older ([Lehtola et al., 2006](#)) [PE](#) mains can potentially release phosphorus, a key nutrient, to the bulk water.

2.1.16.5 Pipe hydraulics

Flow magnitude Under steady state hydraulic conditions (and a constant temperature of $8^{\circ}C$) the magnitude of the invariable shear stress has been shown to be inversely related to the amount of discolouration material at the ([HDPE](#)) pipe wall ([Sharpe, 2013](#)). This is unsurprising given that weaker material can bind and remain bound under lower prevailing shear stresses. [Stoodley et al. \(2001a\)](#) showed that biofilms grown under greater shear stresses have greater shear strengths i.e. are conditioned by hydraulic forces.

Greater flow rates may, in addition to imposing greater shear stresses, also provide a greater supply of suspended particulate discolouration material and biofilm precursors and nutrients.

Discolouration material accumulates within the laminar sublayer at the pipe wall ([Pothof and Blokker, 2012](#)); [Sly et al. \(1990\)](#) hypothesise that this offers the material some protection from erosion. This suggests that the thickness of the laminar sublayer might bound material accumulation under constant flow conditions and an increase in flow may cause erosion through material then projecting out beyond the extent of the laminar sublayer.

Flow variability [Sharpe \(2013\)](#) also studied the influence of steady state and two variable diurnal flow patterns (imitation domestic demand profiles) on material accumulation. All three patterns had the same daily mean shear stress but different standard deviations and were

imposed in near-identical realistic-scale laboratory HDPE pipe systems over twenty eight days at 8°C. The magnitude of the daily peak stress was found to limit discolouration risk far more than the daily mean shear stress or the minimum (night-time) shear stress, complementing the objectives of the Dutch self-cleaning network design strategy (see §2.1.15). In addition, biofilms grown under the variable flow conditioning patterns were found to be less susceptible to erosion from flushing than biofilms grown under a steady state flow with the same daily mean, corroborating Stoodley et al. (2001a)'s aforementioned findings.

Diameter-influenced mixing There may be less mixing in large-diameter, low-velocity trunk systems than DMAs. This may result in a lower take-up of discolouration material and nutrients to the pipe wall. Dispersion effects and incomplete mixing reactions are rarely modelled in potable water networks (Andrade et al., 2010); under the turbulent conditions typically found in DWDS there is negligible dispersion (Tzatchkov et al., 2009) so it is usually appropriate to assume plug flow (no dispersion) and complete, instantaneous transverse mixing in pipes and complete, instantaneous mixing at junctions.

Flow reversals Frequent flow reversals in pipe loops could influence accumulation rates as material may settle out at a 'tidal point' due to low bulk velocities (Vreburg and Boxall, 2007).

2.1.16.6 Network topologies

Discolouration materials appear to develop along the length of pipes rather than only in certain areas (Boxall et al., 2003b), but in a single length of trunk main several kilometres long with a constant pipe diameter, age and material it is possible that concentration (disinfectant, nutrient and metals) and temperature gradients may cause variation in material accumulation rates along its length. Oxidant and disinfectant concentrations have been shown to influence the spacial variability of Mn deposition in DWDS (Sly et al., 1990). Verberk et al. (2009) reported that PaVoC (see §2.1.8), TSS and volatile suspended solids (VSS) concentrations increased as water moved from a WTW past two downstream monitoring points within the supplied DWDS. A stream-wise increase in Fe concentrations in the absence of any Fe pipework was indicative of iron oxidation, precipitation and flocculation. Vreburg and Beverloo (2011) noted that the ratio of Fe to Mn varied over the lengths of two trunk mains, which may indicate that regeneration rates will differ along those mains. Increases in the concentrations of other inorganics such as Si are suggestive of chemical interactions between the bulk water and the surrounding AC mains. Discolouration modelling work by Husband et al. (2011) indicated that more material had accumulated at the upstream end of a flushed English trunk main. Similarly, Vreburg and Beverloo (2011) noted that more material appeared to have collected at the upstream end of a flushed Dutch trunk main (although, if erosion is not instantaneous, it is not always possible to say with confidence where the mobilised material that emerges from a pipe originated from).

Discolouration material can collect around fixtures and fittings. Husband et al. (2010b) and Neilands et al. (2012) related turbidity spikes seen during flushing to features such as right-angled bends and sumps using estimates of advection times. However, this approach is not

always accurate (Schaap and Blokker, 2013) and requires the assumption that all material affected by a shear stress increase is mobilised very quickly (to allow the response at the downstream end of the pipe at a given time to be associated with a specific location upstream). The assets upstream of a pipe of interest could influence accumulation rates, in particular whether there is a Fe source upstream (Husband and Boxall, 2007). In discussions with the author, a UK water provider representative noted that, for one particular network, a particular measure of discolouration risk was much greater for plastic pipes that are downstream rather than upstream of unlined Fe pipes. This is presumably due to corrosion of the Fe pipes providing a supply of iron oxide particles.

The presence of service reservoirs upstream of a pipe of interest is suspected to affect accumulation rates (Husband et al., 2010b). Reservoirs could act as buffers of discolouration material and could provide environments within which dissolved material precipitates and particles flocculate and settle. Temperature variations in reservoirs could influence biofilm growth rates. Also, service reservoirs are exposed to the atmosphere so may experience thermal stratification and temperature fluctuations throughout the year; these in turn may influence the temperature-dependent biological and chemical reactions associated with material accumulation (e.g. biofilm growth; Sharpe, 2013).

2.1.16.7 Water age

Water age has been shown to be a useful proxy metric for water quality as it can be linked to microbial growth rates, residual disinfection concentrations, bulk water temperature and phosphate concentrations (Machell et al., 2009; Machell and Boxall, 2012). Mounce et al. (2014b) have correlated stagnation (old water) in DWDS with high Fe concentrations. Water age is of course a function of the DWDS topology (§2.1.16.6) and hydraulics (§2.1.16.5).

2.1.16.8 Combined effects

The processes that contribute to accumulation rates may not be entirely separable, such as in the cases of microbially-expedited Mn deposition (Sly et al., 1990) and Fe corrosion (McNeill, 2000) and of the adhesion of oxide flocs to EPS.

2.1.16.9 Season

Many of the factors that may influence accumulation rates such as hydraulics (demand), temperature, the quality of surface water supplies (O'Brien et al., 2008) and coagulant doses are very likely to exhibit seasonal trends. It is anticipated that accumulation rates will reflect these trends to some extent.

2.1.17 Discolouration risk management

2.1.17.1 Quantifying the current discolouration risk

As the accumulation of discolouration material is typically a continual and ubiquitous process (§2.1.13.3), there is risk of discolouration in most DWDSs. This risk is a function of the strength

versus quantity profile of wall-adhered material (§2.1.12.2; §2.1.13.4; §2.1.15; §2.1.16.5) and of the likelihood of an unanticipated change in hydraulics due to a burst, the uncontrolled use of a hydrant or sudden network reconfiguration (Randall-Smith et al., 2011). The discolouration potential will vary considerably between systems as many factors influence the discolouration material strength versus quantity relationship and how this changes over time (§2.1.15; §2.1.16). If risk is not pro-actively quantified and managed then both controlled and unexpected increases in system hydraulics have the potential to cause unacceptably turbid water to pass through to consumers, which may result in complaints and regulatory penalties.

At present, discolouration potential is commonly assessed using flushing in one or more representative pipes. A stepped flushing flow profile permits the study of the strength versus quantity relationship (§2.1.13.4) whereas to opt for a constant flushing flow profile instead yields worst case maximum turbidities. Flushing regimes can be designed and the results interpreted using the RPM method (§2.1.13.2) so as to condense complex turbidity data down to simple risk scores for use in prioritising and evaluating interventions, although the method's use of velocity thresholds may not allow risk to be accurately compared between pipes with the same flushing bulk velocity but different flushing shear stresses as discolouration material is held at the wall in cohesive layers (§2.1.12).

In DWDSs and pipes with high accumulation rates, water companies may want to measure the current discolouration potential every few months. However, although customers are typically informed before flushing is undertaken, valving operations and the temporary supply of turbid water may cause inconvenience for customers. Also, frequent flushing may be undesirable because of the volumes of water wasted, the risk of icing up impermeable surfaces and in larger-diameter mains the costs/efforts associated with substantially elevating the shear stress. Thirdly, and possibly most importantly, flushing is a point-in-time measure of discolouration risk: flushing with a constant shear stress of τ_a at time t yields a discolouration response with a peak turbidity of T and removes all material with a shear strength of less than $\tau_a(t)$; however, the risk will then change with time due to material erosion and accumulation effects.

In addition to flushing, water companies may also reactively estimate discolouration risk using customer contacts. Customers are well distributed throughout DWDS and are frequently sampling water (Whelton et al., 2007). However, regulators may not view a reliance on customers as sensors particularly favourably and contacts regarding discolouration only provide very coarse information on a) the pipes from which discolouration material was eroded, b) the onset time and c) the rate of increase, the maximum and rate of decrease of turbidity.

2.1.17.2 Interventions for risk reduction

Asset replacement Risk assessment programmes will often identify unlined Fe as being associated with the greatest risk. Replacement with plastic mains will eliminate internal corrosion as a discolouration material source, plus accumulation is slower in plastic pipes and the discolouration potential is far less than in unlined Fe mains of the same geometry (Husband and Boxall, 2011). However, the replacement of old Fe mains is expensive, disruptive and according to the water industry regulator for England and Wales “unlikely to be economically justifiable on the basis of a sole, aesthetic driver such as discolouration” (OFWAT (2000), cited

by [Boxall et al. \(2003b\)](#)). In addition, the like-for-like replacement of non-corrodible mains will not reduce discolouration risk.

Asset refurbishment Interventions such as the scraping, lining or pressure washing of Fe mains ([Ratnayaka et al., 2009](#), ch. 14) may bring about a temporary reduction in discolouration risk but the associated network activities may cause disturbances that increase the discolouration risk elsewhere in the network [Husband et al. \(2010b\)](#). Scale and tubercles can be physically removed from mains by pigging, which increases hydraulic capacity and reduces friction losses but may also expedite corrosion due to the exposure of unoxidised iron ([Vreeburg and Boxall, 2007](#)). As with pipe replacement, relining with plastic, resin or cement may reduce but will not entirely eliminate discolouration potential.

Quality of WTW finals Tuning processes such as coagulant dosing, filter backwashing ([Vreeburg et al., 2008](#)), phosphate dosing or pH control at the WTW could potentially decrease accumulation rates but such measures are unlikely to reduce the current discolouration potential.

Mains flushing If numerous customer contacts or the flushing of a DMA main at a low flow indicate a notable discolouration risk and it is anticipated that the main may be subjected to high flows from a burst or network reconfiguration, it can then be flushed at the maximum of those anticipated flows to remove all material that is weaker than the corresponding shear stress. However, as mentioned earlier in this section this may inconvenience customers and risk will increase following the intervention due to material accumulation.

An alternative to flushing to waste is what shall henceforth be referred to as **flow conditioning**, where flows are periodically raised to erode discolouration material in a controlled fashion. This material then passes out of the system at supply points in acceptable concentrations. These ideas are central to the Dutch idea of self-cleaning distribution pipes (§2.1.15) where pipe diameters have been chosen to allow natural variations in demand to remove most material from smooth, non-Fe pipes. In the last decade an interest has developed in whether flow conditioning could be used to manage risk in mains of much larger diameter (e.g. [Husband et al., 2011](#)), as is discussed in the following section. Flow conditioning is in keeping with regulators' desire to bring about a better balance between capital and operational expenditure (e.g. [DWI, 2002](#)).

Trunk mains: systems of concern *Trunk* or *transmission* mains, usually $\varnothing > 200\text{ mm}$, convey water from WTW to distribution networks of narrower diameter pipes, typically via pumping stations and/or service reservoirs. These critical assets are of concern to water providers with regards to discolouration risk as they lie upstream of large numbers of customers. They may also only experience low conditioning shear stresses, meaning there may be much weakly-bound material adhered to their large internal surfaces; the low shear stresses are due to their large diameters and in many cases to DWDS buffering (such as balancing tanks and water towers) further downstream limiting the impact of diurnal domestic demand variation on pipe flow. However, in some parts of the world, such as the Netherlands, trunk mains may

typically experience moderate to high shear stresses regularly due to diurnal demand and there being little if any buffering. There are additional reasons for *certain* trunks being of less a concern than local distribution pipework: (a) hydraulic constraints may not allow the shear stress to be increased much beyond prevailing levels and (b) eroded material is passing into a larger volume of water, resulting in smaller increases in turbidity.

There has been relatively little research published regarding discolouration processes in trunk mains and the extent to which discolouration risk in trunks and in the [DMAs](#) they supply are related ([Seth et al., 2009](#); [Husband et al., 2010b,a](#); [Saldarriaga et al., 2010](#); [Vreeburg and Beverloo, 2011](#); [Husband et al., 2011, 2013](#)). The limited understanding of discolouration processes in trunks is in part because these processes can only be observed through disturbing the system, which in trunk mains is difficult for several reasons:

- Very long turnover times: it may not be possible to operate mains at higher flows for prolonged periods without draining down upstream/downstream tanks;
- Experimental designs may require flows that are not easily achievable without network reconfiguration and/or upping pumping levels;
- It is only occasionally possible to flush trunk mains to waste (e.g. [Husband et al., 2010a](#)) due to the volumes of water involved and access points being few and far between. More commonly, trunk mains are flow conditioned; it may therefore be necessary to discharge flushing flow into just one compartment of a downstream reservoir in case regulatory turbidity limits are exceeded. This compartment could then be drained to waste whilst the other is used to ensure continuity of supply, although such wastage would not be desirable.

The difficulty in increasing flows in trunk mains regularly impacts on the ability to characterise material accumulation processes. Also, the few available trunk flow conditioning datasets all relate to systems sufficiently disparate that it is difficult to generalise what is occurring within them. On the other hand, it is in some ways easier to study discolouration processes in trunk mains than [DMA](#) pipework as trunk systems have simpler topologies and in some senses can be much better characterised i.e. hydraulic roughnesses are typically known at pipe level). Also, several trunks now feature far more accurate and reliable turbidity and flow monitoring than would be found in distribution systems.

The impact of accumulation in trunk mains can be estimated by studying if/how increases in customer contacts and/or background turbidity levels are seen over a large area at approximately the same time. [Husband et al. \(2010b\)](#) found that 41% of all customer contacts regarding discolouration could be assigned to clusters spanning multiple [DMAs](#); material erosion in trunk mains was considered the most likely cause. In another study, turbidity was monitored at hydrants in five [DMAs](#) over a multi-month period ([Cook, 2007](#)). 50% of the discolouration incidents that occurred during that time appeared to have been caused by events upstream of the [DMA](#) boundaries (in trunks) as the incidents did not correspond with increases in flow at the [DMA](#) inlets and the incidents were detected at multiple locations within the [DMAs](#).

2.1.17.3 The use of flow conditioning in trunk mains

[Husband et al. \(2011\)](#) reported on how flow conditioning could be used to regularly mobilise discolouration material within a particular trunk in a controlled manner so that it passes through to points of supply at acceptably low levels. This required accurate, [WTW](#)-style permanent turbidity monitoring and fine-grained flow control (an actuated flow control valve ([FCV](#))), to control the turbidity response at highly elevated flows where layer-eroding forces are most sensitive to flow. Investing in such instrumentation may prove to be greatly worthwhile: [Husband et al. \(2011\)](#) found that the cost of instigating a flow conditioning strategy was of orders of magnitude less than asset replacement or refurbishment.

At present the use of flow conditioning in trunks is limited by the system constraints listed in [§2.1.17.2](#) but also by a lack of understanding: of how discolouration risk develops over time; of how frequently to raise flows above prevailing levels; of the necessary magnitude of the flow increases for conditioning; and of the resulting turbidity response. Studies and models that help improve on this understanding would greatly aid the uptake and execution of flow conditioning for trunks.

2.1.17.4 Whole-life costing of flow conditioning

Improved understanding along with predictive capabilities would also allow the whole-life costs of this management strategy to be much better quantified. The capital costs associated with the installation of the accurate flow and turbidity metering required for flow conditioning can easily be ascertained but the long-term costs associated with the maintenance strategy are not yet known.

Whole-Life Costing, also known as Life Cycle Costing (LCC) and Total Cost of Ownership (TCO), is a means for estimating the expenditure associated with the whole life of an asset, not just the cost of initiation ([Skipworth et al., 2002a](#)): the capital expenditure associated with initialising an asset is often only a small proportion of the costs accrued over the life of that asset ([Read and Vickridge, 1997](#)). [Engelhardt et al. \(2002\)](#) present a detailed method for calculating whole-life costs and exploring the effects of undertaking different intervention strategies.

There is a need to consider a number of factors when estimating the whole-life cost associated with a high-discolouration-risk water main:

- Pumping (installation, operation and maintenance);
- Fines for exceeding regulatory water quality limits;
- Reservoir cleaning;
- Flow and turbidity monitoring (installation, operation and maintenance) inc. [SCADA](#) links;
- Treatment of wasted water (if any).

For a given main, the quantification of the above aids with the selection of the most appropriate maintenance strategy, the development of asset management plans and in part determines the cost to the consumer of supplied water. Other non-fiscal costs that may require consideration:

- The damage to a water provider's reputation if discoloured water is supplied to customers;
- The carbon cost associated with pumping at elevated levels and of treating any water that flows to waste;
- The volumes of water that flow to waste.

2.1.18 Discolouration processes: a summary

Over time particulate material accumulates on the internal surfaces of DWDS in cohesive layers. The rate of accumulation on pipe walls is driven by several processes including corrosion, adhesion and microbial growth and is an unknown, complex function of many parameters. The material bound to the pipe wall is conditioned by the prevailing shear stress and is mobilised when the shear stress is elevated. A drop in shear stress allows material to accumulate on the pipe wall over the corresponding strength range.

In smooth DMA pipes, frequent fluctuations in demand can prevent a significant amount of material from accumulating but in rougher, corrodible mains the variability of the shear stress is insufficient to keep pipes 'clean' and discolouration risk due to material accumulation can increase over time.

Unless this risk is pro-actively managed, increases in flow resulting from structural failures, scheduled flow increases and network reconfiguration could cause notable material mobilisation and the supply of water that is unacceptably turbid and has high metals concentrations. Operational flexibility is paramount for the management of trunk systems: water companies need to be able to conduct scheduled maintenance of WTW, pumping stations and service reservoirs but also react promptly when such assets fail, all the while striving not to negatively impact on water quality through generating discolouration. It is therefore important that discolouration processes are well characterised in trunk mains to aid with the latter's management but efforts to this end are impeded by the heterogeneity of such systems and by the practical requirements of trunk-specific discolouration studies.

Discolouration risk is difficult to measure directly due to it being associated with buried infrastructure and because the system cannot easily be characterised without disturbing it (by measuring the turbidity response from a controlled flow increase and assuming turbidity \propto TSS). Water companies need a means for assessing the current discolouration risk, not as complex, noisy flow and turbidity time-series but as a simple, comparable metric (such as a shear-stress-driven version of the RPM) but also need a way of assessing how discolouration risk and turbidity would change in the short, medium and long term under predicted future hydraulic conditions and also under alternative scenarios.

2.2 Modelling environmental systems

2.2.1 Introduction

Modelling provides a way of encoding the system state information and transaction rules in an abstract form that allows for extrapolation. For this reason the following section covers the present state of discolouration modelling including the degree to which current approaches are capable of reflecting reality and meet the needs of industry such as supporting the design of flow conditioning programmes. Strategies for developing environmental models and for modelling DWDS hydraulics and water quality in general are also discussed.

There are numerous reasons for wanting to abstractly characterise complex systems using models (Epstein, 2008); these include but are certainly not limited to:

- The ability to extrapolate the current system state in space and/or time and explore alternative scenarios;
- Being able to plan and make decisions based on evidence rather than just intuition;
- Gaining an understanding of core dynamics;
- Refining data requirements and data collection processes;
- Testing and challenging current theories and prevailing wisdom;
- Exploring trade-offs and efficiencies.

“Essentially, all models are wrong, but some are useful” (Box, 1987): no model fully represents reality nor should it. A 1:1-scale map would be unwieldy and offer few insights over observing the system/world directly (Borges, 2000). What is required is the simplest representation of reality that provides the required insight, Ockham’s Razor being a particularly important philosophy in environmental science where all intricate, inter-dependent physical, chemical and biological interactions and spatio-temporal heterogeneity cannot and most likely should not be comprehensively represented.

This section begins by exploring the various types of environmental models and means for developing them before moving on to consider how models are fitted to data and how quality of fit can be assessed. Some coverage is given to how features can be extracted from complex, noisy time-series datasets and also to software tools for developing and testing models. The second part of this section considers current methods for modelling spatio-temporal variations in water quality within DWDSs and how discolouration risk models have been developed by building on these ideas. The workings, merits and deficiencies of these models are then discussed.

2.2.2 Types of environmental models

Approaches to environmental modelling lie in a spectrum between entirely mechanistic, knowledge/theory-driven process modelling (*white box*) and fully empirical data-driven modelling (*black box*).

Process modelling is a 'bottom-up' approach in which formulations are derived from first principles i.e. from one's understanding of the underlying physics, chemistry, microbiology and/or economics (in the case of environmental engineering). It can explain the relationships between underlying factors but is often non-trivial as the numerous influential factors and interactions not being fully understood ([Giustolisi and Savic, 2006](#)).

By contrast, data-driven models are developed in a 'top down' fashion through analysing inputs and outputs so as to develop a mechanism that maps the former to the latter ([Solomatine and Ostfeld, 2008](#)). Data-driven "models can complement and sometimes replace physically-based models" ([Solomatine, 2002](#)). Knowledge of the underlying processes can help inform construction and testing but is not required; this can be liberating when trying to model systems in which there are many complex interactions and not all of the contributing processes are understood as there is no need to explicitly model the underlying chemistry, microbiology etc. However, for that reason it cannot often offer the same degree of insight into the underlying processes that link inputs to outputs as process modelling. Empirical models can be 'black box' (entirely empirical) in design but are more often 'grey box' i.e. empirical constructs whose design has been influenced to some extent by theory and observations.

Traditionally the structures of data-driven models have been anthropogenic. Formulations are developed using complementary input and output datasets and the resulting models are then used to make predictions given new, possibly unprecedented inputs. Some examples of anthropogenic data-driven models:

- Chlorine decay: [Biswas et al. \(1993\)](#)
- Erosion of estuarine muds: [Parchure and Mehta \(1985\)](#); [Mehta and Lee \(1994\)](#); [Sanford and Maa \(2001\)](#)
- Sewer sediment erosion: [Skipworth et al. \(1999\)](#)
- DWDS discolouration material erosion: [Boxall and Saul \(2005\)](#)
- Impact of free chlorine and temperature on biofilm growth: [Hallam et al. \(2001\)](#)

In recent years, machine-learned data-driven models have become increasingly common, no doubt partly because of Moore's law ([Schaller, 1997](#)). Here, input to output mappings are learned by automated processes, although expert opinion can be used to bound the search domain and determine the predictor variables and target. Common uses for data-driven models include regression, classification and clustering. An overview of data-driven modelling using machine learning, relevant pitfalls, typical methods and some example applications is given by [Solomatine and Ostfeld \(2008\)](#). Other examples of this approach:

- Mathematical expressions of relative discolouration material accumulation rates, found by applying [EPR](#) to [DWDS](#) flushing and asset data from various sites (§2.1.16.1).
- Genetic-algorithm-optimised 'model trees' of biofouling rates in [DWDS](#) pipes ([Opher and Ostfeld, 2011](#)).
- The use of artificial neural networks for modelling the quality of raw (untreated), mid-treatment and treated water has increased in recent years ([Wu et al., 2014](#)).

- Multilayer perceptron artificial neural networks and M5 model trees for estuarine sedimentation prediction ([Bhattacharya and Solomatine, 2006](#)).
- Semblance analysis for quantifying the correlation between flow and turbidity ([Mounce et al., 2014a](#)).

Sometimes a single model may not be sufficiently accurate given the diversity of possible inputs and the complexity of the system(s) that are to be modelled. Here it may be necessary to chain together different sub-models in parallel or series, each of which could be entirely data-driven, based on empirical formulations, process-aware or derived from expert knowledge. The aggregate of these sub-models is known as a *hybrid model* (or a committee of models or a modular model when different sub-models are trained using different regions of the data ([Solomatine and Ostfeld, 2008](#))).

Uncertainty is an important consideration when designing and working with environmental models. There can be numerous sources of uncertainty in a model ([Walker et al., 2003](#)):

- Inputs (imprecision; inaccuracy; missing data);
- Parameters and boundary conditions (imprecision; inaccuracy);
- Structure (e.g. certain effects are only crudely characterised);
- Numerical error

Having an understanding of the possible impact of these uncertainties can help determine whether the model structure ought to allow for stochastic inputs and/or parameters. Model structures may be sufficiently simple for key parameters to be modelled as density functions or sufficiently complex that uncertainty propagation is only feasible using a Monte Carlo approach (running a deterministic model numerous times with perturbed inputs).

Many environmental models require one or more time-series within their sets of inputs, have an internal time-dependent state (i.e. are stateful) and return time-series as outputs. Takens' Theorem ([Takens, 1981](#)) provides one approach to developing empirical stateful, time-series models: *dynamical systems* (each of which are comprised of an input vector, output vector, state vector and first-order differential transition functions) can be approximated using a sliding window of scalar time-series data. [Deyle and Sugihara \(2011\)](#) generalised Takens' theorem to allow the latent system state to be 'recovered' from not only one but also multiple related time-series inputs. A limiting factor is that this approach requires non-noisy data inputs, a pre-requisite that often cannot be satisfied when observing and studying environmental systems.

2.2.3 Workflows for developing environmental models

Although the requirements for and types of environmental models vary enormously, there is some consensus as what makes for a sound model development workflow ([Jakeman et al., 2006](#); [Wainwright and Mulligan, 2004](#); [van Waveren et al., 1999](#)). The first task is to **identify the reasons for modelling and the modelling context**. Reasons may be: to gain a better understanding of a system, to estimate latent quantities, to extrapolate, to test alternative

scenarios and/or to develop control systems. An associated task is to define the **modelling context**, which includes the 'use cases' (needs/requirements) of the various stakeholders (such as the model developer and the model user) and also the scope/extent of the model.

A **conceptual model** can then be developed which defines the inputs, dependencies, outputs, prior knowledge and core assumptions. It is also necessary to consider whether inputs and outputs will be linked by low-level representations of real-world processes or by high-level empirical mechanisms and also decide on the forms of those interactions.

The next stage is to develop a **model structure** that reflects that conceptual model. The structure may be derived from an understanding of the underlying processes, may be anthropogenic empirical formulae or may be learned by automated means from data. There may be a need to encode the model structure as **software** to allow a model instance to be efficiently evaluated. Care needs to be taken when deciding on the number and nature of parameters: an overparameterised model will overfit to noise in calibration data and will not generalise well (Hsieh, 2009).

After this, the model implementation needs to be **verified** to ensure that it is representative of the conceptual model. A **method for calibrating** the model also needs to be developed and validated. The model **sensitivity** to its inputs and parameters needs to be explored.

Finally, the model should be evaluated (**validated**) to ensure that it can represent reality and has sufficient predictive power. The testing of predictive capability should be done using data not used for model calibration, which, if it is time-series data, should contain similar events to the calibration data (Solomatine and Ostfeld, 2008). In simple cases fitting and predictive performance can be quantified using an error metric; in more complex cases it may be necessary to conduct additional/alternative testing such as analysing the ability to accommodate unprecedented scenarios.

2.2.4 Methods for fitting models to data: model fitting as an optimisation problem

2.2.4.1 Introduction

The calibration of environmental models can be posed an optimisation problem in which the aim is to find a set of model parameter values that give a satisfactory correspondence between a model prediction and real-world, noisy data. This approach is attractive and in some cases necessary when calibrating empirical models as parameters cannot be set through taking measurements or making calculations due to them being intangible quantities with no real-world counterparts. Simple empirical models can often be calibrated using trial and error, but this would not be feasible for more complicated models with many parameters and/or complicated, poorly understood parameter spaces due to it being very expensive to thoroughly search the problem space for the optimal parameters. Parameter covariance makes fitting models 'by hand' even more difficult if the modeller can only change and evaluate the effect of one parameter at a time.

Finding a solution to an optimisation problem typically requires the following (Engelbrecht, 2005):

- An *objective function* that is to be minimised (minimisation rather than maximisation is the norm in the optimisation techniques community). For the aforementioned example this function could be the 'Bray Curtis distance' (Scipy, 2012) between a time-series of predictions and a time-series of real-world observations.
- A set of *variables* for which optimal values are sought e.g. model coefficients $m_1, m_2 \dots m_n$.
- A set of variable *constraints* e.g. $0 < m_1 < 1000$.

Optimisation problems can be characterised by the following traits (Engelbrecht, 2005):

- The number and type (real-valued, integer-valued or categorical) of variables (parameters);
- The types of constraints used (just variable boundary conditions or additional equality and/or inequality constraints);
- The number of optima (one (uni-modal) or more than one (multi-modal e.g. fig. 2.16));
- The number of objective functions (uni-objective or multi-objective);

If the objective function is differentiable and uni-modal then optima can be found using gradient-based methods such as Brent's algorithm (Brent, 1972). However, if it is not differentiable and/or has local optima then stochastic optimisation techniques, also known as *metaheuristics*, may be necessary. These use global search algorithms to compare the fitness (objective function value) corresponding to various candidate solutions (sets of parameter values) from throughout the problem space (the domain of all possible candidate solutions). Through such algorithms exploring the entire problem space in a non-deterministic manner rather than searching around a single current best candidate solution for a better one such algorithms are better at finding global optima and avoiding 'getting stuck' in local optima (such as those shown fig. 2.16) (Engelbrecht, 2005). Metaheuristics cannot guarantee finding the true optimum of a problem space but typically return acceptable results, confidence in which can be increased if the stochastic optimisation method is repeatedly applied to the same problem.

The more prevalent stochastic optimisation methods can be classified as being either evolutionary or swarm-based.

2.2.4.2 Evolutionary optimisation and Genetic Algorithms

Evolutionary optimisation techniques such as genetic algorithms (Holland, 1992) mimic the *competition* inherent in natural selection to iteratively move towards an optimal candidate solution. A genetic algorithm typically includes the following steps (Nicklow et al., 2010; Opher and Friedler, 2009):

1. For each of an initial set of candidate solution the corresponding set of variables (model parameters) are encoded as a string to form virtual chromosomes.
2. The objective function is then evaluated for each chromosome to give each a measure of fitness/performance.

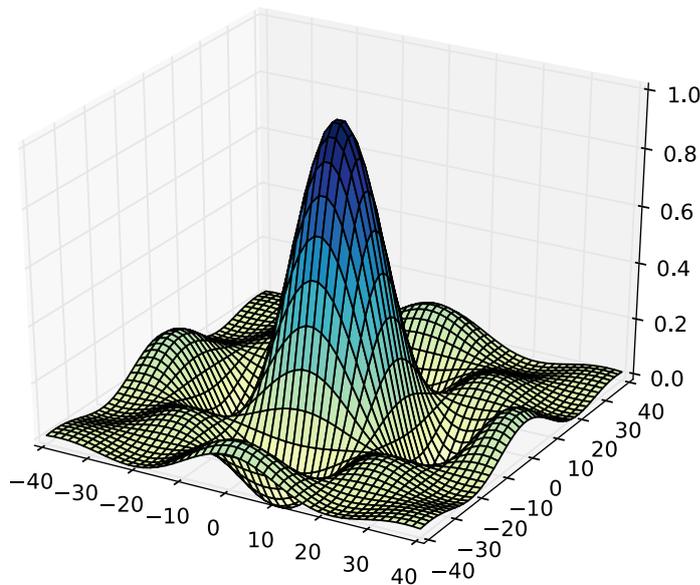


Figure 2.16: A two dimensional sinc function features many local optima and so are often used to test the performance of optimisation algorithms i.e. verify that an algorithm almost invariably finds the single global maximum rather than a local optimum.

3. A new generation of candidate solutions is then generated from the existing generation. Chromosomes with greater fitness values are more likely to be selected in this evolutionary process. To ensure that the problem space of all possible candidate solutions is well explored there is 'crossover' (exchange) of genetic material between certain selected chromosomes whilst other selected chromosomes are subjected to random mutations.
4. The process then repeats until convergence tolerances are met or a maximum number of generations is reached.

Discrete, ordinal variables are most naturally suited to being encoded in strings of chromosomes although real numbers can be encoded using *direct value encoding* (Sivanandam and Deepa, 2007).

2.2.4.3 Swarm optimisation and the particle swarm optimisation algorithm

Swarm optimisation methods are also stochastic but use collaboration rather than competition between candidate solutions to converge on the global optima.

One of the more commonly-used techniques is particle swarm optimisation (PSO) (Kennedy et al., 2001), which can be seen as being analogous to the way a flock of birds swarm to find the optimal concentration of flies in three-dimensional space.

Each dimension corresponds to an optimisation problem variable (model parameter). For the bird flocking problem, the objective function that needs to be minimised could take a set of coordinates in three-dimensional space as its three parameters and return the reciprocal of the concentration of flies at those coordinates. Initially a number of birds or *particles* (typically between 20 and 50 (Eberhart and Shi, 2001)) are assigned random positions and velocities within the problem space. The coordinates of each bird represent a candidate solution. The

birds then move about this problem space over a number of timesteps. At each timestep each bird calculates a new velocity based on three factors (Floareano and Mattiussi, 2008):

- **Inertia:** its current velocity vector
- **Nostalgia:** the problem space position in the bird's memory since $t = 0$ where the fitness value was greatest;
- **Societal influence:** the position of the neighbouring bird with the best fitness at the current timestep (birds may use calls to communicate their current fitness values to each other).

These factors are illustrated in fig. 2.17 using a sinc objective function (note that this is a 2D rather than 3D objective for ease of representation on the page).

Each bird then moves to a new position given its new velocity vector. If at the end of a timestep the best candidate solution in the swarm is good enough to allow some convergence criteria to be satisfied the algorithm exits. A maximum number of iterations is typically specified to ensure the algorithm terminates.

Two key advantages of PSO over evolutionary optimisation approaches such as genetic algorithms are that PSO is purportedly much more efficient than genetic algorithms and the key algorithmic components can be encapsulated in very little code (Kennedy et al., 2001). A possible disadvantage of PSO is that it is not naturally suited to optimising problem involving categorical variables (Floareano and Mattiussi, 2008).

Velocity component weights Explicit weightings and random variables are used to determine the relative influence of the three velocity components. The inertia weight is often decreased as the simulation progresses to ensure that the particles explore the full extent of the problem space in the early timesteps then collectively hone in on one or more optima as the simulation progresses (Xu and Rahmat-Samii, 2007). Typical weight values are given by Eberhart and Shi (2001), Xu and Rahmat-Samii (2007) and Floareano and Mattiussi (2008). The random variables take values in the range $[0,1]$ and are only used to scale the nostalgia and societal influence velocity components.

Particle neighbourhoods The neighbourhoods used to communicate fitness between particles at a given timestep can be either *geographic* or *societal*. The set of particles that comprise the societal neighbourhood of a distinct particle is invariant throughout a simulation (Floareano and Mattiussi, 2008). The use of small neighbourhoods, such as the *ring* social topology in which each particle has exactly two neighbours, are less likely to result in the algorithm getting stuck in local minima than the *global topology* in which all particles communicate with all others. However PSO may take a long time to converge if using smaller neighbourhoods.

Kennedy and Mendes (2002) studied the impact of using different topologies on convergence speed and the accuracy of the result (the ability to avoid getting stuck in local minima). The *Von Neumann lattice* in which each particle has exactly four neighbours (the number of particles must therefore be a square) was recommended as a general-purpose topology.

Neighbourhood sizes of $\approx 15\%$ of the swarm size are common (Eberhart and Shi, 2001).

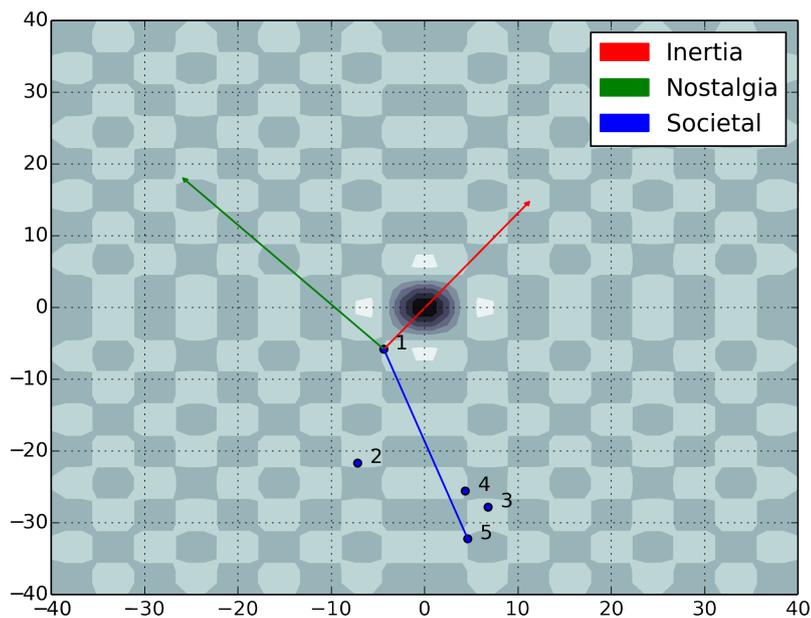


Figure 2.17: A view of particle positions from five timesteps into an attempt to find the global minima of a 2D objective function using *PSO*. The objective function, depicted using a heatmap for illustrative purposes, is a 2D sinc function. Knowing that darker areas correspond to lower ('better') objective function values, one can see that there is a global optimum at $(0,0)$ but there are also many local optima. Only a few timesteps have passed since the start of the simulation so there is limited evidence of the particles converging on the global optimum; however, the optimum is eventually found after 68 timesteps to within the required accuracy of $(0.001, 0.001)$. The arrows show the vectors that are to be used for calculating a new velocity and position for particle 1: there is the current velocity vector ('inertia'), the vector to the position where particle 1 encountered the best objective function value to date ('nostalgia') and the vector to the position of the swarm neighbour who currently has the best objective function value ('societal').

Problem space boundaries If a particle's position is updated following the calculation of a new velocity vector then there is a need to check that the new coordinates correspond to valid values of the optimisation problem's variables. It may not be possible to evaluate the objective function using invalid values (Xu and Rahmat-Samii, 2007).

During particle relocation, particles can either be *restricted* or *unrestricted* by the problem space boundaries. In the case of unrestricted relocation particles are allowed to move outside the problem space but are assigned highly anomalous fitness values (without the objective function being evaluated). Xu and Rahmat-Samii (2007) found that the restrictive, damping boundary condition provides "robust and consistent performance". This condition requires that particles are moved back to the problem space boundary if they leave it; the sign of the velocity component that caused the particle to move outside the problem space is then flipped ('reflection') the particle and the component's corresponding magnitude is reduced ('damping') by multiplying a factor in the range $[0, 1]$ (a loss of momentum).

Velocity limits To ensure that particles do not move too quickly past 'good' regions of the problem space the maximum particle velocity can be set to a user-specified level or more simply to the absolute size of the region bounded by the used problem space box bounds (Eberhart and Shi, 2001).

Algorithm performance and implementations Although PSO is a simple and efficient stochastic optimisation algorithm, execution times can still be notable for certain optimisation problems and PSO configurations, particularly if the objective function is expensive to evaluate. At each iteration during an optimisation run the objective function must be evaluated per particle and hundreds of iterations may be required before a satisfactory optimum is found. As PSO is a stochastic process, multiple runs may be required to increase confidence that the global optimum has been found and that the results are sufficiently repeatable.

PSO execution times can be significantly reduced using parallel computing. The objective function can be evaluated for multiple particle simultaneously to expedite single optimisation runs. Also, multiple independent optimisation runs can be conducted in parallel to more thoroughly explore the problem space and determine the reproducibility of results without notably increasing execution times (Izzo, 2012). These two options are not mutually exclusive and are trivial to implement as there are very few dependencies between parallel threads of execution (it is an *embarrassingly parallel* problem (Wilkinson, 1999))

Existing implementations of the PSO algorithm include *PyGMO* (Izzo, 2012) and *hydroPSO* (Zambrano-Bigiarini and Rojas, 2013).

2.2.4.4 Handling non-trivial constraints when using metaheuristics

There is no de-facto standard for implementing non-trivial parameter constraints (constraints other than independent box bounds per parameter) when using evolutionary optimisation algorithms or PSO. One approach is to lessen the influence of invalid candidate solutions by including a penalty function within the objective function (Parsopoulos and Vrahatis, 2002). Consideration needs to be given to the form of this function: penalise too much and

metaheuristics are more likely to get 'trapped' in local minima; also, dynamically changing the penalty values during a metaheuristic run can yield better results but requires yet more design decisions and fine tuning (Parsopoulos and Vrahatis, 2002; Fuentes Cabrera and Coello Coello, 2007). An alternative approach when using PSO is to reinitialise the positions of invalid particles but then simply ignore invalid candidate solutions when determining the best historical, current social (neighbourhood) and current swarm candidate solutions (Hu and Eberhart, 2002); this approach may be computationally simple but it can significantly reduce the number of candidate solutions that are evaluated concurrently and so negatively impacts on the exploration of the problem space. A third approach is to 'repair' invalid candidate solutions by updating them to the nearest valid position in the problem space (Nicklow et al., 2010), but this introduces complexity by requiring localised searches around invalid solutions. In summary, methods for handling non-trivial constraints when using metaheuristics exist but all are problematic in some respect (robustness, complexity, computational expense and or manual configuration/tuning requirements).

2.2.4.5 Examples of the use of stochastic optimisation techniques in water engineering

Savic and Walters (1997) demonstrated the efficacy of genetic algorithms for determining least-cost designs for DWDSs, an optimisation problem made difficult by the use of discrete variables (e.g. pipe diameter), hydraulic constraints, local optima and multiple objectives. Yagi and Shiba (1999) demonstrated how optimal fuzzy rule sets could be derived using genetic algorithms for controlling sewer pumps. Nicklow et al. (2010) detail many examples of genetic algorithms being used for parameter identification in the varied fields of "water distribution systems, urban drainage and sewer systems, water supply and wastewater treatment, hydrologic and fluvial modeling [and] groundwater systems". Opher and Ostfeld (2011) used genetic algorithms to optimise model-tree regression methods for learning pipeline biofouling rates from a large number of predictor variables. In an article particularly relevant to this thesis McClymont et al. (2013) built on the existing body of work relating to the use of genetic algorithms for DWDSs pipe sizing by developing a multi-objective method for designing networks where the hydraulic regimes are prohibitive to discolouration material accumulation. Farley et al. (2013) showed how genetic algorithms could optimally subdivide DMAs for the purposes of rapidly detecting bursts using pressure sensors. As mentioned in §2.1.16.1, Mounce et al. (2014c) demonstrated how evolutionary polynomial regression (EPR) can be used to learn equations for discolouration material regeneration rates from subsets of possible predictor variables; here a genetic algorithm is used to refine a numerical regression model.

PSO has been used to design distribution networks with optimal topologies and pipe diameters (Suribabu and Neelakantan, 2006), develop optimal reservoir management strategies (Reddy and Nagesh Kumar, 2007), learn the optimal parameters of a lake water quality model (Campbell and Phinn, 2009), parameterise hydrological and hydrogeological flow models (Zambrano-Bigiarini and Rojas, 2013) and learn formula parameters for solute mixing at DWDS junctions (Yu et al., 2014).

2.2.5 Quantifying the dissimilarity between observations and predictions: correlation metrics

To fit an empirical model to data using an optimisation method (§2.2.4) one needs a suitable objective function that includes a measure of the dissimilarity/anti-correlation between some observations and predictions⁴. In environmental and water engineering, model outputs are often vectors that indicate how a quantity varies with time and/or space, so the fitting objective is often to minimise the dissimilarity between 1D vectors of observations and model predictions that correspond to the same temporal/spatial domain. An example of this is calibrating hydraulic models so that time-series pressure predictions are within a tolerance of observed pressures at hydrants monitored over a specific period (Walski et al., 2003).

The choice of dissimilarity metric for comparing series depends on the application and data: metrics vary in their sensitivity to scale, shape and outliers. For example, the *Chebyshev distance* is highly sensitive to outliers as it is solely a function of the greatest pairwise distance between two vectors, whereas the *Cosine distance* and *Pearson product-moment correlation coefficient* (*Pearson's r*) are insensitive to differences in scale between their two inputs. For an analysis of these sensitivities for fourteen dissimilarity metrics (many of which are commonly used in environmental and water engineering) see Sonnenwald et al. (2013).

Dissimilarity metrics can also be classified by whether they are dimensional. When comparing model fits for two systems only relative measures of error such as R^2 (the *Coefficient of Determination*) or the Nash Sutcliffe (Model) Efficiency Index (*NSEI*) (Nash and Sutcliffe, 1970) are appropriate. However, when comparing fits for a model of a single system, as required when fitting model parameters through iterative optimisation, dimensional similarity metrics such as the integral of squared errors (*ISE*) can be more powerful than relative measures (Sonnenwald et al., 2013).

It is also important to be aware that different dissimilarity metrics have different ranges. For example, the *NSEI* has a range of $(-\infty, 1)$ whereas R^2 has a range of $(-1, 1)$ (with 1 being the ideal value for both) (Bennett et al., 2013).

2.2.6 Methods for signal decomposition and filtering

When analysing time-series data it is sometimes necessary to isolate different components of the signal based on frequency. For example, one may wish to filter out low-frequency components, leaving just high-frequency components relating to the process of interest. Alternatively, it may be that the signal includes high-frequency noise that needs to be filtered out before further analysis can be performed.

One of the simplest and better approaches to removing irregular noise from a time-domain-encoded signal is the moving average filter (Smith, 1997), where:

1. A window is defined that contains a fixed number of samples at the start of the time-series (the window must contain a proper subset of the time-series);
2. All data within the window is then aggregated as say the mean of that data;

⁴Dissimilarity rather than similarity as the convention with optimisation algorithms is to minimise.

3. This aggregate value is assigned to the filter output at the start, middle or end of the window period;
4. The window is then moved along one data point in the series (but retains its size) and the process repeated until the window cannot be moved any further through the series.

Various aggregate functions can be used with moving average filters. The most common is the arithmetic mean but the median or a quantile may perform better if the data within the window is not normally distributed.

A more complex set of alternative means for removing unwanted components of a signal are auto-regressive moving average (ARMA) filters. 'Autoregressive' means that the filter output is self-referentially used as a filter input during the filtering process. One example of an ARMA filter is a Butterworth filter, which has the attractive quality of providing a flat response over most of the 'passband' (the range of frequencies not removed from the signal; (Smith, 1997)).

A (digital rather than analogue) Butterworth filter is parameterised by:

- The absolute cut-off frequency: the frequency where the power (amplitude squared) of a filtered signal is half the power of the input signal;
- The Nyquist frequency, which is half the sampling frequency;
- The filter 'order';
- The type of filter e.g. 'high-pass' or 'low-pass', which retain only higher or lower frequency components respectively.

There are numerous examples of the use of ARMA filters in water and environmental engineering. For example, Misiunas et al. (2005) used a Butterworth filter to remove high-frequency noise from pipe transient monitoring data whereas McKay et al. (2013) used a low-pass Butterworth filter to isolate the frequency range relating to the processes of interest in river shear stress data.

2.2.7 Modelling water quality in DWDSs

2.2.7.1 Importance of accurately characterising DWDS hydraulics

Simulating how the amount of discolouration material at the pipe wall and in the bulk water varies in space and time can be thought of as a DWDS water quality modelling problem. A sound understanding of the hydraulics of the system is a fundamental prerequisite for time-series discolouration modelling and water quality modelling in general:

- The shear stress at the pipe wall drives material erosion;
- Advection is a function of the bulk water propagation delay;
- The degree of mixing at junctions is a function of the magnitude and concentrations of inflows;

- Bulk water concentrations/turbidity are a function of material erosion rates from the wall *and* flow.
- Material accumulation rates are most likely a function of nutrient, carbon and disinfectant residual supply rates along with temperature, which also may be a function of hydraulics in DWDSs with long residence times.

To characterise their hydraulics, DWDSs are typically approximated as networks of vertices (representing junctions, inflows and outflows and tanks) and one-dimensional edges (representing pipes, valves and pumps) then solved for flow and pressure (Walski et al., 2003) at a series of discrete times. Hydraulic DWDS solving software is readily available (DNV GL, 2010; Schneider Electric, 2012; Bentley Systems Inc, 2013; Innovyze Ltd, 2014), including the open-source EPANET application and software library (USEPA, 2008).

2.2.7.2 Fundamentals of modelling water quality in DWDSs

DWDS water quality solvers take the calculated time-series solution to hydraulic network as input, either reading it in all at once or on a per-timestep basis. The majority of water quality solvers (including EPANET) allow for the modelling of a single chemical species. Concentrations in mass per unit volume at nodes and average concentrations for pipes are calculated by considering:

- mass injection;
- growth/decay due to interaction with the bulk water;
- growth/decay due to interaction with pipe and/or tank walls;
- advection;
- mixing at junctions;
- mixing in tanks.

2.2.7.3 Advection and mixing

With most DWDSs water quality modelling applications it is assumed that Reynolds numbers are high enough for the effects of diffusion and dispersion to be negligible. This allows advection to be modelled relatively simply by assuming *plug flow* i.e. that the water moving through a (one-dimensional) pipe is comprised of parcels each of which has a specific concentration and does not mix/interfere with its neighbours. This is typically implemented by tracking the movement of parcels through the network (Lagrangian transport; see fig. 2.18), migrating mass in and out of nodes around junctions as parcels move into or out of them, but it is also possible to instead track concentrations for fixed pipe sub-divisions (Euclidean transport) (see Rossman and Boulos (1996), for more information on transport mechanisms).

Note that when modelling the advection of solutes and solute-like substances one may want factor the instantaneous concentration at an observation point by the flow at that time; this

yields a mass flux that, when integrated over time, gives the amount of mass that has passed an observation point over a given period (as per §2.1.13.2).

Mixing models are usually simplistic. It is assumed that complete mixing occurs at pipe junctions. Tank/service reservoir mixing models typically allow for complete mixing, last-in-first-out (LIFO), first-in-first-out (FIFO) and/or simple 2-3-part compartmental models (for modelling mixing and semi-stagnant zones).

2.2.7.4 Sources of error

An accurate understanding of contact times and water age is needed when modelling time-dependent water quality processes in DWDSs but travel times predicted by existing hydraulic models may be inaccurate. Such models are typically built to satisfy flow and pressure tolerances through determining internal pipe roughnesses (Walski et al., 2003) that ensure accurate headloss between calibration locations. To also calibrate for bulk velocities one needs to adjust both roughness *and* internal diameter to account for a lack of hydraulic capacity due to the build-up of scale/tubercles on the pipe wall (Boxall et al., 2004; Walski, 2004). However, measuring velocities is difficult to do in practise due to the systems of interest being buried infrastructure and utility companies are understandably reluctant to introduce tracers into their networks (although this is sometimes permitted; see Skipworth et al. (2002b)).

Another source of error in water quality modelling comes from flow meter readings not always being precise or accurate, which impacts on velocity and therefore solute/wash-load travel time estimates. Even if the meters are accurate through being correctly sized for the adjoining mains, aggregate flow data measurements are typically only transmitted via SCADA every fifteen minutes (Burnell, 2003; Mounce et al., 2012b) so minimums and maximums may not be captured. This could be significant for certain modelling applications (e.g. discolouration) but not others (e.g. chlorine residual decay). Uncertainty in bulk velocities and travel times may also arise from estimates of customer demand, which can be highly variable over a range of timescales. However, stochastic demand modelling can help account for these uncertainties (Blokker, 2010).

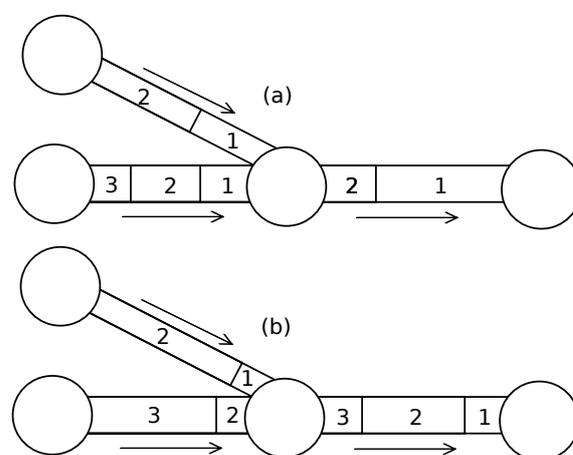


Figure 2.18: Behaviour of segments in Lagrangian methods: (a) Time t ; (b) Time $t + \Delta t$ (after Rossman and Boulos, 1996).

2.2.7.5 Modelling the growth/decay of a single-specie

With single-specie DWDS water quality modelling, growth or decay are usually accounted for by assuming zero, first or second-order reaction kinetics, that all other reactants are in abundance and that temperature is invariant. In reality, concentrations of other reactants may be limiting factors in the myriad of occurring reactions (Walski et al., 2003) and the spatio-temporal variation in temperature in DWDSs may be significant (§2.1.16.2).

Bulk water reaction rate coefficients can be found relatively easily using bottle tests (Walski et al., 2003) but quantifying wall coefficients can be more difficult. Firstly, the wall reaction coefficient for a given pipe can only be quantified (at a moment in time) by measuring growth/decay over its length and discounting the growth/decay due to bulk reaction(s); the error in wall coefficient calculations therefore includes the error associated with the bottle tests used to find the bulk reaction coefficient (Walski et al., 2003; Savic et al., 2009). Secondly, finding the wall reaction coefficient *in-situ* for all required permutations of pipe type, water source and environmental conditions requires extensive field work.

2.2.7.6 Modelling multiple reactants

EPANET MSX (Shang and Uber, 2008) allows for the modelling of interactions between *multiple* chemical and biological species at the pipe wall and suspended/dissolved in the bulk flow (Shang et al., 2007), thus making it easier to account for limiting reactants (Uber, 2009). The dynamics of a specie can be expressed using mathematical formulae as either a) equilibria (for irreversible or near-instantaneous reactions), b) rate laws involving the concentrations of other species or c) formulas for the concentration at a given time that reference the concentrations of other species. Species concentrations are tracked in the bulk water and at pipe walls over time.

Attempts have been made to use EPANET MSX to model complex sets of reactions such as those relating to bacterial regrowth or monochloramine decay (Shang et al., 2007) but such efforts can be limited by:

- The lack of an exhaustive understanding of the types of reactions that occur in distribution systems or a lack of the resources required to calibrate a model containing all interdependent reactions.
- Uncertainty regarding equilibrium constants and rate coefficients for reactions involving bulk and/or wall species. As with single-specie water quality modelling, values must typically be acquired through field tests and are highly site-specific (Shang et al., 2007).

EPANET MSX may provide more power and flexibility than single-species modelling but such multi-species modelling is also more complicated and prone to error.

2.2.8 Discolouration modelling: the need for a data-driven approach

2.2.8.1 Introduction

Many different factors contribute to discolouration material accumulation and erosion in DWDSs (e.g. §2.1.16). One could in theory model all the pertinent processes using a bottom-up,

process-driven, white-box approach (§2.2.2) and a tool such as EPANET MSX (§2.2.7.6) but in practise the number of interactions that would need to be characterised would make this unwieldy if not infeasible. A more practical approach would be to attempt to relate several key predictor variables to variations in turbidity to form a grey-box model that incorporates calibrated empirical water quality models and hydraulic network solutions. Examples of this approach are presented in the following sections.

2.2.8.2 The Prediction of Discolouration in Distribution Systems (PODDS) model

The empirical PODDS discolouration model (Boxall et al., 2001) does not attempt to model corrosion processes, biofilm development, suspended iron flocs and chunks of biofilm directly; instead, material erosion and accumulation are considered to be mechanisms by which turbidity-causing homogeneous material can move between being bound to the pipe wall and being a wash load suspended in the bulk water. The model tracks both the *discolouration potential* of a pipe (the amount of material bound to the pipe wall at a given time) and the *turbidity response* that would result from material erosion.

As turbidity had been shown to be approximately linear with TSS in many cases (§2.1.8) and discolouration material only settles out under entirely quiescent conditions (§2.1.9), turbidity was assumed to be a pseudo-concentration (and so the product of turbidity and bulk water volume gives a material quantity in the abstract units of $NTU \cdot m^3$).

2.2.8.3 PODDS concepts for material mobilisation

Discolouration material typically binds to the pipe wall in cohesive layers around the pipe circumference (§2.1.10; rather than accumulates by gravitational settling along the pipe invert, §2.1.12) and is eroded when the material's cohesive strength is exceeded by the shear stress at the pipe wall (§2.1.12). The mechanism within PODDS that codes for this conceptualisation was inspired by mathematical models of the physical characteristics and erosion of estuarine muds (Parchure and Mehta, 1985) and sediments in sewers (Skipworth et al., 1999).

Under PODDS, the (scalar) shear strength of the cohesive material bound to the pipe wall is a function of the discolouration potential:

$$\tau_c = \frac{C^b - C_{max}}{k} \quad (2.15)$$

where

τ_c The current material strength [Pa];

C The ability to increase turbidity per unit volume of erodible material i.e. the amount of material at the pipe wall [$NTU m^3$];

C_{max} An upper limit for C [$NTU m^3$];

b A power term that has been removed since the model was first proposed as “no additional quality of fit was obtained by allowing values of b other than unity” (Husband, 2010) [–]

k A negative gradient term that relates yield strength and ability of layer to increase turbidity [$NTU\ m^3 \cdot Pa^{-1}$].

Should there be an *excess shear stress* (meaning that the shear stress τ_a (eq. 2.9) exceeds the layer shear strength, τ_c) as a result of an anomalous increase in flow then the rate of supply to the bulk flow is a non-linear function of that excess shear stress.

$$R = P (\tau_a - \tau_c)^n \quad (2.16)$$

R rate of supply from a m^2 of wall area [$NTU \cdot m \cdot s^{-1}$]

P linear coefficient of eroding force [$NTU \cdot m \cdot Pa^{-1} \cdot s^{-1}$]

n exponential coefficient of eroding force [–]

The n term prevents the units from balancing but the calibration of the model using field data frequently requires it to be non-unity.

The rate of supply then determines both the increase in turbidity and the decrease in discolouration potential (eqs. 2.17 and 2.18 respectively).

$$\Delta T = \frac{R \cdot A_s}{Q} \quad (2.17)$$

T turbidity [NTU] (note that Boxall et al. (2001) used the symbol N rather than T but the latter has been used here for consistency throughout the thesis)

A_s area of pipe wall swept [m^2]

Q bulk flow [$m^3 \cdot s^{-1}$]

$$\Delta C_e = R \cdot \Delta t \quad (2.18)$$

ΔC_e change in discolouration potential per m^2 of wall area due to erosion [$NTU \cdot m$]

The above equations assume that the material layers have reached a steady state prior to a notable increase in flow i.e. the layers are as maximal as the prevailing hydraulics will allow (Boxall and Saul, 2005). This assumption is invalid for situations where layers have been eroded by a hydraulic event and are therefore regenerating prior to the modelled period of interest. Previous data analysis has provided estimates of how long it might take the quantity of material bound to the pipe wall to reach maximal levels (§2.1.16.3).

2.2.8.4 Calibrating and validating the PODDS erosion model

The parameters C_{max} , k , P , n and τ_c (and b) are empirical and therefore a PODDS model of a pipe network needs to be calibrated either by hand or less commonly using optimisation techniques (see §2.2.4) such as genetic algorithms (Boxall and Saul, 2005); calibration requires

the minimisation of the dissimilarity of turbidity predictions and observations (fig. 2.19). Flow trials (flushing DMA pipework or raising flows in trunks) are used to generate the turbidity, flow and headloss data required for model calibrations. Headlosses can either be estimated (using eqs. 2.10, 2.11 and 2.12) or measured directly. Parameter sets are transferable between sites to a degree. Parameter values from calibrated network models can be used to initialise calibrations of similar networks prior to the values being adjusted by hand (Husband and Boxall, 2010a). For example, knowledge of the parameters found through calibrating a model of a trunk system (Husband et al., 2010a) has greatly aided the calibration of models of other comparable trunk networks.

Sensitivity analysis of the PODDS parameters have shown them to be highly sensitive to pipe diameter, pipe material and source water (Boxall and Saul, 2005). The fitting the model to data has been shown an indeterminate problem i.e. there is no unique set of model parameters that gives the optimal fit. Aisopou et al. (2010b) found that two calibrations of the same system (using a single flow and turbidity dataset that corresponded to a single discolouration event) yielded very different parameter values. The ways in which the parameters covary have not been investigated in the literature. The calibration of a hydraulic DWDS model can also be indeterminate: there are an infinite number of diameter, roughness, flow and local loss combinations that give the measured headloss and determining the 'true' values of those parameters is not often possible given that they relate to buried infrastructure and roughness and local losses are latent quantities (see fig. 2.20, Boxall et al. (2004) and Walski (2004)).

This indeterminacy can make model fitting difficult, particular as all PODDS parameters are empirical. To 'bootstrap' the calibration of a given model, a sparse look-up table (LUT), indexed by diameter and pipe material, has been compiled, containing values of k , P , n and C_{max} found from calibrating other models. This LUT is considered the intellectual property of the PODDS project stakeholders⁵ and so has not been published (Husband and Boxall, 2010b).

The PODDS model of cohesive layer erosion has been validated for DMAs (Boxall and Saul, 2005; Husband and Boxall, 2010a) and trunk systems (Husband et al., 2010b,a, 2011): the application of a stepped shear stress profile during flushing has been shown to yield discolouration material at each stress level (Husband and Boxall, 2011).

⁵See <http://www.podds.co.uk>

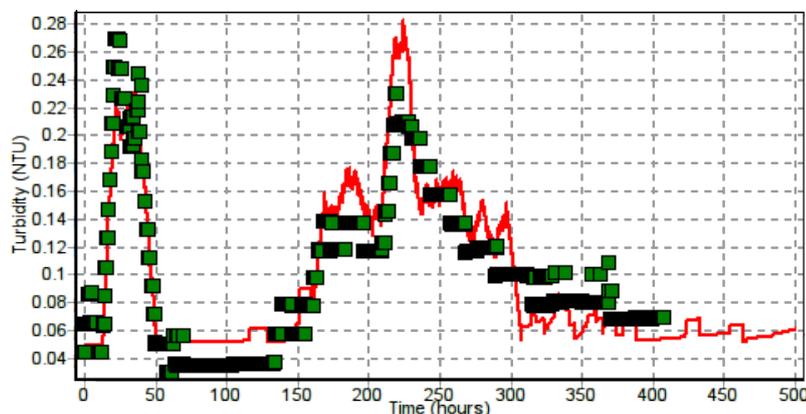


Figure 2.19: Graphical output from the PODDS software, a modified version of EPANET (Rossman, 2000), showing the strong similarity after calibration of turbidity observations (green squares) and predictions (red line) for a calibrated PODDS model.

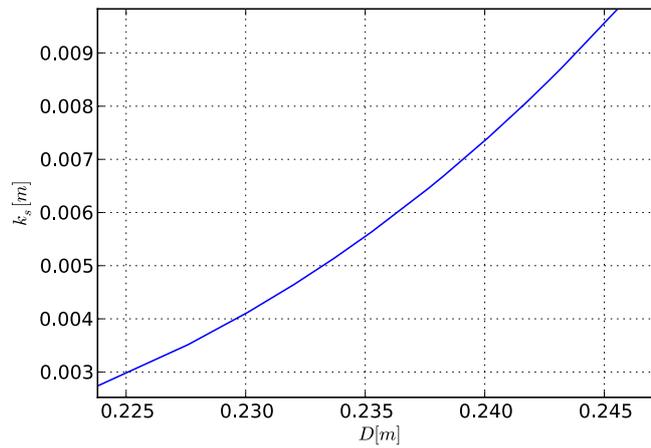


Figure 2.20: The calibration of hydraulic distribution network models is an indeterminate problem. For a measured flow Q an infinite number of pairings of diameters (D) and roughnesses (k_s) give the observed headloss h_L , as illustrated in this quantitative example.

Modelling work has been shown to support the evidence discussed in §2.1.15 of the material bound to the walls of plastic pipes having a maximum strength. Husband and Boxall (2010a) found that their ability to fit the PODDS model to flushing data from English plastic DMA pipes decreased as the flushing shear increased. Better fits were obtained through limiting the maximum material strength to $1.2 Pa$ in certain pipes and $1.8 Pa$ in others. The PODDS formulations ensure that these ultimate strengths and material depletion are approached asymptotically for a step increase in shear stress.

From discussions with the authors of PODDS it was learnt that the b parameter, the exponent of the discolouration potential C , has never needed to be anything other than unity when calibrating models. This and allows the units of eq. 2.15 to balance. In contrast, the n parameter in eq. 2.16 often needs to be set to non-unity.

PODDS has been implemented by revising the water quality engine (Boxall and Saul, 2005) included with the open-source EPANET DWDS modelling software (§2.2.7): the discolouration material release from the pipe wall per timestep is first calculated using the PODDS formulae then existing EPANET code handles the advection and mixing of material that is suspended in the bulk water, including the ‘importing’ of material from upstream of the model boundary. To assess discolouration risk using the software, one typically begins with conditioning the cohesive layers by running a discrete time-series simulation with a flow profile that is representative of the prevailing hydraulics. One then takes the final layer state from this simulation as the starting state for a second simulation. This second simulation is used to predict the turbidity response that will result from a flow time-series profile of interest (Boxall and Husband, 2007).

2.2.8.5 Limitations of the PODDS model

Representation of the material accumulation process The PODDS model was developed to model material erosion from the pipe wall but a temperature-dependent accumulation mechanism was included with the other formulations in recognition of importance of accumulation over timescales longer than those considered when only studying erosion. This mechanism was not derived from lab or field studies or from physical principals but was instead a product of

the authors' engineering experience and intuition; it assumed that material layers regenerate from strongest to weakest following their erosion (Boxall et al., 2001, eq. 9). However, field and laboratory studies have shown that discolouration material accumulates on pipe walls at all strengths over a range simultaneously (§2.1.13.4). If material accumulates at all strengths simultaneously but then the weaker material erodes much quicker than the stronger (as indicated by the validated PODDS erosion mechanism, eq. 2.16) then continual erosion and accumulation will lead to there being different amounts of material at the wall with different strengths. This situation is not something that PODDS can represent as it assumes the amount of material at the wall is inversely proportional to a single scalar strength (eq. 2.15). The PODDS material accumulation model is therefore incapable of accurately representing reality and cannot model discolouration processes other than short-term erosion.

Modelling erosion due to unsteady hydraulics Under the PODDS model, material erosion is only driven by the steady-state (Darcy Weisbach) shear stress (eq. 2.16) and yet under transient conditions the unsteady shear stress could be much larger, even when there is little net pipe flow (§2.1.12.3).

Attempts to couple the PODDS material erosion model to unsteady shear stress and advection models have proved inconclusive:

- Naser et al. (2006) tried to represent the system hydraulics and calculate the turbulent shear stress using a five-region 2-D model. For a presented case study their approach was limited by their shear model not accounting for pipe roughness and the unsteady shear stress being dominated by the roughness-dependent Darcy Weisbach shear stress.
- In contrast, the Vardy-Brown shear stress model (Ghidaoui and Mansour, 2002) used by Aisopou et al. (2010a) for assessing the impact on discolouration material erosion was dependent on roughness. For a case study hydraulic event the unsteady shear was momentarily much larger than the steady state shear however there was no difference in bulk water turbidity prediction between using the sum of the Darcy-Weisbach and Vardy-Brown shear stresses or just the Darcy-Weisbach shear stress.

The insensitivity of the discolouration modelling conducted by Aisopou et al. (2010a) to transients could have been due to the PODDS erosion rate factor P (found through trial and error) being sufficiently slow and/or the duration and magnitude of the transient being such that the erosion due to the transient was negligible. However, Aisopou et al. (2010a) only presented one case study; the unsteady shear may cause significant erosion under other conditions.

The data requirements for 2D unsteady friction models such as those used by Aisopou et al. (2010a) along with the complexity of the friction models are both much greater than their 1D steady-state counterparts. Aisopou et al. (2010b) used a timestep of 0.007 s but pressure logging is rarely performed in DWDS at such fine temporal resolution and considerable effort is required to keep loggers running due to the large volumes of data captured. Given these issues and the infrequency of transient pressure events (between 0 and 2.2 per year in the systems studied by Besner et al. (2010)) it may be difficult to justify the expense of modelling systems using anything but 1D steady state hydraulics.

The discrepancies between steady state versus dynamic hydraulics in the context of discolouration risk have also been modelled without using **PODDS**. [Pothof and Blokker \(2012\)](#) modelled a looped part of a Dutch distribution network using three different approaches: steady-state, 'rigid column' (including inertia effects) and water-hammer (modelled a compressible fluid and the Young's modulus of the pipe, allowing for accurate transient pressure wave propagation). The aim was to assess whether different hydraulic solutions better correlated with the amount of material that had accumulated in each of several pipe lengths. Hydraulic simulation results were very similar between the steady-state and rigid column models, whereas oscillations of $\cong 0.02 m \cdot s^{-1}$ around the steady-state bulk velocity were predicted by the water-hammer model. High-temporal-resolution observations flow measurements were not available to validate this discrepancy.

Model sensitivity to the excess shear exponent term n The use of an exponent term in the expression for the material erosion rate (eq. 2.16) causes an imbalance of units. This term is also undesirable when fitting the model to data as small errors in n during calibration can result in much larger errors when predicting the turbidity response for a subsequent large increase in excess shear (fig. 2.21).

Accounting for longitudinal variation in discolouration processes Under the **PODDS** model it is assumed that material accumulation and accumulation are longitudinally homogeneous along each pipe. However, there is some empirical evidence that indicates this may not always be true (§2.1.16.6). The assumption might still be valid for the short pipe lengths found in **DMA**s but may be less appropriate when modelling lengthy trunk mains, which can be of greater interest when considering discolouration risk (§2.1.17.2).

To circumvent the limitations of this assumption, less-coarse spatial discretisation can be used: by dividing what had assumed to be a homogeneous length of $\varnothing = 344 mm$ **DI** trunk main into four sections and assigning each different **PODDS** parameters [Husband et al. \(2011\)](#) improved the model fit during calibration. However, this partitioning was arrived at by trial and error, presumably because the underlying software (EPANET's water quality engine) would not without modification allow for a more granular, gradient-based approach to modelling material accumulation and erosion.

2.2.8.6 Other discolouration models

The Particle Sediment Model (PSM) The **PSM** ([Ryan et al., 2008](#)) was designed to model a) sedimentation and resuspension from different points along pipe and b) accumulation of material around the pipe circumference due to Van der Waals forces (comparable to the cohesive layers in **PODDS**). However, as of 2008 only a) was implemented in software and the model has not been referenced in the literature since. The inability of the reference implementation to model the development and erosion of cohesive layers (as is required; see §2.1.9, §2.1.10, §2.1.12, §2.1.13.4) limits its applicability.

The Discolouration Risk Model (DRM) and Discolouration Propensity Model (DPM)

The **DRM** ([Dewis and Randall-Smith, 2005](#)) was a framework for assessing the impact of

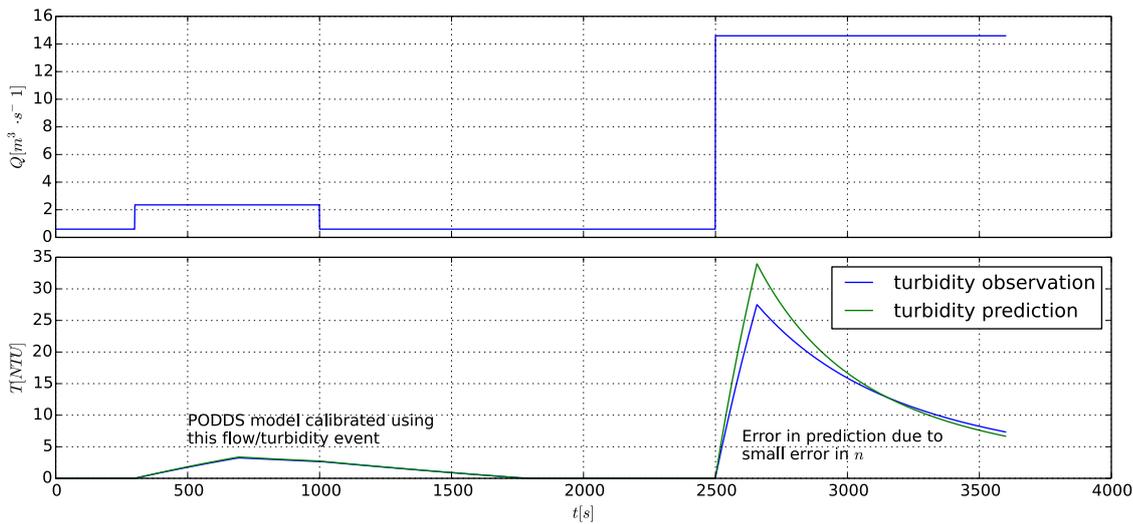


Figure 2.21: A small error in *PODDs* n during calibration can yield much larger errors when making subsequent predictions (here $Q \in [0.002, 0.01]$, $\varnothing = 100 \text{ mm}$, $k_s = 5 \text{ mm}$, $L = 200 \text{ m}$, *PODDs* $k = -0.3$, *PODDs* $P = 0.0001$, *PODDs* $n = 1.5$ (artificial turbidity 'observation'), $n = 1.6$ (turbidity prediction))

not only hydraulic increases on discolouration potential but also structural failures on turbidity throughout a network. Little information has been published on how discolouration potential was to be quantified.

The *DRM* has been superseded by the *DPM* (Randall-Smith et al., 2011; McClymont et al., 2013), which brings together the risk framework and software of the former with the shear-stress-driven mobilisation mechanism of the *PODDs* model. The daily maximum shear stress, as calculated from a potentially unrepresentative hydraulic model, along with background water quality data and pipe material are to be used to estimate the maximum discolouration potential per pipe, although it is not clear how the background water quality data is incorporated into the analysis. The proposed method for calibrating the per-pipe *PODDs* parameters is using discrete samples and colour cards or turbidity data (Randall-Smith et al., 2011), although in associated work McClymont et al. (2013) assumed that the empirical coefficients k and b are constant over all pipes (with values of 2 and 1 taken from Boxall and Saul (2005)). The model outputs are as follows:

- The discolouration potential given the pipe material and daily maximum shear stress;
- A burst score, which is the average impact that a burst in each pipe has on mobilisation in all others;
- A closure score: the average impact that the closure of a pipe has on the shear stress and therefore on mobilisation in all other pipes in the network;
- An 'affected' score: the average impact on a main from the failure (burst and closure) of all other mains in the system).

The *DRM* has significant potential for managing discolouration risk within *DMAs* on a much wider scale than just *PODDs* in isolation (particularly with regards to structural failure/pipe closure impact assessments and scoring metrics for identifying the pipes that need to be kept

clean). However, as the [DPM](#) is understood to be based on [PODDS](#), it is subject to the same limitations and therefore also lacks a valid mechanism for modelling material accumulation.

2.2.8.7 Data-driven models of material accumulation

The [EPR](#) technique has been applied to [DWDS](#) flushing and asset data to derive expressions of relative discolouration material accumulation rates ([Mounce et al. 2014c](#); first mentioned in §2.1.16.1). The dataset that was studied featured values for up to thirteen predictor variables plus an estimate of the relative accumulation rate for 67 pipe lengths, each of which had been flushed more than once. The technique identifies polynomial expressions of a user-defined complexity that most accurately map a subset of predictor variables to relative accumulation rates. The variables found to have the greatest impact on accumulation rates were the background iron concentration, the pipe material, whether the pipe lies within a loop and the pipe volume, in that order. The accuracy of these models is believed to be a function of the size and selectivity of the studied dataset, the number of predictor variables for which data was consistently available, the degree to which the predictor variables are subjective (e.g. 'loopedness' or the presence of unlined [CI](#) upstream).

These results would be of greatest value to water providers if used as inputs to a modelling framework for estimating the impact of erosion and accumulation over time.

2.3 Conclusions

At present water companies are able to quantify discolouration risk in mains that are not self-cleaning at any given point in time using a scoring method such as the [RPM](#) or through calibrating a discolouration model such as [PODDS](#) or the [DPM](#). Short term predictions can then be made of how the system will respond to an increase in hydraulic conditions, although the use of exponent terms in [PODDS](#) can affect the accuracy of the model output. However, at present no discolouration model includes a valid mechanism for modelling the accumulation of cohesive layers, which has been shown to occur simultaneously over a range of shear strengths.

A model of both discolouration material erosion and accumulation has the potential to reflect changes in the state of the material at the pipe wall and in the bulk water over periods of weeks or months and therefore could be used to design, execute and potentially automate flow conditioning regimes, which would be of greatest value in critical trunk mains. Such a model may benefit from:

- Being built atop a 2D unsteady hydraulic model rather than a 1D steady state hydraulic model although the necessity for this has yet to be fully demonstrated;
- Allowing for gradated longitudinal variation in discolouration potential and accumulation rates;

Chapter 3

Aims and objectives

The principle aim of this PhD was to develop a model that reflects the latest understanding of how heterogeneous discolouration material accumulates on and erodes from the walls of DWDS pipes over time. Such a model would be of value to water providers for the medium- to long-term assessment and management of discolouration risk, with possible uses including predicting the turbidity response from future flow profiles and designing non-invasive cleaning programmes ('flow conditioning'). This model needed an accumulation mechanism that reflects the observation that material accumulates at shear strengths greater than the shear stress in a way that is independent of shear strength. This needed to be coupled with a compatible erosion mechanism that has the characteristics of the validated PODDS erosion model.

The objectives were as follows:

1. Identify the important behaviours that a new discolouration model should have, then use this information, along with the process knowledge presented in §2.1, to develop a conceptual model (presented in §4).
2. Identify a suitable approach that captures the conceptual model as mathematical formulations, then encode as software (presented in §4).
3. Verify that this implementation reflects the conceptual model and explore the model's sensitivity to the most pertinent factors and input data features (§5).
4. Explore methods and requirements for calibrating the model (§5).
5. Validate key new model assumptions and investigate the validity of the model as a whole and how well it generalises (§6).
6. Evaluate the practical functionality of the model and its potential value for DWDS management (§7).

Chapter 4

Model development

“Essentially, all models are wrong, but some are useful.”

Box (1987, p. 424)

4.1 Introduction

§2 presented the details of how discolouration modelling could be improved to better describe the relevant processes as we now understand them, in particular with regards to material accumulation on the pipe wall. Attention now shifts to the functional requirements for a new discolouration model. Different approaches that may conceptually satisfy those requirements are evaluated using the evidence and theories covered in §2. A conceptual model of discolouration material state, erosion and accumulation is then constructed from the most appropriate of these approaches. This is then encapsulated as a set of formulations, which are in turn, along with a water quality advection and mixing model, encoded as software. The chapter concludes by summarising how the proposed model can conceptually satisfy a set of ‘use cases’.

4.2 Requirements for a new discolouration model

4.2.1 Use cases

A *use case* is a systems and software engineering term for stakeholder goals that should be realisable upon completion of a project. Different stakeholders have differing needs: in environmental modelling, researchers and modellers are particularly interested in issues pertaining to model sensitivity and calibration whereas system operatives and strategic planners are interested in how such predictive models can be used for medium to long term environmental/asset management and in decision support.

An ideal set of use cases for a new discolouration model are listed below:

1. The model developer and modeller should be able to characterise the sensitivity of the bulk water turbidity in [DWDS](#) pipework to the continual processes of discolouration material

accumulation and erosion. For example, in fig. 4.1, the flow in a pipe is initially high then drops and remains at a low level for a period before quickly increasing in magnitude (first subplot). The amount of material that accumulates on the pipe wall (second subplot) whilst the flow is low depends on the accumulation rate, so the amount of erosion that results from the flow increase and therefore the resulting turbidity response also depends on the accumulation rate (the turbidity responses for several different accumulation rates are shown in the third subplot). Fig. 4.2 has a similar layout of subplots but shows how for a step increase in flow (first subplot) the turbidity response depends on the rate of erosion (third subplot);

2. The model developer and modeller should be able to use this knowledge to characterise the sensitivity of the latent state of the material at the pipe wall to accumulation and erosion, to refine understanding of the relationships between key variables and assess the spatio-temporal heterogeneity in discolouration processes. The second subplots in figs. 4.1 and 4.2 show how the state (in this case the amount) of material at the pipe wall compares for a range of different accumulation (fig. 4.1) and erosion (fig. 4.2) rates;
3. The model developer should be able to consolidate this understanding by developing and validating a time-series model of wall state (the amount and shear strength of material on the wall) and bulk water state;
4. The modeller and/or software should be able to calibrate this model for a given system. For example, in fig. 4.3, the pipe flow $Q(t)$ for the period t_0 to 'now' has been perturbed in steps (first subplot) so that over that period the amount of material at the pipe wall (second subplot) increases due to accumulation (when τ_a is lower) then decreases due to erosion (when τ_a increases). This results in an observed turbidity response ($T_{ds}(t)$ 'Obs' in the third subplot) that is a function of the amount of material at the wall at t_0 and the characteristics of the accumulation and erosion processes. Here the discolouration model can be calibrated at 'now' by minimising the dissimilarity between turbidity observations and time-series predictions ('Pred A' in the third subplot) over the period t_0 to 'now';
5. The modeller and/or software should be able to update a calibrated model given new time-series observations and make predictions and test scenarios given hypothetical time-series inputs. For example, given a calibrated discolouration model that describes the system state at 'now' in fig. 4.3 and two possible flow profiles for the period 'now' to ' t_2 ' (black and blue lines in the first subplot), the model should be able to accurately predict how the state of the material at the pipe wall (black and blue lines in the second subplot) and the bulk water turbidity (third subplot) differ for those two inputs. Here, the black future flow profile is predicted to cause a turbidity failure whereas the blue profile is not.
6. The modeller and/or software should be able to design optimal, ongoing flow conditioning (§2.1.17.2) programmes for maintaining an acceptable level of discolouration risk, particularly for trunk mains. For example, in fig. 4.4, discolouration modelling predictions and sensitivity analysis have informed the design of a flow profile (first subplot) where

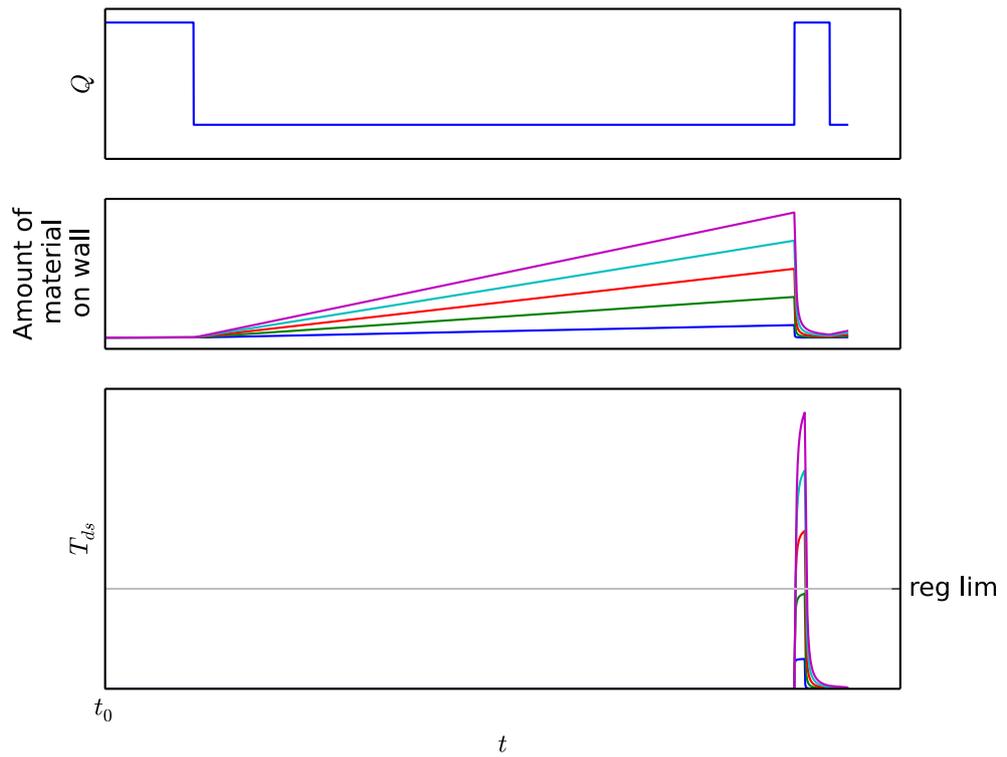


Figure 4.1: Exploring the sensitivity of the discolouration model output and latent internal state to the rate of material accumulation of discolouration material at the pipe wall. This is an illustrative figure produced using the model that is developed later in this chapter.

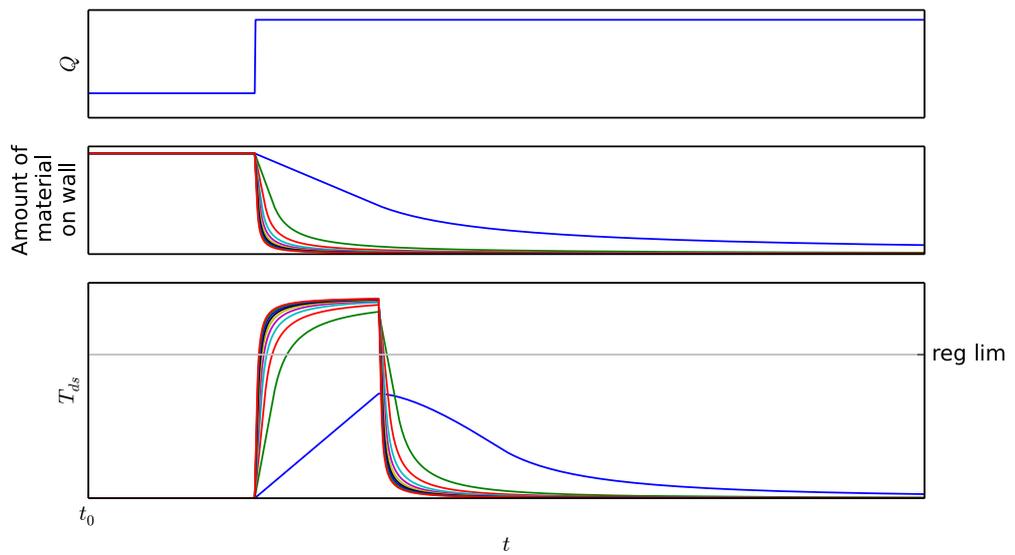


Figure 4.2: Exploring the sensitivity of the discolouration model output and latent internal state to the rate of material erosion of discolouration material from the pipe wall. This is an illustrative figure produced using the model that is developed later in this chapter.

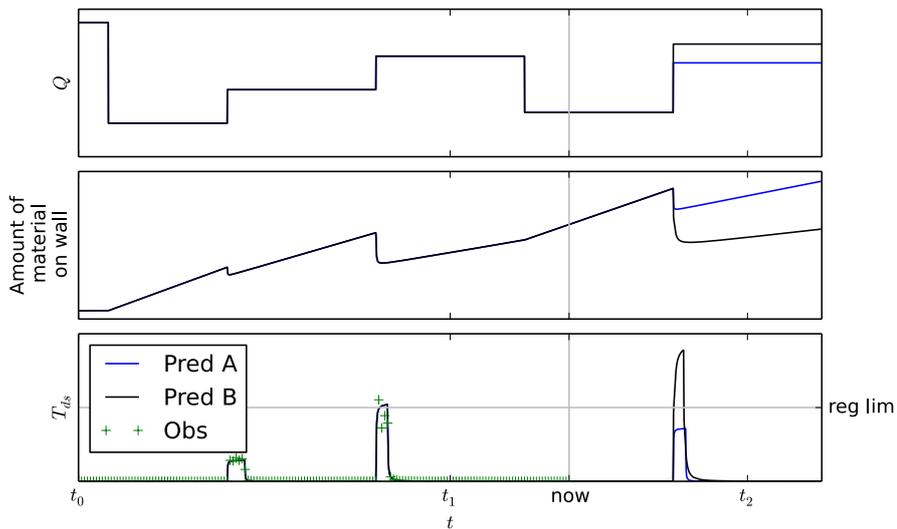


Figure 4.3: A flow time-series $Q(t)$ for a hypothetical pipe, beneath which is shown how the material erosion and accumulation that result from corresponding fluctuations in shear stress $\tau_a(t)$ (which is non-linear with flow) influence the total amount of material at the pipe wall (a latent quantity). This material erosion causes increases in turbidity $T_{ds}(t)$ that are detected at the downstream end of the pipe as shown; an illustrative regulatory limit 'reg lim' is depicted as a reminder that increases in turbidity are not always unacceptable. This is an illustrative figure produced using the model that is developed later in this chapter.

flow (and therefore shear stress) is periodically increased well above the prevailing level for a short duration. The material that accumulates whilst the flow is at the lower level is eroded from the pipe wall when the flow is increased to the higher level (second subplot). This erosion causes increases in turbidity (third subplot) that are always well below the regulatory limit. The magnitude and frequency of the regular flow increases periodically remove material from the pipe wall so as to ensure that neither those increases nor any other deliberate or unanticipated increases of up to the same magnitude can cause regulatory failures with regards to turbidity.

4.2.2 The ability of **PODDS** to satisfy the use cases in §4.2.1

Use case #1: the way in which **PODDS** predicts changes in bulk water turbidity given a calibrated model and a time series shear stress profile has successfully been validated for over periods short enough for the effects of regeneration to be negligible (§2.2.8.4). Erosion process parameters have been shown to be transferable between sites within similar characteristics, although transferability is limited by the site-specific relationships between turbidity and **TSS**. However, the need for and sensitivity to a power term (n ; §2.2.8.5) that prevents dimensional analysis indicate that the conceptual representation of the relationship between turbidity and eroding forces could be improved upon.

The ability of **PODDS** to accurately model over longer periods is limited by its material accumulation model not reflecting the current understanding of accumulation processes (§2.2.8.5).

Use case #2: **PODDS** can model the change in bulk water turbidity given a series of step changes in shear stress, indicating that the encapsulated relationship between internal state

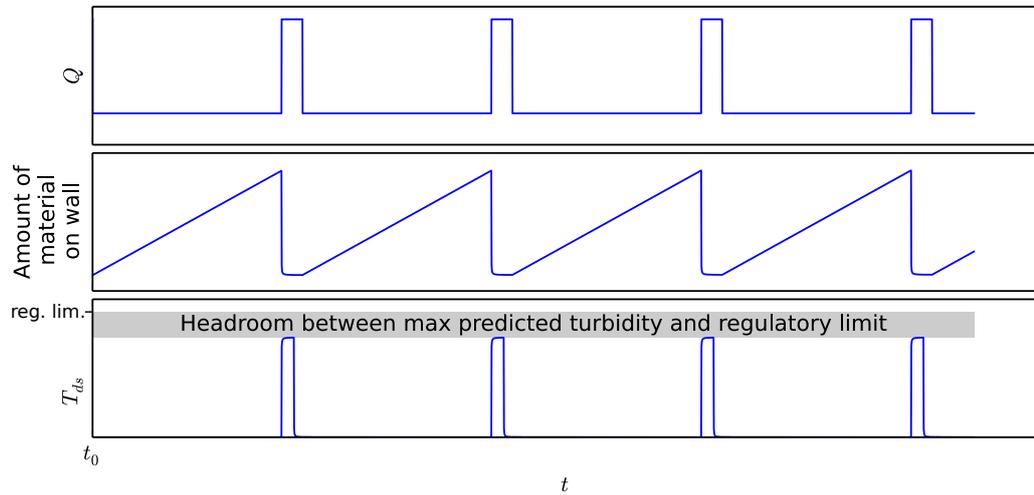


Figure 4.4: A flow time-series $Q(t)$ for a hypothetical pipe, beneath which is shown how the material erosion and accumulation that result from these fluctuations in flow influence the total amount of material at the pipe wall. Here the periodic flow profile has been designed to limit the risk of unanticipated flow increases causing unacceptable increases in turbidity $T_{ds}(t)$. This is an illustrative figure produced using the model that is developed later in this chapter.

and eroding forces is valid (at least over periods of hours or days where the effect of material accumulation is negligible). The key concept here is that the scalar shear strength of material on the wall increases linearly (or near-linearly) whilst there is an excess shear stress (fig. 4.5). It should be noted that more work could be done to assess whether updates to the amount and strength of material at the pipe wall are still valid if increases in shear stress are not totally conditioning i.e. if an increase in shear stress is sufficiently short-lived so as not to remove all material weaker than that shear strength.

The sensitivity of the **PODDS** model's internal state to material accumulation has not been validated. **PODDS** assumes that the scalar shear strength of material on the wall increases linearly whilst there is an excess shear stress (fig. 4.6) i.e. stronger material accumulates before weaker material. However, observations indicate that material accumulates at all shear strengths beneath τ_a *simultaneously* (figs. 4.6b and 4.6c; see also §2.1.13.4), hence the need for a more accurate means for modelling accumulation.

This discrepancy between model predictions and observed behaviour is illustrated in fig. 4.7.

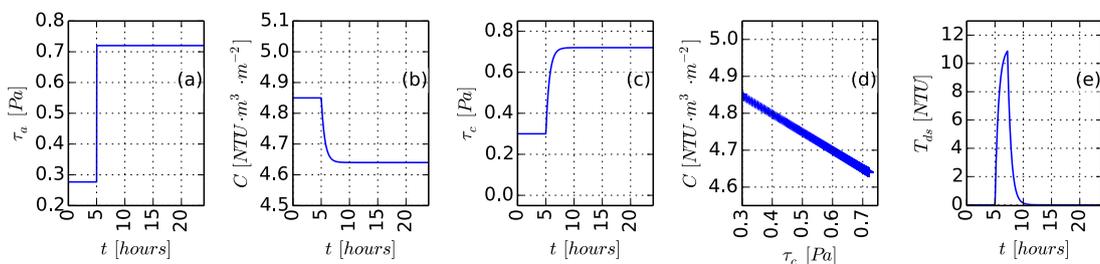


Figure 4.5: An illustration of the **PODDS** erosion model: (a) A step increase in shear stress τ_a at $t = 5$ hours within a particular pipe from 0.30 Pa to 0.72 Pa ; (b) The amount of material C on the pipe wall then decreases asymptotically whilst (c) the scalar shear strength of that remaining material τ_c increases asymptotically; (d) the relationship between C and τ_c remains linear throughout this period of erosion; (e) The resulting turbidity response as seen at the downstream end of the pipe. This figure was produced using a software implementation of the **PODDS** model.

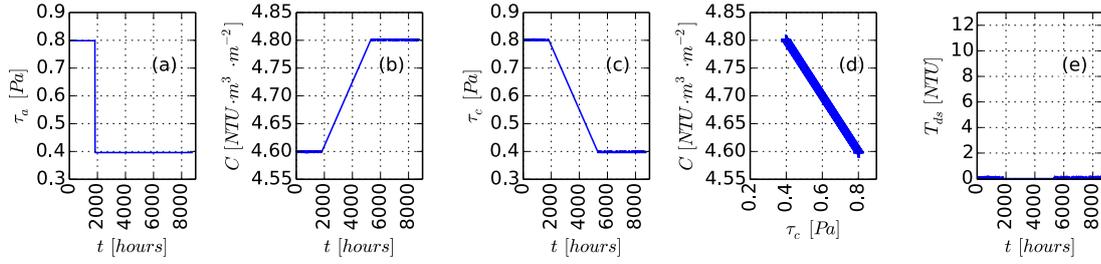


Figure 4.6: An illustration of the **PODDS** material accumulation model: (a) A step decrease in shear stress τ_a at $t = 0$ within a particular pipe from 0.80 Pa to 0.40 Pa (flow Q drops from $1.36 \text{ L} \cdot \text{s}^{-1}$ to $0.96 \text{ L} \cdot \text{s}^{-1}$); (b) The amount of material C on the pipe wall then increases linearly and (c) the scalar shear strength of that remaining material τ_c decreases linearly to τ_a -limited levels; (d) as per the erosion case, the relationship between C and τ_c is linear (in the **PODDS** formulations they are related by scalar gradient term k); (e) No change in turbidity is seen at the downstream end of the pipe (any reduction in turbidity due to material accumulation is assumed to be negligible). This figure was produced using a software implementation of the **PODDS** model.

Here we see a short-lived increase in τ_a which erodes all material with $\tau < \tau_4$. τ_a then drops to and remains at τ_1 for some time, allowing for the accumulation for material with $\tau \in [\tau_1, \tau_4]$. Field and lab experiments indicate that material accumulates at all shear strengths simultaneously yet **PODDS** predicts that stronger material will accumulate before weaker material (see the plots of material quantity versus τ for t_2 and t_3 in fig. 4.7). This difference in behaviour is not always noticeable: if a maximal material quantity has accumulated and τ_a is then increased in steps of sufficient duration to exhaust the corresponding shear strength ranges then the validated **PODDS** erosion model can predict how material erodes (see $t \in [t_4, t_7]$); however, if these increases in τ_a had occurred whilst material with $\tau \in [\tau_1, \tau_4]$ was in the process of accumulating then the prediction per τ_a step would be far less accurate, as can be seen in fig. 4.8.

To date, much of the understanding of the relationships between key variables relating to discolouration processes has been gained from the analysis of field data rather than **PODDS** modelling; this is not unexpected given the complexity of the underlying processes and the use of a grey-box empirical model. Regeneration rates are anticipated to vary with many time-dependent factors but to date furthering the understanding of the time-dependence of regeneration using modelling rather than just labour-intensive fieldwork has been limited by **PODDS** lacking a valid material accumulation mechanism.

Another potentially influential factor that could be explored more easily using modelling rather than just fieldwork is longitudinal distance: there is evidence to suggest that there is some longitudinal heterogeneity in discolouration processes (§2.2.8.5) but the reference **PODDS** implementation does not allow for this to be modelled directly as pipe properties and model parameters are assumed to be constant with length and with unit area of pipe wall respectively. However, longitudinal heterogeneity in model parameters and state along a pipe can be approximated by ‘daisy-chaining’ a sequence of models (§2.2.8.5).

Use case #3: as mentioned previously **PODDS** provides a validated means for modelling the change in wall state and bulk water turbidity due to material erosion but not accumulation. It can therefore only be used for modelling one more short periods over which accumulation is assumed to be negligible and/or the amount of material at the pipe wall is maximal and has

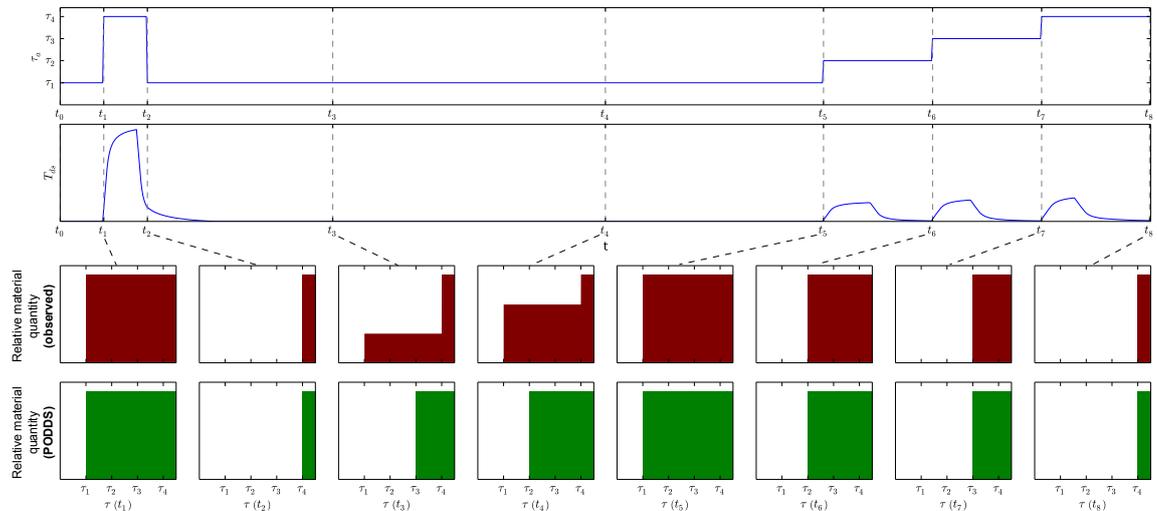


Figure 4.7: An illustration showing how, for a given shear stress (τ_a) profile, the relationship between the amount of material per shear strength τ differs over the period t_2 to t_5 between the **PODDS** model (shown in red) and the understanding garnered through field and lab experiments (shown in green) due to the way in which **PODDS** models material accumulation. However, after t_5 the validated **PODDS** erosion model produces predictions that tally with observations. Note that the time axis is not to scale.

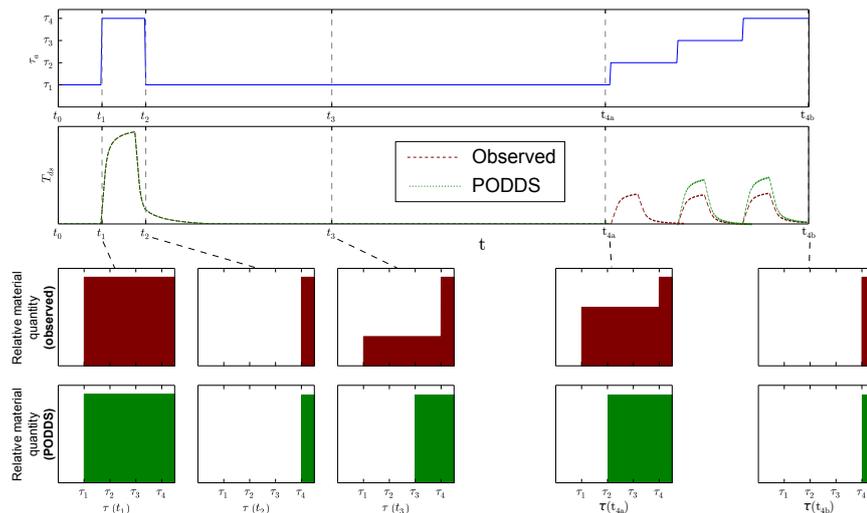


Figure 4.8: If the step increases in flow shown in fig 4.7 had occurred sooner during the modelled time period then, unlike in fig 4.7, the observed turbidity response differs from the turbidity response predicted by **PODDS** due to the way in which material accumulation at the pipe wall was modelled over the period t_2 to t_{4a} . Again, note that the time axis in this illustrative figure is not to scale.

reached a steady state.

Use case #4: **PODDS** is often partially calibrated using parameters from similar systems then a comparison of the turbidity response from a small flow perturbation to turbidity predictions is made to refine those parameter values (§2.2.8.4). Any successor to **PODDS** may be easier to calibrate if its erosion model is not highly sensitive to a power term (such as **PODDS** n) and if the new model is not under-determined.

Use cases #5 and #6: a new model of material accumulation is required to accurately model using a single parameter set a number of discolouration events over a continuous period long enough for regeneration effects to be significant. To accurately model a notable increase in flow following a period of regeneration, **PODDS** first requires recalibration using turbidity and flow data from during a small controlled increase in flow due to the **PODDS** regeneration mechanism not being valid: accurate medium- to long-term modelling is not possible using a single parameter set. An improved regeneration model would greatly benefit the design of flow conditioning programmes, particularly if coupled to an erosion model that is not highly sensitive to a power term (such as n in **PODDS**).

In summary, what is required is a discolouration model with an erosion model that is similar to that of **PODDS** but coupled with a material accumulation model where material accumulates at all shear strengths below τ_a simultaneously and ideally without an equivalent to the n power term.

4.3 Conceptual model

4.3.1 Introduction

This section presents a conceptual design for a model that satisfies the majority of use cases listed in §4.2.1 by drawing upon theories and evidence from previous discolouration process studies (§2.1) and modelling work (§2.2.7) and by providing possible mechanisms for addressing the issues listed in §4.2.2. The conceptual model is comprised of two sub-models:

Wall state sub-model Describes how the amount of material at the pipe wall varies over a range of shear strengths and changes over time due to material erosion and accumulation;

A material transport sub-model Describes how material that is eroded from the pipe wall becomes entrained as a fine suspension in the bulk water and undergoes advection and mixing.

4.3.2 Treating shear strength and material quantities as distributions and functions of distributions rather than simple scalars

If the amount of material bound to the wall at a given time varies with strength in a non-linear, potentially non-differentiable way (illustrative fig. 4.7) then this suggests that material quantity should be a function of a *range of shear strengths* rather than a scalar shear strength, with erosion and accumulation determining how this function changes with time. Treating shear

strength as a scalar which is linearly (and inversely) proportional to the amount of material at the wall (e.g. as per §2.2.8.3) results in a model of shear strength versus quantity that is at times invalid and can yield inaccurate turbidity predictions (again, see fig. 4.7).

The erosion of wall-bound material has been shown to occur when and only when the shear stress, τ_a , exceeds the shear strength, τ , of some of that material (§2.1.12.2); this knowledge has been encoded in PODDS' validated material erosion model (§2.2.8.3). If the material at the wall has a *shear strength range* (rather than a scalar shear strength) at any given time, then at any moment there could be some material at the wall that is weaker than the shear strength and some that is stronger, hence PODDS's rule for deciding if and when erosion is occurring must be revised. Here it is postulated that the wall-bound material for which $\tau \leq \tau_a$ is *eroded* whilst simultaneously material for which $\tau_a < \tau < \tau_{max}$ *accumulates* on the pipe wall, where τ_{max} is the maximum possible τ with which material binds to the wall). Finite values of τ_{max} have been detected in certain pipe types such as PE pipes but not others such as unlined CI (§2.1.15). τ_{max} is comparable with the PODDS parameter $\tau_{ultimate}$ (§2.2.8.4).

Material erosion is assumed to be driven solely by the steady shear stress as it has yet to be conclusively demonstrated that the unsteady shear stress has a predictable and significant impact on discolouration material erosion (§2.2.8.5). Also, the data requirements for establishing a one-dimensional steady-state hydraulic model are far less than for a two-dimensional truly dynamic model (§2.2.8.5).

The model proposed in this chapter shall henceforth be known as the *Variable Condition Discolouration Model (VCDM)* as it tracks the amount of material at the pipe wall for a vector of shear strengths (rather than a scalar shear strength).

4.3.3 Modelling material accumulation

Accumulation has been shown to be a process that occurs throughout all DWDS pipework (§2.1.13.3) for $\tau_a < \tau < \tau_{max}$. As mentioned in §2.1.14, there is uncertainty as to whether the accumulation of material with shear strengths within this range continues until the pipe is hydraulically impassible or whether a steady state is reached before then. A material accumulation model should either assume that accumulation is truly continual or is bounded without there being further evidence to support greater model complexity. Evidence for and against bounded and continual accumulation is limited: many of the field and laboratory studies of material accumulation that used flushing have been of insufficient duration to see if/when the quantity of wall bound-material eventually reached the level it was at prior to the first flush. However, continuous monitoring has shown that in a number of cases the magnitude of background turbidity variations and frequency of turbidity events decreased following flushing and invasive cleaning operations, suggesting that material levels may have been maximal prior to the interventions and so following cleaning interventions material can start passing from the bulk water the the pipe wall again. Also, no DWDS mains have been found to be completely impassible due to the build-up of erodible discolouration material (scale may accumulate over time but is most likely sufficiently consolidated as to not be erodible by achievable hydraulic forces, plus the shear stress would increase as the diameter decreases, limiting accumulation). Given this evidence, it can be reasonably assumed for the purposes of this model that the

amount of discolouration material that can accumulate with a given shear strength has an upper bound.

This bound is assumed to be pipe-dependent as hydraulics have been shown to be one of the most influential factors that controls the amount of material at the wall (§2.1.16.1). However, there is insufficient evidence to assume that this threshold is time-varying. If the amount of material with shear strength τ has a time-invariant upper bound then the accumulation process can be considered to have a *relative* rather than absolute impact on that amount. Henceforth, the term *regeneration* is used in preference to *accumulation* to emphasise the relative nature of the accumulation process, with a pipe-specific yet time-invariant, linear scaling factor being used to translate between relative quantities of discolouration material and quantities in units that are a function of *NTU*.

Note that **PODDS** also places a calibration-specific upper bound on the amount of material that can bind to the pipe wall (C_{max} ; see eq. 2.15) but this is a slightly different idea to that proposed here as under **PODDS** the amount of material at the wall C can only reach C_{max} if the shear stress remains at 0 *Pa* for some time. Here, the amount of material over a high strength range could be maximal whilst the shear stress is far from zero.

The (relative) regeneration rate is a function of many factors (§2.1.16), with many of these being inter-dependent (fig. 4.9). Modelling this rate could in theory be achieved using a process model of regeneration, but this would require a sufficiently-detailed understanding of these inter-dependencies (a set of physical, chemical and biological reactions per distinct environment, with the stoichiometry and kinetics/equilibria being characterised for each biological and chemical reaction) so as to be difficult and impractical (§2.2.7.6). For the purposes of this model regeneration is initially considered to be a simple, empirical rate that is a scalar and time-invariant (eq. 4.1).

$$\beta_r \left[s^{-1} \right] \quad (4.1)$$

As further data and process knowledge becomes available, this rate-based model can then be refined through decomposition into empirical and process-aware sub-models of accumulation mechanisms (such as corrosion and biofilm growth; Husband and Boxall (2011)) using a hybrid modelling approach (fig. 4.10 and §2.2.2) and data mining techniques such as **EPR** (§2.2.8.6), with some of the sub-models being time-dependent.

In this model, the regeneration rate is initially assumed to be invariant with respect to:

Time Field and laboratory studies have shown material accumulation to be continuous and repeatable (§2.1.13.3). There is mixed evidence of regeneration being notably non-linear with time, although it is expected that, as temperature has a significant effect on growth rates (§2.1.16.2), that further evidence of seasonal variation in regeneration rates will result from future field studies. The regeneration rate could be made temperature-dependent if it can be proven to be necessary, but should initially be assumed to be constant for the sake of parsimony. Note that the regeneration of a given shear strength is considered to be linear up to the point where the quantity of material at that shear strength is maximal, after which time there can be no further material accumulation without there first being some erosion.

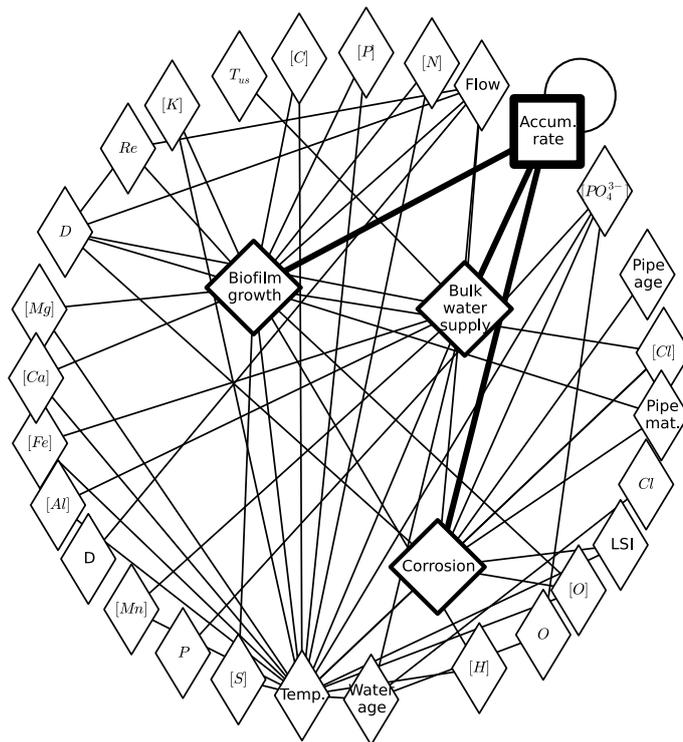


Figure 4.9: A dependency graph for a subset of the myriad factors that influence discolouration material accumulation rates. Here, D is internal pipe diameter and LSI is the Langlier Saturation Index (a measure of calcite concentration; see McNeill (2000)).

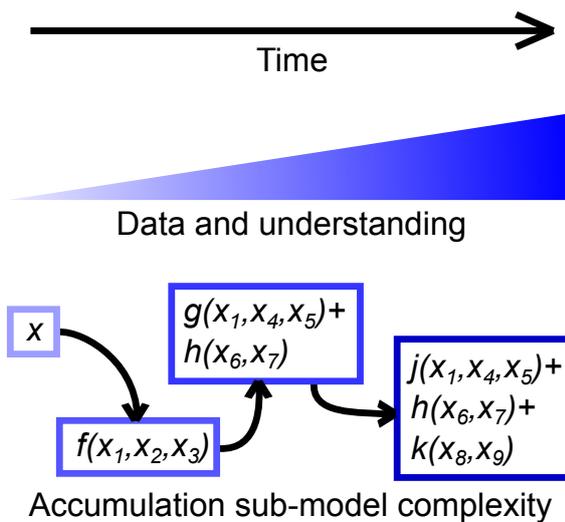


Figure 4.10: A model of the material accumulation rate can be refined over time as more data is analysed, as illustrated in this conceptual diagram; increasing complexity should only be added when it can be justified,

Material shear strength The amount of regeneration over a given time period was believed to be constant with shear strength in one field study (§2.1.13.4; based on visual interpretation of bar graphs rather than statistical analysis) but the evidence from other studies is inconclusive. However, without having significant evidence to the contrary, the regeneration rate must be assumed to be constant with shear strength.

Angle to the pipe invert The regeneration rate is assumed to be constant around the pipe circumference as a) the accumulation of discolouration material cannot principally be driven by gravitational settling, b) no significant variation in biofilm quantity around the pipe circumference has been detected and c) generic discolouration material has been directly observed around the circumference of a transparent pipe that was operated at a moderate flow rate (§2.1.9). There is limited evidence of more material collecting at the invert rather than the soffit under certain conditions such as low flows (§2.1.9) but this material is most likely comprised of atypically large particles and so will have a lesser impact on discolouration given that turbidity is most sensitive to small particles (§2.1.8).

Longitudinal distance Material accumulation is not highly localised: material has been shown to develop along the length of all mains (§2.1.16.6). There is some evidence of material quantity varying along lengthy mains (§2.1.16.6), possibly due to faster regeneration rates, but this is difficult to model directly due to the complexity of the inter-related processes and the number of influential variables (chemical and microbial oxidation/deposition; concentration gradients; temperature). Also, more material passing from the bulk water to the wall means that particulate concentrations will be lower downstream, so regeneration rates may be lower. However, it has not yet been demonstrated that regeneration rates vary *significantly* with distance along the length of a main with otherwise uniform properties, so material regeneration is therefore to be assumed to be constant with longitudinal distance along each pipe and, as it is also invariant around the pipe circumference, constant per unit of wall area.

The shear strength of a quantum of wall-bound material, like the regeneration rate, is assumed to be time-invariant.

The behaviour of the described regeneration model over a time period Δt is shown in fig. 4.11. No material accumulates over Δt with a shear strength less than the shear stress, the latter being constant over Δt . The amount of material at all shear strengths greater than the shear stress is not maximal, so material over that shear strength range can accumulate over Δt . The amount of accumulation over Δt is constant for that shear strength range.

4.3.4 Modelling erosion

4.3.4.1 The need for a new erosion model

The proposed discolouration model requires a material erosion mechanism to complement the regeneration model described in §4.3.3. PODDS' validated material erosion model (§2.2.8.3) cannot be coupled to this regeneration model due to the latter having the unprecedented requirement that the quantity of wall-bound material be tracked for a range of shear strengths i.e. neither quantity nor shear strength are scalars.

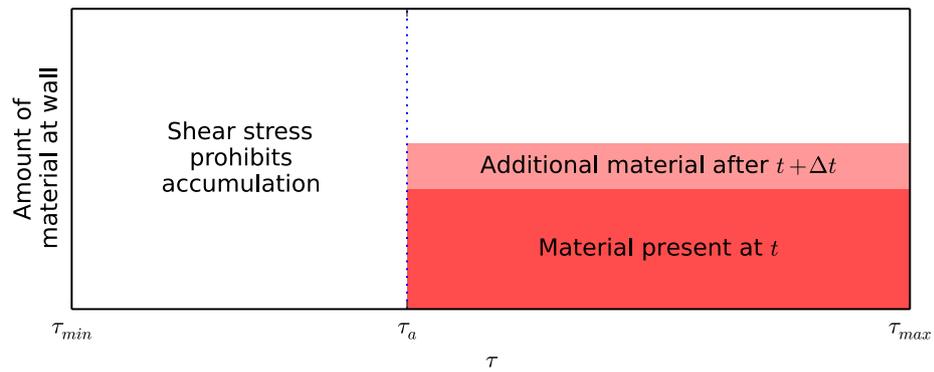


Figure 4.11: An illustration of the proposed model of discolouration material regeneration, showing the accumulation per shear strength over time period Δt . Note that here the shear stress τ_a is constant over Δt .

4.3.4.2 Shear stress as an eroding force

It is assumed that discolouration material has a fairly narrow and consistent particle size distribution (PSD) and density distribution at a site-specific level; this argument is supported by the findings of a number of (but not all) relevant field studies (§2.1.4). If discolouration material were loose sediment then it would be re-suspended when a critical shear stress were exceeded.

It is also assumed that discolouration material cohesively binds to the pipe wall and none settles on the pipe invert as studies have shown that a portion of material is eroded from around the circumference each time the shear stress τ_a is increased in steps above the prevailing level (§2.1.12.2 and figs. 2.11 and 2.12), which would not be the case if the material was comprised of loose deposits (§2.1.9). The amount of erosion is therefore driven by the *excess shear stress* (the amount by which τ_a exceeds the shear strength of some material). This behaviour is illustrated in fig. 4.12; note that in this figure it is assumed that τ_a does not exceed the magnitude required to remove all material from the pipe wall (see §2.1.15) over the shown time period.

A consequence of the amount of erosion being driven by the *excess* shear stress rather than *absolute* shear stress is that a large increase in shear stress above prevailing levels has the potential to cause more material mobilisation than a smaller increase above a much higher prevailing shear stress (as seen in field data (§2.1.12.2 and fig. 2.9) but here also illustrated conceptually in fig. 4.13).

The erosion *rate* is also considered to be a function of the excess shear stress: the validated **PODDS** material erosion model along with validated models of the erosion of estuarine muds all treat the erosion rate as a function of the excess shear stress (§2.2.8.3), although in the latter field the linearity of this relationship varies between models.

Note that flow not only influences shear stress and material erosion but also the bulk velocity and material advection.

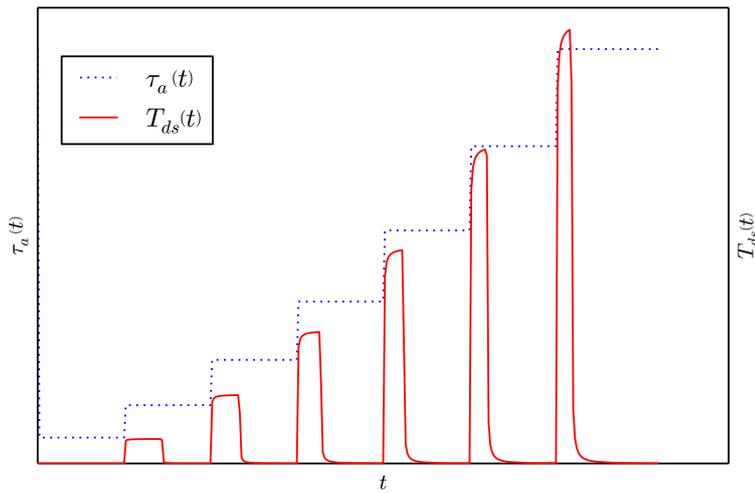


Figure 4.12: Discolouration material must cohesively bind to the pipe wall given the way erosion is repeatedly seen (as downstream turbidity, T_{ds} , when the shear stress, τ_a , is increased above the prevailing level in steps. This is an illustrative figure produced using the model that is developed later in this chapter.

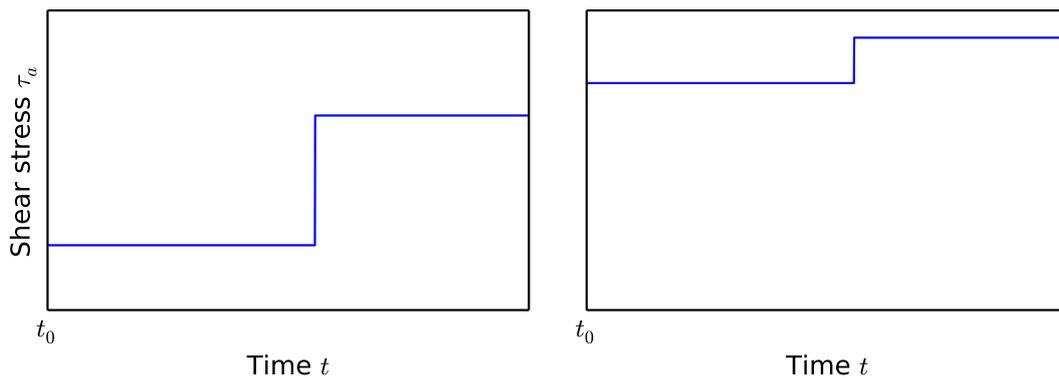


Figure 4.13: If material erosion is driven by the excess shear stress rather than absolute shear stress then the temporary increase in shear stress shown on the left-hand side of this conceptual diagram would likely result in more mobilisation than the smaller increase at higher values shown on the right (assuming that the only difference between these two environments are the prevailing and temporary shear stresses).

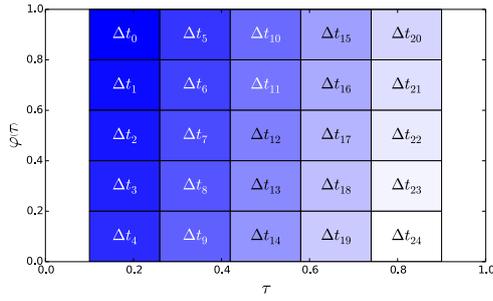
4.3.4.3 Sequential or simultaneous or layer erosion

The erosion of material for which there is an excess shear stress could be modelled in several different ways should the quantity of material at a given time be tracked over time for each of a range of shear strengths. One possibility is that for a large step increase in shear stress, the material weaker than the new shear stress could erode sequentially from weakest-bound to strongest-bound. This concept is illustrated for discrete quanta of material in fig. 4.14a. Here the rate of erosion is assumed to be constant with time t . An alternative 'variable-strength' erosion model is one where all material with shear strengths between the old and new values of shear stress (following a step increase) erode *simultaneously* at a time-invariant rate (fig. 4.14b).

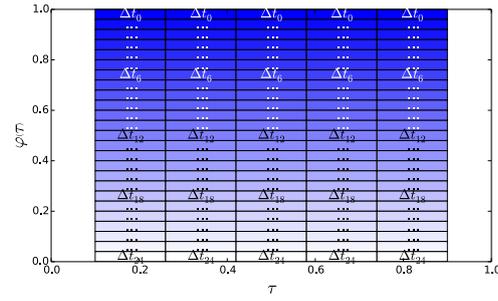
The sequential erosion model shown in fig. 4.14a is arguably more intuitively accurate for inorganic material at the micro scale: it could be conceptualised that the inorganic material furthest from the pipe's internal surface is most weakly bound and so erodes before more strongly-bound material. The sequential model assumes a homogeneity of discolouration composition over the pipe's internal surface, with weaker material always atop stronger material. In reality, the thickness of discolouration material has been seen to vary at the macro scale and biofilms are understood to vary spatially in quantity and structure at the micro scale (§2.1.11): chunks of biofilm can slough off into the bulk water (§2.1.10), indicating that cohesiveness varies in a complex way with depth into the biofilm and the weakest parts of biofilms are not always furthest from the pipe wall. The quantity versus shear strength profile therefore most likely also varies across the pipe wall and so erosion is unlikely to be sequential with respect to shear strength (from weakest to strongest) when considering the total material mobilised from say 1 m^2 of wall.

Further evidence of spatial heterogeneity in biofilms was found by undertaking new statistical analysis on biofilm distribution data from Fish (2014). In that study, three near-identical recirculating pipe rigs (HDPE; $\varnothing = 79.3\text{mm}$) were operated with differing diurnal flow patterns for 28 days, with replaceable sections of pipe wall being extracted for confocal laser scanning microscope (CLSM) imaging of what was growing there at the start and end of that period. By using multiple stains and imaging over a range of focal distances, Fish (2014) produced three-dimensional representations of cell, carbohydrate and protein distributions. In total, twenty-five fields of view were imaged at each end of that period and for each a measure of volume per biofilm component was produced (see Fish (2014) for further details of the method).

To assess whether the spatial heterogeneity of biofilms increases over time, the distributions of 25 volumes for each of cells, carbohydrates and proteins have been compared between the start and end of the 28-day period for one of the three pipe loops (which had been operated at a steady 0.4L/s over that time). The Brown-Forsythe test (NIST, 2013, §1.3.5.10: Levene Test for Equality of Variances) was used to test the null-hypothesis that the two (non-normally distributed) samples had the same variance. After removing an outlying biological sample (the data for one field of view from the end of the period) the hypotheses that the variances of the 'day 0' and 'day 28' samples were the same for cells (Brown-Forsythe W statistic = 9.021), carbohydrates ($W = 6.102$) and proteins ($W = 5.455$) were all rejected as the W statistics exceeded the critical value of 4.042 at the $P = 0.05$ level. Histograms of the samples showed that spatial variance in biofilm volume had increased in time as the biofilm had grown. The model of material erosion where different shear strengths can erode simultaneously is therefore



(a) A model of material erosion where material is eroded sequentially from the weakest to the strongest for at all shear strengths $\tau < \tau_a$, at a rate that is constant with time t .



(b) A model of material erosion where the erosion occurs at all shear strengths $\tau < \tau_a$ simultaneously at a rate that is constant with time t and shear strength τ .

Figure 4.14: Discolouration material erosion could be modelled by considering material to erode from the weakest to the strongest sequentially or erode at all shear strengths simultaneously. The simplest instances of both approaches are presented in these conceptual diagrams: it is assumed that there is initially maximal accumulation at $\tau > 0.1$ then there is a step increase in τ_a from 0.1 to 0.9. Here, both the label in each rectangle and the rectangle shading indicate in which timestep that the corresponding relative portion of material with a corresponding shear strength erodes, with darker shading denoting that a portion erodes earlier following the step increase in shear stress.

intuitively preferable as it can account for this heterogeneity.

Designing tests to differentiate the aforementioned simultaneous and sequential erosion models and validate one of them is non-trivial as in many situations the model outputs should be similar: for situations where an increase in τ_a always results in complete erosion over the affected τ range the models are undistinguishable (fig. 4.15, example A). However, rapid variation in τ_a may only cause the partial erosion of material within a τ range; subsequent erosion could allow the models to be distinguished, either by studying the shape or number of *material flux* responses (fig. 4.15, examples B and C respectively). Here, material flux is the product of sampled turbidity and flow at each timestep (as defined in eq. 2.14 and justified in §2.1.13.2).

In practise these distinctions may not be easy to observe:

- Short lived changes in τ_a that result in partial rather than exhaustive erosion of the material within a shear strength band might not be easy to engineer, particularly if the erosion rate is sufficiently fast that partial erosion is difficult to achieve: network operators must transition between flow levels in a controlled manner to avoid generating transient pressure waves (§2.1.7.3). For this reason, a sequence of rapid changes in τ_a is unlikely to occur under standard operational conditions and therefore it is less important to distinguish the two models in this case. However, such rapid changes in flow may result from a structural failure and the subsequent valve operations required for isolation, ensuring continuity of supply and repair, so there is still a need to select the most accurate erosion model.
- Advection and mixing processes may have a dominating effect on the observed turbidity response, obscuring the aspects of the signal that could differentiate the two models.
- τ_a must be controlled so that maximal regeneration can be reached over a wide shear strength range prior to the start of the tests, something that is not easy to achieve given that it may take up 1.5 to 4 years to reach maximal regeneration in DMA pipes (§2.1.14).

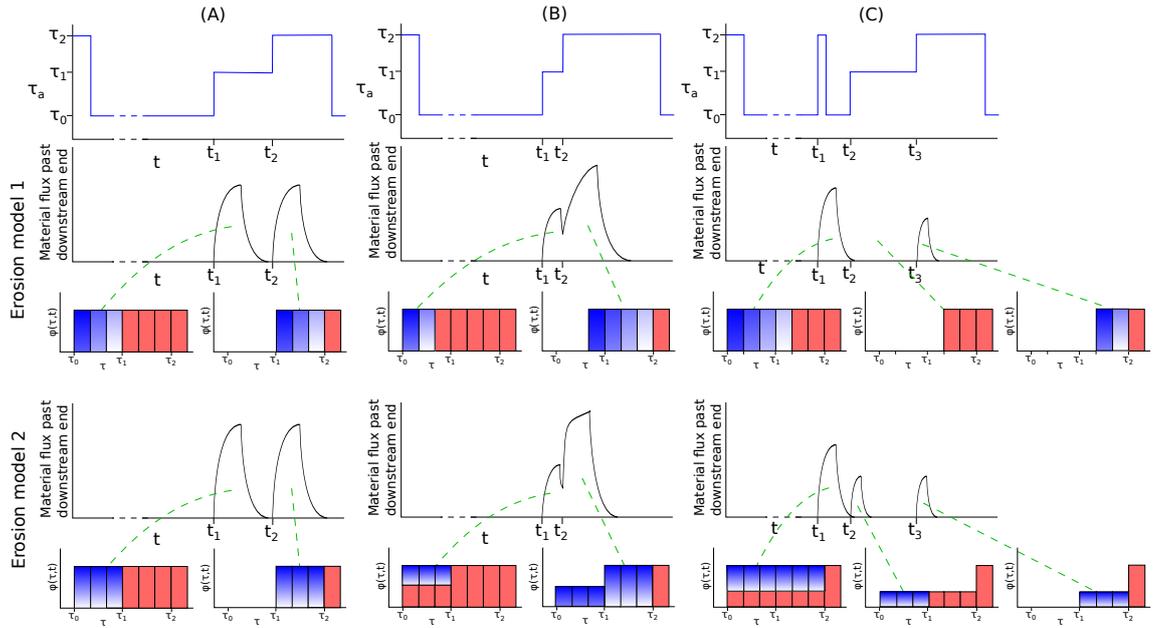


Figure 4.15: Three example shear stress $\tau_a(t)$ profiles (A, B, C) along with illustrations of the likely material flux response and change of state wall predicted by two possible material erosion models. Here, material flux is the product of sampled turbidity and flow at each timestep (as defined in eq. 2.14 and justified in §2.1.13.2). In the illustrations of wall state at particular times, red shear strength bands are not affected by the most recent increase in τ_a but blue ones are, with darker blue material eroding more quickly than lighter blue material. Erosion model 1 erodes material sequentially from weaker τ to stronger (see also fig. 4.14a) whilst erosion model 2 erodes all material for which there is an excess shear stress at the same rate, regardless of τ (see also fig 4.14b). In case A, the shear stress initially drops significantly to allow regeneration to occur at weaker τ ; maximal regeneration is then reached before τ_a increases in steps of equal magnitude back up to the original level, where each step is of sufficient duration for complete erosion to occur over the affected τ range. Case B is similar but the first of the two τ_a step increases is too short to fully erode the affected τ range. Case C is also similar to case A but the first τ_a increase is prefixed by a high-magnitude but short increase in τ_a that only partly erodes the affected τ range.

4.3.4.4 More complicated erosion models

An erosion model where material erodes at all shear strengths simultaneously as opposed to sequentially from weak to strong may be intuitively preferable (§4.3.4.3) but there are several forms that such a model could take:

- The simplest form is one where the erosion rate is constant with shear strength and time, as described previously (fig. 4.14b);
- The erosion rate could be variable with shear strength but constant with time. This permits the weaker-bound material to erode more quickly (fig. 4.16), in keeping with the interest in the rate of erosion being a function of the excess shear stress.
- Conversely, the erosion rate could be variable with time but constant with shear strength (e.g. fig. 4.17, where the erosion rate per shear strength is actually a function of the relative amount of material with that shear strength).

4.3.4.5 Chosen erosion model

These three erosion models are preferable to the simpler ‘simultaneous’ erosion model (fig. 4.14b) as the aggregate rate of material erosion slows over time in a similar manner to an existing validated erosion model (§2.2.8.3) rather than ceasing abruptly. However, the shear-strength-variable yet time-invariant erosion model (fig. 4.16) has been chosen for use with the regeneration model presented in §4.3.3 for two reasons:

- It allows for non-serial layer erosion so can account for spatial heterogeneity in material composition;
- It can ensure that the erosion rate is a function of the excess shear stress.

Although the chosen erosion and regeneration conceptual models best encode the current understanding of material erosion and accumulation they may not actually describe the workings

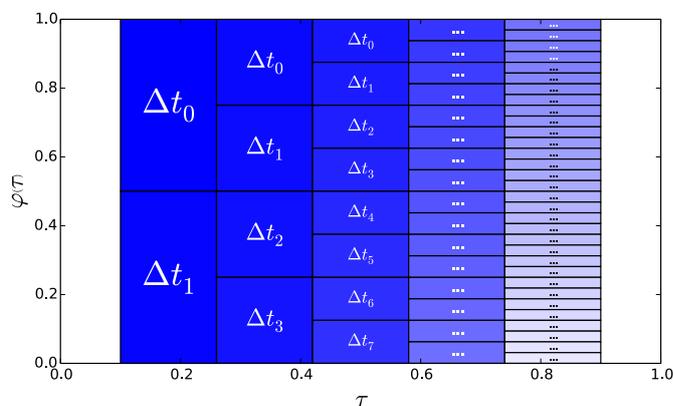


Figure 4.16: Illustration of an erosion model where all shear strengths erode simultaneously at a rate that is constant with time for each shear strength but, unlike the model shown in fig. 4.14b, varies between shear strengths, with the weaker material eroding more quickly.

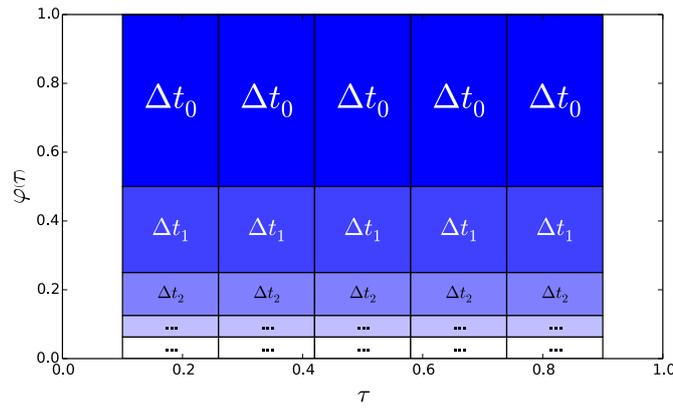


Figure 4.17: Illustration of an erosion model where all shear strengths erode simultaneously at a rate that is constant with shear strength but, unlike the model shown in fig. 4.14b, varies in time, with the rate of erosion for all shear strengths slowing over time.

of the underlying latent physical, chemical and biological processes particularly accurately. However this is unimportant if the model generates acceptable outputs for the range of inputs of interest and is of minimal complexity.

Finally, as with the conceptual model of material regeneration, the erosion process is assumed to be homogeneous with respect to longitudinal distance and the angle to the pipe invert for a given pipe.

4.3.5 A complex system state

Combining the chosen material erosion and regeneration conceptual models results in a model where the material per shear strength can erode or regenerate independently. Over time, this can lead to a complex, non-linear relationship developing between material quantity and shear strength as a result of fluctuations in the shear stress (e.g. fig. 4.18). Henceforth, the term *wall state* will be used to mean the relative amount of wall-bound material at each tracked shear strength.

An empirical model is preferable to a process-aware model given the large number of factors that influence discoloration-related processes. The material quantity versus shear strength relationship is considered to be sufficiently complicated for it not be possible to learn model structure solely from data through empirical machine learning. Using Takens' theorem for 'recovering' latent model state using time-series inputs (§2.2.2) is also inappropriate as observed flow and turbidity signals are not free of noise. The most suitable approach is believed to be a grey-box (§2.2.2) empirical model consisting of anthropogenic formulations that have been informed by field and lab studies.

The proposed conceptual model tracks the relative amount of material at the pipe wall per shear strength. Without placing bounds on the modelled shear strength there would be a need to maintain an understanding of relative quantity over an infinite shear strength range, which is neither computationally practical nor does it have any meaning in terms of the underlying processes.

There is no need for the upper bound of the modelled shear strength range to be greater than the maximum shear stress that the main is anticipated to experience: material stronger than

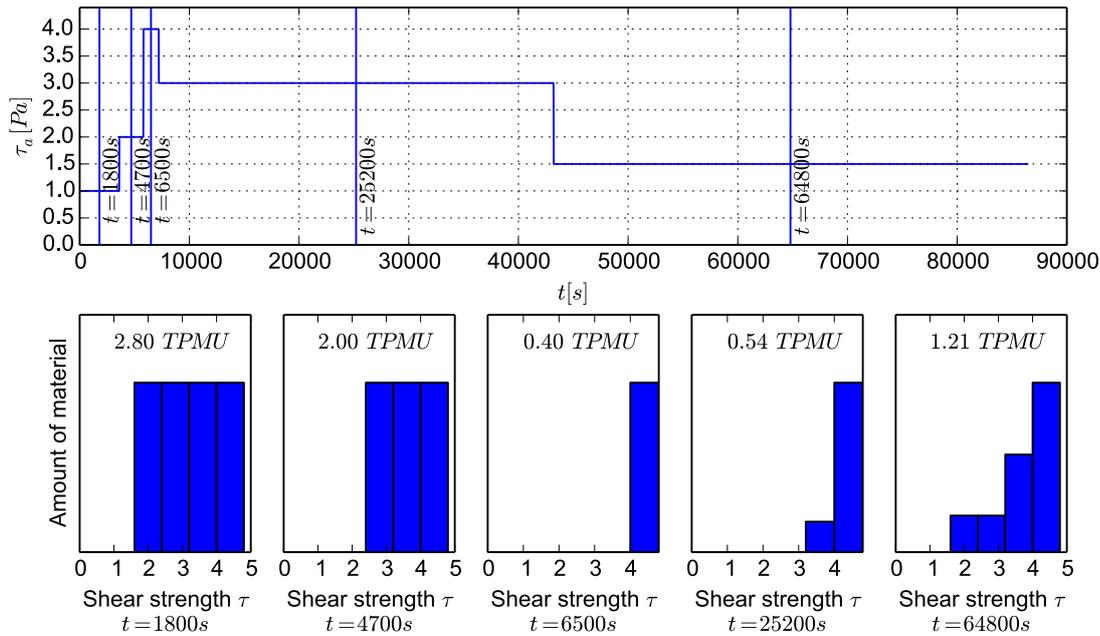


Figure 4.18: Top: an artificial shear stress time-series profile. Bottom: a conceptual illustration of the relationship between material quantity and shear strength at five points along that time-series. Here, over short time-scales material quantity appears to be linear with shear strength yet this is not the case over longer time-periods due to regeneration effects.

this maximum shear stress will never be eroded so cannot influence discolouration risk. The lower bound of the shear strength range could conservatively be set to approach a limit of $0 Pa$ (as per [PODDS](#)) to allow material regeneration at very weak strengths during past/future periods of extremely low flow.

4.3.6 Modelling discolouration material in the bulk water: advection, mixing and reactions

Models of the wall state and material erosion/regeneration are only useful for assessing and predicting discolouration risk if they are coupled to a model of how material propagates through the [DWDS](#). This is because the driver for discolouration modelling is water quality at the locations at which customers are sampling and regulatory samples are taken. Several processes relating to mass transport need to be considered: advection, mixing and bulk reactions (§2.2.7). These are relatively simple to model if the following assumptions, which were considered true for [PODDS](#), remain valid given the current understanding of the relevant processes:

4.3.6.1 Mobilised material does not settle

The majority of discolouration material remains suspended as a wash load once mobilised under all but quiescent conditions. Settling velocity analysis given the size and density distributions of discolouration material particles from various [DWDSs](#) indicates that this is valid (§2.1.9), at least in the regions where these studies were conducted.

4.3.6.2 Turbidity is a pseudo-concentration; pseudo-masses of material

If turbidity is treated as a concentration over the range of turbidities typically seen in DWDSs then existing DWDS water quality models (§2.2.7) can be used to model advection, mixing and, potentially, reactions too. Also, if turbidity is a concentration then discolouration material can be quantified by multiplying turbidity by a dilution volume. The change of wall state scaling parameter mentioned in §4.3.3 can then be set to translate a relative change in the amount of material bound to a unit area of pipe wall to an absolute quantity in the same pseudo-mass units as the product of turbidity and dilution volume.

Here it is assumed that turbidity *can* be treated as a pseudo-concentration: several (but not all) studies show that turbidity is approximately linear with TSS for large, aggregated datasets (§2.1.8). However, that relationship can be much stronger if data are grouped by water source or pipe material, which suggests that the assumption that turbidity can be treated as a pseudo-concentration is only valid at a site-specific level and turbidity ‘concentrations’ cannot be compared between sites. At certain sites turbidity was not found to be linear with turbidity but these are a minority of the sites studied to date.

It is preferable to model turbidity as a concentration and quantify wall-bound material in turbidity pseudo-mass units rather than model TSS directly as:

- Although TSS is a more transferable measure of material quantity, conversion to/from TSS introduces additional error and complexity and TSS cannot be measured as cheaply or continually as turbidity (§2.1.8);
- Turbidity is good measure of the public perception of water quality, which is one of the main drivers for managing discolouration risk.

A pre-requisite of treating turbidity as proxy for TSS and therefore a pseudo-concentration is that turbidity is directly proportional to the volume fraction of particles in the fluid. Turbidity must therefore be linear with particle volume and so linear with D_p^3 yet there is an intuitive argument for turbidity, a measure of light scattering, being more linear with D_p than D_p^3 . However, for three different theoretical expressions for turbidity presented in §2.1.8 turbidity would only be proportional to D_p under very specific conditions. This does not indicate that turbidity is more linear with D_p^3 than D_p , only that it is unlikely to be linear with D_p .

Here it is assumed that the site-specific relationship between turbidity and TSS is invariant with time and is also invariant for ‘background’ and (re)mobilised material.

4.3.6.3 Dispersion and diffusion are negligible; lateral mixing is instantaneous

There is insufficient evidence that dispersion and/or diffusion have a significant impact on the distribution of suspended discolouration material (§2.1.16.5), most likely as DWDS flows are rarely non-turbulent. Dispersion may occur at low Reynold’s numbers in large-diameter trunk mains but this has yet to be conclusively demonstrated. Advection can therefore be modelled using one-dimensional plug flow (§2.2.7.3) and mixing at junctions and transverse mixing in pipes can be assumed to be instantaneous.

4.3.6.4 The upstream turbidity is known or can be estimated

Turbidity as observed at the downstream end of a pipe is not only a function of the release of material from the pipe wall but also of the material that is passing into the upstream end of the main (fig. 4.19). Therefore, to model how the turbidity at the downstream end of a main changes over time one must know or be able to estimate the turbidity at the upstream end. At the upstream extent of a discolouration model the turbidity profile of the inflow can be treated as a boundary condition by common DWDS water quality modelling methods in a similar manner to how chlorine dosing is typically modelled.

4.3.6.5 The modelling scope is the extents of a single pipe length

For the purposes of developing the model it is useful to restrict the model scope to that of a single pipe length (e.g. fig. 4.19). There is then no need to couple the discolouration material release model to a fully-featured mass transport system that can iteratively solve for flows and pressures in looped systems and includes reservoir routing models; instead a simple plug-flow advection model can be used, allowing for easier prototyping. A series of pipe lengths with differing hydraulic or asset characteristics can collectively be modelled in a ‘daisy-chained’ fashion: the turbidity predictions from modelling the most-upstream pipe are supplied as input when modelling the change of state and turbidity for the next pipe downstream, and the output from that sub-model is used as input when considering the next pipe downstream etc. (fig. 4.20). Losses and eddies at pipe joints/enlargements are to be ignored as the mobilisation of material from fixtures and fittings are to be ignored as the relationship between these localised effects and discolouration material accumulation have not yet been well characterised (although in some cases have been observed; see §2.1.16.6).

Reducing the system complexity to a single pipe length does not necessarily impede model validation as there is greater interest in the long-term modelling of discolouration risk in trunk mains than in narrower distribution pipes (at least in the UK; some reasons given in §2.1.17.2) and trunk systems feature relatively few take-offs and are rarely looped (apart from at large scales) so can often be approximated as series of pipe lengths.

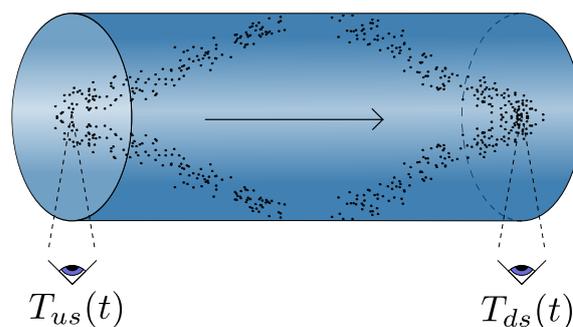


Figure 4.19: The turbidity at the downstream end of a main, T_{ds} , is a function of material erosion, material accumulation (although to a possibly negligible extent) and the turbidity of the water flowing into the upstream end, T_{us} .

Another advantage of limiting the model scope to one or more contiguous pipe lengths is that there is then no need for a storage mixing model; to date the propagation of discolouration material through reservoirs and tanks of different structures has not been characterised.

4.3.7 Summary

A conceptual model of discolouration material erosion, regeneration and pipe wall state has been presented. The output from this model, the rate of material release from the pipe wall, can along with an understanding of system hydraulics be fed as input into an existing distribution network water quality model so as to account for advection and mixing and to calculate turbidity at locations of interest/concern.

By novelly treating the discolouration material bound to the pipe wall as a relative quantity per shear strength, at any given time material within a lower shear strength range can erode whilst stronger material can accumulate. This continuous and concurrent erosion and accumulation has the potential to allow the accumulation process to be more accurately modelled and permit the relationship between the total amount of material and shear strength to be non-linear, which observations indicate is required (see §2.1.13, conceptually illustrated using figs. 4.8 and 4.7).

4.4 Model formulations

In this section a set of terms and formulations are defined that encapsulate the conceptual model described in §4.3, after first defining the units used by several of those terms.

4.4.1 Units for quantifying discolouration material: TPMU

In previous work (e.g. Boxall et al., 2001, 2003a), volume and area featured in the units of wall state, model parameters and model output, such as:

- the amount of material on the wall of a given pipe in $NTU \cdot m^3$ (turbidity as a pseudo-concentration (justification for this assumption is given in §4.3.6.2) multiplied by dilution volume of fluid);
- the amount of material per m^2 of wall in $NTU \cdot m$;
- shear stress in $N \cdot m^{-2}$.

Factorising with respect to length obfuscated the nature of those units. Shear stress is henceforth to be expressed in Pascals (Pa ; $1 Pa = 1 N \cdot m^{-2}$) and quantity of material in *turbidity pseudo-mass units* ($TPMUs$) (a pseudo-mass), where

$$1 TPMU = 1 NTU \cdot m^3 \quad (4.2)$$

Not explicitly including m in units of material quantity and parameters that were derived from turbidity measurements ensures that those units are more tangible.

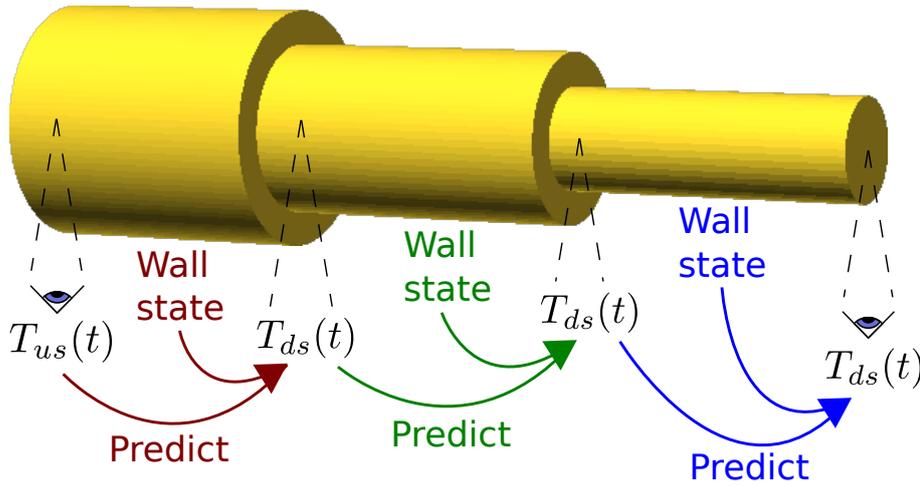


Figure 4.20: Discolouration material release and propagation can be modelled for a heterogeneous series of pipes by conducting a simulation per pipe using a simple advection model then supplying the output of each sub-model as input to the next sub-model downstream.

4.4.2 Formulations for state change and material release from the pipe wall

One of the key concepts of the model is that the amount of wall-bound discolouration material is tracked per shear strength over time (§4.3.2). Let $\varphi(\tau, t)$ be the quantity of discolouration material bound to a unit area of pipe wall with strength τ at time t . Let $\varphi(\tau, t) \in [0, 1]$ be the amount of material present at the wall with a given shear strength; the range $[0, 1]$ denotes that this is a relative rather than absolute quantity (§4.3.3).

Let the rate of change of $\varphi(\tau, t)$ with time be as per eq. 4.3.

$$\frac{\partial \varphi(\tau, t)}{\partial t} = \begin{cases} -\beta_e \cdot (\tau_a(t) - \tau) & \tau_a \geq \tau \\ \beta_r & \tau_a < \tau \end{cases} \quad (4.3)$$

The first case of the case statement in eq. 4.3 corresponds to material erosion: material with a given shear strength erodes if there is an excess shear stress at the given time ($\tau_a(t) \geq \tau$) and there is some material at the wall with that shear strength ($\varphi > 0$). β_e is an erosion rate coefficient with units $Pa^{-1}s^{-1}$. This expression is negated to ensure erosion results in a decrease in the amount of material φ .

Eq. 4.3 satisfies the conceptual requirements that:

- All shear strengths erode simultaneously rather than sequentially (§4.3.4.3) but
- the weaker material erodes more quickly following a step increase in τ_a as the rate of erosion per τ is a function of excess shear stress $\tau_a(t) - \tau$ and
- the total rate of erosion (over the τ range) decreases over time (§4.3.4.4 and fig. 4.16) rather than is linear and then stops abruptly.

The second case of eq. 4.3 relates to the material accumulation process: for all shear strengths for which there is no excess shear stress ($\tau_a(t) < \tau$) and the material quantity has not reached a maximal level, material regenerates at a rate $\beta_r [s^{-1}]$ that is constant with time and shear strength (as per the conceptual specification in §4.3.3).

The amount of material is limited to the range $[0, 1]$ as the amount of material per shear strength is to be modelled as a relative quantity. This clipping requires that the partial derivative in Eq. 4.3 be evaluated iteratively (eq. 4.4) over a discrete mesh bounded by $\tau \in [\tau_{min}, \tau_{max}]$ and $t \in [t_{min}, t_{max}]$. Here τ_{min} is $0 Pa$ and τ_{max} is case-specific (e.g. it could be the maximum shear stress that a main could ever experience as per §4.3.5).

$$\varphi(\tau_i, t) = \varphi(\tau_i, t - \Delta t) + \left(\frac{d\varphi(\tau, t)}{dt} \cdot \Delta t \right) \quad (4.4)$$

Note that eq. 4.3 allows the material per shear strength to erode and regenerate independently of other shear strengths.

To determine the *total rate* of material release from a unit of wall area, the change in the amount of material over the entire range $[\tau_{min}, \tau_{max}]$ must be aggregated. This is shown in eq. 4.5, where the result is then converted from relative to absolute units of material quantity using a linear scaling parameter α (introduced in §4.3.3) with units of $TPMU \cdot Pa^{-1} \cdot m^{-2}$. The negation ensures that there is a positive material release rate $dN(t)/dt$ (with units of $TPMU \cdot m^{-2} \cdot s^{-1}$) for a decrease in the amount of wall-bound material ¹.

$$\begin{aligned} \frac{dN(t)}{dt} &= -\alpha \cdot \sum_i^{\tau_a \geq \tau_i} \frac{\varphi(\tau_i, t) - \varphi(\tau_i, t - \Delta t)}{\Delta t} \cdot \Delta \tau \\ \left[\frac{TPMU}{m^2 \cdot s} \right] &= \left[\frac{TPMU}{m^2 \cdot Pa} \right] \left[\frac{-}{s} \right] [Pa] \end{aligned} \quad (4.5)$$

In eq. 4.5, the summation is only over the *eroding shear strengths* at time t ; material accumulation is assumed to have a negligible effect on the bulk water turbidity (and vice versa) so shear strengths that are accumulating are not considered in this summation. The rate of accumulation is therefore only modelled (indirectly) using the β_r parameter. For example, in fig. 4.21 the calculation of the material release rate from the pipe wall over the time-step Δt only considers the mobilisation of material with $\tau \in [0.6, 2.1] Pa$ but not how the accumulation over $\tau \in (0.6, 2.1] Pa$ slightly reduced the net material release rate.

The three model parameters β_e , β_r and α are treated as scalars. There is insufficient evidence for any being *significantly* variable with respect to t or (absolute) τ ; good justification is needed for making say β_r temperature- and therefore time-dependent as this will increase the model complexity.

4.4.3 Sub-models and formulations for material transport

The output of the wall state change and material release model, $dN(t)/dt$, is supplied as an input to a mass transport model. As mentioned in the conceptual model section (§4.3.6), this mass transport model needs to advect both material passing into the upstream end of a given pipe length and material released from the wall into the bulk water.

¹The model formulations presented in §4.4 were previously published in Furnass et al. (2014b). Note that the units of α and $dN(t)/dt$ as shown in Furnass et al. (2014b) are incorrect. The definitions and units in §4.4 should be considered to be definitive.

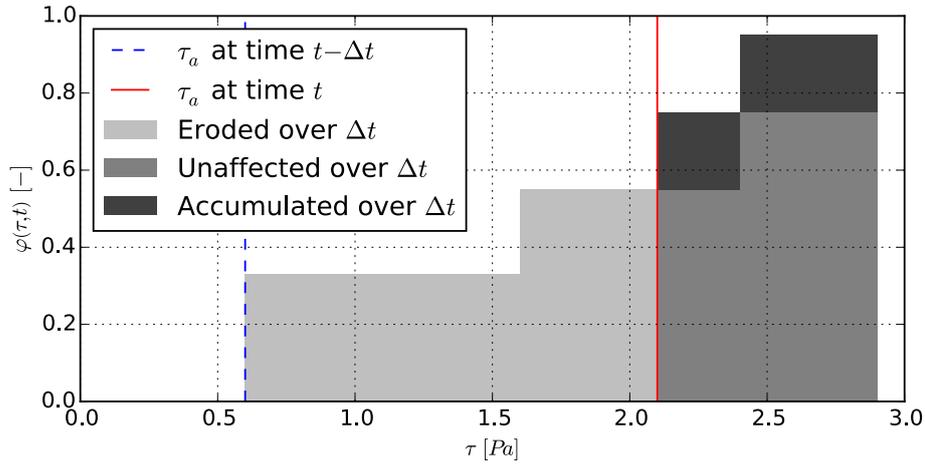


Figure 4.21: The model formulations allow material for which there is an excess shear stress to erode whilst simultaneously the material with shear strengths greater than the shear stress can accumulate up to a maximal level (illustrative diagram).

An analytical transport solution is not appropriate: an iterative method is required as the input $dN(t)/dt$ cannot be evaluated analytically and another key input, pipe flow, is a non-differentiable function of time. A Lagrangian transport model, where discrete water parcels/segments of constant concentration are tracked as they move down a pipe (further described in §4.3.6), satisfies the conceptual requirements outlined earlier in §4.3.6 (one-dimensional plug flow and a means for defining an upstream concentration profile as a boundary condition, but no immediate requirement for a storage mixing model). More specifically, a *time-driven* Lagrangian transport model has been chosen over other Lagrangian and Eulerian models that offer similar functionality as it is computationally efficient and is less affected by numerical dispersion (Rossman and Boulos, 1996). With this model, parcel positions and concentrations are updated at regular time intervals.

The migration of mobilised material from the pipe wall into well-mixed Lagrangian water parcels/segments can be calculated as per the pseudocode shown in alg. 4.1 (near identical to the transport algorithm used in EPANET (Rossman, 2000, p. 41)). Note that in this pseudocode the symbol \leftarrow represents variable assignment. At each time t each parcel has a known volume, concentration and position relative to all other parcels, with the combined volume of all parcels equalling that of the pipe.

At the upstream end of the modelled pipe, a volume of water is added at each timestep that is equal to the flow Q integrated over the timestep i.e. $Q(t) \cdot \Delta t$. For the purposes of efficiency new water segments are not added the upstream end at every timestep: the incoming turbidity, $T_{us}(t)$, at a given time t is used to either set the concentration of a new segment with that volume or the most upstream parcel is augmented by that volume, and $T_{us}(t)$ along with the

Algorithm 4.1 Calculating material release into the bulk water for a pipe p at time t .

```

procedure PIPE_WALL_REACT( $t, p, D, L, dN(t)/dt, \Delta t$ )
  for all  $s \in p.segments$  do
     $s.wall\_area \leftarrow \pi \cdot D \cdot L \cdot s.volume/p.volume$ 
     $s.concentration \leftarrow s.concentration + dN(t)/dt \cdot \Delta t \cdot s.wall\_area/s.volume$ 
  end for
end procedure

```

segment's current concentration are used to assign a volume-weighted average.

In a complementary way, water is removed from the downstream end of the pipe at each timestep. The 'mass' (material) that passes out of the pipe during the timestep is then divided by the total pipe discharge over that period to give a downstream (nodal) concentration, which in this case is the downstream turbidity $T_{ds}(t)$.

For each timestep the water that passes into and is discharged from the pipe only spends half of that time within the pipe. The material release into these two volumes of water therefore needs to be halved to account for this.

These three processes, material being released from the pipe wall to the bulk water, water and material passing into the upstream end of the pipe and water and material passing out via a 'monitoring point', are calculated in series for each timestep of a simulation and all timesteps are solved iteratively (assuming that the time-series solutions of the discolouration material release model (§4.4.2) and the single pipe hydraulic model are already available).

4.4.4 Boundary conditions

The temporal boundary condition for the pipe wall model is the relative amount of material per shear strength band at the start of the simulation i.e. $\varphi(\tau_i, t = 0)$. The model formulations do not allow the amount of material at a given shear strength band to exceed that of a band of the same width but greater magnitude. The wall state and therefore the wall boundary condition $\varphi(\tau_i, t = 0)$ must therefore be a *monotonically increasing function* with respect to shear strength τ e.g. fig. 4.22.

The temporal boundary condition for the bulk transport model is the array of concentrations (here, turbidities) of *each* Lagrangian water parcel at $t = 0$. All parcel concentrations need to be initially set to a constant value as there is insufficient information to do otherwise. This means that if the turbidity at the upstream end of the modelled pipe is non-zero at the start of the simulation then that background turbidity signal will only be reflected in the model output after the duration of one pipe turnover (allowing for variability in the flow rate). It is not uncommon for other water quality models such as those for age to also require a 'bedding in' period at the start of each simulation before the output becomes representative (Machell and Boxall, 2012).

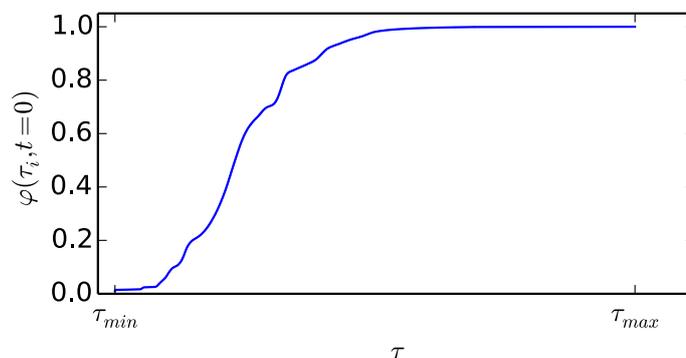


Figure 4.22: The boundary condition at the pipe wall is $\varphi(\tau_i, t = 0)$, which is a monotonically increasing function with respect to τ . An illustrative example of a monotonically increasing pipe wall boundary condition is shown here.

The spatial boundary condition of the bulk transport model is the time-series turbidity profile of water coming in to the modelled pipe, as mentioned in §4.3.6.4 and §4.4.3.

4.4.5 Variable dependencies

The relationships between all quantities explicitly referenced by the proposed model are shown in fig. 4.23 (with the addition of temperature and fluid density, which are not explicitly included in the model as the impact of variation in those factors in real DWDS is considered to be negligible). Note that flow and pipe diameter influence *both* the material erosion process (by affecting shear stress) and material transport.

4.4.6 A note on parsimony

The formulations and algorithms presented in this section were made only as complex as could be justified given the available field and laboratory observations and intuitive understanding of discolouration and other water quality and hydraulic processes. An over-parameterised, overly-complex mode is not only more difficult to fit but may also be too flexible and therefore prone to over-fitting (§2.2.3).

4.5 Implementation as software

4.5.1 Requirements

The proposed formulations (§4.4) and underlying conceptual model (§4.3) need to be encoded as software in a way that allows for the exploration of sensitivity with respect to parameters and data inputs and for model verification and validation.

This reference implementation should be quick to develop and easy to comprehend and modify, at least until the model has been successfully validated. An application programmer's interface (API) plus additional user interface components offers more versatility than a monolithic application with a rich graphical user interface.

The software should also be efficient with respect to CPU time, memory usage and computational complexity i.e. the execution time increases at an acceptable rate as the problem size, a function of the cardinality of the input vectors, increases. Efficient execution along with the ability to execute multiple simulations in parallel greatly expedites the model fitting process when using methods such as swarm optimisation or genetic algorithms (§2.2.4) to find the parameters that minimise the dissimilarity between predictions and observations.

The software needs to solve a numerical problem where the rate of change of inputs, state and outputs varies with time (it may be months between short, sharp discolouration incidents): it is therefore more computationally efficient for the software to allow for and use a variable rather than fixed resolution mesh w.r.t. time when modelling using historical data. Being able to vary the mesh resolution with respect to shear strength is also advantageous for similar reasons.

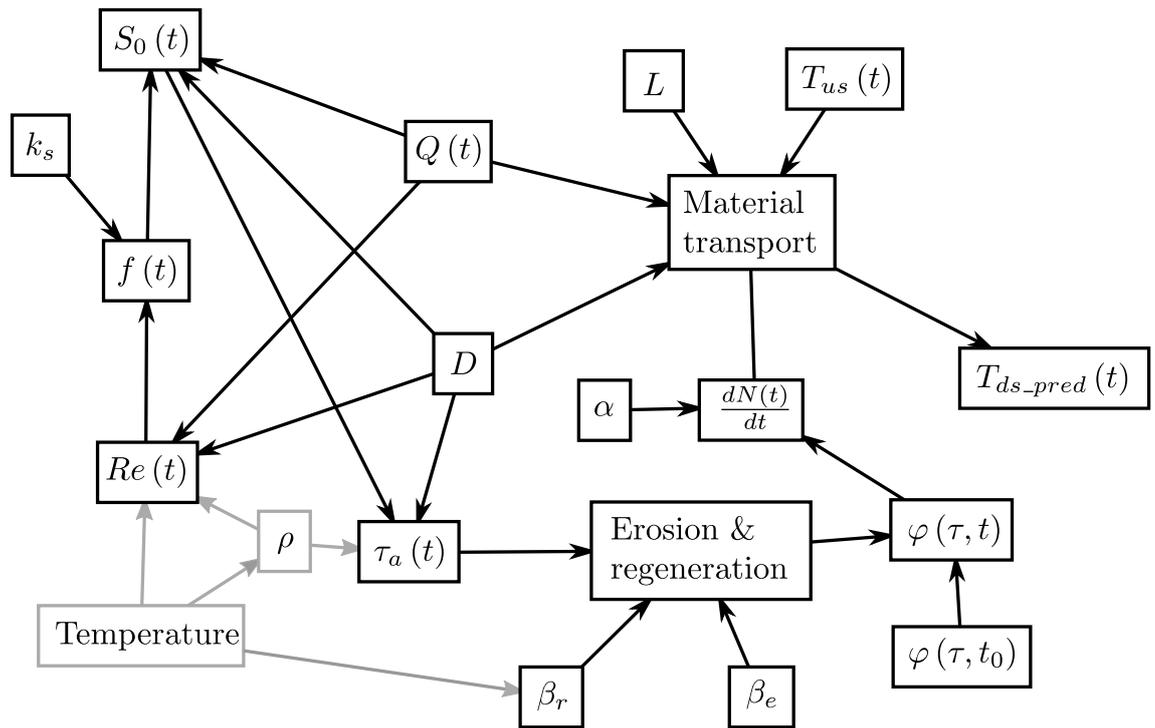


Figure 4.23: Interdependencies between quantities relating to the VCDM model. k_s : pipe roughness; S_0 : hydraulic gradient; τ_a : shear stress; t : time; f : Darcy Weisbach friction factor; Re : Reynolds number; Q : flow; D : internal pipe diameter; ρ : fluid density; β_e , β_r and α : VCDM parameters; $\varphi(\tau, t)$: relative amount of material on the pipe wall with strength τ at time t ; $\varphi(\tau, t_0)$: relative amount of material on the pipe wall with strength τ at time t_0 ; $dN(t)/dt$: rate of material release per unit area of pipe wall; L : pipe length; T_{us} : turbidity at the upstream end of the pipe of interest; T_{ds_pred} : predicted turbidity at the downstream end of the pipe of interest. Variation in the parameters in grey boxes is not directly modelled by the proposed model.

4.5.2 Implementation options

EPANET (§2.2.7), the open-source hydraulic and water quality modelling library, command-line application and graphical application, is commonly used as a base for investigating variations on traditional DWDS modelling approaches (e.g. Andrade et al., 2010). It cannot be used in its original form for modelling discoloration as its single-specie water quality engine does not track the quantity of specie at the pipe wall: when PODDS was integrated into EPANET (§2.2.8.4) the water quality engine, input and output functions, API and graphical user interface (GUI) were all modified to allow additional pipe and node parameters to be set, to track the wall state and to allow material to be released into the bulk water, which was then advected and mixed using in-built functions. Let us label the option of implementing the proposed model by *forking* (modifying) the PODDS version of EPANET as *PODDS_C*.

Alternative means for encoding the model as software were to use the multi-species extension (MSX) to EPANET (§2.2.7.6), either through supplying models using input files (henceforth referenced as *MSX_F*) or through modification of its source code (*MSX_C*).

The other option that was considered was to develop a bespoke software library rather than fork an existing project (*BESP*). The programming language Python (Lutz, 2011) is a good choice for this task as it is interpreted and dynamically typed, facilitating rapid prototyping, data exploration and experimentation, yet can interface with statically typed, compiled *extension modules* for expediting computationally-intensive tasks (e.g. via Cython (Behnel et al., 2011)).

4.5.2.1 Encoding state change, material release and material transport

For *PODDS_C*, the PODDS + EPANET water quality engine and API would need to be modified so that firstly the wall state were a vector rather than a scalar and secondly the model initialisation and state change rules reflected the needs and formulations of the proposed model. As with PODDS, the advection and mixing of released material would be simulated using a single-specie water quality model, as if discoloration material were an inert solute.

MSX-derived solutions (*MSX_F* and *MSX_C*) offer in-built support for tracking the quantity of one (or more) species at the pipe wall and in the bulk water. MSX input files allow water quality equilibrium, rate and formula expressions to be defined for pipes and tanks but the grammar of these expressions is sufficiently limiting that it is difficult to encode the formulations of the proposed model in this way (*MSX_F*), particularly in a manner that allows for the tracking of material quantities over a large range of shear strength bands. However, the MSX expression parsing and evaluation rules could be rewritten to allow for this by modifying the MSX source code (*MSX_C*).

Another disadvantage of the *MSX_F* approach is that the predicted turbidity response lags by one timestep behind flow changes as dN/dt is updated outside of the MSX timestep code. This error may be significant if the timestep duration is as long as fifteen minutes unless it is accounted for.

A bespoke solution coded in a language that allows for rapid prototyping (*BESP*) would allow for easy, transparent encoding of the proposed model's formulations. However, it would require

functions for estimating shear stress and for the advection and mixing of a water quality specie; this functionality is already provided for PODDS_C, MSX_C and MSX_F.

4.5.2.2 Model scalability

The four implementation options differ in their ability to scale to meet potential future modelling needs.

Rapid prototyping The ease of development afforded by Python (BESP) could make sensitivity analysis and extensions to the model (e.g. a transport mechanism that can model dispersion) much easier than for PODDS_C or MSX_C.

Multi-specie sub-models of material regeneration The regeneration rate parameter β_r is to be decomposed into regeneration sub-models if/when there is sufficient data and understanding to allow those sub-models to be justified and mathematically articulated (fig. 4.10). In future the dependencies of β_r may include parameters that in reality significantly vary over short spatio-temporal scales. It would therefore be desirable (yet not currently essential) to allow these parameters to be accurately represented within the model as bulk water or wall species. Such multi-species, multi-phase modelling would be possible under MSX_F and MSX_C but not PODDS_C. Support for this could be included within BESP, albeit with additional time requirements.

Multi-pipe models Working with multi-pipe models would be much easier if using software that is based on EPANET (inc. MSX) as EPANET can iteratively solve for flows and pressures in large, complex, looped networks containing devices such as pumps and regulating valves. Enabling the modelling such systems using BESP would require considerable development effort; BESP is therefore seen as being a simple means for verifying and validating the proposed model, after which the formulations can be incorporated into a more complex, powerful modelling framework e.g. PODDS_C or MSX_C. However, strictly dendritic multi-pipe networks could be modelled using BESP by modelling each pipe in turn using a breadth or depth first traversal starting from the most upstream node: the output from each sub-model could be supplied as an input to the next pipe(s) downstream (as per fig. 4.20).

4.5.2.3 Performance scalability

Variable mesh resolution EPANET/MSX-based implementations require that the duration of hydraulic and water quality timesteps be constant throughout a simulation, whereas BESP permits the timestep duration to vary, reducing the computational requirements over parts of the input time-series where there was little temporal variation (§4.5.1).

Efficient CPU and memory usage EPANET and MSX are both statically typed and compiled so are fast and make efficient use of memory. Statically typed, compiled software such as code written in C or Fortran is particularly well suited to numerical computing such as solving

DWDS models as all arrays and matrices are homogeneously typed and so elements can be iterated over very efficiently in memory as the size in bytes of each element is known in advance. However, care is required when making changes to software such as EPANET that is written in C as memory management, and therefore avoiding memory leaks, is entirely the responsibility of the programmer.

Conversely, interpreted Python is slow for numerical work as its native enumerable data types are heterogeneous collections of objects and so iterating over a 'list' in Python requires that the size and position of each object be determined at run-time, which greatly limits performance for enumerables containing many elements. However, Python has libraries that allow the more numerical parts of a given application to be stored and manipulated using datastructures and routines written in low-level, compiled programming languages such as C:

- *NumPy* (van der Walt et al., 2011) provides functions for creating homogeneous C arrays/matrices and for fast vectorised operations on those arrays. This is achieved through much of NumPy being compiled Python *extension modules*.
- *pandas* (McKinney, 2013) does the same but for tabular *labelled data* (similar to the way that tables in spreadsheets have column and row labels). Both it and NumPy have functions for efficiently reading and writing numerical data from/to files and for timestamp and time delta manipulation.
- *Cython* (Behnel et al., 2011) allows Python code to be compiled as a Python extension module, with a user-defined amount of static typing, allowing complicated algorithms that cannot be expressed in terms of vectorised operations to be compiled as a library then accessed from Python. Cython can also be used to create Python wrappers around libraries that were originally written in C or C++ such as EPANET/MSX.

Multiple concurrent simulations It is desirable to run many simulations simultaneously when exploring the parameter space of a model using Monte Carlo simulations (Metropolis and Ulam, 1949) or trying to fit the model to data using an automated method such as PSO (§2.2.4). EPANET and MSX are not well suited to such tasks: the extensive use of global variables means they are not *thread-safe* unless statically compiled into an executable; multiple instances of that program could then be executed concurrently but this approach is convoluted, particularly with regards to inputs and extracting outputs from the various program instances. With a Python-based approach sub-processes could be spawned in a way that is more easy for the programmer to manage. Unfortunately, as with PODDS_C, MSX_F and MSX_C, there is no way of running simulations entirely concurrently using threads (which could be far more memory efficient than using parallel processes) as Python has a Global Interpreter Lock that prevents multiple threads from executing simultaneously, even on a multi-core machine².

4.5.2.4 Data input, output and user interface

The EPANET and MSX user interfaces (EPANET: GUI, command-line and API; MSX: command-line and API only) are well understood within academia and industry. However, these interfaces

²Note that the Global Interpreter Lock can be temporarily be released for certain functions within Python extensions written directly in C or using Cython, which permits parallel execution those functions.

along with file parsing and file generating functions would have to be modified for PODDS_C, MSX_F and MSX_C to allow for the input of an array shear strength bands, an array of the initial amount of material per band and for the output of the final amount of material per band. BESP affords far simpler means for inputting and outputting numerical data. NumPy and pandas can be used to read and write delimited text files or efficient binary files in formats such as *HDF5* (The HDF Group, 2014; McKinney, 2013). Complicated plots can be created using the *matplotlib* library Hunter (2007). The IPython Notebook (McKinney, 2013) is a *literate programming* environment where blocks of code, formatted text along with outputs such as code-generated data and graphs can be brought together within an interactive document, allowing analysis and demonstrations to be more easily reproduced, shared and comprehended (through the use of in-line documentation).

4.5.2.5 Dependencies

PODDS_C, MSX_X and MSX_F have no dependencies (other than MSX on EPANET), whereas BESP requires NumPy, matplotlib³, a library of hydraulic functions and optionally depends on Cython, pandas and IPython.

4.5.2.6 Chosen approach

For the purposes of validating the proposed model for a single pipe, BESP was considered to offer the best balance of efficiency and ease of development and the greatest flexibility with respect to model fitting and sensitivity analysis. These benefits were seen to more than compensate for having to develop code to calculate shear stresses and model advection and mixing.

4.5.3 Overview of implementation

The `pyvcdm` Python package provides `VCDM` class with which discolouration material release models can be instantiated then executed for a given time-series before then modelling advection and mixing for that process i.e. the wall state change and material release code is largely decoupled from the advection and mixing code.

A model can be instantiated using this Python code:

```
my_model = pyvcdm.VCDM(times_v,
                        strength_v, init_cond_v,
                        applied_shear_v,
                        removal_rate, regen_rate, alpha)
```

The parameters of these methods are as follows:

times_v A 1D array (of type `numpy.ndarray`) of time offsets since a datum t_0 , in s , over which to simulate. Allows for variable mesh resolution with respect to t ;

³It would be impractical to code array manipulation and plotting libraries from scratch.

strength_v A 1D array of the extents (in Pa) of the shear strength bands at which the relative amount of material will be tracked over time. Allows for variable mesh resolution with respect to τ ; note that the excess shear stress per shear strength band is calculated with respect to the *mid-points* of the bands defined by `strength_v`;

init_cond_v $\varphi(\tau_i, t = t_0)$; a 1D array of the relative amount of material, a unit-less quantity, of all layers at $t = t_0$. Each element must be in $[0, 1]$, with 1 indicating maximal material present. Must have one fewer element than `strength_v` i.e. have as many elements as there are shear strength band mid-points;

applied_shear_v A vector of values of shear stress per timestep [Pa]. Must be the same cardinality as `times_v`.

removal_rate β_e ; Scalar coefficient, in $Pa^{-1} \cdot s^{-1}$, that describes the ease with which the relative amount of material per shear strength decreases upon erosion as a function of the excess shear strength;

regen_rate β_r ; Scalar coefficient, in s^{-1} , that describes the rate at which the relative amount of material per shear strength increases due to material regeneration;

alpha α ; Scalar coefficient, in $TPMU \cdot m^{-2} \cdot s^{-1}$, that is used to convert the relative amount of material mobilised to absolute units.

Once a model has been instantiated the material release can then be simulated for the previously-supplied time-series:

```
my_model.sim(cython=True)
```

Here, the iterative change of wall state and material release calculations can be performed using a Cython function for far greater performance than could be achieved using native Python code.

The material release simulation results can then be advected to the downstream end of a single pipe using a time-driven Lagrangian mass transport model (again written in Cython for performance). The turbidity at the upstream end of the pipe of interest can be supplied as an input to the transport model (§4.4.3).

```
my_model.advect(Q_v, D, L, init_num_segs,
               upstream_conc_v=None,
               conc_tol=0.05, cython=True)
```

The parameters of the `advect` method are as follows:

Q_v A 1D array of pipe flow in $m^3 \cdot s^{-1}$, with the same cardinality as `times_v`. This should be consistent with the `applied_shear_v`, `k_s` and `D` function arguments given that shear stress is a function of flow, diameter and pipe roughness;

D Uniform pipe diameter in m ;

L Pipe length in m ;

`init_num_segs` At the start of the advection process the pipe volume is initially divided into a number of equal-sized segments of zero turbidity;

`upstream_conc_v` A 1D time-series array of turbidities in NTU corresponding to the upstream end of the pipe. The array should have the same cardinality as `times_v`;

`conc_tol` Tolerance in NTU used to determine whether a new segment should be added to the upstream end of the pipe at a particular timestep (see §4.4.3).

Functions for hydraulic calculations have been encapsulated within a separate Python package, `pyhyd`⁴, which includes functions for calculating hydraulic quantities from both scalar and vector (array) inputs. The most useful function for modelling material erosion is the following for calculating the shear stress at the pipe wall, which solves eqs. 2.9, 2.10, 2.11 and 2.12:

```
shear_stress(D, Q, k_s, T=10.0, den=1000.0)
```

where

D Internal diameter in m ;

Q Flow in $m^3 \cdot s^{-1}$;

`k_s` Nikuradse roughness height in m ;

T Temperature of the bulk water, which defaults to $10^\circ C$. Dynamic viscosity is a function of temperature and therefore the Reynolds number and friction factor are too;

`den` Bulk water density, which defaults to $1000 \text{ kg} \cdot m^{-3}$.

4.6 Summary of model assumptions

The assumptions made in §4.3 and §4.4 have been collated and enumerated here for quick reference.

Assumptions particular to just the VCDM

1. At a given time, all wall-bound material with a shear strength less than the shear stress is subject to erosion whilst simultaneously all material with a shear strength greater than the shear stress is subject to accumulation i.e. weaker material may erode whilst at the same time stronger material accumulates.
2. Weaker-bound material erodes more quickly than stronger material: all strengths for which there is an excess shear erode simultaneously at rates that are constant with time but variable with excess shear.

⁴<https://github.com/willfurnass/pyhyd>

3. The relative regeneration rate, β_r , is pipe specific but constant with shear strength and time (so is not significantly affected by rapid changes in bulk water quality such as discolouration material erosion events).
4. The other two model parameters, α and β_e , are also constant with shear strength and time.
5. The shear strength of a quantum of wall-bound discolouration material is constant with time.

Assumptions particular to both [PODDS](#) and the [VCDM](#)

1. The wall state (quantity of wall-bound material versus shear strength profile), erosion rates and regeneration rates are invariant with longitudinal distance and angle to the invert along a pipe length with constant asset characteristics (diameter, roughness and material) and longitudinally homogeneous hydraulics (bulk flow).
2. The amount of material that can accumulate on the pipe wall with a shear strength in a given range has a time-invariant upper bound, allowing erosion and regeneration to be considered in relative terms.
3. The maximum shear strength with which discolouration material can bind to the wall is pipe-dependent and invariant with time.
4. The minimum shear strength of wall bound material is time-invariant and approaches the limit of $0 Pa$.
5. The relative rate of erosion is solely a function of the excess shear stress, which is the amount by which the steady-state shear stress exceeds the shear strength for a quantum of material.
6. The relationship between turbidity and [TSS](#) at a given site is linear and invariant of time and also, for eroded material, the shear strength and longitudinal position along the homogeneous pipe prior to erosion. Turbidity can therefore be treated as a concentration and modelled as a water quality species. The effects of coagulation, flocculation and precipitation are negligible.
7. Bulk water turbidity is not significantly reduced by material accumulating on the pipe wall.
8. Material remains suspended once mobilised.

Assumptions commonly made when modelling [DWDS](#) hydraulics and water quality

1. The dispersion and diffusion of solutes and suspensions that behave like solutes are negligible and lateral mixing is instantaneous.
2. Dynamic hydraulic effects do not significantly contribute to the shear stress at the pipe wall.

4.7 Conclusions

In this chapter a conceptual model was proposed for more accurately modelling the shear strength profile of the discolouration material that accumulates around the circumference of DWDS pipe walls. The key advance of this Variable Condition Discolouration Model (VCDM) is that the state of the pipe wall at a given time is a relative quantity of material for each of a range of shear strengths, as opposed to all material bound to the pipe wall only having a scalar strength. Each of the tracked shear strengths can be eroded or accumulate more material independently of all other shear strengths. This permits the accumulation process to be modelled in a distinct manner from the erosion process rather than simply being the inverse: weaker material can erode more quickly than stronger material for shear strengths less than the shear stress, yet shear strengths above the shear stress can accumulate at all shear strengths simultaneously and at the same rate, in concordance with field and laboratory observations.

Tracking material quantity over time for a range of shear strengths rather than for a single shear strength requires a novel approach to modelling excess-shear-stress-driven material erosion. It was not possible to distinguish the several possible conceptual models of erosion that were considered in this chapter given the available historical data so the most intuitive option was selected: all shear strengths for which there is an excess shear stress erode simultaneously but at a linear rate that is proportional to the excess shear stress.

The VCDM model has fewer parameters than PODDS, the parameters being: a relative erosion rate factor β_e , a relative regeneration rate β_r , and a scaling factor α for converting from relative to absolute quantities of discolouration material. The wall state boundary condition for $t = 0$ is the relative material quantity per tracked shear strength, which is a non-trivial monotonically-increasing function of τ . The complexity of the VCDM wall state model is greater than that of the PODDS model but only as complex as is required to represent the observed characteristics of material erosion and regeneration processes.

The proposed VCDM formulations calculate the change in relative material quantity at the pipe wall due to material erosion and regeneration, determine the total material release into the bulk water from each unit area of pipe wall then this material is advected to the observation point at the downstream end of a single pipe by a time-driven Lagrangian transport model.

The wall state and mass transport formulations need to be evaluated using an iterative and discretised rather than analytical approach. These formulations have been encoded as Python software.

At the start of the chapter several 'use cases' for a new discolouration model were defined (§4.2.1). As has been demonstrated, the proposed model has the potential to characterise the sensitivity of both the bulk water turbidity and the latent wall state to the continual processes of discolouration material and erosion (use cases #1, #2 and #3). In the next chapter, §5, the mesh independence of this numerical model is assessed and sensitivity analysis is undertaken on the mesh resolution, model parameters and the boundary conditions. It concludes with analysis of how the empirical model can be calibrated (use case #4) using artificial data. The following chapter, §6, explores how successfully the model can be calibrated and validated using data from real DWDS (use cases #4 and #5). The final use case, #6, is the ability to use the

discolouration model to design optimal mains flow conditioning programmes; this is discussed in §7.

Chapter 5

Model verification, sensitivity analysis and fitting to data

“Model verification is substantiating that the model is transformed from one form into another, as intended, with sufficient accuracy. Model verification deals with building the model *right*. ”

Balci (1997)

5.1 Introduction

This chapter consists of three sections:

- A demonstration of the functionality of the model proposed in §4, including the verification of model formulations and implementation (confirming that the model satisfies the conceptual specification);
- An exploration of the sensitivity of the model output to the model parameters, inputs (hydraulics and asset characteristics) and boundary conditions (initial system state).
- The development and validation of a method for fitting the model to data.

5.2 Demonstrations of model functionality

Five demonstrations of the model are presented in this section, which collectively cover various aspects of model functionality:

1. The impact of a simple step increase in shear stress on the model state, the rate of material release from the pipe wall and the turbidity at the downstream end of the pipe;

2. The rate of material release time-series is partially obscured when transformed by bulk water transport processes into the downstream turbidity time-series;
3. Fluctuations in shear stress can result in the partial erosion of shear strength bands. This demonstration explores firstly the impact of partial erosion on the model state and outputs, and secondly how, under certain circumstances, distinct releases of material from the pipe wall cannot easily be distinguished from the turbidity signal;
4. The **VCDM** and **PODDS** models can differ following periods of material regeneration;
5. Erosion is driven by the excess shear stress and the importance of quantifying discolouration material using the time-integral of material flux (not just of observed turbidity).

5.2.1 First example: model response to a simple, stepped shear stress profile

The functionality of the model concepts and formulations proposed in §4.3 and §4.4 and implemented as per §4.5 is demonstrated in fig. 5.1 using artificial inputs. Figures depicting simulation data in a similar manner appear throughout the remainder of this thesis, hence there is a need to explain this graphical representation.

The four subplots of fig. 5.1 are time-series plots corresponding to the same time period (in this case 1000 s):

- The uppermost subplot, fig. 5.1-i, shows the variation in shear stress at the wall of a single pipe; this was calculated from a diameter D , roughness k_s and flow profile Q . In this example, the flow (not shown here) and shear stress profiles increase and decrease in a simple stepped fashion;
- The second subplot shows how the relative amount of wall-bound material, $\varphi(\tau_i, t)$, varies for three different shear strength (τ) bands during the simulation. A value of 1 for a given shear strength band indicates that there is maximal material present with that shear strength and a value of 0 indicates that band is fully depleted;
- The third subplot, fig. 5.1-iii, shows the profile of the rate of material release, dN/dt , from each m^2 of pipe wall due to *all* eroding shear strength bands (not just the three shown in the second subplot);
- The fourth subplot, fig. 5.1-iv, shows the turbidity that was predicted at the downstream end of the modelled pipe. Note that here it is assumed that the turbidity of water that enters the upstream end of the pipe is zero, so the downstream turbidity profile, T_{ds} , shown in this subplot is solely a function of material erosion from within the pipe.

In fig. 5.1-ii, the relative amount of wall-bound material $\varphi(\tau_i, t)$ was plotted for just three shear strengths for the purposes of concisely illustrating how erosion and accumulation are conceptually being modelled; however $\varphi(\tau_i, t)$ was tracked for a total of 250 discrete shear strengths over a contiguous range.

In fig. 5.1-ii, all three of the shown shear strengths are simultaneously eroded by the first step increase in τ_a , with the weakest material being eroded more quickly (as per the conceptual

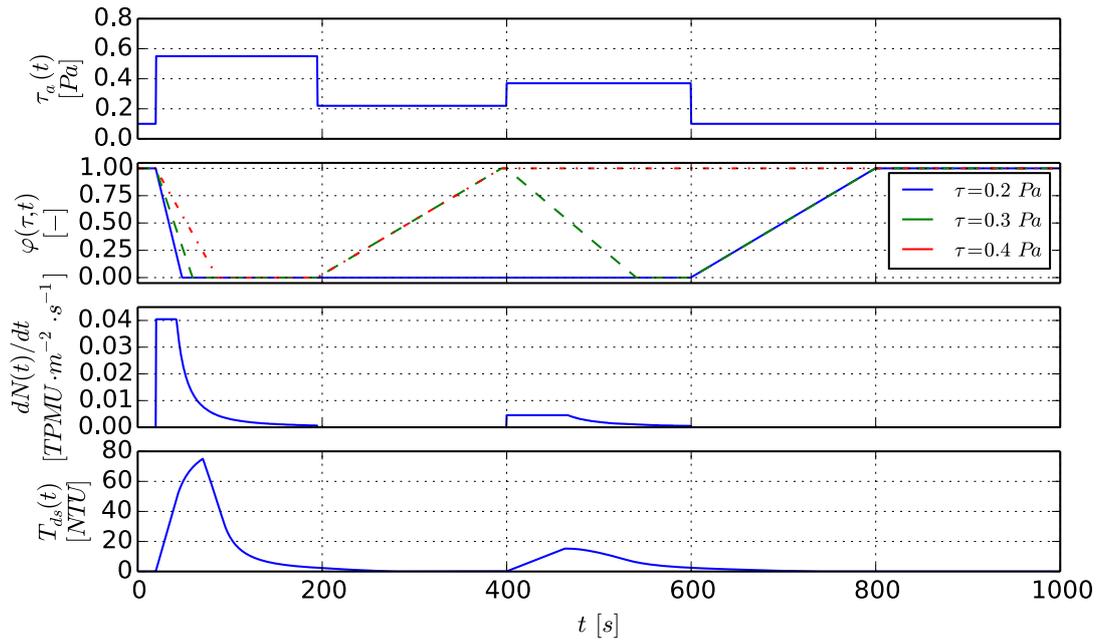


Figure 5.1: A demonstration of how the proposed model generates a turbidity response at the downstream end of a single pipe. From top to bottom: (i) A time-series shear stress profile (one of the model inputs); (ii) the relative material quantity $\varphi(\tau_i, t_j)$ for an illustrative selection of shear strengths $\tau_i \in \{0.2, 0.3, 0.4\}$, which varies with time due to erosion and regeneration; (iii) the rate of material release from the pipe wall $dN(t_j)/dt$ due to the erosion of all layers weaker than the shear stress; (iv) the turbidity response (one of the model outputs). Pipe diameter $D = 75 \text{ mm}$, roughness $k_s = 0.1 \text{ mm}$ and length $L = 20 \text{ m}$; timestep $\Delta t = 0.5 \text{ s}$; 250 shear strength bands ($\Delta\tau = 0.0026 \text{ Pa}$); boundary conditions: if $\tau_i < \tau_a(t=0)$ then $\varphi(\tau_i, t_j=0) = 0$ else $\varphi(\tau_i, t_j=0) = 1$ (see eqs. 4.3 and 4.4); erosion rate factor $\beta_e = 0.1 \text{ Pa}^{-1} \cdot \text{s}^{-1}$; regeneration rate $\beta_r = 0.005 \text{ s}^{-1}$; scaling factor $\alpha = 4.0 \text{ TPMU} \cdot \text{Pa}^{-1} \cdot \text{m}^{-2}$; flow $Q \in [0.68, 1.75] \text{ L} \cdot \text{s}^{-1}$. Note that in this example the rate of material regeneration has been set to a high value to permit the effects of erosion and regeneration to be clearly visualised on a single plot.

requirements in §4.3.4). Material with shear strengths of $0.2 Pa$ and $0.3 Pa$ then regenerates simultaneously at the same rate (as per the requirements in §4.3.3) to a maximal level (see the requirements in §4.3.3) when τ_a drops to below $0.2 Pa$. Upon a second step increase in τ_a , only the strongest of the three explicitly depicted shear strengths withstands the eroding forces. All three identified strengths then subsequently regenerate simultaneously at the same rate again when τ_a drops down to $0.1 Pa$. Figures 5.1-iii and 5.1-iv show two spikes in the material release rate, dN/dt , and the downstream turbidity, T_{ds} , caused by the two periods when τ_a was most elevated. Note that in this example the regeneration rate is extremely fast to allow the effects of erosion and regeneration to be visualised using a single scale.

Time-series figures such as fig. 5.1 can benefit from being paired with *heat maps* where the relative quantity of material is rendered as colour for *all* modelled shear strength bands and every timestep, not just for a select few shear strengths. An example heatmap is shown in fig. 5.2, which corresponds to the simulation depicted in fig. 5.1. Such plots allow temporal changes in material quantity over the entire modelled shear strength range to be visualised.

5.2.2 Second example: signal transformation due to advection and mixing; the peak turbidity occurring one turnover after a shear stress increase

This example further demonstrates how advection and mixing transform the material release signal into the turbidity signal detected at the downstream end of the pipe. A step increase in shear stress results in an increase in the rate of mass release from the pipe wall (fig. 5.3). Observe that the downstream turbidity profile is visually quite different from the material release rate profile; this is due to the latter being transformed into the former by advection and mixing.

The change in concentration of each Lagrangian water parcel over time is shown in fig. 5.4. One can see that, following the increase in shear stress, parcel concentrations increase as they move down the pipe due to the material release rate being finite and them continually incorporating more material for a period. This explains why the peak turbidity occurs exactly one turnover after the step increase in shear stress (see fig. 5.3): the water parcel that is at the upstream end of the modelled pipe at the time of the step increase has received more eroded material than any other by the time it reaches the downstream end.

5.2.3 Third example: partial erosion and compounded turbidity responses

In a third example, shown in the time-series fig. 5.5 and heatmap fig. 5.6, one can see that short periods of elevated shear stress can result in the incomplete erosion of material (see fig. 5.5-ii where $t \in [1500, 3100] s$; this requires that the erosion rate factor β_e is not low relative to the duration of those periods). As per the first example (§5.2.1), the regeneration rate β_r was made very fast to allow the effects of erosion and regeneration to be visualised using a single scale. Also, note that all depicted material shear strengths regenerate simultaneously at the same rate up to maximal level over the period $t \in [3100, 7500] s$; these shear strengths are then eroded by the step increases in τ_a over $t \in [7500, 10000] s$. Thirdly, observe that there were several separate releases of material from the pipe wall over that period (fig. 5.5-iii) yet these responses

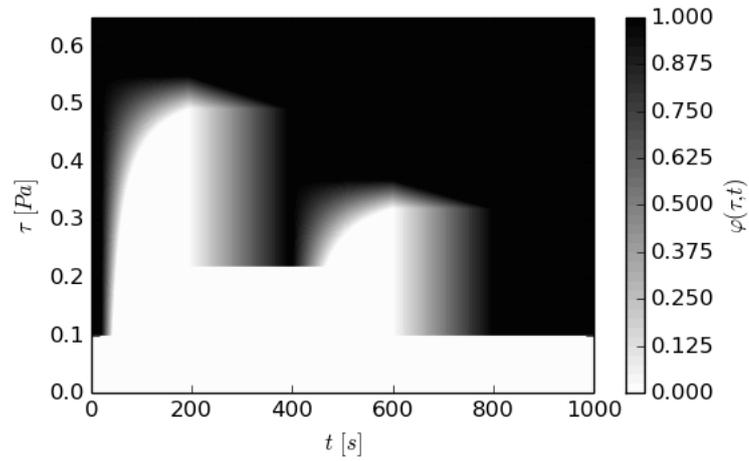


Figure 5.2: Variation in the relative amount of wall-bound material $\varphi(\tau_i, t_j)$ with respect to discrete shear strength τ_i and time t_j for the simulation results shown in fig. 5.1. White indicates shear strengths that are completely depleted (i.e. $\varphi(\tau_i, t_j) = 0$); black indicates maximal material at a shear strength (i.e. $\varphi(\tau_i, t_j) = 1$).

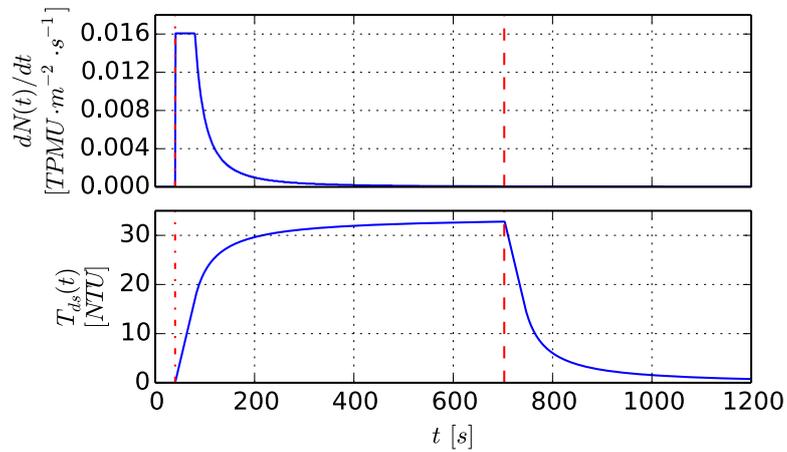


Figure 5.3: The material release dN/dt and downstream turbidity time-series T_{ds} profiles that result from a step increase in shear stress τ_a in a single pipe (where all shear strengths affected by this increase were maximally regenerated over $t \in [0, 40]$). The step increase occurs at 40 s, shown by a dashed and dotted line, and that time plus the duration of one pipe turnover is shown with a dashed line.

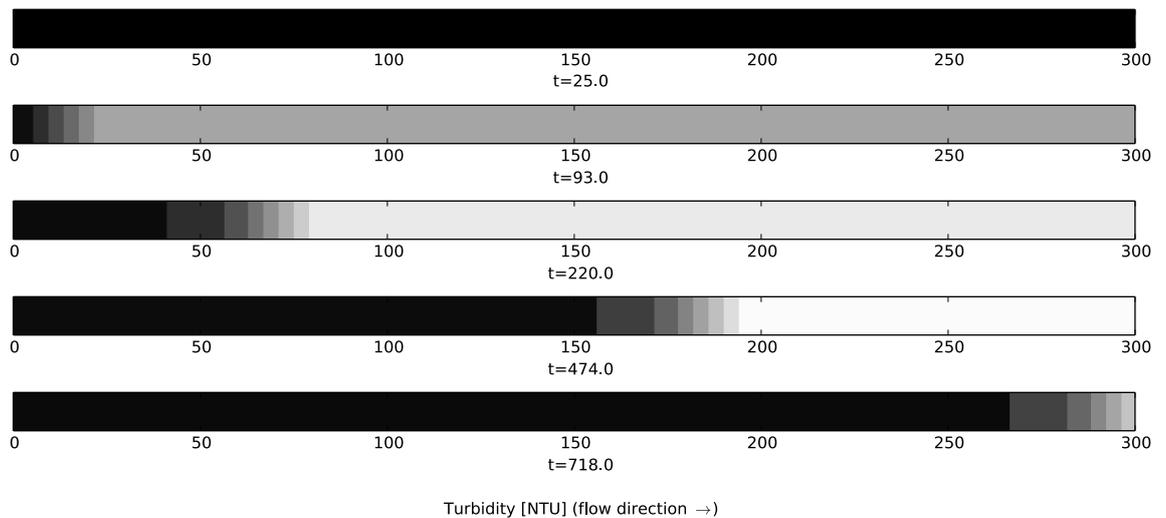


Figure 5.4: Variation in turbidity with distance along the single pipe referenced in fig. 5.3 during a material erosion event. Here the discretely tracked Lagrangian water parcels have been shaded to indicate their turbidity. Black corresponds to zero turbidity whereas white corresponds to the maximum turbidity experienced during the whole event (the same level as the peak in the second subplot of fig. 5.3).

had become far less distinct by the time they had been advected to the downstream end of the modelled pipe (fig. 5.5-iv). This illustrates how the processes of advection and mixing obscure and limit our understanding of state and change at the pipe wall.

5.2.4 Fourth example: a comparison of the **VCDM** and **PODDS** models

The simulation results shown in fig. 5.7 demonstrate how the proposed **VCDM** model and the **PODDS** model represent the wall state and the regeneration process differently and, importantly, illustrates how the turbidity predictions of these two models differ given the same observational data as inputs.

The simulation begins with a ‘flushing phase’, during which the shear stress increases in steps of equal magnitude. Both models show material being mobilised for each step increase, illustrating the desired characteristic that the material at the pipe wall has a strength profile. Note that, when modelling the erosion of maximally regenerated material, the output of the two models is similar.

During a subsequent, lengthy ‘growth phase’, the shear stress remains at a low, constant value, permitting the regeneration of shear strengths that had been eroded during the preceding flushing phase. Finally, there is a second flushing phase; this has an identical shear stress profile to the first flushing phase.

The models are distinguished by their turbidity predictions for flushing phase 2. The **PODDS** model predicts that only the stronger shear strengths affected by erosion during flushing phase 1 will have been able to notably regenerate by the start of flushing phase 2. The **VCDM**, however, predicts a similar turbidity spike for each of the step increases in shear stress during flushing phase 2, demonstrating that different shear strengths have been regenerating simultaneously; it is this response that more accurately reflects observations on material accumulation (§2.1.13.3) and encapsulates the conceptual requirements for regeneration presented earlier (§4.3.3).

The changes of wall state (i.e. the strength versus relative material quantity profile) are shown for the **VCDM** and **PODDS** in figs. 5.8 and 5.9 respectively. Fig. 5.8 shows the weaker material regenerating at the same rate at all strengths simultaneously, as per the **VCDM** conceptual specification in §4.3. The **PODDS** regeneration mechanism shown in fig. 5.9 differs in that, at the start of the growth phase, the material has a greater shear strength than at the end, therefore material cannot be regenerating at the same rate at all shear strengths simultaneously.

Note that the **PODDS** simulation results shown in figs. 5.7 and 5.9 were not calculated using the EPANET-based implementation (§2.2.8.4) but using functionality included in the `pyvcdm` Python module (§4.5.3) for ease of simulation and sensitivity analysis.

5.2.5 Fifth example: Erosion being driven by the excess shear

This example demonstrates how the model responds to equal step increments in shear stress, in terms of both the amount and rate of erosion per step.

In the simulation depicted in fig. 5.10, a hypothetical pipe of interest has been operating under a constant flow, Q_{bg} , for long enough prior to $t = 0$ such that the material at the wall at $t = 0$

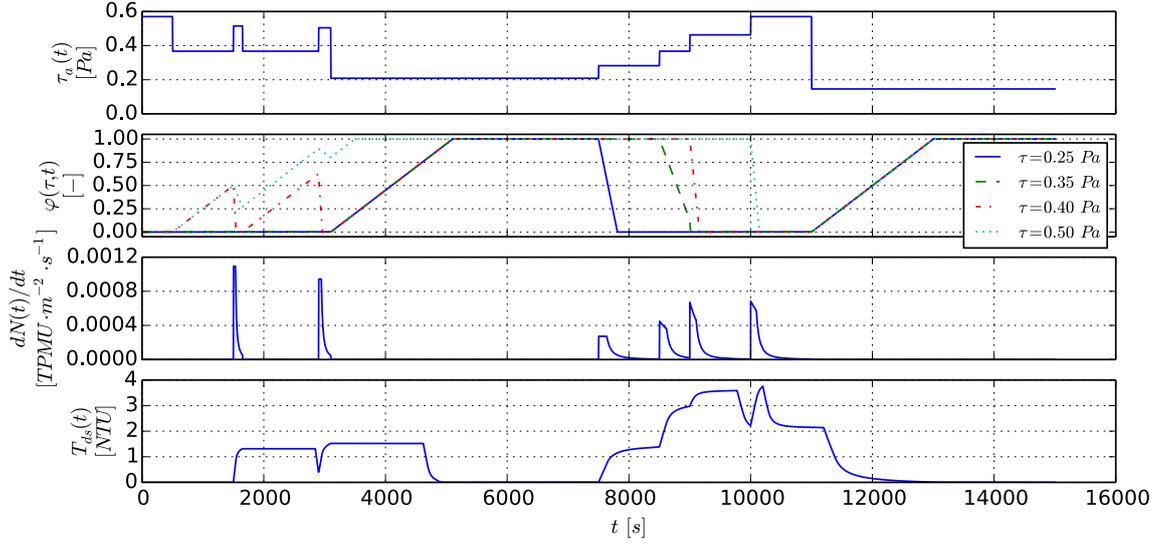


Figure 5.5: A third demonstration of the proposed model using artificial data. From top to bottom: (i) Time-series shear stress profile for a single pipe; (ii) relative material quantity $\varphi(\tau_i, t_j)$ for an illustrative selection of shear strengths $\tau_i \in \{0.25, 0.35, 0.40, 0.50\}$ Pa; (iii) rate of material release from the pipe wall $dN(t_j)/dt$; (iv) turbidity response at downstream end of pipe. Diameter $D = 200$ mm, roughness $k_s = 3.0$ mm and length $L = 350$ m; timestep $\Delta t = 1.0$ s; 1000 shear strength bands ($\Delta\tau = 0.0006$ Pa); if $\tau_i < \tau_a(t = 0)$ then $\varphi(\tau_i, t_j = 0) = 0$ else $\varphi(\tau_i, t_j = 0) = 1$ (see eqs. 4.3 and 4.4); erosion rate factor $\beta_e = 0.1$ Pa $^{-1} \cdot$ s $^{-1}$; regeneration rate $\beta_r = 0.005$ s $^{-1}$; scaling factor $\alpha = 1.0$ TMPU \cdot Pa $^{-1} \cdot$ m $^{-2}$; flow $Q \in [5.0, 10.0]$ L \cdot s $^{-1}$. As per fig. 5.1 the rate of material regeneration has been set to a high value to permit the effects of erosion and regeneration to be clearly visualised on a single plot.

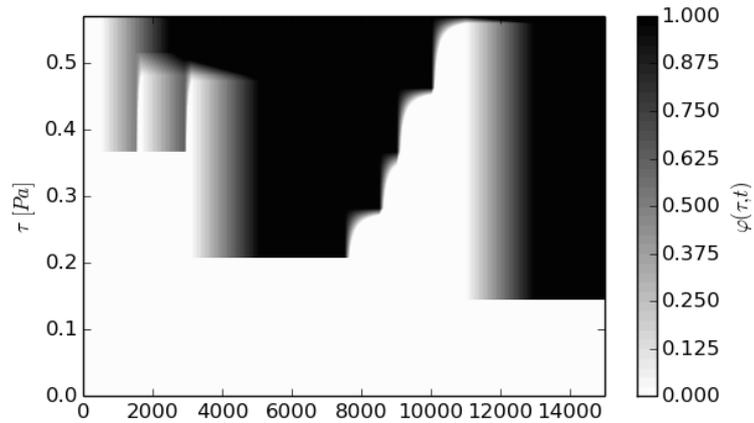


Figure 5.6: Variation in relative material condition $\varphi(\tau_i, t_j)$ with respect to discrete shear strength τ_i and time t_j for the simulation results shown in fig. 5.5. White indicates shear strengths that are completely depleted (i.e. $\varphi(\tau_i, t_j) = 0$); black indicates maximal material at a shear strength (i.e. $\varphi(\tau_i, t_j) = 1$).

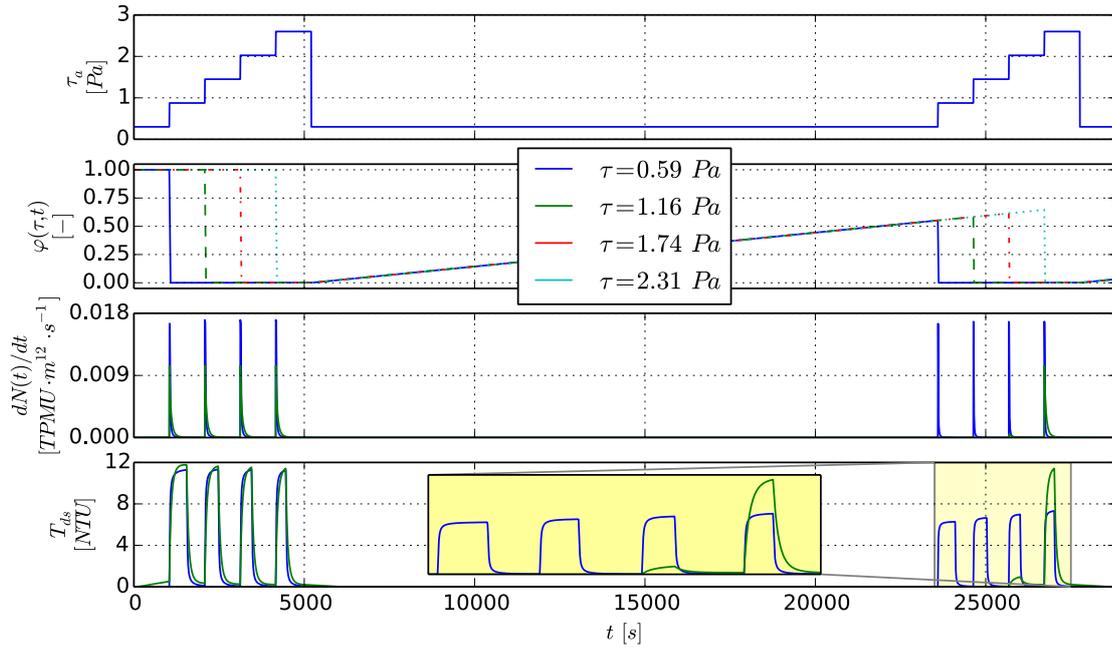


Figure 5.7: Simulations comparing the outputs from **VCDM** and **PODDS** models for a given set of pipe parameters ($D = 200 \text{ mm}$; $L = 200 \text{ m}$; $k_s = 3 \text{ mm}$) and a given hydraulic time-series profile ($dt = 3.47 \text{ s}$; $Q \in [7.2, 21.5] \text{ L} \cdot \text{s}^{-1}$; $\tau_a \in [0.3, 2.6] \text{ Pa}$). Details of the **VCDM** and **PODDS** model configurations shown in the respective captions of figs. 5.8 and 5.9. The subplot layout is as per fig. 5.1, albeit with subplots iii and iv including **PODDS** simulation outputs (green lines) in addition to those from the **VCDM** simulation (blue lines).

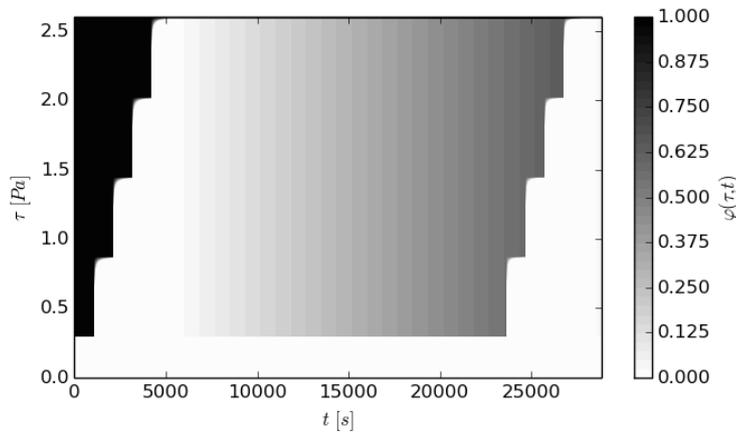


Figure 5.8: Change of the relative material quantity (wall state) with time for the **VCDM** simulation depicted in fig. 5.7. Model configuration: 1000 strengths tracked ($\Delta\tau = 0.0026$); boundary conditions: if $\tau_i < \tau_a(t = 0)$ then $\varphi(\tau_i, t_j = 0) = 0$ else $\varphi(\tau_i, t_j = 0) = 1$; $\beta_e = 0.1 \text{ Pa}^{-1} \text{ s}^{-1}$; $\beta_r = 0.00003 \text{ s}^{-1}$; $\alpha = 1.0 \text{ TPMU} \cdot \text{Pa}^{-1} \cdot \text{m}^{-2}$.

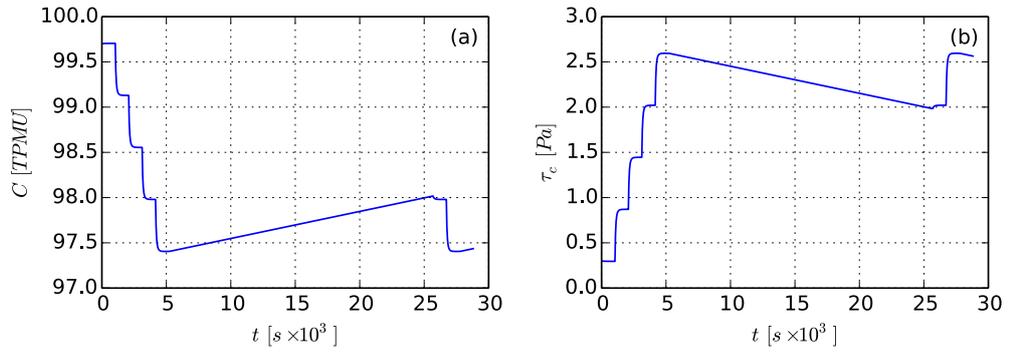


Figure 5.9: Change of (a) the total amount and (b) the scalar shear strength of discolouration material at the wall with time for the *PODDS* simulation depicted in fig. 5.7. Model configuration: boundary condition $\tau_c(t=0) = \tau_a(t=0) \text{ Pa}$; $k = -1.0 \text{ TPMU} \cdot \text{Pa}^{-1}$, $P = 0.02 \text{ TPMU} \cdot \text{m}^{-2} \cdot \text{Pa}^{-1} \cdot \text{s}^{-1}$; $n = 1.2$; $C_{max} = 100.0 \text{ TPMU}$; regeneration parameters: $P' = 0.00003 \text{ TPMU} \cdot \text{m}^{-2} \cdot \text{Pa}^{-1} \cdot \text{s}^{-1}$ but temperature shear strength power terms were both zero (see *Boxall and Saul (2005)* for details of the *PODDS* regeneration model's formulations and parameters).

is in a steady state and all shear strengths are either fully depleted or maximally regenerated. More concisely, if the shear stress at Q_{bg} is τ_{a_bg} then

$$\varphi(\tau, t=0) = \begin{cases} 0 & \tau \leq \tau_{a_bg} \\ 1 & \tau > \tau_{a_bg} \end{cases} \quad (5.1)$$

During the simulation the shear stress is increased in steps of equal magnitude above τ_{a_bg} , with each step causing the complete depletion of the corresponding shear strength range. This in turn results in a turbidity spike being seen at the downstream end of the hypothetical pipe. The turbidity response per step increase is partly characterised by the amount of material mobilised per step increase, which should be a constant between all shear strengths if:

- τ_a was increased in equal increments i.e. there was a constant excess stress for each step;
- All shear strengths affected by the step increase were maximally regenerated beforehand;
- There was negligible regeneration over the duration of the simulation;

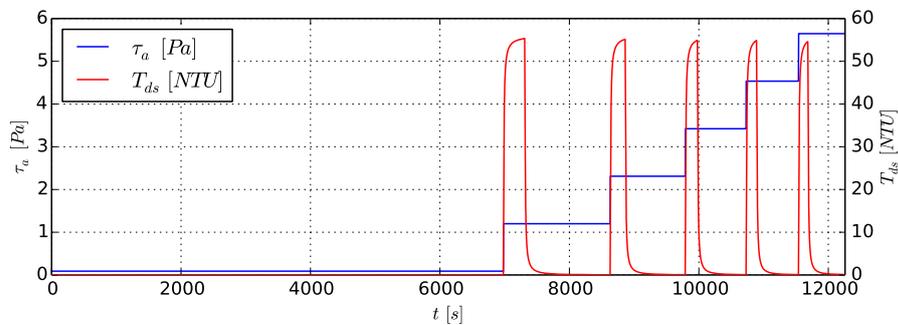


Figure 5.10: The turbidity response predicted by the *VCDM* for a series of step increases in shear stress of equal magnitude. $D = 79.3 \text{ mm}$; $L = 203.799 \text{ m}$; $k_s = 0.075 \text{ mm}$; $Q(t) \in [0.72, 7.00] \text{ L} \cdot \text{s}^{-1}$; $\tau_a(t) \in [0.09, 5.64] \text{ Pa}$; $\Delta\tau_a = 1.11 \text{ Pa}$; $\tau \in [0.09, 5.64] \text{ Pa}$; $\varphi(\tau_i, t_0) = 1$; five pipe turnovers per τ_a level; $\beta_e = 0.1 \text{ Pa}^{-1} \text{ s}^{-1}$; $\beta_r = 0.0 \text{ s}^{-1}$; $\alpha = 1.0 \text{ TPMU} \cdot \text{Pa}^{-1} \text{ m}^{-2} \text{ s}^{-1}$. The water passing into the upstream end, T_{us} , had a turbidity of 0 NTU .

- Each shear stress level was sustained until all affected shear stresses had been fully depleted and the turbidity reduced to background levels again;
- The model implementation is valid with regards to this test.

The step duration should be a linear function of the turnover time to ensure that each step is of sufficient duration for the turbidity to return to background levels before its end. This is because the peak turbidity and the rate of decay are functions of turnover time (and therefore flow) as seen in §5.2.2, which varies between steps.

Intuition suggests that the amount of discolouration material that is mobilised per shear stress step can be quantified by integrating turbidity with respect to time over the step duration. However, as mentioned in §2.1.13.2, if we are assuming that turbidity is proportional to TSS and so is a site-specific pseudo-concentration (§4.3.6.2) then this is an invalid measure of material quantity. Instead, the pseudo-concentration should be *multiplied by flow* to give the *material flux* past the turbidity measurement point per second; the integral of this derived time-series over the duration of each shear stress step is then a measure of the total amount of material mobilised over that step (eq. 2.14). Table 5.1 presents both metrics for the simulation depicted in fig. 5.10; the integral of turbidity with respect to time varies between shear stress steps but the integral of the product of turbidity and flow does not, as should be the case if the same amount of material was mobilised per shear stress step.

It is useful to be able to quantify discolouration material using the integral of material flux when working with observational data, such as when assessing regeneration rates by repeatedly flushing a given pipe length (§2.1.13.2). However, the VCDM model output should be turbidity rather than material flux as material flux is not a regulated parameter and turbidity is a better proxy for water consumers' perception. Turbidity is therefore presented in preference to material flux in the majority of the relevant modelling output plots throughout the remainder of this thesis.

The treatment of turbidity as a pseudo-concentration is not novel (the assumption was made in PODDS; see §2.2.8.2). However, previous attempts to quantify material regeneration rates using DMA flushing data have not always recognised that this assumption requires quantities of eroded material to be quantified using material flux when the flow varies within or between flushes (§2.1.13.2).

Returning to the simulation shown in fig. 5.10, an additional model characteristic that can be demonstrated is the similarity of material erosion rates following each shear stress step increase

Table 5.1: The amount of material mobilised per shear stress step in the VCDM simulation depicted in fig. 5.10 is a constant. The integral over the duration of each step of the product of flow and 'observed' turbidity reflects this (third column), but the integral over each step of just turbidity (second column) does not as it is not a valid metric for comparing quantities of discolouration material.

τ_a step	$\int_{step_start}^{step_end} T_{ds}(t) dt$	$\int_{step_start}^{step_end} T_{ds}(t) \cdot Q(t) dt$
1	18468.63	56.46
2	13002.01	56.61
3	10548.70	56.64
4	9113.00	56.83
5	8107.19	56.76

of constant magnitude above τ_{a_bg} . These consistent shear stress steps all give rise to the same excess shear stress, which should be directly proportional to the rate of erosion (eq. 4.3) for a verified model implementation.

If the turbidity responses from all step increases are offset so they all start from the same moment in time then one can see from (fig. 5.11) that:

- All responses follow a common upward-trending curve;
- Each response reaches a step-specific peak at the step-specific turnover time;
- Each response decays by following a curvature that is common to all steps but has a different temporal offset per step.

Discounting the effects due to differing rates of advection, the erosion rate can be seen to be constant for a unit of excess shear stress.

5.3 Sensitivity of the model to the mesh resolution

5.3.1 Introduction

Exploring the sensitivity of the output(s) of an empirical model to variation in inputs is important for several reasons. Firstly, there is a need to verify that the model output(s) do not scale aggressively/erroneously with the mesh resolution should the model require a discretised solution. The **VCDM** mesh resolution is defined by Δt and $\Delta \tau$.

Secondly, there is a need to determine an acceptable trade-off between mesh resolution and model accuracy. For example, a **VCDM** simulation featuring too few τ bands will be inaccurate (fig. 5.12); having a greater number of τ bands allows for a better assessment of the sensitivity of the model to variables that influence the shape of turbidity predictions. However, if $\Delta \tau$ is too small then τ bands will be nearly indistinguishable and computational requirements may be excessive.

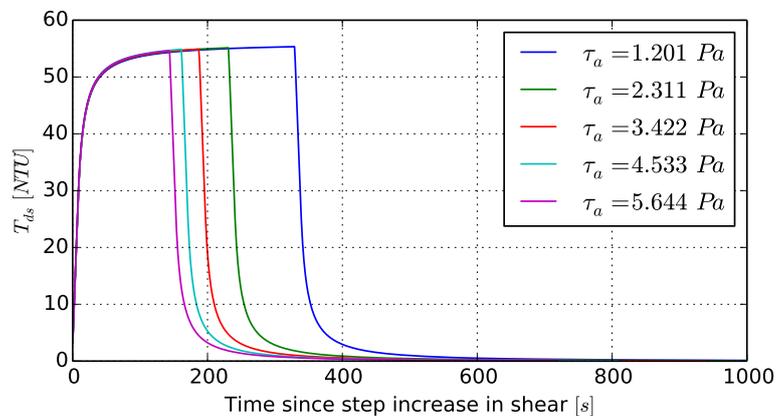


Figure 5.11: The set of turbidity responses shown in fig. 5.10 but offset so they all start from the same moment in time.

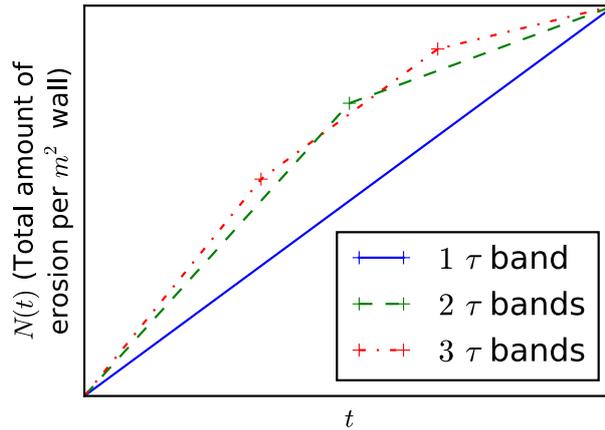


Figure 5.12: A *VCDM* simulation in which the relative amount of material at the pipe wall is tracked over time for only a small number of τ bands could be quite inaccurate, as illustrated in this diagram. The output will be sensitive to the discretisation interval for larger values of $\Delta\tau$.

Sensitivity analysis is also valuable for qualifying and quantifying how discretisation intervals (Δt ; $\Delta\tau$), model parameters (β_e ; β_r ; α), time-varying inputs (hydraulic quantities Q , τ_a and the excess shear stress τ_e ¹) and time-invariant inputs (pipe characteristics D , k_s and L) influence each other and the model outputs. The identification of parameter/input covariance is of importance as if say two parameters have a similar impact on model output(s) then the fitting of the model to data may be an indeterminate problem.

In this section, the relationships between model state/outputs and the mesh resolution are first qualified by analytical means, then these relationships are quantified and satisfactory mesh resolutions identified through numerical analysis. The expressions that result from the former are not only useful for studying the influence of mesh resolution but more generally explain how the various quantities pertaining to the model relate and also help identify characteristic times and shear strengths of the model.

5.3.2 Scalar measures of model outputs

The relationship between model inputs, parameters and outputs can be most easily qualified and quantified if vector and matrix outputs are aggregated as one or more scalars, each describing one characteristic of the outputs. For example: one metric might be a measure of scale, another might be a measure of rate and a third might be an integral over time.

The *VCDM* outputs have been reduced to the following four metrics, two of which are functions of the wall state sub-model and the others are functions of both the wall state *and* material transport sub-models:

- The total amount of material to be released from the pipe wall during a simulation can be determined directly from the output of the wall state sub-model (eq. 5.2, in which A_w is the pipe's internal surface area [m^2]). This allows the sensitivity of the total amount

¹Henceforth the excess shear stress shall be defined as $\tau_e(\tau, \tau_a(t)) = \min(\tau_a(t) - \tau, 0)$, abbreviated as $\tau_e(\tau, \tau_a)$ if t is obvious from the contents or just as τ_e if both τ and τ_a are obvious from the context.

of erosion to the mesh resolution to be explored independently of the material transport sub-model.

$$A_w \cdot \int_0^{t_{max}} \frac{dN(t)}{dt} dt \quad [TPMU] \quad (5.2)$$

- An expression for the total amount of material mobilised during a simulation can also be derived from the output of the material transport model (eq. 5.3), which in turn requires a solution to the wall state sub-model. Eqs. 5.2 and 5.3 are conceptually equivalent for a given simulation if the simulation is continued for long enough that all mobilised material passes out of the end of the pipe and the material transport sub-model is not a source of error.

$$\int_0^{t_{max}} Q(t) \cdot T_{ds}(t) dt \quad [TPMU] \quad (5.3)$$

- The maximum rate of material release from the pipe wall over the duration of the simulation, which is solely a product of the wall state sub-model:

$$\max_{t \in [0, t_{max}]} \frac{dN(t)}{dt} \quad [TPMU \cdot m^{-2} \cdot s^{-1}] \quad (5.4)$$

- The maximum turbidity predicted at the downstream end of the modelled pipe over the duration of the simulation, which is a function of both the wall state and material transport sub-models and, for non-trivial flow profiles, cannot be derived from just the wall state sub-model output, $dN(t)/dt$. The error associated with this metric is of particular interest to model users (see use case #1 in §4.2.1).

$$\max_{t \in [0, t_{max}]} T_{ds}(t) \quad [NTU] \quad (5.5)$$

5.3.3 Resolving the depletion of a shear strength band within a timestep for a simple modelling case

Consider a simple [VCDM](#) model where:

- The flow increases from one steady level to another instantaneously and immediately at $t = t_0$, causing an increase in shear stress, τ_a , from τ_{a_lo} to τ_{a_hi} ;
- The shape of the boundary condition profile (the relative amount of wall-bound material per shear strength at $t = t_0$) is very simple (eq. 5.6);

$$\varphi(\tau, t_0) = \begin{cases} 0 & \tau < \tau_{a_lo} \\ 1 & \tau \geq \tau_{a_lo} \end{cases} \quad (5.6)$$

- The result of the increase in τ_a is therefore that all material at the pipe wall where $\tau_{a_lo} \leq \tau \leq \tau_{a_hi}$ is completely eroded over the course of the simulation if the simulation is long enough i.e. the relative material quantity, φ , per shear strength $\tau \in [\tau_{a_lo}, \tau_{a_hi}]$ decreases from 1 to 0 during the period $t \in [t_0, t_{max}]$.

To be able to accurately resolve the erosion of a τ band τ_i from a maximal level to full depletion (i.e. $\Delta\varphi(\tau_i) = -1$) this erosion must occur over at least one simulation timestep. Given that the amount of erosion at τ_i over Δt is $\tau_e \cdot \beta_e \cdot \Delta t$ (eq. 4.3), Δt therefore needs to be small enough to satisfy the following inequality:

$$\begin{aligned} \tau_e(\tau_i, \tau_a) \cdot \beta_e \cdot \Delta t &\leq |\Delta\varphi(\tau_i, t_0)| \\ &\leq 1 \end{aligned} \quad (5.7)$$

This can be rearranged to give an expression for the maximum Δt for which the complete erosion of τ band τ_i can be resolved in Δt :

$$\Delta t \leq \frac{1}{\tau_e(\tau_i, \tau_a) \cdot \beta_e} \quad (5.8)$$

For the previously-defined modelling case, the Δt required for accurately resolving the erosion of discrete τ bands is most sensitive to the shear strength of the *weakest* τ band for which $\tau_e > 0$ and $\varphi(\tau, t_0) > 0$ (let this be referenced here as τ_{min}). This is because τ_e is largest for this shear strength, causing it to erode faster than all others:

$$\Delta t \leq \frac{1}{\tau_e(\tau_{min}, \tau_a(t_0)) \cdot \beta_e} \quad (5.9)$$

The right-hand-side of eq. 5.9 can be considered a characteristic time of the wall state sub-model. Let this quantity be known as t_d (see eq. 5.10). The d subscript serves as a reminder that this characteristic time relates to material *depletion*. Here $\tau_i = \tau_{min}$ and $t = t_0$.

$$t_d(\tau_i, \tau_a) = \frac{1}{\tau_e(\tau_i, \tau_a) \cdot \beta_e} \quad (5.10)$$

5.3.4 Sensitivity w.r.t. Δt : more complex boundary conditions

Let us briefly digress from considering the simple boundary conditions of eq. 5.6. To resolve the erosion of *partially regenerated* material when the shear stress increases, it is not sufficient for Δt to satisfy the inequality of eq. 5.9. For example. to resolve the erosion of the lowest, non-depleted shear strength band, τ_{min} , that is 15% regenerated at $t = t_0$ then $|\Delta\varphi|$ for τ_{min} can be no more than 0.15 during Δt (whereas for fully-regenerated bands $|\Delta\varphi|$ can be no more than 1). Therefore, in this case Δt must satisfy:

$$\Delta t \leq 0.15 \cdot t_d(\tau_{min}, \tau_a) \quad (5.11)$$

This upper bound for Δt may be accurate but could be overly sensitive to the erosion of a very small (potentially negligible) amount of weakly-bound material (see the shaded region

$\varphi(\tau_i \leq \tau \leq \tau_j, t_0)$ in fig. 5.13a). This is due to φ being a monotonically increasing function with regards to τ (§4.4.4). The result of this could be an unnecessarily-small upper bound to Δt . Also, this approach requires an accurate estimate of $\varphi(\tau_{min}, t_0)$, which may not be trivial to estimate from historic shear stress data for $t < t_0$.

This approach could be made less sensitive to the erosion of a small amount of material if the critical relative amount of material that must be resolved is the mean over the shear strength range from the minimum τ where $\varphi(\tau, t_0) > 0$ up to the minimum τ where $\varphi(\tau + \Delta\tau, t_0) = 1$ (dotted line in fig. 5.13b):

$$\Delta t \leq \frac{\int_0^{\tau_j + \Delta\tau} \varphi(\tau, t_0) d\tau}{(\tau_j + \Delta\tau) \cdot \tau_e(\tau_i, \tau_a(t_0)) \cdot \beta_e} \quad (5.12)$$

$$\leq \frac{\int_0^{\tau_j + \Delta\tau} \varphi(\tau, t_0) d\tau}{\tau_j + \Delta\tau} \cdot t_d(\tau_i, \tau_a(t_0)) \quad (5.13)$$

However, now an understanding of how the historic τ_a profile has impacted on the boundary conditions $\varphi(\tau, t_0)$ is required for an even greater τ range than for the previous approach.

A third option for determining a rule for an upper bound to Δt is to simply strive to resolve the depletion of shear strength bands for which the relative amount of material at the wall is greater than a threshold γ i.e. $\varphi(\tau, t_0) \leq \gamma$ where $\gamma \in [0, 1]$. For example, one may wish to resolve the erosion of all bands that are at least to 10% regenerated (so $\gamma = 0.1$; see fig. 5.13c).

An advantage of this option is that finding an upper bound for Δt does not depend on already having an accurate estimate of $\varphi(\tau, t_0)$; one only requires an estimate of $\varphi(\tau_{min}, t_0)$, where τ_{min} is the lowest shear stress that is not depleted at t_0 . Using this approach the upper bound for Δt becomes:

$$\Delta t \leq \gamma \cdot t_d(\tau_{min}, \tau_a(t_0)) \quad \gamma \in [0, 1] \quad (5.14)$$

This third approach is to be the one taken here due to its pragmatic simplicity. It has the additional advantage of placing limits on the computational effort required to conduct a simulation: approaches (a) and (b) are more likely to identify a low upper bound for Δt that results in a unwieldy number of calculations being required over a lengthy time series simulation. For the modelling case outlined at the start of §5.3.3, γ is of course 1 (eq. 5.9).

Note that although the inequalities in eqs. 5.9 and 5.14 describe upper bounds of Δt that afford the accurate resolution of the erosion of τ bands, numerical sensitivity analysis is still required to determine how model outputs vary with Δt and decide upon heuristics for selecting an appropriate Δt for a given modelling scenario.

5.3.5 Sensitivity w.r.t. Δt : dependence on β_e and τ_e

When considering how the VCDM state and output metrics (§5.3.2) vary with the simulation timestep, Δt , for the modelling case described in §5.3.3, one must be aware that these

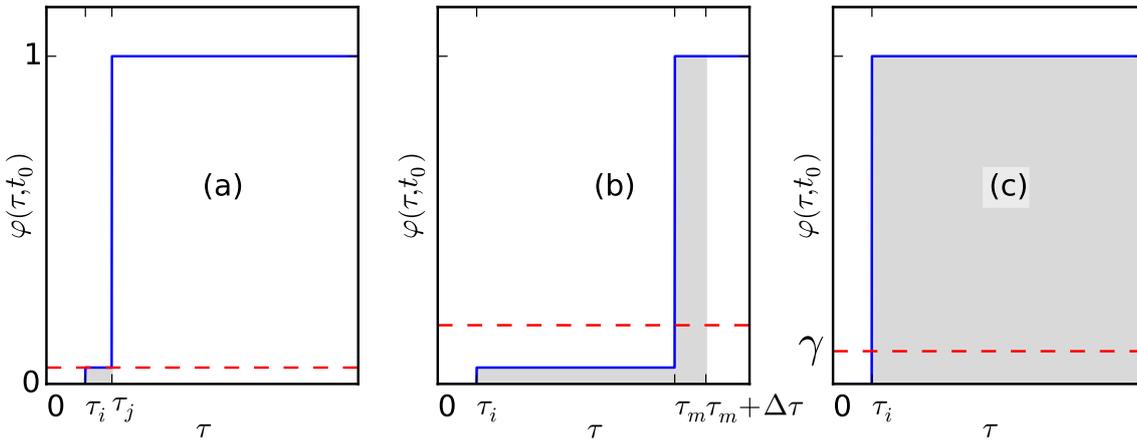


Figure 5.13: When deciding which simulation timestep values allow for the adequate resolution of the depletion of shear strength bands, there are several ways of defining the critical relative material quantity (dashed red line) that informs this rule. Three options are shown in this conceptual diagram. Options a and b derive this critical quantity using estimates of the wall state boundary conditions $\varphi(\tau, t_0)$ arrived at by studying historic shear stress time-series. Alternatively, this critical quantity can just be a threshold γ e.g. $\gamma = 10\%$ of the maximum material quantity per shear strength band, as shown in (c).

relationships are dependent upon:

- β_e , as this is a rate parameter (with units of $[Pa^{-1} \cdot s^{-1}]$);
- τ_e , the excess shear stress, as a larger τ_e results in more rapid erosion (eq. 4.3) and therefore a smaller Δt is needed to accurately represent the steeper material erosion curve.
- β_r , but this is largely irrelevant as a value of Δt that provides sufficient accuracy for erosion will almost certainly offer accurate temporal resolution for modelling regeneration given that the latter is a much slower process (erosion to depletion occurs in only a few pipe turnovers (§2.1.12.2; §2.1.13.4) yet regeneration takes 1.5 to 4 years (§2.1.14)).

Any rules or heuristics for selecting a value of Δt that gives a desired level of model accuracy should therefore be independent of β_e and τ_e to ensure that those rules are generic. The impact of Δt can be studied independently of β_e and τ_e if t and Δt are expressed relative to the wall state characteristic time t_d (i.e. non-dimensionally) rather than as absolute values. This independence can be demonstrated by plotting the material release rate versus absolute and non-dimensional time for the modelling scenario presented at the start of §5.3.3. The differing curves in fig. 5.14a show that dN/dt versus t varies with β_e , but these curves all collapse to a single curve if dN/dt is normalised by $\max(dN/dt)$ and t is normalised by t_d (i.e. time is plotted as $t_d(\tau_{a_hi}, \tau_{a_lo}) = t \cdot (\tau_{a_hi} - \tau_{a_lo}) \cdot \beta_e$; fig. 5.14b). Curves of dN/dt versus t for various τ_e (achieved by varying τ_{a_hi} between simulations; fig. 5.15a) also collapse to a single curve if normalised in the same manner (fig. 5.15b).

A consequence of Δt being dependent on β_e and τ_e is that to appraise or select a value of Δt for a given modelling scenario one must know β_e and τ_e . The implications of this will be discussed subsequently.

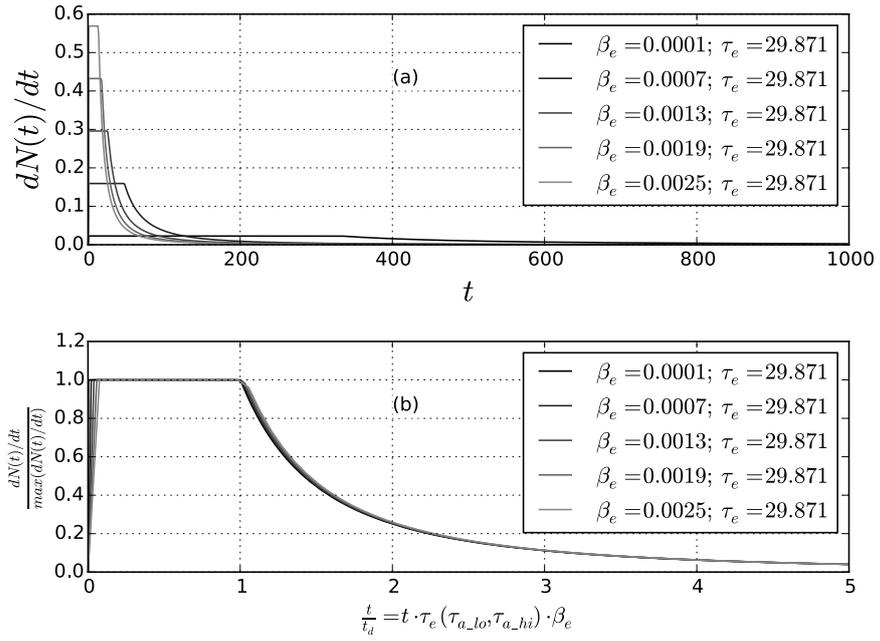


Figure 5.14: For a step increase in τ_a from τ_{a_lo} to τ_{a_hi} at $t = t_0$ and simple **VCDM** boundary conditions ($\varphi(\tau < \tau_{a_lo}, t_0) = 0$ and $\varphi(\tau \geq \tau_{a_lo}, t_0) = 1$) (a) shows curves of dN/dt resulting from simulations with various values of β_e that collapse to a single curve (b) if dN/dt is normalised by $\max(dN/dt)$ and t is normalised by $t_d(\tau_{a_lo}, \tau_{a_hi})$ (b). Note that in the legend $\tau_e \equiv \tau_{a_hi} - \tau_{a_lo}$.

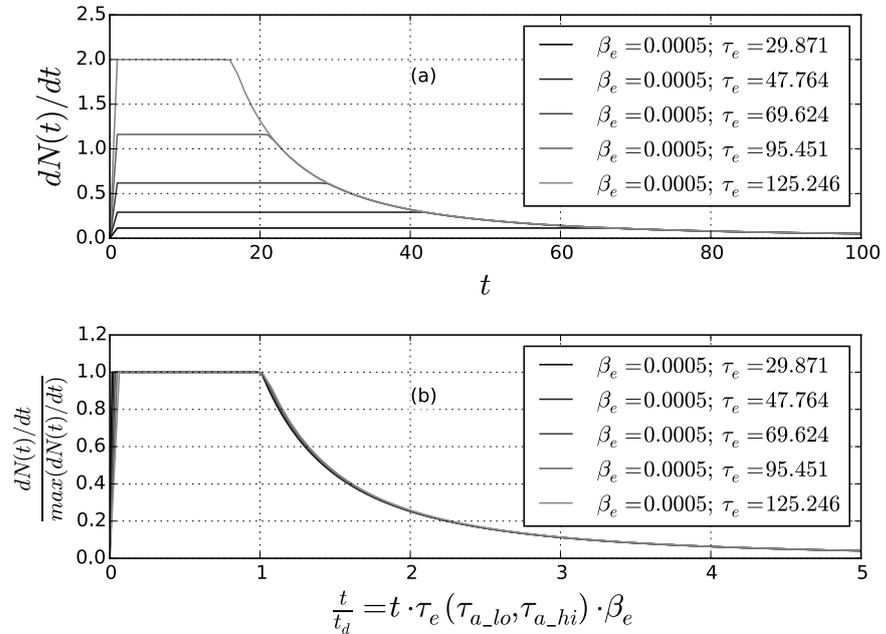


Figure 5.15: For a step increase in τ_a from τ_{a_lo} to τ_{a_hi} at $t = t_0$ and simple **VCDM** boundary conditions ($\varphi(\tau < \tau_{a_lo}, t_0) = 0$ and $\varphi(\tau \geq \tau_{a_lo}, t_0) = 1$) (a) shows curves of dN/dt resulting from simulations with various values of τ_e at t_0 that collapse to a single curve (b) if dN/dt is normalised by $\max(dN/dt)$ and t is normalised by $t_d(\tau_{a_lo}, \tau_{a_hi})$ (b). Note that in the legend $\tau_e \equiv \tau_{a_hi} - \tau_{a_lo}$.

5.3.6 Sensitivity w.r.t. Δt : complex flow profiles

So long as Δt is chosen/appraised using the maximum possible/achievable shear stress for the pipe of interest as an estimate of τ_e (§5.3.5) then it does not matter if or how pipe flow (and therefore shear stress) vary with t as the most limiting case has been considered.

5.3.7 Relationship between the mesh discretisation steps

Before progressing to numerically exploring the sensitivity of the model with respect to Δt or $\Delta\tau$, it is first necessary to consider whether the mesh discretisation intervals Δt and $\Delta\tau$ are *covariant* with regards the model outputs. If they are independent, then the sensitivity to Δt can be assessed independently of $\Delta\tau$ and vice versa.

For a given discrete shear strength band τ_i , bounded by τ_{i_lb} and τ_{i_ub} (with difference $\Delta\tau$), the excess shear stress is calculated as follows:

$$\tau_e(\tau_i, \tau_a) = \max\left(\tau_a - \frac{\tau_{i_lb} + \tau_{i_ub}}{2}, 0\right) \quad (5.15)$$

If, for the modelling case at the start of §5.3.3, Δt must be $\leq t_d = 1/\tau_e(\tau_i, \tau_a) \cdot \beta_e$ for the complete erosion of a shear strength band τ_i to be resolvable in Δt (eq. 5.9) then the upper bound for Δt (given this expression) is most restrictive for larger value of τ_e . The bounds of the weakest shear strength band, τ_1 , (with width $\Delta\tau$) are therefore the most critical for determining Δt .

Looking at fig. 5.16 one can see that for the subplot where the τ discretisation is more fine-grained (lower subplot) the τ_e for the weakest shear strength band τ_1 is larger. τ_e therefore varies with $\Delta\tau$. An implication of this is that Δt is dependent on $\Delta\tau$: the minimum Δt for the resolution of complete shear strength band erosion depends on τ_e (eq. 5.9) and τ_e varies with $\Delta\tau$. This relationship is shown in eq. 5.16, where eq. 5.15 has been substituted into eq. 5.9.

$$\Delta t = \frac{1}{\tau_e(\tau_i, \tau_a(t)) \cdot \beta_e} = \frac{1}{\min\left(\tau_a(t) - \frac{\tau_{i_lb} + \tau_{i_ub}}{2}, 0\right) \cdot \beta_e} \quad (5.16)$$

However, this relationship is weak if $\Delta\tau$ is small as there is little change in τ_e as the number of τ bands is increased for $\Delta\tau \ll \tau_a$. The mesh discretisation intervals Δt and $\Delta\tau$ can therefore be considered to be independent if $\Delta\tau \ll \tau_a$.

5.3.8 Smaller values of $\Delta\tau$ result in lengthier material release profiles

A second way in which model outputs are dependent on $\Delta\tau$ is that for smaller values of $\Delta\tau$ the material release rate (and therefore the downstream turbidity) takes more time to decay to zero following a step increase in shear stress. This can be a source of error when evaluating metrics by integrating over time (i.e. eqs. 5.2 and 5.3) for simulations of a finite duration.

The relative erosion rate, $\beta_e \cdot \tau_e$, is lowest for a shear strength band τ_j with a mid-point only just less than τ_a i.e. as τ_e approaches zero. The time required for the complete depletion of

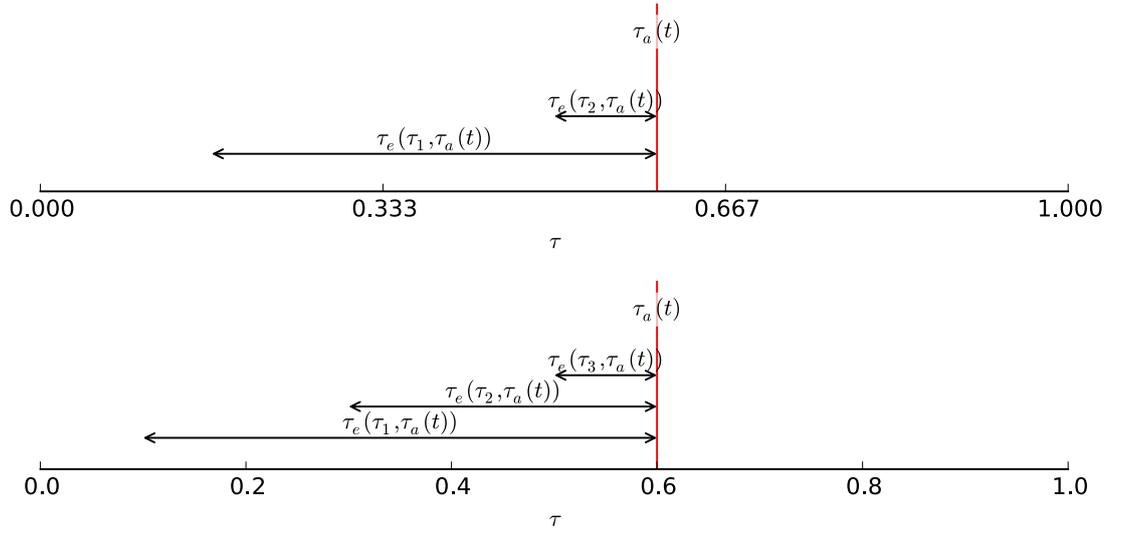


Figure 5.16: For a given shear stress τ_a at time t the excess shear stress $\tau_e(\tau_i, \tau_a)$ varies depending how coarsely the modelled shear strength range is discretised, as can be seen in this illustrative diagram.

band τ_j can be found as follows (note that τ_a is invariant after the step increase at t_0). Firstly, given:

$$|\Delta\varphi(\tau_j, t \in [t_0, t])| = 1 = \beta_e \cdot \tau_e(\tau_j, \tau_a) \cdot t \quad (5.17)$$

This can be rearranged to give:

$$t = \frac{1}{\beta_e \cdot \tau_e(\tau_j, \tau_a)} \quad (5.18)$$

If τ_j is the shear strength band for which the mid-point is closest to but less than τ_a then that mid-point is:

$$\tau_j = \text{floor}\left(\frac{\tau_a}{\Delta\tau} - \frac{1}{2}\right) \cdot \Delta\tau + \frac{\Delta\tau}{2} \quad (5.19)$$

This is because:

- $\text{floor}(\tau_a/\Delta\tau)$ is the number of whole τ bands below τ_a ;
- $\text{floor}(\tau_a/\Delta\tau) \cdot \Delta\tau$ is the upper bound of the corresponding shear strength range;
- the $-1/2$ and $+\Delta\tau/2$ are necessary as we want to find the τ band *mid-point* closest to τ_a .

Substituting into the previous equation gives the time required for shear strength band τ_j to become fully depleted:

$$t = \frac{1}{\beta_e \cdot \tau_e\left(\text{floor}\left(\frac{\tau_a}{\Delta\tau} - \frac{1}{2}\right) \cdot \Delta\tau + \frac{\Delta\tau}{2}, \tau_a\right)} \quad (5.20)$$

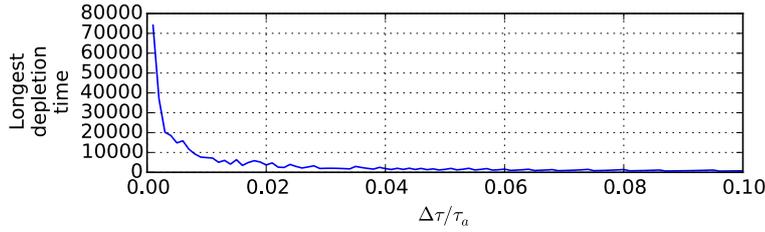


Figure 5.17: The time (in s) required for the complete depletion of the slowest eroding shear strength band given different shear strength discretisation step sizes (see eq. 5.20; here simulations were conducted with $\beta_e = 0.002 \text{ Pa}^{-1} \cdot \text{s}^{-1}$, $\tau_a = 4.5 \text{ Pa}$ and $\Delta\tau \in [\tau_a/1000, \tau_a/10]$).

This time will be greater for smaller values of $\Delta\tau$, as can be seen in fig. 5.17.

For a simulation over the period t_0 to t_{max} the model output metrics eqs. 5.2 and 5.3 will both be under-estimates if t , calculated using eq. 5.20, exceeds t_{max} .

A simple yet not entirely satisfactory method to reduce the impact of $\Delta\tau$ on eqs. 5.2 and 5.3 is to only evaluate the integrals in those equations for times where dN/dt and T_{ds} (respectively) are greater than a relative threshold, such as $(\max(dN/dt))/1000$ or $(\max(T_{ds}))/1000$, over t_0 to t_{max} . This will eliminate the long tails that lengthen as $\Delta\tau$ gets smaller. This treatment of long tails is primarily of interest when conducting numerical sensitivity analysis; it is far less relevant to the modeller when studying real pipes as the signal-to-noise ratio (SNR) is likely to dominate the turbidity due to erosion when the latter drops by a couple of orders of magnitude following an erosion event.

5.3.9 Sensitivity w.r.t. Δt and $\Delta\tau$: numerical exploration

5.3.9.1 General approach

The aim here is to qualify and quantify how evaluations of the model output metrics listed in §5.3.2 change as Δt and $\Delta\tau$ vary in size between justifiable upper bounds and lower bounds that approach the infinitesimal.

The identified, quantified relationships need to ideally be independent of simulation inputs and parameters (e.g. pipe length or the α parameter) so that this knowledge can be used to estimate the discretisation error associated with *any* model configuration.

To achieve this, one can first non-dimensionalise the numerical metrics listed in §5.3.2 by normalising by entirely analytical expressions that are conceptually equivalent for a particular model configuration. These non-dimensionalised terms can then be plotted against a non-dimensional measure of Δt or $\Delta\tau$ (e.g. $\Delta t/t_d$, $\Delta t/t_t$ or $\Delta\tau/\tau_e(\tau_{min}, \tau_a(t))$) to elucidate the (ideally linear but potentially non-linear) relationship between the relative error in that term and the non-dimensional mesh resolution dimension. This error relationship is independent of the model configuration (inputs and parameters) so long as the conditions of conceptual equivalence between the two components of the non-dimensional model output metric remain valid. For example:

$$\frac{\text{numerical output metric}}{\text{equivalent analytical metric}} = \text{relative err func} \left(\frac{\Delta t}{\text{a characteristic time}} \right)$$

given conditions $\{A, B, C\}$

5.3.9.2 Error in the total amount of material released from the pipe wall

The metric eq. 5.2 must be evaluated numerically and is dependent on both Δt and $\Delta \tau$. This metric can be non-dimensionalised by considering the simple model configuration described in §5.3.3.

Let there be a step increase in shear stress from τ_{a_lo} to τ_{a_hi} that, over the period $t \in [t_0, t_{max}]$, completely erodes only the material that had previously been maximally regenerated. The relative change in material quantity per shear strength τ_i over $t \in [t_0, t_{max}]$ is:

$$\Delta \varphi (\tau_i, t \in [t_0, t_{max}]) = 1 \quad (5.21)$$

Therefore the absolute change in material quantity over $\tau \in [\tau_{a_lo}, \tau_{a_hi}]$, given eqs. 4.3, 4.4 and 4.5, is:

$$\Delta N = \alpha \int_{\tau_{a_lo}}^{\tau_{a_hi}} \Delta \varphi (\tau_i, t \in [t_0, t_{max}]) \Delta \tau \quad (5.22)$$

$$= \alpha (1 \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})) \quad (5.23)$$

This is the total amount of material released per m^2 of wall area during $t \in [t_0, t_{max}]$. The error that results from the numerical evaluation of eq. 5.2 (using, say, the trapezoid rule) can therefore be non-dimensionalised as follows:

$$\frac{A_w \cdot \int_0^{t_{max}} \frac{dN(t)}{dt} \Delta t}{A_w \cdot \alpha \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})} = \frac{\int_0^{t_{max}} \frac{dN(t)}{dt} \Delta t}{\alpha \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})} \left[\frac{TPMU \cdot m^{-2} \cdot s^{-1} \cdot s}{TPMU \cdot Pa^{-1} \cdot m^{-2} \cdot Pa} \right] = [-] \quad (5.24)$$

This expression is only valid if t_{max} is large enough so that all shear strengths affected by the increase in τ_a are fully depleted before $t = t_{max}$. The peak turbidity always occurs at the turnover time, t_t , following a step increase in τ_a , so t_{max} should therefore here be a large multiple of t_t .

Results For a model configuration as per §5.3.3, the scalar eq. 5.24 was first plotted against the non-dimensional simulation timestep, $\Delta t/t_d$, when evaluated for a variety of timestep sizes t_d/i where $i \in \{16, 15, \dots, 2, 1, 0.5\}$. Here $\tau_e = \tau_{a_hi} - \tau_{a_lo}$. This process was then repeated for a variety of multipliers for τ_e , L , D , k_s , β_e and α (all results shown in fig. 5.18a) to confirm non-dimensionality. τ_e was varied by scaling $\tau_{a_lo} = \tau_a (t < t_0)$. The integral in eq. 5.24 was evaluated using the trapezoid method over $t \in [t_0, 30 \cdot t_t]$, ignoring times where

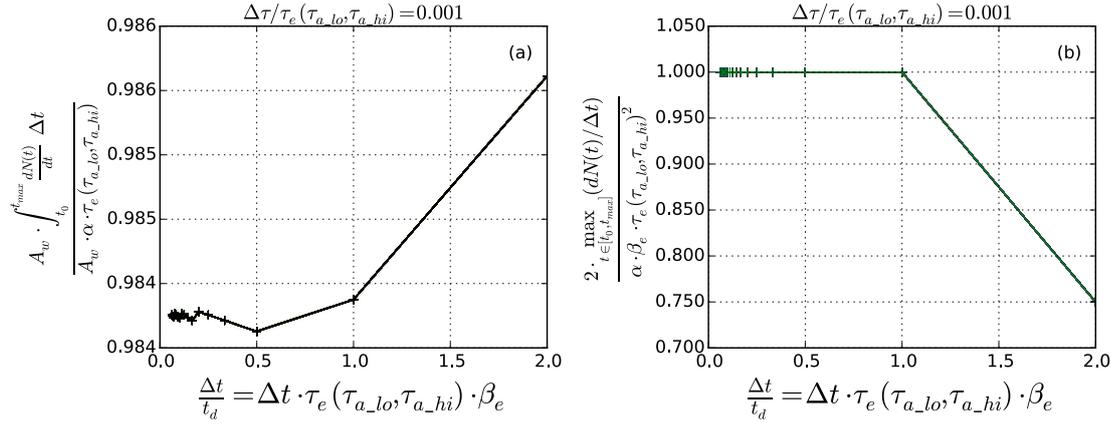


Figure 5.18: For a simple simulation where the shear stress increases from τ_{a_lo} to τ_{a_hi} at $t = 0$ (a) shows the non-dimensionalised (relative) error from discretely calculating the total amount of material to have been released from the pipe wall over the simulation (with caveats detailed in the body of the text) whereas (b) shows the relative error from discretely calculating the maximum rate of material release from the pipe wall during the simulation. The sensitivity of these two relative metrics was explored with regards to a normalised measure of timestep whilst the shear strength step was held constant. Simulations were conducted for various multipliers of $\tau_e (\tau_{a_lo}, \tau_{a_hi})$, L , D , k_s , β_e (all results plotted here) to confirm that the identified error relationship was not dependent on any of those five factors.

$dN/dt < \left(\max_{t \in [t_0, t_{max}]} dN/dt \right) / 1000$ for the reason outlined in §5.3.8. $\Delta\tau$ was fixed at $\tau_e/1000$.

Δt was found to have a very small influence on the discretisation error in eq. 5.2, with this error being larger for $\Delta t/t_d > 1$ (fig 5.18a). The effort to non-dimensionalise the model output metric were successful as all lines for different values of τ_e , L , D , k_s , β_e and α collapsed onto the same line. As can be seen, there was a constant error of $\approx 1.6\%$ due to ignoring the ‘long tail’ when evaluating the integral.

The sensitivity of eq. 5.24 was also explored with respect to $\Delta\tau \in \{\tau_e/i\}$ where $i \in \{5, 10, 20, 50, 100, 200, 400, 600, 800\}$ i.e. τ_e was considered to be a characteristic shear strength of the system and sensitivity was investigated for values of $\Delta\tau$ that were whole fractions of the excess shear stress. The timestep was set to a fixed value $\Delta t = t_t/256$. When evaluating T_{ds} , times when $dN/dt < \left(\max_{t \in [t_0, t_{max}]} dN/dt \right) / 1000$ were again ignored when integrating over time.

The error in eq. 5.24 was found to be low (1.7%) for $\Delta\tau/\tau_e < 0.05$ and negligible for greater values of $\Delta\tau/\tau_e$ (fig. 5.19a).

5.3.9.3 Error in the maximum rate of material release from the pipe wall

The other output metric that is solely a function of the wall release sub-model and not of the material transport sub-model is eq. 5.4. The error associated with evaluating this metric numerically for a given mesh resolution can also be non-dimensionalised in the context of the simple model configuration in §5.3.3.

If $t < t_d$ (again, $\tau_e = \tau_{a_hi} - \tau_{a_lo}$) then for $t \in [t_0, t_0 + dt]$ the change in the relative amount of material at the pipe wall is the dark blue shaded area in fig. 5.20.

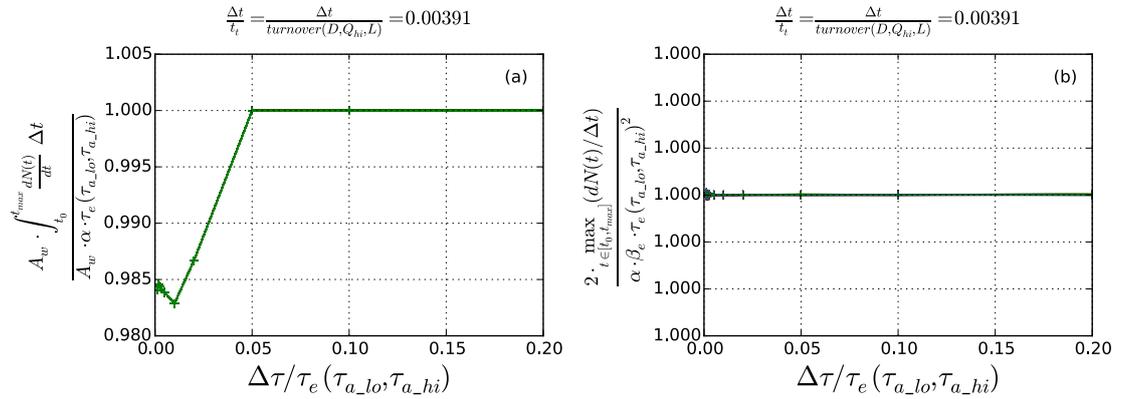


Figure 5.19: For a simple simulation where the shear stress from τ_{a_lo} to τ_{a_hi} at $t = 0$ (a) shows the non-dimensionalised (relative) error from discretely calculating the total amount of material to have been released from the pipe wall over the simulation (with caveats detailed in the body of the text) whereas (b) shows the relative error from discretely calculating the maximum rate of material release from the pipe wall during the simulation. The sensitivity of these two relative metrics was explored with regards to a normalised measure of shear strength step whilst the timestep was held constant. Simulations were conducted for various multipliers of τ_e (τ_{a_lo}, τ_{a_hi}), L , D , k_s , β_e (all results plotted here) to confirm that the identified error relationship was not dependent on any those five factors.

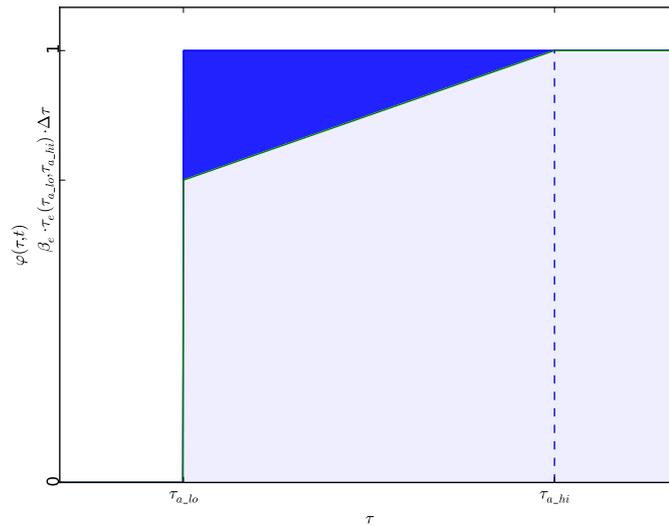


Figure 5.20: The dark blue area in this illustrative diagram is the relative amount of material eroded over dt following a step increase in shear stress from τ_{a_lo} to τ_{a_hi} at $t = 0$.

The material for which $\tau = \tau_{a_hi}$ experiences negligible erosion, whereas the most erosion over dt will be for $\tau = \tau_{a_lo}$. From fig. 5.20 and eq. 4.3, it can be seen that the rate of change in the amount of wall-bound material (the area of the dark blue triangle) is:

$$\int_0^{\tau_{a_hi}} \frac{\partial \varphi(\tau, t)}{\partial t} d\tau = \frac{\beta_e \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})^2}{2} \quad (5.25)$$

The maximum material release rate per m^2 of pipe is therefore:

$$\begin{aligned} \frac{dN(t)}{dt} &= \frac{\alpha \cdot \beta_e \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})^2}{2} \\ \left[\frac{TPMU}{m^2 \cdot s} \right] &= \left[\frac{TPMU}{Pa \cdot m^2} \right] \left[\frac{1}{Pa \cdot s} \right] \left[\frac{Pa}{1} \right]^2 \end{aligned} \quad (5.26)$$

The error from numerically evaluating eq. 5.4 can be conceptually non-dimensionalised by normalising it by eq. 5.26, as shown in eq. 5.27. Here, Δt is used to indicate that the numerator is evaluated numerically.

$$\frac{2 \cdot \max_{t \in [t_0, t_{max}]} \left(\frac{dN(t)}{\Delta t} \right)}{\alpha \cdot \beta_e \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})^2} \left[\frac{TPMU \cdot m^{-2} \cdot s^{-1}}{TPMU \cdot m^{-2} \cdot s^{-1}} \right] = [-] \quad (5.27)$$

Results Varying Δt was found to have no impact on the numerical evaluation of eq. 5.4 (fig. 5.18b) for smaller values of $\Delta t/t_d$. For $\Delta t/t_d > 1.0$, the discretisation error increased linearly.

Varying $\Delta \tau$ had no impact on the numerical evaluation of eq. 5.4 (fig. 5.19b).

5.3.9.4 Error in the total amount of eroded material as observed at the downstream end of the pipe

The error from the numerical evaluation of eq. 5.3, which is a function of both the wall state and material transport models, can be non-dimensionalised in the same way as eq. 5.2, only here the wall area does not cancel between numerator and denominator:

$$\frac{\int_0^{t_{max}} Q(t) \cdot T_{ds}(t) \Delta t}{A_w \cdot \alpha \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})} \quad (5.28)$$

Again, this expression is only valid if t_{max} is large enough so that all shear strengths affected by the increase in τ_a are fully depleted before $t = t_{max}$.

The error in eq. 5.28 was quantified for different values of Δt normalised by t_t i.e.

$$\{\Delta t / (t_t/i) \mid i \in \{16, 15, \dots, 2, 1\}\}$$

rather than by t_d as the intention was to quantify one measure of error in the material transport model, of which t_t (but not t_d) is a characteristic time.

The discretisation error in eq. 5.28 is a function of both the wall state and material transport sub-models. The portion of this error due that stems solely from discretisation within the transport model can be quantified by normalising the total amount of material detected at the end of the pipe during the simulation by the total amount of material released from the pipe wall (see eq. 5.29).

$$\frac{\int_0^{t_{max}} Q(t) \cdot T_{ds}(t) \Delta t}{A_w \cdot \int_0^{t_{max}} \frac{dN(t)}{dt} \Delta t} \quad (5.29)$$

Results

The value of Δt impacts considerably on the numerical error in eq. 5.3. Performing linear regression on all lines in fig. 5.21a results in a mean gradient $m = -0.488$ [–], a mean y-axis intercept $c = 0.996$ [–] and a mean R^2 of 1.000. The relative error in eq. 5.3 for a given Δt is therefore:

$$0.996 - 0.488 \cdot \frac{\Delta t}{t_t} \quad (5.30)$$

Note that when evaluating the numerator in eq. 5.3 that the long ‘tail’ of the dN/dt profile was not truncated as this resulted in a y-axis intercept that was closer to the true value of $c = 1$ ($c = 0.984$ if the tail of dN/dt is truncated). Also, note that model sensitivity was explored for $\Delta t < t_t$ to ensure that water parcels did not pass through the pipe undetected.

The attempt to non-dimensionalise the error in eq. 5.28 with respect to Δt was not entirely successful: the lines in fig. 5.21a corresponding to different values of τ_e , L , D , k_s , β_e and α almost but not quite collapsed onto a single line. This discrepancy can be seen to be due to the wall state model as if one plots eq. 5.29 against $\Delta t/t_t$ all lines collapse onto one (fig. 5.22).

The similarity of figs. 5.21a and 5.22 indicates that the error in eq. 5.28 is almost all due to the discrete mass transport model solution. Here the long tail of dN/dt was not truncated when calculating both the numerator and denominator of eq. 5.29.

The error in eq. 5.29 was found to be variable but low (< 1.9) for $\Delta\tau/\tau_e < 0.05$ and a constant 0.2% for greater values of $\Delta\tau/\tau_e$ (fig. 5.23a).

5.3.9.5 Error in the maximum turbidity experienced at the downstream end of the pipe

The maximum turbidity following a step increase in τ_a at t_0 occurs at the turnover time t_t (§5.2.2), when material mobilised from along the entire length of the pipe wall at different times all converges at the downstream end of the pipe at the same time. The concentration at the downstream end of the pipe at t_t can be thought of being a function of the concentration of each of a number of water parcels where the mobilised material in each has been in transit for different durations. The concentration of each parcel is the material release rate at a time t (dN/dt , where $t \in [t_0, t_t]$), factored by Δt and the internal surface area to volume ratio, $4/D$. Integrating this expression over $t \in [t_0, t_t]$ gives the total concentration at the downstream end of the pipe due to the concurrent arrival of all of these parcels. The error that results

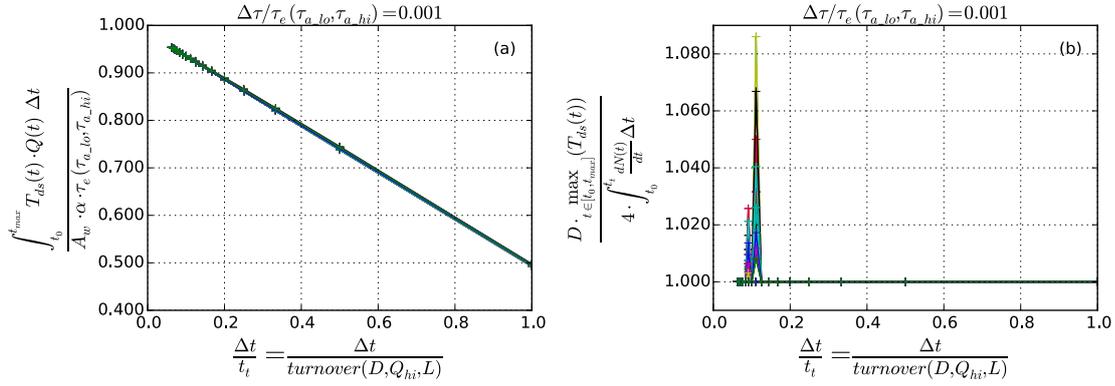


Figure 5.21: For a simple simulation where the shear stress increases from τ_{a_lo} to τ_{a_hi} at $t = 0$ (a) shows the non-dimensionalised (relative) error from discretely calculating the total amount of material to have been mobilised and passed out of the pipe during the simulation (with caveats in the body of the text) whereas (b) shows the relative error from discretely calculating the maximum turbidity predicted at the downstream end of the pipe during the simulation. The sensitivity of these two relative metrics was explored with regards to a normalised measure of timestep whilst the shear strength step was held constant. Simulations were conducted for various multipliers of τ_e (τ_{a_lo}, τ_{a_hi}), L , D , k_s , β_e (all results plotted here) as a check for the identified error relationship being independent of those five factors; this was found not to be the case for (b).

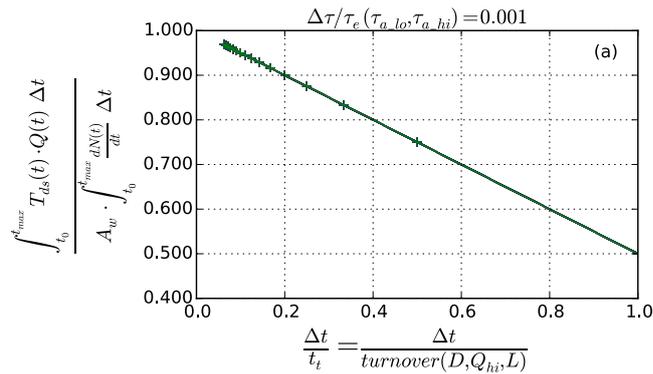


Figure 5.22: For a simple simulation where the shear stress increases from τ_{a_lo} to τ_{a_hi} at $t = 0$, this figure shows the non-dimensionalised (relative) error from discretely calculating the total amount of material to have been mobilised and passed out of the pipe during the simulation (with caveats in the body of the text). It is conceptually similar to fig. 5.21a but here the error metric is calculated in a different manner (in a way that ensures the error is solely a function of the transport sub-model rather than from the wall state sub-model too).

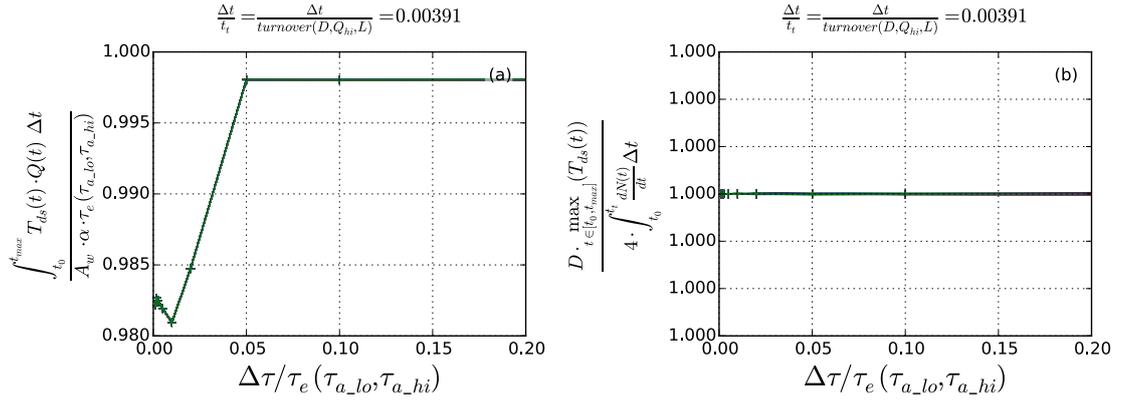


Figure 5.23: For a simple simulation where the shear stress increases from τ_{a_lo} to τ_{a_hi} at $t = 0$ (a) shows the non-dimensionalised (relative) error from discretely calculating the total amount of material to have been mobilised and passed out of the pipe during the simulation (with caveats in the body of the text) whereas (b) shows the relative error from discretely calculating the maximum turbidity predicted at the downstream end of the pipe during the simulation. The sensitivity of these two relative metrics was explored with regards to a normalised measure of the shear strength step whilst the timestep was held constant. Simulations were conducted for various multipliers of τ_e , L , D , k_s , β_e (all results plotted here) to confirm that the identified error relationship was not dependent on any those five factors.

from numerically evaluating the expression for peak turbidity in eq. 5.5 (a function of both the material transport and wall state sub-models) can be non-dimensionalised by dividing by this integral (a function of just the wall state sub-model) to give eq. 5.31.

$$\frac{\max_{t \in [t_0, t_{max}]} T_{ds}(t)}{\frac{4}{D} \cdot \int_{t_0}^t \frac{dN(t)}{dt} \Delta t}, \left[\frac{TPMU \cdot m^{-3}}{m^{-1} \cdot TPMU \cdot m^{-2} \cdot s^{-1} \cdot s} \right] = [-] \quad (5.31)$$

Note that in this case the denominator is not an analytical expression; it requires the numerical evaluation of the wall state sub-model so includes the error from that sub-model. However, this error should cancel between numerator and denominator, resulting in an expression for the error that arises solely from the numeric evaluation of the material release model.

Results

Fig. 5.21b indicates that efforts to non-dimensionalise the error with regards to Δt were not entirely successful: the lines corresponding to simulations run with different values of τ_e , L , D , k_s , β_e and α did not collapse to a single line in this figure (for $\Delta t/t_t < 0.1$). One explanation for this is that the attempt to normalise eq. 5.31 so that it was solely a function of the transport sub-model were not entirely successful and so it was therefore inappropriate to plot eq. 5.31 against the timestep normalised by t_t as t_t is solely a characteristic time of the transport sub-model. If the error is also a function of the wall state sub-model then eq. 5.31 should ideally be plotted against a composite of t_t and the wall state characteristic time t_d . However, it is not known if it is possible to combine these terms.

Eq. 5.31 was found to be invariant with regards to $\Delta\tau/\tau_e$ (fig. 5.23b).

Note that the long tail of dN/dt was again clipped when evaluating the numerator and denominator of eq. 5.31 and exploring the sensitivity with regards to Δt and $\Delta\tau$.

5.3.9.6 Repeating the numerical sensitivity analysis using refined mesh resolution ranges and constants

The numerical sensitivity analysis presented so far in §5.3.9 indicates that the sensitivity of certain model output metrics to $\Delta\tau$ may have been undertaken for sub-optimal values of Δt and vice versa. An iterative approach is therefore required to find regions of the mesh resolution space that are stable in both dimensions i.e. with regards to both Δt and $\Delta\tau$.

When exploring the sensitivity of model output metrics to Δt , the value of $\Delta\tau$ was previously fixed at $\Delta\tau \geq \tau_e/1000$ yet figs. 5.19 and 5.23 indicate that the model output to be more stable for $\Delta\tau \geq \tau_e/20$. Also, the wall state sub-model outputs are potentially more stable when $\Delta t \leq t_d(\tau_{a_lo}, \tau_{a_hi})$ (fig. 5.18) and the material transport metrics are more stable for $\Delta t \geq t_t/5$ (fig. 5.21b). However, the significance of these lower bounds for $\Delta t/t_t$ and $\Delta\tau/\tau_e$ should not be overstated for two reasons: firstly, the results suggesting that Δt should be $> t_t/5$ were derived from partially unsuccessful non-dimensionalisation (fig. 5.21b) and, secondly, reducing $\Delta\tau$ below $\tau_e/20$ only has a small (%2) effect on results (figs. 5.19a and 5.23a).

Repeating the sensitivity analysis for non-dimensional measures of Δt that satisfy these constraints and for $\Delta\tau$ fixed at $\tau_e/20$ confirms that the outputs from the wall state sub-model and material transport sub-model (figs. 5.24 and 5.25 respectively) are less sensitive to Δt for this $\Delta\tau$ value.

The wall state sub-model was also found to be less sensitive to $\Delta\tau$ after refining the mesh resolution space, as can be seen in fig. 5.26. Note that here Δt was normalised by t_t (rather than t_d as had been done previously in fig. 5.19).

Refining the mesh resolution space appeared to reduce the maximum error in the total amount of material to be mobilised as quantified using the transport model output and various $\Delta\tau/\tau_e$ (fig. 5.27a). However, all lines plotted in fig. 5.27a do not collapse down to a single line. One possible reason for this is that shear strength has not been non-dimensionalised in an appropriate way, yet fig. 5.27b shows that the peak turbidity *is* invariant with this form of non-dimensionalised shear strength for the refined mesh resolution space.

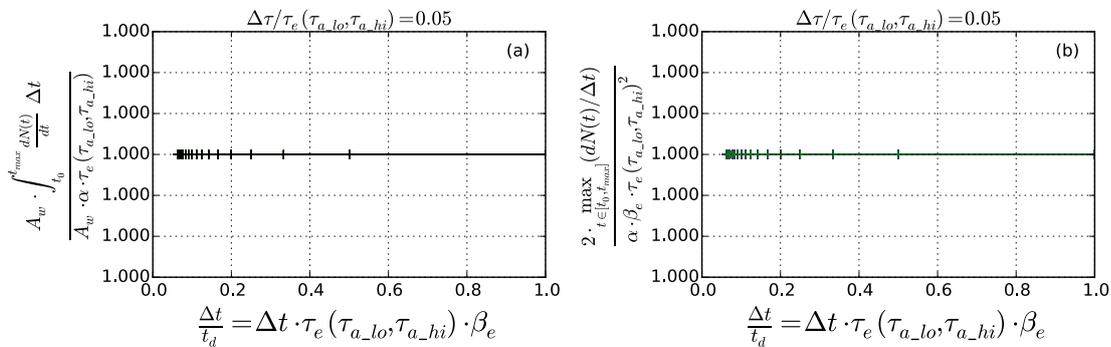


Figure 5.24: The sensitivity of the wall state sub-model to a non-dimensionalised measure of Δt given a normalised value of $\Delta\tau$. This sensitivity analysis is identical to that presented in fig. 5.18 but here output metrics have been calculated for a narrower $\Delta t/t_d$ range (≤ 1) and $\Delta\tau/\tau_e$ was fixed at a higher value (1/20 rather than 1/1000).

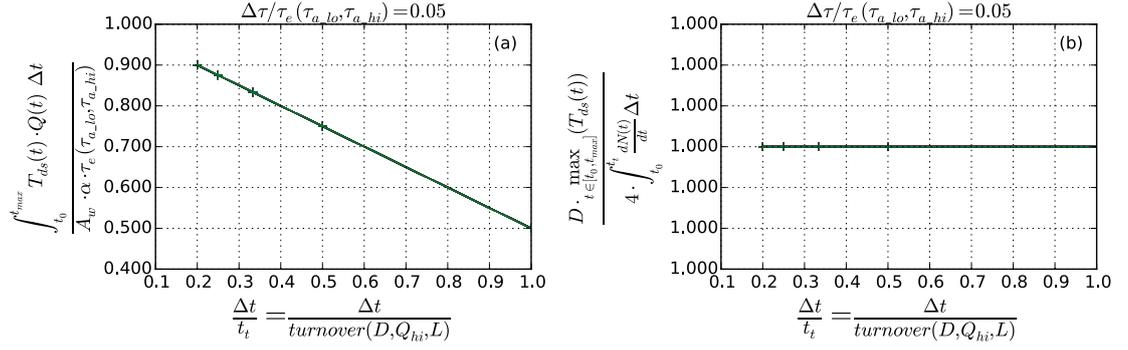


Figure 5.25: The sensitivity of the material transport sub-model to a non-dimensionalised measure of Δt given a normalised value of $\Delta\tau$. This sensitivity analysis is identical to that presented in fig. 5.21 but here output metrics have been calculated for a narrower $\Delta t/t_t$ range ($0.2 \leq \Delta t/t_t \leq 1$) and $\Delta\tau/\tau_e$ was fixed at a higher value ($1/20$ rather than $1/1000$).

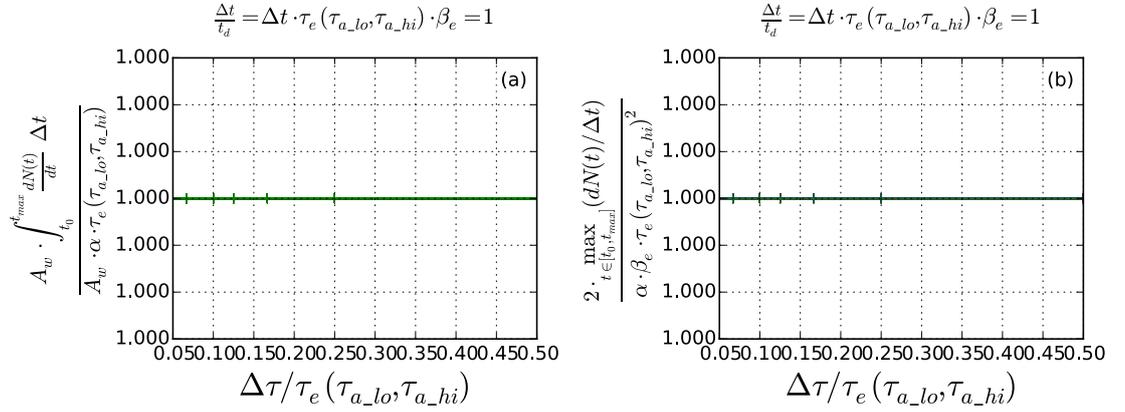


Figure 5.26: The sensitivity of the wall state sub-model to a non-dimensionalised measure of $\Delta\tau$ given a normalised value of Δt . This sensitivity analysis is identical to that presented in fig. 5.19 but here output metrics have been calculated for a narrower $\Delta\tau/\tau_e$ range ($1/20 \leq \Delta\tau/\tau_e \leq 1/2$) and Δt was constrained by $\Delta t/t_d = 1$ rather than $\Delta t/t_t = 1/256$.

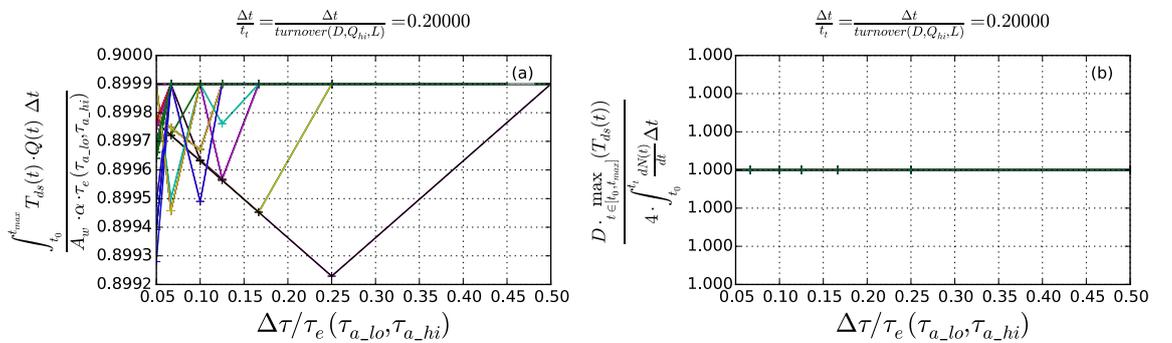


Figure 5.27: The sensitivity of the material transport sub-model to a non-dimensionalised measure of $\Delta\tau$ given a normalised value of Δt . This sensitivity analysis is identical to that presented in fig. 5.23 but here output metrics have been calculated for a narrower $\Delta\tau/\tau_e$ range ($1/20 \leq \Delta\tau/\tau_e \leq 1/2$) and Δt was constrained by $\Delta t/t_t = 1/5$ rather than $\Delta t/t_t = 1/256$.

5.3.9.7 Discussion and Summary

The results from investigating the sensitivity of the proposed model to the mesh resolution (given the resolution ranges/values from §5.3.9.6) are summarised in table 5.2.

From this analysis, it can be concluded that the mesh resolution for a simulation should satisfy the following constraints (for $\tau_e = \tau_{a_hi} - \tau_{a_lo}$ and $t_d = ((\tau_{a_hi} - \tau_{a_lo}) \cdot \beta_e)^{-1}$):

$$\Delta t \geq t_d/16 \quad (5.32)$$

$$\leq t_d \quad (5.33)$$

$$< t_t \quad (5.34)$$

$$\Delta \tau \leq \tau_e/2 \quad (5.35)$$

Numerical results indicate that the model is less sensitive to mesh resolution if the following constraints are also satisfied, although these are of lesser significance for the reasons given in §5.3.9.6. In particular, although table 5.2 indicates that the error associated with Δt is more stable for $\Delta t \geq t_t/5$, this should not be used as a lower bound for Δt as this relationship is not non-dimensional (see fig. 5.21) and so is not transferable between model configurations.

$$\Delta t \geq t_t/5 \quad (5.36)$$

$$\Delta \tau \geq \tau_e/20 \quad (5.37)$$

Importantly, mesh resolution constraints have been found that minimize the discretisation error in the peak turbidity prediction. It is this output metric that is of greatest importance to water users and therefore to network operatives when assessing the effect of a real or hypothetical flow increase.

In practise, these constraints need to be used with a safety margin when modelling real DWDS pipes to allow for imprecise and/or inaccurate model inputs (such as systematic flow measurement error or inaccurate roughness estimates (if used for calculating shear stress)). It is suggested that lower and upper bounds be increased and decreased by 10% respectively.

When evaluating these constraints, there is a need to scale $t_d(\tau_{a_lo}, \tau_{a_hi})$ by $0 < \gamma \leq 1$ as per §5.3.4 when $\varphi(\tau_{a_lo}, t_0) < 1$. Secondly, it is preferable to select a $\Delta \tau$ value from the lower end of the range enclosed by the constraints to minimise the dependence of $\Delta \tau$ on Δt (see §5.3.7).

Several factors impede the use of these constraints for identifying a suitable mesh resolution prior to solving a particular model instance. Firstly, due to a dependence on t_d , several constraints depend on β_e , which is a quantity that cannot be easily be determined precisely before model calibration. However, prior model calibrations of similar mains may give an imprecise understanding of β_e). Alternatively, β_e could be found in advance using an iterative method (find the β_e that gives the best fit to observations given an assumed Δt then find the best Δt given the found β_e until the process converges on acceptably good values of Δt

Table 5.2: Summary of the ways in which the mesh resolution impacts on the model outputs. Δt has been substituted for dt in the expressions referenced in the first column to indicate that it is the numerical evaluation of the expressions that is of interest. Here, $\tau_e = \tau_{a_hi} - \tau_{a_lo}$ so $t_d = ((\tau_{a_hi} - \tau_{a_lo}) \cdot \beta_e)^{-1}$. Also, \max is over the range $t \in [0, t_{max}]$.

Output metric	Error metric	Studied mesh resolution range	Comments re error
$A_w \cdot \int_0^{t_{max}} \frac{dN(t)}{dt} \Delta t$ (eq. 5.2)	$\frac{\int_0^{t_{max}} \frac{dN(t)}{dt} \Delta t}{\alpha \cdot \tau_e}$ (eq. 5.24)	$1/16 \leq \Delta t/t_d \leq 1$ $\Delta\tau/\tau_e = 1/20$	Invariant
		$1/20 \leq \Delta\tau/\tau_e \leq 1/2$ $\Delta t/t_d = 1$	Invariant
$\max(dN(t)/\Delta t)$ (eq. 5.4)	$\frac{2 \cdot \max(\frac{dN(t)}{\Delta t})}{\alpha \cdot \beta_e \cdot \tau_e^2}$ (eq. 5.27)	$1/16 \leq \Delta t/t_d \leq 1$ $\Delta\tau/\tau_e = 1/20$	Invariant
		$1/20 \leq \Delta\tau/\tau_e \leq 1/2$ $\Delta t/t_d = 1$	Invariant
$\int_0^{t_{max}} Q(t) \cdot T_{ds}(t) \Delta t$ (eq. 5.3)	$\frac{\int_0^{t_{max}} Q(t) \cdot T_{ds}(t) \Delta t}{A_w \cdot \alpha \cdot \tau_e (\tau_{a_lo}, \tau_{a_hi})}$ (eq. 5.28)	$1/5 \leq \Delta t/t_t \leq 1$ $\Delta\tau/\tau_e = 1/20$	Inversely proportional to $\Delta t/t_t$
		$1/20 \leq \Delta\tau/\tau_e \leq 1/2$ $\Delta t/t_t = 1/5$	Invalid non-dimensionalisation.
$\max(T_{ds}(t))$ (eq. 5.5)	$\frac{\max T_{ds}(t)}{\frac{4}{D} \cdot \int_0^{t_t} \frac{dN(t)}{dt} \Delta t}$ (eq. 5.31)	$1/5 \leq \Delta t/t_t \leq 1$ $\Delta\tau/\tau_e = 1/20$	Invariant
		$1/20 \leq \Delta\tau/\tau_e \leq 1/2$ $\Delta t/t_t = 1/5$	Invariant

and β_e). Secondly, several constraints are dependent on τ_e , which is undesirable as τ_e is a time-dependent quantity. For certain constraints, τ_e could be approximated by $\max(\tau_a)$, the maximum shear stress that the pipe is anticipated to experience, but this is only appropriate for situations where substituting $\max(\tau_a)$ for τ_e results in a more restrictive constraint. Given the difficulty of estimating β_e and τ_e prior to deciding upon a mesh resolution, it is instead recommended that only eq. 5.35 be used in advance of deciding upon a mesh resolution and the other constraints be used to *appraise* a mesh resolution for an existing, calibrated model.

This sensitivity analysis has focussed on the sensitivity of model outputs to mesh resolution for *just erosion events*. Similar analysis could have been conducted for regeneration in which expressions of model output were non-dimensionalised using $\Delta t/\beta_r$, the characteristic time required for complete regeneration in one timestep. However, such analysis would be redundant given that complete erosion occurs over several turnovers (§2.1.12.2; §2.1.13.4) yet regeneration takes years (§2.1.14) and so the model is far more sensitive to the characteristic time for material depletion, t_d , than that characteristic time of regeneration.

It should be noted that this sensitivity analysis is subject to several limitations. Firstly, the error in the total amount of material to pass out of the end of pipe increases linearly with Δt (fig. 5.25a). Secondly, attempts to identify a non-dimensional relationship between this model output metric and $\Delta\tau$ were unsuccessful (fig. 5.27a). These issues are thought to be due to an artefact of the transport model's implementation (the precise nature of which is unknown).

5.4 Sensitivity of the model to its parameters

In this section, the influence of the model parameters (β_e , β_r and α) on the model outputs is considered.

When considering what factors most significantly influenced the mesh resolution (§5.3) there was little need to consider material regeneration given that material erosion is a much faster and therefore more sensitive process. The sensitivity to the mesh resolution was therefore investigated for just a short-term erosion event and regeneration was assumed to be negligible. Such a basic model configuration is of course insufficient for studying the sensitivity of the model to the regeneration rate β_r ; in this section the model configuration is therefore as follows:

- No material was bound to the pipe wall at the start of the simulation i.e. $\varphi(\tau, t_0) = 0$;
- Following t_0 , there were three months of constant low flow (with shear stress $\tau_a = \tau_{a_lo}$), allowing discolouration material to accumulate on the pipe wall;
- The shear stress then instantaneously increased from τ_{a_lo} to τ_{a_hi} , causing material erosion, as a result of a trebling of pipe flow;
- The simulation ended one hour after the flow increase.

A series of twenty simulations was conducted for each of the three parameters. Within each of these three sets of simulations the parameter of interest was linearly varied over a specified range whilst the other two parameters were held constant. The bounds of each parameter range was determined through speculative exploration of the problem space: it was not possible to make more justifiable choices given the empirical nature of the model and parameters. The constant parameter values were the means of the ranges used in the sensitivity analysis.

5.4.1 Sensitivity to α and β_r

The α parameter simply scales both the material release rate, dN/dt , and turbidity prediction, T_{ds} , linearly over the studied parameter range (figs. 5.28a and 5.28b).

The β_r parameter has what initially appears to be a similar effect on the turbidity prediction (fig. 5.29b) but is actually quite distinct: it linearly scales the peak turbidity (but not the entire turbidity profile) over a certain range then converges on an upper limit (fig. 5.29b). Unlike α , it also changes the *shape* of the turbidity response. This difference in effect can be explained by looking at how the parameters influence the rate of material release from the pipe wall:

- If β_r is small then partially-regenerated shear strengths are depleted in one timestep when τ_a increases to τ_{a_hi} (see fig. 5.29a). A slug of turbid water then passes down and out of the pipe without it notably increasing in turbidity, resulting in a downstream turbidity profile that is almost constant from soon after the τ_a increase up to one turnover time later (when the turbid slug begins to pass out of the pipe).
- If β_r is larger then the shear strengths affected by the τ_a increase take multiple timesteps to erode and so the maximum dN/dt is constant over this range (see fig. 5.29a). The

turbidity prediction is more arched and scales more linearly with β_r over this range of β_r values.

- If β_r is larger still then maximal regeneration is achieved before τ_a increases to τ_{a_hi} . The effect of this is that β_r ceases to scale the turbidity output, as can be more clearly seen in fig. 5.30, where the total amount of material that is mobilised (calculated using the integral of material flux as per eq. 2.14) is plotted against the values of β_r used in the sensitivity analysis.

The α and β_r parameters could in principle be distinguished using dN/dt . However, as dN/dt is non-trivial to measure in practise, understanding of the system behaviour must be entirely deduced from the turbidity response observed at a specific location along the pipe of interest: certain flow profiles, such as that in fig. 5.31 where particular a flushing profile is imposed at irregular intervals, could allow the parameters to be more easily distinguished using T_{ds} , whereas others do not (figs. 5.28b and 5.29b).

5.4.2 Sensitivity to β_e

The parameter β_e has a different effect on the turbidity response to α and β_r . Like β_r , it influences the scale and shape of both the dN/dt and T_{ds} responses (fig. 5.32) but in a more notable way. Note that larger (faster) values of β_e have little impact on the scale of the T_{ds} profile (fig. 5.32b).

Much smaller values of β_e have a very different effect on the T_{ds} profile (fig. 5.33): the rise in turbidity is linear rather than curved. This is due to the rate of erosion being sufficiently slow that a number of timesteps pass before any eroding shear strength is depleted, with the net erosion rate from the pipe wall therefore being constant for a period as the erosion rate per eroding shear strength is not time-dependent (see fig. 5.33a and the heatmaps in fig. 5.34).

5.4.3 Implications for fitting the model to data

The most significant model output for water consumers is the maximum turbidity at points of supply; all three parameters are therefore of importance in this regard as they all scale the turbidity prediction at the downstream end of the modelled pipe.

Ideally, one would want each parameter of an empirical model to have an easily distinguishable effect on the model output(s); however, if all three parameters of the proposed model influence the scale of the downstream turbidity prediction then this is not the case here. The proposed model must therefore be fitted to data using a method that is robust in the face of such indeterminacy and/or be fitted to datasets containing features that allow the effects of the parameters to be distinguished (e.g. fig. 5.31).

5.5 Sensitivity to the model boundary conditions

To assess the sensitivity of the model to the boundary conditions defined in §4.4.4, one *only* needs to consider the influence of the initial wall state (the relative amount of material per

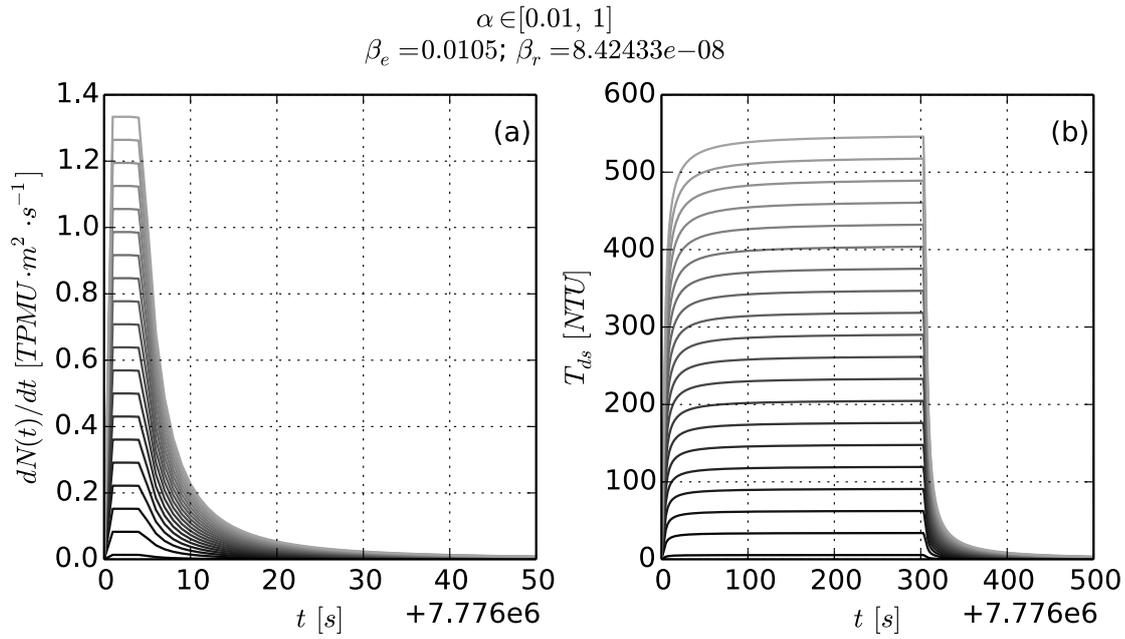


Figure 5.28: Exploring the sensitivity of the turbidity output to the α parameter. Larger values of the parameter under examination are depicted in lighter shades of grey. $\varnothing = 76 \text{ mm}$; $L = 400 \text{ m}$; $k_s = 5 \text{ mm}$; Q over first 90 days $= 2 L \cdot s^{-1}$ ($\Delta t = 900 \text{ s}$); Q over following 1 hour $= 6 L \cdot s^{-1}$ ($\Delta t = 900 \text{ s}$); $\tau \in [\tau_a(\varnothing, 2 L \cdot s^{-1}, k_s), \tau_a(\varnothing, 6 L \cdot s^{-1}, k_s)]$ (1000 τ bands); $\varphi(\tau, t = 0)$.

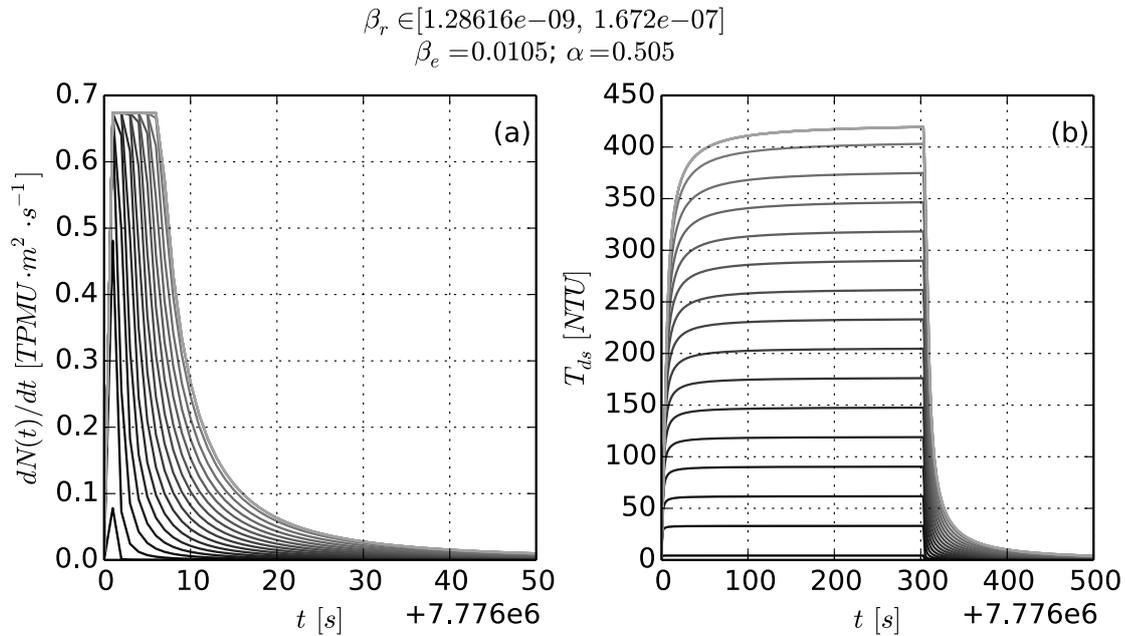


Figure 5.29: Exploring the sensitivity of the material release from the pipe wall and the downstream turbidity to the β_r parameter. Larger values of the parameter under examination are depicted in lighter shades of grey. Simulation configuration as per fig. 5.28.

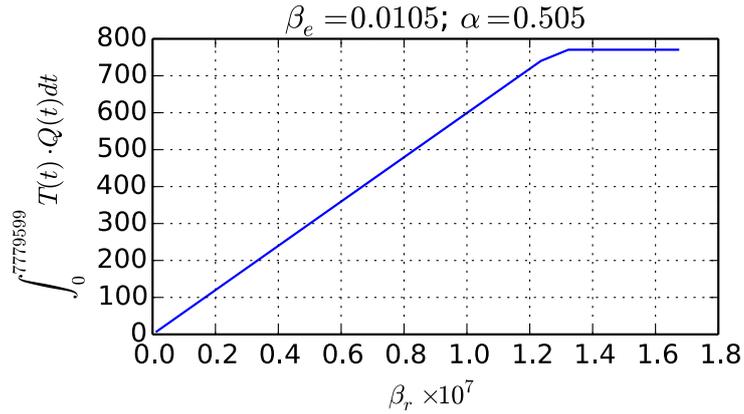


Figure 5.30: Exploring the sensitivity of the total amount of material mobilised (quantified using eq. 2.14) to regeneration rate (VCDM β_r parameter). Simulation configuration as per fig. 5.28. The amount of material mobilised from the wall during the simulation increases with the regeneration rate until the regeneration rate is $\approx 1.3 \times 10^{-7} \text{ s}^{-1}$; increasing the regeneration rate above this value has no impact on the total amount of material that is eroded.

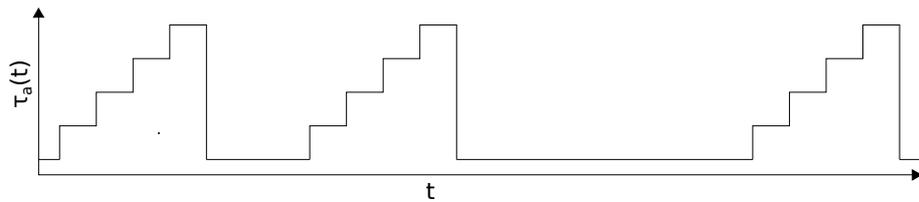


Figure 5.31: Particular τ_a profiles such as the one illustrated above can potentially allow the effects α and β_r to be distinguished more easily.

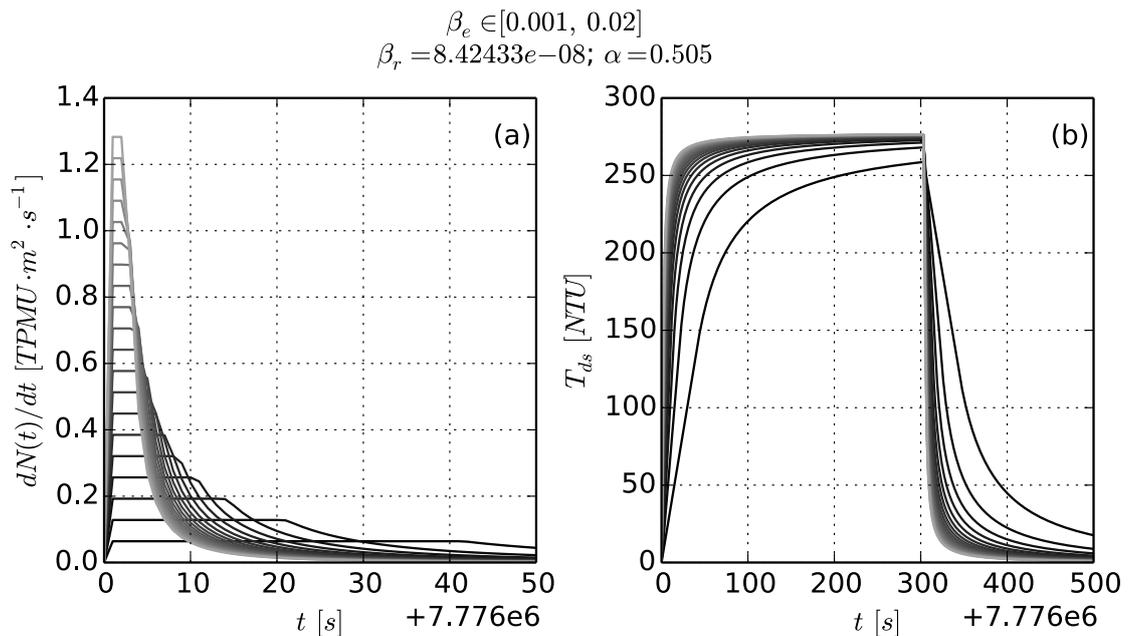


Figure 5.32: Exploring the sensitivity of the turbidity output to the β_e parameter. Larger values of the parameter under examination are depicted in lighter shades of grey. Simulation configuration as per fig. 5.28.

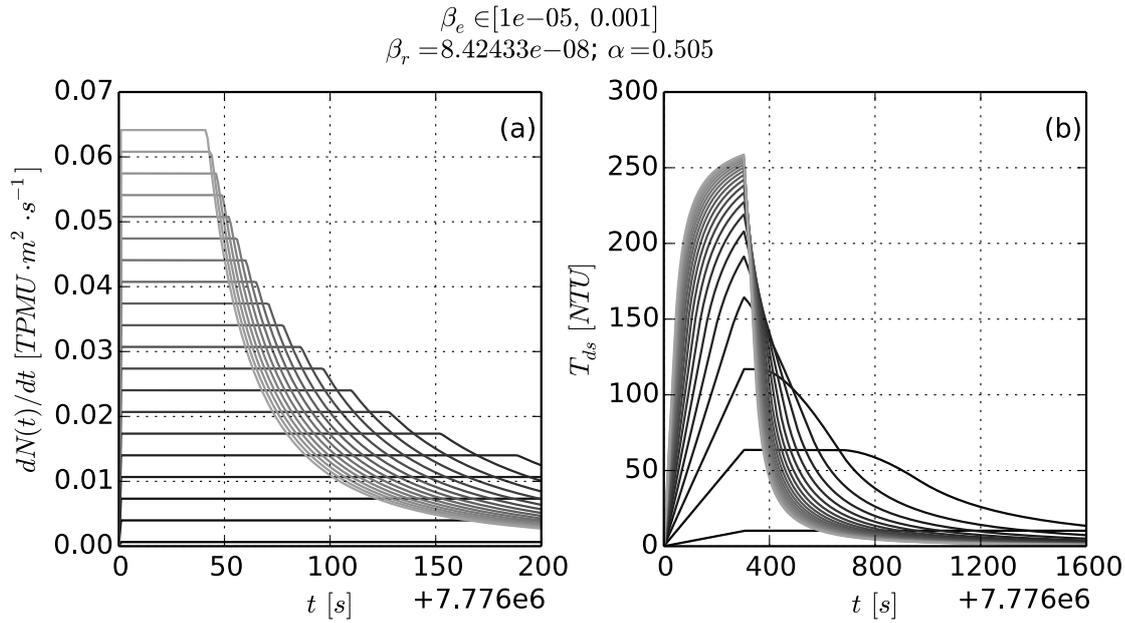


Figure 5.33: Exploring the sensitivity of the turbidity output to the β_e parameter using lower β_e than in fig. 5.32. Larger values of the parameter under examination are depicted in lighter shades of grey. Simulation configuration as per fig. 5.28. Note that the time axis scales differ to figs. 5.28 and 5.29 and 5.32).

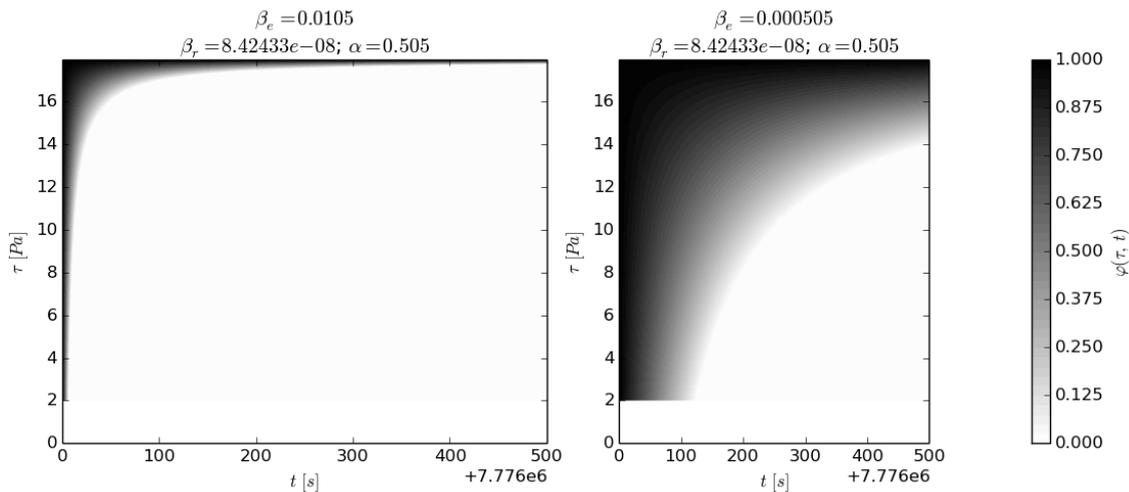


Figure 5.34: Exploring the sensitivity of the wall state, $\varphi(\tau, t)$, to faster values (left-hand subplot) and slower values (right-hand subplot) of β_e . The β_e values used in the left-hand and right-hand subplots were the means of the ranges explored in figs. 5.32 and 5.33 respectively. Simulation configuration as per fig. 5.28.

shear strength at t_0) as the initial bulk water state is ideally known and invariant between model calibrations.

A series of five simulations were conducted in which there was a step increase in shear stress from $1.0 Pa$ to $3.0 Pa$. Each simulation was initialised with a different boundary condition shape ($\varphi(\tau, t_0)$; see fig. 5.35a): these defined the relative amount of material present at the pipe wall solely for $\tau \in [1.0, 3.0] Pa$ and were ramped, stepped or constant with respect to τ . The amount of material eroded from the increase in shear stress was the same for all five simulations as all boundary condition shapes had the same area beneath them and the elevated shear stress level was sustained for long enough that all material with $\tau < 3.0 Pa$ was eroded. The shape of the wall state boundary condition influences both the shape and scale of the resulting material release rate and turbidity predictions (figs. 5.35b and 5.35c). As mentioned previously, β_e and β_r do too (see figs. 5.29, 5.32 and 5.33). To distinguish the effects of $\varphi(\tau, t_0)$, β_e and β_r when fitting the model to data one therefore needs to either:

- Have a good understanding of $\varphi(\tau, t_0)$ from knowing and studying historical data ($\tau_a(t < t_0)$) and/or
- Simulate over a long enough period for the effect of an estimate of the boundary condition $\varphi(\tau, t_0)$ to influence only the first part of the simulation. In the worst case, this duration would need to be at least β_r^{-1} , the time required for complete regeneration.

5.6 Model sensitivity to hydraulics and pipe attributes

5.6.1 Sensitivity to flow magnitude

The pipe flow profile, Q , influences both the release of material from the pipe wall and the advection of suspended material:

- The eroding force that drives the rate of material release from the pipe wall is the excess shear stress, which is a function of the absolute shear stress at a particular time (eq. 4.3). The absolute shear stress is a non-linear but monotonically increasing function of absolute flow (fig. 2.7).
- The advection of suspended material within a pipe is linearly related to flow if 'plug flow' is assumed (§2.2.7.3). The time of the peak turbidity following a step increase in shear stress is a function of flow as it occurs one turnover after that increase (§5.2.2).

Fig. 5.36 demonstrates how uncertainty in flow impacts on turbidity predictions. If the measured flow is greater than the true flow then this impacts on both the eroding force and advection i.e. (the effects do not negate each other):

- Material is predicted to be mobilised quicker as the excess shear stress is greater.
- The peak turbidity and total amount of material that is mobilised are both greater as a larger excess shear stress results in a greater range of shear strengths being eroded.

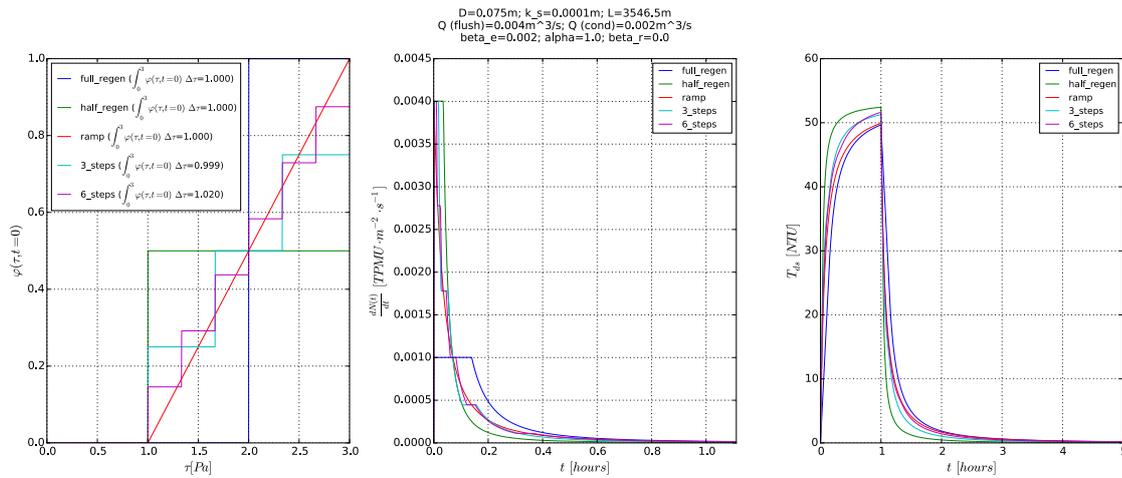


Figure 5.35: (a) shows a variety of wall state boundary condition ($\varphi(\tau, t_0)$) shapes, all of which bound the same area i.e. the amount of material at the pipe wall is the same for all (assuming a common value of α). (b) and (c) shows the simulated effect of a step increase in erosion that removes all material within the strength range shown in (a). The shape of the wall state boundary condition influences both the shape and scale of the predicted material release rate, dN/dt and turbidity response, T_{ds} .

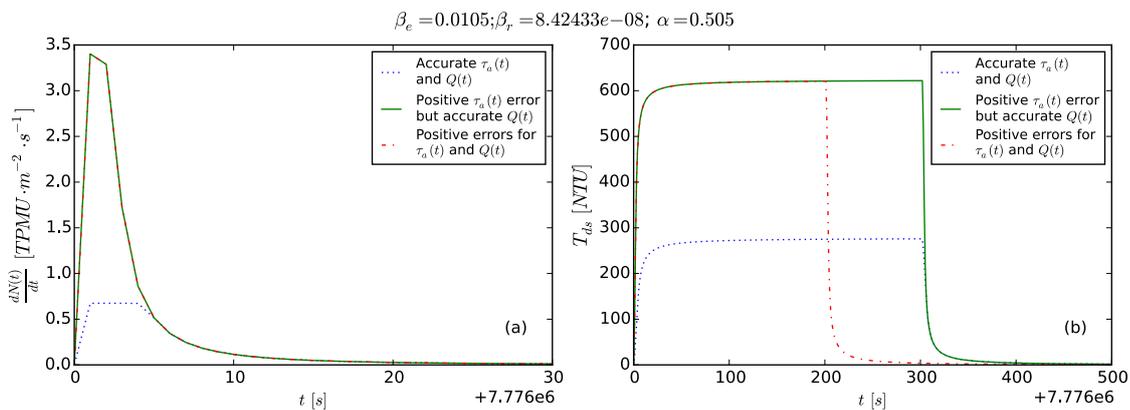


Figure 5.36: Exploring the sensitivity of (a) the material release rate and (b) downstream turbidity prediction to flow error. For the blue dotted turbidity profile the simulation configuration is as per fig. 5.28. For the green solid line turbidity profile the configuration was the same but the τ_a profile and the limits of the studied τ range were derived after multiplying the flow profile by 1.5; however, the original flow profile was supplied to the advection routine. For the red dotted and dashed turbidity profile the simulation configuration was as per the green solid line profile but the scaled-up flow profile was supplied to the advection routine. Note that a 50% error in flow is unrealistic but here clearly demonstrates the sensitivity of the model outputs to flow.

- The peak turbidity is predicted to occur sooner as the turnover time is shorter (fig. 5.36b).

In the related fig. 5.37b, one can see that:

- The magnitude of the turbidity response is non-linear with flow;
- The duration between the flow increase and the peak of the resulting increase in turbidity is non-linear with flow;

Flow, like the model parameters (§5.4), influences both the shape and scale of turbidity response. However, this is not thought give rise to a greater degree of indeterminacy and therefore be an impediment for fitting the model to data as both flow and shear stress can be measured/estimated and therefore accurately input into the model (unlike the empirical parameters).

5.6.2 Sensitivity to pipe length

The length, L , of the modelled pipe influences the turnover time for a given flow and therefore the time of the peak T_{ds} following a step increase in τ_a (see §5.2.2 for an explanation as to why) i.e. it impacts on how the dN/dt profile is transformed into T_{ds} (fig. 5.38). However, it does not influence the upward curve of T_{ds} to the peak nor the downward curve following the peak as it does not influence the rate of material release along the pipe length nor the rate of regeneration.

5.6.3 Sensitivity to the pipe internal diameter

The magnitude of the predicted turbidity profile, T_{ds} , scales in a negative, non-linear way with pipe diameter, D , (fig. 5.39) due to the relationships between D and τ_a (which contributes to the eroding force) and D and the bulk velocity \bar{u} for a given flow Q (material transport) also being non-linear. In addition, D influences the turnover time for a given \bar{u} so, like Q , also influences the time of the peak turbidity.

5.6.4 Sensitivity to the pipe roughness

The pipe roughness, k_s , also scales the T_{ds} profile (fig. 5.40) but in a positive way that is more linear² than for D .

5.6.5 Summary of sensitivity analysis and implications for model fitting

The sensitivity of the proposed model to the model parameters, boundary conditions, hydraulics and measurable pipe properties and are summarised in table 5.3. Mesh discretisation intervals Δt and $\Delta \tau$ have not been included as the sensitivity of model outputs to both was previously summarised in table. 5.2.

²The slight non-linearity in the referenced figure is possibly due to inaccuracies from calculating shear stress using the Swamee-Jain approximation to the Colebrook-White formula.

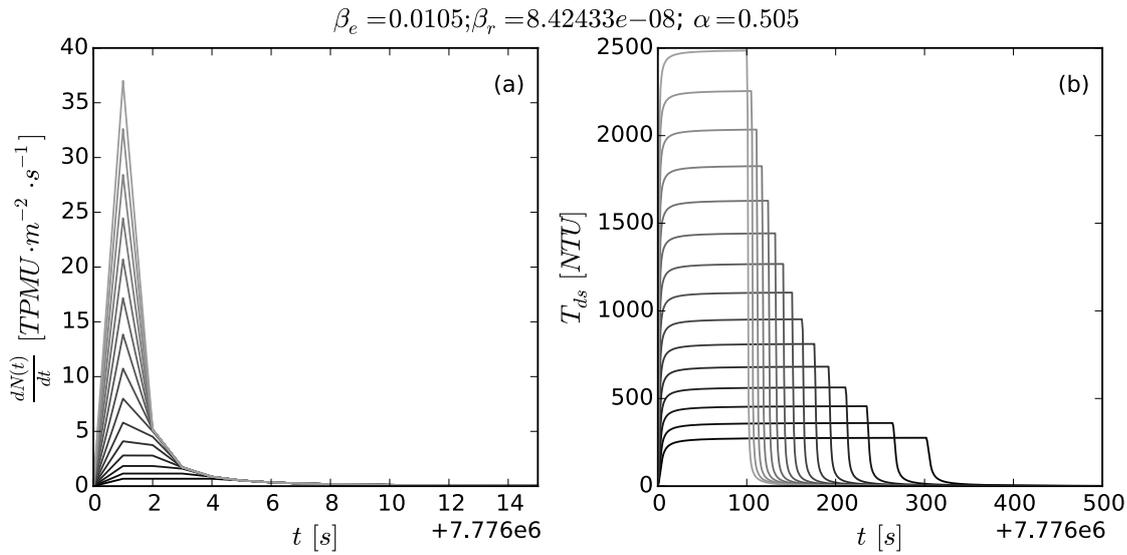


Figure 5.37: Exploring the sensitivity of (a) the material release rate and (b) downstream turbidity prediction given a range of different flow profiles. The simulation configuration is as per fig. 5.28 but the fifteen different simulations were conducted; these used the flow profile corresponding to that of fig. 5.28 multiplied by a factor between one and three (factors were linearly spaced over that range). The darkest and lightest lines shown correspond to flow flow factors of one and three respectively.

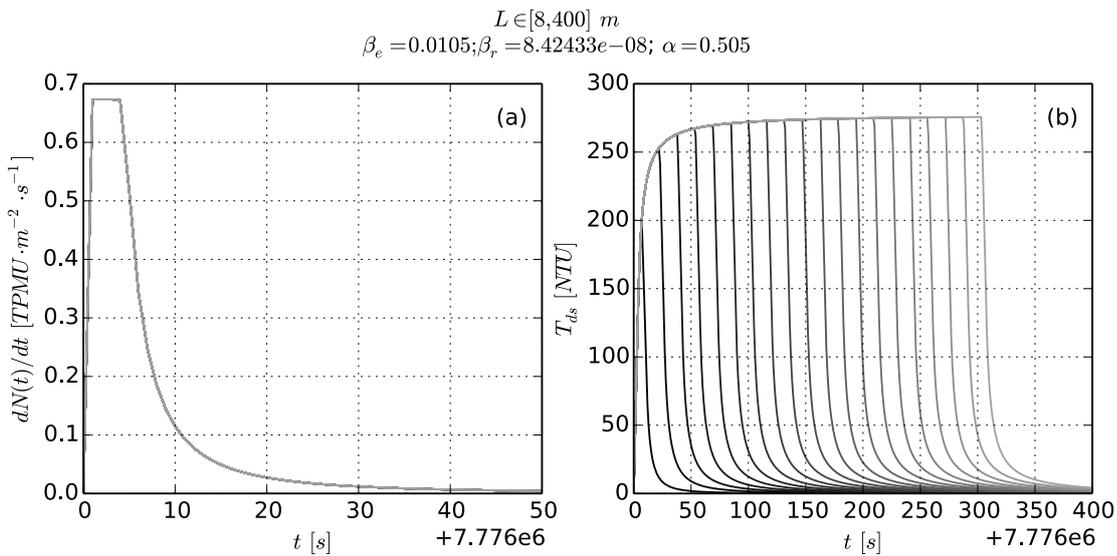


Figure 5.38: Exploring the sensitivity of the turbidity output to the length, L , of the modelled pipe. Simulation configuration is as per fig. 5.28. Simulations were conducted for each of twenty L values, with these values being linearly spaced over the range shown in the figure title. The results for larger values of L are shown in lighter shades of grey.

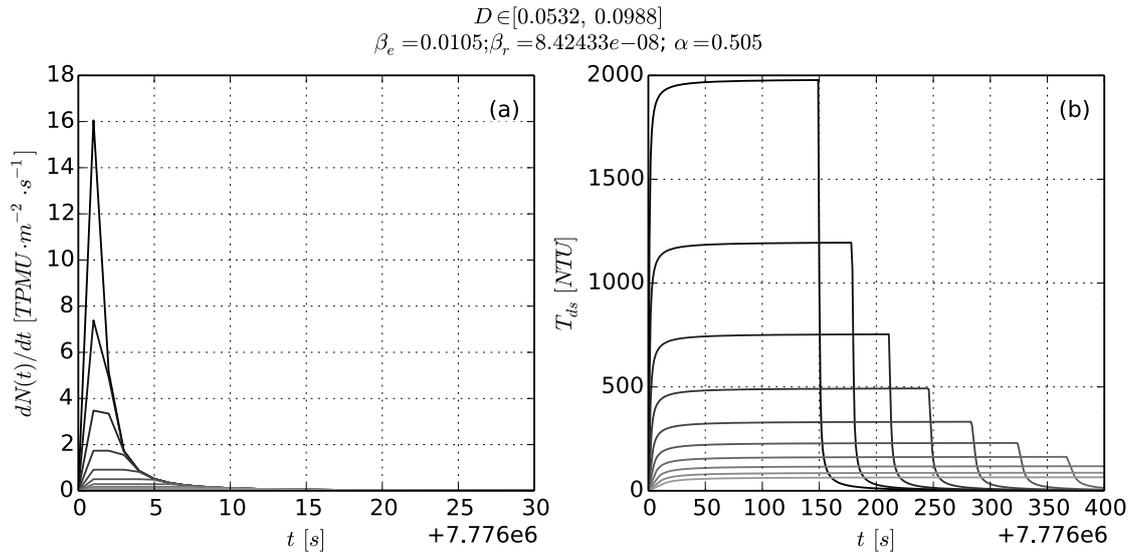


Figure 5.39: Exploring the sensitivity of (a) the material release from the pipe wall and (b) the turbidity output to the internal diameter D of the modelled pipe. Simulation configuration is as per fig. 5.28. Simulations were conducted for each of ten D values, with these values being linearly spaced over the range shown in the subplot titles. The results for larger values of D are shown in lighter shades of grey.

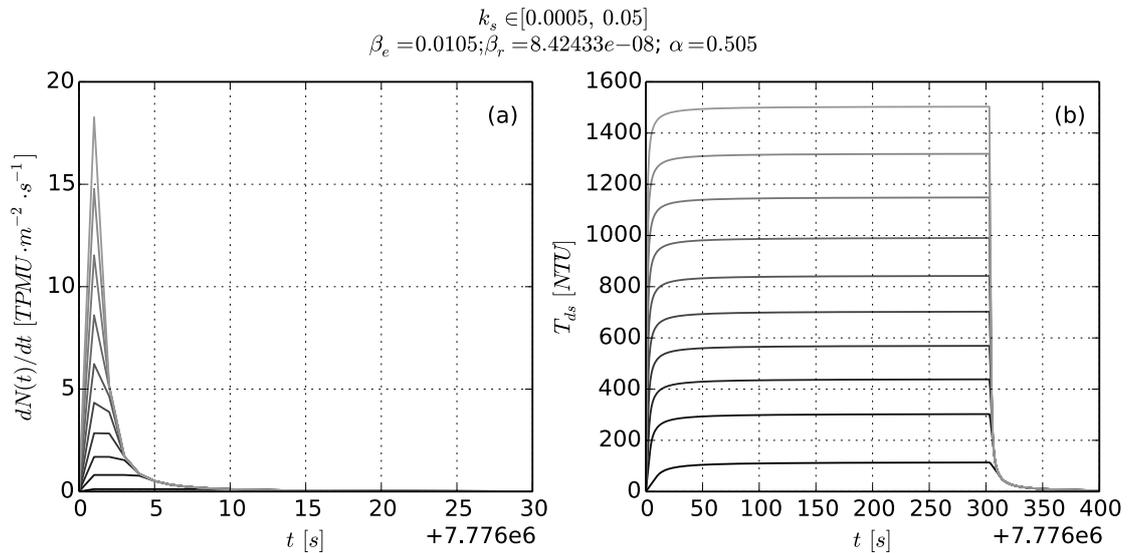


Figure 5.40: Exploring the sensitivity of (a) the material release from the pipe wall and (b) the turbidity output to the roughness k_s of the modelled pipe. Simulation configuration is as per fig. 5.28. Simulations were conducted for each of ten k_s values, with these values being linearly spaced over the range shown in the subplot titles. The results for larger values of k_s are shown in lighter shades of grey.

Table 5.3: A summary of how model parameters, boundary conditions, pipe hydraulics and measurable pipe attributes influence the model outputs. Here '+ve' and '-ve' are abbreviations for positive and negative respectively and 'lin' stands for linear.

Output Influence on	$dN(t)/dt$		$T_{ds}(t)$	
	Scaling	Shape	Scaling	Shape
α	+ve lin. w.r.t. dN/dt	No	+ve lin. w.r.t. T_{ds}	No
β_r	+ve non-lin. with dN/dt for small values only	Yes	+ve lin. scaling of peak and +ve non-lin. scaling of entire profile until β_r nears an upper bound;	Yes (for β_r values below a upper bound)
β_e	+ve lin. w.r.t. dN/dt ; -ve non-linear with t	Yes	+ve non-lin. with T_{ds}	Yes
$\varphi(\tau, t=0)$	Yes	Yes	Yes	Yes
Q	+ve non-lin with dN/dt	Yes	+ve non-lin. with T_{ds} ; -ve lin. with t	No
D	-ve non-lin. with dN/dt	Yes	-ve non-lin. with T_{ds} ; +ve non-lin. with t	No
k_s	+ve non-lin. with $dN(t)/dt$	Yes	+ve, almost-lin. with $T_{ds}(t)$	No
L	No	No	Yes, for short pipes only	The time of the peak (falls at the turnover time) but not the curvature to/from the peak

5.6.6 Assessing the need for an equivalent term to the **PODDS** excess shear power term

The parameters β_e and α are thought to offer similar functionality to the **PODDS** n parameter (§2.2.8.3) by being able to collectively alter both the shape and scale of the predicted turbidity output for a step increase in shear stress during a simulation short enough for the effects of regeneration to be negligible. See figs. 5.28, 5.32 and 5.33 which show the sensitivity of the **VCDM** to β_e and α and fig. 5.41 which shows the sensitivity of **PODDS** to n . This suggests that there is no need for an equivalent to the undesirable **PODDS** n term in the **VCDM**, which is beneficial for the reasons given in §2.2.8.5.

The initial wall state boundary conditions also influence the shape and scale of the turbidity response (§5.5) so also contribute to the **VCDM**'s ability to reproduce the effect of the **PODDS** n parameter without needing to include a similar power term in the **VCDM** formulations.

5.7 Model calibration: fitting the model parameters

5.7.1 Introduction

Without calibration (being fitted to real-world data), the model proposed in §4 cannot be used to simulate how bulk water turbidity and the properties of wall-bound discolouration material change over time within a given pipe. Every combination of pipe, source water and environmental conditions could result in different material accumulation and erosion characteristics (§2.1.12.2; §2.1.14; §2.1.15; §2.1.16) which need to be represented by different parameter values. Therefore, when constructing a model of a pipe (a model *instance*), the modeller needs to find parameter values and boundary conditions that ensure the instance state and outputs are sufficiently representative of reality over a period of interest. For the most accurate representation, the calibration process should be repeated for every distinct pipe length, although this requirement could be relaxed if pipes are deemed to have sufficiently similar properties (as is often done when calibrating **PODDS**).

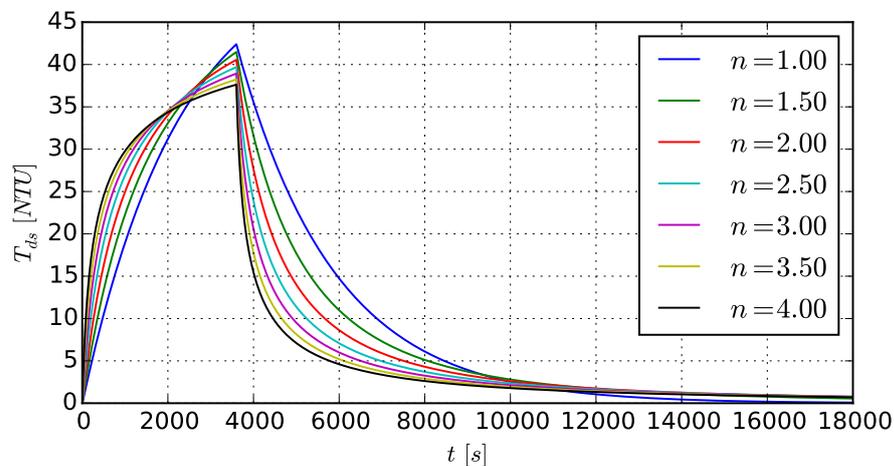


Figure 5.41: The **PODDS** turbidity response to an instantaneous step increase in shear stress τ_a (at t_0) for different values of the n parameter.

It is challenging if not impossible to obtain satisfactory parameter and boundary condition values for a given pipe by *direct* measurement of system properties: the parameters β_e , β_r and α are empirical constructs that have no physical or biochemical meaning, whereas the initial shear strength profile of material bound to the wall, $\varphi(\tau, t_0)$, cannot currently be directly measured in a timely, representative and non-destructive way (§2.1.18). Sharpe (2013) and Fish (2014) (see §2.1.9, §2.1.11, §2.1.13.4) were able to study biofilm structures in a largely non-destructive way in realistic-scale pipe systems but discolouration material is not solely comprised of biofilm material and their microscopy methods are too labour-intensive to easily be regularly repeated.

The *inversion* of mathematical models can provide another means for finding appropriate parameter values: the expressions that comprise the model can be reformulated to allow model parameters and/or boundary conditions to be calculated using observed outputs (and measurable system properties). However, this approach is also not applicable in this case as the material transport sub-model cannot be inverted: given time-series of flow and turbidity measurements, one cannot determine where and when the observed suspended discolouration material left the pipe wall upstream. One therefore cannot calculate when the state of the wall-bound material changed and so cannot calculate quantities relating to this state such as the wall state boundary conditions and the three main parameters. The inability to derive a generic, inverted form of the model from its original formulations is principally due to flow and shear stress being non-trivial functions of time. This makes it impossible to separate the turbidity signal observed at the downstream end of a pipe into the system responses that resulted from distinct increases in shear stress. Advection processes therefore partly obscure the observer's view of if, when and how the wall state changed as a result of variation in pipe flow.

A more viable means for calibrating the proposed model is to fit a model instance by minimising the dissimilarity between observed data and model predictions (§2.2.4). This involves:

1. Taking a consistent set of observations regarding the pipe(s) of interest;
2. Assigning values to model parameters and boundary conditions (either arbitrarily or using some prior knowledge of the studied system);
3. Making a model prediction corresponding to part or all of the observation period;
4. Quantifying the similarity of observations and predictions;
5. Then either:
 - (a) Accepting the model is sufficiently calibrated or
 - (b) Reiterating steps 3-5 after refining the parameter and boundary conditions.

Empirical models can be fitted to data 'by hand' (without automation) through trial and error where the *parameter space* (the potentially infinite set of all possible parameter values) is relatively small and/or well understood. For example, the PODDS parameter space has previously been explored using Monte Carlo simulations (Boxall and Saul, 2005) and also by fitting the model to many different pipe lengths (e.g. Husband and Boxall, 2010a). This resulted in knowledge of how parameters were somewhat transferable between pipe lengths and

of the effect of changing different parameters, allowing models to presently be calibrated by hand. At present, a modeller can often calibrate a [PODDS](#) model instance relatively quickly by *bootstrapping* the fitting process by starting with parameter values from previously successful calibrations of models of similar systems (§2.2.8.4).

A less manual approach is believed to be required for calibrating the model proposed in §4. The rationale for this is as follows:

- The parameter space is (presently) far less well known than that of [PODDS](#), although some understanding of likely regeneration rates can be gained from studying data from repeated mains flushing (§2.1.13.2, §2.1.14) and the shear strength profile at the start of a simulation can be estimated using the prior shear stress history (subsequently explored in §5.8). Calibrating the proposed model for a given system could therefore involve a fairly exhaustive search of the parameter space.
- An exhaustive search could prove costly given the dimensionality of the problem space (the number of modeller-tunable values: β_e , β_r , α and the values of the boundary condition array) and that each dimension is a continuous range rather than a set of discrete values.
- Multiple parameters have similar effects on the shape and scale of model outputs (as can be seen in table 5.3), suggesting that the fitting of the proposed model is an indeterminate problem i.e. a potentially infinite number of parameter sets can offer a similar model fit (as is also the case for [PODDS](#); see §2.2.8.4). The instinctive approach to hand calibration of adjusting one parameter at a time may therefore not be appropriate; a holistic approach to refining parameter values is required to lessen the chance of finding a fit that is locally but not globally optimal.

There are other reasons for the model being non-trivial to fit to data. Firstly, the proposed model and [PODDS](#) can only be fitted by comparing turbidity predictions from model instances to turbidity observations. The state of discolouration material bound to the internal surfaces of buried pipes and other assets cannot easily be directly measured (§2.1.18). There is a need to derive much information about the system of interest from a signal (bulk water turbidity) whose information content is limited by the spatial and temporal sampling resolution (observation via a limited number of sampling points is depicted in fig. 4.19). For example, the hydrants available for sampling bulk water turbidity along the length of a main may be far enough apart that pipe properties such as material and diameter notably vary over that distance, making the interpretation of turbidity observations more difficult. Hydrants could be particularly widely spaced in trunk systems (§2.1.17.2). The temporal resolution may also be low relative to model characteristic times such as t_t and t_d .

Secondly, and on a related note, certain latent quantities that define a calibrated model are more difficult to derive from turbidity observations than others. It is more challenging to characterise the highly latent material accumulation process (including the regeneration rate) for a pipe than to quantify the material erosion model (including the wall state and erosion rate) as, given flow and turbidity signals, information regarding accumulation can only be deduced from two or more characterised erosion events. The number, nature and spacing of events in the input signals and the [SNR](#) of those signals are important for calibrating the latent model quantities with sufficient accuracy (e.g. fig. 5.31).

For example, certain events in the turbidity signal may allow the erosion process and shear strength profile at the pipe wall to be characterised but there is insufficient information to quantify the regeneration rate. Pipe flow may therefore need to be deliberately increased in a controlled manner on one or more occasions to generate enough observational data for accurate calibration with regards to regeneration. Such interventions are often undertaken when calibrating the [PODDS](#) erosion model (§2.2.8.4). Note that manipulating flow for model calibration not only produces useful observations but also disturbs the system in a destructive manner, hence the need for such flow (and shear stress) increases to be precisely and accurately controlled.

5.7.2 An idealised method for fitting the proposed model to data

Returning to the question of how to fit the model to data, an automated fitting method is considered to be far preferable to hand fitting for the reasons outlined in the previous section. To this end, the calibration of model parameters can be formulated as an optimisation problem (as per §2.2.4) that can be solved via automated means:

- The objective function to be minimised is an appropriately sensitive scalar measure of the dissimilarity between turbidity observations and model instance turbidity predictions over the calibration period;
- The parameters to be optimised are β_e , α and β_r (unless one is confident that the duration of the simulation period is negligible compared to the time required for full regeneration);
- The optimisation constraints are independent parameter bounds for β_e , α and β_r : all three parameters must be positive and finite.

The wall state boundary conditions can also be fitted by optimisation, although this is not the only way of arriving at values for them. Options for fitting/estimating the wall state boundary conditions are discussed in more detail in §5.8.

5.7.3 Objective function

The dissimilarity metric used in the objective function for this optimisation problem needs to satisfy the following requirements:

- The metric value should approach a known finite limit such as 1 if the observation and prediction vectors are very similar;
- The objective function is to be used to fit parameters that collectively influence both shape and scale (table 5.3); the dissimilarity metric should therefore be sensitive to differences in both shape and scale between input vectors;
- The metric should not be overly sensitive to outliers as the [SNR](#) of turbidity instrumentation can be affected by degassing and moving/knocking;

- The metric should be trivial to calculate as it may be necessary to execute an objective function thousands of times if fitting the model using an optimisation method similar to those detailed in §2.2.4.

The dissimilarity metric should also be dimensional rather than relative as this should aid fitting (§2.2.5).

The integral of squared errors (ISE) was chosen as it is moderately sensitive to the magnitude of transformations, is very insensitive to noise, is simple to calculate (eq. 5.38) and is dimensional (Sonnenwald et al., 2013).

$$ISE(T_{obs}, T_{pred}) = \sum (T_{obs}(t) - T_{pred}(t))^2 \quad \forall t \in t_{compare} \quad (5.38)$$

where $t_{compare}$ is the set of times at which errors are to be calculated (which do not need to be contiguous).

It may be necessary to pre-process one or both of the time-series inputs to the dissimilarity metric:

- One of the time-series may need to be resampled to ensure a one-to-one pairing of 'observation' and prediction values (i.e. a common time-series 'index'). This may result in 'missing' (N/A) values;
- Missing values in turbidity observation time-series may need to be infilled using say linear interpolation, back-filling or forward-filling. Alternatively, times when one of the two metric inputs is missing can be ignored when calculating the metric.

It may be necessary to ignore certain timesteps/period when calculating the dissimilarity metric, such as:

- Periods when a turbidity meter was known to be generating erroneous data (when the flow was zero or the instrument was moved/connected/disconnected);
- Periods between turbidity 'events' when the background turbidity is not negligible and cannot easily be accounted for.

5.7.4 Fitting by optimisation

5.7.4.1 Approaches for finding a global optimum

If the problem space is complicated and multi-dimensional then one needs a way of searching for the global optimum in an efficient manner. Exhaustive search for the best fitting parameters (given wide possible ranges for each parameter) is not appropriate, even if the process is automated.

It has conservatively been assumed that the problem space for fitting the proposed model to data is complicated (non-convex and multi-modal) as multiple parameters have a similar effect on predictions over certain temporal scales that could be difficult to distinguish by trial and

error (§5.7.1). Also, testing for a convex problem space given the variety of possible time-series model inputs is non-trivial. If the problem space is not convex and uni-modal then simple deterministic gradient descent methods are unsuited to this optimisation problem (§2.2.4): such methods are very sensitive to how they are initialised and may struggle to find the global optimum (the parameters that give lowest dissimilarity metric value and therefore the best model fit).

A better approach is to use a stochastic optimisation method that intelligently explores different parts of the problem space simultaneously so as to increase the probability of finding the global optimum (§2.2.4). Also, the stochasticity makes such methods more robust when presented with noisy data. Two methods that have been used to success in environmental engineering (§2.2.4.5) are genetic algorithms and particle swarm optimisation (PSO), which mimic biological evolution (§2.2.4.2) and the swarming of fish and birds (§2.2.4.3) respectively so as to (be likely to) find global optima.

5.7.4.2 Chosen optimisation algorithm: PSO

PSO is considered preferable to a genetic algorithm for fitting the proposed model to data. Genetic algorithms have the advantage over PSO of being able to optimise both numerical and categorical parameters (§2.2.4.2). However, no input parameters are categorical so PSO is a more attractive method in this case due to its simplicity and efficiency of design, implementation and execution (§2.2.4.3).

The PSO configuration options used when fitting the proposed model to data are specified in table 5.4 (bar the specific parameter space bounds and convergence tolerance values as these more case-specific).

5.7.4.3 Chosen PSO implementation

Although various implementations of PSO exist (§2.2.4.3), the version used in this project, pyshoal³, was developed by the author to provide the features and satisfy the requirements listed in table 5.4. pyshoal is a Python package that provides a simple API. Once instantiated, an optimisation run continues until convergence tolerances are satisfied for a contiguous number of iterations or a maximum number of iterations is reached. The outputs following a run are the best objective function parameter values found during the run, the value of the objective function given those parameter values and the number of iterations. The software allows for the optimisation of arbitrary objective functions (which can take any number of parameters) and can execute objective functions in parallel using multiple processes for reducing the execution time of a single optimisation run (§2.2.4.3).

5.7.5 Optimising model parameters

The process of developing and verifying a method for fitting the proposed model to data can be made simpler by first considering how best to fit just the model parameters to observations

³<https://github.com/willfurnass/pyshoal>

Table 5.4: Generic *PSO* configuration used when fitting the proposed model to data.

Option	Choice	Reason/reference
Number of particles	25	Smallest square number in the recommended range of 20-50 (§2.2.4.3). Must be a square number if using Von Neumann neighbourhoods (see 'particle neighbourhoods' below).
Parameter space box bounds (valid parameter ranges)	(No generic values as case specific)	-
Particle initialisation	Uniformly distributed within parameter space box bounds	No justification for using more complicated methods such as Latin Hypercube sampling (Wymer, 2007).
Velocity component weights	Inertia weight decreases linearly from 0.9 to 0.4 over the total number of optimisation run iterations; nostalgic and societal weights are both 2.1	References listed in §2.2.4.3.
Problem space boundary handling	Restrictive, damping boundary conditions	See §2.2.4.3.
Maximum particle velocity	The absolute size of the region bounded by the box bounds.	See §2.2.4.3.
Particle neighbourhoods	Social rather than geometric; exactly four neighbours per particle (Von Neumann lattice)	See §2.2.4.3.
Maximum number of iterations	500	-
Convergence tolerance check	The algorithm terminates before the maximum number of iterations if the 'best' position in the swarm changes by less than a similarly-dimensioned tolerance threshold over five <i>PSO</i> iterations. Tolerance thresholds not specified here as they are case-specific.	-

given known boundary conditions then progressing on to consider how to fit/estimate the boundary conditions. Both 'sub-methods' can then be combined to fully calibrate a model.

5.7.5.1 Fitting model parameters simultaneously

One approach to calibrating the model parameters is to fit all three parameters simultaneously given static and time-series observations of pipe and bulk water properties, a PSO configuration (including parameter bounds) and known wall state boundary conditions ($\varphi(\tau, t_0)$). This approach, illustrated in fig. 5.42, shall henceforth be referred to as FIT_SIMULT.

5.7.5.2 Fitting model parameters in a nested fashion

The efficacy of FIT_SIMULT for accurately calibrating VCDM model parameters may be limited by those parameters affecting turbidity predictions in similar ways under certain conditions (table 5.3). A more robust approach could be to fit parameters in a nested fashion.

Let us momentarily consider the case of material erosion without regeneration. The α parameter has a much simpler effect on the model output than β_e as it influences the scale but does not influence the shape of turbidity predictions (fig. 5.28), whereas β_e influences both scale and shape (figs. 5.32 and 5.33). The effects of the two parameters can be separated to some extent by searching for just the optimal β_e value for an erosion event using the PSO objective function but then *within that function* finding the α value that best complements the current candidate β_e value. This 'inner' optimisation is relatively trivial: α can be found by a deterministic, gradient-based method such as Brent's algorithm⁴ (§2.2.4), as α only linearly influences the scale of output. The objective function of this inner optimisation can again be the ISE between time-series turbidity observations and predictions.

The two nested objective functions required for this approach (henceforth referred to as FIT_NESTED) can be seen in fig. 5.43.

5.7.5.3 Fitting a subset of model parameters after calculating others

An alternative approach to finding the values of α and β_e that give the best model fit for erosion events is to first *calculate* α directly from data then subsequently fit β_e using PSO. Again,

⁴Implemented in Python as `scipy.optimize.brent`.

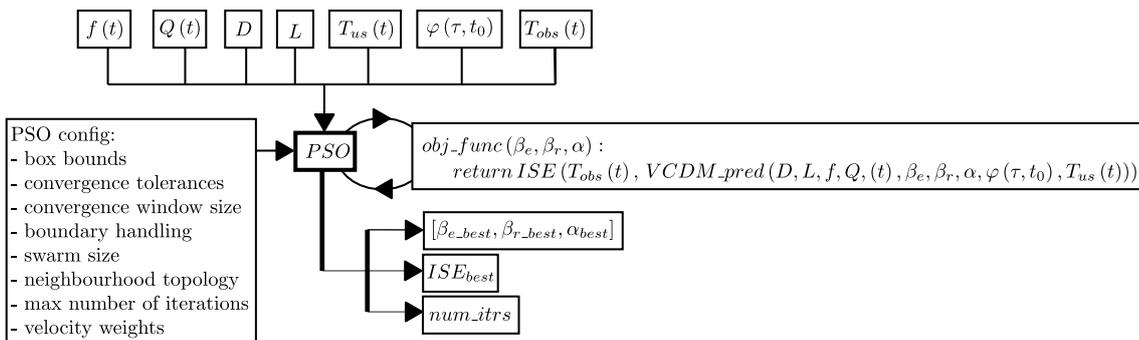


Figure 5.42: A data flow chart showing how all three VCDM parameters can be fitted to data simultaneously using PSO (FIT_SIMULT method).

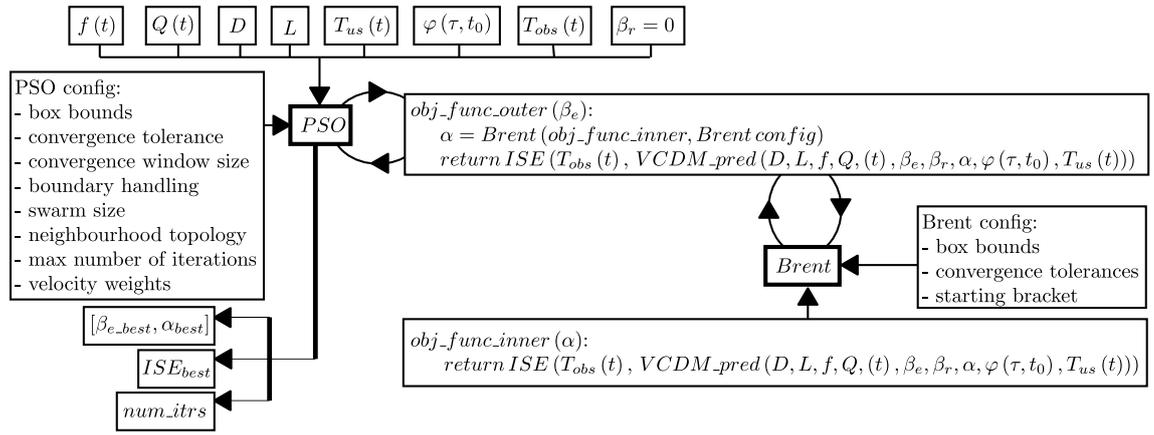


Figure 5.43: A data flow chart showing how the **VCDM** erosion parameters α and β_e can be fitted to data in a nested fashion using **PSO** and **Brent's algorithm** (Brent, 1972) (**FIT_NESTED** method). Here, the bracket is the range of starting values used to initialise **Brent's algorithm** (not the parameter bounds).

like **FIT_NESTED** this method (henceforth be referred to as **CALC_AND_FIT**) attempts to separate the effects of β_e and α . What follows demonstrates how α can be calculated from data.

Let us assume that over a time period $t \in [t_0, \infty]$ that regeneration is negligible ($\beta_r = 0 \text{ s}^{-1}$), that the shear stress τ_a increased from τ_1 to τ_2 over that period and that all material with a strength $\tau \in [\tau_1, \tau_2]$ was fully regenerated at t_0 and was fully depleted at $t = \infty$. Let us also assume that the bulk water turbidity, T_{ds} , was solely a function of the erosion of this material and not of the influx of material from upstream of the pipe of interest.

The total amount of material mobilised as a result of the excess shear stress $\tau_e = \tau_2 - \tau_1$ can be determined by integrating the flux of material exiting the pipe over $t \in [t_0, \infty]$ (eq. 5.3). The total amount of material to be mobilised can also theoretically be calculated by integrating the material release per m^2 of wall area over $t \in [t_0, \infty]$ then multiplying by the pipe wall area (eq. 5.2). Therefore:

$$\int_{t_0}^{\infty} Q(t) \cdot T_{ds}(t) dt = A_w \cdot \int_{t_0}^{\infty} \frac{dN(t)}{dt} dt \quad (5.39)$$

Although the execution of a **VCDM** simulation requires solving of discrete integrals, here continuous integrals can and have been used for simplicity. Substituting in a concise form of eq. 4.5 yields:

$$\int_{t_0}^{\infty} Q(t) \cdot T_{ds}(t) dt = A_w \cdot \int_{t_0}^{\infty} (-\alpha) \int_{\tau_1}^{\tau_2} \frac{\partial \varphi(\tau, t)}{\partial t} d\tau dt \quad (5.40)$$

By factoring $-\alpha$ outside the integral then reversing the order of integration, it can be seen that the integration and differentiation with respect to time can be simplified like so:

$$\int_{t_0}^{\infty} Q(t) \cdot T_{ds}(t) dt = -\alpha \cdot A_w \cdot \int_{t_0}^{\infty} \int_{\tau_1}^{\tau_2} \frac{\partial \varphi(\tau, t)}{\partial t} d\tau dt \quad (5.41)$$

$$= -\alpha \cdot A_w \cdot \int_{\tau_1}^{\tau_2} \int_{t_0}^{\infty} \frac{\partial \varphi(\tau, t)}{\partial t} dt d\tau \quad (5.42)$$

$$= -\alpha \cdot A_w \cdot \int_{\tau_1}^{\tau_2} \Delta \varphi(\tau, t) d\tau \quad (5.43)$$

If the relative amount of material at the wall $\varphi(\tau, t)$ goes from being maximal (1) to nothing (0) over the period $t \in [t_0, \infty]$ for all $\tau \in [\tau_1, \tau_2]$ then $\Delta \varphi(\tau, t) = -1$ and so

$$\int_{t_0}^{\infty} Q(t) \cdot T_{ds}(t) dt = -\alpha \cdot A_w \cdot \int_{\tau_1}^{\tau_2} (-1) d\tau \quad (5.44)$$

Evaluating the integral gives

$$\int_{t_0}^{\infty} Q(t) \cdot T_{ds}(t) dt = \alpha \cdot A_w \cdot (\tau_2 - \tau_1) \quad (5.45)$$

$$= \alpha \cdot A_w \cdot \tau_e \quad (5.46)$$

Therefore, α can be calculated using eq. 5.48 given the conditions specified earlier in this subsection. An optimal, complementary β_e value can subsequently be found using PSO (as illustrated in fig. 5.44).

$$\alpha = \frac{\int_{t_0}^{\infty} Q(t) \cdot T_{ds}(t) dt}{A_w \cdot \tau_e} \quad (5.47)$$

$$\approx \frac{\sum_{i=0}^{t_i \leq t_{max}} Q(t) \cdot T_{ds}(t_i) \Delta t}{A_w \cdot \tau_e} \quad (5.48)$$

This expression can be used to find optimal values for α and β

The conditions required for CALC_AND_FIT to be a valid method can be summarised as follows:

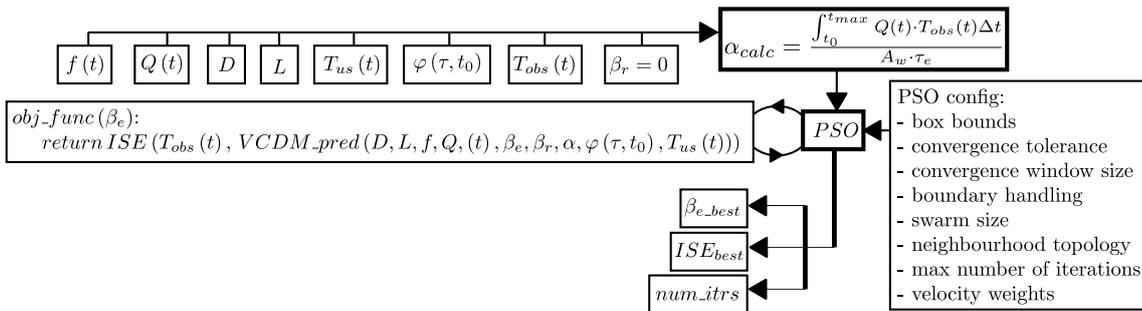


Figure 5.44: A data flow chart showing how optimal VCDM erosion parameters α and β_e be found during model calibration by first calculating α directly from data (using eq. 5.48 that is independent of β_e) then finding the best β_e value using PSO (CALC_AND_FIT method).

- τ_a increases from τ_1 to τ_2 during the period $[t_0, t_{max}]$;
- This causes the complete erosion of the material for which $\tau_1 \leq \tau \leq \tau_2$ i.e. $\Delta\varphi = -1$ over $[t_0, t_{max}]$;
- To accurately quantify this complete erosion using just $T_{ds}(t)$, $T_{ds}(t)$ must be effectively zero by t_{max} ;
- Complete erosion requires that the amount of material at the pipe wall must initially have been maximal i.e. the boundary conditions must be $\varphi(\tau_1 \leq \tau \leq \tau_2, t_0) = 1$.
- Regeneration is negligible over $[t_0, t_{max}]$.

The requiring of very simple boundary conditions and for the turbidity response to drop to negligible levels at the end of the calibration period means this method is applicable for far fewer model configurations and input datasets than FIT_SIMULT and FIT_NESTED.

5.7.5.4 Estimating/bounding the α parameter using known values of the **PODDS** k parameter

The fitting or calculating of α can be aided by comparing that parameter to the k parameter in **PODDS**. The parameters have been found to be equivalent given the pre-conditions required for the use of the CALC_AND_FIT method. The value of α can therefore be estimated and/or bounded using the k values from successfully calibrated **PODDS** models of pipes with similar pipe materials and diameters.

Martin Jackson of Wessex Water and the **PODDS** Consortium⁵ noted⁶ that the total amount of material mobilised from a step increase in shear stress (quantified as the material flux integrated over time) divided by the pipe surface area and excess shear stress is the negative of the **PODDS** k parameter (eq. 5.49). Both this expression and k have units of $TPMU \cdot Pa^{-1} \cdot m^{-2}$. This equivalence was validated by making a **PODDS** prediction using a known value of k then 'recovering' that value using eq. 5.49.

$$k \approx - \frac{\sum_{i=0}^{t_i \leq t_{max}} Q(t) \cdot T_{ds}(t_i) \Delta t}{A_w \cdot \tau_e} \quad (5.49)$$

Note that eq. 5.49 is identical to that for calculating α (eq. 5.48) (bar the negative sign), so here the absolute values of α and k are equivalent. Previously identified values of k can therefore be used to bound α when using Brent's method as part of FIT_NESTED or to directly estimate α when using CALC_AND_FIT.

5.7.5.5 Validating the three parameter estimation methods

Fitting artificial erosion events given known simple boundary conditions The ability of FIT_SIMULT, FIT_NESTED and CALC_AND_FIT to fit the proposed model to a variety of short-term erosion events was explored using artificial **PODDS** and **VCDM** 'target' predictions,

⁵<http://www.podds.co.uk>

⁶In personal correspondence on 24th April 2013.

generated using known parameters and boundary conditions. The set of **VCDM** or **PODDS** model parameters, flow levels and pipe attributes associated with each of these target predictions shall henceforth be known as a *target attribute set*.

For each erosion event, the flow increased instantaneously at $t = t_0$. For the **VCDM** target predictions, the relative amount of material at t_0 was 1 for shear strengths above the prevailing, constant shear stress before t_0 (corresponding to flow of Q_{cond}), and was 0 for shear strengths below the prevailing shear stress. For the **PODDS** target predictions, the shear strength before t_0 was the shear strength at Q_{cond} . No regeneration occurred over the simulations. The **VCDM** and **PODDS** target attribute sets are shown in tables 5.5 and 5.6 respectively; asset attributes were chosen to be representative of a variety of pipes. The simulation timesteps were from t_0 to five hours in 5 s increments and pipe lengths were chosen so that the turnover time at Q_{flush} (and peak turbidity) occurred at $t_0 + 3600$ s.

Each of the eight erosion event models were used to generate two predictions: one with the idealised step increase in flow and a second after having added noise to the flow profile to mimic natural fluctuations in flow⁷. Each of the methods FIT_SIMULT, FIT_NESTED AND CALC_AND_FIT was then used to try to recover the β_e and α **VCDM** parameter values (or **PODDS** k value) for each of the resulting 8×2 model predictions, resulting in $8 \times 2 \times 3$ model fits and **ISE** values. The boundary conditions used to generate the predictions were provided as inputs to the fitting processes.

The **PSO** configuration used for testing these model fitting methods was as per table 5.4 with the following parameter bounds and convergence tolerances:

- Parameter box bounds:
 - $\alpha \in [1 \times 10^{-7}, 10]$; the bounds of α were set by acknowledging the conceptual equivalence with the **PODDS** k parameter (§5.7.5.4; under certain conditions). The bounds encompass the (negative of the) range of k values found by successfully calibrating a variety of **PODDS** models by hand (Husband and Boxall, 2010b);
 - $\beta_e \in [1 \times 10^{-7}, 10]$; there is no known direct equivalent of β_e in the **PODDS** model so, without a posteriori knowledge, the bounds of β_e were simply set to the same range as α .
- Convergence tolerances (all need to be satisfied over five iterations for the algorithm to terminate; all are an order of magnitude smaller than the corresponding lower bounds):

⁷The noise was generated by finding the moving average of a time-series of uniformly sampled values.

Table 5.5: Target attribute sets used for generating target **VCDM** predictions given known simple boundary conditions ($\varphi(\tau \leq \tau_a(D, Q_{cond}, k_s), t_0) = 0$; $\varphi(\tau > \tau_a(D, Q_{cond}, k_s), t_0) = 1$) and negligible regeneration ($\beta_r = 0$ s⁻¹). The target predictions were used for testing the efficacy of several **VCDM** parameter fitting methods.

Quantity	D	k_s	Q_{cond}	Q_{flush}	L	β_e	α
Units	mm	mm	$\frac{L}{s}$	$\frac{L}{s}$	m	$\frac{1}{Pa \cdot s}$	$\frac{TPMU}{Pa \cdot m^2}$
VCDM-1	75	0.1	1	4	3259	0.002	1.0
VCDM-2	203	5.0	1	6	667	0.002	1.5
VCDM-3	400	0.3	4	40	1146	0.005	1.0

Table 5.6: Target attribute sets used for generating target **PODDS** predictions given known simple boundary conditions ($C_{max} = 100 \text{ TPMU} \cdot \text{m}^{-2}$; $\tau_c(t_0) = \tau_a(D, Q_{cond}, k_s)$) and negligible regeneration ($\beta_r = 0 \text{ s}^{-1}$). The target predictions were used for testing the efficacy of several **VCDM** parameter fitting methods.

Quantity	D	k_s	Q_{cond}	Q_{flush}	L	P	n	k
Units	mm	mm	$\frac{L}{s}$	$\frac{L}{s}$	m	$\frac{\text{TPMU}}{\text{Pa} \cdot \text{m}^2 \cdot \text{s}}$	—	$\frac{\text{TPMU}}{\text{Pa} \cdot \text{m}^2}$
PODDS-1	150	0.1	1	9	1833	0.02	1.2	-3.5
PODDS-2	100	0.1	1	6	2750	0.0002	2.5	-0.5
PODDS-3	50	0.1	4	9	16501	0.0002	3.0	-0.5
PODDS-4	89	5.0	1	6	3480	0.00022	3.0	-0.5
PODDS-5	700	0.25	15	500	4677	0.005	1.0	-3.0

– $\alpha : 1 \times 10^{-8}$;

– $\beta_e : 1 \times 10^{-8}$.

The results were as follows:

All three fitting methods were successful at recovering β_e and α for the three **VCDM** predictions with non-noisy flow profiles. The accuracy of the **CALC_AND_FIT** method was worse for **VCDM** predictions where the turbidity was notably greater than zero by the end of the modelled time period (e.g. fig. 5.45); this is because the error in the calculated estimate in α is greater for $t_{max} \ll \infty$ (see eqs. 5.48 and 5.48).

The flow-noise-free **PODDS** target predictions were slightly more difficult to fit; the target prediction and the best turbidity predictions identified by the fitting methods often differed slightly in shape (e.g. fig. 5.46). This was not due to the **VCDM** not having a power term like **PODDS'** n , which is known to influence shape (fig. 5.41), as fits were slightly inaccurate even for target predictions made using $n = 1$. The difference in shape and scale between target and prediction profiles was not particularly large, indicating that the **VCDM** was able to provide similar functionality to the validated **PODDS** erosion model.

The predictions generated using noisy flow profiles proved more difficult to fit: on several occasions the **FIT_SIMULT** and **FIT_NESTED** methods both returned values of β_e that were far too fast (e.g. fig. 5.47), yielding low **ISE** values but finding model fits that were unlikely to offer accurate predictions for $t > t_{max}$. The **CALC_AND_FIT** method proved more robust, suggesting that only having to fit one of the two erosion parameters that influence the scale of

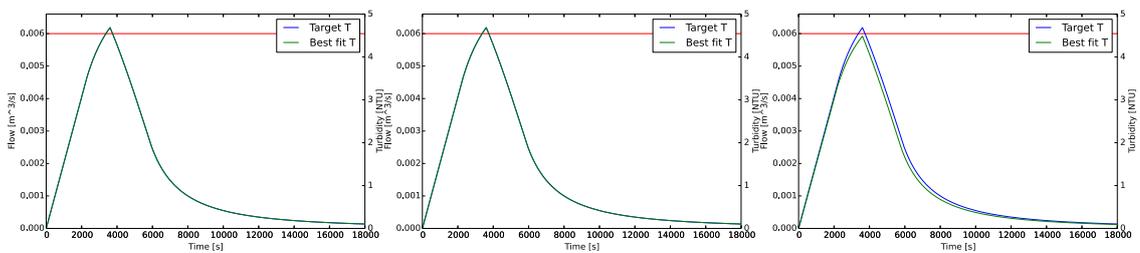


Figure 5.45: Attempts to fit the **VCDM** to target prediction **VCDM-2** (see table 5.5; without flow noise) using the following methods: (a) **FIT_SIMULT** ($\beta_e = 0.002$; $\alpha = 1.50$; $ISE = 1.6 \times 10^{-14}$; converged after 141 iterations); (b) **FIT_NESTED** ($\beta_e = 0.002$; $\alpha = 1.50$; $ISE = 4.4 \times 10^{-9}$; 45 iterations); (c) **CALC_AND_FIT** ($\beta_e = 0.002$; $\alpha = 1.41$; $ISE = 39$; 52 iterations). The red line is pipe flow, which increased at $t = 0 \text{ s}$ (see table 5.5).

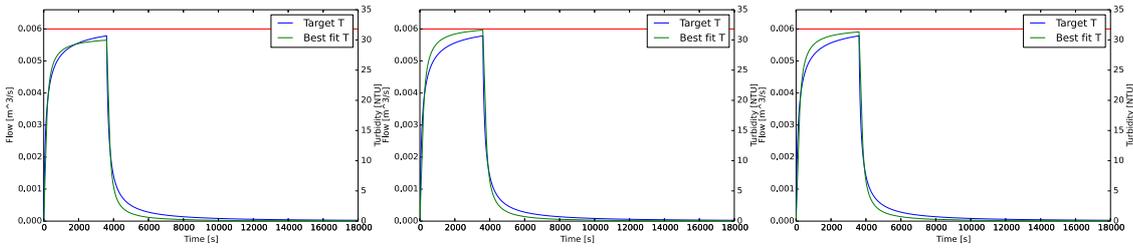


Figure 5.46: Attempt to fit the **VCDM** to prediction **PODDS-2** (see table 5.6 without flow noise) using the following methods: (a) **FIT_SIMULT** ($\beta_e = 0.0037$; $\alpha = 0.46$; $ISE = 2374$; converged after 163 iterations); (b) **FIT_NESTED** ($\beta_e = 0.0033$; $\alpha = 0.49$; $ISE = 4129$; 64 iterations); (c) **CALC_AND_FIT** ($\beta_e = 0.0033$; $\alpha = 0.48$; $ISE = 3563$; 66 iterations). The red line is pipe flow, which increased at $t = 0$ s (see table 5.5).

turbidity prediction is easier when the flow is unstable, most likely a result of both β_e and α influencing the scale of turbidity predictions.

Fitting artificial erosion events given known non-trivial boundary conditions Given that the ability to track the amount of material at different shear strengths over time is a key feature of the proposed model it was important to verify that the proposed fitting methods could recover β_e and α given more complicated boundary conditions i.e. non-trivial values of $\varphi(\tau, t_0)$. This was important as in practise it is unlikely that the amount of material at the pipe wall is constant with shear strength.

The **FIT_SIMULT** and **FIT_NESTED** methods were used to fit the proposed model to:

- target predictions made using three different **VCDM** target attribute sets (table 5.5),
- which were made with and without noisy flow profiles and
- six different profiles for $\varphi(\tau, t_0)$ (fig. 5.48), which were known (and not fitted/estimated) by the fitting methods.

The **CALC_AND_FIT** method is not appropriate when $\varphi(\tau, t_0)$ is non-trivial; it has the precondition that $\varphi(\tau, t_0)$ must be polarised around a particular shear strength (§5.7.5.3).

The error between target predictions and ‘fitted’ predictions was negligible for the majority of the $3 \times 2 \times 6 = 36$ optimisation runs. However, for a particular target prediction made using a noisy flow profile, both **FIT_SIMULT** and **FIT_NESTED** found near-identical but erroneous

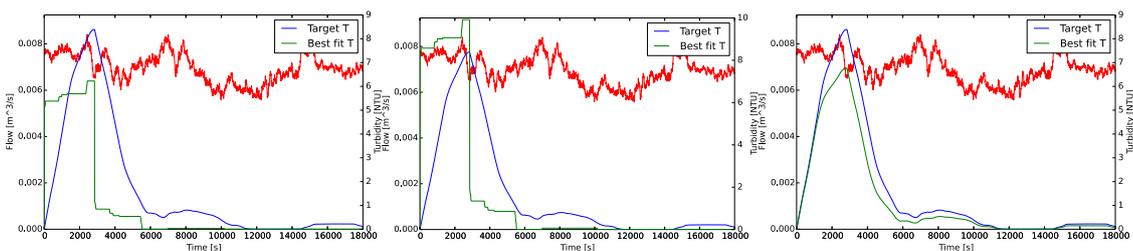


Figure 5.47: Attempt to fit the **VCDM** to prediction **VCDM-1** (with flow noise) using the following methods: (a) **FIT_SIMULT** ($\beta_e = 10.0$; $\alpha = 0.71$; $ISE = 12746$; 89 iterations); (b) **FIT_NESTED** ($\beta_e = 9.85$; $\alpha = 1.14$; $ISE = 19281$; 144 iterations); (c) **CALC_AND_FIT** ($\beta_e = 0.0025$; $\alpha = 1.13$; $ISE = 1508$; 63 iterations). The red line is pipe flow, which increased at $t = 0$ s (see table 5.5).

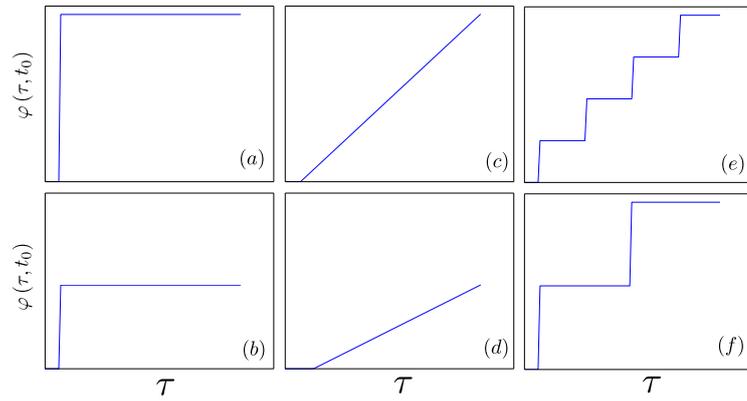


Figure 5.48: Fitting artificial erosion events given known non-trivial boundary conditions: $\varphi(\tau_i, t_0)$ shapes of target predictions are (a) Full regeneration; (b) half regeneration; (c) steep ramp; (d) shallow ramp; (e) four steps; (f) two steps.

fits; the β_e value was far too fast in both cases. These results could not be reproduced when repeating these two fitting attempts with different noise, suggesting that the fitting process input data contained insufficient information to distinguish β_e and α , or possibly that there is a need to take the best results of more than one fitting attempt for each PSO configuration. A second FIT_SIMULT result also had an overly quick β_e but otherwise the results for both fitting methods were very similar.

Fitting both erosion and regeneration for artificial observations given non-trivial boundary conditions

It is important to be able to reliably fit the parameters relating to *regeneration* as well as those relating to erosion. A third set of fitting tests were performed to see how well FIT_SIMULT and FIT_NESTED could recover β_e , α and β_r . Target predictions were generated for all 36 permutations of:

- Three sets of asset characteristics and VCDM parameters;
- Non-noisy versus noisy flow (with the added noise here being generated using a random walk);
- Six different $\varphi(\tau, t_0)$ profiles (fig. 5.48).

Table 5.7: VCDM target attribute sets used to generate VCDM predictions used for fitting both erosion and regeneration for artificial observations given non-trivial boundary conditions. Timesteps per turnover: 200 during each of the two flushes, 10 during the intervening period.

	D	k_s	Q_{cond}	Q_{flush}	L	β_e	α	β_r^{-1}	Turnovers per flush	Time between flushes
Units	mm	mm	$\frac{L}{s}$	$\frac{L}{s}$	m	$\frac{1}{Pa \cdot s}$	$\frac{TPMU}{Pa \cdot m^2}$	years	-	years
VCDM-4	75	0.1	0.5	4	300	0.002	1.0	0.2	5	0.2 / 4 = 0.04
VCDM-5	203	5.0	1	6	1000	0.002	1.5	0.5	3	0.5 / 12 = 0.04
VCDM-6	400	0.3	4	40	3000	0.005	2.0	1.5	3	1.5 / 15 = 0.1

The flow profile of each of these predictions consisted of two periods of higher flow (emulating pipe flushes) separated by a long period of low flow (e.g. fig. 5.49), allowing regeneration to occur between the two erosion events. The regeneration rate parameter, β_r , and duration of the low flow period were chosen so that maximal regeneration had not occurred by the start of the second flush, the effect of this being that the both the amount of material mobilised and the shape of the resulting turbidity response differed between the two flushes.

The PSO configuration was as per the first set of fitting tests described in this section but with bounds and a convergence tolerance defined for the β_r parameter. Complete regeneration has previously been estimated by linear extrapolation to take 1.5 to 4 years in DMA mains (§2.1.14). As the reciprocal of β_r is the minimum time required for complete regeneration to occur following complete material depletion, the bounds of β_r were set to the reciprocals of 0.2 to 10 years i.e. $\beta_r \in [1/(86400. \times 365 \times 10), 1/(86400. \times 365 \times 0.2)] s^{-1} = [3.17 \times 10^{-9}, 1.58 \times 10^{-7}] s^{-1}$. The convergence tolerance was set to $1 \times 10^{-10} s^{-1}$, one order of magnitude less than the lower bound.

The ability of the FIT_SIMULT and FIT_NESTED methods to recover the parameters of each of the 36 target datasets was investigated, requiring 72 distinct fitting tests. As PSO is a non-deterministic *metaheuristic* that is likely but not guaranteed to find the global optimum (§2.2.4), each of this set of fitting tests was attempted up to three times to achieve a satisfactory fit. This required the qualification of fit in a generic way so as to define a numeric threshold for what is considered acceptable. Absolute measures of the dissimilarity between observations and predictions (e.g. ISE) are of no use here as they are test-specific (§2.2.5). The *Coefficient of Determination* (R^2 ; see eq. 5.50) was used as a non-dimensional means for quantifying and comparing model fits in a test-independent way. This relative dissimilarity metric was chosen because it is sensitive to the type and magnitude of signal transformations but not to the number of data points in the time-series nor to signal noise (Sonnenwald et al., 2013). A given fitting test was reattempted up to two times if $R^2 < 0.95$ (where 1 is ideal).

$$R^2 = \left[\frac{\sum ((a_i - \bar{a}) \cdot (b_i - \bar{b}))}{\sqrt{\sum (a_i - \bar{a})^2 \cdot \sum (b_i - \bar{b})^2}} \right]^2 \quad (5.50)$$

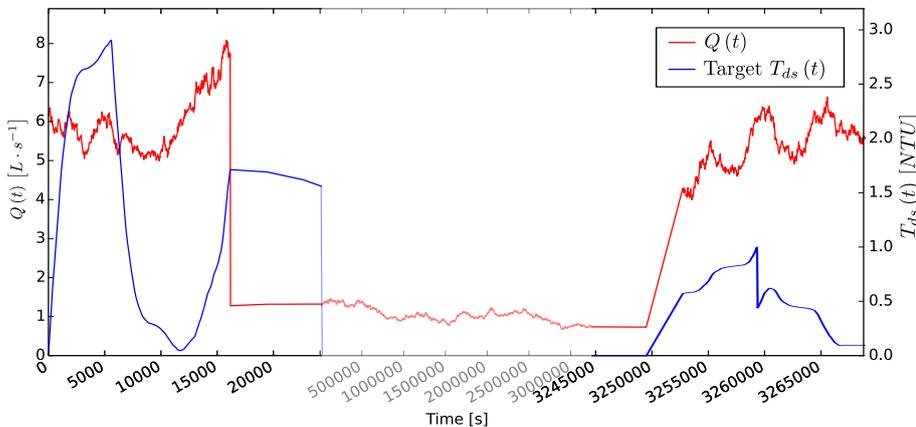


Figure 5.49: Target VCDM dataset VCDM-5 (see table 5.7), generated using a noisy flow profile and the ‘two steps’ boundary condition ($\varphi(\tau, t_0)$) shape. Time has been condensed in the middle of the time axis for the purposes of visualising all data within a single figure.

All attempts at all 72 fitting tests converged in less than 500 PSO iterations. For all fitting tests where noise had not been added to the flow signal, the value of R^2 was effectively 1.00 and the ratio of the fitted parameters to the target parameters was unity (to 2 d.p.) other than in one case where those ratios were 99.77%, 100.23% and 92.77% for β_e , α and β_r respectively. Fitting target datasets with fluctuating flow profiles proved more difficult. More than one attempt was required for twelve of the 36 corresponding fitting tests. For seven of those tests, a good fit ($R^2 \geq 0.95$) was found on the second or third attempt⁸, indicating that, for these cases, satisfactory model fits could be achieved given enough PSO runs. For another of the 36 tests, R^2 increased over three fitting attempts but no satisfactory solution was found. It is not known whether a good fit would have been found through conducting more than three fitting attempts. However, for four tests, the same R^2 was found for each of the three fitting attempts, suggesting that there was insufficient information in the inputs to the optimisation method for truly optimal parameter values to be easily found.

Common to all unacceptable fits was a very high β_e value (the fitted value was on average 2221 times greater than the target value amongst these fitting attempts). Note that β_e was the only one of the three parameters for which there was no prior information to inform the bounds during model fitting. These results indicate that the upper bound of β_e should be reduced from 10 to say 1.0 or 0.1 to avoid local optima.

One fitting method was not obviously preferable to the other: four of the FIT_SIMULT tests required multiple attempts whereas eight of the FIT_NESTED tests required multiple attempts. However, of the subset of tests where no acceptable solution was found after three attempts, only two tests were associated with FIT_SIMULT but three tests were associated with FIT_NESTED. As both methods perform similarly well, FIT_SIMULT is the preferred method as FIT_NESTED is less computationally efficient.

There were no obvious correlations between the shape of $\varphi(\tau_i, t_0)$ and the ease of fitting.

For the satisfactory model fits, it was the β_r parameter that was associated with the greatest relative error. The mean and minimum of the ratios of the fitted parameter value to the target value were 0.92 and 0.38 (fit shown in fig. 5.50), indicating that the fitting methods are prone to underestimating β_r and that the fitting method inputs do not contain sufficient information to accurately characterise regeneration using the chosen dissimilarity metric (which could result in unacceptable predictions following model calibration). This raises questions about whether the target/observed data needs to feature more than two erosion events and/or less rapidly fluctuating flow for parameter errors to be detectable using ISE and for regeneration to be well characterised.

In summary:

- 500 PSO iterations is sufficient for convergence;
- FIT_SIMULT and FIT_NESTED perform equally well but FIT_SIMULT is computationally more efficient;
- Flow profiles with both large infrequent plus small and rapid changes result in datasets that are more difficult to fit than those for which flow changes in simple steps;

⁸The flow profile was constant across attempts.

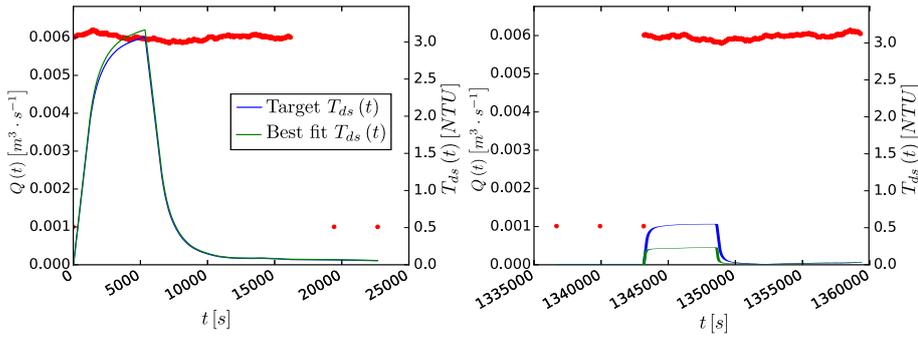


Figure 5.50: Attempt to fit the **VCDM** to target dataset VCDM-5 (see table 5.7; generated using a noisy flow profile and the ‘half ramp’ $\varphi(\tau_i, t_0)$ profile) using the FIT_NESTED method. Fitted parameter values: $\beta_e = 1.981 \times 10^{-3}$, $\alpha = 1.548$, $\beta_r = 2.44 \times 10^{-8} \text{ s}^{-1}$ ($ISE = 20.4$, $R^2 = 0.977$, 303 iterations). Note that β_r of VCDM-5 was considerably faster: $\beta_r = 2.44 \times 10^{-8} \text{ s}^{-1}$ for a β_r^{-1} of 0.5 years.

- Fitting attempts for which $R^2 < 0.95$ should be retried up to two times;
- The shape of the (known) $\varphi(\tau_i, t_0)$ profile had no notable impact on the difficulty of fitting;
- Empirical evidence suggests that the upper bound of β_e should be reduced to less than 10 (e.g. 0.1) so as to avoid local optima.
- Certain datasets contain insufficient information to accurately calibrate the β_r parameter using either the FIT_SIMULT or FIT_NESTED method.

5.8 Fitting/estimating the wall sub-model boundary conditions

5.8.1 Introduction

Calibrating the proposed model requires the fitting to data of not only the three empirical parameters but also the initial amount of material per shear strength at the pipe wall, $\varphi(\tau_i, t_0)$. Different options for fitting these boundary conditions are explored in this section, as are suggestions for using the most appropriate of these options in conjunction with the FIT_SIMULT parameter fitting method.

5.8.2 Fit the initial amount of material at every modelled shear strength band using PSO

One option is to consider $\varphi(\tau_i, t_0)$ as being a vector of model parameters, with m being the number of modelled shear strength bands and $i \in \{1, 2, \dots, m\}$. These plus the parameters β_e , α and β_r could be fitted through optimisation using a variant of the FIT_SIMULT (or FIT_NESTED) method.

This approach may be intuitively appealing but it increases the complexity of the fitting optimisation problem. Firstly it increases the dimensionality of the problem space from 3 to $3 + m$, where sensitivity analysis has indicated that m should be up to 20 for a given step increase in τ_e (table 5.2). Finding an optimal solution to the fitting problem in this space will be more difficult and computationally expensive. Secondly, it requires non-trivial metaheuristic

constraints. **PSO** by default assumes that all parameters are independent so would allow $\varphi(\tau_i, t_0)$ to freely vary with respect to τ_i for each candidate solution. However, if $\varphi(\tau_i, t_0)$ must be monotonically increasing (§4.4.4) then inter-parameter constraints are required to ensure that $\varphi(\tau_i, t_0) \leq \varphi(\tau_{i+1}, t_0)$. Such constraints are non-trivial to enforce when using **PSO**: penalty functions require manual configuration and tuning, ignoring invalid candidate solutions limits the problem space exploration and replacing invalid solutions with valid ones requires that **PSO** be augmented with a local search method (see §2.2.4.4).

5.8.3 Reducing the complexity of the wall state boundary condition vector by approximation using cusps

The dimensionality and therefore the complexity of fitting both the three model parameters and the wall state boundary conditions can be reduced if those boundary conditions are treated not as a freely-varying but monotonically increasing vector, $\varphi(\tau_i, t_0)$, but as a simple shape.

For example, one could assume that at $t = t_0$ all shear strengths less than a ‘cusp’ value are fully exhausted while all shear strengths greater than that (but less than τ_{max}) are maximally regenerated ($\varphi(\tau \leq \tau_{cusp}, t_0) = 0$ and $\varphi(\tau > \tau_{cusp}, t_0) = 1$; see fig. 5.51a). This requires the fitting/estimating of just $3 + 1$ parameters rather than $3 + m$. Note the conceptual similarity between τ_{cusp} and the **PODDS** (scalar) material shear strength (τ_c).

Introducing a second cusp and linearly interpolating the relative quantity of material between the two cusps (fig. 5.51b) offers greater flexibility whilst only increasing the number of parameters (and the number of dimensions of the fitting problem space) by one. Additionally, the relative amount of material at strengths lower than the first cusp could be set to > 0 and/or the amount greater than the higher cusp to < 1 (fig. 5.51c) at a cost of one or two more parameters/dimensions (taking the total to ≤ 9). One issue with using two cusps is that, like the method presented in §5.8.2, complexity is introduced by inter-parameter constraints being required to ensure the monotonicity of $\varphi(\tau, t_0)$ (specifically to ensure $\tau_{low_cusp} \leq \tau_{high_cusp}$ and $\varphi_{low_cusp} < \varphi_{high_cusp}$), although the constraints are only one or two in number rather than the $m - 1$ needed by the method in §5.8.2.

An issue with all forms of this method is that in some cases it may not be flexible enough to approximate the wall state boundary condition vector with sufficient accuracy: shear stress

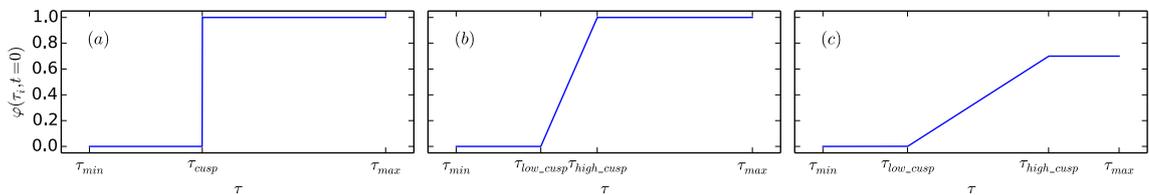


Figure 5.51: This illustrative diagram shows different methods of reducing the complexity of the wall state boundary conditions $\varphi(\tau_i, t_0)$ to a small number of vertices. (a) all shear strengths less than a shear strength could be assumed to be fully depleted whilst all others are maximally regenerated. (b): all shear strengths less than a shear strength τ_{low_cusp} could be fully depleted, all shear strengths greater than a second shear strength τ_{high_cusp} (where $\tau_{high_cusp} > \tau_{low_cusp}$) could be maximally regenerated and linear interpolation could be used to determine the quantity of material over the range $[\tau_{low_cusp}, \tau_{high_cusp}]$. In certain circumstances there may be sufficient evidence for the shear strengths outside $[\tau_{low_cusp}, \tau_{high_cusp}]$ not to be minimal or maximal (e.g. (c)) i.e. to be limited to $[\varphi_{low_cusp}, \varphi_{high_cusp}]$.

levels in trunk mains (the pipes of greatest concern; §2.1.17.2) may change level multiple times over the period required for full regeneration (potentially 1.5 - 4 years; §2.1.14), which should result in boundary vector profiles that are more stepped than the simple profiles shown in fig. 5.51. Also, for the simplest single-cusp model the relative amount of material greater than τ_{cusp} may be an over-estimate.

There are however situations where the simplest single-cusp method may be appropriate e.g. where:

- the conditioning shear stress is known to have been largely invariant before the period of interest or
- there is insufficient information, time-series or otherwise, to justify a more complex approach.

5.8.4 Estimate the wall state boundary conditions using the shear stress history

Fitting the wall state boundary condition vector by optimisation is clearly non-trivial; primarily due to the difficulty in finding an appropriate trade-off between flexibility and complexity when forming the optimisation problem.

A different approach is to use knowledge of system dynamics and variable dependencies to estimate the boundary condition vector (as opposed to treating the system to be calibrated as a black box). In particular, the shear stress history up to a given moment in time can be used to estimate/bound that vector at that moment. The model parameters can then be fitted independently of that vector using data from a time period starting at that moment (see fig. 5.52).

During a time period $t \in [t_L, t_0]$ of say one week (here t_L is the start of a) there will be temporary and non-temporary increases in shear stress. If the maximum shear stress over this period is τ_{a_crit} then this indicates that there will be little if any material for which $\tau < \tau_{a_crit}$ by the time $t = t_0$. For the case of simplicity one can assume that the erosion rate factor β_e is sufficiently fast relative to the duration of the shear stress increases and so all shear strengths less than τ_{a_max} are *fully depleted* by t_0 . Also, one can assume that regeneration is negligible.

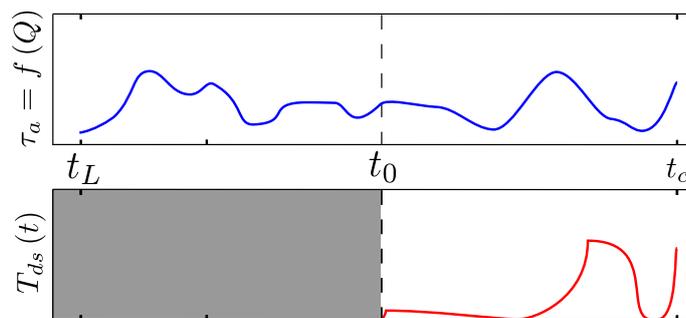


Figure 5.52: The wall state boundary condition vector, $\varphi(\tau, t_0)$ can be estimated using shear stress (τ_a) data from one time period, t_L to t_0 . The model parameters β_e , α and β_r can then be fitted using shear stress and turbidity (T_{ds}) time-series data for t_0 to a future time, t_c , as illustrated here.

The shear stress signal during $[t_L, t_0]$ conveys very little information about the amount of material for which $\tau > \tau_{a_crit}$ at t_0 . It is therefore sensible to assume the worst case scenario, which is that the shear strengths for which $\tau > \tau_{a_crit}$ are maximally regenerated at t_0 .

This approach to estimating $\varphi(\tau_i, t_0)$ is illustrated in fig. 5.53. Note that with this method it is assumed that, as well as erosion being instantaneous, regeneration is negligible.

Here, the length of the period t_L to t_0 is crucial. The estimate of the wall state vector may be inaccurate if the period is too short to contain much information about fluctuations in shear stress. The chance of a high peak in shear stress occurring just *before* t_L but its impact on the wall state vector not being accounted for is greater for shorter periods. However, estimating $\varphi(\tau_i, t_0)$ using a longer period may also result in inaccurate estimates of the wall state vector as the assumption that regeneration is negligible over t_L to t_0 is less valid as $t_0 - t_L$ increases. Say that the peak shear stress, τ_{a_crit} , over t_L to t_0 occurs just after t_L and that the shear stress is then lower for some time until t_0 . This method would estimate that all shear strengths for which $\tau < \tau_{a_crit}$ are fully depleted at t_0 yet if $t_0 - t_L$ is long enough some of them will have significantly regenerated so $\varphi(\tau_i < \tau_{a_crit}, t_0)$ will be an underestimate.

The method is very sensitive to variation in shear stress: if erosion is assumed to be instantaneous then a single high data point could significantly affect the wall state estimate. However, in reality that momentary increase in shear stress could be too short-lived for much erosion to occur, so the calculated wall state at t_0 would undesirably be an underestimate.

5.8.5 Estimate the wall state boundary conditions using the shear stress history and an approximate regeneration rate

The method presented in §5.8.4 can be made more accurate by estimating the wall state vector using not only the shear stress profile calculated from the flow history over the period t_L to t_0 but also an estimate of (the reciprocal of) the regeneration rate, for the reason given at the end of §5.8.4.

The reciprocal of the regeneration rate, β_r^{-1} , is the time required for material at any strength to fully regenerate following complete depletion (assuming no erosion occurs). Estimates of

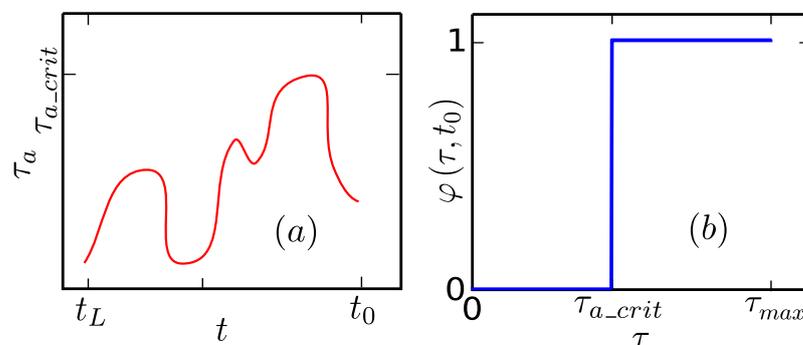


Figure 5.53: The maximum shear stress, τ_{a_crit} , over a period t_L to t_0 (a) can be used to directly estimate $\varphi(\tau_i, t_0)$, the amount of material per shear strength at t_0 (b), as illustrated here. It is assumed that erosion is instantaneous, regeneration is negligible and so at t_0 shear strengths less than τ_{a_crit} are fully depleted and those greater than τ_{a_crit} are conservatively assumed to be fully regenerated. Here the time t_L is the start of a period that leads up to time t_0 .

this *regeneration time* have been estimated for DMA pipes of a variety of pipe materials and supplied with various source waters (§2.1.14).

Let us assume that our estimate of β_r^{-1} is longer than the period t_L to t_0 for which we can calculate the shear stress profile from flow data. An example shear stress profile is shown in black in fig. 5.54a. Let the duration before t_0 since a particular material shear strength τ_i was exceeded by the shear stress τ_a be represented by the function

$$exe(\tau_i, \tau_a(t), t_L, t_0) \quad (5.51)$$

which evaluates to ∞ if τ_i was never exceeded by $\tau_a(t)$ between t_L and t_0 . Again, the material erosion rate factor, β_e , is assumed to be fast enough so that complete material erosion occurred when each material strength was last eroded i.e.

$$\varphi(\tau_i, t = t_0 - exe(\tau_i, \tau_a(t), t_L, t_0)) = 0 \quad (5.52)$$

The relative amount of material at t_0 is therefore the time over which a shear strength has been regenerating prior to t_0 divided by the reciprocal of the estimated regeneration rate, all clipped to the range $[0, 1]$ (eq. 5.53). The result of this is the red line in fig. 5.54b.

$$\varphi(\tau_i, t_0) \cong clip\left(\frac{exe(\tau_i, \tau_a(t), t_L, t_0)}{\beta_r^{-1}}, [0, 1]\right) \quad (5.53)$$

Here it is assumed that β_r is constant w.r.t. t over the period t_L to t_0 (for the reasons given in §4.3.3), although this method could be adapted should a model of the relationship between β_r and recorded time-varying parameters such as temperature be incorporated into the proposed model.

It is thought that the time required for complete regeneration, β_r^{-1} , will often be longer than the period t_L to t_0 for which there is data available for calculating the shear stress: regeneration times are estimated to be up to four years (§2.1.14) yet reliable flow data for calculating the shear stress may not be available over such durations due to flow meter and/or SCADA issues. Under these circumstances a variant of the above method is required.

Let the maximum shear stress between t_L and t_0 again be τ_{a_crit} . The relative amount of material at t_0 with strengths less than τ_{a_crit} can be unambiguously estimated using eq. 5.53

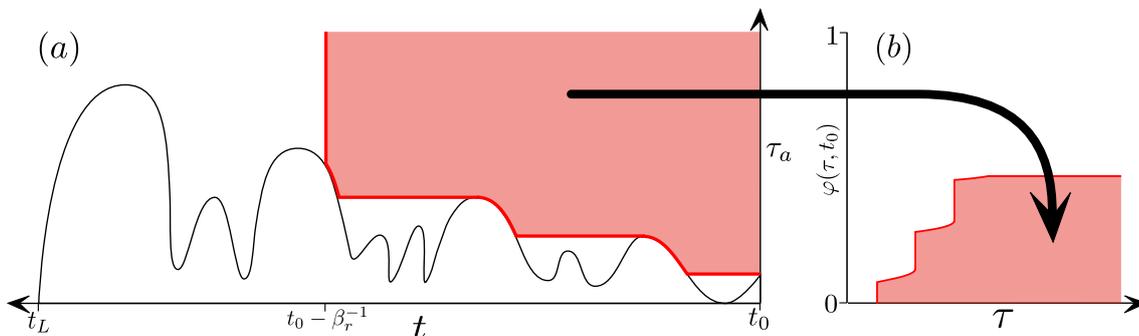


Figure 5.54: Illustration showing how one could calculate the wall state boundary condition vector, $\varphi(\tau_i, t_0)$, from a historical shear stress profile, τ_a , that is longer than the estimated time required for complete regeneration, β_r^{-1} .

(see the red line in fig. 5.55a). However, the amount of material at t_0 with strengths greater than τ_{a_crit} cannot be calculated as it could have been regenerating

1. since as long ago as $t_0 - \beta_r^{-1}$ (blue line in fig. 5.55a; $\varphi(\tau_i, t_0)$ is therefore maximal i.e. 1),
2. just since t_L if there was a large shear stress spike immediately before t_L (green line in fig. 5.55b; $\varphi(\tau_i, t_0)$ is therefore $(t_0 - t_L) / \beta_r^{-1}$) or
3. for a given strength τ_i it could have been regenerating since some time *between* $t_0 - \beta_r^{-1}$ and t_L . This time is 'unknowable', therefore one can only calculate upper and lower bounds for $\varphi(\tau_i, t_0)$ as per 1 and 2 above.

The more historical flow data that is available and the more the duration $t_0 - t_L$ approaches β_r^{-1} , the more similar the lower and upper bounds of $\varphi(\tau_i, t_0)$ are.

Note that it is not appropriate to refer to the upper and lower bounds of $\varphi(\tau_i, t_0)$ as worst and best cases as, after estimating $\varphi(\tau_i, t_0)$ and then fitting the parameters to data (at $t = t_c$ in fig. 5.52), what was thought of as the 'best case' for $\varphi(\tau_i, t_0)$ may be associated with a higher α value (output scaling factor) which may result in a greater turbidity response being predicted.

If applying the method prior to fitting β_r (using FIT_SIMULT) then there is a need to acknowledge the uncertainty in the estimation of β_r used here. This understanding of the expectation of / variance in β_r for a given system will be informed by β_r values found through the prior analysis/modelling of similar systems. Prior information on regeneration rates exists for DMAs but not for trunk mains, for which a high variance in β_r may need to be assumed. The conclusion from studying the effect of using the fastest and slowest estimates of β_r on the upper and lower bounds of $\varphi(\tau_i, t_0)$ for a trunk main model may be that the 'upper upper' bound and 'lower lower' bound may be sufficiently dissimilar for the method to have negligible benefit in estimating/bounding $\varphi(\tau_i, t_0)$.

Another unattractive quality of this method is that it requires the assumption that if a particular shear stress was reached at any point between t_L and t_0 then all material weaker than that strength was instantly eroded i.e. β_e is assumed to be effectively infinite prior to t_0 . This makes the method very sensitive to short-lived shear stress spikes that in reality only cause partial conditioning.

A method of yielding upper and lower estimates of boundary conditions is in theory valuable but it could prove difficult to couple it to the FIT_SIMULT parameter fitting method. Decisions

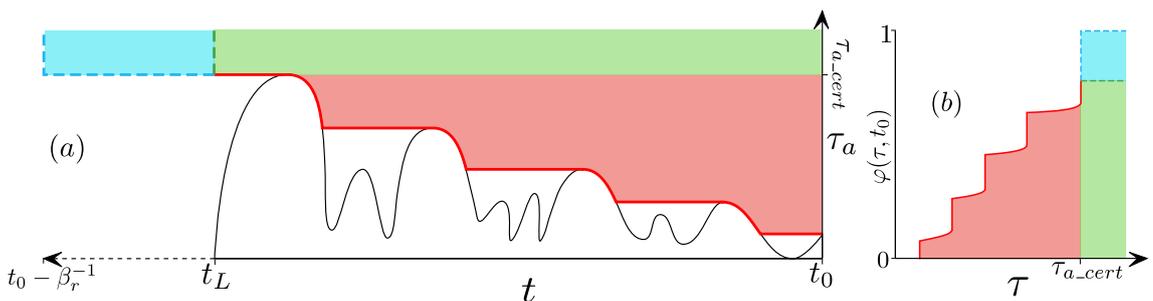


Figure 5.55: Illustration showing how one could estimate the wall state boundary condition vector $\varphi(\tau_i, t_0)$ from a historical shear profile, τ_a , that is shorter than the time required for complete regeneration, β_r^{-1} .

would need to be made regarding if/how to propagate the upper and lower boundary condition estimates through a PSO run and if/how the β_r PSO bounds should be informed by the fastest and slowest β_r values used to estimate the boundary conditions.

5.8.6 Estimate the wall state using VCDM conditioning

The wall state can more easily be estimated using historical time-series data and without the need for assumptions regarding erosion rates and regeneration rates by using the VCDM itself rather than an additional method.

First, identify a period over which to calibrate the model that starts with flows being particularly high. Next, set $\varphi(\tau_i, t_0)$ to trivially be 1 for all modelled shear strength bands. Start a VCDM simulation and allow the shear stress profile to erode material naturally i.e. erosion is not assumed to be instantaneous as β_e has a finite value. The initial turbidity response will be erroneous as material erodes but over time the model wall state will converge on the true wall state.

With this approach the wall state can be estimated and the parameters fitted simultaneously using FIT_SIMULT. This will work best if, within the PSO objective function, the timesteps corresponding to the most erroneous parts of the turbidity prediction (the initial conditioning) are ignored when calculating the correlation between predictions and observations.

The extents of the calibration period are important: if the flow is low at the start of the period but is high(er) shortly before then the modelled wall state could take months to converge to the true value. However, if the highest flow during the calibration period occurs early on then convergence would be much quicker and when calculating the correlation of observations and predictions within the PSO objective function one may simply need to ignore say five turnovers after this high flow period to allow the material that occurred during the wall state conditioning to pass out of the modelled pipe.

This approach to estimating the wall state is simple, requires no more assumptions regarding processes/parameter values than the VCDM (unlike that of §5.8.4 and §5.8.5) and neatly integrates with the FIT_SIMULT method. One disadvantage is that, like §5.8.4 and §5.8.5, it requires months if not years of historical flow data. Another is that it does not provide any measure of the uncertainty in the estimate of the wall state. A third is that knowing when to ignore turbidity responses when calculating the ISE may in cases where the flow profile is highly variable not be obvious. However, if the calibration period flow profile is of sufficient duration and begins with a large flow event then it may be the most attractive option.

5.8.7 Estimating the wall state boundary condition using generalisations/ observations of the data

Another simple approach is to use high-level knowledge of the flow history and or relationship between flow and turbidity to arrive at simple boundary conditions (simple as per the single-cusp shape described in §5.8.3). For example, if it is known that the flow has been constant for many months before and at the start of the calibration period then one might assume that relationship between the relative amount of material at the pipe wall and shear strength is

polarised around the average shear stress during the first few days of the calibration period. Alternatively, one may observe that a main only produced a turbidity response when the shear stress exceeded a particular level and so use that level to polarise the material quantity versus shear strength relationship.

This is a subjective, potentially overly-simplistic approach but it may be useful when fitting to esoteric datasets. A similar approach is often taken by modellers to decide on an initial τ_c value when calibrating [PODDS](#) models.

5.8.8 Discussion and summary

In this section several possible ways for fitting or otherwise estimating the wall state are presented. These vary in flexibility, complexity, their assumptions and the ease of integration with the chosen parameter fitting method; none are ideal for all situations and choosing one over the another depends on factors such as the size and timing of flow events within a calibration period, the quantity of available historic flow data and the accuracy of flow data.

The method where the initial amount of material at every modelled shear strength band is fitted using [PSO](#) (§5.8.2) thought to be too complicated to be practical (due to the increase in fitting problem dimensionality and the need for inter-dimensional fitting constraints). Estimating the wall state boundary conditions using the shear stress history (§5.8.4) and possibly an estimate of the regeneration rate (§5.8.5) requires the assumption that erosion is instantaneous, which is not believed to be justifiable. This leaves the methods where simple boundary conditions are fitted (§5.8.3) or estimated (§5.8.7) or alternatively the wall state is simply allowed to converge to an accurate solution over time ([VCDM](#) conditioning; see §5.8.6)

5.9 Determining the ‘net turbidity response’ within a pipe

When attempting to calibrate a [VCDM](#) model for a particular pipe, the simplest, ideal scenario is that downstream turbidity observations, T_{ds} , are solely a function of discolouration processes within that pipe, as the turbidity of water passing into that pipe is negligible. In practise, as will be seen in the next chapter, T_{ds} is often also shaped by T_{us} (previously mentioned in §4.3.6.4). The two most likely reasons for this are:

- Fluctuations in source water turbidity (e.g. due to treatment processes and interactions in reservoirs and at blending junctions) are of notable magnitude compared to turbidity response due to material erosion. The impact of this is illustrated using artificial data in [fig 5.56](#);
- Increases in flow may cause erosion in the pipe of interest but also in pipes further upstream. Note that this is less likely to be the case if:
 - The asset(s) upstream of the pipe of interest is a service reservoir or pipework of greater diameter or lesser roughness, so the an increase in flow only causes a significant increase in shear stress in the pipe of interest;

- Unidirectional flushing operations have been sequenced so that the pipes upstream of the pipe of interest are flushed first, ensuring a clear water front (Vreeburg and Boxall, 2007) when flushing that pipe;

If T_{us} is not negligible then one needs to discount T_{us} before or during the VCDM fitting process so that only the *net turbidity response*, T_{net} , is modelled by the VCDM. T_{us} can be discounted before the fitting process by removing its effect from the T_{ds} signal to yield T_{net} . The model is then fitted to the T_{net} signal. It is trivial to derive T_{net} by subtracting T_{us} lagged by the turnover time from T_{ds} if the bulk velocity, \bar{u} , is time-invariant (eq. 5.54). However, this is unlikely to be the case during turbidity events given that changes in shear stress drive erosion. The simple expression in eq. 5.54 cannot be directly calculated for cases where bulk velocity (and turnover time) is time-varying as the effect of T_{us} on $T_{ds}(t)$ is not just a function of the turnover time at time t but also the (changing) turnover times as material passes down the pipe of interest. A tractable approach is to supply T_{us} to a Lagrangian transport model to model advection to the end of the pipe then subtract this signal from T_{ds} to give T_{net} .

$$T_{net}(t) = T_{ds}(t) - T_{us}\left(t - \frac{L}{\bar{u}}\right) \quad (5.54)$$

Alternatively T_{us} can be supplied as a turbidity boundary condition *during* the model fitting process: it can be used as input to the VCDM's Lagrangian transport function (the `upstream_conc_v` parameter in `pyvcdm` (§4.5.3)).

Both approaches are obviously most accurate if T_{us} was directly observed using a second turbidimeter but when analysing historical datasets or in situations where a second turbidimeter was not available, could not be fitted or malfunctioned it may be necessary to estimate T_{us} using signal processing such as those described in §2.2.6.

5.10 Conclusions

The model formulations and implementation were verified. Demonstrations of model functionality confirmed that the model's formulations and software were a satisfactory implementation of the conceptual model. In particular, discolouration material was shown to regenerate at different shear strengths simultaneously, in keeping with observed behaviour.

The (numerical) evaluation of the proposed model requires the discretisation of both time and shear strength; the sensitivity to mesh resolution (Δt and $\Delta \tau$) was therefore explored both analytically and numerically. Several characteristic values of the model were identified through the analytical studies, such as the minimum timestep required for resolving the erosion of a shear strength (t_d) and the minimum timestep required to prevent a water parcel from passing through a pipe undetected (the turnover time, t_t).

$\Delta \tau$ was found to be analytically dependent on Δt but it can be considered independent if $\Delta \tau$ is notably less than the shear stress. Also, one effect of the chosen erosion model is that $\Delta \tau$ analytically influences the rate at which turbidity decays following an erosion event: in theory finer discretisation results in longer turbidity 'tails', but in practise this is unlikely to be relevant as such tails are anticipated to be dominated by the system SNR.

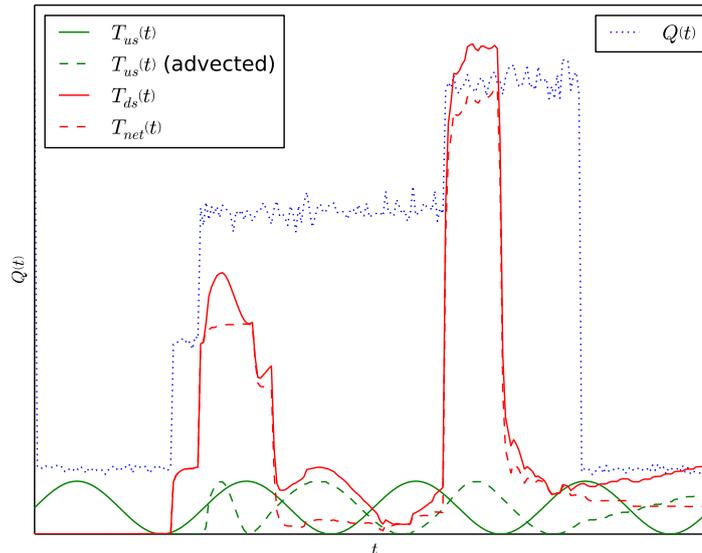


Figure 5.56: The turbidity response detected at the downstream end of a pipe of interest, $T_{ds}(t)$, is a function of the turbidity of water passing into that pipe, $T_{us}(t)$, after that material has been advected to the downstream end, and the latent turbidity signal due to the erosion of discolouration material from the pipe wall (the net turbidity response, $T_{net}(t)$). Here, $T_{us}(t)$ and, by extension, $T_{net}(t)$, were produced using the mass transport model coded as part of the reference software implementation of the [VCDM](#).

The sensitivity to the mesh resolution was investigated numerically with regards to four scalar measures of model output, resulting in a set of bounds for Δt and $\Delta \tau$ that minimise the effect of discretisation error due on the model outputs.

The sensitivity of the model to inputs and parameters was explored. The model may be non-trivial to fit to data and the fitting process may be indeterminate for certain model instances: sensitivity of the model parameters and boundary conditions showed that there are multiple ways to change the shape and scale of turbidity predictions (for erosion events) by changing these quantities. The sensitivity to the magnitude of flow changes, pipe length, internal pipe diameter and pipe roughness were also investigated; these also notably affected the shape and size of turbidity responses but this is less problematic for the purposes of model fitting as those quantities are easier to measure/estimate than the latent parameters and boundary conditions.

A method for fitting instances of the model to data was developed and verified under idealised conditions. Methods for fitting the three model parameters and boundary conditions were developed independently then combined. Several options were considered for fitting the parameters. The chosen 'FIT_SIMULT' option uses [PSO](#) to fit all three parameters simultaneously, the objective function here being the [ISE](#) of turbidity predictions and observations.

The performance of FIT_SIMULT and of alternative methods was explored by trying to 'recover' the parameters used to generate artificial turbidity predictions. The chosen method was no less suited to fitting the model to erosion events and regeneration periods than any other method but was simpler and more efficient than the alternatives. For certain model configurations,

multiple fitting attempts were required to achieve $R^2 > 0.95$ due to the stochastic nature of PSO; in others the fitting process appeared to contain insufficient information to recover the 'target' model parameters.

Methods for fitting or estimating the wall state boundary conditions. Several different methods were considered. These differ in their assumptions (regarding parameter values) and requirements (regarding the availability and variability of historical flow data). Three of these methods were considered to be suitable for developing a VCDM model of a real pipe.

Methods for isolating/fitting just the net turbidity signal In situations where the incoming turbidity is not negligible compared to the turbidity response due to erosion, the incoming turbidity signal must be discounted to prevent it from limiting the quality of model fit.

Chapter 6

Model validation

“Model validation is substantiating that the model, within its domain of applicability, behaves with satisfactory accuracy consistent with the [modelling and simulation] objectives. Model validation deals with building the *right* model.”

Balci (1997)

6.1 Introduction

In chapter 5, the proposed model was verified i.e. the [VCDM](#) formulations and reference implementation were shown to satisfy the conceptual specification set out in the earlier parts of chapter 4. However, this model is of little practical value unless it can be demonstrated that the assumptions that underpin the model are valid and that the model can be calibrated to reflect and predict reality with acceptable accuracy.

The first part of this chapter identifies the unvalidated assumptions made when developing the model that are most critical regarding its usefulness. The assumptions that could feasibly be tested were investigated under controlled laboratory conditions using a pipe rig of realistic proportions.

In the second part, the validity of the model as a whole is explored by calibrating model instances of several [DMA](#) and trunk mains using historical datasets and the calibration method developed in [§5.7](#) and [§5.8](#), the latter having been adapted to accommodate the non-ideal characteristics of real world data and systems. The accuracy of predictions following calibration is also tested.

6.2 Testing model assumptions

6.2.1 Introduction

Not all of the model assumptions listed in §4.6 could/needed to be tested. Those common to the proposed model and the field-validated [PODDS](#) erosion model or common approaches to [DWDS](#) hydraulic or water quality modelling did not need to be re-evaluated.

The following two assumptions, both specific to the proposed model, **were tested**:

1. **The model parameters are invariant with (absolute) shear strength.** The motivation for investigating this assumption was that the model would be far more complex if β_e , β_r and/or α were a function of shear strength (fig. 6.1): the dependency on shear strength would need to be parameterised, yielding more model variables, a greater parameter space and a model that is more difficult to calibrate and use. To test for differing erosion and regeneration characteristics between several shear strength bands, an experiment was conducted in a system of realistic scale but under laboratory conditions (for accurate flow/shear stress control over a multi-month period); this is covered subsequently in §6.2.2.
2. **The model parameters are invariant with time.** Again, model complexity would be far greater if β_e , β_r and/or α were found to be time-varying by a significant amount. However, testing for time-variance in discolouration processes under controlled laboratory conditions is non-trivial due to the required time-scales. For example, it is not practical to conduct a laboratory experiment once per quarter for one or two years to study seasonal variation in regeneration rates. Instead, the assumption has been implicitly tested by calibrating and validating real trunk main models over multi-year time periods (§6.3.4 and §6.3.5).

Several other assumptions were also specific to the proposed model but **testing them was not possible**, at least within the scope of this project:

1. **Holistically, weaker and stronger shear strength bands erode simultaneously.** This assumption, illustrated in fig. 4.16, is difficult to validate as different conceptual models of material erosion are only distinguishable under certain conditions such as partial erosion (see figs. 4.15 and 4.15) that will be difficult to establish in practise. However, it may not matter if the erosion model does not accurately reflect the true nature of underlying physical processes if it produces satisfactory results.
2. **The rate of material erosion is constant with time but not with shear strength.** As per the above.
3. **A quantum of wall-bound material has an invariant shear strength.** This is unlikely to be true but it cannot be disproved. It is not possible to monitor the strength of a quantum of discolouration material over time: one cannot observe the state of a quantum of material without altering it.

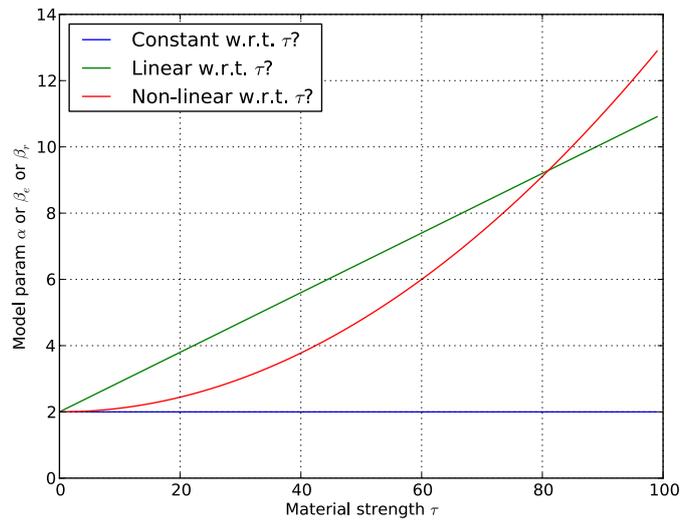


Figure 6.1: The assumption made in §4 was that the model parameters β_e , α and β_r do not vary with the shear strength of wall-bound discolouration material. Alternatively, they could vary linearly or non-linearly with shear strength, as illustrated here.

4. **Bulk water turbidity is *not* significantly reduced by material regeneration.** With perfectly accurate and calibrated turbidity instrumentation upstream and downstream of a length of main it would be possible to test this assumption by looking for a notable drop in turbidity over the main under invariant flow conditions (no erosion). However, in practise it would be difficult ensuring a constant enough flow to test this in a pipe long enough for any turbidity drop to be notable relative to signal/instrument noise.

The following assumptions were **also not tested** as they are common to the **VCDM** model and the validated **PODDS** erosion model:

1. Discolouration processes (and model parameters) are invariant with the angle to the pipe invert.
2. Discolouration processes (and model parameters) are invariant with longitudinal distance along a pipe.
3. Dispersion effects are insignificant.
4. Transient pressure waves and unsteady shear stresses do not impact on material erosion.
5. Material does not settle as sediment except under quiescent conditions.
6. Turbidity is proportional to **TSS**

6.2.2 Experimentally verifying the assumption that model parameters are invariant with shear strength¹

6.2.2.1 Theory

If/how the model parameters vary with shear strength can be tested by conducting operations similar to those described in the demonstration of **VCDM** capabilities presented in §5.2.5.

¹This section contains text that has been quoted and adapted from [Furnass et al. \(2014a\)](#).

Take a pipe where the diameter, roughness, bulk velocity and flow are invariant along its length and its internal surfaces are entirely devoid of discolouration material and biofilm. If this is supplied with water of a consistent quality and experiences low shear stresses for several months, then discolouration material will accumulate over a wide range of shear strengths.

Following this 'growth phase', the variation in the characteristics of discolouration material with shear strength can be studied by increasing the shear stress in steps and studying the turbidity response per step. The model formulations ensure that both the rate and amount of erosion are functions of the *excess shear stress*; therefore to study whether parameters vary with absolute shear strength one needs to increase the shear stress during the experiment in steps of equal magnitude to control for excess shear stress (i.e. each step gives the same excess shear stress).

During this 'flushing phase', let each shear stress step be sustained for a sufficient number of turnovers for all affected shear strengths to be fully depleted and all mobilised material discharged from the pipe (see fig. 6.2). Given these conditions the total amount of material mobilised per step can be quantified as the observed material flux integrated over the duration of each flow step (see §2.1.13.2, §5.2.5 and eq. 5.3; assuming that the background turbidity is negligible).

Here the amount of material mobilised per step is a function of how long the corresponding shear strength band had been regenerating for when it was eroded (and depleted), the relative rate of material accumulation (encoded as β_r) and the maximum material capacity and light scattering properties of that shear strength band (encoded as α). If the quantity of mobilised material is invariant between shear stress steps and all corresponding shear strengths have been accumulating for the same duration then *both* β_r and α are therefore invariant with shear strength. Note that flushing phase is assumed to be sufficiently short relative to the growth phase that each of the shear strength bands mobilised during the flushing phase has had the same amount of time to regenerate when they are successively eroded. Another assumption is that discolouration processes are longitudinally homogeneous along a pipe length of consistent diameter, roughness and pipe material.

Quantifying the amount of material mobilised per shear stress step as the integral of material flux is preferable to quantifying it using the peak turbidity (eq. 5.5) per step:

- The latter is not independent of the erosion rate factor, β_e , as that parameter also influences the scale of the turbidity response (§5.4) yet we wish to separately test that β_e is invariant with shear strength;
- Material flux integrated over time will be less sensitive to flow and turbidity measurement error.

It is not possible *using this test* to distinguish the effects of β_r and α ; one can only say both or neither are invariant with shear strength. A lengthier experiment involving multiple successive regeneration periods of differing durations (e.g. fig. 5.31) would be required to differentiate them; this would be of interest if it could be demonstrated that at least one of the two parameters varies with shear strength.

To test for erosion rate factor β_e being invariant with shear strength the turbidity responses from all steps should follow the same upward curve to the turnover time (fig. 5.11). This is

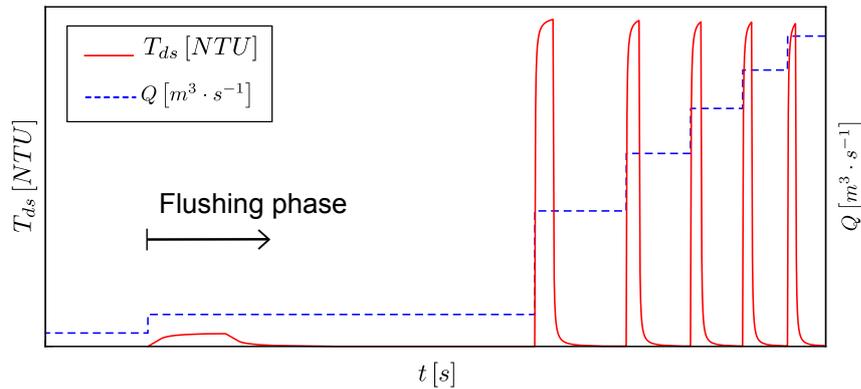


Figure 6.2: Under controlled conditions, increasing the shear stress in equal increments following a lengthy growth phase allows the variation in erosion and regeneration processes with absolute shear stress to be explored. This requires uneven flow increments due to the non-linear relationship between flow and shear stress. The step increase in shear stress at the very start of the simulated flushing phase shown here is somewhat sacrificial: it sets a minimum shear strength that is just above the maximum shear stress of the preceding growth phase, reducing the uncertainty in the state of the weakest material that results from slight variations in shear stress over the growth phase.

only true if all steps mobilise the same amount of material, so β_r and α must already been shown to be invariant with shear strength before this test can be performed.

Let us define a null hypothesis and two alternative hypotheses:

H_0 All three model parameters, β_r , α and β_e , are constant with shear strength.

H_{A1} The parameters β_r and α (the perceived amount of accumulated material after a given time) are constant with shear strength but β_e is not.

H_{A2} One or both of β_r and α vary with shear strength and it cannot be determined whether β_e does.

6.2.2.2 Experiment

These hypotheses were tested under controlled laboratory conditions but using two independent pipe systems of a scale partially representative of DWDS pipework. There are several reasons why it is advantageous to conduct such an experiment using a dedicated experimental pipe rig rather than in the field:

- Greater control of variations in water quality including temperature (which can have a domineering impact on regeneration rates; §2.1.16.2);
- Accurate, high-resolution, monitoring of turbidity and flow (particularly if there are no branches);
- Accurate, fine-grained flow control, which is important when moving between higher flows due to the non-linear relationship with shear stress (see fig. 2.7). Also, flow can be controlled over lengthy periods as it is not demand-driven;
- The ability to erode discolouration material and increase bulk water turbidity without there being a risk of affecting customers.

The pipe systems, A and B, were each structured as per fig. 6.3 and table 6.1, with the length of HDPE pipe per system being sufficient so that discolouration processes are dominated by what occurs at the pipe wall and not erosion and accumulation within the tank and around hydraulic devices such as pumps and valves. Pipe diameters were typical for DMA mains in the UK, ensuring that system hydraulics are largely representative of distribution system mains. The pipe system configurations were similar to those used by Sharpe (2013) and Fish (2014) (previously referenced in §2.1.13.4), the main differences being the type of turbidity monitoring hardware and the monitoring/control software.

Table 6.1: Dimensions of pipe systems A and B

System	Capacity [m^3]		Distance [m]				
	Tank	Tank & pipes	i-ii	ii-iii	iii-iv	iv-v	v-vi
A	0.486	1.527	8.46	181.76	0.31	9.30	16.50
B	0.486	1.493	9.10	181.76	0.31	9.30	9.00

The internal surfaces of the pipe systems were initially cleaned by dosing with sodium hypochlorite to $20 \text{ mg} \cdot \text{L}^{-1}$ then pumping at the maximum attainable flow rate ($7 \text{ L} \cdot \text{s}^{-1}$) for 24 hours. During the subsequent growth phase water partially recirculated within each system but the outflow to waste from each tank and the inflow from supply were controlled to give a 24-hour hydraulic residence time. This prevented stagnation and provided a constant renewal of free chlorine and a background concentration of inorganic and organic matter. Supply was via an unlined CI trunk main from a works that treats peaty moorland run-off using iron coagulant. The room temperature was controlled at 16°C as fluctuations in temperature can have a dominant effect on accumulation rates (Sharpe, 2013).

System monitoring, control and data logging were all automated. Proportional integration (PI) controllers, implemented using Labview (Elliott et al., 2007) software and hardware, varied the pump speed to maintain 20 m of hydrostatic pressure at monitoring point P_2 . The same system continually monitored flow using Flownetix ultrasonic meters and manipulated an actuated valve so as to follow user-specified flow profiles.

During the growth phase (85 and 86 days in pipe systems A and B respectively) the diurnal hydraulic profiles shown in fig. 6.4 were imposed. These profiles differed for the purposes of a coincident experiment that is not discussed here. This phase was followed by a flushing

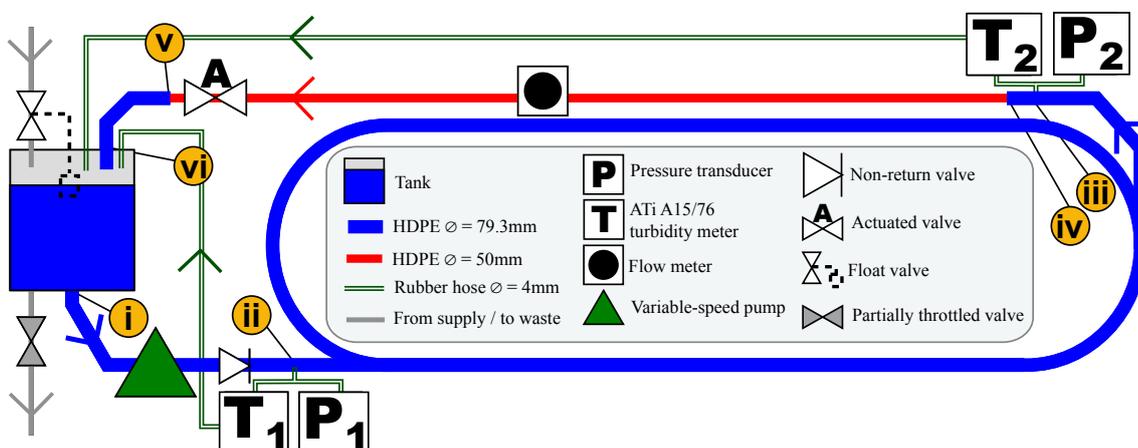


Figure 6.3: Schematic of each pipe system.

phase where the shear stress was increased in both systems in steps of $1.11 Pa$ up to the maximum flow as per table 6.2. Each step was sustained for five turnovers to ensure complete erosion over the affected shear strength range. The shear stress during Step 0 was just higher than the shear stress experienced in either system during the growth phase, thus ensuring a consistent minimum shear strength across both systems at the start of the flushing phase. Flow and pressure set-points and measurements along with turbidity readings were logged every $5 s$ during this phase (and every 5 minutes during the growth phase).

Table 6.2: Flow rates and corresponding shear stresses imposed during the flushing phase.

Flow step	0	1	2	3	4	5
Flow $Q [L \cdot s^{-1}]$	0.72	3.06	4.35	5.37	6.23	7.00
Shear stress $\tau_a [Pa]$	0.09	1.12	2.31	3.42	4.54	5.65

Shear stress, τ_a , was determined using a roughness ($0.075 mm$) previously calculated by Husband et al. (2008) for the same pipe diameter and material plus eqs. 2.9, 2.10, 2.11 and 2.12. For these calculations, the pipe diameter was assumed to be a constant $79.3 mm$ as the length of the section where it was $50 mm$ (see table 6.1) was assumed to be negligible w.r.t. the total length of pipework. Also, this narrower section lay downstream of turbidity meter T_2 , so the shear stress only changed (increased) after the second turbidity monitoring point, before water returned to and was mixed within the tank.

During each system's flushing phase, turbidity was monitored at locations T_1 and T_2 . Turbidity meter optics were cleaned at the end of the growth phase. The turbidity within each system increased over the flushing phase due to the recirculation of mobilised material, making H_0 more difficult to test. The intention had been to measure the net response between T_1 and T_2 and discount for recirculation effects by modelling the advection of the signal detected at T_1 to the position of T_2 then subtracting this from the signal recorded at T_2 ; however this was not possible as T_1 in both systems failed and so only the T_2 signal could be analysed.

Each system was hydraulically isolated from mains supply during its flushing phase; because of this, the outflows from turbidity meters T_1 and T_2 were returned to the relevant tank rather than allowed to go to waste so as to prevent tank depletion and pump trips. The impact of these flows on the tank turbidities has been ignored in the following data analysis due to the negligible magnitude of the flows compared to those in the $\varnothing = 79.3 mm$ pipes ($0.85 L \cdot min^{-1}$; found by solving the Bernoulli and Darcy Weisbach equations for flow). The signal propagation delay due to flow from the monitored pipe section to the turbidity meters is said to be negligible

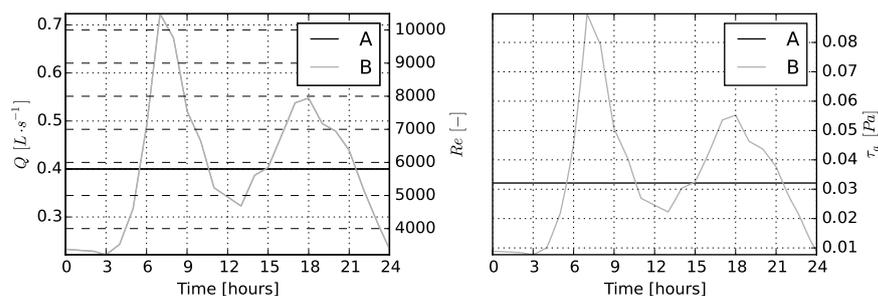


Figure 6.4: Flow, Reynolds number and shear stress diurnal profiles imposed over the 85 to 86-day growth phase.

for the same reason.

The control system increased the pump speed quickly but asymptotically at the start of each flow step to lessen the risk of transient shear stresses impacting on the wall-bound discolouration material. The rate of increase was not constant with excess shear stress but was a function of measured flow. The method for increasing flows that was programmed into the LabView control software was to allow 15 s per $1 L \cdot s^{-1}$ increase with a minimum of one minute for the operation. There was therefore concern that the characteristics of this control system could limit the ability to study the relationship between the erosion rate and shear stress, particularly at higher flow rates where the time taken to reach the set point was a greater proportion of the turnover time. However, the effects of the way in which flows were increased between steps was found to be negligible (discussed subsequently in §6.2.2.7).

6.2.2.3 Tests adapted for a recirculating system with only one turbidity meter

To test H_0 , there was a need to distinguish the material mobilised by each τ_a step. This was done by iteratively subtracting the mean of $T_2(t)$ over the *final turnover* of a particular step from $T_2(t)$ over the *entire* following step (whilst ensuring the result was non-negative). Let this derived signal be $T(t)$. Although this data pre-processing partially accounted for recirculating by removing the 'background' turbidity from each step, it did not account for signal obscuration due to the recirculation of mobilised material *during* a single step, which will be seen following one turnover time into each step. The tests outlined in §6.2.2.1 therefore cannot be applied without modification, as material recirculation affects the integral of the product of $T(t)$ and flow over each step (and therefore the perception of material accumulation) and the upward curve of the $T(t)$ spike per step (and the perception of material erosion). Those theoretical tests can however be adapted to accommodate these constraints:

- The perceived accumulation (and β_r and α) are constant with shear strength if the integral of the material flux over the *last turnover* per step is constant (assuming that after five turnovers erosion has reached steady state and the system is well-mixed).
- The erosion rate is constant with shear strength if the perceived accumulation is constant with and the curvature of $T(t)$ is the same for the *first turnover* of each step (before unaccountable-for recirculation effects are first seen).

6.2.2.4 Data analysis results

Fig. 6.5 shows the flow and turbidity per system during the flushing phase along with the derived turbidity signal $T(t)$ for which the 'background' turbidity due to recirculation has been partially removed. An erosion response per τ_a step can be seen in $T(t)$; no upper limit to τ (i.e. τ_{max}) was found. Note that the drop in turbidity after Step 2 is due to 6-9 minutes of data being cut from the flow and turbidity time-series: the pumps were stopped over this time for microbial sampling (for a coincident experiment and not discussed further here).

The turbidity signal $T(t)$ over the last turnover per τ_a step (fig. 6.6) indicates that the same amount of material accumulated per unit τ during the growth phase over most of the studied

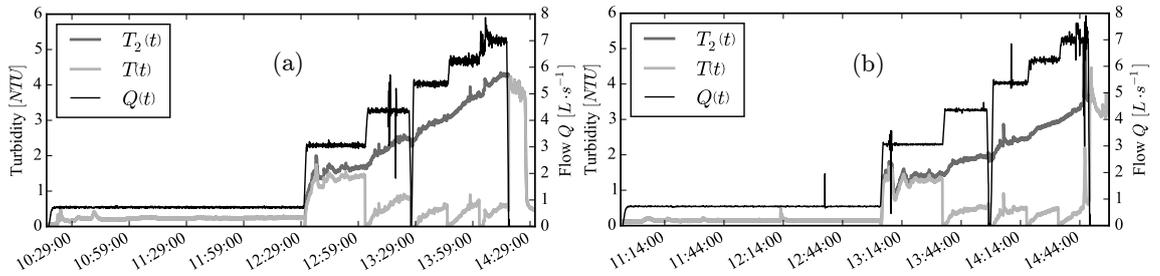


Figure 6.5: Flow $Q(t)$ and measured turbidity $T_2(t)$ over the flushing phases of systems A (subplot a) and B (b). $T(t)$ is the turbidity per τ_a step with the background turbidity of the previous step subtracted.

strength range (Steps 3-5 in system A; Steps 2-5 in system B). However, proportionally more material accumulated at weaker strengths. Here the amount of material mobilised per step, $A_w \cdot N$, was calculated as the time-integral of material flux and is expressed in TPMU (eq. 4.2). The assumption that erosion would have reached a steady state and the system be well mixed before the fifth and final turnover per τ_a step appears to be valid: there is little if any rise in turbidity over the fifth turnover. Note that the drop in turbidity during Step 5 in system B was due to a momentary pump trip during the previous turnover. Perceived accumulation appears to be constant with τ for all but the weakest strengths, allowing the relationship between the erosion rate factor (β_e) and absolute τ to subsequently be assessed for all but the weaker end of the studied shear strength range.

Fig. 6.7 shows $T(t)$ for the first turnover of each τ_a step bar Step 1. The SNR prevents a conclusion from being drawn as to whether the curvature of all responses is the same and the erosion rate (and therefore β_e) is constant with shear strength. The SNR is not surprising: the ATi A16/76 turbidity meter has an accuracy of the greater of $\pm 5\%$ and ± 0.02 NTU when the output range is set to 0 – 40 NTU. Systems A and B both show a slower rise in $T(t)$ over the first turnover of Step 3; the reason for this is not known.

6.2.2.5 Data analysis discussion

The perceived material accumulation was not found to be constant with shear strength over the entire studied τ range in pipe systems A and B; H_0 has therefore been rejected in favour of H_{A2} . The implication of this is that β_r and α are not constant with τ and the variance of β_e with τ cannot be conclusively decided. However, the results also indicate that all three parameters are constant with τ for all but the weakest-bound material.

It is not possible using the data from this experiment to determine whether it was β_r and/or α varying with τ that caused more material to be detected at the lower end of the studied strength range. These effects could be distinguished in future by repeating the experiment with differing growth phase durations (somehow controlling for seasonal variation in source water quality) and/or using particle counting equipment to study if and how particle diameters and therefore light-scattering properties vary with τ .

There is a question of whether the amount of perceived material is greater for shear strengths above the maximum conditioning shear strength of the growth phase, or whether the perceived amount of material is greater over an *absolute* shear strength range. The two pipe systems

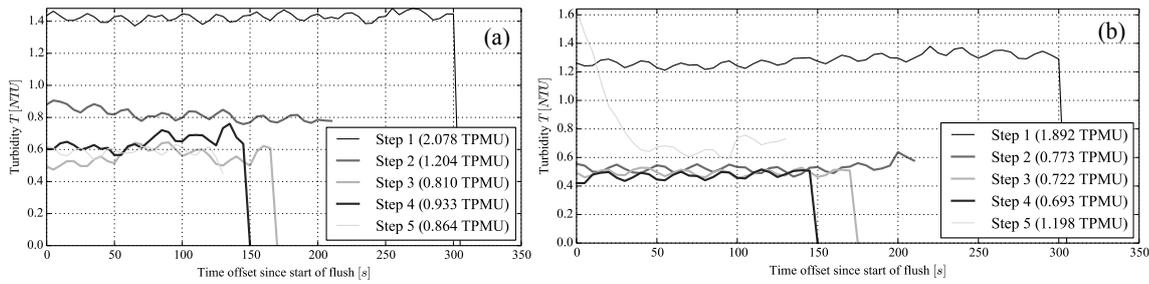


Figure 6.6: Processed turbidity $T(t)$ for the last turnover at each τ_α level, which is a measure of the amount of material per discrete τ range. (a) and (b): pipe systems A and B respectively. Figures in parentheses are the integral of material flux over the last turnover of each flow step (after removing the ‘background’ turbidity from previous flow steps).

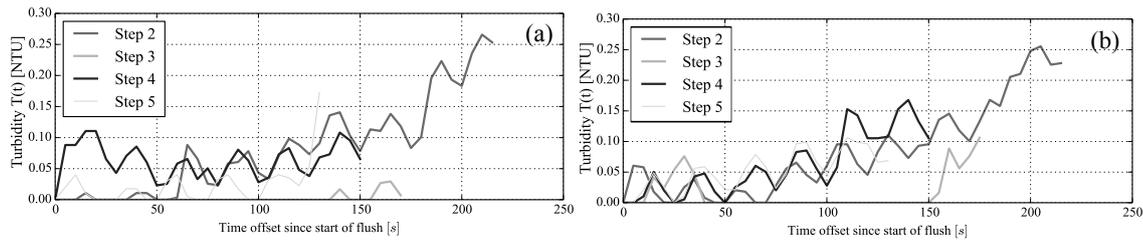


Figure 6.7: Processed turbidity $T(t)$ for the first turnover at each τ_α level bar Step 1. The curvature of $T(t)$ is a measure of how the erosion rate and amount of accumulated material vary with absolute τ . (a) and (b): pipe systems A and B respectively.

experienced different daily maximum shear stresses over the growth phase yet exhibited similar responses during the flushing phase, suggesting that distinct turbidity responses for Step 1 were due to there being more material over an absolute shear strength range. Overall, more material was mobilised in system A, which could be due to differences in growth or flushing conditions (e.g. different nutrient fluxes and/or growth phase shear stress profiles) but the differences were slight.

6.2.2.6 Modelling method

In addition to testing the hypotheses listed in §6.2.2.1 using an entirely data-driven approach, the flushing phase experimental results were also modelled numerically using the developed **VCDM** implementation. Modelling the experiment has the advantage over the data-driven analysis of the **VCDM** implementation of being able to account for recirculation effects by modelling advection and mixing. This negates the need for ad-hoc methods for removing the background turbidity at the start and during each flow step. However, the modelling approach has a number of disadvantages compared to the data-driven analysis: the model requires the calibration of parameters and assumptions need to be made with regards to boundary conditions and tank mixing models.

The implemented material transport method required some adaptation to allow it to model a recirculating system featuring a tank. For each simulation timestep other than the first:

1. React all water parcels in the pipe length comprised of sections iii to iv, iv to v and v to vi (see fig. 6.8); let this pipe length be known as `SENSOR_TO_TANK`;
2. Determine the turbidity at the tank inlet;

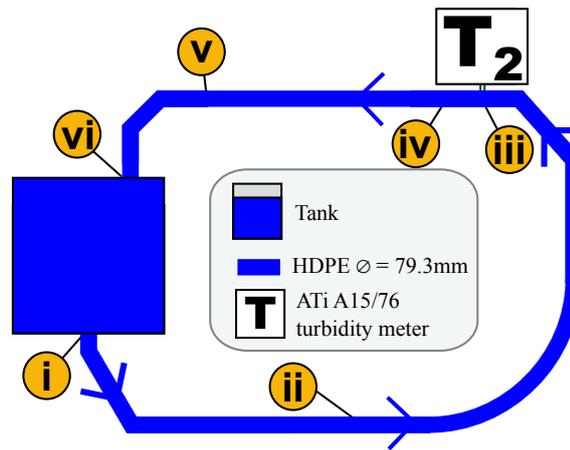


Figure 6.8: The schematic of each pipe system as modelled. Roman numerals correspond to the locations shown in fig. 6.3.

3. Calculate a new tank turbidity, assuming complete, instantaneous mixing and a full tank;
4. React all water parcels downstream of the tank and upstream of the turbidity sensor (pipe sections i to ii and ii to iii in fig. 6.8);
5. React all water parcels in the pipe length comprised of sections i to ii and ii to iii (see fig. 6.8); let this pipe length be known as TANK_TO_SENSOR;
6. Advect material from the tank into the TANK_TO_SENSOR pipe length.
7. Advect material out of the TANK_TO_SENSOR pipe length and calculate a nodal turbidity at the turbidity meter tapping point.
8. Determine the material that has passed in to SENSOR_TO_TANK from TANK_TO_SENSOR.

The following assumptions were made when fitting the VCDM to the flushing phase data from pipe systems A and B:

- Regeneration was negligible;
- Tank mixing was instantaneous and negligible;
- The change in diameter downstream of the turbidity sensor tapping point was not important and so for ease of modelling all pipework could be treated as having an internal diameter of 79.3 mm for the reasons given previously in §6.2.2.2.
- The shear strength range of interest, $[\tau_{min}, \tau_{max}]$, was from the lower to upper bound of the modelled flushing phase shear stress profile. A turbidity response is seen from every step increase in shear stress so there was no evidence of a 'cleaning' shear stress (see §2.1.15) and so no justification for setting τ_{max} to any less than the maximum shear stress. The range $[\tau_{min}, \tau_{max}]$ was discretely modelled as 1000 shear strength bands.
- The relative amount of material at the start of the flushing phase was (at least initially) assumed to be maximal (i.e. 1) over the modelled shear strength range. Flow step 0 then served the purpose of conditioning the wall state (like the VCDM wall state conditioning

method described in §5.8.6). Any shear strength initially being maximally regenerated was unlikely given that full regeneration is estimated to take 1.5 to 4 years (§2.1.14) yet the duration of the experiment was approximately 90 days. However, this inaccuracy was not relevant to the testing of hypothesis H_0 ; what was important was that at the start of the flushing phase all shear strengths greater than the shear stress of step 0 were considered to have the same relative amount of material as they had been regenerating for the same duration during the growth phase and there was insufficient evidence to say that material regeneration varies with shear strength (§2.1.13.4).

- The background turbidity at the start of the flushing phase was 0.06 NTU , which was the average turbidity during the growth phase.
- The time period during which there was no flow for several minutes (for biofilm sampling; see §6.2.2.4) was removed from the modelling input time-series to simplify the modelling process as the turbidity observations during this period were erroneous (due to settling effects).

Both the ideal (stepped) and measured (asymptotically increasing) flow profiles were used to generate **VCDM** predictions to assess if/how non-instantaneous flow increases and small, momentary flow fluctuations influenced predicted turbidities.

The **VCDM** was fitted to data iteratively by hand rather than by using **PSO**.

6.2.2.7 Modelling results

No values of β_e and α could be found that gave an accurate turbidity prediction over all shear stress steps: it was only possible to achieve good fits to the earlier (fig. 6.9) or later (fig. 6.10) part of the turbidity profile for each pipe system. As can be seen from these figures, two sets of flow profiles were used to generate **VCDM** predictions: Q_{ideal} is the desired (set point) flow whereas Q_{high_sens} is the measured flow. The differences between the two sets of predictions were small, even though the measured flow asymptotically approaches the set point at the start of each flow step and the measured flow includes several rapid fluctuations due to pump and control system instabilities.

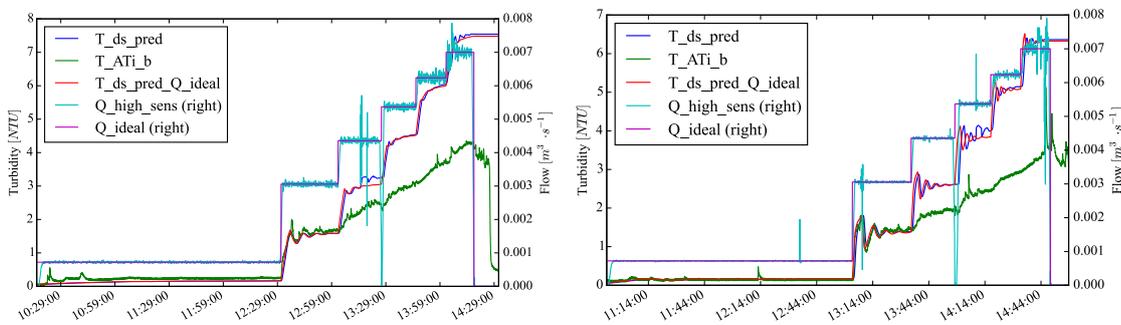


Figure 6.9: The **VCDM** was fitted to just earlier turbidity observations from pipe systems A (left subplot; $\beta_e = 0.005$; $\alpha = 0.038$) and B (right subplot; $\beta_e = 0.015$; $\alpha = 0.032$). Here T_{ATi_b} is the series of turbidity observations (T_2). For each pipe system the model was first fitted using the ideal flow profile (Q_{ideal} , yielding $T_{ds_pred_Q_ideal}$) and then again using the measured flow (Q_{high_sens} , yielding T_{ds_pred}).

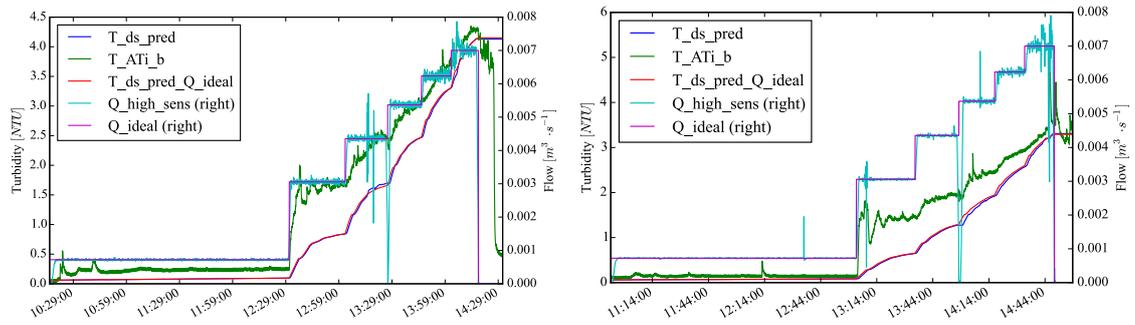


Figure 6.10: Similar to fig. 6.9 but here the *VCDM* was fitted to just later turbidity observations from pipe systems A (left subplot; $\beta_e = 0.0014$; $\alpha = 0.022$) and B (right subplot; $\beta_e = 0.001$; $\alpha = 0.018$).

The data-driven analysis of §6.2.2.4 indicated that the material at weaker strengths (material eroded by Step 1) has different characteristics. If the turbidity responses from Steps 0 and 1 are removed from the turbidity profile that is supplied to the model then far better model fits can be achieved. (fig. 6.11).

Here background turbidity response was removed by

1. Calculating the mean turbidity over the last turnover of Step 1;
2. Subtracting this value from the entire turbidity profile
3. Clipping all turbidities less than the ‘background turbidity’ (0.06 *NTU*) to that value.

Also, the boundary conditions were changed so that all shear strengths less than the (ideal) shear stress during Step 2 were initially depleted; this ensured that the model did not predict a turbidity response for Steps 0 and 1.

6.2.2.8 Conclusions and implications

Both sets of results indicated that the characteristics of discolouration material are not entirely consistent with shear strength, despite the entirely data-driven analysis and semi-empirical modelling differing in how rigorously recirculation effects were accounted for and in data analysis relying on fewer assumptions regarding discolouration processes. More material appears to be

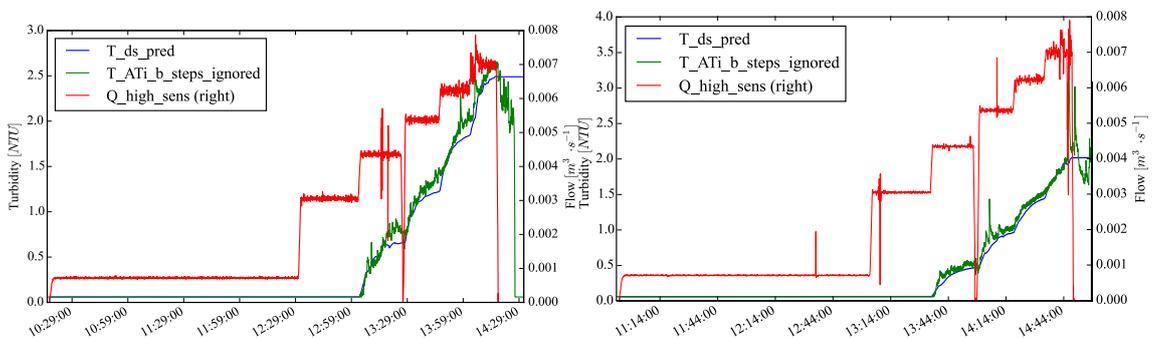


Figure 6.11: Similar to figs. 6.9 and 6.10 but here the *VCDM* was fitted to turbidity observations from pipe systems A (left subplot; $\beta_e = 0.0025$; $\alpha = 0.016$) and B (right subplot; $\beta_e = 0.0017$; $\alpha = 0.013$) after discounting the turbidity responses from Steps 0 and 1. These predictions were generated using just the measured flow profile (not the ideal (set-point) flow profile).

present at the weakest shear strengths following a period of regeneration. It is not possible to say with confidence whether the erosion rate also differs for weaker shear strengths.

An implication of these findings is that a **VCDM** model that is calibrated using shear-strength-invariant parameters and time-series data from a period where the shear stress varied widely may overestimate the amount of material present at the weakest shear strengths. Alternatively, the amount of material at all but the weakest strengths may be underestimated if the model is calibrated by imposing only small increases in shear stress. It is worth noting that the results indicate that the assumption that $\varphi(\tau, t)$ is a monotonically increasing function (§4.4.4) may not always be valid (at least for the studied pipe rig).

It is not possible to directly compare the experimental results presented here with those from previous stepped flushing studies conducted under laboratory conditions. Although there is a strong argument for treating turbidity as a concentration (§4.3.6.2) and therefore quantifying the amount of material mobilised by integrating material flux over the duration per flow step (to give a pseudo-mass; §2.1.13.2), previous studies (Husband et al., 2008; Sharpe, 2013) quantified material mobilised per step solely as a function of turbidity (a pseudo-concentration).

It should be noted that these findings pertain to two plastic pipes of the same diameter, roughness and material; the experiment should be repeated for a variety of pipes and environmental conditions to confirm whether the findings are universal. Also further experiments comprised of multiple regeneration periods of different durations could allow for the effects of β_r and α to be distinguished and so determine whether one or both vary with shear strength. In addition, the experimental setup could be improved by (a) using two working turbidity meters (as per fig. 6.3) for easier and more accurate recirculation effect compensation and (b) using two working pressure transducers per pipe system to calculate shear stress directly from measured headloss (rather than by using a previously calculated roughness).

Experimental work and **PODDS** modelling has indicated that discolouration material (particularly in plastic pipes) has a maximum strength (which may or may not be pipe/environment-specific; see the references to $\tau_{ultimate}$ in §2.1.15). Given that these studies have not explored the relationship between shear strength and material quantity with much granularity, it may be the case that the results of those studies support and are explained by the results of this experiment e.g. the **PODDS** modelling in Husband and Boxall (2010a) could be interpreted as indicating that the distribution of discolouration material prior to flushing was weighted towards the weaker end of the modelled shear strength range.

The fitted parameter values presented in §6.2.2.7 differed between pipe systems A and B, particularly in the case of β_e (see figs 6.9, 6.10 and 6.11). The reason for this is not known, but it may have been difficult to fit β_e (a rate factor parameter) consistently due to recirculation effects (and the data pre-processing steps used to compensate for those effects) causing turbidity to appear to rise more or less rapidly than it did in reality. Alternatively, fits may have been sub-optimal and inaccurate as in this case they were found by hand rather than using **PSO**. A third possibility is that the growth phase shear stress profile influencing how quickly material with much higher shear strengths is eroded given a unit excess shear stress. If this hypothesis could be proven, it would be in conflict with the **VCDM** assumption that variations in shear stress at lower levels do not influence the behaviour/state of discolouration material properties with higher shear strengths (so long as there is no excess shear stress).

The parameter values of the fitted model instances presented in §6.2.2.7 should not be considered accurate: the relative amount of material per shear strength at the start of the flushing phase was erroneously assumed to be maximal for simplicity (it is unlikely full regeneration occurred in three months), so the fitted α parameter values are overly small to compensate for this.

Further work is required to determine whether the identified difference in the accumulation process over different τ bands is actually *significant* for managing discolouration risk and proactive network management: if it is not then accumulation could be assumed to be invariant with τ , so as not to increase the complexity of the parameters of the discolouration model.

6.3 Model validation using field datasets

6.3.1 Requirements and potential issues when calibrating models of real mains

In §6.2.2.6 and §6.2.2.7 the developed model was shown to largely be capable of representing discolouration processes in a particular pipe rig. However, this does not constitute validation as:

- It was a single system and is not representative of all source waters, diameters, pipe materials, roughnesses etc.
- It was not possible to accurately calibrate model parameters (§6.2.2.8) β_r and α could not be distinguished nor could β_r be calculated.
- At the end of the growth period material on the walls of the pipe loops was only partly regenerated.

It is therefore necessary to show that the model can represent discolouration processes in a variety of real DWDS pipes.

Calibrating models of real DWDS pipes is typically more complicated than exploring model behaviour using artificial data (as was done in §5). Some of the key issues are as follows:

- The modelled main(s) may not be homogeneous with respect to diameter, pipe material, roughness and flow, particularly if the main is several kilometres in length. Ideally, turbidity should be monitored at locations where there is a known change in pipe properties to provide boundary conditions that allow the model to be fitted to pipe sections with longitudinally homogeneous properties.
- Ideally turbidity will have been continuously sampled at the upstream and downstream ends of the modelled main(s) over a period of interest. However, if only the downstream turbidity could be / has been sampled then signal processing techniques (§2.2.6) may be needed to distinguish the turbidity response due to erosion within the main from the signal due to imported material.
- To model both shear-stress-driven erosion and advection one requires a hydraulic model of the mains(s) of interest. The relative roughness of a hydraulic model may accurately predict headlosses but the absolute roughness and diameter may be inaccurate (§2.2.7.4),

resulting in erroneous estimations of water quality travel times. Ideally one would calibrate DWDS models using both flow and velocity data but measuring both requires the use of tracers, which is problematic for network operatives so is rarely done in practice. However, occasionally temporal offsets between observed and predicted turbidity spikes can help determine how much usable internal diameter has been lost since a (corrodible) pipe was installed or last lined and/or since a hydraulic model was last calibrated.

- Given that full regeneration is estimated to take at least 1.5 years (§2.1.14) it is desirable to calibrate a discolouration model using at least a couple of months worth of flow and turbidity data (ideally at least two years worth of data if one wishes to investigate whether seasonal/temperature effects influence regeneration rates). However, the more data that is captured, the more likely the calibration dataset is to have been affected by issues such as instrument drift, error or failure. Data pre-processing is required to remove/interpolate over periods of bad data (e.g. outliers, series of suspiciously constant values, negative values). Also signal processing techniques (§2.2.6) may be required to account for changes in the turbidity baseline that are due to instrument drift.
- To confidently calibrate a model of material erosion and regeneration one needs a dataset containing a number of shear stress events and for those events to have resulted from turbidity responses significantly greater than the background turbidity fluctuations (i.e. a high SNR is required). For a time-series dataset of several years in duration a shear stress event detection method may be useful for determining the periods over which it is particularly important that observations and model predictions tally.
- The temporal resolution of flow and turbidity data should be high relative to the turnover time (§5.3.9.7) (and ideally t_d , but this cannot be determined in advance of calibration due to t_d being dependent on the β_e parameter).
- Time-series inputs to the model may need to be lagged to account for inaccurate time-stamps and may need to be re-sampled to ensure all time-series that comprise a calibration dataset share a single, consistent time-series index (although the model allows for a time-varying Δt).

Strategies for addressing these are presented in the following model calibration case studies. The first of these studies validates the proposed erosion model using data from the flushing of four DMA mains. The validity of the regeneration model is then explored using long-term monitoring data from three trunk mains.

6.3.2 Validating the VCDM erosion model by fitting to DMA flushing data

6.3.2.1 Objectives

The proposed model of material erosion differs from the validated PODDS erosion mechanism so it must be independently validated. This is easiest to achieve using field data obtained during the flushing of previously undisturbed mains:

- The relative material quantity at strengths greater than a threshold such as the estimated daily maximum can be assumed to be maximal and the relative material quantity at strengths less than that threshold can be assumed to be minimal. The boundary conditions of the wall state sub-model are therefore trivial to define.
- The duration of mains flushes are sufficiently short relative to the time required for full regeneration (particularly for DMA mains) that regeneration can be assumed to be negligible over duration of a flush. The β_r parameter can therefore be set to $0 s^{-1}$ in any model of a single flush, simplifying the fitting process.

By 'previously undisturbed' it is meant that prior to flushing a main experienced a regular (most likely daily) flow pattern over the course of many months (ideally years). Husband and Boxall (2010a) presented PODDS model calibrations for four such flushes and quantified the fit using a relative correlation metric.

This section explores the accuracy and precision (repeatability) with which the VCDM can be fitted to several of the isolated flushing event datasets presented in Husband and Boxall (2010a) and compares the quality of model fit (using a relative dissimilarity metric; see §2.2.5) and the values of fitted parameters to the corresponding PODDS calibrations. Here model fits to different datasets are quantified using the NSEI metric Nash and Sutcliffe (1970) (rather than R^2 using during the third set of fitting tests in §5.7.5.5) for comparison with results presented in Husband and Boxall (2010a). The NSEI is calculated as per eq. 6.1, where m_i is measurement t of n , p_i is prediction i of n , and \bar{m} is the mean of all measurements.

$$NSEI = 1 - \frac{\sum_{i=1}^n (m_i - p_i)^2}{\sum_{i=1}^n (m_i - \bar{m})^2} \quad (6.1)$$

The range of the NSEI is $[-\infty, 1]$, with 1 corresponding to maximum correlation between observations and predictions.

6.3.2.2 Characteristics of flushed reaches

The VCDM was fitted to each of the four isolated flushing events listed in table 6.3; these datasets were chosen for validating the VCDM erosion model as they collectively feature a variety of pipe materials, roughnesses and source waters.

The shear stress profile during each flush had been calculated by Husband and Boxall (2010a) using roughness values that were either estimated (PODDS-PE1 and PODDS-PE2) or calculated using headloss measurements (PODDS-CI1 and PODDS-CI2). To successfully hydraulically calibrate PODDS-CI2, Husband and Boxall (2010a) needed to reduce the internal diameter by 25 mm and use a roughness of 12.5 mm, indicating that this pipe is in poor condition and is heavily tuberculated. It was important that the roughnesses of the unlined CI pipes were determined by hydraulic calibration as the relative roughness of these pipes was fairly high and shear stress is a non-linear function of relative roughness.

Table 6.3: Several isolated flushing events for which the **PODDS** model has been successfully calibrated (details in [Husband and Boxall \(2010a\)](#)). The flush references given in the first column are specific to this document (for cross-referencing purposes) and were not used by [Husband and Boxall \(2010a\)](#). The figure references in the final column relate to [Husband and Boxall \(2010a\)](#), not this document. Note that **PODDS-PE1** and **PODDS-PE2** were previously referenced in §2.1.12.2 (figs. 2.8c and 2.8d respectively). In [Husband and Boxall \(2010a\)](#) the length of **PODDS-CI2** was incorrectly stated as 230 m and the length of **PODDS-PE2** was incorrectly stated as 280 m in figs. 5 and 11 respectively. The diameter shown in parentheses is the value found through hydraulic calibration (see [Husband and Boxall, 2010a](#)). The **NSEI** values shown in brackets are for all but the last flow step per flush.

ID	Source	Material	D [mm]	L [m]	k_s [mm]	PODDS fit (NSEI)
PODDS-CI1	Groundwater	Unlined CI	95.49	660	3	0.86 (fig. 3)
PODDS-CI2	River water	Unlined CI	76 (51)	93	12.5	0.72 (fig. 5)
PODDS-PE1	River water	PE	89	380	0.01	<0 (0.35) (fig. 9)
PODDS-PE2	River water	PE	72	272	0.01	<0 (0.45) (fig. 11)

6.3.2.3 Fitting method

The **VCDM** was calibrated using each of the flushing datasets listed in table 6.3 by using the **FIT_SIMULT** method (5.7) to fit the erosion parameters β_e and α but not the regeneration rate, as regeneration was assumed to be negligible during each modelled flush (i.e. $\beta_r \approx 0 \text{ s}^{-1}$).

The wall state boundary conditions were fitted using a single cusp approach (§5.8.3): the relative amount of material at the start of the simulation was polarised around a shear strength τ_c , which, like β_e and α , was fitted by **FIT_SIMULT**. Here τ_c is conceptually similar to the eponymous term in the **PODDS** model. It is reasonable to assume simple, polarised wall sub-model boundary conditions given that the studied **DMA** flushes were in pipes that were believed to have been undisturbed for many months (or possibly years) beforehand (§6.3.2.1). This boundary condition estimation method had several advantages over others proposed in §5.8:

- It requires little historical data: little if any historical data is available for **DMA** flushes;
- It does not require an accurate estimate of the prevailing flow: this is useful as prior information of the prevailing shear stress in **DMA** mains (such as from hydraulic models) may not be particularly accurate; Identified τ_c values can be compared τ_c values used in **PODDS** simulations.

A fourth quantity was fitted by **FIT_SIMULT**, this being the maximum shear strength, τ_{max} , with which material can be bound to the walls of pipes (see §4.4.2). This term is comparable to the **PODDS** $\tau_{ultimate}$ term (§2.2.8.4). τ_{max} was only fitted for the two **PE** pipes as no evidence of a maximum shear strength has been found in **CI** pipes (§2.1.15). Attempts to fit **PODDS-PE1** and **PODDS-PE2** using **PODDS** indicated that the ability to fit the two datasets was greatly increased by ensuring $\tau_{ultimate} < \max(\tau_a)$ ([Husband and Boxall, 2010a](#)).

The **PSO** configuration used for model fitting was as per table 5.4, with the **PSO** problem space bounds defined in table 6.4. The upper bound for β_e was set to 1 rather than 10 as used in §5.7.5.5 to steer **PSO** away from parameter values that are obviously erroneous but are associated with low **ISE** values. For each parameter fitted by **PSO**, the convergence tolerance was a tenth of the lower bound of that variable. Each **FIT_SIMULT** run was executed five times to determine the repeatability of model fitting.

Table 6.4: *PSO* Problem space boundaries used when validating the *VCDM* using the datasets listed in table 6.3.

Quantity	Lower bound	Upper bound
β_e	0.0001	1
α	0.0001	2000
τ_{max}	0.0001	$\max(\tau_a) \cdot 1.1$
τ_c	0.0001	$\max(\tau_a)$ (2nd elevated shear stress level for PODDS-CI1)

6.3.2.4 PODDS-CI1 results (downstream-most reach)

In Husband and Boxall (2010a) PODDS-CI1 was modelled two reaches of unlined CI due to these reaches having slightly different D and k_s values. The *VCDM* was only fitted to the downstream-most of the two reaches (370 m in length) as the upstream and downstream turbidity observations required to determine the net turbidity response (§5.9; required for accurate modelling) were only available for this reach. This also simplified the modelling process: if modelling two reaches with differing static properties then two daisy-chained models (potentially with different α and β_e parameters) would be required (§4.3.6.5), but if just modelling the downstream reach then only a single *VCDM* model is required as D and k_s are consistent over that reach.

Turbidity observations showed that the turbidity response from the second flow increase was much larger than that from the first. The *PSO* bounds for τ_c were therefore set so that it could only take values in the range 0.0001 Pa up to the shear stress following the second flow increase.

Fig. 6.12 shows the turbidity observations and the five turbidity predictions that resulted from the multiple FIT_SIMULT runs. Fitted quantities per FIT_SIMULT run along with *NSEI* values and the number of *PSO* iterations are listed in table 6.5.

Here the FIT_SIMULT process proved highly repeatable: the fitted β_e and α parameters are consistent over the five *PSO* runs, as is the *NSEI*. The slightly greater variability of τ_c , along with the maximum number of iterations being reached for the first three runs, indicate that τ_c was the most difficult quantity to fit given the specified tolerance; however the magnitudes of the fitted τ_c values are plausible given the prevailing and elevated shear stress levels.

The *NSEI* is comparable to that achieved for PODDS (table 6.3), but it is believed that a better model fits could have been achieved if the unprocessed flow or turbidity observation time-series had been lagged so that increases in flow and turbidity commenced from the same moments

Table 6.5: Results of fitting the *VCDM* to the PODDS-CI1 dataset five times using *PSO*. The fitted variables were the β_e and α *VCDM* parameter values, the cleaning shear strength, τ_{max} , and the polarising shear strength, τ_c . Also shown is the relative similarity between prediction and observation time-series (*NSEI*) and the number of iterations the *PSO* algorithm took to converge.

Run	β_e	α	τ_{max}	τ_c	<i>NSEI</i>	Iterations
0	0.000640	0.593	N/A	1.087	0.889	213
1	0.000641	0.593	N/A	1.058	0.889	134
2	0.000641	0.593	N/A	1.084	0.889	163
3	0.000641	0.593	N/A	1.080	0.889	93
4	0.000641	0.593	N/A	1.024	0.889	219

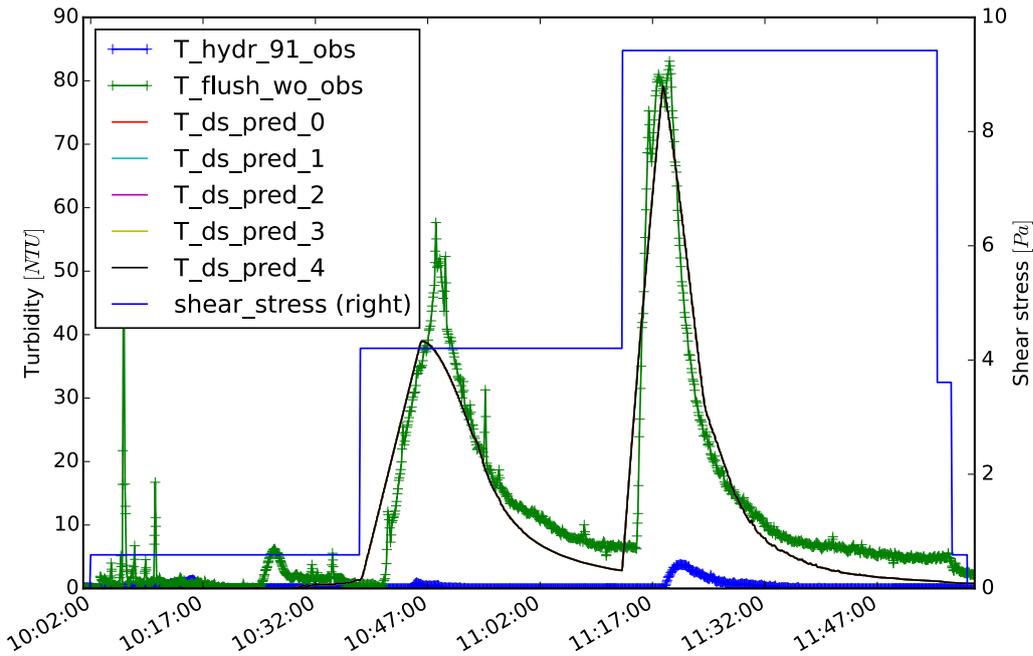


Figure 6.12: Results of fitting the **VCDM** to the **PODDS-CI1** dataset. $T_{hydr_91_obs}$ and $T_{flush_wo_obs}$ are the turbidity observed at the upstream and downstream end of the modelled 370 m reach respectively; $T_{ds_pred_0}$, $T_{ds_pred_1}$... $T_{ds_pred_4}$ are the turbidity predictions generated by five different **PSO** runs, which are sufficiently similar that only $T_{ds_pred_4}$ appears to have been plotted here.

in time. Such lagging appears to have been performed by **Husband and Boxall (2010a)** but details of this are not provided in that article.

6.3.2.5 PODDS-CI2 results

The results of the five **FIT_SIMULT** runs for fitting the **VCDM** to the **PODDS-CI2** dataset are presented in fig. 6.13 and table 6.6.

The β_e parameter values appear less consistent than for **PODDS-CI1** but are all within the convergence tolerance of each other. The α values are very similar but the τ_c value varies considerably; however the **NSEI** is very consistent over the five runs, indicating that this particular fitting problem is not particularly sensitive to τ_c . Again, the τ_c values appear to be plausible given the understanding of variation in τ_a . The **NSEI** is again comparable to that achieved for **PODDS**.

Note that shape of the turbidity response during the first flow step of the flush differs between the unprocessed data plotted in figs. 6.13 and the relevant figure in **Husband and Boxall (2010a)**: the first six-point turbidity spike was removed from the latter. This spike corresponds

Table 6.6: Results of fitting the **VCDM** to the **PODDS-CI2** dataset five times using **PSO**.

Run	β_e	α	τ_{max}	τ_c	NSEI	Iterations
0	0.000501	0.582	N/A	4.076	0.740	154
1	0.000507	0.587	N/A	4.393	0.740	146
2	0.000487	0.573	N/A	3.612	0.739	190
3	0.000509	0.590	N/A	4.503	0.740	219
4	0.000498	0.582	N/A	3.990	0.740	154

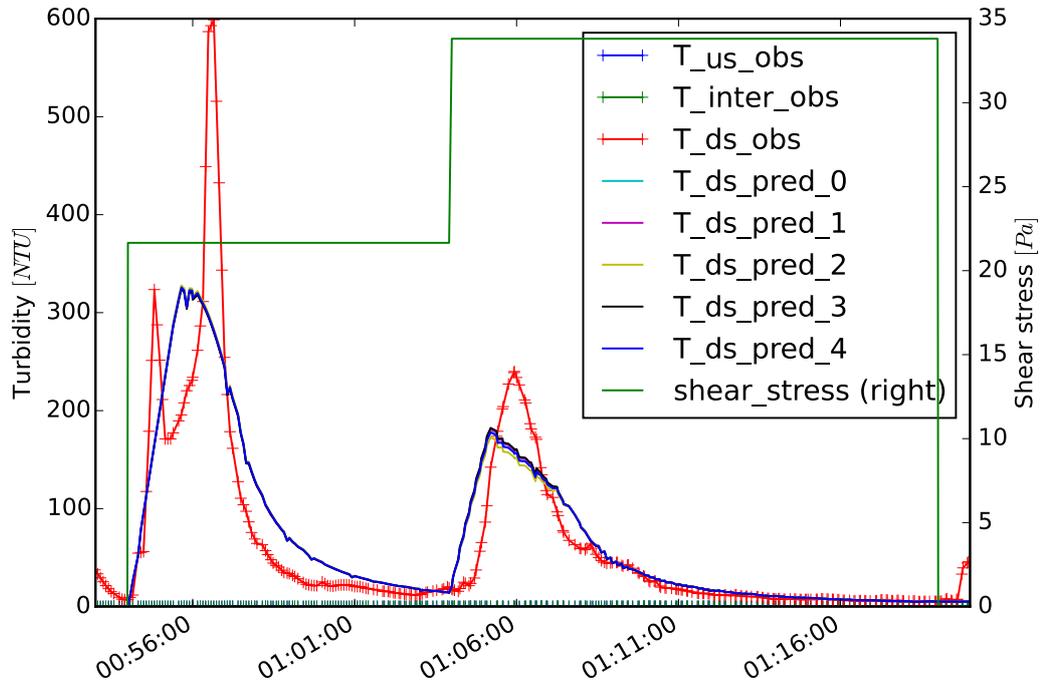


Figure 6.13: Results of fitting the **VCDM** to the **PODDS-CI2** dataset. T_{inter_obs} is the turbidity observed at an intermediate point between the upstream (T_{us_obs}) and downstream (T_{ds_obs}) monitoring points. Again, $T_{ds_pred_0}$, $T_{ds_pred_1}$... $T_{ds_pred_4}$ are the turbidity predictions generated by five different **PSO** runs.

to multiple data points and so it was not considered justifiable to remove it prior to this analysis.

6.3.2.6 PODDS-PE1 results

The results of the five **FIT_SIMULT** runs for fitting the **VCDM** to the **PODDS-PE1** dataset are presented in fig. 6.14 and table 6.7.

Four out of the five fits were fairly consistent in terms of both parameter values and **NSEI**, with the **NSEI** being similar to that achieved by **Husband and Boxall (2010a)** for **PODDS**. However, the third **FIT_SIMULT** run yielded a very high β_e value and a negative **NSEI**. This serves as a reminder of the importance of conducting multiple runs when using a metaheuristic such as **PSO** that is not *guaranteed* to return optimal parameter values (§2.2.4). Both the τ_{max} and τ_c values are plausible given the understanding of τ_a variation. Aside from the obviously erroneous fit, the most difficult value to fit is again τ_c .

Table 6.7: Results of fitting the **VCDM** to the **PODDS-PE1** dataset five times using **PSO**.

Run	β_e	α	τ_{max}	τ_c	NSEI	Iterations
0	0.013771	0.465	1.903	0.002	0.363	296
1	0.011269	0.504	1.779	0.007	0.385	127
2	0.753448	0.688	4.148	3.331	-2.571	176
3	0.011269	0.504	1.779	0.002	0.385	157
4	0.011269	0.504	1.779	0.007	0.385	124

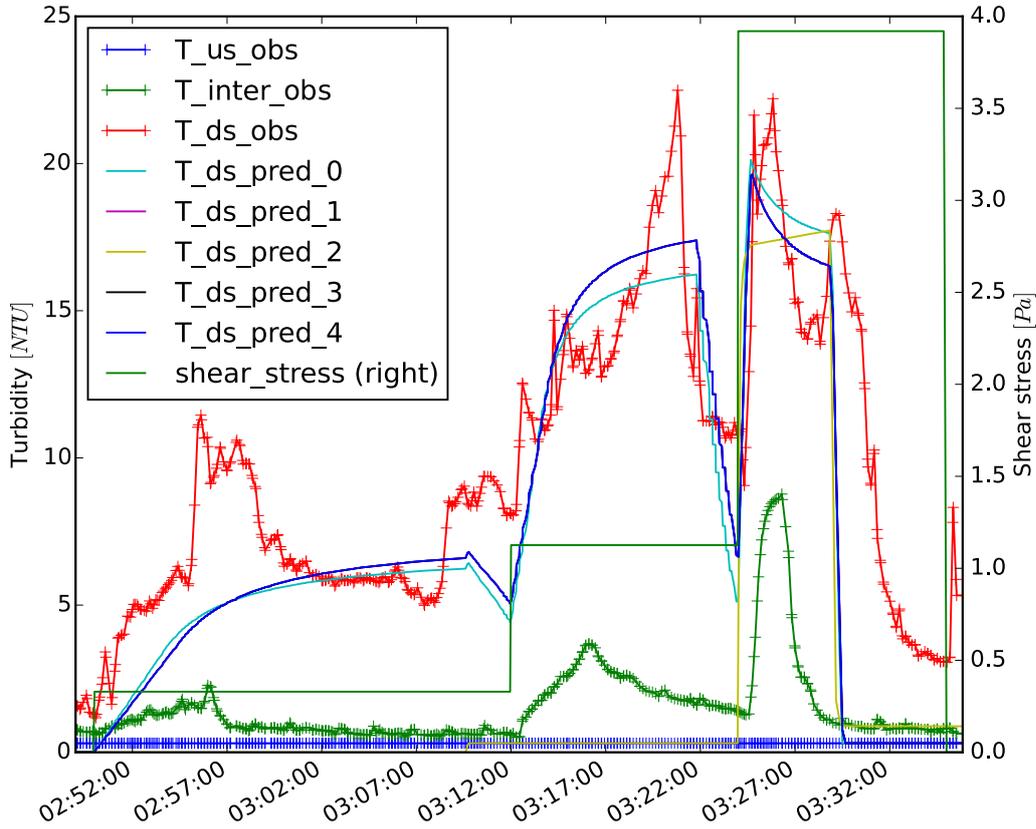


Figure 6.14: Results of fitting the **VCDM** to the **PODDS-PE1** dataset. Again, T_{obs_inter} is the turbidity observed at an intermediate point between the upstream (T_{obs_us}) and downstream (T_{obs_ds}) monitoring points. Again, $T_{ds_pred_0}, T_{ds_pred_1} \dots T_{ds_pred_4}$ are the turbidity predictions generated by five different **PSO** runs.

6.3.2.7 PODDS-PE2 results

The results of the five **FIT_SIMULT** runs for fitting the **VCDM** to the **PODDS-PE2** dataset are presented in fig. 6.15 and table 6.8.

This dataset proved far more difficult to fit than the other three: parameter values and the **NSEI** varied considerably between **PSO** runs, the **NSEI** was negative for all but one run, β_e was the lower bound of its range after one run, yet all runs converged before the maximum number of iterations (500).

The **NSEI** is greater than that of the **PODDS** fit in **Husband and Boxall (2010a)** over the full duration of the flush.

One possible explanation for this flush being difficult to fit is that the downstream turbidity observations were incomplete and/or partially erroneous. The background turbidity at the

Table 6.8: Results of fitting the **VCDM** to the **PODDS-PE2** dataset five times using **PSO**.

Run	β_e	α	τ_{max}	τ_c	NSEI	Iterations
0	0.008173	7.488	8.323	8.196	-0.778	215
1	0.001210	0.272	9.608	0.023	-0.037	187
2	0.000936	0.310	3.541	0.016	-0.382	296
3	0.000100	350.392	5.703	8.457	-2.173	77
4	0.000879	0.544	3.954	0.006	0.181	154

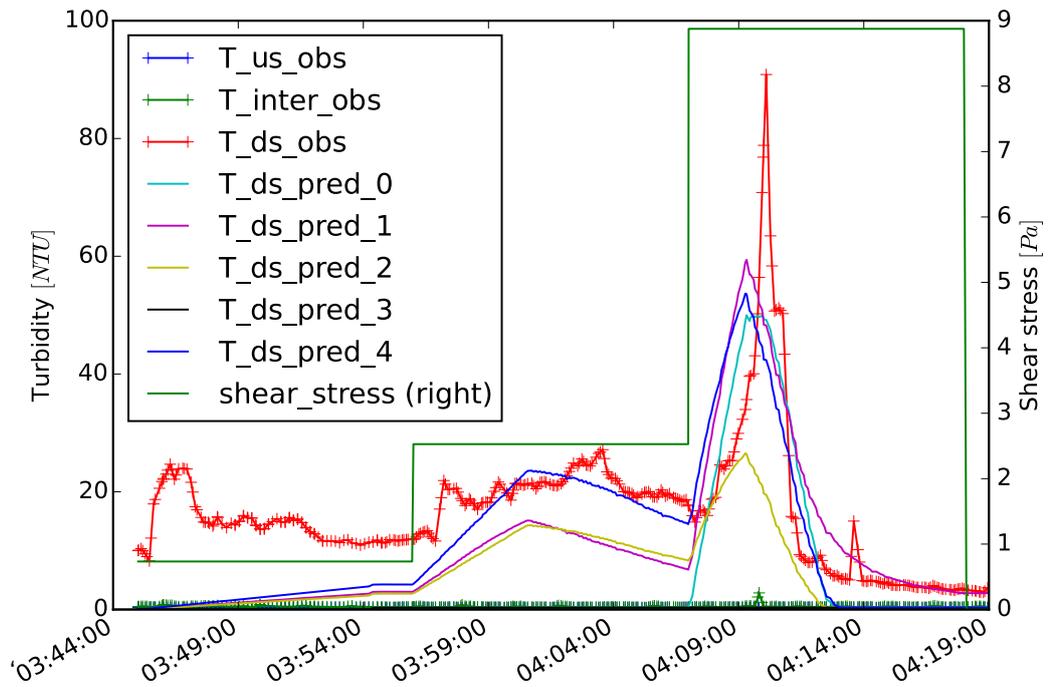


Figure 6.15: Results of fitting the **VCDM** to the PODDS-PE2 dataset.

downstream monitoring location before the flush was greater than the background after the flush yet the turbidity at the upstream end of the pipe reach and at the intermediate monitoring point varied little during the flush. This indicates that the change in background levels was due to error in the downstream monitoring rather than a change in the turbidity of water flowing in to the main.

6.3.2.8 Conclusions

- The **VCDM** erosion model is able to fit the first two datasets well. The best correlations between predictions and observations were similar for the **VCDM** and **PODDS**. The ability to fit the fourth dataset, PODDS-PE2, is believed to have been hampered by data quality. There is therefore no strong evidence for making changes to the design of the **VCDM** erosion model.
- The fitting process is mostly repeatable, although it is necessary to run FIT_SIMULT multiple times to confirm this repeatability (see the single erroneous fit identified when fitting PODDS-PE1).
- The ability to fit the model is impacted by data quality, hence certain reasonable data requirements should be met:
 - The timestamps associated with flow and turbidity time-series observations should be accurate.
 - The magnitude of turbidity observations at both the upstream and downstream ends of the reach of interest should be accurate to allow the ‘net’ turbidity response from the pipe of interest to be determined.

- Fitting the (simple) boundary conditions using FIT_SIMULT yielded plausible results. Converting the mean τ_c of the five PSO runs² yielded conditioning flows that were of sensible magnitude for DMA pipes: $2.7 L \cdot s^{-1}$, $0.8 L \cdot s^{-1}$, $0.2 L \cdot s^{-1}$ and $0.2 L \cdot s^{-1}$ for PODDS-CI1, PODDS-CI2, PODDS-PE1 and PODDS-PE2 respectively. For each dataset the spread in the values of τ_c found by PSO indicated that the τ_c convergence tolerance was overly tight given the input data, model structure and/or calibration method.
- The VCDM modelling provides mixed evidence for the cleaning shear stress within the two PE pipes being less than the maximum shear stress experience over each flush.
 - PODDS-PE1: τ_{max} was less than $\max(\tau_a)$ for four of the five PSO runs;
 - PODDS-PE2: τ_{max} was less than $\max(\tau_a)$ for three of the five PSO runs.

6.3.3 Trunk main case study TM-YR: fitting to flow trials³

6.3.3.1 Motivation and system configuration

The proposed model is only of greater value than PODDS if both its erosion *and* regeneration mechanisms can be shown to be valid. The simplest way of validating the regeneration mechanism is to assess whether it can sufficiently accurately represent the turbidity response during a series of flow trials where the response is primarily due to erosion within a homogeneous pipe. Ideally:

- The flow trials would be separated by periods of lower, recorded flow during which time eroded shear strengths can regenerate.
- The boundary conditions could be estimated with confidence as either (a) the shear stress during the first trial is greater than during any subsequent trial or (b) there is a sufficient duration of accurate flow data before the first trial for any inaccuracies in the boundary condition estimate to be irrelevant by the time of the first trial (§5.8.6).

Three flow trials that satisfy these criteria were conducted over 72.2 days in a trunk main in Northern England. Here this main is referred to as TM-YR. Treated impounding reservoir water is pumped from the WTW up TM-YR to a SR 3.7 km away. TM-YR is a concrete-lined DI main with an internal diameter of 440 mm. Discolouration issues associated with the main persisted following the upgrade of the WTW filters, suggesting that turbidity events were due to material mobilisation within the main. Three flow trials were then conducted in 2009 to qualitatively assess the erosion and discolouration characteristics of the main. Fig. 6.16 is a schematic of the system. Note that further details of the system configuration and the first of these trials can be found in Seth et al. (2009).

The shear stress prior to, during and between the three trials (2009-04-08, 2009-04-22 and 2009-06-18) can be seen in the first subplot in fig. 6.17; upstream and downstream turbidity

²Bar runs 0 and 3 for PODDS-PE2 as the τ_c values for these runs were unrealistically high so were considered outliers.

³This section contains text that has been quoted and adapted from Furnass et al. (2014b).

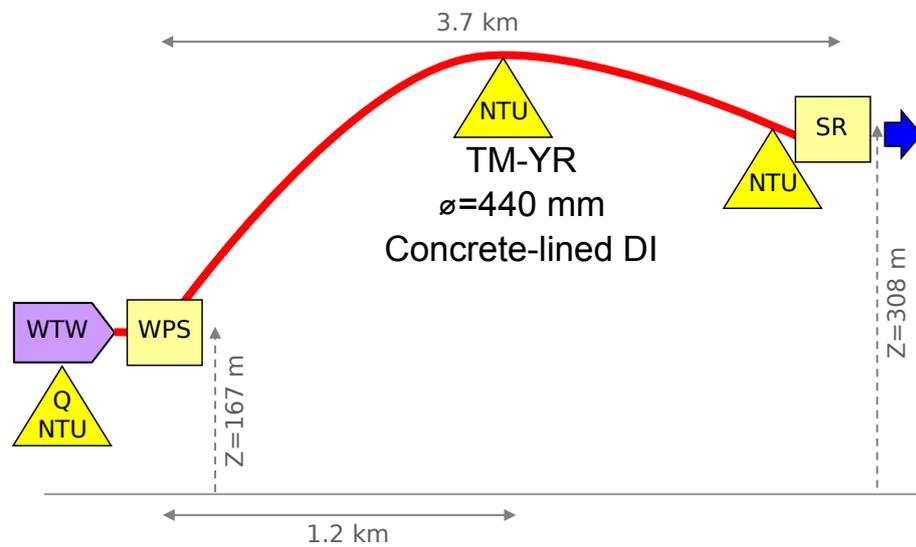


Figure 6.16: A schematic of trunk main TM-YR. Water passes through a WPS at the outlet of a WTW then uphill to a SR. Turbidity and flow at the WTW /WPS are continually logged every 15 minutes via SCADA. Additional temporary, higher-resolution turbidity and flow monitoring was undertaken during the three flow trials (although not all of the three locations were monitored at a higher temporal resolution during every trial).

observations are shown in the subplots below that. Shear stress was estimated by converting a Hazen Williams C value of 120 (taken from a calibrated hydraulic model) and converting it to a Darcy Weisbach roughness of 0.1 mm.

6.3.3.2 Data collation and pre-processing

The trunk main was modelled as a single length as it is of a consistent diameter and pipe material, does not feature any branches and there is limited evidence for longitudinal variation in material accumulation and erosion processes (§2.1.16.6).

The following time-series data were collated:

- **Pipe flow:** $\Delta t = 1 \text{ min}$ for flow trials 1 and 3; $\Delta t = 15 \text{ mins}$ for trial 2 and between trials.
- **Turbidity at the downstream end of the main:** $\Delta t = 10 \text{ s}$ for trials 1 and 3; $\Delta t = 5$ to 25 mins for trial 2 (spot sampling); no monitoring between trials.
- **Turbidity at the upstream end of the main:** i.e. WTW finals; $\Delta t = 15 \text{ mins}$. Trial-specific monitoring was performed during trials 1 and 3 with $\Delta t = 10 \text{ s}$.

These time-series were combined as a single dataset (fig. 6.17) with a variable timestep duration. This involved:

1. Lagging the per-trial time-series to account for perceived inaccuracy in the recorded timestamps. The following lags gave the maximum cross-correlation between per-trial flow data and the lower-resolution long-term flow observations recorded at the WTW/WPS via SCADA:

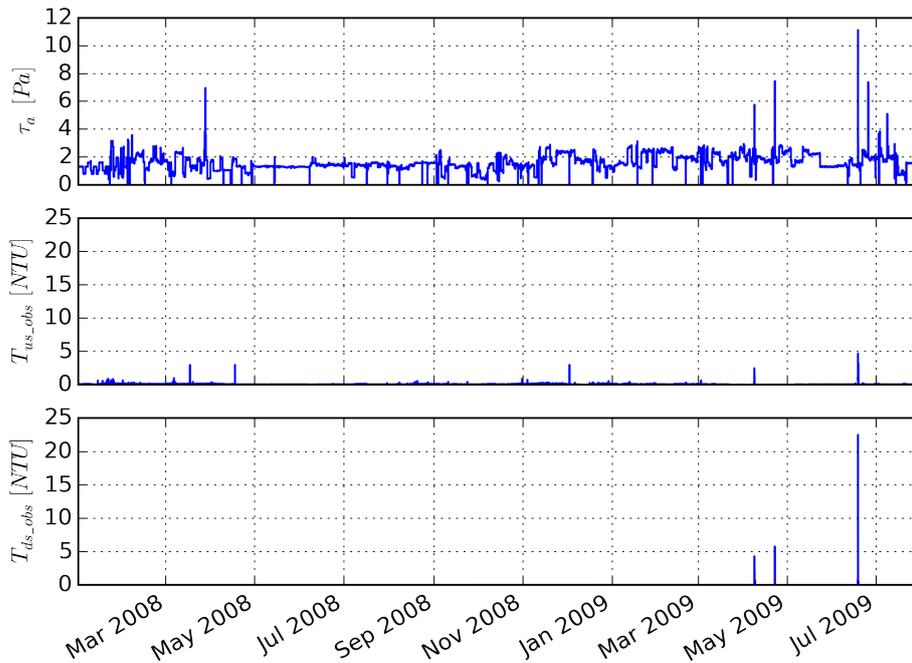


Figure 6.17: Three flow trials were conducted in trunk main TM-YR on 2009-04-08, 2009-04-22 and 2009-06-18. As can be seen from the first subplot, the flow and shear stress were also briefly elevated on three other occasions including on 2008-03-28, which is over a year before the first flow trial. The τ_a and T_{us_obs} profiles shown here are composites of low-temporal-resolution long-term monitoring data and higher-resolution monitoring during trials. As can be seen, the downstream SR inlet was only monitored during trials 1, 2 and 3. The mean and median of T_{us_obs} were 0.10 NTU and 0.06 NTU respectively, whereas the mean and median of T_{ds_obs} (just during trials 1 to 3) were 2.40 NTU and 0.65 NTU respectively.

Trial	1	2	3
Lag [minutes]	-49	-60	-54

2. Removing and linearly interpolating over outlying turbidity values e.g. a single high upstream turbidity value during the first trial that did not correlate with a flow increase.
3. Removing all but the first of duplicate timestamps.
4. Resampling the trial data at a resolution of exactly 10 s.
5. Ensuring the flow and turbidity data for each trial are all valid by infilling any resulting missing values using either forward filling (flow data) or time-aware linear interpolation (turbidity data) and clipping any negative values to a minimum of zero (needed for the turbidity data during trials). Linear interpolation was used for the turbidity data as a more sophisticated interpolation mechanism could not be justified given the non-smooth nature of turbidity signals.
6. Splicing the high-temporal-resolution per-trial data into the long-term flow and turbidity data, then infill any resulting missing values using linear interpolation. A shear stress profile was then calculated from this composite flow profile.

6.3.3.3 Modelling TM-YR: Fitting attempt 1

During the first flow trial, T_{ds_obs} does not start to increase (and therefore no material is mobilised) until the shear stress exceeds 4.28 Pa (fig. 6.18). This information was used to set

the boundary conditions⁴: for the first attempt at modelling TM-YR, the model was fitted from the start of trial 1 to the end of the third trial and it was assuming that all tracked shear strengths less than or equal to $4.28 Pa$ were fully depleted at the start of the simulation and all other tracked shear strengths at that time were fully regenerated (as per the second example in §5.8.7).

The amount of material at the wall was tracked for 100 shear strength bands covering the range 0 to $11.13 Pa$. The three model parameters were fitted using the FIT_SIMULT method (see fig. 5.42). The parameter box bounds were:

- $\beta_e : [1 \times 10^{-7}, 0.1] Pa^{-1} \cdot s^{-1}$
- $\beta_r : [3.169 \times 10^{-9}, 3.169 \times 10^{-7}] s^{-1}$ (equivalent to $[1/10, 1/0.1]$ years)
- $\alpha : [0.01, 100] TPMU \cdot Pa^{-1} \cdot m^{-2}$

The convergence tolerances were a tenth of the lower box bounds.

The upstream turbidity profile was used as an input to the Lagrangian transport mechanism to model the influx of suspended material from upstream.

As in §5.7, the objective function was the ISE of downstream turbidity observations and predictions. This dissimilarity metric was only calculated for timesteps where the observed turbidity exceeded $0.5 NTU$: preliminary investigations indicated that this prevented the fitting process being inadvertently steered by low magnitude turbidity variation and possibly measurement inaccuracy before and after trial flow increases.

Results of the fitting attempt are shown in table. 6.9 and figs. 6.19, 6.20 and 6.21. Again, the fitting process was repeated three times to account for the non-deterministic nature of PSO. The quality of fit was quantified using the R^2 correlation metric, in keeping with the FIT_SIMULT validation tests (§5.7.5.5) but not the validating of the erosion model (§6.3.2), where the NSEI was used instead (to allow for a comparison with previously published results). The R^2 metric was calculated for the same subset of timesteps as the ISE. Note that the y-axis scales of figs. 6.19, 6.20 and 6.21 all differ.

⁴Note that this threshold was previously incorrectly given in Furnass et al. (2014b) as $3.15 Pa$ rather $4.28 Pa$.

Table 6.9: Results of the first attempt to fit the VCDM to the TM-YR trunk main dataset. β_r^{-1} , (the reciprocal of the fitted relative regeneration rate parameter) expressed in years, is an estimation of how quickly a shear strength band could reach maximal regeneration after being fully depleted.

Config	Run	β_e	β_r	α	ISE	PSO itrs	β_r^{-1} [yrs]	R^2
1	0	0.033083	6.642×10^{-8}	0.544	30153.361	283	0.477	0.702
1	1	0.001744	5.308×10^{-7}	0.658	22040.922	241	0.597	0.775
1	2	0.000333	1.033×10^{-7}	0.705	16236.069	81	0.307	0.825

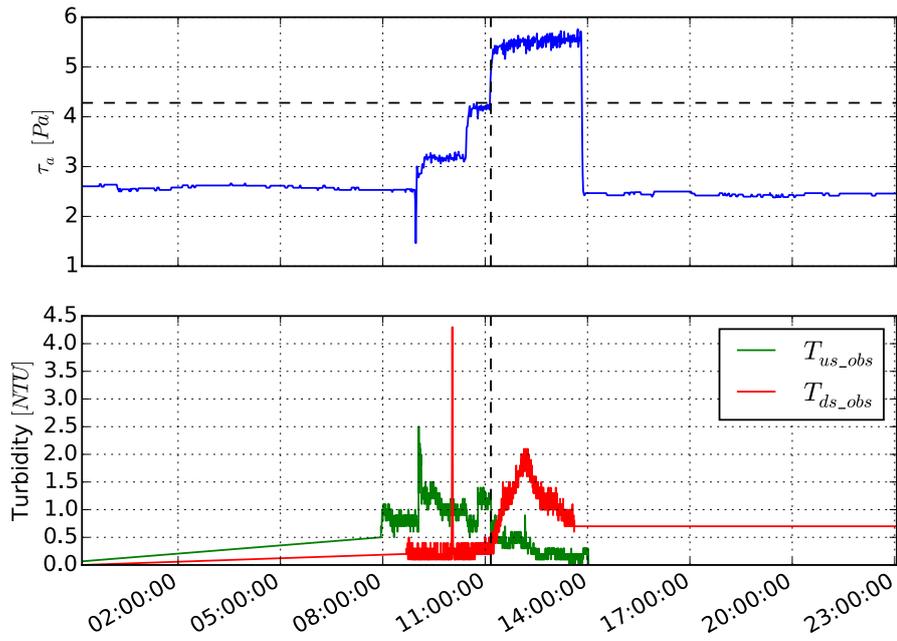


Figure 6.18: During trial 1 there was negligible erosion until the shear stress exceeded 4.28 Pa.

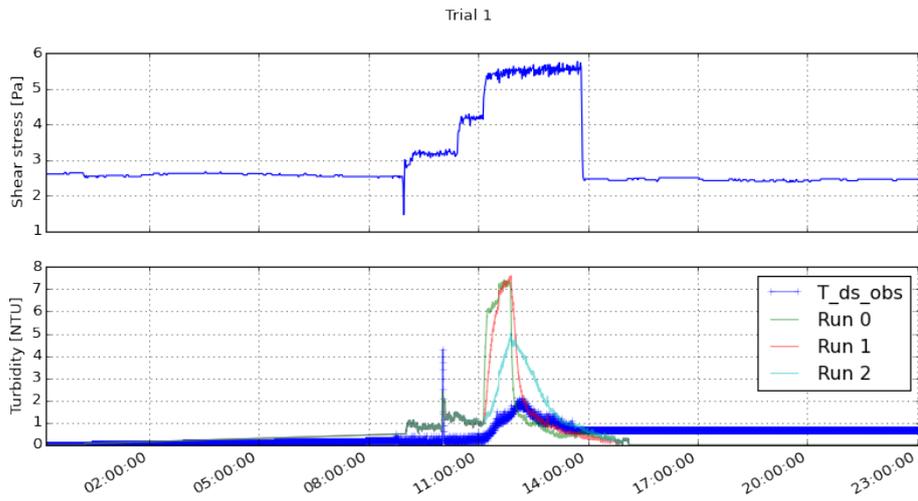


Figure 6.19: Comparing turbidity observations and three predictions made during the first attempt to fit the model to the TM-YR dataset. Only trial 1 is shown here.

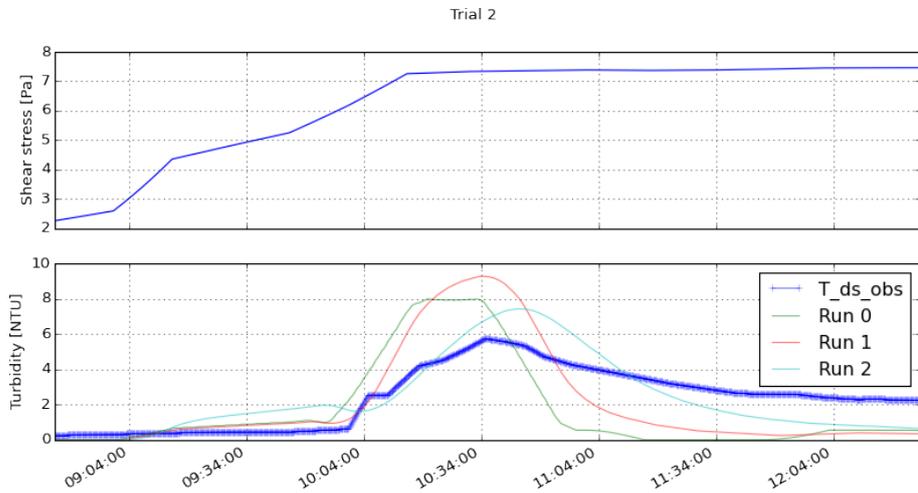


Figure 6.20: Comparing turbidity observations and three predictions made during the first attempt to fit the model to the TM-YR dataset. Only trial 2 is shown here.

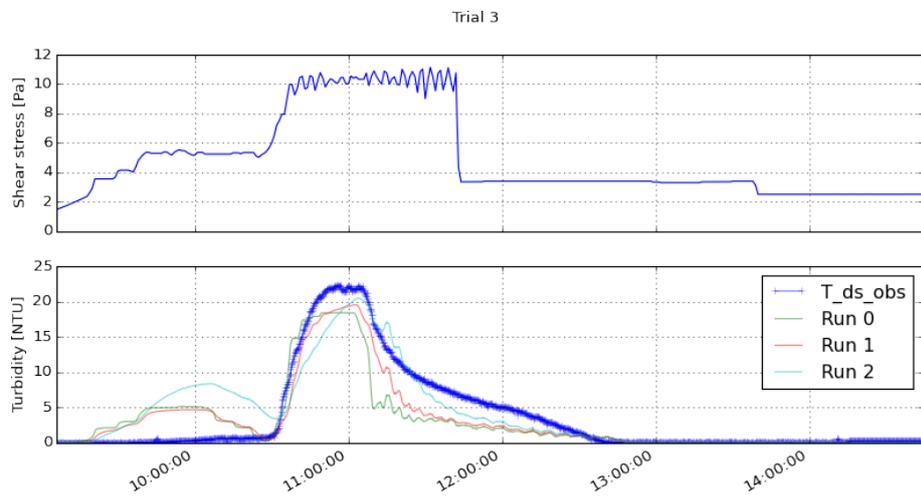


Figure 6.21: Comparing turbidity observations and three predictions made during the first attempt to fit the model to the TM-YR dataset. Only trial 3 is shown here.

Although the R^2 values of the three fitting runs were reasonable, the model over-predicted the turbidity response from all flow trials apart from during the second flow step of trial 3. The first shear stress step during trial 3 does not exceed the maximum shear stress of trial 2; the affected shear strengths are therefore very likely to be in the earlier stages of regeneration at the start of trial 3 yet much mobilisation is predicted by the three runs for the first shear stress step of trial 3. This is believed to be due in part to the overly fast regeneration rate values identified by FIT_SIMULT: 0.3 to 0.6 years for full regeneration is unlikely given the evidence from numerous previous field trials (see §2.1.14).

Also note the variability of the fitted parameters between the three FIT_SIMULT runs: there is two orders of magnitude between the lowest and highest β_e values.

6.3.3.4 Modelling TM-YR: Fitting attempt 2

A second attempt was made to fit the model to the TM-YR dataset, this time using more historical shear stress information to shape the estimate of the material quantity versus strength relationship ($\varphi(\tau, t)$) at the onset of trial 1.

The model was fitted to data from the period starting at 2008-03-27 (the first period of elevated flow shown in fig. 6.17), up until the end of the collated, pre-processed dataset (2009-07-31). The wall state was estimated using VCDM conditioning (§5.8.6): the amount of material at the pipe wall was assumed to be maximal for all tracked shear strengths at midnight on 2008-03-27, then the shear stress increase on that day and the shear stress variation over the subsequent 376 days before trial 1 all informed the $\varphi(\tau, t)$ estimate at the start of trial 1 (by which time the modelled wall state will more accurately reflect reality). Erroneous turbidity predictions during the model ‘bedding in’ period during 2008 do not affect the fitting process as, within the FIT_SIMULT objective function, the ISE was again only calculated for times during trials when $T_{ds_obs} > 0.5 NTU$.

Results of the second fitting attempt are shown in table. 6.10 and figs. 6.22, 6.23 and 6.24. Again, note that the y-axis scales of these figures differ.

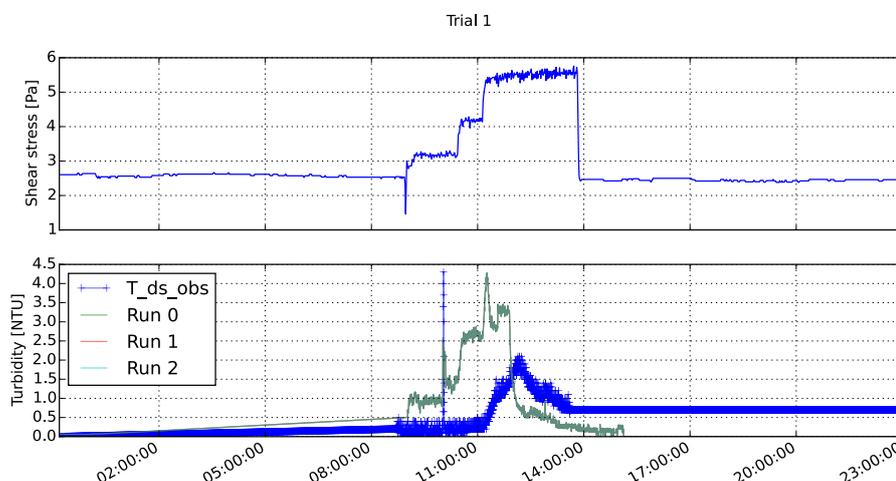


Figure 6.22: Comparing turbidity observations and three predictions made during the second attempt to fit the model to the TM-YR dataset. Only trial 1 is shown here.

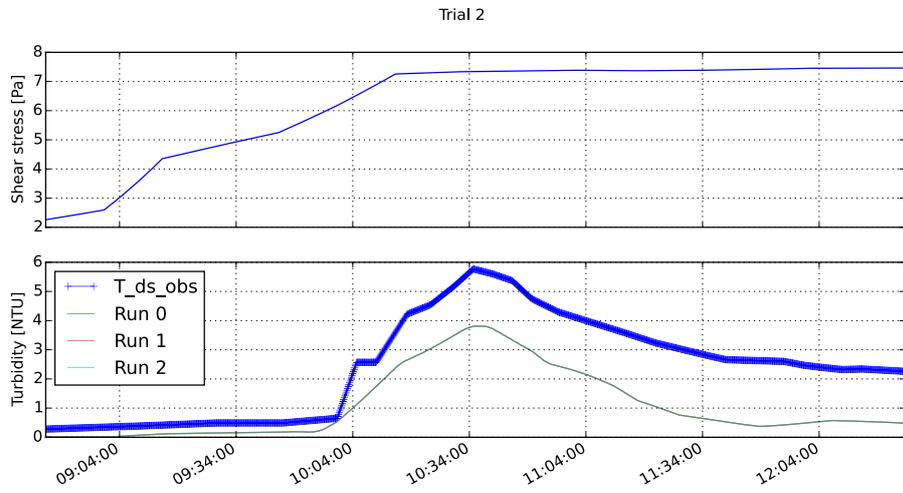


Figure 6.23: Comparing turbidity observations and three predictions made during the second attempt to fit the model to the TM-YR dataset. Only trial 2 is shown here.

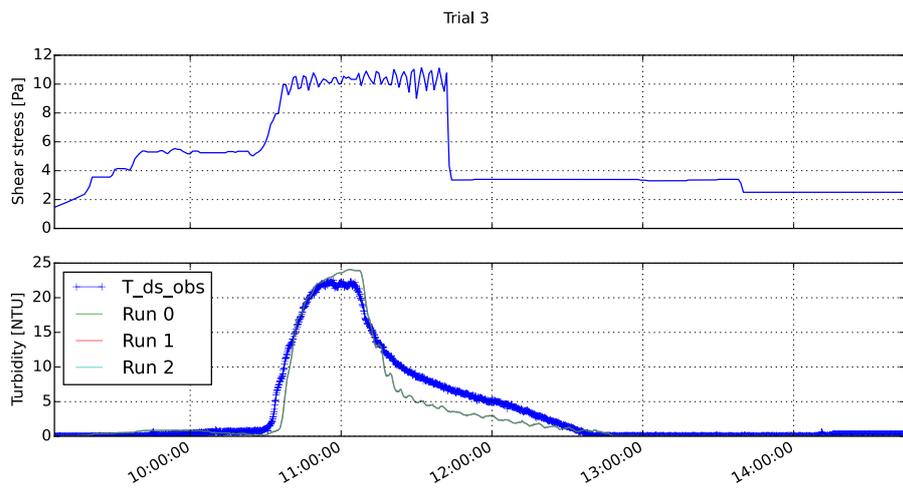


Figure 6.24: Comparing turbidity observations and three predictions made during the second attempt to fit the model to the TM-YR dataset. Only trial 3 is shown here.

Table 6.10: Results of the second attempt to fit the *VCDM* to the *TM-YR* trunk main dataset.

Config	Run	β_e	β_r	α	ISE	PSO itrs	β_r^{-1} [yrs]	R^2
2	0	0.001807	6.135×10^{-9}	0.944	10728.279	108	5.166	0.930
2	1	0.001807	6.137×10^{-9}	0.944	10728.279	96	5.164	0.930
2	2	0.001807	6.138×10^{-9}	0.944	10728.279	87	5.163	0.930

The three sets of parameters found using the FIT_SIMULT method are far more consistent than those that resulted from the first fitting attempt and the quality of fit (R^2) is notably greater. However, the model fits over-predict the amount of material mobilised at weaker shear strengths during trial 1 and slightly under-predicts the turbidity response during trial 2.

Interestingly, the three FIT_SIMULT runs during the second fitting attempt do not predict the amount of material per shear strength at the start of trial 1 to be polarised around $4.28 Pa$ (fig. 6.25). This explains why the predicted turbidity responses were inaccurate during the first part of trial 1 and indicates that this fitting attempt also used inaccurate estimations of the boundary conditions.

6.3.3.5 Modelling TM-YR: Discussion and conclusions

The second method produced a better overall fit, but the turbidity response predicted for trial 1 was inaccurate for both fitting attempts; neither method of estimating the boundary conditions proved to be entirely suitable.

The inability to fit the model accurately to data could be due to the regeneration model being invalid and/or one of the following:

- Comparing the shear stress and turbidity response profiles during the three trials indicated that the shear strengths that had been eroded during trials 1 and 2 were not further eroded during trial 3. The same is likely to be true of trial 2 as the separation of trials 1 and 2 is less than that of 2 and 3 (but this cannot be proven due to the way the flow during trial 2 was not increased in steps). The turbidity responses were therefore primarily due to the erosion of long-undisturbed shear strengths and the dataset therefore contained little information regarding regeneration to calibrate to.

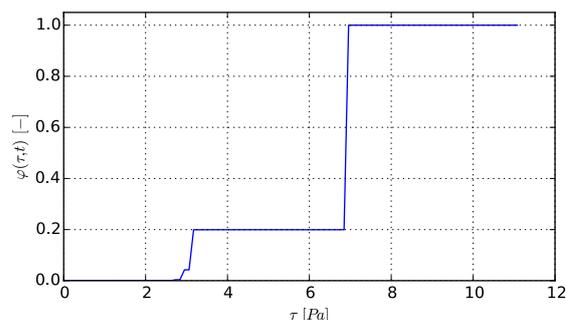


Figure 6.25: The predicted relative quantity of wall-bound material versus shear strength relationship at 2009-04-08 08:56, which was the start of flow trial 1 (fitting attempt 2).

- Also, Q and T_{us_obs} were not measured at a high temporal resolution during trial 2 so may not have been accurately accounted for. Likewise, short-lived flow and turbidity spikes may have occurred before or between the trials that were not accurately captured by the fifteen-minute-averaged logging.
- During trial 3, T_{us_obs} only started to increase *after* T_{ds_obs} did, indicating that the attempts to account for inaccurate per-trial timestamps (by lagging so as to maximise cross-correlation of flow series) were not entirely successful.

Lessons learned from these fitting attempts:

- Most importantly, the ability to validate the [VCDM](#) may depend on the quantity and quality of information describing regeneration processes: it may be easier to fit the model to systems where historical flow increases were separated by longer periods or systems with faster regeneration kinetics (such as unlined iron pipes).
- The accurate timestamping of long-term and short-term monitoring data reduces the need for lagging to ensure maximum consistency of a dataset and other forms of data pre-processing: if the flow and turbidity data are misaligned then this limits the extent to which the model can be fitted to the combined dataset.
- The ability to fit the model to data can sometimes be improved if the dissimilarity metric used within the FIT_SIMULT objective function is only calculated for times when the observed turbidity is great enough to indicate that material erosion has been occurring.

6.3.4 Trunk main case study TM-NR: fitting to a longer, continual dataset

6.3.4.1 Motivation and system configuration

In this case study, the model was fitted to a long time-series dataset where the spacing between certain shear stress increases was such that significant regeneration could occur (unlike for TM-YR; see [6.3.3.5](#)). The dataset was also longer (22 months in duration) and featured more monitored turbidity events than that available for TM-YR. To the author's knowledge, turbidity has only been monitored over multi-month periods in a very small number of trunk mains, and turbidity meters were only installed in these systems in the last five years. The analysis and modelling of this trunk main therefore presents a novel opportunity to study continual erosion and regeneration over an extended period.

Trunk main TM-NR is a large-diameter (800 mm), long (15 km), lined steel trunk main that is situated in the north-east of England. The main runs between a break pressure tank ([BPT](#)) and a [SR](#) (fig. [6.26](#)), conveys treated impounding reservoir water and has historically conveyed flows in the range $8 \times 10^3 \text{ m}^3 \cdot \text{d}^{-1}$ to $24 \times 10^3 \text{ m}^3 \cdot \text{d}^{-1}$.

Flow and turbidity monitoring have been deployed at the downstream end of trunk main TM-NR for several years at the time of writing. To the author's best knowledge, turbidity has been continually monitored for more than a few days in only five or six trunk mains in the UK and no trunk main has been monitored for as long as TM-NR. The dataset used for calibrating the

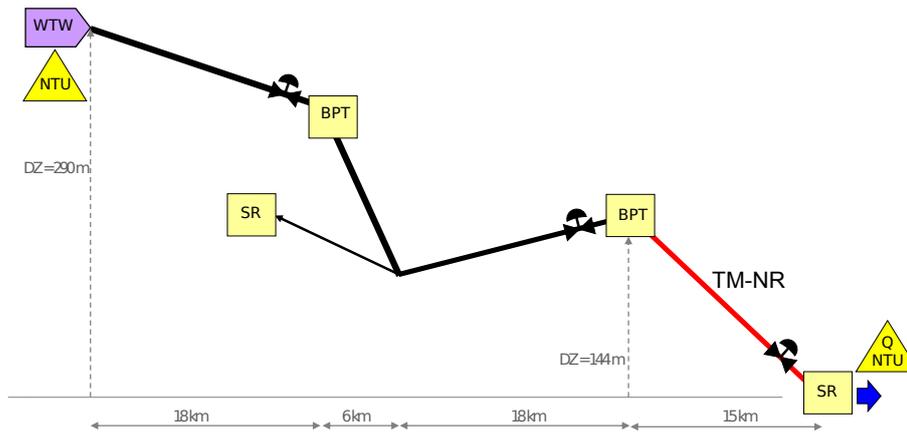


Figure 6.26: A network schematic showing the location of trunk main TM-NR in relation to the supplying WTW and various SRs and BPTs. TM-NR is steel, is 15 km long and has an internal diameter of 800 mm. The schematic also shows the locations of continuous flow (Q) and turbidity (NTU) monitoring equipment.

VCDM corresponded to the period 2010-01-31 at 15:30 to 2011-11-03 at 09:15 (61512 records at 15 minute intervals).

Unfortunately, the turbidity at the upstream end of TM-NR was not monitored over this period, which impacts on the ability to distinguish turbidity due to erosion from the walls of TM-NR from turbidity due to the importing of particulate material from upstream (§5.9). It was therefore necessary to calibrate and validate the VCDM using an estimate of T_{us} . Although T_{us} had not been directly observed, turbidity observations from the supplying WTW (see fig. 6.26) provided some information on the baseline turbidity of water passing into the trunk main system. However, this turbidity signal will most likely have been transformed due to material accumulation/mobilisation within the 42 km of trunk main and the BPTs that lie between the WTW and the upstream end of TM-NR.

Several data pre-processing operations were performed:

- Outer-joining the Q , T_{ds_obs} and T_{wtw_obs} time-series to form a single dataset, using the timestamp as the joining key;
- Converting flows to SI units ($m^3 \cdot s^{-1}$)
- The timestep was typically 15 minutes but with some steps of 30 minutes, so the dataset was resampled at 15-minute resolution and the resulting null values were in-filled using linear interpolation.
- A shear stress profile was calculated using the flow profile, diameter (800 mm) and a roughness of 0.5 mm (taken from a calibrated hydraulic model).

The resulting dataset is shown in fig. 6.27 and summarised in table 6.11. Note that the turbidity instruments/loggers appear to have been set to overly restrictive ranges as the maximum values of T_{ds_obs} and T_{wtw_obs} are exactly 5.0 NTU and 1.0 NTU respectively. Turbidity spikes that exceed these thresholds will therefore be poorly represented in the captured data.

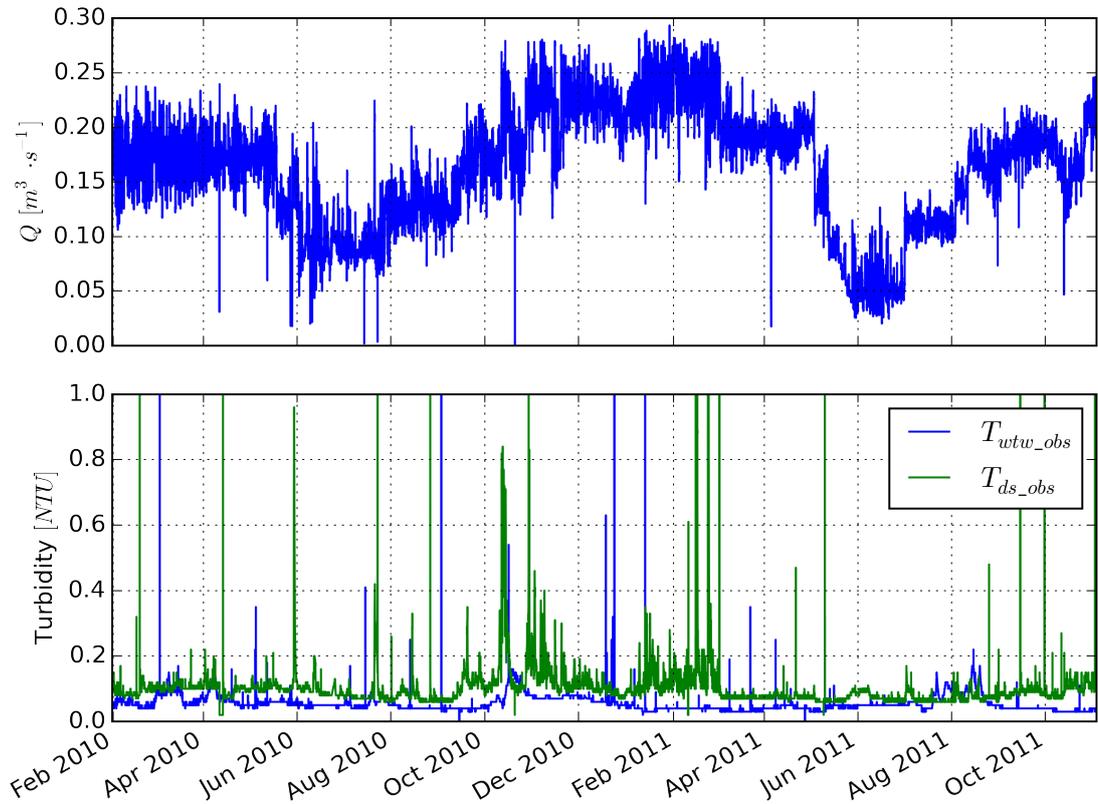


Figure 6.27: Flow and turbidity dataset from trunk main TM-NR used for calibrating and validating the VCDM. Note that the upper bound of the turbidity y-axis scale has been set to 1 NTU to show the variation between erosion events of the turbidity of the WTW finals (T_{wtw_obs}) and the turbidity at the downstream end of TM-NR (T_{ds_obs}).

Table 6.11: Basic statistics of the TM-NR dataset that was used for calibrating and validating the VCDM.

	Q	$t_t (Q > 0)$	T_{ds_obs}	T_{wtw_obs}
count	61512	61512	61512	61512
mean	0.161	15:51:05	0.103	0.054
std	0.055	11:32:40	0.108	0.025
min	0.000	07:08:07	0.020	0.000
25%	0.122	10:30:30	0.070	0.040
50%	0.170	12:17:35	0.090	0.050
75%	0.199	17:09:07	0.110	0.060
max	0.294	47 days 02:58:24	5.000	1.000

6.3.4.2 Shear stress event detection: identifying the times in lengthy datasets when increases turbidity are likely to occur

Given a dataset of shear stress and turbidity time-series pertaining to a single DMA flushing operation (e.g. the observations shown in fig. 6.12), it is usually trivial to identify the times at which there would be a net turbidity response: such responses should follow an increase in shear stress. For continuous shear stress and turbidity datasets covering periods of many months or years (e.g. fig. 6.27), identifying the times when turbidity events are most likely is less easy without automation. The motivation for identifying the times of likely events is that the data over these periods best describes the discolouration characteristics of the main of interest. Model validation can therefore be tailored to better qualify/quantify the ability of the VCDM to characterise the processes of interest by focussing on these disjoint subsets of the time-series dataset.

Turbidity events due to erosion within a pipe are most likely when the pipe shear stress exceeds the recent maximum. The set of likely turbidity events of a given dataset can be identified using the following automatable method:

1. For each timestamp t , mark t as being an event time if $\tau_a(t)$ exceeds the maximum of the previous t_{win} seconds;
2. (Optional) All identified events that correspond to just a single timestamp should be classified as not being events as they could just be due to measurement error;
3. For all remaining events, consider all timestamps within $n_{turnovers} \times t_t$ after each event to be part of that event, thus allowing for discolouration material propagation following the flow increase.

This method has been implemented using several functions that have been packaged with the Python reference implementation of the VCDM (`pyvcdm`). Given the number of times the data at each timestamp is visited by the method, the functions have been implemented as a compiled Cython (Behnel et al., 2011) module for efficiency. The relevant functions are:

- `pyvcdm.eventdetect.is_event(time_offset_array, shear_stress_array, t_win, ignore_start=False)`: returns an array of boolean values, where 'true' indicates that an event occurred. `time_offset_array` is the time in seconds since the start of the dataset. `t_win` is the duration in seconds over which to calculate the recent maximum shear stress. `ignore_start` indicates whether to consider the start of the time-series to be an event by default.
- `pyvcdm.eventdetect.ignore_single_point_events(event_array)`: this is typically called with the output of the first function as input and again returns an array of boolean values, where 'true' indicates that an event occurred and that event is not just an isolated data point.
- `pyvcdm.eventdetect.is_propagated_event(time_offset_array, flow_array, event_array, n_turnovers, D, L)`: typically called with the output of the first or second functions.

The value of τ_{win} obviously has a considerable impact on the number of events that are identified: if τ_{win} is set too low, then many events are identified but a notable proportion of those may correspond to negligible increases in turbidity as insufficient time has passed since the last event for much regeneration to have occurred. If τ_{win} is set too high, then valuable information regarding erosion and regeneration may be ignored through not being associated with an event period.

The value of $n_{turnovers}$ is also a trade-off: if it is too low, then the tail or even the peak of the majority of turbidity responses may not fall within identified events, but if it is too high then distinct turbidity responses may be inadvertently amalgamated as a single event.

This event detection method was applied to the TM-NR dataset. Single point events were *not* ignored as single higher flow values may be significant when flow is only sampled every 15 minutes. $n_{turnovers}$ was set to 5, the same number of turnovers that each flow step was sustained for in the rig experiment detailed in §6.2.2. τ_{win} was set $14 \times 86400 \text{ s}$ (fourteen days), a value that resulted in good correlation between event periods and turbidity increases. The result was the identification of 37 distinct events. Fig. 6.28 shows a portion of the TM-NR dataset with event times superimposed on the T_{ds_obs} profile as red crosses.

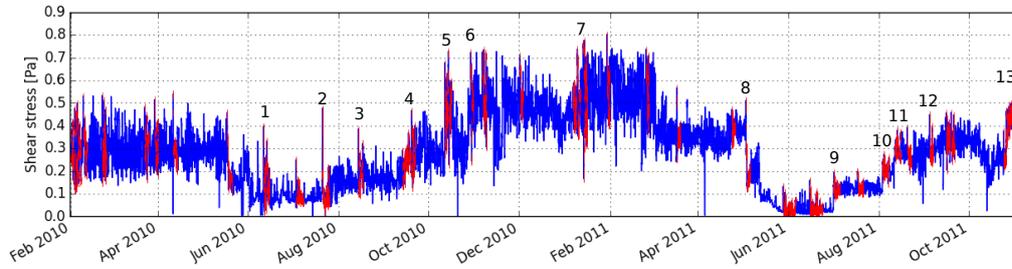
6.3.4.3 Estimating the net turbidity response from trunk main TM-NR

The VCDM can only model the net turbidity response, T_{net} , for a given pipe if the turbidity of water passing into the upstream end, T_{us} , is negligible or observations/estimations of it can be supplied as boundary conditions to the advection mechanism (§5.9). Without this compensation the ability to fit the model by minimising a time-series dissimilarity metric will be significantly impacted: the offset between predictions and observations over a long time-series that is due to T_{us} will dwarf the differences between predictions and observations that are due to material erosion. Plotting T_{ds_obs} for TM-NR (see fig. 6.27) shows that T_{us} is not negligible neither is it constant: higher frequency, higher magnitude components that appear to correspond to material erosion sit atop a lower frequency baseline, where the duration of variations relative to the turnover time indicate that baseline fluctuations are due to changes in the incoming water or potentially to instrument drift.

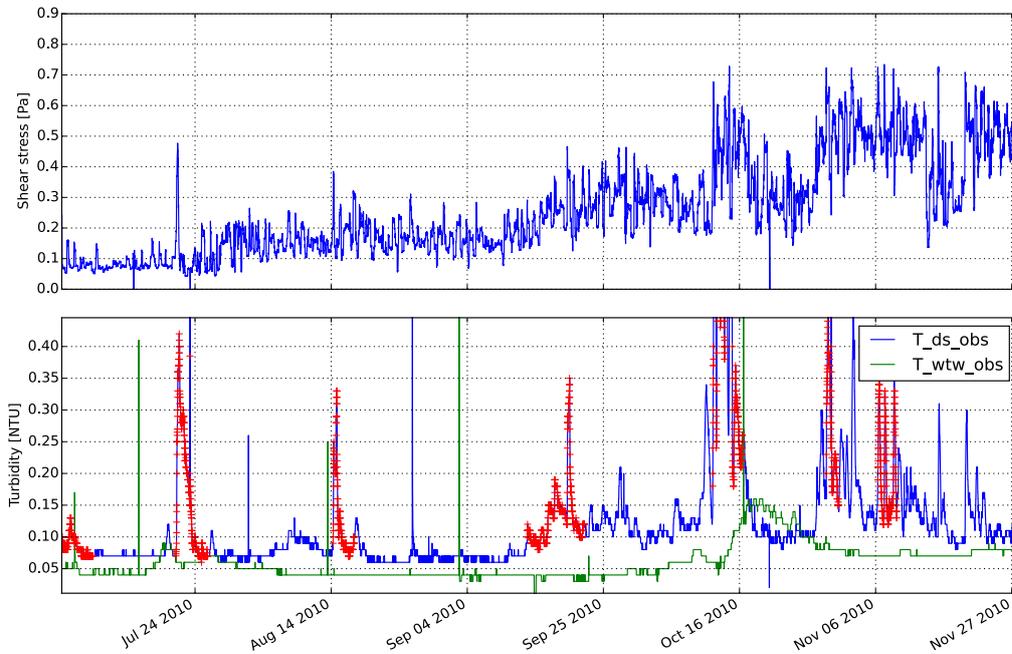
As mentioned previously, T_{us} was not measured directly during the study period. Turbidity was recorded further upstream at the supplying WTW, T_{wtw_obs} , but this monitoring location is 42 km upstream of TM-NR and there are two BPTs en-route, so the unprocessed T_{wtw_obs} signal was not considered to be representative of T_{us} . It is therefore necessary to compensate for T_{us} by transforming either T_{wtw_obs} or T_{ds_obs} into an estimate of T_{us} . Five options for achieving this were explored.

Option 1: Estimating T_{net} by applying a low-pass Butterworth filter to T_{wtw_obs} A low-pass Butterworth filter (§2.2.6) was applied to T_{wtw_obs} in an attempt to remove higher frequency components from the signal and provide an estimate of the background turbidity signal at the downstream monitoring point.

$$T_{net_est} = T_{ds_obs} - low_pass_butterworth(T_{wtw_obs}) \quad (6.2)$$



(a) All 39 identified shear stress events are shown in red. A subset of these events have been enumerated for purposes discussed later in this section.



(b) The correlation between shear stress and turbidity for a small number of shear stress events (these events being shown in red).

Figure 6.28: Demonstration of detecting likely turbidity events from just a shear stress time-series. Here the method presented in §6.3.4.2 was applied to the TM-NR dataset using a time window of fourteen days and a propagation duration of five pipe turnovers.

Here it is assumed that the true, latent T_{us} signal was somewhat periodic, allowing it to be isolated using filters. Also, the time required for turbid water to travel from the **WTW** to the downstream monitoring point was assumed to be sufficiently small relative to the cut-off period (the reciprocal of the cut-off frequency) for (a) advection time to be negligible and so (b) there not to be a need to lag the Butterworth input signal or filtered signal in the following.

In fig. 6.29, T_{net_est} has been plotted along with T_{ds_obs} . The red crosses superimposed on T_{net_est} indicate the previously identified event times (the times at which T_{ds_obs} is most likely to be driven by the effects of material erosion within the modelled pipe; see §6.3.4.2).

For successful removal of the background turbidity, one wants a filter that results in:

- The net turbidity response estimate, T_{net_est} , being above zero for the majority of events;
- The baseline T_{net_est} between events being close to zero (although there may be periods following identified events when there are turbidity responses due to erosion if the identified event was short and sharp enough not to cause full conditioning).

The effect of using different filter cut-off periods was investigated using an IPython Notebook (Perez and Granger, 2007; Shen, 2014): the impact of setting the cut-off period to a value in 5 to 720 days could be instantly visualised using an IPython Notebook 'slider widget'⁵. The filter order was set to 2.

No cut-off period was found for which both of the above criteria were satisfied: certain values resulted in periods where T_{net_est} would be considerably negative whereas other values did remove enough of the background turbidity signal, resulting in T_{net_est} being > 0 NTU between most events (see fig. 6.29).

Options 2 and 3: Estimating T_{net} by applying a moving average filter to T_{wtw_obs} Two other approaches that were investigated involved estimating T_{us} by applying a *rolling mean* or *rolling 5th percentile* to T_{wtw_obs} to remove the undesirable higher-frequency components. T_{net} could then be estimated by subtracting the result, T_{us_est} , from T_{ds_obs} :

$$T_{net_est} = T_{ds_obs} - \text{rolling_mean}(T_{wtw_obs}, \text{window_size}) \quad (6.3)$$

$$T_{net_est} = T_{ds_obs} - \text{rolling_percentile}(T_{wtw_obs}, 5^{th}, \text{window_size}) \quad (6.4)$$

Both rolling aggregate functions are types of moving average filter (§2.2.6). 'Centred' moving averages were produced: the value calculated for each position of the moving window was associated with position at the centre of the window. This approach resulted in no filter output values being assigned to the positions within half the window size of the extremities of the time-series. These positions were infilled by repeating the first and last non-null values out to the very start and end of the time-series.

The reason for comparing these two moving average filters is that the rolling mean is simpler (requires one fewer parameter) whereas the rolling percentile can be more robust when the

⁵See https://ipython.org/_static/sloangrant/sloan-grant.html. Last accessed 2015-03-30.

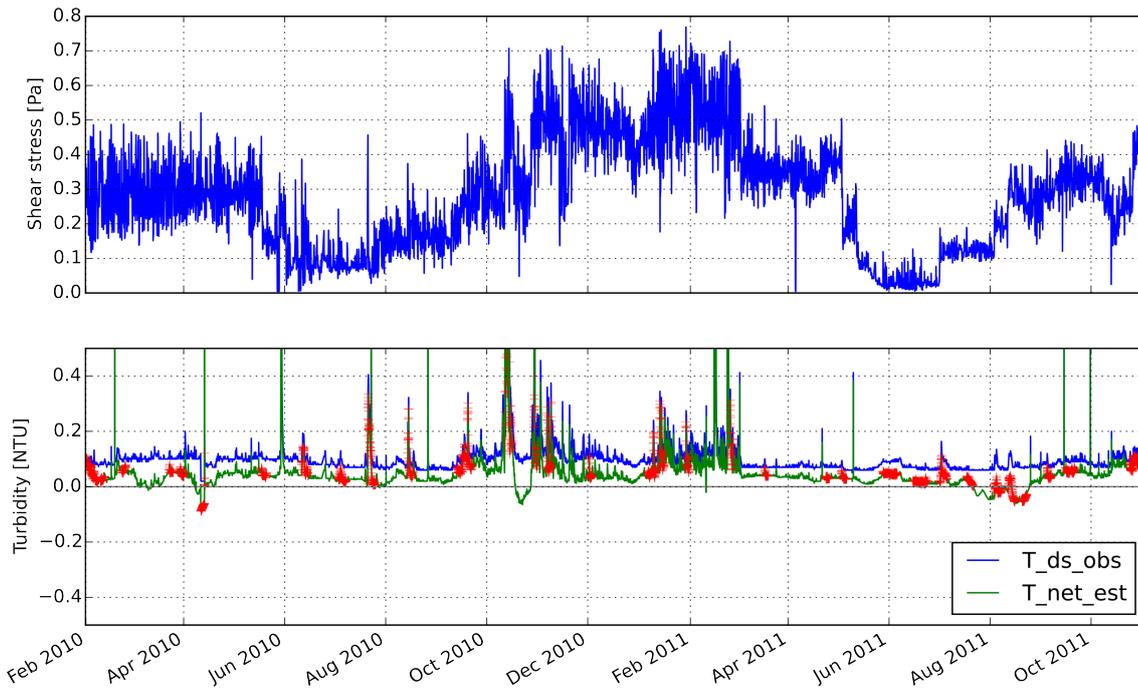


Figure 6.29: Attempting to remove the background turbidity by applying a low-pass Butterworth filter to T_{wtw_obs} then subtracting the result from T_{ds_obs} . Filter order was 2. Cut-off period was 10 days. Red crosses are times at which a ‘shear stress event’ occurred ($t_{win}=14$ days; see §6.3.4.2).

data within the moving average window is not normally distributed, possibly as a result of it containing short-lived spikes (as can be the case with turbidity data).

The impact of the moving average window size was explored using interactive widgets within the IPython Notebook but again it was not possible to identify a moving average window size that ensured most spikes in T_{net_est} were above zero *and* the majority of the signal at non-event times was close to zero.

Options 4 and 5: Estimating T_{net} by applying a Butterworth filter or moving average filter to T_{ds_obs} The conclusion that was drawn from trying to estimate T_{net} by applying filters to T_{wtw_obs} was that T_{wtw_obs} is sufficiently distinct from T_{ds_obs} , possibly due advection lag or to discolouration processes along the mains upstream of the pipe of interest, that T_{wtw_obs} cannot be used to estimate T_{net} .

The ability to estimate T_{net} using just T_{ds_obs} (rather than T_{wtw_obs}) was then explored. Four different approaches were investigated:

- Estimate T_{net} by applying a *high-pass* filter to T_{ds_obs} to remove the (*lower-frequency*) background turbidity:

$$T_{net_est} = \text{butterworth}(T_{ds_obs}, \text{cutoff_period}, \text{type} = \text{high_pass}, \text{order} = 2) \quad (6.5)$$

- Estimate T_{us} by applying a *low-pass* filter to T_{ds_obs} to remove the *higher-frequency* components then subtracting the result from T_{ds_obs} to give T_{net_est} :

$$T_{net_est} = T_{ds_obs} - \text{butterworth}(T_{ds_obs}, \text{cutoff_period}, \text{type} = \text{low_pass}, \text{order} = 2) \quad (6.6)$$

- Estimate T_{us} by finding the rolling mean or rolling 5th percentile of T_{ds_obs} then subtracting the result from T_{ds_obs} to give T_{net_est} :

$$T_{net_est} = T_{ds_obs} - \text{rolling_mean}(T_{ds_obs}, \text{window_size}) \quad (6.7)$$

$$T_{net_est} = T_{ds_obs} - \text{rolling_percentile}(T_{ds_obs}, 5^{th}, \text{window_size}) \quad (6.8)$$

Again, the efficacy of these methods for estimating a suitable T_{net} signal was explored using IPython Notebook interactive widgets. Only the last of these methods, eq. 6.8, yielded an acceptable result: this was using a window size of approximately 10 days (fig. 6.30).

Conclusions For the TM-NR dataset, the net turbidity response from the pipe of interest can be most rationally estimated from observations of the turbidity passing out of the pipe of interest, T_{ds_obs} , by subtracting the rolling 5th percentile (with a window size of ten days) of T_{ds_obs} from T_{ds_obs} . It is not possible to properly validate this estimate without measurements of the turbidity passing into that pipe, but turbidity at the upstream extent of the pipe has not been directly observed. However, the chosen filtering mechanism results in a signal that has the desired qualities of the magnitude between identified shear stress events being near zero and the magnitude during identified shear stress events being much higher. The rolling percentile is believed to be effective here as it is not overly affected by sudden sharp spikes in the filter input data (unlike the rolling mean).

6.3.4.4 Model fitting

Fitting attempt 1 The model was first fitted to the entire dataset bar the first three days. The amount of material at the wall was tracked for 100 shear strength bands covering the range of shear stresses experienced over the fitting period. The shear stress was fairly constant over the first few months of the dataset, allowing the amount of material per shear stress at the start of the fitting period to be polarised around the mean shear stress of the three previous days (similar to the first example in §5.8.7). The relative amount of material present at shear strengths greater than this polarising shear stress at the start of the fitting period was simply assumed to be zero but this may not have been the case in reality.

Again, the three model parameters were fitted using the FIT_SIMULT method (see fig. 5.42). The parameter box bounds were the same as used in §6.3.3 apart from the lower bound for β_e . This was a factor of two greater than that of §6.3.3 so as to steer the fitting process away from regions of the parameter space associated with poor fits and unlikely β_r values, and to increase the repeatability of the fitting process. Again, the convergence tolerances were a tenth of the lower box bounds. Here the objective function was the ISE of estimations (based on observations) and predictions of the net turbidity response (not the downstream turbidity, so there was not a need to supply an upstream turbidity profile to the Lagrangian transport mechanism). To focus the fitting mechanism on the processes of interest, the ISE and R^2 were only calculated from after five turnovers into the fitting period (to discard erroneous turbidity predictions during the model ‘bedding-in’ period) during shear stress events using the method of §6.3.4.2. Results are shown in table 6.12.

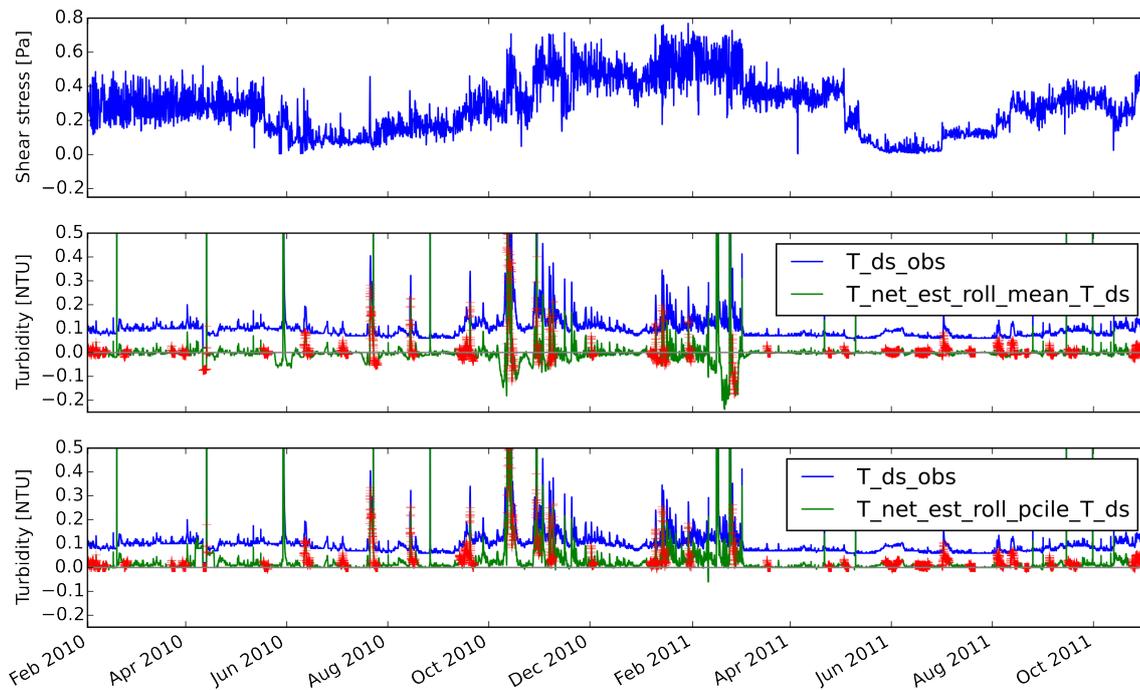


Figure 6.30: Attempting to remove the background turbidity by applying a centred rolling mean (middle subplot) and centred rolling 5th percentile to T_{ds_obs} then in each case subtracting the result from T_{ds_obs} . The window size here was 10 days. Red crosses are times at which a ‘shear stress event’ occurred ($t_win=14$ days; see §6.3.4.2).

Fitting attempt 2 The fitting approach was identical apart from the period of the dataset that was used, which was from midnight on the day of the highest shear stress (2011-01-29) to the end of the dataset, and the boundary conditions. The wall state was estimated using VCDM conditioning (§5.8.6) to see if this resulted in a better fit than using generalisations based on information from network operators and initial flow trends (§5.8.7). The advantage to focussing on this period was that any observed erosion response could safely be assumed to be due to regeneration and the initial wall state could be polarised with greater confidence (assuming that the highest shear stress was fully conditioning).

Fit attempt 3 This attempt was identical to attempt 1 apart from allowing the regeneration rate to vary over time as a function of temperature. Including such a dependency in the fitting process allows the impact of temperature-driven variations in source water characteristics and changes in soil/backfill temperature to be accounted for; both factors have been shown or are expected to influence material accumulation processes (§2.1.16). Seasonal/temperature-driven changes in source water are more likely to be significant in mains supplied with surface water (such as TM-NR). β_r was made temperature-dependent (and therefore time-varying) by:

1. Acquiring monthly mean atmospheric temperature data for the relevant region of the country from the UK Met Office⁶;
2. Resampling the temperature data using the time-series index of the fitting period dataset, using linear interpolation to infill null values;

⁶http://www.metoffice.gov.uk/pub/data/weather/uk/climate/datasets/Tmean/date/England_E_and_NE.txt

3. Within the objective function, calculating a time-dependent regeneration rate using eq. 6.9, where m is a scaling factor and θ is the temperature in $^{\circ}C$.

$$\beta_r(t) = \bar{\beta}_r + m \cdot (\theta(t) - \bar{\theta}) \quad (6.9)$$

4. When fitting β_e , α , $\bar{\beta}_r$ and m using FIT_SIMULT, m was restricted to the following range to ensure that β_r could not be negative at any time.

$$\left[10^{-12}, \frac{\text{lower_bound}(\bar{\beta}_r)}{|\min_t(\theta(t) - \bar{\theta})|} \right] \quad (6.10)$$

Fit attempt 4 This was similar to attempt 2 but used the temperature-dependent regeneration rate of attempt 3. Note that making β_r temperature-dependent increases the complexity of the model (§4.4.2) as here an additional parameter is required. The purpose of exploring this temperature dependence in fitting attempts 3 and 4 was to see if this additional complexity was justifiable and beneficial.

Results The parameters identified using FIT_SIMULT for each of the four fitting approaches are shown in table 6.12, as are the ISE and R^2 correlation metrics and the reciprocal of the (mean or constant) β_r , expressed in years.

The following observations were made:

- The longer fitting period used in attempts 1 and 3 resulted in a better fit than for attempts 2 and 4, although no fitting attempt yielded a high-quality fit (R^2 never exceeded 0.311).
- After increasing the lower PSO bound of β_e from 10^{-7} (results not shown here) to 10^{-5} all fitting attempts bar attempt 3 were consistent over three runs.
- The regeneration rates found during attempts 2 and 4 are unlikely: they are far faster than rates found by linear extrapolation during previous (DMA) studies (§2.1.14).

Table 6.12: Fitting the VCDM to data from trunk main TM-NR: parameters and correlation metrics resulting from the application of the FIT_SIMULT fitting method.

Config	Run	β_e	$\bar{\beta}_r$	m	α	ISE	PSO itrs	β_r^{-1} [yrs]	R^2
1	0	7.650×10^{-5}	1.717×10^{-8}	N/A	5.472	131.751	122	1.846	0.311
1	1	7.650×10^{-5}	1.717×10^{-8}	N/A	5.472	131.751	86	1.846	0.311
1	2	7.650×10^{-5}	1.717×10^{-8}	N/A	5.474	131.751	75	1.845	0.311
2	0	4.005×10^{-5}	3.169×10^{-7}	N/A	0.858	4.722	69	0.100	0.177
2	1	4.002×10^{-5}	3.170×10^{-7}	N/A	0.852	4.722	87	0.100	0.177
2	2	4.005×10^{-5}	3.169×10^{-7}	N/A	0.861	4.722	101	0.100	0.177
3	0	7.650×10^{-5}	1.717×10^{-8}	1.003×10^{-12}	5.471	131.751	102	1.846	0.311
3	1	1.062×10^{-5}	5.307×10^{-8}	2.130×10^{-10}	7.755	135.531	88	0.597	0.274
3	2	1.062×10^{-5}	5.352×10^{-8}	3.691×10^{-11}	7.702	135.522	205	0.592	0.274
4	0	3.980×10^{-5}	3.169×10^{-7}	4.862×10^{-10}	0.847	4.723	110	0.100	0.177
4	1	3.952×10^{-5}	3.168×10^{-7}	3.157×10^{-10}	0.842	4.722	161	0.100	0.177
4	2	3.952×10^{-5}	3.165×10^{-7}	4.090×10^{-11}	0.846	4.721	84	0.100	0.177

- Allowing β_r to vary with temperature did not improve the model fit (R^2). This is potentially due to (a) the form of the relationship between temperature and β_r and/or (b) the PSO bounds of the m scaling parameter not being appropriate.

Quality of best T_{net} predictions during shear stress events Looking closer at the time-series predictions associated with the fitting attempt with the highest R^2 (fitting attempt 1) provides further information about how and when the fitting was successful/unsuccessful. Thirteen of the 39 shear stress events identified using the method outlined in §6.3.4.2 were studied (these events are enumerated in fig. 6.28a).

The fit to the earlier events in the dataset was often reasonable but several increases in turbidity were not and could not have been predicted as they do not correspond to flow increases within TM-NR e.g. during event 1 on 2010-06-12 (fig. 6.31) and during event 3, late on 2010-08-14 (fig. 6.32). These increases in turbidity are most likely due to the importing of material from upstream TM-NR or potentially in the case of event 3 from more material having accumulated at the upstream end of TM-NR. Having T_{us} observations for defining turbidity boundary conditions would therefore have been beneficial when modelling TM-NR. Note that although there is a tank at the upstream end of TM-NR that it is unlikely to smoothen/dampen the upstream turbidity signal as it is a BPT and will therefore be small with a short residence time and limited mixing.

The efforts to discount the unobserved T_{us} signal using the centred, rolling 5th percentile of the T_{ds} observations were not always successful. For example, the baseline turbidity during event 4, shown in fig. 6.33, is elevated in a way that will have a notable impact on time-series correlation metrics such as the ISE. It will also impact on the scale-insensitive R^2 metric if the efficacy of the background turbidity removal method varies over the duration of the fitting period. However, it should be noted that the relative changes in the net turbidity during event 4 were reasonably accurately predicted and, for the most part, the method for turbidity background compensation was effective e.g. during event 6 (fig. 6.34).

The SNR during much of the dataset following the peak shear stress (i.e. the fitting period for fitting attempts 2 and 4) was very low, principally because the observed turbidity signal was also very low (but still within the range of values over which light scattering intensity is proportional to turbidity (§2.1.5)). For a number of shear stress events (e.g. event 8, shown in fig. 6.35), the magnitude of turbidity responses was similar to the turbidity digitisation resolution. The low resolution of the turbidity observations during this period could explain

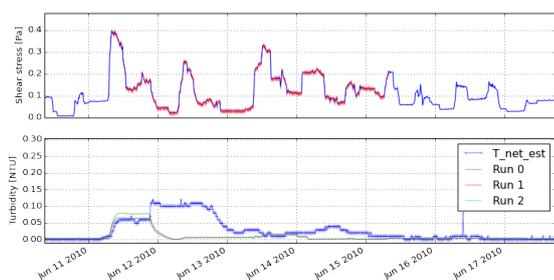


Figure 6.31: Best model fit to TM-NR: fitting attempt 1, event 1 (see also fig. 6.28a)

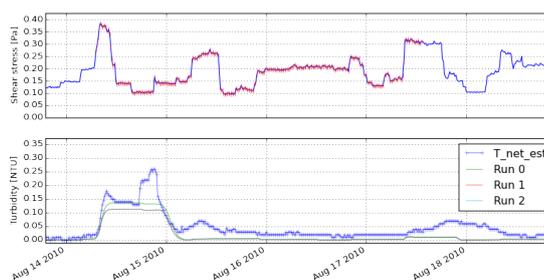


Figure 6.32: Best model fit to TM-NR: fitting attempt 1, event 3 (see also fig. 6.28a)

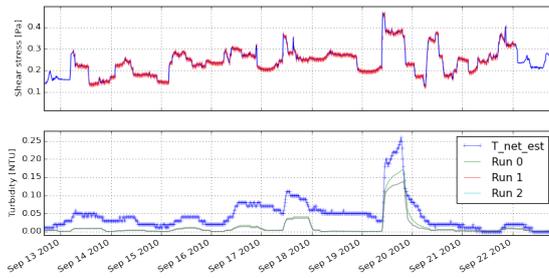


Figure 6.33: Best model fit to TM-NR: fitting attempt 1, event 4 (see also fig. 6.28a)

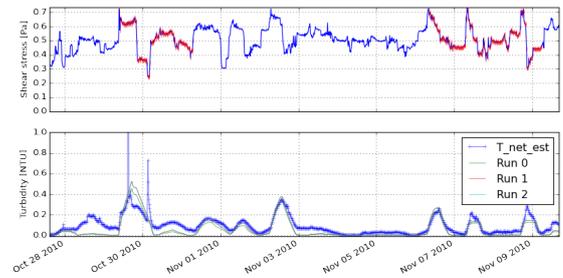


Figure 6.34: Best model fit to TM-NR: fitting attempt 1, event 6 (see also fig. 6.28a)

why the R^2 from fitting attempts 2 and 4 was low (as the signal was less accurate) and the identified parameters were less plausible than for fitting attempts 1 and 3. Also, events during the later part of this period such as events 10, 11 (fig. 6.36) and 13 were over-predicted. This is surprising given that these events mobilised material that had been regenerating over the summer, when one might expect regeneration rates to be higher, yet one interpretation of the time-series plots is that less material had regenerated than anticipated. Another is that the erosion rate over this period was an over-estimate, resulting in turbidity increases that were too short and sharp and a third is that short, sharp increases in flow were not captured by the fifteen-minute-averaged flow observations.

6.3.4.5 Discussion and summary

Although the R^2 of the best model fits was relatively low, visual assessment of turbidity predictions during the major shear stress events indicated that the model was representing the key discolouration processes to an extent. There were however periods where the best model fits correlated less with the estimate of the net turbidity response and there were turbidity events that the model could not account for. There are several suspected reasons for this, including: the net turbidity response being estimated using T_{ds_obs} rather than calculated using both T_{us_obs} and T_{ds_obs} , so the importing of material from upstream of TM-NR only being crudely accounted for; the flow observations being unrepresentative of short, sharp spikes in flow due to each data point being the average of a fifteen-minute period; the turbidity responses during a large portion of the dataset being small relative to the turbidity digitisation step.

The period used for model calibration was found to have an impact on the quality of model fit and plausibility of identified parameters. Attempts to fit the model to just the second part

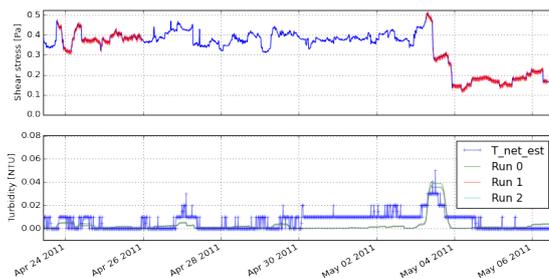


Figure 6.35: Best model fit to TM-NR: fitting attempt 1, event 8 (see also fig. 6.28a)

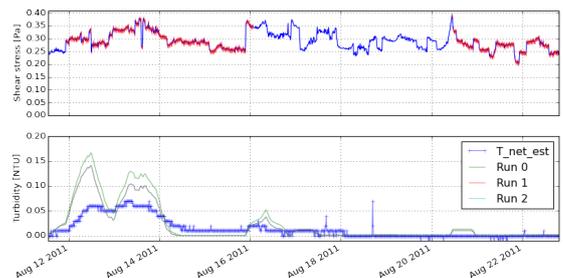


Figure 6.36: Best model fit to TM-NR: fitting attempt 1, event 11 (see also fig. 6.28a)

of the dataset (attempts 2 and 4) yielded poorer results, most likely due to the paucity of shear-stress-driven turbidity events significantly larger than turbidity digitisation step. This highlights the conflicting needs of network operatives and modellers: the former want to manage mains so that flow increases are carefully controlled and turbidity is always low, whereas modellers need turbidity events (below the regulatory limit) to calibrate and validate their discolouration risk models.

The regeneration rate associated with the best model fits was ≈ 1.8 years, a value that is plausible given regeneration rates derived during previous studies (§2.1.14).

The more complex model forms where β_r was permitted to vary with temperature surprisingly did not allow for more accurate model fits. Either the inaccuracies stemming from data availability/quality issues were too dominant, the form of the temperature dependent relationship was not appropriate or regeneration rates vary less with season/temperature than was anticipated.

The amount of mobilisation being over-estimated at the end of summer 2011 was unexpected. As stated previously, hypotheses for this are that the FIT_SIMULT-identified value of β_e was too high or eroding shear stress spikes during the summer were not captured by the fifteen-minute-averaged flow data.

It is unknown to what extent the overly-narrow ranges used when logging turbidity have limited the ability to fit the model to the dataset. Ideally, all turbidity logging should be able to record values from effectively zero to hundreds of NTU, with the number of bits per sample being chosen to provide a high (e.g. 0.01 – 0.02 NTU) resolution.

6.3.5 Trunk main case study TM-WB: fitting to a long dataset containing flow trials then testing predictive performance

6.3.5.1 Site description

The third trunk main case study, like the study of TM-NR, made use of a multi-year dataset for model calibration but this dataset featured deliberate flow trials. These resulted in notable increases in turbidity and provided valuable information for characterising erosion and regeneration processes.

Trunk main TM-WB (see fig. 6.37) is supplied with groundwater and runs between two SRs in the south-west of England (SR1 and SR2). The reach referred to as TM-WB is a series of mains of varying materials (AC and unlined DI), diameters, and flows, due to there being various take-offs along its length.

Network operators want to be able to rapidly increase flows in TM-WB from $\approx 1100 \cdot m^3 \cdot day^{-1}$ to $\lesssim 2800 \cdot m^3 \cdot day^{-1}$ in the event of a WTW fail, SR fail or major burst without there being a negative impact on water quality. However, several deliberate, controlled flow increases ('flow trials') indicated that there is a significant risk of discolouration if flows are elevated above typical levels.

On 2008-08-06, flows were increased whilst monitoring turbidity at two locations: where the material changes from AC to DI (node N5 in fig. 6.37) and at the inlet to the downstream

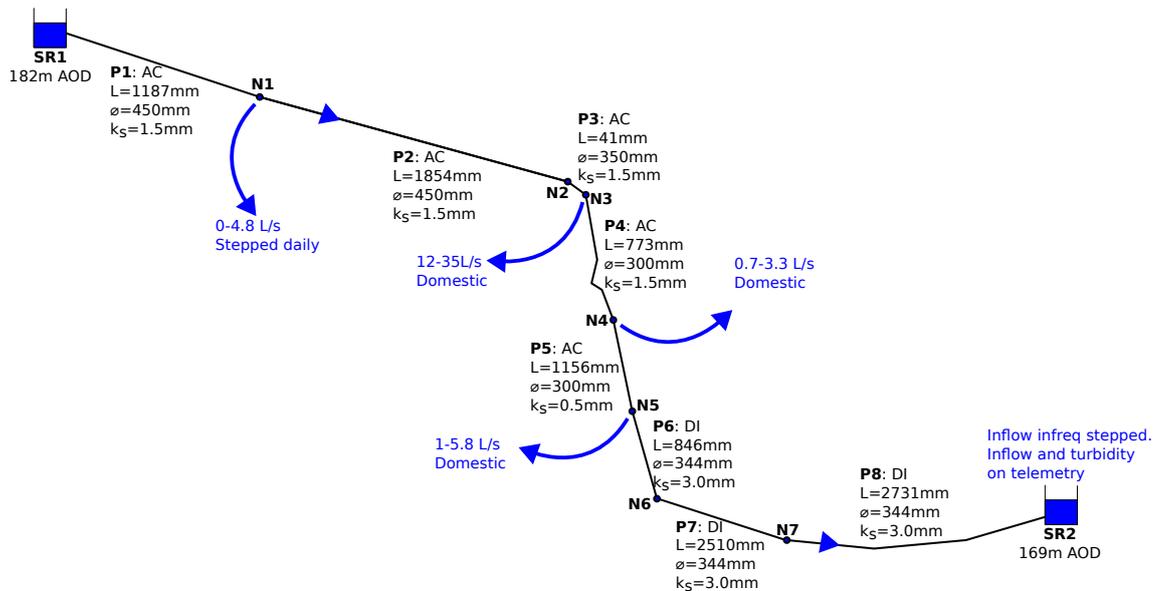


Figure 6.37: A network schematic showing the key attributes of the pipes (P1-P8) of trunk main TM-WB, including which sections are AC and which are unlined DI, the length (L) per section, internal diameter (D) and the roughness (k_s ; taken from a calibrated hydraulic model). Flow and turbidity are logged via telemetry every 15 minutes at the inlet to the downstream SR (permanent turbidity meter installed in March 2012). The schematic also shows the take-offs along the length of the main as they are represented in the hydraulic model of TM-WB, including the range of demand flows per node, the type of flow profile (domestic or stepped) and the duration over which the flow profile patterns repeat. The shown internal diameters are from the calibrated hydraulic model of TM-WB.

SR (SR2) (Husband et al., 2011). The increase in turbidity detected at SR2's inlet (fig. 6.38) exceeded the regulatory limit, but the response detected at N5 was much less, indicating that most of the material eroded during the trial had been mobilised within the length of unlined DI.

During each of a number of subsequent flow trials, the maximum flow was raised closer to the target conditioning flow of $2800 \text{ m}^3 \cdot \text{day}^{-1}$, with the intention that, once the target flow was attained, the flow in the main could be periodically elevated back up to that level (i.e. flow conditioning; see §2.1.17.2) to ensure any necessary rezoning and routing around failed treatment/storage facilities would not cause discolouration. An actuated valve and permanent turbidity meter were then installed at the inlet to SR2 to facilitate accurate flow control and remotely-managed flow conditioning, with a turbidity observation being logged (almost) every 15 minutes via SCADA from the date of commissioning (March 2012) onwards. Some of the key events in the history of TM-WB regarding discolouration risk management are shown in table 6.13.

The water company responsible for TM-WB are currently managing discolouration risk using regular flow trials to condition the main to flows approaching $1600 \text{ m}^3 \cdot \text{day}^{-1}$ above typical demand. However, there is considerable interest and value in modelling how the amount of material on the pipe wall develops between trials, as this would allow the risk associated with postponing a trial to be quantified and permit the optimisation of the flow conditioning programme with regards to discolouration risk, energy, discharged water, fiscal cost and/or man hours.

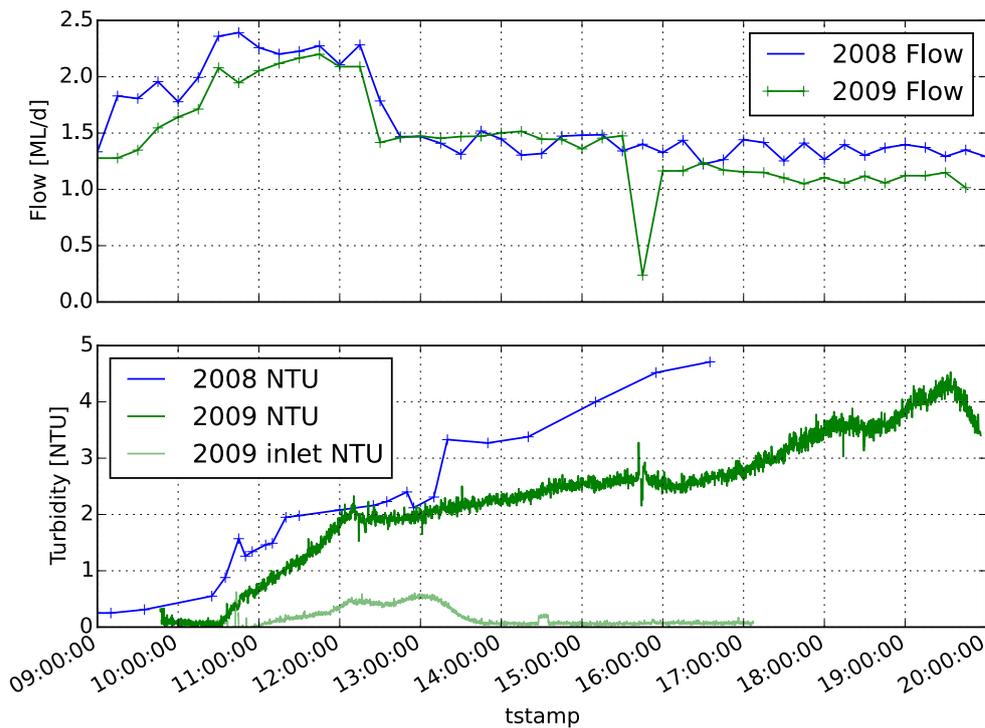


Figure 6.38: Superimposed turbidity and flow data for 6 km DI section from first two trials (after Husband et al., 2011). Trials were separated by 351 days.

Table 6.13: Key events in the history of trunk main TM-WB regarding discolouration risk. Note that during the flow trials turbidity was often recorded at a high temporal resolution but the flow measured by the meter at the downstream SR meter continued to be recorded only every 15 minutes (via SCADA).

Date	Comment	T_{ds_obs} monitoring
2008-08-06	Flow trial 0	Discrete sampling at SR2 sampling tap
2009-09-27	Flow trial 1	$\Delta t = 10 s$; N5 and SR2 sampling tap)
2010-12-07	Flow trial 2	$\Delta t = 900 s$
2010-12-14	Flow trial 3	$\Delta t = 10 s$
2011-01-18	Flow trial 4	$\Delta t = 11 s$
2011-02-01	Flow trial 5	$\Delta t = 10 s$
2011-06-21	Flow trial 6	$\Delta t = 11 s$
2012-03-08	Permanent turbidity meter commissioned for SR2 inlet	$\Delta t = 60 s$
2012-03-13	Permanent turbidity meter connected to telemetry system	$\Delta t = 900 s$ via SCADA from that date onwards
2012, mid-Sept	Main drained down to repair a leak. This permitted the inspection of the internal surface (fig. 6.39).	N/A



Figure 6.39: The TM-WB main was drained down during September 2012 to repair a leak. This allowed for the visual inspection of condition of the pipe wall. Part of the DI section is shown here.

6.3.5.2 Previous PODDS modelling

Previous attempts to model the first two TM-WB flow trials using PODDS were presented by Husband et al. (2011). Only the DI section was modelled as the AC was not considered to pose much discolouration risk. SR inflow data and turbidity data from the upstream end of the DI, N5, were used as model inputs. Parameters from a calibrated model of a different (steel) trunk main were used to 'bootstrap' the calibration process then parameter values were manually adjusted (not using PSO) to increase the similarity of turbidity observations and predictions.

Husband et al. (2011) found through trial and error that the model fit could be improved if the unlined DI section of TM-WB were split into three subsections ($846\text{ m} + 2510\text{ m} + 2731\text{ m} = 6087\text{ m}$), each with a distinct roughness or PODDS regeneration rate (previously mentioned in §2.2.8.5). Modelling the upstream-most subsection as having a higher roughness or non-zero PODDS regeneration rate in effect emulated there being more material at the upstream end of the main (given the nature of the PODDS model). No longitudinal variation in roughness could be found using pressure monitoring and fig. 6.39 suggests that there was little tuberculation, so the modelling results were believed to provide some evidence for discolouration processes varying along the length of pipes that are believed to have uniform characteristics (see also §2.1.16.6).

6.3.5.3 Objectives for modelling TM-WB using the VCDM

The principle objective was to model the impact of erosion and regeneration on the amount of material within the DI reach over a multi-year period. Field data has indicated that the discolouration risk associated with the DI reach is notably greater than for the AC reaches so only the 6087 m DI reach was explicitly modelled using the VCDM. This simplifies the modelling process as the DI reach, unlike the AC reaches, has a consistent nominal diameter along its length and is also assumed to be consistent with regards to flow (i.e. leakage is negligible and there are no take-offs) and roughness (Husband et al., 2011, could not detect variation in roughness with distance along the DI using headloss measurements). The modelled 6087 m pipe was assumed to have an internal diameter of 344 mm and a Darcy Weisbach roughness of 3 mm (values taken from a calibrated hydraulic model).

A second objective was to test the accuracy of the calibrated model by comparing turbidity observations and predictions for a period following the calibration period.

6.3.5.4 Data pre-processing

The following steps were taken to pre-process the time-series model input data:

1. **Import the flow and turbidity time-series SCADA data.** Express the flow data in $m^3 \cdot s^{-1}$. Replace turbidity data where the meter was obviously offline or reading exactly zero with null values. Clip the minimum flow to zero.
2. **Import per-trial, high-resolution downstream turbidity data.** Clip the minimum to $0 NTU$, interpolate over null values using time-aware linear interpolation, remove records with repeated timestamps, combine with T_{ds_obs} SCADA data to form a single time-series.
3. **Combine all flow and turbidity time-series as a single dataset** with a common time-series index (which is the mathematical union of the indexes of all combined time-series). This requires some up-sampling and results in null values being associated with some timestamps.
4. **Infill null values.** Fill forwards over one or more null values in the flow series. Infill missing turbidity values during trials and after the permanent turbidity meter was installed using time-aware linear interpolation.
5. **Estimate T_{net}** by subtracting the centred, rolling 5th percentile of T_{ds_obs} (with a time window of ten days) from T_{ds_obs} itself, as was also done for trunk main TM-NR (§6.3.4.3). The application of a centred, moving average function results in null values at the extremities of the output series: nulls at the end of the output series were infilled by filling forwards but the values at the start of the series were left null as the first part of the moving average input, T_{ds_obs} , was also null (for months before the first flow trial).
6. **Identify the times when increases in flow are likely to cause turbidity events** using the method presented in §6.3.4.2. Here τ_win was 28 days, single-data-point events were not ignored and $n_turnovers$ was 5.
7. **Estimate T_{net} .** As with TM-NR this could only be done using T_{ds_obs} (although T_{us_obs} data was captured during flow trials it was not available at the time of modelling/writing): the centred, rolling, 5th percentile of T_{ds_obs} was subtracted from the unprocessed T_{ds_obs} signal to give T_{net_est} . A rolling percentile window size of 10 days offered the best correlation between times when flow increases were likely to cause turbidity increases in TM-NR and observed turbidity increases (see §6.3.4.2 for details of the method). Prior to applying the moving average function, null turbidity values between trials were infilled with zeros to ensure that the removal of the background turbidity signal was not overly aggressive. Again, null values at the extremities of T_{net_est} resulting from the application of the moving average function were infilled by filling backwards and forwards using the first and last non-null values.

8. Calculate the shear stress profile.

The resulting dataset is shown in fig. 6.40.

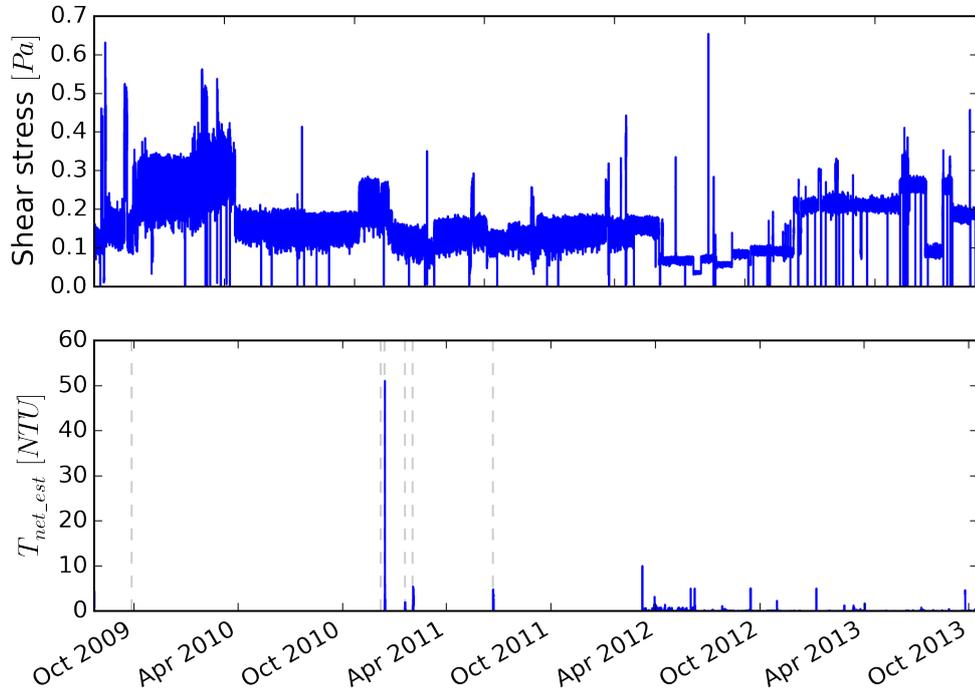


Figure 6.40: Shear stress and turbidity dataset from (the *DI* part of) trunk main TM-WB used for calibrating and validating the *VCDM*. The days on which flow trials 1 to 6 (see table 6.13) were conducted are shown as dashed lines. One can see that continuous turbidity data was available from March 2012, when a permanent turbidity meter was installed and connected to *SCADA*.

6.3.5.5 Model fitting method

The model was first fitted over a ≈ 1633 -day period up to the end of trial 6 (after which there were significant turbidity fluctuations due to the commissioning of the permanent turbidity meter; many of these were due to instrument error during/after installation rather than material erosion within TM-WB). Basic statistics describing this portion of the dataset are shown in table 6.14.

Table 6.14: Basic statistics of the part of the TM-WB dataset that was used for calibrating and validating the VCDM. Note that flows and turnover times relate to the DI section of TM-WB and that turnover times were only calculated for strictly positive flows.

	Q	$t_t(Q > 0)$	τ_a	T_{ds_obs}	T_{us_est}	T_{net_est}
count	187048	187048	187048	30342	187048	30342
mean	0.017	09:21:01	0.180	1.697	0	1.697
std	0.004	09:00:37	0.078	1.290	0	1.290
min	0.000	04:34:58	0.000	0.000	0	0.000
25%	0.016	07:52:37	0.134	0.519	0	0.519
50%	0.017	09:21:34	0.155	1.496	0	1.496
75%	0.020	10:03:38	0.217	2.668	0	2.668
max	0.034	109 days 03:07:31	0.632	51.066	0	51.0659

VCDM conditioning (§5.8.6) was used to estimate the wall state. The highest shear stress over the calibration period was near the beginning, some time before trial 1 (see fig. 6.40). It was therefore thought that by setting the initial amount of material per tracked shear strength to 1, the wall state would quickly converge on accurate values during this early period of elevated shear stress. When fitting and assessing the quality of fit, T_{net_est} and T_{net_pred} were only compared during trials; erroneous turbidity predictions during the ‘bedding-in’ period before the first trial therefore had no impact on the fitting process. During simulations, 100 strength bands were tracked. These corresponded to the range $[0, 0.654]$ Pa, with the upper bound being the maximum shear stress over the duration of the entire dataset (not just the fitting period).

The model was fitted to the dataset using the same PSO box bounds and convergence tolerances used when fitting the TM-YR dataset (see §6.3.3). Again, the FIT_SIMULT method was executed three times with the same inputs to assess the repeatability of the PSO fitting process.

6.3.5.6 Fitting results

The parameters identified by the three repeat runs of FIT_SIMULT were very consistent (table. 6.15) and the correlation with (slightly transformed) observations was reasonably good ($R^2 \approx 0.7$). The estimated duration required for full regeneration (β_r^{-1}) was shorter than the 1.5 years estimated for unlined Fe mains in DMAs (§2.1.14); this is not implausible given that regeneration processes in trunk mains have not previously been characterised and may have different kinetics to DMAs given the proximity to treatment works, different flow regimes, higher oxidant concentrations etc.

Table 6.15: Results of the attempt to fit the VCDM to the TM-WB trunk main dataset.

Run	β_e	β_r	α	ISE	PSO itrs	β_r^{-1} [yrs]	R^2
0	0.000341	3.5×10^{-8}	8.251	15965.608	104	0.905	0.709
1	0.000341	3.5×10^{-8}	8.250	15965.608	99	0.905	0.709
2	0.000341	3.5×10^{-8}	8.251	15965.609	109	0.905	0.709

Per-trial estimations and three predictions of the net turbidity (§5.9) are shown in figs. 6.41, 6.42, 6.43, 6.44, 6.45 and 6.46. For several trials, the discrepancies between observations and predictions appear to be a result of T_{us} not having been sufficiently discounted: the baseline T_{net_est} value before any flow increase is notably greater than 0 NTU. It may have been possible to achieve a better model fit had accurate T_{us} observations been available for use as model inputs, which would have negated the need to estimate T_{us} using rolling 5th percentile of T_{ds_obs} .

Another potential source of error is the timestamping of flow and/or turbidity data during trial 6: the temporal offset between turbidity observations and predictions is approximately an hour and the trial is the only one of the six that was conducted when the UK was operating on British Summer Time rather than Greenwich Mean Time, suggesting that the timezone of the flow or turbidity data timestamps is incorrect.

As with previous **PODDS** modelling (§6.3.5.2), the **VCDM** was unable to represent the unusual shape of T_{net} during trial 1, which may have been due to T_{us} observations not having been available for use as model inputs in this case, or may be due to regeneration processes not being consistent with length along the **DI** section of TM-WB. β_r , α , β_e or $\varphi(\tau, t_0)$ could potentially vary with longitudinal distance to give the impression that more material was mobilised from the upstream end of the **DI** reach (variation in roughness and diameter were previously ruled out (§6.3.5.2)).

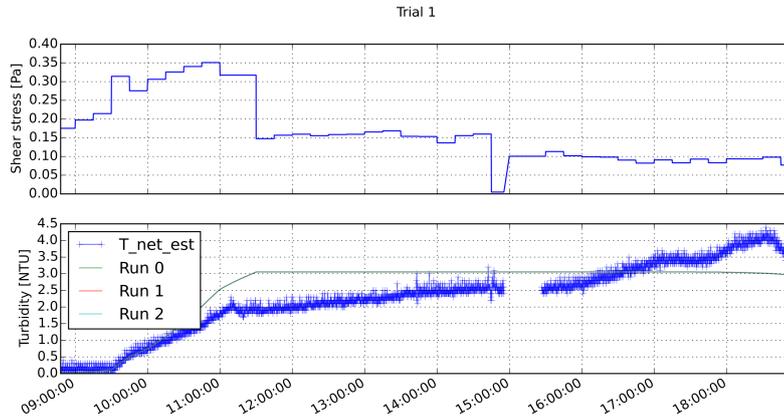


Figure 6.41: Trunk main TM-WB, **VCDM** fitting attempt, flow trial 1.

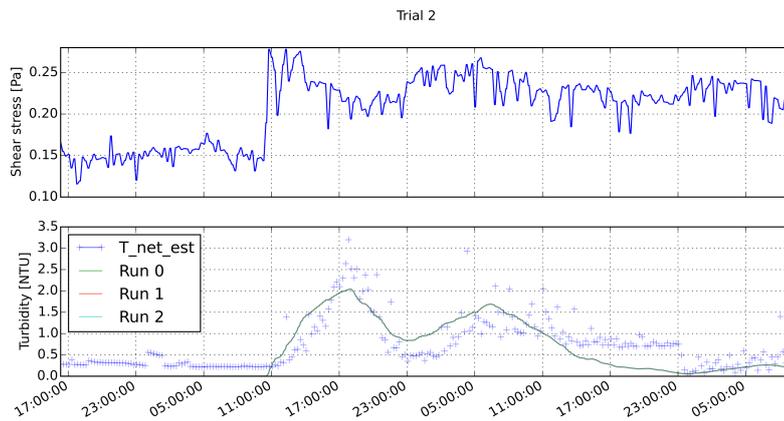


Figure 6.42: Trunk main TM-WB, **VCDM** fitting attempt, flow trial 2.

6.3.5.7 Accuracy of predictions using the calibrated model

The accuracy/validity of the calibrated model was explored by predicting the turbidity response over the part of the dataset not used for calibration (from after trial 6 on 2011-06-22 at 09:12:17 to 2013-10-17 at 13:00:00, which is approximately 848 days). The values of $\varphi(\tau, t)$ from the final timestep of the calibration simulations (see fig. 6.47; this varied by very little between runs) was used as the wall state boundary condition when simulating from 2011-06-22 at 09:12:17 onwards. Note that the relationship between shear strength and relative material quantity at that time was not trivial and not something that the **PODDS** model could represent; this gives weight to the **VCDM** design decision of allowing relative material quantity to vary freely (or at least monotonically increasingly) with shear stress.

The correlation between T_{net_est} and T_{net_pred} over this period was not as strong as for the

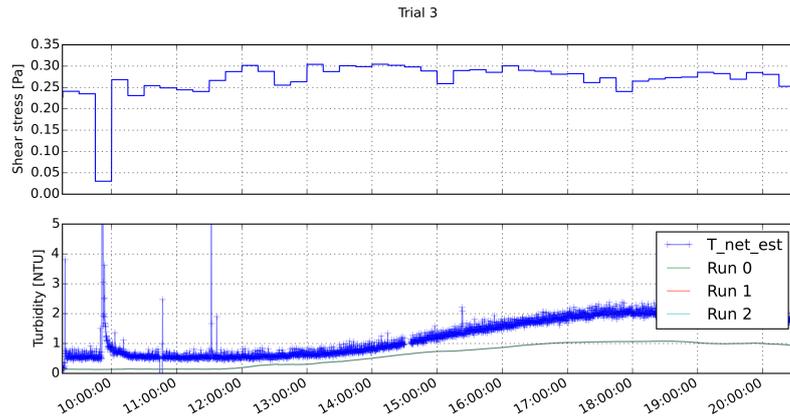


Figure 6.43: Trunk main TM-WB, VCDM fitting attempt, flow trial 3. The turbidity observations (and therefore the T_{net_est} profile) featured large turbidity spikes between 09:30 and 12:00: these are believed to be due to the stopping of flow (and settling/resuspension within the turbidimeter) and/or degassing rather than a discolouration response from TM-WB.

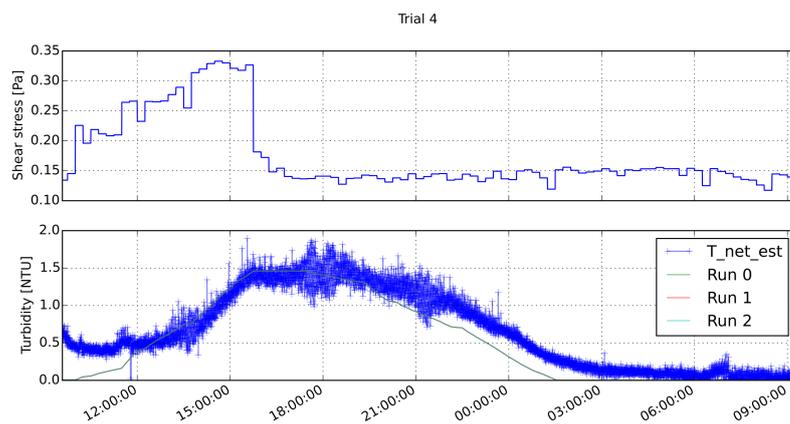


Figure 6.44: Trunk main TM-WB, VCDM fitting attempt, flow trial 4.

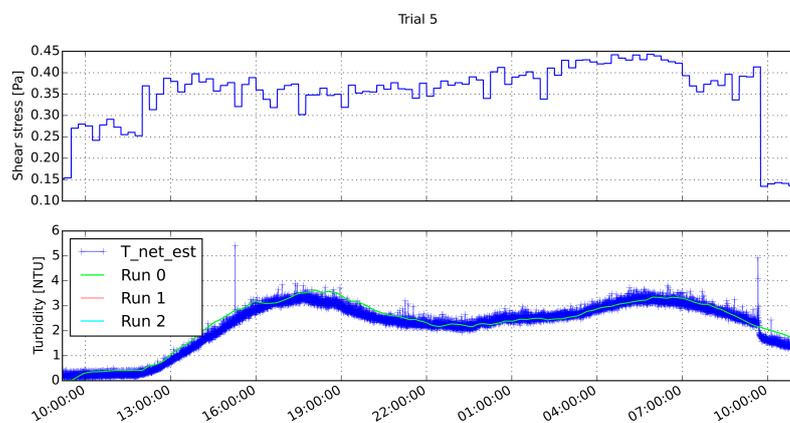


Figure 6.45: Trunk main TM-WB, VCDM fitting attempt, flow trial 5.

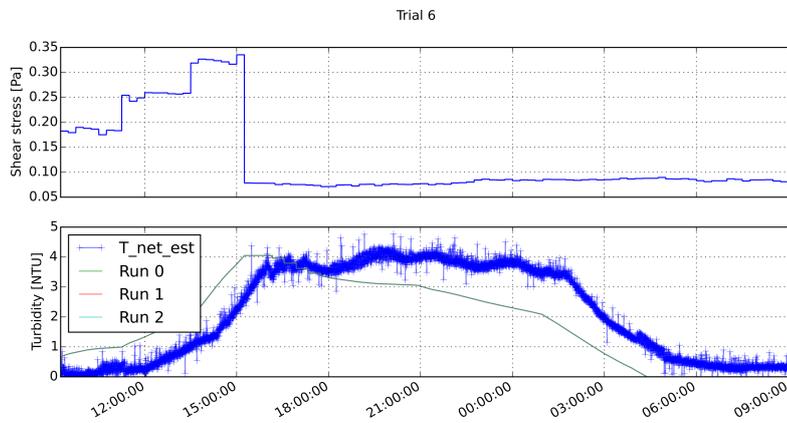


Figure 6.46: Trunk main TM-WB, *VCDM* fitting attempt, flow trial 6.

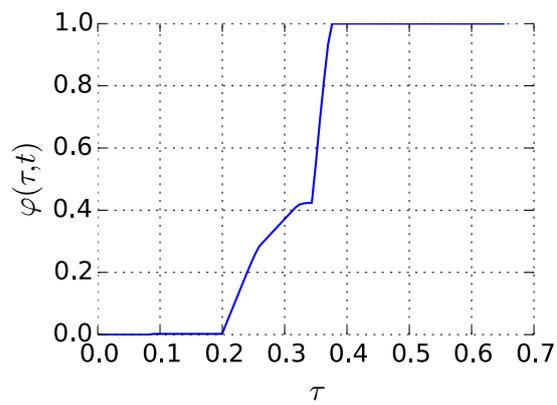


Figure 6.47: Trunk main TM-WB: the relative material quantity versus shear strength relationship at the end of the *VCDM* fitting attempt.

period used for fitting the model. R^2 was:

- 0.041 if only considering timestamps where T_{net_est} was not null;
- 0.015 if considering timestamps where T_{net_est} was not null and the aforementioned event detection method has flagged the timestamp as belonging to an event. It is surprising that this dissimilarity value is slightly lower.

Of the fourteen identified flow events where the increase in flow above prevailing levels was notable:

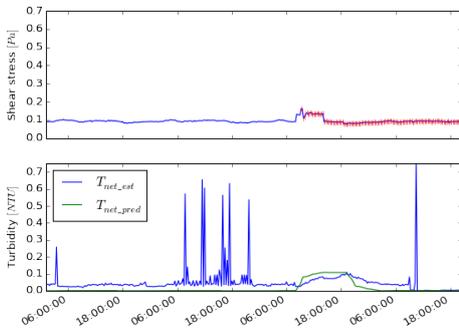
- 2 were fairly accurate (figs. 6.48a and 6.48b), including the very last event on 2013-09-24 (fig. 6.48b), ~824 days after the calibration period;
- 9 turbidity responses were over-predicted (e.g. on 8th, 18th/19th and 31 May 2012 in fig. 6.48c and 11th September 2012 in fig. 6.48d) , some significantly so;
- 1 was under-predicted;
- 2 predictions were too lengthy and low in magnitude (i.e. initially under-predicting but then over-predicting).

The T_{net_est} signal included two classes of feature that could not be explained by the model:

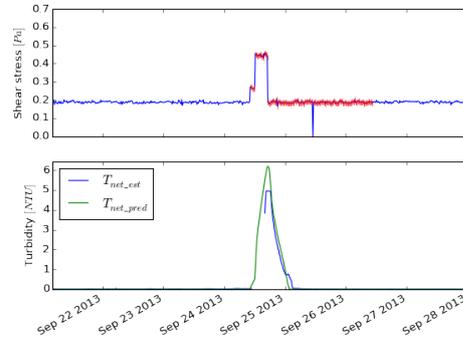
- Multi-point turbidity spikes at the end of a period when flow was temporarily zero (e.g. on 13th/14th September 2012 in fig. 6.48d); these spikes are most likely due to the re-suspension of settled material in the turbidity meter's flow cell;
- Clusters of single-point increases above the turbidity baseline with no corresponding flow variations (e.g. prior to the flow increase in fig. 6.48a and during 9th-16th May 2012 in fig. 6.48c), most likely due to measurement noise and/or degassing;
- Multi-point increases above the turbidity baseline with no corresponding flow variations, which could be due to the importing of discolouration material into the upstream end of TM-WB. Without turbidity monitoring data from the upstream end of TM-WB, it is not possible to test this theory.

6.3.5.8 Discussion and conclusions

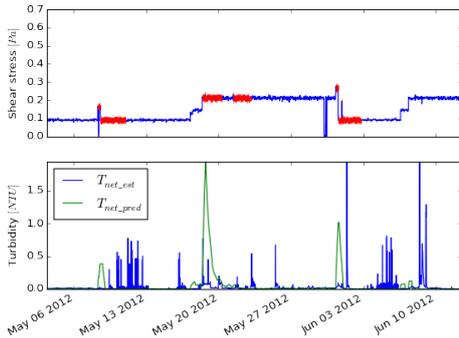
When fitting the model to six distinct flow trials in TM-WB, the FIT_SIMULT method proved to be highly repeatable, both in terms of identified parameter values and net turbidity profiles. The identified maximum regeneration duration of 0.9 years is plausible given understanding gained from previous field studies. It is believed that the quality of model fit, although reasonable, could have been improved had upstream turbidity observations been available for use as model inputs, and had timestamps been accurately recorded for all modelled flow trials. The turbidity response to the first trial corroborated Husband et al. (2011)'s theory that, immediately prior to that trial, discolouration process characteristics varied with distance along TM-WB.



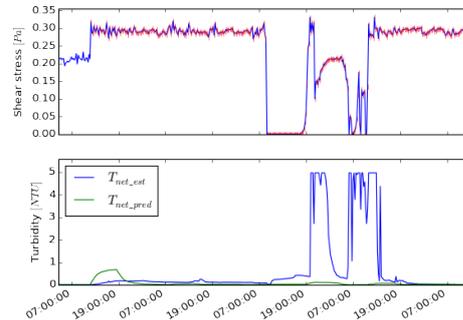
(a) Predicting turbidity in trunk main TM-WB: first strongly-correlating turbidity increase on 24th/25th April 2015. Note the T_{net_est} noise prior to the flow increase; this cannot be explained by the recorded flow observations and may be due to instrument error and/or degassing.



(b) Predicting turbidity in trunk main TM-WB: second strongly-correlating turbidity increase on 24th September 2013.



(c) Predicting turbidity in trunk main TM-WB: the turbidity responses on 8th, 18th/19th and 31 May 2012 were over-predicted by the calibrated model. The T_{net_est} noise during the period 9th-16th May 2012 cannot be explained by the recorded flow observations and may be due to instrument error and/or degassing.



(d) Predicting turbidity in trunk main TM-WB: the turbidity response from the flow event on 11th September 2012 was over-predicted by the calibrated model. Note the subsequent large spikes in turbidity (clipped to a maximum recordable value of 5 NTU) after the flow in TM-WB was raised above zero: these are most likely due to the resuspension of material within the turbidity meter flow cell and not to material erosion within TM-WB.

Figure 6.48: Predicting turbidity in trunk main TM-WB using the model calibrated during the fitting attempt. Red crosses are flow events, identified by an automated method, that are likely to cause material erosion.

Predicting forwards from the end of the calibration period yielded mixed results: the turbidity increases associated with nine of fourteen flow increases were over-predictions, although another two of the fourteen turbidity increases were reasonably accurately predicted, including one over *two years* after the end of the calibration period. Several issues made assessing the accuracy of the prediction more difficult: the resuspension of material within the turbidity meter flow cell following the restarting of flow in the main, rapid turbidity fluctuations with no corresponding flow variations and longer/larger turbidity increases with no corresponding T_{net} increase.

Upstream turbidity observations would have been useful for determining which turbidity increases were due to imported material. Their availability would also have negated the need to estimate T_{net} using the rolling fifth percentile of T_{ds_obs} , a process that could have yielded different results for the calibration period of short, discrete trials versus the subsequent period of continual monitoring due to the window size used with that moving average function. Any resulting difference in T_{net_est} characteristics between calibration and prediction period may explain why the predictions after the calibration period were often under-estimates.

6.4 Conclusions

- The assumption that certain discolouration process characteristics are constant with shear strength were tested using a realistic-scale laboratory pipe rig. The amount of accumulated discolouration material was found to be consistent for all but the weakest shear strengths (§6.2.2.8). An implication is that a calibrated model may over-predict the amount of material at weaker shear strengths or under-predict for greater shear strengths. It is not yet known whether this finding applies to all pipes, whether it is significant for discolouration risk management nor whether the erosion *rate* is constant with shear strength (this assumption could not be independently tested). However, it should be noted that the model was able to accurately represent many (but not all) turbidity events during four DMA case studies (erosion only; §6.3.2) and three trunk main case studies (erosion and regeneration; §6.3.3, §6.3.4, §6.3.5) whilst retaining the assumption that the three model parameters are constant with shear strength. Parameter values from fitted models differed somewhat between the two pipe systems, more notably for β_e . This could be due to recirculation effects and imprecise attempts to compensate for them, to fitting by hand rather than using PSO or to material eroding in one system faster than the other. There is insufficient information to prove any of these suggestions, but the third suggestion warrants further investigation as if the erosion rate of stronger material is dependent on the magnitude and fluctuation of a much lower shear strength during a preceding growth phase then this is in conflict with the VCDM assumption that the erosion and regeneration of different shear strength bands can be calculated independently.
- The developed model could represent three out of four DMA stepped flushing events as accurately as the validated PODDS model, with ability to fit the model to the fourth event believed to be limited by data accuracy. This indicates the chosen empirical erosion model is capable of generating results that reflect reality, even though internally it may not be based on physical principles.

- Three additional case studies explored the ability of the developed model to simulate both erosion and regeneration in trunk mains. The model fits produced using the FIT_SIMULT method were able to represent a number of (but not all) turbidity events over multi-month calibration periods in pipes of different materials and diameters and supplied with different source waters. The way in which well-fitted events and less-well-fitted events were interspersed indicates that the model is partially capable of representing discolouration processes. A number of data availability / quantity issues limited the ability to determine the extent to which the model was accurate / valid:
 - For trunk mains TM-NR and TM-WB no upstream turbidity observations were available so the net turbidity response (the signal to be modelled) had to be estimated by taking the centred, rolling 5th percentile of the downstream turbidity observations. This worked well for discounting imported material and instrument drift in the majority of cases. However there were several turbidity events (particularly in TM-NR) where the model could not represent the net turbidity estimate due to the latter featuring increases that appear to be the result of mobilisation upstream of the modelled mains.
 - In TM-NR the model under-predicted regeneration after a long period of relatively low flow. This could potentially be due to short, sharp flow spikes that were not captured or accurately represented by the fifteen-minute-averaged flow observations. Alternatively this could be due to the material at lower strengths having different properties in this case; if so, this is different behaviour to that seen in the pipe rig as less rather than more material was predicted at lower strengths. Another possibility is that the regeneration rate decreased during the period of low flow; however this is unlikely as flow was lower over summer, a time when one would expect regeneration to be quicker if it is a temperature-dependent process. To ensure that flow spikes are detectable in long term flow monitoring data it is recommended that the 95th percentile be recorded every fifteen minutes instead of/as well as the mean. Also, in future the extent to which the rolling 5th percentile of the downstream turbidity observations is representative of the upstream turbidity could be validated using data from one or more trunk mains that were continuously monitored at both ends.
 - The timestamps associated with the higher-frequency time-series data captured during several flow trials were found to be inaccurate. It is important that flow and turbidity timestamps are accurate as the model cannot represent observations if increases in turbidity appear to occur too early relative to the causal increase in shear stress.
- The model fitting process yielded more accurate predictions, more plausible β_r values and was more repeatable if the wall state at the time one wants to start modelling turbidity responses accurately was informed by more antecedent flow data.
- The quality of model fit was also dependent on the calibration dataset featuring multiple pairs of shear stress and corresponding turbidity events events, with these being greater than the baseline turbidity signal and the digitisation threshold. Ideally these events

need to be separated by a period of say three months so that a detectable amount of regeneration can occur over that period.

- Regeneration rates were plausible for the best fits of model to data from trunk mains TM-YR, TM-NR and TM-WB.
- The fitting process was aided by only calculating correlation metrics when turbidity exceeded a threshold (TM-YR), during shear stress events (TM-NR) or during flow trials (TM-WB). A method developed for identifying shear stress events (and accounting for discolouration material propagation delay) proved useful here.
- For trunk main TM-NR a fourth parameter was fitted using FIT_SIMULT to allow for temperature/seasonal variation in β_r . Interestingly, the best model fits for the assumed relationship between temperature and β_r were when β_r was effectively time-invariant and independent of temperature. However, this case study did demonstrate how the regeneration rate parameter could be substituted for a more complex sub-model, functionality that may become more important if future field/lab studies provide data that supports such complexity.
- Although the model of TM-WB fitted the calibration data well, attempts to predict just over two years forwards in time yielded mixed results, with the majority of the shear stress events during that time being over-predictions. However, an event near the end of that period was fitted very accurately so the prediction error might not be the result of a steady drift of the represented latent state but may instead be due to the model capturing some but not all of the key characteristics of discolouration processes.
- There was some limited evidence of discolouration processes varying with distance along TM-WB, although this variation has not been quantified.

Chapter 7

Discussion and recommendations for future work

7.1 Holistic appraisal of model performance

Details of the best calibration fits for all seven DWDS pipe lengths modelled in §6 are shown in table 7.1.

Overall quality of calibrations The model was able to represent DMA erosion events in unlined CI pipes well. The fits to erosion events in plastic DMA pipes were not as good, even if the maximum shear strength is explicitly limited, but the same was also true for the PODDS model (Husband and Boxall, 2010a). This indicates that one or more of the VCDM / PODDS erosion models and the $\tau_{max} / \tau_{ultimate}$ concepts need to be refined.

The best quality of fit metrics for the two medium-diameter trunk mains (TM-YR and TM-WB) were good and were notably higher than for the largest diameter trunk main (TM-NR). This could be due to the TM-NR dataset lacking higher resolution flow and turbidity data during flow events (i.e. 15-minute-averaged data under-representing peak shear stresses), to discolouration processes being slightly different in TM-NR and/or to the T_{net} estimation method being less effective for TM-NR. A useful test of the impact of using low-temporal-resolution data would be

Table 7.1: Details of the best model calibration fits for all DWDS pipe lengths modelled in §6. The first set of parameters for each pipe system is the result of trying to fit to the earlier part of the flushing phase and the second set from trying to fit to the later part.

Pipe(s)	Material	D [mm]	Fit config	β_e	β_r	α	β_r^{-1} (yrs)	τ_{max}	$\hat{\tau}_c$	R^2	NSEI
PODDS-CI1	Unlined CI	95	N/A	6.41×10^{-4}	N/A	0.593	Unk.	N/A	1.058	-	0.889
PODDS-CI2	Unlined CI	76 (51)	N/A	5.038×10^{-4}	N/A	0.585	Unk.	N/A	4.241	-	0.74
PODDS-PE1	PE	89	N/A	1.123×10^{-2}	N/A	0.504	Unk.	1.779	0.005	-	0.385
PODDS-PE2	PE	72	N/A	8.79×10^{-4}	N/A	0.544	Unk.	3.954	0.006	-	0.181
TM-YR	Concrete-lined DI	440	2	1.807×10^{-3}	6.137×10^{-9}	0.944	5.164	Unk.	N/A	0.930	-
TM-NR	Lined steel	800	1	7.650×10^{-5}	1.717×10^{-8}	5.472	1.846	Unk.	N/A	0.311	-
TM-WB	Unlined DI	344	N/A	3.41×10^{-4}	3.500×10^{-8}	8.251	0.905	Unk.	N/A	0.709	-
Rig sys A	HDPE	79.3	N/A	5×10^{-3}	N/A	0.038	Unk.	Unk.	0.09	-	-
				1.4×10^{-3}	N/A	0.022	Unk.	Unk.	0.09	-	-
Rig sys B	HDPE	79.3		1.5×10^{-2}	N/A	0.032	Unk.	Unk.	0.09	-	-
				1×10^{-3}	N/A	0.018	Unk.	Unk.	0.09	-	-

to log at a high rate such as once per 10 s then compare best model fits between that dataset and a subsampled version where each datapoint represented the mean of multiple datapoints in the original dataset.

Clearly, the accuracy and precision of the magnitudes and timings of time-series model inputs can be limiting factors for the quality of model fit.

Predictive power The prediction made using the calibrated TM-WB model did not correlate strongly with observations: many turbidity events were over-predicted but a small number of events were predicted accurately (note that it is preferable to over-predict turbidity events than under-predict). As this was only one prediction, there is insufficient evidence to dismiss the model as having no potential for extrapolation, testing scenarios or designing flow conditioning programmes. As stated in §6.3.5, the quality of prediction was limited by the method for estimating T_{net} (in the absence of T_{us} observations) being only partly effective and also by not being able to model resuspension within the instrument head after the restarting of flow. Other *possible* reasons for the low correlation of observed and predicted turbidity:

- The model was over-fitted to the calibration period.
- The model formulations including the regeneration sub-model may be insufficiently representative of reality (although the fit to the calibration period was fairly good).
- Discolouration process characteristics could have differed between the calibration period and the prediction period. For example, a change in WTW processes could have brought about a change in β_r .

For the prediction made using the TM-WB it was not possible to quantify how the error in the latent wall state increased over time. However, it is noteworthy (or possibly coincidence) that one of the two accurately predicted turbidity events occurred at the end of the two-year prediction, after a number of over-predicted events.

Plausibility of fitted β_r values The (fitted) annual regeneration rate for the TM-YR trunk main (§6.3.4) is 19.4%, which is far lower than the expectations for a DMA pipe with similar properties (not unlined Fe: 28%; supplied with surface water: 55%; Fe coagulation used: 49%; no unlined Cl upstream: 27% (Husband and Boxall, 2011, ignoring variances)). This could be due to the proximity to the supplying WTW (see §2.1.16.7 and §2.1.16.6), the pipe being relatively new so the concrete lining being in better condition, the diameter being much larger than for the pipes flushed by Husband and Boxall (2011) (§2.1.16) or to the dataset the model was fitted to containing little information about regeneration processes (§6.3.3.5), resulting in an inaccurate estimate of β_r .

The annual regeneration rate in TM-NR (54.2%) was very much in keeping with the expectations for DMA pipes also supplied with surface water (55%) and with some unlined Fe pipes upstream (54%). However, although TM-NR is purportedly lined, its regeneration was closer to the expectation for unlined Fe pipes (52%) than other types of pipes (28%) (although the low R^2 indicates that it may be inaccurate). One explanation for this is that the lining of this trunk

main has deteriorated. It should also be noted that here the regeneration rate is not atypical for DMA pipes, even though the internal diameter of TM-NR is approximately a factor of ten greater.

For TM-WB, the annual regeneration rate was found to be 110.5%, which is far faster than the expectations for DMAs (Husband and Boxall, 2011) and is particularly surprising given that the main is supplied with groundwater (§2.1.16.3). Through studying mains with fast regeneration rates such as TM-WB, it may be possible to better validate the assumption that material accumulation over a given shear strength range is bounded (see §4.3.3 and §2.1.14).

Ultimately, the sample of three trunk mains does not allow any strong conclusions to be drawn about if/how regeneration rates differ between trunk mains and DMA pipes. Calibrated VCDM models of a greater number and variety of trunk mains (e.g. large-diameter AC and PE trunk mains and pipes with diameters larger than 1 m) are required to further the understanding of how the potential factors listed in §2.1.16 influence regeneration rates.

Plausibility of fitted α values For trunk mains TM-WB, TM-NR and TM-YR, the best fitted α parameter was factors of ten, five and two (respectively) greater than for the four narrower flushed DMA pipes referenced in §6.3.2. The PODDS k parameter was shown to be conceptually equivalent to the VCDM α parameter (§5.7.5.4) and the absolute magnitude of k for calibrated PODDS models is also often larger, or at least more variable, for $D \geq 150 \text{ mm}$ (unpublished calibration parameters obtained from the PODDS project steering group¹). The best fitted α value is similar between the four DMA pipes, regardless of diameter and whether the pipes are PE or unlined CI. This could indicate that large diameter differences have a more significant impact on α than pipe material. However, the fit to PE pipes is poor so the α identified for those pipes could be inaccurate.

The chosen T_{net} estimation method (§6.3.4.3) has been shown to be only partly effective; any error resulting from this method may cause turbidity spikes to be over/under-represented, leading to error in the fitted value of α (the amount of material per unit pipe wall and per unit shear strength) and the other two parameters.

Plausibility of fitted β_e values The highest (fastest) fitted erosion rate factor was associated with one of the PE DMA pipes. There may be a valid relationship between diameter and β_e as, although the quality of the VCDM fits for DMA PE pipes was poor, the PODDS P parameter, which is conceptually similar and has the same units, was previously also found to be greater for narrow plastic pipes (unpublished calibration parameters obtained from the PODDS project steering group). Similarly, the lowest (slowest) β_e value was associated with the largest diameter trunk main, TM-NR, although it may be that the β_e values found through applying FIT_SIMULT to the TM-NR data are subject to error as β_e is a rate-like parameter and the sampling resolution was only fifteen minutes.

Plausibility of fitted τ_{max} values The best values of τ_{max} identified by fitting the model to flushes in PE DMA pipes PODDS-PE1 and PODDS-PE2 were lower than $\max(\tau_a)$ during the

¹<http://www.podds.co.uk>

flushes, indicating that the flushes may have removed all material from the pipe walls. These values were lower than the $\tau_{ultimate}$ values found by Husband and Boxall (2010a) when they calibrated PODDS models of the same two datasets: the VCDM τ_{max} values were 1.78 Pa and 3.96 Pa for PODDS-PE1 and PODDS-PE2 and the PODDS $\tau_{ultimate}$ values were 1.2 Pa and 1.8 Pa (§2.2.8.4). This indicates that τ_{max} and $\tau_{ultimate}$ may not be conceptually equivalent and adds to the uncertainty regarding the magnitude and factors of ‘cleaning’ shear stresses in DMA pipes (§2.1.15).

A reason for the poorer fits in plastic pipes could be that the maximum shear strength in either all pipes or just plastic pipes (τ_{max} in the VCDM; $\tau_{ultimate}$ in PODDS) is not a crisp value. A smoother transition is more plausible but it may be difficult to demonstrate this by experimental means. Also, the VCDM would need to be made more complicated if $\partial\varphi/\partial t$ were dependent on shear strength at higher shear strengths in plastic pipes. Another possible explanation for the behaviour of discolouration material in plastic pipes is that $\tau_{ultimate}$ and τ_{max} are conceptually invalid but their use allow models to be fitted to pipes where there was more material present at the weakest shear strengths, as was observed in the pipe rig experiment presented in §6.2.2 and is supported by visual interpretation of the results for plastic pipes shown in Husband and Boxall (2010a), fig. 6.14 and fig. 6.15. More material being present at the weaker shear strengths is not something that the VCDM can represent if $\varphi(\tau, t)$ is required to be a monotonically increasing function (§4.4.4). However, $\tau_{ultimate}$ and τ_{max} are conceptually consistent with the proven idea of self-cleaning, entirely dendritic networks of self-cleaning pipes (§2.1.15). A third option is that both of the other two possibilities are true: more material accumulates at the weakest shear strengths, at least in plastic pipes, and there is a shear stress above which more material can be removed from the wall, although this may not have been reached in the modelled DMA mains.

An implication of these hypotheses is that the material strength versus shear strength relationship may not be a monotonically increasing function (as had been originally assumed; see §4.22).

Plausibility of fitted τ_c values The fitted τ_c values were plausible (§6.3.2.8) but cannot be compared to the τ_c values used in the PODDS modelling presented by Husband and Boxall (2010a) as the latter are not available at the time of writing.

Repeatability The repeatability of the calibration process can be assessed in different ways for a given pipe:

- Across multiple PSO runs given the same boundary conditions and regeneration sub-model;
- Between different fitting configurations (using different boundary configurations and/or regeneration models).

In the first of those two cases, the FIT_SIMULT (PSO) fitting approach was repeatable for the two CI DMA mains flushes (§6.3.2), for trunk main TM-YR fitting attempt 2 (§6.3.3), for all TM-NR fitting attempts bar the third one (§6.3.4) and for trunk main TM-WB (§6.3.5). The PODDS-PE1 DMA flush case study and the FIT_SIMULT performance tests in §5.7.5.5 demonstrated the importance of running FIT_SIMULT multiple times given the same inputs to

account for PSO being a non-deterministic process that is likely but not guaranteed to find the global optimum. The calibration of PODDS-PE2 (§6.3.2) was not repeatable. This is believed to be due to data quality and potentially to assumptions regarding the maximum shear strength being overly simplistic.

The first attempt to fit TM-YR (§6.3.3) was also not repeatable, indicating that the choice of boundary conditions and how the wall state is conditioned at the start of the simulation impact on the ease of fitting. This is also true of TM-NR (§6.3.3): inconsistencies between the four fitting attempts demonstrate that the fitting process is sensitive to the choice of boundary conditions.

In this thesis the model was only calibrated for a small number of pipes (nine in total). A better understanding of the repeatability of the calibration process will come through fitting the model to a greater variety of distinct systems.

It should be noted that the fitted parameters from the two rig experiment pipe systems, A and B, should have been very similar but were not. This is discussed more in §7.5.

Model error Several mesh resolution constraints for minimising model error were presented in §5.3.9.7. The first of these, eq. 5.34, was that Δt should always be less than the turnover time; this was true for all pipe rig and field study models presented in §6. A second constraint was that the shear strength step, $\Delta\tau$, should be much smaller than the shear stress (eq. 5.3.7). This was found not to be true for any of the three trunk mains²: $\Delta\tau$ was less than the minimum strictly positive shear stress in all three systems. Another was that the timestep should always be less than t_d (eq. 5.32); this was true for the the models of all three trunk mains². Here $\max(\tau_a)$ was used instead of τ_e to simplify the calculation of t_d , which results in a slightly more restrictive constraint. The β_e used to calculate the t_d was taken from the best fit (table 7.1).

Other mesh resolution constraints were not tested because either they were nonsensical for cases where the most restrictive case was very restrictive for cases where τ_e approaches zero or because the constraints were not valid due to the attempts to base them on non-dimensional expressions being unsuccessful.

7.2 Value of the modelling framework including practical applications

7.2.1 Designing optimal flow conditioning programmes

Once validated, the model is anticipated to be of greatest use for designing flow conditioning programmes for trunk mains (§2.1.17.2; §2.1.17.3) due to the facility for modelling accumulation and erosion over periods of months/years. At present, flow conditioning is used to manage discolouration potential in a small number of UK trunk mains but these programmes have been developed using trial and error, have not been optimised with regards to metrics such as cost and energy and may not be flexible i.e. the system state may not be known/quantifiable if

²Not checked for the pipe rig nor the DMA flushing trials.

there is a change to the schedule. In this sub-section, the outline of a method is presented for designing optimal flow conditioning programmes using the [VCDM](#) (use case 6 in §4.2.1).

To begin with, the process requires a [VCDM](#) model instance that is representative of reality, either through it having been calibrated very recently or from the system state having been updated using flow data from after the end of the calibration period.

The modeller needs to have a good understanding of how flow is expected to vary due to typical operation in the coming weeks/months/years. For example, this may be a constant level, a daily pattern or a weekly pattern. Flows are expected to follow this pattern for long enough to warrant developing a flow conditioning programme. The modeller also needs an estimate of the largest flow magnitude that the modelled main may experience in the coming weeks/months. There is no need to anticipate the rise up to this elevated flow level; assuming an instantaneous increase up to this level gives a worst case estimate of discolouration risk. This flow magnitude may be:

- A function of the maximum pumping rate for rising mains;
- The maximum flow that may result from a burst. This could be estimated using the standard orifice equation ([Pudar and Liggett, 1992](#));
- The maximum pipe flow that would be required if a [WTW](#) with a specific outflow or a parallel main or [SR](#) were taken out of service;
- A maximum derived from a combination of the above.

This 'target flow' corresponds to the maximum shear stress and therefore the maximum shear strength that needs to be modelled. Ideally the model should have been calibrated over a period where the shear stress reached this level to ensure that estimates of the amount of material at higher strengths are accurate.

For a main conveying a diurnally-varying flow, a flow conditioning programme might take the form of [fig. 7.1](#). Given that the daily peak flow is short-lived, it may be necessary to boost that peak over several consecutive days during each flow conditioning programme cycle. Increasing the additional flow in steps up to the target flow level will yield smaller turbidity responses than if the target flow is achieved on all of those consecutive days. The principle aim when designing the programme is to guard against the worst case shear stress increase, which is an instantaneous increase up to the target level just before the start of the next flow conditioning cycle.

The design of this programme can be considered an optimisation problem. One of several ways of posing this problem is presented here. The optimisation variables are:

- t_s : The number of seconds into each flow pattern after which extra flow is added.
- t_e : The number of seconds extra flow is then added for.
- r_p : The number of flow pattern repeats in a flow conditioning programme cycle (see [fig. 7.1](#)).

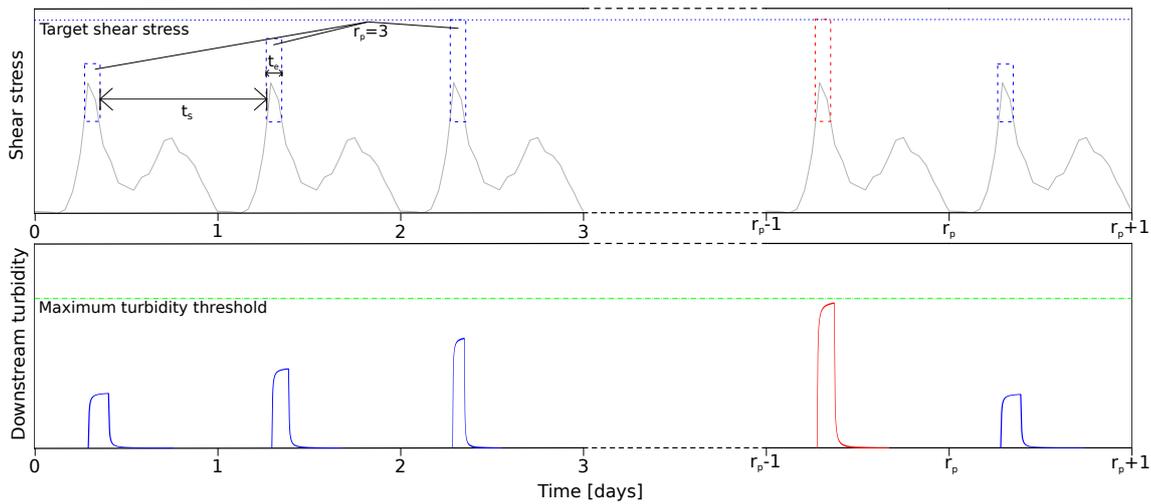


Figure 7.1: An illustrative design for a flow conditioning programme. The flow within the main varies diurnally, so additional conditioning flow must be added to the peak daily flow for each of several consecutive days. Additional conditioning flows are shown as blue dashed lines. The worst case erosion event that the programme is designed to protect against is the instantaneous increase in shear stress (red dashed line) up to the target level (green dotted line). The aim is to design conditioning programme so that the resulting turbidity response is always less than a threshold (green dashed line). Symbols are explained in the text.

- r_q : The number of flow pattern repeats at the start of that cycle when extra flow is added to the peak of the pattern e.g. 3 repeats of a daily flow pattern.
- Q_1 : The magnitude of the first extra flow period during each flow conditioning programme. The shear stress could be linearly increased to the target shear stress over r_q repeats of the flow pattern (e.g. over 3 days, as shown in the example).

The objectives are to:

- Minimise r_q .
- Minimise pumping/energy costs.
- Minimise the number of changes of pump status/speed. This minimises the risk of generating transients, pump wear (if applicable) and minimises man hours (if flow changes must be manually controlled, either locally or remotely using an actuated valve).

While the constraints are that:

- $r_q < r_p$
- There is no material with a strength less than the target shear stress bound to the pipe wall from r_q flow patterns into a flow conditioning cycle. This ensures that the flow conditioning programme is a statistically stationary process and is therefore repeatable.
- $\max(T_{ds}) < \text{a particular threshold}$ throughout the flow conditioning programme, even if the flow instantaneously increases to the target level during the last flow profile peak before the start of the next conditioning cycle.
- SR levels do not drop below minimum specified levels.

- $t_s + t_e <$ the flow pattern duration.

If instead the main of interest terminates at a [SR](#) then the optimisation problem can be expressed differently, as there is no need to ensure that the periods of elevated flow occur at the same time as the short-lived daily peak. Fig. 7.2 shows an example programme for a main with an effectively constant background flow. Such a programme could be designed by solving an optimisation programme where the variables are:

- s_q : The number of flow steps during each flow conditioning programme.
- t_q : The duration of each flow step relative to the turnover time.
- t_{rech} : The duration between flow steps, during which upstream and/or downstream [SRs](#) can be recharged.
- Q_1 : The magnitude of the first extra flow period during each flow conditioning programme. The shear stress could be linearly increased to the target shear stress over the s_q flow steps.
- t_p : Duration of a flow conditioning programme cycle.

The objectives are to:

- Minimise the combined duration of all flow steps and intervening t_{rech} periods.
- Minimise pumping/energy costs.
- Minimise the number of changes of pump status/speed.

The constraints are that:

- The combined duration of all flow steps and intervening t_{rech} periods is less than t_p .
- There is no material with a strength less than the target shear stress bound to the pipe wall after the combined duration of all flow steps and intervening t_{rech} periods. Again, this is to ensure that the conditioning programme is a stationary and so repeatable process.
- $\max(T_{ds}) <$ a particular threshold throughout the flow conditioning programme, even if the flow instantaneously increases to the target level just before the start of the next conditioning cycle.
- [SR](#) levels do not drop below minimum specified levels.

For all cases, a turbidity threshold of say 1 *NTU* would offer a reasonable margin of error below the typical regulatory limit of 4 *NTU* (§2.1.6). Note that in certain areas it may be necessary to use a lower turbidity threshold to prevent metals concentrations from exceeding regulatory limits; for example, [Cook et al. \(2015\)](#) estimated using linear regression that Fe would exceed the regulatory limit of 200 $\mu\text{g} \cdot \text{L}^{-1}$ for turbidities greater than 0.6 *NTU* (note that this not a

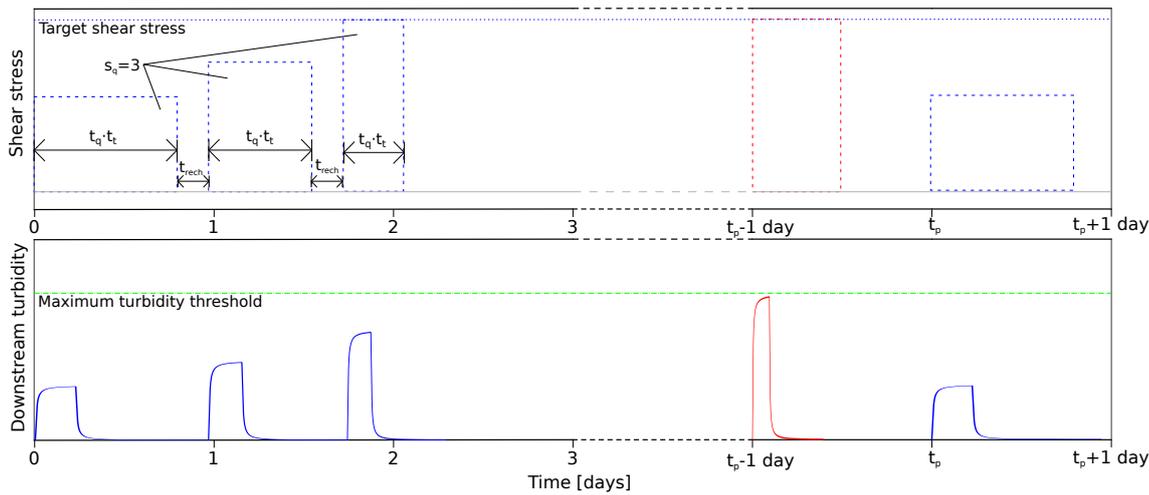


Figure 7.2: An illustrative design for a flow conditioning programme. Here the background flow in the main is effectively constant. The flow is periodically increased in steps (dashed blue lines) to remove material weaker than the target shear strength from the pipe wall. Gaps between these steps allow upstream and downstream *SR* to recharge. The worst case erosion event that the programme is designed to protect against is the instantaneous increase in shear stress (red dashed line) up to the target level (green dotted line). The aim is to design conditioning programme so that the resulting turbidity response is always less than a threshold (green dashed line). Symbols are explained in the text.

universal rule as this relationship depends on source water characteristics and the treatment scheme).

The optimisation problem is better attempted using a genetic algorithm (§2.2.4.2) than PSO as genetic algorithms are well suited to multi-objective problems with non-trivial constraints. A solution to the optimisation problem, plus capital costs associated with the commissioning of monitoring and control hardware, could inform an assessment of the whole life cost (§2.1.17.4) of the programme. This may be of value for comparing operational and capital expenditure between flow conditioning programmes and other asset management schemes.

Note that if *PODDS* was used for designing programmes like those presented here, then it would overestimate the amount of material mobilised from later, higher shear stress steps during each cycle (see §5.2.4).

Once a programme has been designed and is being implemented, accurate on-line turbidity monitoring would be useful for the continual validation of the model. If scheduled conditioning flow increases must be brought forwards or postponed, then the optimisation process could be re-visited to accommodate this.

Some water providers such as Northumbrian Water Ltd in the UK have started using real-time monitoring and automated control systems such as Aquadapt (Bunn and Reynolds, 2009) for maintaining optimal pumping schedules. Such systems continually solve complex optimisation problems and can adjust pumps and actuated valves to reflect the solutions to these problems. The design and automated execution of flow conditioning programmes could potentially be incorporated within such control systems for integrated network management.

7.2.2 Scenario testing and smart alarms

A real-time model, continually updated using flow data from [SCADA](#) and periodically validated using turbidity data, also has potential for providing real-time operational guidance and alarms.

Operators may want to know what the impact of a near-term future flow increase would be on bulk water turbidity (use case #5 in §4.2.1). Operators may also want to know by how much they could increase the flow in the near future without exceeding a turbidity threshold. This could be posed as a simple optimisation problem. Again, assuming an instantaneous increase in flow gives a worst case estimate. Alternatively, operators could be presented with a real-time operational management chart that demonstrates the maximum turbidity that will result from any number of instantaneous flow increases over a range (illustrative example shown in fig. 7.3).

Another use for near-term predictions using a calibrated [VCDM](#) models is for alerting operators when a recent increase in flow is predicted to cause the bulk water turbidity to exceed a threshold. Promptly decreasing the flow may then prevent this turbidity threshold from being breached. An illustration of this idea is shown in fig. 7.4. Here, the flow in the target pipe doubled at a particular time (subfig. a). If simulations are continually being performed using the current model state and extrapolations of flow, then the model may predict that the bulk water turbidity will exceed the regulatory limit (subfig. c). If this raises an alarm and operators are then able to reverse the flow increase within a short period (e.g. 15 minutes; see subfig. b), then this may prevent turbidities from increasing to those levels (subfig. d).

This method will be particularly effective in trunk mains with long turnover times, where there is considerable delay between an increase in flow and the peak turbidity after one pipe turnover. More sophisticated methods of flow extrapolation could be used e.g. using the first- or second-order derivative of the flow time-series.

7.2.3 Revisiting the modelling use cases

In §4.2.1, it was specified that the model developer should be able to characterise the sensitivity of bulk water turbidity (use case #1) and latent wall state (#2) to continual erosion and regeneration, then develop a time-series model of the pipe wall state and bulk water state (#3). The current understanding of the sensitivity of the material quantity versus strength to erosion

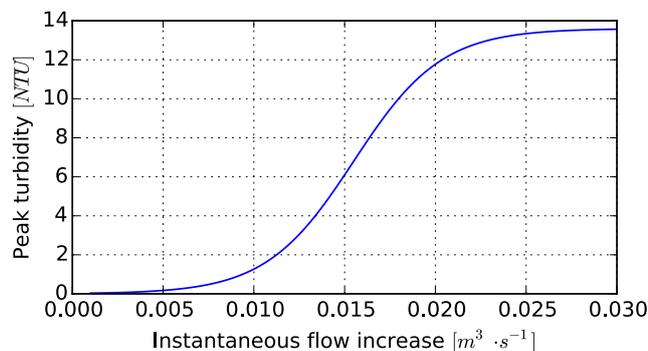


Figure 7.3: An illustrative operational management chart showing for a particular main and moment in time the impact of various instantaneous flow increases on the peak bulk water turbidity.

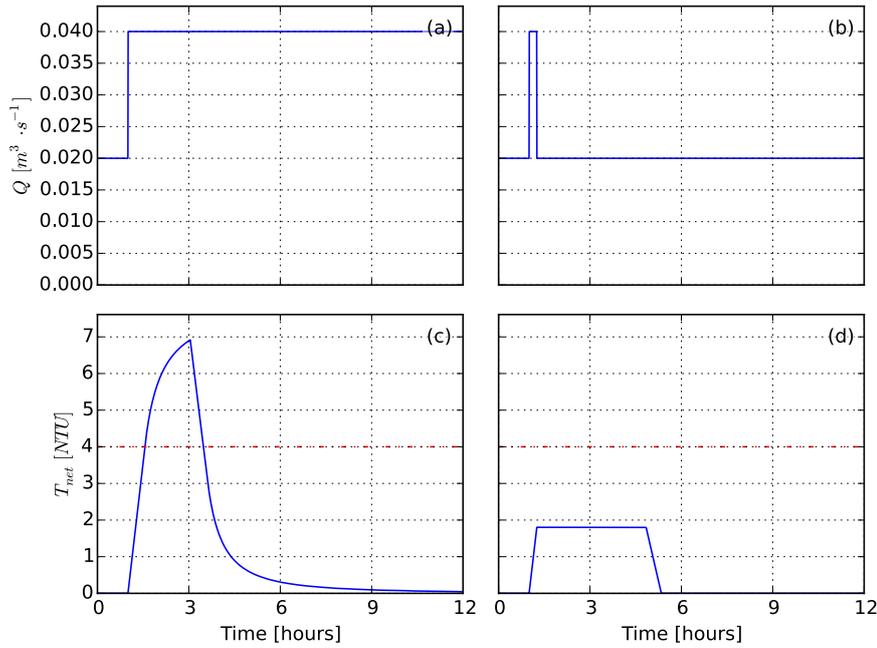


Figure 7.4: If, following a flow increase (a) the turbidity is predicted to exceed a threshold (c), operators may then be able reduce the flow (b) in response to a **VCDM** alarm. This could prevent the turbidity threshold from being exceeded (d) if the operators respond quickly enough (in this case < 1 hour). The example here was produced using the following **VCDM** configuration: $D = 250 \text{ mm}$; $k_s = 3 \text{ mm}$; $L = 6 \text{ km}$; $Q_{lo} = 0.02 \text{ m}^3 \cdot \text{s}^{-1}$; $Q_{hi} = 0.04 \text{ m}^3 \cdot \text{s}^{-1}$; $\beta_e = 0.002 \text{ Pa}^{-1} \cdot \text{s}^{-1}$; $\beta_r = 0 \text{ s}^{-1}$; $\alpha = 0.2 \text{ TPMU} \cdot \text{Pa}^{-1} \cdot \text{m}^2$; $\varphi(\tau < \tau_{a_lo}, t = 0) = 0$; $\varphi(\tau_{a_lo} \leq \tau < \tau_{a_hi}, t = 0) = 1$.

and regeneration was encapsulated in the **VCDM** conceptual model (§4.3) and associated, verified formulations and software (§4.4; §4.5). Sensitivity analysis showed that different parameters and physical quantities could have similar effects on the bulk water turbidity, hence the importance of modellers using an automated, multi-variate optimisation method when fitting the model to data. The modelling framework also offers model developers the ability to further the understanding of the sensitivity of discolouration processes to factors such as bulk water temperature and longitudinal distance: models can be fitted to different lengths of a network reach or different periods/seasons then parameters and model states compared between these models. Alternatively, β_r could be replaced with a sub-model that is temperature/season dependent to see if/how this improves model fit.

Use case #4 was that the modeller should be able to calibrate a model of a particular pipe and #5 was that he/she could then use this calibrated model to make predictions and test scenarios. Using the developed fitting process, accurate model fits were achieved for two out of four studied **DMA** flushing events (i.e. erosion but no regeneration) and two out of three long-term trunk main datasets (erosion and regeneration). This fitting process proved robust in certain cases but is sensitive to the availability and accuracy of upstream turbidity data, to the availability of information for informing the wall state boundary conditions and to the number, magnitude and spacing of hydraulic events during the calibration period. As shown in §6.3.5.7, functionality exists for taking a calibrated **VCDM** model then predicting forwards in time using measured or hypothetical flow data. For the one case where predictive capability was tested the overall accuracy was low but a couple of events were predicted accurately. Improved data quality and more granular regeneration sub-models may result in more accurate predictions.

The final use case, #6, was the modeller is able to design optimal flow conditioning programmes. A framework for this is presented in §7.2.1 that, once the model has been shown to be capable of making accurate predictions, would allow discolouration potential to be managed over extended periods whilst minimising costs.

7.3 Comparing the VCDM to other discolouration models / risk assessment methods

The VCDM was able to represent isolated erosion events as well as the PODDS model when fitting the erosion parameters and fitted, polarised boundary conditions (§6.3.2). Also, the VCDM was also able to closely approximate various artificial PODDS predictions in §5.7.5.5. The model fits presented in §6.3.2 and §5.7.5.5 were achieved without the need for an equivalent to PODDS' n parameter. In §5.6.6, sensitivity analysis showed for a step increase in shear stress the VCDM erosion parameters can transform scale and shape of the turbidity response in a similar way to n . The VCDM not needing an equivalent power term is an attractive quality as the PODDS model was very sensitive to n and its inclusion in PODDS prevented dimensional analysis (§2.2.8.5). Another advantage of not requiring an n equivalent is that erosion can be modelled using two empirical parameters rather than three. However, these two parameters, α and β_e , cannot always be determined independently (see §5.7.5.3).

Both PODDS and the VCDM have means for representing the material bound to the wall having a maximum shear strength. With PODDS this is achieved by specifying C_{max} (then calculating $\tau_{ultimate}$ from it using k ; see §2.2.8.4), with the VCDM this is done by limiting the upper bound of the tracked shear strength range, τ_{max} . The VCDM is more flexible as it offers a cleaner separation of quantities that code for the maximum amount of material on the pipe wall and the maximum shear strength.

It is recommended that the VCDM be used in preference to PODDS for modelling erosion events as the VCDM is as representative of reality when using polarised boundary conditions and $\beta_r = 0 \text{ s}^{-1}$ but has fewer parameters. In addition, the VCDM has a regeneration model that reflects characteristics observed in numerous field studies so it has potential for long-term modelling, although the predictive capability has presently only been tested for one case study and has not yet been validated.

If all material in a shear strength range $[\tau_1, \tau_2]$ is eroded after which there is a period of regeneration then:

- The models will predict the *same turbidity response* for each of a number of subsequent flow steps if the regeneration period is long enough for all to have reached maximal levels before the flow increases again.
- The models will predict *differing responses* for the step increases if the corresponding shear strength bands are only partially regenerated at the end of that period, with the behaviour of the VCDM more closely reflecting observed regeneration process characteristics.

Since the start of this PhD, both PODDS and the VCDM have now been implemented within the (commercial) Synergi DWDS modelling software (DNV GL, 2010), making it easier for

companies to switch to using the [VCDM](#) and saves having to convert models from commercial formats to EPANET formats for use with [PODDS](#). However, at present the FIT_SIMULT calibration is not implemented within Synergi, so calibrating [VCDM](#) models using Synergi may be non-trivial.

Too little is known about the [PSM](#) model (§2.2.8.6) to draw many comparisons. If the proposed van der Waals wall attraction model is ever implemented then it may provide similar functionality to the [VCDM](#).

The [DPM](#) (§2.2.8.6) provides an automated means for exploring the impact of hydraulic scenarios (failures and pipe 'closures') on turbidity, with turbidity responses being predicted using [PODDS](#). The [DPM](#) could be of greater value if coupled to an accurate [VCDM](#) model that is continually being updated using flow data from [SCADA](#). The [DPM](#) could continually and exhaustively calculate the impact of closures/breaks in any pipe in an 'all-mains'³ model on a subset of those trunk/critical pipes for which 'on-line' [VCDM](#) models were available (calibrating the [VCDM](#) for all pipes in an all-mains model would be infeasible).

7.4 Ideal field test for validating and demonstrating the advantages of the new model

There is relatively little high-quality data available for validating the [VCDM](#) and proving that it can represent material regeneration more accurately than [PODDS](#). Presented in this section are requirements for an ideal field test that should clearly demonstrate the differences between the two models. The pre-requisites for the test are as follows:

- The pipe length should be homogeneous with regards to pipe wall material, diameter and age (and therefore relative roughness). There should be no take-offs, so the pipe is also homogeneous with regards to flow and shear stress (assuming negligible leakage). This allows the [VCDM](#) to be tested with the simplest system configuration: a single pipe length over which the model parameters are hypothesised to be constant
- Flow monitoring in the pipe is via an in-pipe meter (not via a standpipe and hydrant/wash-out, as has been done during [DMA](#) flushing fieldwork (§6.3.2)) and so provides absolute measures of flow and is accurate (even at low flows).
- Reservoirs upstream and downstream of the pipe length allow the flow in the pipe to be controlled independently of diurnal variation in customer demand and allow the flow to be raised considerably above typical levels. There are no significant constraints on how/when water is transferred from the upstream to the downstream reservoir. Reservoir capacities are such that they can be considered just physical boundaries rather than constraining elements with finite capacities. These pre-requisites offer flexibility when designing a field test flow profile. Flow control should be facilitated by a [FCV](#) to ensure changes are prompt and precise.

³'All mains' models feature all [DWDS](#) pipes in a certain area bar customers' service pipes i.e. trunk and [DMA](#) pipework.

- Historic flow data should be studied to determine the typical operational maximum, Q_1 . A higher but achievable 'target' flow level, Q_2 , should be identified through hydraulic modelling of the system. The shear stresses at each of these levels (τ_1, τ_2), along with all subsequent shear stress estimations should be calculated using a roughness value derived from headloss measurements over the reach A to B during pre-test flow increases.
- The pipe length should have access points (e.g. hydrants, wash-outs, quadrina points) to allow for turbidity monitoring.
- Two of these monitoring points, A and B , should be separated by enough distance for the corresponding turnover time at Q_2 in that main to be much larger (e.g. a factor of two larger) than the logging timestep (§5.3.9.7). Requirements for the timestep are specified below.
- There should be few if any fixtures/fittings between A and B that might cause local accumulation or local headlosses (§2.1.16.6).
- Turbidity should be logged at each of A and B . This should be done using instruments of the same make/model using the ISO 7027 measurement standard (§2.1.5) and the same logging discretisation regime. Both should be well-calibrated before the start of the test.
- The timestamps of all logged data are unambiguously expressed in terms of Coordinated Universal Time (UTC) in the ISO 8601 format. All loggers should have synchronised clocks.
- The shear stress step should be much less than the largest excess shear stress (e.g. $< 0.01 \cdot \max(\tau_e)$) as indicated by model sensitivity analysis (§5.3.9.7). Here, τ_2 should be used as an approximation of $\max(\tau_e)$.
- If the aforementioned criteria regarding the timestep is satisfied, there is no justification prior to starting the test for the logging frequency to be any greater than once every 10 s, which is a common, achievable logging frequency during discolouration field trials. Flow and turbidity should be logged at the same temporal resolution and at the same timestamps. The flow data should be the 95th percentile rather than the volumetric mean of each timestep to ensure that peak shear stresses are represented in the aggregated data.

Given these prerequisites, the test can be conducted as follows:

1. Elevate the shear stress at the pipe wall to τ_2 , thus removing all material at the pipe wall with a shear strength in $[0, \tau_2]$.
2. Decrease the shear stress to a lower level, τ_1 , and wait for the material with shear strengths in $[0, \tau_2]$ to partly regenerate. This period should be long enough so that there is significant regeneration but to distinguish the **PODDS** and **VCDM**, it should not be so long that maximal regeneration has occurred. Also, the test is simplest if the regeneration period is short enough for the bulk water to experience little seasonal

variation in temperature, allowing β_r to be considered constant with time. As it was estimated that material can fully regenerate in as short a period as 0.9 years (table 7.1; §2.1.16.3), the regeneration period should be three months. Over this period, the shear stress can fluctuate but should not be permitted to exceed τ_1 .

3. At the end of this period, elevate the flow in steps where each is sustained for long enough for all material with a shear strength lower than the corresponding shear stress to be eroded and pass out of the modelled pipe. The first step should increase the flow to just above τ_1 , if the shear stress was lower than that at that time; this ensures that prior to the remaining flow increases that all material has been regenerating for the same duration and has a known minimum strength. Four subsequent flow increases up to Q_2 will provide measures of the perceived amount of accumulation and the rate of erosion for each of four contiguous shear strength bands.
4. Attempts should then be made to fit the **VCDM** and **PODDS** to the logged data using the FIT_SIMULT method (§5.7.5.1). The models should be fitted from the start of the regeneration period where the shear strength profile of the wall-bound material is known, allowing model boundary conditions $\varphi(\tau, t_0)$ (**VCDM**) and τ_c (**PODDS**) to be estimated with confidence. The models should be fitted to T_{net_obs} , which can be estimated using the upstream and downstream monitored turbidity data, the logged flow data and a Lagrangian transport model (§5.9).

Under **PODDS**, the weakest-bound material regenerates before the stronger-bound material (2.2.8.5) yet the **VCDM** reflects the observation that all shear strengths appear to regenerate simultaneously. Therefore, the automated FIT_SIMULT method will not find a combination of **PODDS** model parameters that results in as good a fit as can be achieved using the **VCDM**. The test described above could be repeated for different pipe materials and source waters to confirm that this result is universal.

This test could also be useful for investigating:

- The impact of using 15-minute-averaged data as model inputs and calibration references, by subsampling higher frequency turbidity and flow signals (§7.1);
- Whether the assumption that accumulation reduces the bulk water turbidity is valid, by studying the discrepancy between the downstream turbidity measurements and the upstream turbidity measurements (lagged using a Lagrangian transport model to account for material advection along the main) under non-eroding conditions.
- The experiment presented in §6.2.2 could be repeated in a non-recirculating system to generate further evidence for more material appearing to have accumulated at weaker strengths over the flushing phase (§6.2.2.8) and $\varphi(\tau, t)$ not increasing monotonically with τ (§7.1).

7.5 Revisiting the model assumptions and future research opportunities

The conceptual model assumed that all three **model parameters were time-invariant**. In three trunk main studies, the model was accurately calibrated ($R^2 > 0.7$) for two out of three mains whilst retaining that assumption. However, the limited quality of the best fit for TM-NR and the limited prediction for TM-WB could indicate that allowing the regeneration rate to vary with time may lead to improved accuracy/performance. The way in which β_r can be substituted for a (slightly) more complex, time-varying expression was demonstrated in §6.3.5, but the form of the modelled relationship between regeneration and atmospheric temperature was possibly too simple as it didn't improve the calibration and the relationship between atmospheric temperature and bulk water temperature is known to be non-trivial (Blokker and Pieterse-Quirijns, 2013). Quantifying the regeneration rate for a greater variety of systems through VCDM modelling could further the understanding of how the many factors listed in §2.1.16 influence regeneration rates. This in turn could inform the PSO bounds when fitting. A more directed approach could be to use machine learning techniques such as EPR (§2.2.8.7) to learn from paired flushing data not only the likely parameters but also the *structure* and *key factors* of an improved regeneration sub-model. However, this is weighted towards narrower distribution pipes as there is presently insufficient information on regeneration in trunk mains for either data-driven analysis or determining the degree of commonality with regeneration in DMAs. Another way in which the regeneration sub-model could be refined would be to drive it using Blokker and Pieterse-Quirijns (2013)'s bulk water temperature model, which would account for heat transfer from subsurface to advected water (§2.1.16.2). Blokker and Pieterse-Quirijns (2013) found that bulk water temperature was more influenced by residence times than atmospheric temperature, so making β_r a function of the model output could introduce longitudinal variation in regeneration rates for lengthy trunk mains. Any increase in the complexity of the regeneration model must be justifiable as it will make the model more difficult to fit and increases the risk of over-fitting.

The assumption that the three **model parameters were constant with shear strength** had to be rejected (§6.2.2.8), although it has yet to be demonstrated that the differing behaviour of the weakest-bound material is significant for the management of real DWDS. Further laboratory studies involving multiple regeneration periods (see fig. 5.31) are required to determine which of β_r and α (or both) vary with shear strength. Also, it would be useful to know if the finding that more material accumulates at the weakest shear strengths is reproducible in pipes of different diameters (e.g. large trunk mains) and materials (e.g. AC and unlined Fe). Thirdly, further experiments could confirm whether more material accumulates within an absolute shear strength range or whether more material accumulates at shear strengths just above the prevailing shear stress.

The fitting (by hand) of the VCDM to the data collected for each of rig subsystems A and B resulted in two different parameter sets, with β_e being more noticeably different than α (β_r was not fitted). Possible causes are (a) recirculation effects and attempts to compensate for them impacting on signal quality, (b) fitting by hand rather than using PSO and (c) the growth phase shear stress profile, which differed between the subsystems, resulting in different characteristics

at higher shear strengths (§6.2.2.8). This hypothesis, if true, would conflict with the the assumption that the **regeneration and erosion history of weaker shear strengths does not directly influence the current state/behaviour at greater shear strengths** (§4.3.5, §4.4.2). Further experimental evidence should be captured to determine whether this experimental result is repeatable. It is recommended that this evidence is produced by re-running the experiment described in §6.2.2 to save having to find two identical pipes as part of a field test (such as that in §7.4).

Two assumptions particular to the erosion sub-model were that material **erodes at a rate that is constant with time** but **variable with shear stress** (with weaker material eroding more quickly than stronger material). It is non-trivial to validate these assumptions and they may not accurately reflect the underlying physical processes. However, the quality of fits for the majority of modelled DMA erosion events and trunk main datasets suggest that they may be sufficiently valid.

Another core assumption was that the **shear strength of a quantum of material is time-invariant**. This is unlikely to be true as, for example, biofilms undergo a cycle of primary adhesion, secondary adhesion and death and/or detachment (§2.1.10). An interesting study could be to develop an entirely distinct wall erosion/regeneration sub-model that more closely represents this cycle, then compare the accuracy and ease of use to the VCDM. This alternative approach could model material attaching with a certain strength then the strength of that material changing over time, functionality that has the potential to be considerably more complex than the VCDM if material quantities, initial binding shear strengths and current shear strengths are all tracked over time.

Both the VCDM and PODDS make the assumption that **turbidity is always proportional to the intensity of light scattering** yet the ISO 7027 and USEPA 180.1 measurement standards only consider the relationship between the two to be linear over the range [0.012, 40] NTU (§2.1.5). For some of the case studies presented in §6, the maximum observed turbidity exceeded the upper limit of this range, these being TM-WB (51 NTU; §6.3.4) PODDS-CI1 (82 NTU; §6.3.2), PODDS-CI2 (600 NTU; §6.3.2), and PODDS-PE2 (92 NTU; §6.3.2). The turbidity in TM-NR over the studied period may also have exceeded 40 NTU but the recorded values had a maximum of exactly 5.0 NTU, indicating that instrument output may have been clipped to a maximum of 5 NTU and the true maximum over that period was not recorded. The linearity/non-linearity of turbidity and light scattering does not explain why more material was detected at weaker shear strengths (§6.2.2.8), why some pipes seem to have an ultimate shear strength (which may be the same or a related phenomenon; §2.1.15 and §6.3.2) nor why the VCDM and PODDS do not fit data from plastic pipes as well as for corrodible pipes:

- The model fit was quite good for pipes PODDS-CI1 and PODDS-CI2 where the maximum turbidity was > 40 NTU;
- The model was unable to accurately fit the turbidity data from all shear stress steps in the two pipe rig subsystems (§6.2.2), even though the maximum turbidity was < 40 NTU
- The VCDM and PODDS model fits for PODDS-PE1 were poor yet the maximum turbidity < 40 NTU.

The significance with respect to the **VCDM** of turbidity potentially being non-linear with light scattering at higher turbidities is therefore not known.

On a related matter, **turbidity was assumed to be linear with TSS**. The results presented in §6 do not explicitly contradict this. However, if the constant of proportionality differed between the weaker wall-bound material and the stronger-bound material, due to for example biofilm composition and so potentially light scattering properties varying with depth (Fish et al., 2015), then this might explain more material being detected at weaker shear strengths in §6.2.2 and/or also might partly explain why the amount of material in plastic pipes appears to tail off as one approaches higher shear strengths (§2.1.15). Turbidity was selected over **TSS** as the model output as it more closely relates to customer perceptions, although the model would be no more complex if it predicted **TSS** rather than turbidity (but it would be more difficult to collate calibration data).

Further work could improve the understanding of the relationship between turbidity, **TSS** and shear strength. The turbidity versus **TSS** relationship has been quantified through much fieldwork yet these experiments have not focussed on whether this relationship is constant over a series of contiguous shear strength bands. To investigate this, complementary turbidity and **TSS** measurements could be taken during each step of a stepped flushing experiment. The pipe length would need to be long enough to capture enough water and material per shear stress step for the sample drying and weighing required for **sTSS** measurement to be accurate. Ideally, particle count measurements (see §2.1.5) should also be taken during each flow step to confirm that the **PSD** is consistent over the flushing experiment, as light scattering and therefore turbidity are a function of particle diameter (§2.1.5). Although there are numerous advantages to performing stepped flushing experiments in a laboratory rig (§6.2.2.2, it may be preferable to conduct this experiment in a real **DWDS** where there is no recirculation of material as extending the background turbidity preprocessing described in §6.2.2.3 To **TSS** data and count data for multiple particle size bins may be non-trivial. If turbidity is found to be non-linear with **TSS** over a wide range of shear strengths then the parameter that codes for turbidity/light scattering, α , may need to be made a function of τ in the **VCDM**.

Discolouration processes were assumed to be **homogeneous with longitudinal distance**. This could not be validated for the presented case studies due to the lack of intermediary sampling (there was an intermediary sampling point in TM-YR but the data obtained from this location was limited and of poor-quality). If the assumption is invalid then this will be most significant for trunk mains, the pipes of greatest concern, due to their length. It may not be a coincidence that, for the three trunk mains modelled in §6, the quality of the best fit negatively correlated with pipe length. Further field studies could qualify and quantify longitudinal heterogeneity by calibrating a **VCDM** model of each of several contiguous lengths of a trunk main and comparing fitted parameter values between the lengths. This would require a trunk main of constant diameter and material with no take-offs to be monitored at least three (ideally more) locations. In addition to turbidity, pressure should also be sampled for the purpose of testing whether the headloss and therefore the roughness varies between monitored lengths. If longitudinal heterogeneity is detected, then the significance of this needs to be quantified. If it is significant then the developed modelling approach should be augmented, either by adapting the model formulations (which may make the model much more complex, or, more pragmatically, by

finding the optimal number of daisy chained sections a pipe needs to be modelled as to provide a sufficiently accurate result.

Material accumulation was assumed to have a **maximal level**. In theory, it may be possible to validate this assumption by studying a fast-regenerating main such as TM-WB, but in practise it may not be possible to control the necessary factors over a sufficient duration to conduct such a test.

The model was able to represent discolouration processes using just steady-state hydraulics. However, the **unsteady shear stress** has the potential to influence shear stress and previous experimental results have showed that the bulk water turbidity increased due to transient pressure waves (§2.1.12.3). To date, no attempts at modelling erosion using the unsteady shear stress have been successful and experiment work has not quantified both transient pressure waves and the resulting impact on material at the wall. There is therefore an opportunity to characterise that relationship through controlled laboratory/field work, which in turn could lead to improved understanding of how (in)significant unsteady hydraulics are to discolouration modellers.

The aim of this thesis has been to model the discolouration risk associated with pipes i.e. it has been implicitly assumed that impact of other classes of assets on bulk water turbidity is negligible. However, a calibrated **VCDM** model instance of a trunk main that supplies a **SR** might overestimate the turbidity at the customers' tap as the **SR** will provide some attenuation. A study of how discolouration material **propagates through SRs**, considering both residence time distributions and settling velocities, could be useful for determining if, when and how **SRs** should be treated when modelling discolouration in **DWDS**.

The **VCDM** uses the same material transport model as **PODDS**, which was developed for use in **DMA** pipes. However, given the current interest in operational management strategies for trunk mains up to 1.8 m in diameter (Husband et al., 2010b), it may be necessary in future to revisit some of the associated assumptions. In very large diameter trunk mains, the effects of **dispersion and diffusion** may no longer be negligible and **lateral mixing** may not be instantaneous (§4.3.6.3). Also, the prevailing bulk velocities in such mains may be sufficiently low that material may **not always remain suspended** as a wash load once mobilised: a member of the **PODDS** steering group has reported that, in one large main with a very low velocity profile, material appears to be easily resuspended following a period of deliberately elevated flow. Through further studies, the nature of settling, resuspension, transport and mixing in large-diameter trunk mains can be investigated to determine whether those effects are detectable and significant. Such investigations may be driven by 'top-down' **VCDM** modelling using the simplest set of justifiable assumptions or, in contrast, by 'bottom-up' physical analysis using, for example, measured **PSDs** and settling/resuspension velocity calculations.

7.6 Future work

The opportunities for future related work, many of which were discussed in more detail earlier in the chapter, are summarised below:

- Further modelling to validated the **VCDM**'s predictive functionality, using new and

extended trunk main monitoring datasets as they become available (ideally with upstream turbidity monitoring data) (§7.4).

- Develop and justify more granular, time-dependent regeneration rate sub-models.
- If the model can be shown to produce accurate predictions when used with more accurate data and possibly a refined regeneration sub-model then design and implement flow conditioning programmes for real DWDS and quantify the whole life costs of these programmes (§7.2.1). Also, implement a turbidity smart alarm using a real-time VCDM model.
- Further investigate more material accumulating at weaker shear strengths: test for universality, absolute versus relative shear strength ranges and try to confirm which of β_r and α (or both) varying with shear strength is the cause of the identified behaviour. Collecting stepped flushing data from water companies could help with this.
- Develop and test an alternative discolouration model where the strength of material changes over time, emulating biofilm growth and decay.
- Perform field studies to qualify and quantify longitudinal heterogeneity in discolouration processes, updating the VCDM and/or modelling guidelines if any additional complexity can be justified.
- Confirm whether turbidity is proportional to TSS for material eroded from contiguous shear strength bands by measuring turbidity, particle counts and TSS during stepped flushing experiments.
- Repeat the rig experiment and model fitting to the resulting flow and turbidity data to provide further evidence for the growth phase shear stress profile influencing (or not) the fitted model parameters.
- Validate the use of the rolling 5th percentile of the downstream turbidity observations as a proxy for the upstream turbidity using data from one or more trunk mains that were continuously monitored at both ends.
- Confirm the impact of using fifteen-minute-averaged flow and turbidity data for model instantiation and calibration using sub-sampling.
- Experimentally quantify the effect of unsteady hydraulics including transient pressure waves on discolouration potential.
- Ascertain if/how propagation through SRs affects the turbidity signal, developing a SR-level discolouration model if necessary.
- Quantify whether the VCDM will scale to very large diameter trunk mains without needing to revise any model assumptions.
- As mentioned in §7.3, demand from the UK water industry has encouraged DNV GL to implement the PODDS and VCDM formulations within an add-on module for their Synergi DWDS modelling software DNV GL (2010). This module could not be used

in this project as it was only very recently released. However, now it can be used by water providers to assess the potential of the [VCDM](#) and help further validate its potential. All the same, it would be of benefit to the research community to have an open source implementation of the [VCDM](#) that includes greater hydraulic modelling and water quality modelling functionality than `pyvcdm`. Following on from the possibilities discussed in §4.5, a good solution could be to incorporate the [VCDM](#) within EPANET MSX, therefore allowing for the solving of multi-pipe dendritic and looped models. Another benefit of this solution would be that the regeneration rate could be a function of the concentrations/densities of multiple bulk water and pipe wall species. Complementing the [VCDM+MSX API](#) with a [GUI](#) could facilitate high-level access to the software, allowing models to be built more quickly and easily.

Chapter 8

Conclusions

This thesis presents a semi-empirical model of how discolouration material accumulates on and is eroded from the walls of DWDS pipes, for the purpose of tracking discolouration potential over an extended period (weeks or months). The motivation for this research is to equip water providers with a tool for assessing how different interventions and incidents will impact on bulk water turbidity and for developing and optimising flow conditioning programmes for the non-invasive cleaning of critical pipes.

The key advance of the Variable Condition Discolouration Model (VCDM) is the representation of the latent material quantity versus shear strength relationship of heterogeneous, wall-bound discolouration material as a relative amount for each of a number of contiguous shear strength bands, where each band can erode or regenerate independently (§4). This empirical construct permits the erosion and accumulation mechanisms to function differently, in keeping with observed behaviour: field and laboratory studies have previously shown that material weaker than the shear stress erodes from weakest to strongest but, for shear strengths greater than the shear stress, material appears to accumulate in a way that is independent of shear strength (§2.1.13.4).

The model has three scalar parameters: a relative accumulation rate (β_r), a relative erosion rate factor (β_e) and a scaling factor (α) that converts the rate of material released from the pipe wall to units that are a function of the industry-standard NTU (§4.4; §4.4.5). The variability of the three parameters with shear strength was tested by allowing material to accumulate on the walls of two realistic pipe rigs over three months then increasing the shear stress in equal increments whilst monitoring turbidity (§6.2.2). Analysis and modelling of the flow and turbidity data showed that the model assumption β_r and α are constant with shear strength was valid for all but the weakest studied shear strengths. It is not known if this result is common to all DWDS pipe materials and source waters. Test dependencies meant that the variability of β_e with shear strength could not be confirmed.

The β_r parameter is assumed to be constant with time, as at present there is insufficient quantified understanding of the factors that influence accumulation processes for them to be represented by anything more complex. However, the VCDM can be considered a framework where β_r could be made dependent on asset and bulk water characteristics if further laboratory and field work is conducted to develop a more granular accumulation sub-model. The potential for switching between accumulation sub-models has here been demonstrated for a trunk main

where the effect of making β_r dependent on atmospheric temperature, a likely factor, was investigated (§6.3.4).

Models of four DMA flushes show that the erosion sub-model of the VCDM can represent erosion events as well as the validated PODDS discolouration model but using just two parameters to describe erosion rather than four. These results, along with verification tests where the VCDM was fitted to PODDS predictions, do not disprove the chosen empirical structure of the erosion sub-model, this being that each shear strength band erodes at a rate that is constant with time (until the band is depleted) and weaker bands erode more quickly than stronger bands. This permits non-serial layer erosion, so can account for spatial heterogeneity in material composition, and can ensure that the erosion rate is a function of the excess shear stress (a key characteristic of the validated PODDS erosion model). It is proposed that the VCDM replace the PODDS discolouration model: it provides the same erosion functionality with fewer parameters, makes the same assumptions regarding the relationship between turbidity and TSS, uses the same tested mass transport model, but it also includes an accumulation mechanism that far more closely resembles observed behaviour. The combined effect of these erosion and regeneration mechanisms is that, as the shear stress varies over time, the material quantity versus shear strength relationship can become a non-trivial, monotonically increasing function (§4.3.5; fig. 6.47); this latent state is not something that the PODDS formulations can represent (§4.2.2).

The validity of the model for representing both erosion and regeneration was explored using three multi-month/year trunk main monitoring datasets (§6.3.3; §6.3.4; §6.3.5). This was a novel opportunity as at present very few long-term DWDS turbidity monitoring datasets exist and until now there has been no modelling framework suitable for representing such data. However, the small number of available datasets limit the assessment of the generality of the VCDM.

For each of the three modelled trunk mains and for the four DMA flushes, model parameters were fitted using the robust, automatable particle swarm optimisation (PSO) metaheuristic, as sensitivity analysis showed that there were many similar ways in which model parameters and measurable/calculable properties, such as flow and pipe diameter, could influence the shape and scale of time-series model outputs (§5.4; §5.4.3; §5.5; §5.6). The PSO objective function was the error between time-series turbidity observations and predictions, optionally only calculated at times when erosion events were likely. A novel method was developed to estimate when those events would occur (§6.3.4.2). Model boundary conditions (the upstream turbidity time-series and the initial quantity per shear strength band) were fitted/estimated using one of three methods depending on factors such as the availability of historic flow data (§5.8).

The quality fit was high for two of the three modelled trunk mains (§6.3.3; §6.3.5) and was found to depend on the boundary condition estimation method, the fitting period and the presence and characteristics of flow events during that period. Fitted regeneration rates were plausible given the results of previous studies (§7).

The model is of value to water providers if it can make accurate turbidity predictions once calibrated. This capability was tested for one trunk main dataset (§6.3.5.7). Turbidity was over-estimated for the majority of fourteen shear stress events during the prediction and two

events were predicted accurately. Calibration and predictive errors are in part due to a lack of upstream turbidity data but could also be due to over-fitting, to model formulations being partly invalid or to how flow data was aggregated when first recorded.

Although the predictive capability has not yet been conclusively proven, the developed model has been shown to capture much of the known behaviour of discolouration material and has been validated as far as possible given the presently available data.

the model was recently implemented in commercial software (DNV GL, 2010, latest release as of May 2015). The VCDM software used during model development (§4.5) was designed for rapidly validating the model for simple trunk main systems (§2.1.17.2) and so could only model a single pipe (§4.3.6.5; §4.5.2.2) or multiple pipes in series (§4.3.6.5); the commercial implementation is not subject to this limitation.

This model that encapsulates the latest understanding of discolouration processes (5.2) can in future be refined using knowledge gained from calibrating models of a variety of different systems and by using machine learning techniques to build regeneration sub-models (§7.5; §2.2.8.7). The resulting improvements in predictive accuracy should allow the development of asset management strategies where shear stresses are periodically increased so that anticipated/unexpected flow increases do not cause unacceptable turbidities. Using a calibrated VCDM model, the magnitude, duration and spacing of controlled shear stress increases could be optimised to keep the discolouration potential below a threshold whilst minimising operational expenditure and energy (§7.2.1). Additionally, a model instance could be continually updated using data streams for SCADA so as to allow operators to promptly explore the impact of changes in flow and to provide pre-emptive alarms for turbidity failures (§7.2.2). These proposed developments would further enhance the potential of the model presented in this thesis for the medium term, pro-active management of discolouration potential.

References

- ABB, 2012. 4690 Series Turbidity systems (Data sheet DS/4690–EN Rev. A).
URL [http://www05.abb.com/global/scot/scot203.nsf/veritydisplay/453f2b1a92e27ad4c1257a460031049b/\\$file/DS_4690-EN%20Rev.%20A.pdf](http://www05.abb.com/global/scot/scot203.nsf/veritydisplay/453f2b1a92e27ad4c1257a460031049b/$file/DS_4690-EN%20Rev.%20A.pdf)
- Abe, Y., Skali-Lami, S., Block, J.-C., Francius, G., 2012. Cohesiveness and hydrodynamic properties of young drinking water biofilms. *Water Research* 46 (4), 1155–1166, [doi:10.1016/j.watres.2011.12.013](https://doi.org/10.1016/j.watres.2011.12.013).
- Ackers, P., White, W. R., 1973. Sediment Transport: New Approach and Analysis. *Journal of the Hydraulics Division* 99 (11), 2041–2060.
- Aisopou, A., Stoianov, I., Arora, A., O'Hare, D., Graham, N., 2010a. Multi-parameter water quality sensors for water supply systems. In: Boxall, J., Maksimovic, C. (Eds.), *Integrating water systems: Proceedings of the Tenth International Conference on Computing and Control for the Water Industry*. Sheffield, UK.
- Aisopou, A., Stoianov, I., Graham, N., 2010b. Modelling Discolouration in WDS Caused by Hydraulic Transient Events. In: *Water Distribution System Analysis 2010*. Tucson, AZ, USA.
- Aisopou, A., Stoianov, I., Graham, N. J., 2011. In-pipe water quality monitoring in water supply systems under steady and unsteady state flow conditions: A quantitative assessment. *Water Research* 46 (1), 235–246, [doi:10.1016/j.watres.2011.10.058](https://doi.org/10.1016/j.watres.2011.10.058).
- Andrade, M. A., Rojano, F., Romero-Gomez, P., Choi, C. Y., 2010. Integrated Water Quality Modeling of Water Distribution Systems. In: *WDSA2010*. Tucson, AZ, USA.
- Archuleta, E. G., Manwaring, J. F., May 2002. *Guidance Manual for Monitoring Distribution System Water Quality*, pap/cdr Edition. American Water Works Association, US.
- ATi, 2011. A15/76 Turbidity Monitor O&M Manual. Tech. Rep. Rev-H (8/11), Analytical Technology Inc, Collegeville, USA.
- AWWA Research Foundation, 1996. *Characterization and modeling of chlorine decay in distribution systems*. AWWA Research Foundation : American Water Works Association, Denver, CO.
- Balci, O., 1997. Verification validation and accreditation of simulation models. In: *Proceedings of the 29th conference on Winter simulation*. IEEE Computer Society, pp. 135–141.

- Barbeau, B., Gauthier, V., Julienne, K., Carriere, A., 2005. Dead-end flushing of a distribution system: Short and long-term effects on water quality. *Journal of Water Supply: Research and Technology - AQUA* 54 (6), 371–383.
- Batté, M., Appenzeller, B., Grandjean, D., Fass, S., Gauthier, V., Jorand, F., Mathieu, L., Boualam, M., Saby, S., Block, J., 2003. Biofilms in Drinking Water Distribution Systems. *Reviews in Environmental Science and Biotechnology* 2 (2), 147–168, doi:10.1023/B:RESB.0000040456.71537.29.
- Behnel, S., Bradshaw, R., Citro, C., Dalcin, L., Seljebotn, D., Smith, K., 2011. Cython: The Best of Both Worlds. *Computing in Science Engineering* 13 (2), 31 –39, doi:10.1109/MCSE.2010.118.
- Bennett, N. D., Croke, B. F., Guariso, G., Guillaume, J. H., Hamilton, S. H., Jakeman, A. J., Marsili-Libelli, S., Newham, L. T., Norton, J. P., Perrin, C., Pierce, S. A., Robson, B., Seppelt, R., Voinov, A. A., Fath, B. D., Andreassian, V., 2013. Characterising performance of environmental models. *Environmental Modelling & Software* 40, 1–20, doi:10.1016/j.envsoft.2012.09.011.
- Bentley Systems Inc, 2013. WaterGEMS.
URL <http://www.bentley.com/en-GB/Products/WaterGEMS/>
- Besner, M.-C., Broséus, R., Lavoie, J., Giovanni, G. D., Payment, P., Prévost, M., 2010. Pressure Monitoring and Characterization of External Sources of Contamination at the Site of the Payment Drinking Water Epidemiological Studies. *Environmental Science & Technology* 44 (1), 269–277, doi:10.1021/es901988y.
- Besner, M.-C., Ebachner, G., Lavoie, J., Prévost, M., 2007a. Low and Negative Pressures in Distribution Systems: Do They Actually Result in Intrusion? ASCE, pp. 480–480.
- Besner, M.-C., Gauthier, V., Trepanier, M., Martel, K., Prévost, M., 2007b. Assessing the effect of distribution system O&M on water quality. *J. Am. Water Works Ass.* 99 (11), 77–91.
- Besner, M.-C., Modak, P. R., Glause, N., 2012. Extensive Sediment Characterization During Unidirectional Flushing in a Distribution System. ASCE, Adelaide, Australia.
- Bhattacharya, B., Solomatine, D., Mar. 2006. Machine learning in sedimentation modelling. *Neural Networks* 19 (2), 208–214, doi:10.1016/j.neunet.2006.01.007.
- Bhave, P., 1991. Analysis of flow in water distribution networks. Technomic Publishing, Lancaster, PA, USA.
- Biswas, P., Lu, C., Clark, R. M., Dec. 1993. A model for chlorine concentration decay in pipes. *Water Research* 27 (12), 1715–1724, doi:10.1016/0043-1354(93)90108-T.
- Blokker, E. J. M., 2010. Stochastic water demand modelling for a better understanding of hydraulics in water distribution networks. Ph.D. thesis, Water Management Academic Press, Delft University of Technology.
URL <http://resolver.tudelft.nl/uuid:82f6a988-2cef-4eba-aef2-d9790e283f95>

- Blokker, E. J. M., Vreeburg, J. H. G., Schaap, P. G., van Dijk, J. C., 2010. The Self-Cleaning Velocity in Practice. In: *Water Distribution System Analysis 2010*. Tucson, AZ, USA.
- Blokker, E. M., Pieterse-Quirijns, E., 2013. Modeling temperature in the drinking water distribution system. *Journal - American Water Works Association* 105 (01), E19–E28, [doi:10.5942/jawwa.2013.105.0011](https://doi.org/10.5942/jawwa.2013.105.0011).
- Blokker, E. M., Schaap, P. G., Vreeburg, J. H., 2011. Comparing the fouling rate of a drinking water distribution system in two different configurations. In: *Urban Water Management: Challenges and Opportunities*. Exeter, UK.
- Borges, J. L., 2000. On Exactitude in Science. In: *Labyrinths: selected stories and other writings*. Penguin, London.
- Box, G. E. P., 1987. *Empirical model-building and response surfaces*. Wiley series in probability and mathematical statistics. Wiley, New York.
- Boxall, J., Husband, S., 2007. *PODDS II Discolouration Risks Analysis Tool: User Guide (v3.2)*. Tech. rep., Pennine Water Group, University of Sheffield, Sheffield, UK.
- Boxall, J. B., Saul, A. J., 2005. Modeling Discoloration in Potable Water Distribution Systems. *Journal of Environmental Engineering* 131, 716, [doi:10.1061/\(ASCE\)0733-9372\(2005\)131:5\(716\)](https://doi.org/10.1061/(ASCE)0733-9372(2005)131:5(716)).
- Boxall, J. B., Saul, A. J., Gunstead, J. D., Dewis, N., 2003a. Regeneration of Discolouration in Distribution Systems. In: *ASCE EWRI Conference Proceedings*. Philadelphia, USA.
- Boxall, J. B., Saul, A. J., Skipworth, P. J., 2004. Modeling for hydraulic capacity. *Journal / American Water Works Association* 96 (4), 161–169+16.
- Boxall, J. B., Skipworth, P., Saul, A., 2003b. Aggressive flushing for discolouration event mitigation in water distribution networks. *Water Supply* 3 (1-2), 179–186.
- Boxall, J. B., Skipworth, P., Saul, A. J., 2001. A novel approach to modelling sediment movement in distribution mains based on particle characteristics. In: *Water Software Systems: v. 1: Theory and Applications*.
- Brent, R. P., 1972. Chapter 4: An Algorithm with Guaranteed Convergence for Finding a Zero of a Function. In: *Algorithms for minimization without derivatives*. Prentice-Hall series in automatic computation. Prentice-Hall, Englewood Cliffs, N.J.
- Brunone, B., Karney, B., Mecarelli, M., Ferrante, M., 2000. Velocity Profiles and Unsteady Pipe Friction in Transient Flow. *Journal of Water Resources Planning and Management* 126 (4), 236–244, [doi:10.1061/\(ASCE\)0733-9496\(2000\)126:4\(236\)](https://doi.org/10.1061/(ASCE)0733-9496(2000)126:4(236)).
- Bunn, S. M., Reynolds, L., May 2009. The Energy-efficiency Benefits of Pumps-scheduling Optimization for Potable Water Supplies. *IBM J. Res. Dev.* 53 (3), 388–400.
- Burnell, D., Jan. 2003. Auto-validation of district meter data. In: Maksimovic, C., Ali Memon, F., Butler, D. (Eds.), *Advances in Water Supply Management*. Taylor & Francis.

- Campbell, G., Phinn, S. R., 2009. Accuracy and precisions of water quality parameters retrieved from particle swarm optimisation in a sub-tropical lake. In: SPIE 7473. SPIE.
- Chadwick, A. J., 2004. *Hydraulics in civil and environmental engineering*, 4th Edition. Spon Press, London.
- Chanson, H., Takeuchi, M., Trevethan, M., Sep. 2008. Using turbidity and acoustic backscatter intensity as surrogate measures of suspended sediment concentration in a small subtropical estuary. *Journal of Environmental Management* 88 (4), 1406–1416, doi:10.1016/j.jenvman.2007.07.009.
- Clifford, N. J., Richards, K. S., Brown, R. A., Lane, S. N., Dec. 1995. Laboratory and field assessment of an infrared turbidity probe and its response to particle size and variation in suspended sediment concentration. *Hydrological Sciences Journal* 40 (6), 771–791, doi:10.1080/02626669509491464.
- Collins, R., Boxall, J., 2013. Influence of Ground Conditions on Intrusion Flows through Apertures in Distribution Pipes. *Journal of Hydraulic Engineering* 139 (10), 1052–1061, doi:10.1061/(ASCE)HY.1943-7900.0000719.
- Connors, K. A., 1990. *Chemical Kinetics: The Study of Reaction Rates in Solution*. Wiley-VCH.
- Cook, D., 2007. Field investigation of discolouration material accumulation rates in live drinking water distribution systems. Ph.D. thesis, University of Sheffield.
- Cook, D., Boxall, J., 2011. Discoloration Material Accumulation in Water Distribution Systems. *Journal of Pipeline Systems Engineering and Practice* 2 (4), 113–122.
- Cook, D., Husband, P., Boxall, J., Jan. 2015. Operational management of trunk main discolouration risk. *Urban Water Journal*, 1–14 doi:10.1080/1573062X.2014.993994.
- Cooper, G., Cowan, D., 2008. Comparing time series using wavelet-based semblance analysis. *Computers & Geosciences* 34 (2), 95–102, doi:10.1016/j.cageo.2007.03.009.
- Costello, S., Chapman, D., Rogers, C., Metje, N., Sep. 2007. Underground asset location and condition assessment technologies. *Tunnelling and Underground Space Technology* 22 (5-6), 524–542, doi:10.1016/j.tust.2007.06.001.
- Denkhaus, E., Meisen, S., Telgheder, U., Wingender, J., 2007. Chemical and physical methods for characterisation of biofilms. *Microchimica Acta* 158 (1-2), 1–27, doi:10.1007/s00604-006-0688-5.
- Dewis, N., Randall-Smith, M., 2005. Discolouration Risk Modelling. In: *Proceedings of the Eighth International Conference on Computing and Control for the Water Industry*. University of Exeter, Exeter, UK.
- Deyle, E. R., Sugihara, G., 2011. Generalized Theorems for Nonlinear State Space Reconstruction. *PLoS ONE* 6 (3), 1–8, doi:10.1371/journal.pone.0018295.
- DNV GL, 2010. SynerGEE Water.
URL <http://www.gl-group.com/en/water/SynerGEEWater.php>

- Douterelo, I., Husband, S., Boxall, J. B., 2014. The bacteriological composition of biomass recovered by flushing an operational drinking water distribution system. *Water Research* 54, 100–114, doi:10.1016/j.watres.2014.01.049.
- Douterelo, I., Sharpe, R., Boxall, J., 2013. Influence of hydraulic regimes on bacterial community structure and composition in an experimental drinking water distribution system. *Water Research* 47 (2), 503–516, doi:10.1016/j.watres.2012.09.053.
- Dukan, S., Levi, Y., Piriou, P., Guyon, F., Villon, P., 1996. Dynamic modelling of bacterial growth in drinking water networks. *Water Research* 30 (9), 1991–2002, doi:10.1016/0043-1354(96)00021-8.
- Dunlop, E., 1991. WADI Users Manual. Local Government Computer Services Board, Dublin, Ireland.
- Dunne, W. M., 2002. Bacterial Adhesion: Seen Any Good Biofilms Lately? *Clinical Microbiology Reviews* 15 (2), 155–166, doi:10.1128/CMR.15.2.155-166.2002.
- DWI, 2002. Distribution Operation and Maintenance Strategies (DOMS) - DWI Requirements and Expectations. Information Letter 15/2002, Drinking Water Inspectorate, London, UK. URL <http://dwi.defra.gov.uk/stakeholders/information-letters/>
- DWI, 2005. A Brief Guide to Drinking Water Safety Plans. Tech. rep., Drinking Water Inspectorate, London, UK. URL <http://dwi.defra.gov.uk/stakeholders/guidance-and-codes-of-practice/Water%20Safety%20Plans.pdf>
- DWI, Jul. 2009. Bridging the Gaps: deriving knowledge from data for water distribution systems - context, issues and challenges. Meeting organised by the Pennine Water Group, University of Sheffield.
- DWI, 2010. Discoloured water. Tech. rep., Drinking Water Inspectorate, London, UK. URL <http://dwi.defra.gov.uk/consumers/advice-leaflets/discolouration.pdf>
- DWI, 2010. The Water Supply (Water Quality) (England and Wales) Regulations 2000 (unofficial consolidated version). Tech. Rep. SI No.3184, Drinking Water Inspectorate, UK. URL http://dwi.defra.gov.uk/stakeholders/legislation/ws_wqregs2000_cons2010.pdf
- DWI, 2011a. Drinking water 2010 : Southern region of England. Drinking Water Inspectorate, London.
- DWI, 2011b. Drinking water 2010 : Western region of England. Drinking Water Inspectorate, London.
- DWI, 2013a. Annual Report 2012: Letter to Minister - England. Tech. rep., Drinking Water Inspectorate. URL <http://dwi.defra.gov.uk/about/annual-report/2012/letter-english.pdf>
- DWI, 2013b. Drinking Water 2012. Drinking Water Inspectorate, London. URL <http://dwi.defra.gov.uk/about/annual-report/2012/>

- Eberhart, R. C., Shi, Y., 2001. Particle swarm optimization: developments, applications and resources. In: Proceedings of the 2001 Congress on Evolutionary Computation. Vol. 1. pp. 81–86.
- EC, 1998. Drinking Water Directive. Tech. Rep. COUNCIL DIRECTIVE 98/83/EC, European Commission.
URL <http://ec.europa.eu/environment/water/water-drink>
- Elliott, C., Vijayakumar, V., Zink, W., Hansen, R., 2007. National Instruments LabVIEW: A Programming Environment for Laboratory Automation and Measurement. *Journal of the Association for Laboratory Automation* 12 (1), 17–24, doi:10.1016/j.jala.2006.07.012.
- Engelbrecht, A. P., 2005. Fundamentals of Computational Swarm intelligence. Wiley, Chichester, England; Hoboken, NJ.
- Engelhardt, M., Skipworth, P., Savic, D. A., Cashman, A., Walters, G. A., Saul, A. J., Jan. 2002. Determining maintenance requirements of a water distribution network using whole life costing. *Journal of Quality in Maintenance Engineering* 8 (2), 152–164, doi:10.1108/13552510210430026.
- Epstein, J. M., 2008. Why Model? *Journal of Artificial Societies and Social Simulation* 11 (4), 12.
- Farley, B., Mounce, S. R., Boxall, J. B., 2013. Development and Field Validation of a Burst Localization Methodology. *Journal of Water Resources Planning and Management* 139 (6), 604–613, doi:10.1061/(ASCE)WR.1943-5452.0000290.
- Fish, K., 2014. The Impact of Hydraulic Regime upon Biofilms in Drinking Water Distribution Systems. Ph.D. thesis, University of Sheffield, UK.
- Fish, K. E., Collins, R., Green, N. H., Sharpe, R. L., Douterelo, I., Osborn, A. M., Boxall, J. B., Feb. 2015. Characterisation of the Physical Composition and Microbial Community Structure of Biofilms within a Model Full-Scale Drinking Water Distribution System. *PLOS ONE* 10 (2), e0115824, doi:10.1371/journal.pone.0115824.
- Floreano, D., Mattiussi, C., 2008. Bio-inspired artificial intelligence theories, methods, and technologies. MIT Press, Cambridge, Mass.
URL <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=237758>
- Forward, P. D., 2004. Iron Bacteria – Controlling The Scourge of Groundwater Pumping. In: *Engineering Salinity Solutions*. Perth, Australia, pp. 266 – 271.
- Fuentes Cabrera, J. C., Coello Coello, C. A., 2007. Handling Constraints in Particle Swarm Optimization Using a Small Population Size. In: Gelbukh, A., Kuri Morales, A. F. (Eds.), *MICAI 2007: Advances in Artificial Intelligence*. Vol. 4827 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 41–51.
- Furnass, W., Douterelo, I., Collins, R., Mounce, S., Boxall, J., 2014a. Controlled, Realistic-scale, Experimental Study of How the Quantity and Erodibility of Discolouration Material Varies with Shear Strength. *Procedia Engineering* 89, 135–142, doi:10.1016/j.proeng.2014.11.169.

- Furnass, W., Mounce, S., Boxall, J., 2013. Linking distribution system water quality issues to possible causes via hydraulic pathways. *Environmental Modelling & Software* 40, 78–87, doi:10.1016/j.envsoft.2012.07.012.
- Furnass, W. R., Collins, R. P., Husband, P. S., Sharpe, R. L., Mounce, S. R., Boxall, J. B., 2012. Modelling both the continual erosion and regeneration of discolouration material in drinking water distribution systems. In: *New Developments in IT and Water*. IWA, Amsterdam, Netherlands.
- Furnass, W. R., Collins, R. P., Husband, P. S., Sharpe, R. L., Mounce, S. R., Boxall, J. B., 2014b. Modelling both the continual erosion and regeneration of discolouration material in drinking water distribution systems. *Water Science & Technology: Water Supply* 14 (1), 81–90, doi:10.2166/ws.2013.176.
- Gaffney, J. W., Boulton, S., 2011. The Need for and Use of High Resolution Turbidity Monitoring in Managing Discolouration in Distribution. *Journal of Environmental Engineering* 1 (1), 416–416, doi:10.1061/(ASCE)EE.1943-7870.0000521.
- Gauthier, V., Gérard, B., Portal, J.-M., Block, J.-C., Gatel, D., 1999a. Organic matter as loose deposits in a drinking water distribution system. *Water Research* 33 (4), 1014–1026, doi:10.1016/S0043-1354(98)00300-5.
- Gauthier, V., Redercher, S., Block, J.-C., 1999b. Chlorine Inactivation of *Sphingomonas* Cells Attached to Goethite Particles in Drinking Water. *Applied and Environmental Microbiology* 65 (1), 355–357.
- Ghidaoui, M., Mansour, S., 2002. Efficient Treatment of the Vardy–Brown Unsteady Shear in Pipe Transients. *Journal of Hydraulic Engineering* 128 (1), 102–112, doi:10.1061/(ASCE)0733-9429(2002)128:1(102).
- Giustolisi, O., Savic, D., 2006. A symbolic data-driven technique based on evolutionary polynomial regression. *Journal of Hydroinformatics* 8 (3), 207–222, doi:10.2166/hydro.2006.020.
- Graebel, W. P., 2001. *Engineering fluid mechanics*. Taylor & Francis, New York.
- Gray, J., Mar. 2008. Water contamination events in UK drinking-water supply systems. *Journal of Water and Health* 6 (S1), s21, doi:10.2166/wh.2008.123.
- HACH, 2010. Lange 2100q and 2100qis user manual. Tech. Rep. DOC022.52.80041.
URL <http://www.hach.com/2100q-portable-turbidimeter/product-downloads?id=7640450963>
- Hallam, N., West, J., Forster, C., Simms, J., 2001. The potential for biofilm growth in water distribution systems. *Water Research* 35 (17), 4063–4071, doi:10.1016/S0043-1354(01)00248-2.
- Holland, J. H., 1992. *Adaptation in natural and artificial systems : an introductory analysis with applications to biology, control, and artificial intelligence*. MIT Press, Cambridge, Mass.
- Hsieh, W. W., 2009. *Machine learning methods in the environmental sciences: neural networks and kernels*. Cambridge University Press, Cambridge, UK ; New York.

- Hu, X., Eberhart, R., 2002. Solving constrained nonlinear optimization problems with particle swarm optimization. In: Proceedings of the sixth world multiconference on systemics, cybernetics and informatics. Vol. 5. Citeseer, pp. 203–206.
- Hulst, H. C. v. d., 1981. Light scattering by small particles. Dover Publications, New York.
- Hunter, J. D., 2007. Matplotlib: A 2d graphics environment. Computing In Science & Engineering 9 (3), 90–95, doi:10.1109/MCSE.2007.55.
- Husband, P. S., 2010. Discolouration in Water Distribution Systems; Understanding, Modelling and Practical Applications. Ph.D. thesis, University of Sheffield, UK.
- Husband, P. S., Boxall, J. B., 2007. PODDS II Final Report. Tech. rep., University of Sheffield.
- Husband, P. S., Boxall, J. B., 2008. Water Distribution System Asset Deterioration and Impact on Water Quality: A Case Study. ASCE, Hawaii, pp. 516–516.
- Husband, P. S., Boxall, J. B., 2010a. Field Studies of Discoloration in Water Distribution Systems: Model Verification and Practical Implications. Journal of Environmental Engineering 136 (1), 86–94, doi:10.1061/(ASCE)EE.1943-7870.0000115.
- Husband, P. S., Boxall, J. B., 2010b. PODDS III Final Report. Tech. rep., University of Sheffield.
- Husband, P. S., Boxall, J. B., 2011. Asset deterioration and discolouration in water distribution systems. Water Research 45 (1), 113–124, doi:10.1016/j.watres.2010.08.021.
- Husband, P. S., Boxall, J. B., 2012. Misleading Velocity Analysis for Water Quality Management in Transmission Mains. In: World Environmental and Water Resources Congress 2012: Crossing Boundaries. ASCE, pp. 3178–3185.
- Husband, P. S., Boxall, J. B., Saul, A., 2008. Laboratory studies investigating the processes leading to discolouration in water distribution networks. Water Research 42, 4309–4318, doi:10.1016/j.watres.2008.07.026.
- Husband, P. S., Boxall, J. B., Williams, R., 2010a. Discolouration risk management for trunk mains. In: Water Distribution System Analysis 2010. Tucson, AZ, USA, pp. 535–542.
- Husband, P. S., Jackson, M., Boxall, J. B., 2011. Trunk Main Discolouration Trials and Strategic Planning. In: Urban Water Management: Challenges and Opportunities. Exeter, UK.
- Husband, P. S., Whitehead, J., Boxall, J. B., 2010b. The role of trunk mains in discolouration. Proceedings of the ICE - Water Management 163, 397–406, doi:10.1680/wama.900063.
- Husband, S., Mistry, J., Boxall, J., 2013. Modelling and flow conditioning to manage discolouration in trunk mains. In: 12th International Conference on Computing and Control for the Water Industry. Elsevier, Perugia, Italy.
- Innovyze Ltd, 2014. Infoworks WS.
URL http://www.innovyze.com/products/infoworks_ws/
- ISO, 2008. Water quality: Determination of turbidity. Tech. Rep. BS EN ISO 7027:2000.

- Izzo, D., 2012. PyGMO and PyKEP: Open Source Tools for Massively Parallel Optimization in Astrodynamics (the case of interplanetary trajectory optimization).
- Jakeman, A., Letcher, R., Norton, J., May 2006. Ten iterative steps in development and evaluation of environmental models. *Environmental Modelling & Software* 21, 602–614, doi:10.1016/j.envsoft.2006.01.004.
- Jakubovics, N., Shields, R., Rajarajan, N., Burgess, J., Dec. 2013. Life after death: the critical role of extracellular DNA in microbial biofilms. *Letters in Applied Microbiology* 57 (6), 467–475, doi:10.1111/lam.12134.
- Kennedy, J., Mendes, R., 2002. Population structure and particle swarm performance. Vol. 2. IEEE, pp. 1671–1676.
- Kennedy, J. F., Eberhart, R. C., Shi, Y., 2001. *Swarm intelligence*. Morgan Kaufmann Publishers, San Francisco.
- Kirmeyer, G. J., 2000. *Guidance manual for maintaining distribution system water quality*. AWWA Research Foundation and American Water Works Association, Denver, CO.
- Kissa, E., 1999. Dispersions: characterization, testing, and measurement. No. 84 in *Surfactant science series*. M. Dekker, New York.
- Kohl, P. M., Medlar, S. J., 2006. Occurrence of manganese in drinking water and manganese control. AWWA Research Foundation, Denver, USA.
- LeChevallier, M. W., 2005. Microbial Water Quality Within the Distribution System. In: MacPhee, M. J. (Ed.), *Distribution system water quality challenges in the 21st century: a strategic guide*. American Water Works Association, Denver, USA.
- LeChevallier, M. W., Evans, T. M., Seidler, R. J., 1981. Effect of turbidity on chlorination efficiency and bacterial persistence in drinking water. *Applied and Environmental Microbiology* 42 (1), 159–167.
- LeChevallier, M. W., Gullick, R. W., Karim, M. R., Friedman, M., Funk, J. E., 2003. The potential for health risks from intrusion of contaminants into the distribution system from pressure transients. *Journal of Water and Health* 1 (1), 3–14.
- Lee, S. H., O'Connor, J. T., Banerji, S. K., 1980. Biologically mediated corrosion and its effects on water quality in distribution systems. *Journal of the American Water Works Association* 72 (11), 636–645.
- Leeder, A., Mounce, S. M., Boxall, J. B., 2012. Analysis of multi-parameter water quality data using event detection software on laboratory simulated events. Adelaide, Australia.
- Lehtola, M., Miettinen, I., Hirvonen, A., Vartiainen, T., Martikainen, P., 2006. Resuspension of biofilms and sediments to water from pipelines as a result of pressure shocks in drinking water distribution system. In: *Biofilm Systems IV*.
- Lehtola, M. J., Miettinen, I. T., Keinänen, M. M., Kekki, T. K., Laine, O., Hirvonen, A., Vartiainen, T., Martikainen, P. J., 2004a. Microbiology, chemistry and biofilm development

- in a pilot drinking water distribution system with copper and plastic pipes. *Water Research* 38 (17), 3769–3779, [doi:10.1016/j.watres.2004.06.024](https://doi.org/10.1016/j.watres.2004.06.024).
- Lehtola, M. J., Nissinen, T. K., Miettinen, I. T., Martikainen, P. J., Vartiainen, T., 2004b. Removal of soft deposits from the distribution system improves the drinking water quality. *Water research* 38 (3), 601–610, [doi:10.1016/j.watres.2003.10.054](https://doi.org/10.1016/j.watres.2003.10.054).
- Lutz, M., Jan. 2011. *Programming Python*, fourth edition Edition. O'Reilly Media.
- Machell, J., Boxall, J., 2012. Field Studies and Modeling Exploring Mean and Maximum Water Age Association to Water Quality in a Drinking Water Distribution Network. *Journal of Water Resources Planning and Management* 138 (6), 624–638, [doi:10.1061/\(ASCE\)WR.1943-5452.0000220](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000220).
- Machell, J., Boxall, J., Saul, A., Bramley, D., 2009. Improved representation of water age in distribution networks to inform water quality. *Journal of Water Resources Planning and Management* 135 (5), 382–391.
- Mann, A. G., Tam, C. C., Higgins, C. D., Rodrigues, L. C., 2007. The association between drinking water turbidity and gastrointestinal illness: a systematic review. *BMC Public Health* 7 (1), [doi:10.1186/1471-2458-7-256](https://doi.org/10.1186/1471-2458-7-256).
- McClymont, K., Keedwell, E., Savic, D., Randall-Smith, M., 2013. A general multi-objective hyper-heuristic for water distribution network design with discolouration risk. *Journal of Hydroinformatics* 15 (3), 700–716, [doi:10.2166/hydro.2012.022](https://doi.org/10.2166/hydro.2012.022).
- McKay, P., Blain, C. A., Di Iorio, D., Hansell, H., Dec. 2013. Mixing and Turbulence in a Flooding Coastal River. *Journal of Hydraulic Engineering* 139 (12), 1213–1222, [doi:10.1061/\(ASCE\)HY.1943-7900.0000751](https://doi.org/10.1061/(ASCE)HY.1943-7900.0000751).
- McKinney, W., 2013. *Python for data analysis*. O'Reilly, Beijing.
- McNeill, L., 2000. *Water Quality Factors Influencing Iron and Lead Corrosion in Drinking Water*. Ph.D. thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA, chapter 1 of thesis: Review of Iron Pipe Corrosion in Drinking Water Distribution Systems.
- McNeill, L. S., Edwards, M., 2001. The Importance of Temperature in Assessing Iron Pipe Corrosion in Water Distribution Systems. *Environmental Monitoring and Assessment* 77 (3), 229–242, [doi:10.1023/A:1016021815596](https://doi.org/10.1023/A:1016021815596).
- Mehta, A. J., Lee, S.-C., 1994. Problems in Linking the Threshold Condition for the Transport of Cohesionless and Cohesive Sediment Grain. *Journal of Coastal Research* 10 (1), 170–177, articleType: research-article / Full publication date: Winter, 1994 / Copyright © 1994 Coastal Education & Research Foundation, Inc.
- Metropolis, N., Ulam, S., Sep. 1949. The Monte Carlo Method. *Journal of the American Statistical Association* 44 (247), 335–341, [doi:10.1080/01621459.1949.10483310](https://doi.org/10.1080/01621459.1949.10483310).
- Misiunas, D., Vítkovský, J., Olsson, G., Simpson, A., Lambert, M., Jul. 2005. Pipeline Break Detection Using Pressure Transient Monitoring. *Journal of Water Resources Planning and Management* 131 (4), 316–325, [doi:10.1061/\(ASCE\)0733-9496\(2005\)131:4\(316\)](https://doi.org/10.1061/(ASCE)0733-9496(2005)131:4(316)).

- Mounce, S., Gaffney, J., Boulton, S., Boxall, J., 2014a. Automated data driven approaches to evaluating and interpreting water quality time-series data from water distribution systems. *ASCE Journal of Water Resources Planning and Management* In press.
- Mounce, S., Machell, J., Boxall, J., 2012a. Water quality event detection and customer complaint clustering analysis in distribution systems. *Water Science & Technology: Water Supply* 12 (5), 580, [doi:10.2166/ws.2012.030](https://doi.org/10.2166/ws.2012.030).
- Mounce, S., Mounce, R., Boxall, J., 2012b. Identifying Sampling Interval for Event Detection in Water Distribution Networks. *Journal of Water Resources Planning and Management* 138 (2), 187–191, [doi:10.1061/\(ASCE\)WR.1943-5452.0000170](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000170).
- Mounce, S., Sharpe, R., Speight, V., Holden, B., Boxall, J., 2014b. Knowledge discovery from large disparate corporate databases using self-organising maps to help ensure supply of high quality potable water. New York, USA.
- Mounce, S. R., Husband, P. S., Furnass, W. R., Boxall, J. B., 2014c. Multivariate Data Mining for Estimating the Rate of Discoloration Material Accumulation in Drinking Water Systems. *Procedia Engineering* 89, 173–180, [doi:10.1016/j.proeng.2014.11.174](https://doi.org/10.1016/j.proeng.2014.11.174).
- Mustonen, S. M., Tissari, S., Huikko, L., Kolehmainen, M., Lehtola, M. J., Hirvonen, A., 2008. Evaluating online data of water quality changes in a pilot drinking water distribution system with multivariate data exploration methods. *Water Research* 42 (10-11), 2421–2430, [doi:10.1016/j.watres.2008.01.015](https://doi.org/10.1016/j.watres.2008.01.015).
- Naser, G., Karney, B. W., Boxall, J. B., Aug. 2006. Red Water and Discoloration in a WDS: A Numerical Simulation. *ASCE Conference Proceedings* 247 (40941), 152–152, [doi:10.1061/40941\(247\)152](https://doi.org/10.1061/40941(247)152).
- Nash, J., Sutcliffe, J., Apr. 1970. River flow forecasting through conceptual models part I — A discussion of principles. *Journal of Hydrology* 10 (3), 282–290, [doi:10.1016/0022-1694\(70\)90255-6](https://doi.org/10.1016/0022-1694(70)90255-6).
- Neilands, K., Bernats, M., Rubulis, J., 2012. Accumulation and modeling of particles in drinking water pipe fittings. *Drinking Water Engineering and Science Discussions* 5 (1), 139–171, [doi:10.5194/dwesd-5-139-2012](https://doi.org/10.5194/dwesd-5-139-2012), rEVIEW ONLY.
- Nicklow, J., Reed, P., Savic, D., Dessalegne, T., Harrell, L., Chan-Hilton, A., Karamouz, M., Minsker, B., Ostfeld, A., Singh, A., Zechman, E., 2010. State of the Art for Genetic Algorithms and Beyond in Water Resources Planning and Management. *Journal of Water Resources Planning and Management* 136 (4), 412–432, [doi:10.1061/\(ASCE\)WR.1943-5452.0000053](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000053).
- NIST, 2013. NIST/SEMATECH e-Handbook of Statistical Methods. Last accessed: 2014-07-09. URL <http://www.itl.nist.gov/div898/handbook/>
- O'Brien, H., Labadz, J., Butcher, D., Billett, M., Midgley, N., 2008. Impact of catchment management upon dissolved organic carbon and stream flows in the Peak District, Derbyshire, UK. In: *British Hydrological Society's 10th National Hydrology Symposium*. Exeter, UK, pp. 178–185.

- Opher, T., Friedler, E., 2009. A preliminary coupled MT-GA model for the prediction of highway runoff quality. *Science of The Total Environment* 407 (15), 4490–4496, [doi:10.1016/j.scitotenv.2009.04.043](https://doi.org/10.1016/j.scitotenv.2009.04.043).
- Opher, T., Ostfeld, A., 2011. A coupled model tree (MT) genetic algorithm (GA) scheme for biofouling assessment in pipelines. *Water Research* 45 (18), 6277–6288, [doi:10.1016/j.watres.2011.09.037](https://doi.org/10.1016/j.watres.2011.09.037).
- Parchure, T. M., Mehta, A. J., 1985. Erosion of Soft Cohesive Sediment Deposits. *Journal of Hydraulic Engineering* 111 (10), 1308–1326, [doi:10.1061/\(ASCE\)0733-9429\(1985\)111:10\(1308\)](https://doi.org/10.1061/(ASCE)0733-9429(1985)111:10(1308)).
- Parsopoulos, K. E., Vrahatis, M. N., 2002. Particle Swarm Optimization Method for Constrained Optimization Problems. In: *In Proceedings of the Euro-International Symposium on Computational Intelligence 2002*. IOS Press, pp. 214–220.
- Perez, F., Granger, B. E., May 2007. IPython: a System for Interactive Scientific Computing. *Comput. Sci. Eng.* 9 (3), 21–29.
- Pothof, I. W. M., Blokker, E. J. M., 2012. Dynamic hydraulic models to study sedimentation in drinking water networks in detail. *Drinking Water Engineering and Science Discussions* 5 (1), 121–137, [doi:10.5194/dwesda-5-121-2012](https://doi.org/10.5194/dwesda-5-121-2012).
- Potter, S., 2000. Water Temperature Predicts Maintenance Peaks. *Opflow* 26 (7), 6–7.
- Powell, J. C., Hallam, N. B., West, J. R., Forster, C. F., Simms, J., 2000. Factors which control bulk chlorine decay rates. *Water Research* 34 (1), 117–126, [doi:10.1016/S0043-1354\(99\)00097-4](https://doi.org/10.1016/S0043-1354(99)00097-4).
- Prince, R., McManus, K., Goulter, I., 2000. Colour, turbidity levels and dirty water customer complaints-water system performance indicators? In: *10th World Water Congress: Water, the Worlds Most Important Resource*.
- Pudar, R. S., Liggett, J. A., Jul. 1992. Leaks in Pipe Networks. *Journal of Hydraulic Engineering* 118 (7), 1031–1046, [doi:10.1061/\(ASCE\)0733-9429\(1992\)118:7\(1031\)](https://doi.org/10.1061/(ASCE)0733-9429(1992)118:7(1031)).
- Quinlan, J., 1992. Learning with continuous classes. In: *5th Australian joint conference on artificial intelligence*.
- Randall-Smith, M., Collingbourne, J., McClymont, K., 2011. Targeting Water Distribution Network Interventions for the Cost-Effective Mitigation of Discolouration Risk: A Case Study. In: *Urban Water Management: Challenges and Opportunities*. Vol. 2. University, Exeter, UK, pp. 631–633.
- Ratnayaka, D. D., Brandt, M. J., Johnson, M., Feb. 2009. *Water Supply*, Sixth Edition, 6th Edition. Butterworth-Heinemann.
- Read, G. F., Vickridge, I. G. (Eds.), 1997. *Sewers - Rehabilitation and New Construction: Repair and Renovation*. Arnold, London.
- URL <http://www.engineeringvillage.com/controller/servlet/OpenURL?genre=book&isbn=9780340544723>

- Reddy, M. J., Nagesh Kumar, D., 2007. Multi-objective particle swarm optimization for generating optimal trade-offs in reservoir operation. *Hydrological Processes* 21 (21), 2897–2909, doi:10.1002/hyp.6507.
- Reddy, S., Fogler, H., Jan. 1981. Emulsion stability: Determination from turbidity. *Journal of Colloid and Interface Science* 79 (1), 101–104, doi:10.1016/0021-9797(81)90052-7.
- Rossman, L., Boulos, P., 1996. Numerical Methods for Modeling Water Quality in Distribution Systems: A Comparison. *Journal of Water Resources Planning and Management* 122 (2), 137–146, doi:10.1061/(ASCE)0733-9496(1996)122:2(137).
- Rossman, L. A., 2000. EPANET 2 User's Manual. Tech. Rep. EPA/600/R-00/057, U.S. Environmental Protection Agency, Cincinnati, USA.
- Ryan, G., Mathes, M., Haylock, G., Jayaratne, A., Wu, J., Noui-Mehidi, N., Grainger, C., Nguyen, B. V., 2008. Particles in Water Distribution Systems. Tech. Rep. 33, Cooperative Research Centre for Water Quality and Treatment, Salisbury, Australia.
URL <http://www.wqra.com.au/publications/document-search/?download=67>
- Sadar, M., 2003. Turbidity Science. Tech. Rep. 7061, Hach.
URL <http://www.hach.com/quick.search-quick.search.jsa?keywords=lit7061>
- Saldarriaga, J., Nieto, L., Hernandez, F., Osorio, C., 2010. Discolored water analysis and modeling for trunk networks: the case of Bogota, Columbia. In: *Water Distribution System Analysis 2010*. Tucson, AZ, USA.
- Sanford, L. P., Maa, J. P.-Y., 2001. A unified erosion formulation for fine sediments. *Marine Geology* 179 (1–2), 9–23, doi:10.1016/S0025-3227(01)00201-8.
- Savic, D. A., Kapelan, Z. S., Jonkergouw, P. M., 2009. Quo vadis water distribution model calibration? *Urban Water Journal* 6 (1), 3–22, doi:10.1080/15730620802613380.
- Savic, D. A., Walters, G. A., 1997. Genetic Algorithms for Least-Cost Design of Water Distribution Networks. *Journal of Water Resources Planning and Management* 123 (2), 67–77, doi:10.1061/(ASCE)0733-9496(1997)123:2(67).
- Schaap, P., Blokker, M., 2012. Understanding differences in sediment characteristics in drinking water distribution networks. ASCE, Adelaide, Australia.
- Schaap, P., Blokker, M., 2013. Zooming in on network fouling locations. In: *World Environmental and Water Resources Congress 2013: Showcasing the Future*. ASCE, Cincinnati, USA, pp. 1033–1043.
- Schaller, R., Jun. 1997. Moore's law: past, present and future. *IEEE Spectrum* 34 (6), 52–59, doi:10.1109/6.591665.
- Schneider Electric, 2012. Aquis.
URL <http://www.schneider-electric.com/products/ww/en/5100-software/5125-information-management/61417-aquis-software/>

- Scipy, 2012. `scipy.spatial.distance.braycurtis`.
URL <http://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.braycurtis.html>
- Seth, A., Bachmann, R., Boxall, J., Saul, A., Edyvean, R., 2004. Characterisation of materials causing discolouration in potable water systems. *Water Science and Technology* 49 (2), 27–32.
- Seth, A. D., Husband, P. S., Boxall, J. B., 2009. Rivelin trunk main flow test. In: *Integrating Water Systems*. Sheffield.
- Shang, F., Uber, J. G., 2008. EPANET Multi-Species Extension User's Manual. Tech. Rep. EPA/600/S-07/021, U.S. Environmental Protection Agency, Cincinnati, USA.
- Shang, F., Uber, J. G., Rossman, L. A., 2007. Modeling Reaction and Transport of Multiple Species in Water Distribution Systems. *Environ. Sci. Technol.* 42 (3), 808–814, [doi:10.1021/es072011z](https://doi.org/10.1021/es072011z).
- Sharp, E. L., Parsons, S. A., Jefferson, B., Jun. 2006. Seasonal variations in natural organic matter and its impact on coagulation in water treatment. *Science of The Total Environment* 363 (1–3), 183–194, [doi:10.1016/j.scitotenv.2005.05.032](https://doi.org/10.1016/j.scitotenv.2005.05.032).
- Sharpe, R., 2013. Laboratory investigations into processes causing discoloured potable water. Ph.D. thesis, University of Sheffield, UK.
- Shen, H., Nov. 2014. Interactive notebooks: Sharing the code. *Nature* 515 (7525), 151–152, [doi:10.1038/515151a](https://doi.org/10.1038/515151a).
- Siemens, 2013. HydraClam Water Quality Monitor.
URL <http://www.water.siemens.com/SiteCollectionDocuments/Hydraclam%20Water%20Quality%20monitor%20FINAL.pdf>
- Sivanandam, S. N., Deepa, S. N., 2007. *Introduction to genetic algorithms*. Springer, Berlin; New York.
URL <http://public.ebib.com/choice/publicfullrecord.aspx?p=337395>
- Skipworth, P., Engelhardt, M., Cashman, A., Savic, D., Saul, A., Walters, G., 2002a. Whole life costing for water distribution network management. Thomas Telford, London.
- Skipworth, P. J., Machell, J., Saul, A. J., Jan. 2002b. Empirical travel time estimation in a distribution network. *Proceedings of the ICE - Water and Maritime Engineering* 154 (1), 41–49, [doi:10.1680/wame.2002.154.1.41](https://doi.org/10.1680/wame.2002.154.1.41).
- Skipworth, P. J., Tait, S. J., Saul, A. J., Jun. 1999. Erosion of Sediment Beds in Sewers: Model Development. *Journal of Environmental Engineering* 125 (6), 566–573, [doi:10.1061/\(ASCE\)0733-9372\(1999\)125:6\(566\)](https://doi.org/10.1061/(ASCE)0733-9372(1999)125:6(566)).
- Sly, L. I., Hodgkinson, M. C., Arunpairojana, V., 1990. Deposition of manganese in a drinking water distribution system. *Applied and Environmental Microbiology* 56 (3), 628–639.

- Smith, S. W., 1997. The scientist and engineer's guide to digital signal processing, 1st Edition. California Technical Pub, San Diego, Calif.
- Solomatine, D., 2002. Data-driven modelling: paradigm, methods, experiences. In: Proc. 5th international conference on hydroinformatics. pp. 1–5.
- Solomatine, D. P., Ostfeld, A., Jan. 2008. Data-driven modelling: some past experiences and new approaches. *Journal of Hydroinformatics* 10 (1), 3, [doi:10.2166/hydro.2008.015](https://doi.org/10.2166/hydro.2008.015).
- Sonnenwald, F., Stovin, V., Guymer, I., 2013. Correlation Measures for Solute Transport Model Evaluation. In: Rowinski, P. (Ed.), *Experimental and Computational Solutions of Hydraulic Problems*. Poland.
- Soulsby, R., 1997. *Dynamics of marine sands*. T. Telford, London.
- Srinivasan, S., Harrington, G., Xagorarakis, I., Goel, R., 2008. Factors affecting bulk to total bacteria ratio in drinking water distribution systems. *Water Research* 42 (13), 3393–3404.
- Stoodley, P., Jacobsen, A., Dunsmore, B. C., Purevdorj, B., Wilson, S., Lappin-Scott, H. M., Costerton, J. W., 2001a. The influence of fluid shear and AICI₃ on the material properties of *Pseudomonas aeruginosa* PAO1 and *Desulfovibrio* sp. EX265 biofilms. *Water science and technology: a journal of the International Association on Water Pollution Research* 43 (6), 113–120.
- Stoodley, P., Wilson, S., Hall-Stoodley, L., Boyle, J. D., Lappin-Scott, H. M., Costerton, J. W., 2001b. Growth and Detachment of Cell Clusters from Mature Mixed-Species Biofilms. *Applied and Environmental Microbiology* 67 (12), 5608–5613, [doi:10.1128/AEM.67.12.5608-5613.2001](https://doi.org/10.1128/AEM.67.12.5608-5613.2001).
- Suribabu, C. R., Neelakantan, T. R., 2006. Design of water distribution networks using particle swarm optimization. *Urban Water Journal* 3 (2), 111–120, [doi:10.1080/15730620600855928](https://doi.org/10.1080/15730620600855928).
- Takens, F., 1981. Detecting strange attractors in turbulence. In: *Dynamical systems and turbulence*, Warwick 1980. Springer, pp. 366–381.
- Telgmann, U., Horn, H., Morgenroth, E., Oct. 2004. Influence of growth history on sloughing and erosion from biofilms. *Water Research* 38 (17), 3671–3684, [doi:10.1016/j.watres.2004.05.020](https://doi.org/10.1016/j.watres.2004.05.020).
- Tennekes, H., 1972. *A first course in turbulence*. MIT Press, Cambridge, Mass.
- The HDF Group, 2014. Hierarchical Data Format, version 5. [Http://www.hdfgroup.org/HDF5/](http://www.hdfgroup.org/HDF5/).
- Tzatchkov, V. G., Buchberger, S. G., Li, Z., Romero-Gomez, P., Choi, C., 2009. Axial dispersion in pressurized water distribution networks-A review. In: *Int Symp Water Manag Hydraul Eng*. pp. 581–592.
- Uber, J., 2009. Multi-species network water quality modeling: Current examples, future potential, and research needs. In: *Integrating Water Systems*. pp. 13–19.
- USEPA, 1993. Determination of Turbidity by Nephelometry. Tech. Rep. Method 180.1 revision 2.

- USEPA, 2001. Potential Contamination Due to Cross-Connections and Backflow and the Associated Health Risks. Tech. rep., US Environment Protection Agency.
URL http://www.epa.gov/safewater/disinfection/tcr/regulation_revisions.html
- USEPA, 2006. Distribution System Indicators of Drinking Water Quality (draft). Tech. rep., US Environment Protection Agency.
URL http://www.epa.gov/ogwdw000/disinfection/tcr/pdfs/issuepaper_tcr_indicators.pdf
- USEPA, 2008. EPANET.
URL <http://www.epa.gov/nrmrl/wswrd/dw/epanet.html>
- USEPA, 2009. National Primary Drinking Water Regulations. Tech. Rep. EPA 816-F-09-004, US Environment Protection Agency.
URL <http://water.epa.gov/drink/contaminants/upload/mcl-2.pdf>
- USEPA, 2010. Water Quality Event Detection Systems for Drinking Water Contamination Warning 576 Systems: Development, Testing and Application of CANARY. Tech. Rep. EPA/600/R-10/036, Washington, DC.
- Vaerewijck, M. J., Huys, G., Palomino, J. C., Swings, J., Portaels, F., 2005. Mycobacteria in drinking water distribution systems: ecology and significance for human health. FEMS Microbiology Reviews 29 (5), 911–934, doi:10.1017/S1462399403005908.
- van der Walt, S., Colbert, S. C., Varoquaux, G., Mar. 2011. The NumPy Array: A Structure for Efficient Numerical Computation. Computing in Science & Engineering 13 (2), 22–30, doi:10.1109/MCSE.2011.37.
- van Thienen, P., Floris, R., Meijering, S., Dec. 2011a. Application of optical tomography in the study of discolouration in drinking water distribution systems. Drinking Water Engineering and Science 4 (1), 61–69, doi:10.5194/dwes-4-61-2011.
- van Thienen, P., Vreeburg, J., Blokker, E., 2011b. Radial transport processes as a precursor to particle deposition in drinking water distribution systems. Water Research 45, 1807–1817, doi:10.1016/j.watres.2010.11.034.
- van Waveren, R., Groot, S., Scholten, H., van Geer, F., Wosten, J., Koeze, R., Noort, J., 1999. Good modelling practice handbook. No. 99.036. STOWA/RIZA, Dutch Dept. of Public Works, Institute for Inland Water Management and Waste Water Treatment,, Utrecht; Lelystad; Den Haag.
URL harmoniqua.wau.nl/public/Reports/Existing%20Guidelines/GMP111.pdf
- Verberk, J., Vreeburg, J., Rietveld, L., van Dijk, J., 2009. Particulate fingerprinting of water quality in the distribution system. Water SA 35 (2), doi:10.4314/wsa.v35i2.
- Vreeburg, J., 2007. Discolouration in Drinking Water Systems: A Particular Approach. PhD thesis, Delft University of Technology, Netherlands.
- Vreeburg, J., Schaap, P., van Dijk, J., 2004. Particles in the drinking water system: from source to discolouration. Water Science & Technology: Water Supply 4 (5-6), 431–438.

- Vreeburg, J., Schippers, D., Verberk, J., van Dijk, J., 2008. Impact of particles on sediment accumulation in a drinking water distribution system. *Water Research* 42 (16), 4233–4242, [doi:10.1016/j.watres.2008.05.024](https://doi.org/10.1016/j.watres.2008.05.024).
- Vreeburg, J. H., Beverloo, H., 2011. Sediment accumulation in drinking water trunk mains. In: *Urban Water Management: Challenges and Opportunities*. Exeter, UK.
- Vreeburg, J. H. G., Boxall, J. B., 2007. Discolouration in potable water distribution systems: A review. *Water Research* 41, 519–529, [doi:10.1016/j.watres.2006.09.028](https://doi.org/10.1016/j.watres.2006.09.028).
- Wainwright, J., Mulligan, M. (Eds.), 2004. *Environmental modelling: finding simplicity in complexity*. Wiley, Chichester, West Sussex, England ; Hoboken, NJ.
- Walker, W., Harremoës, P., Rotmans, J., van der Sluijs, J., van Asselt, M., Janssen, P., Kreyer von Krauss, M., 2003. Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support. *Integrated Assessment* 4 (1), 5–17, [doi:10.1076/iaij.4.1.5.16466](https://doi.org/10.1076/iaij.4.1.5.16466).
- Walski, T., 2004. Modeling for hydraulic capacity: discussion. *Journal / American Water Works Association* 96 (10), 104–106.
- Walski, T. M., Chase, D. V., Savic, D. A., Grayman, W. M., Beckwith, S., Koelle, E., 2003. *Advanced water distribution modeling and management*. Haestad Press.
- Wang, H., Hu, C., Hu, X., Yang, M., Qu, J., Mar. 2012. Effects of disinfectant and biofilm on the corrosion of cast iron pipes in a reclaimed water distribution system. *Water Research* 46 (4), 1070–1078, [doi:10.1016/j.watres.2011.12.001](https://doi.org/10.1016/j.watres.2011.12.001).
- Whelton, A. J., Dietrich, A. M., Gallagher, D. L., Roberson, J. A., 2007. Using customer feedback for improved water quality and infrastructure monitoring. *Journal / American Water Works Association* 99 (11), 62–76.
- WHO, 2005. *Water Safety Plans: Managing drinking-water quality from catchment to consumer*. Tech. Rep. WHO/SDE/WSH/05.06, World Health Organization, Geneva.
URL http://www.who.int/water_sanitation_health/dwq/wsp0506
- WHO, 2011. *Guidelines for drinking-water quality, 4th Edition*. World Health Organization, Geneva.
URL http://www.who.int/water_sanitation_health/publications/2011/dwq_guidelines/en
- Wilkinson, B., 1999. *Parallel programming: techniques and applications using networked workstations and parallel computers*. Prentice Hall, Upper Saddle River, N.J.
- WRC plc, 1994. *Water Industry Instrument Handbook: Book 4 Turbidity*. Tech. rep., Wiltshire, UK.
- Wu, W., Dandy, G. C., Maier, H. R., Apr. 2014. Protocol for developing ANN models and its application to the assessment of the quality of the ANN model development process in drinking water quality modelling. *Environmental Modelling & Software* 54, 108–127, [doi:10.1016/j.envsoft.2013.12.016](https://doi.org/10.1016/j.envsoft.2013.12.016).

- Wymer, L. J., 2007. *Statistical Framework for Recreational Water Quality Criteria and Monitoring*. Wiley-Interscience.
- Xu, S., Rahmat-Samii, Y., 2007. Boundary conditions in particle swarm optimization revisited. *IEEE Transactions on Antennas and Propagation* 55 (3 I), 760–765.
- Yagi, S., Shiba, S., 1999. Application of genetic algorithms and fuzzy control to a combined sewer pumping station. *Water Science and Technology* 39 (9), 217–224, doi:10.1016/S0273-1223(99)00236-X.
- Yu, T. C., Shao, Y., Shen, C., May 2014. Mixing at Cross Joints with Different Pipe Sizes in Water Distribution Systems. *Journal of Water Resources Planning and Management* 140 (5), 658–665, doi:10.1061/(ASCE)WR.1943-5452.0000372.
- Zacheus, O. M., Lehtola, M. J., Korhonen, L. K., Martikainen, P. J., May 2001. Soft deposits, the key site for microbial growth in drinking water distribution networks. *Water Research* 35 (7), 1757–1765.
- Zambrano-Bigiarini, M., Rojas, R., 2013. A model-independent Particle Swarm Optimisation software for model calibration. *Environmental Modelling & Software* doi:10.1016/j.envsoft.2013.01.004.
- Ziegler, A. C., 2002. Issues Related to Use of Turbidity Measurements as a Surrogate for Suspended Sediment. In: *Proceedings of the Federal Interagency Workshop on Turbidity and Other Sediment Surrogates*. Reno, USA, pp. 16–18, u.S. Geological Survey Circular 1250.