Inherently Robust, Adaptive Model Predictive Control: An Opportunity for Gas Turbines

A thesis submitted to the University of Sheffield for the degree of Doctor of Philosophy

The University Of Sheffield.

Kacper Grzędziński

Department of Automatic Control and Systems Engineering

March 2020

Dla rodziny

Acknowledgements

My profound thanks go to the University of Sheffield Rolls-Royce University Technology Centre for providing me with the opportunity to develop my ideas in parallel with the sponsor's interests. Special mentions go to Paul Trodden, my supervisor and mentor, who introduced me to the wonderfully nuanced field of optimisation-based control. I owe Andrew Hills my sincerest gratitude for his assistance with all things computerised; including this no-nonsense thesis template. Visakan Kadirkamanthan and Andy Mills both found the patience to take me on-board, for which I humbly thank them for.

My parents and teachers deserve the biggest praise - their support in promoting my curiosity in the sciences, throughout my entire life, got me where I am. Sheffield has provided me with experiences and friendships that I will hold dearly in memory; from ludicrous lunch-time discussions to impromptu trips to the Peak District.

I would also like to acknowledge the helpfulness of the colleagues from across the pond at Georgia Tech, for letting me use their engine data to simulate my algorithms. Their kindness in providing permission for use of their rig photographs does a great job to remind us all that this is an engineering thesis, and not just esoteric mathematics.

Abstract

Civil-aero gas turbines are becoming increasingly more complex. A-priori modelling of engine dynamics through first-principles physics and experimental data collection are mandraulic tasks that are mandatory for certifying engine control. There is tremendous economic incentive for developing controllers that are robust, optimal and fast, with known performance guarantees despite the presence of natural engine degradation. Moreover, increasing engine intelligence by utilising coupling between novel control effectors has been identified as a route to increasing capability and efficiency of future civil aero-engine concepts. In this thesis, automatic control for regulation of engine thrust, in a multivariable engine architecture that exhibits degradation, is used as a case-study for developing a novel adaptive control algorithm.

The thesis reviews, critiques and provides a comprehensive summary of contemporary control strategies in the field of aero-propulsion. The result of this review is the selection of model predictive control (MPC) as a candidate control algorithm, for tackling the multivariable and constrained thrust regulation problem. The fundamental theoretical properties of nominal linear quadratic (LQ)-MPC and its limitations motivate the theoretical contribution of this research.

The academic contribution is a chapter-by-chapter derivation of an inherently robust adaptive MPC control law. The dual-control problem of simultaneously regulating a plant and identifying its dynamics is addressed in a linear and constrained MPC framework. The technical chapters describe how robustness properties of the nominal LQ-MPC formulation can be used to overcome the shortcomings of closed-loop identification for learning degraded engine dynamics; without using explicitly robust MPC methods that are known to induce conservatism. Under technical assumptions, the proposed scheme is used to show how the prediction model can be safely updated to provide self-tuning capability.

The efficacy of the inherently robust adaptive MPC is finally demonstrated using experimentally validated linear models of a real desktop sized turboprop engine. Regulation of the engine's shaft speeds, to those at cruise, is used to emulate a typical acceleration command. Despite model mismatch and presence of constraints, the inherently robust adaptive MPC controller is able to learn the true dynamics of the gas turbine whilst regulating the states to the desired cruise operating point. Adapting the controller's parameters to the ones associated with the true dynamics can therefore allow for more economical and safer exploitation of aero-engines in the presence of natural degradation.

Contents

List of Figures

List of Tables

List of Algorithms

Nomenclature

A list of the variables and notation used in this thesis is defined below. The definitions and conventions set here will be observed throughout unless otherwise stated. For a list of acronyms, please consult page xix.

α_u	input constraint tightening factor
α_x	state constraint tightening factor
$\bar{x}^+$	predicted consecutive state
ϵ	excitation rate of change bounding constant
γ	RLS forgetting factor
$\mathbb{P}_1$	regulation optimisation problem
$\mathbb{P}_2$	excitation optimisation problem
$\mathbb{U}$	input constraint set
$\mathbb{X}$	state constraint set
$\mathbb{X}_f$	terminal state constraint set
$\mathcal{X}_N$	feasible state-space region
$\top$	Matrix transpose
Ξ	degradation scalar
ζ	value function decay rate
A	dynamics matrix
B	input matrix
b	parametrisation constant of the largest sub-level set

c_1	value function lower bound constant
c_2	value function upper bound constant
E	excitation magnitude bounding constant
K	linear state feedback matrix
k	current time-step
L	Lipschitz constant
N	control horizon
P	terminal state weighting matrix
Q	state weighting matrix
R	RLS information matrix
R	input weighting matrix
u	measured input at current time-step
$V_N^0(x)$	regulation value function (optimal cost function)
$V_N(x, \mathbf{u})$	regulation cost function
w	unknown disturbance at current time-step
x	measured state at current time-step

Definitions

The following list defines numerous fundamental concepts and notation that are referred to throughout the body of this thesis.

Linear Algebra

- *Positive (Semi-)Definiteness* : $M(\succeq) \succ 0$ denotes positive (semi-)definiteness of the matrix M.

- *Convex Function* : A function f over some domain D is convex if for all pairs $x_1, x_2 \in D$ the following condition holds: $f(\alpha x_1 + (1-\alpha)x_2) \leq \alpha f(x_1) + (1-\alpha)f(x_2)$, where $0 \leq \alpha \leq 1$.

- *Vector Norm* : The operator $||.||_p$ satisfies the properties of a vector norm for $1 \leq p \leq \infty$ and is defined for $x \in \mathbb{R}^n$ as $||x||_p = \sum_{i=1}^{n} \left(|x_i|^p\right)^{\frac{1}{p}}$. If the subscript is omitted for brevity, $||.||$ denotes the euclidean norm *i.e.*, $p = 2$.

- *Weighted 2-norm* The operator $||.||_{2,Q}$ satisfies the properties of a vector norm and is defined for $x \in \mathbb{R}^n$ as $||x||_{2,Q} = \sum_{i=1}^{n} \left(\alpha_i |x_i|^2\right)^{\frac{1}{2}}$ and diagonal weighting matrix Q with $\text{diag}(Q) = [\alpha_1, \ldots, \alpha_N]^\top$.

- *Frobenius Matrix Norm* : The operator $||A||$ applied to a matrix $A \in \mathbb{R}^{n \times m}$ is defined as
$$||A|| = \sqrt{\sum_{i=1}^{m} \sum_{j=1}^{n} |a_{ij}|^2}.$$

- *Linear Mapping* : For the set $X \subset \mathbb{R}^n$, AX denotes the mapping $\{Ax : x \in \mathbb{R}^n\}$.

- *Matrix Inverse* : For a square matrix $A \in \mathbb{R}^{n \times n}$, the inverse A^{-1} satisfies $A^{-1}A \triangleq I$ and holds for the commutation.

- *Minimum Eigenvalue* : For a square matrix $A \in \mathbb{R}^{n\times n}$, the minimum eigenvalue is $\underline{\lambda}(A) \triangleq \min(\Lambda)$ where $\Lambda = \{\lambda_1, \ldots, \lambda_n\}$ is the set of eigenvalues (possibly non-unique) and $\mathcal{V} = \{v_1, \ldots, v_n\}$ are the corresponding eigenvectors that satisfy $Av_i = \lambda_i v_i$ for $i = 1, \ldots, n$.

Set Methods

- *Polyhedra* : A polyhedron is a convex intersection of closed half-spaces, and a polytope is a bounded polyhedron denoted $S = \{x \in \mathbb{R}^n : Ax \leq b\}$.

- *Positive Invariance* : A set $X \subset \mathbb{R}^n$ is positive invariant (PI) for the system $x^+ = f(x)$ (where $f : \mathbb{R}^n \mapsto \mathbb{R}^n$) if $f(x) \in X$ for all $x \in X$.

- *Level Set* : The level set of a function $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ over a domain $X \subseteq \mathbb{R}^n$ is defined as $\Omega_a \triangleq \{x : V(x) \leq a, x \in X\}$ for some level constant $a \in \mathbb{R}_{\geq 0}$.

- *Set Intersection and Union* : For two sets A and B the set intersection and union is denoted $A \cap B$ and $A \cup B$ respectively.

- *Minkowski Addition* : The Minkowski sum of two sets $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^n$ is defined as $A \oplus B \triangleq \{a + b : a \in A, b \in B\}$.

- *Pontryagin Difference* : The Pontryagin difference of two sets $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^n$ is defined as $A \ominus B \triangleq \{a : a + b \in A, \forall b \in B\}$.

- *Cartesian Set Product* : The Cartesian set product between two sets $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$ is denoted $A \times B$ and is defined in the lifted space $\mathbb{R}^{n+m}$.

- *Element-wise Multiplication* : The element-wise multiplication of N elements $a_i \in A$ for $i = 1, \ldots, N$ by a compatible matrix M is defined $M \odot A = \{Ma_1, Ma_2, \ldots, Ma_N\}$.

Stability

- *Radially Unbounded Function* : A function $f : \mathbb{R}^n \to \mathbb{R}^n$ is radially unbounded if $||x|| \to \infty$ implies $f(x) \to \infty$.

- *Local Stability of the Origin* : The origin is locally stable for the autonomous system $x^+ = f(x)$ if for all $\epsilon > 0$, there exists a $\delta > 0$ such that $||x(0)|| \leq \delta$ implies that $||x(k)|| \leq \epsilon$ for all $k \in \mathbb{Z}_{\geq 0}$.

- *Attractivity of the Origin* : The origin is asymptotically attractive for the system $x^+ = f(x)$ with a domain of attraction $\mathcal{X}$ and $0 \in \mathcal{X}$, if for all $x(0) \in \mathcal{X}$, $k \to \infty$ implies $||x(k)|| \to 0$.

- *Asymptotic Stability of the Origin* : The origin is asymptotically stable for the autonomous system $x^+ = f(x)$ with region of attraction $\mathcal{X}$ ($0 \in \mathcal{X}$), if the origin is both locally stable and asymptotically attractive.

- *Exponential Stability of the Origin* : The origin is exponentially stable for the autonomous system $x^+ = f(x)$ with region of attraction $\mathcal{X}$, if there exist two constants $\gamma \in (0,1)$ and $c > 1$ such that for all $x(0) \in \mathcal{X}$, $||x(k)|| \leq c\gamma^k ||x(0)|| \quad \forall k \in \mathbb{Z}_{\geq 0}$.

- *Class $\mathcal{K}$, $\mathcal{K}_\infty$ and $\mathcal{KL}$ Functions* : A function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is classified as $\mathcal{K}$ function if it is continuous, zero-at-zero and strictly increasing. A function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is classified as $\mathcal{K}_\infty$ if it is a $\mathcal{K}$ function and also radially unbounded. A function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{Z}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is classified as $\mathcal{KL}$ if $\beta(.,t) \in \mathcal{K}_\infty$ for every $t \in \mathbb{Z}_{\geq 0}$ and $\beta(r,.)$ is monotonically decreasing for each $r \in \mathbb{R}_{\geq 0}$.

- *Lyapunov Function* : A function $V : \mathbb{R}^n \to \mathbb{R}$ is said to be a Lyapunov function for the autonomous system $x^+ = f(x)$ if there exist functions $\alpha_1(.), \alpha_2(.), \alpha_3(.) \in \mathcal{K}_\infty$ such that for all $x \in \mathbb{R}^n$, the following two conditions are satisfied:

$$\alpha_1(||x||) \leq V(x) \leq \alpha_2(.),$$
$$V\big(f(x)\big) \leq -\alpha_3(||x||).$$

- *Input-to-state Stability* : The system $x^+ = f(x,w)$ is input-to-state (ISS) stable if there exist functions $\beta \in \mathcal{KL}$ and $\rho \in \mathcal{K}$ such that for each input $w \in \mathbb{R}^m$ and each initial state $\xi \in \mathbb{R}^n$

$$||x(k;\xi)|| \leq \beta(||\xi||, k) + \rho(||w||).$$

- *Input-to-state Lyapunov Function* : A function $V : \mathbb{R}^n \to \mathbb{R}$ is said to be an ISS Lyapunov function for the system $x^+ = f(x)$ if there exist functions $\alpha_1(.), \alpha_2(.), \alpha_3(.) \in \mathcal{K}_\infty$ and $\sigma \in \mathcal{K}$ such that for all $x \in \mathbb{R}^n$, the following two

conditions are satisfied:

$$\alpha_1(||x||) \leq V(x) \leq \alpha_2(.),$$
$$V\big(f(x,w)\big) \leq -\alpha_3(||x||) + \sigma(||w||).$$

Miscellaneous Notation

- *Predecessor Value* : Superscript '$-$' denotes a predecessor value for a time dependent variable e.g. $x^- = x_{k-1} = x(k-1)$.

Acronyms

CCI common control invariant. 88

DP dynamic programming. 30

DT discrete time. 62, 88

ENR extended Newton-Raphson. 16, 17

EPR engine pressure ratio. 6, 11, 12

FPGA field programmable gate array. 21

GMV generalised minimum variance control. 13

GPC generalised predictive control. 15, 16

GT gas turbine. 1–4, 6–11, 17, 19, 21–23, 28, 29, 35, 81, 100, 115

HP high pressure. 11, 12, 15, 98, 101, 106

ISS input-to-state. xvii, 48

LMI linear matrix inequality. 42, 43, 88

LMPC learning model predictive control. v, vii, viii, 40, 46, 47, 65, 66, 68, 75, 80–83, 109, 114

LP low pressure. 11, 12

LPV linear parameter varying. 25, 95, 117

LQ linear quadratic. ii, v–x, 16, 17, 35, 36, 39, 47, 53–58, 60–63, 77, 79, 83, 85, 86, 88, 90–93, 98, 100–110, 112, 114–116, 122, 123

LQG linear quadratic Gaussian. 91, 118

LQR linear quadratic regulator. 12, 92

LTI linear time invariant. 35, 92, 93, 115–117

MDT minimum mode-dependent dwell-time. 89

MIMO multi-input multi-output. 11–13, 18

MOGA multiple objective genetic algorithms. 10, 21

MPC model predictive control. ii, v, vi, viii–xi, 3, 4, 7, 15–23, 25–30, 32–36, 39, 40, 43, 45, 52, 55–58, 60, 62–64, 66, 70, 75–93, 95, 97–118, 120–123

NARMAX Non-linear Auto-Regressive Moving Average with eXogenous input. 15, 117

NN neural network. 15

OCP optimal control problem. 33

PE Persistency of excitation. 30–34, 36, 39, 40, 42, 44, 66, 70, 82, 109–111, 113–115

PI positive invariant. xvi, 28, 34, 48, 58, 62, 88, 92

PID Proportional-integral-derivative. 7–10, 13, 15, 18

PL preview learning. v, viii, x, 75–84, 91, 92, 95, 99–111, 114, 115, 117, 118

PRBS pseudo-random binary signal. 66

QCQP quadratically constrained quadratic program. 31

QP quadratic program. 27, 28

RCI robust control invariant. 28, 29

RGA relative gain array. 11

RLS recursive least squares. viii, x, xi, xiii, xiv, 14, 15, 35–38, 44, 45, 53, 65, 74, 75, 77, 82, 90, 102, 108–110, 114, 118, 119

RoA region of feasibility. viii, 78, 80, 81

RPI robust positive invariant set. 28, 30

RRIM Rolls-Royce inverse model. 8

SDP semi-definite program. 41

SFC specific fuel consumption. 1

SISO single-input single-output. 7, 8, 10, 13, 31

SM surge margin. 11, 12

SMC sliding mode control. 13, 14

TRL technology readiness level. 19

VBV variable bleed valve. 12, 18, 20

VGV variable guide vane. 9, 12, 17, 18, 20

Chapter 1

Motivation and Background

1.1 Motivation

All technological advancement in aero-propulsion has been driven by three simple considerations: performance, economics and safety. Commercial competition, compliance to certification regulations and environmental legislation all generate tremendous cost driven incentive for reducing gas turbine (GT) fuel consumption, emissions and unnecessary maintenance. Consequently, increasing engine intelligence and autonomy through novel sensing, computing and decision making have been identified as major technological advancements where untapped cost savings can be made [14, 31]. This observation demonstrates the biggest paradigm shift from traditional GT research which has focused on mechanical innovations. These innovations relied on a brute-force rationale which would reduce weight or permit higher thermodynamic efficiencies by increasing engine pressure ratios and turbine inlet temperatures; in the same-old Brayton cycle [24].

Engine intelligence and autonomy will not only improve the efficiency of existing designs, but may also enable commercial validity of new GT architectures, thermodynamic cycles and even future engine concepts; previously deemed too complex [103]. With the advent of unprecedented reliable and low cost processing power, this research attempts to capitalize on intelligent decision making for *thrust control*; in the context of civil GTs with access to many control degrees of freedom. Prior art utilising only the coupling between fuel-flow and nozzle area in a turbojet engine boasts up to 1% specific fuel consumption (SFC) saving at the cruise condition [86]. The additional complexities of a large civil turbofan case study are therefore used to motivate a complete set of desirable characteristics for a novel multivariable and holistic control algorithm.

1.1.1 Engine Degradation

The GT's underlying non-linear dynamics, caused by highly varied operating conditions, provide a considerable challenge for designing its control and monitoring systems [109]. Traditionally, engine control and health monitoring have been mutually isolated concepts. Target tracking would be performed using fixed control laws based on analytical and experimental results conducted during the engine design phase. These control laws would prove to be reliable, but not necessarily optimal, due to conservative design; uncertainty in working line excursions and surge margins being primary sources of conservatism (engine transients accounting for up to 20% of the margin stack-up [63]). On the other hand, when considering the health monitoring system, sensed engine parameters would have been stored for manual on-the-ground health analysis. This information would give insight into the engine's health but would only be used to at most, re-schedule maintenance or to replace a faulty component.

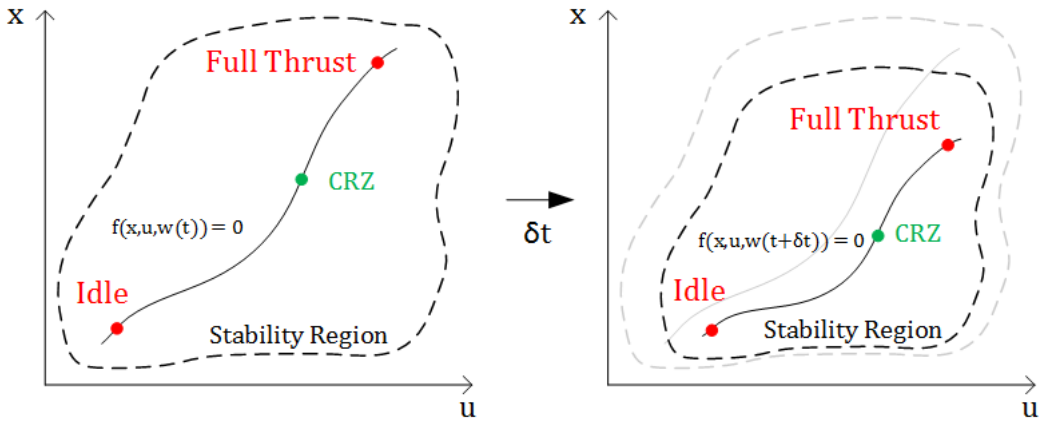

Figure 1.1: A conceptual diagram of a degrading gas turbine operating space, parametrised by a health parameter $w(t)$. CRZ denotes the steady-state operating point for cruise thrust (not to scale).

A natural progression is to integrate GT control with health monitoring by developing algorithms that allow the engine to automatically adapt to its own health condition for improved performance, economy and safety; creating an engine which could be crudely described as self-tuning and perhaps "self-aware". As GTs become more complex, complete a-priori modelling of through-life engine dynamics and health cannot be achieved through sole application of first principles physics. It is clear that for such a demanding task, a degree of supervised autonomy will have to be exercised. Before such a system is designed, let alone certified, rigorous theory and deterministic algorithms have to be developed.

1.1.2 The Theoretical Gap

Fortunately, the intersection of control theory and system identification already provides some rigorous tools for integrating control and health monitoring; adaptive control will be the focus of this thesis [7]. Despite the existing rigour, the aero-propulsion industry still finds it difficult to certify advanced controllers [109]. The shortcoming of existing self-tuning regulators is the lack of explicit constraint handling [6]. Therefore, the contribution of this work will be to provide additional theory to show that, given mild technical conditions are met during control synthesis, the proposed adaptive controller will lead to deterministic and guaranteed performance whilst satisfying ubiquitous GT constraints.

A key observation from literature is the significant interest in the use of MPC due to its explicit and rigorous handling of system constraints. Stability, robustness and recursive feasibility are gold standards when it comes to MPC, yet most real implementations of MPC may only have one, if any, of these mentioned properties. The main cause of this is the hesitance to use hard state constraints, in an attempt to avoid infeasibility of the optimisation problem that has to be solved online. Soft constraints seem to work in practice, however, formulations with guaranteed stability require additional and system-specific technical conditions [139]. Hard state constraints (specifically, positively invariant terminal constraint sets) are the academically well-known and accepted necessities for showing all three theoretical properties [93, 95].

Motivated by the observations from the GT control problem and the reviewed literature, this research proposes extensions of nominal model predictive control; the simplest form of MPC which provides the discussed gold standards *for time-invariant systems* [107]. The work avoids explicit robust formulations to treat the true time-varying dynamics for two reasons. Firstly, computational complexity in computing robust positive invariant sets, which to this date, are difficult to obtain for high dimensional systems. Even the simpler positive invariant sets (used in this work) are computationally challenging, with only recent research showing potential for easy computation [97]. Secondly, robust designs reduce performance depending on the designer's chosen degree of conservatism. Instead, the analysis conducted in this thesis is motivated by the following questions: given the initial uncertainty in the controller's plant description, how robust can the controller be without explicitly accounting for potential disturbance? Hence, can the controller be tuned to give a margin of nominal robustness that permits self-tuning? The contribution of this thesis: inherently robust adaptive model predictive control; provides an answer to this question in a gas turbine setting.

1.2 Thesis Overview

This section outlines the structure of the thesis. The first two chapters form the non-technical review of prior art which underpins the technical contributions of chapters 3, 4 and 5.

Chapter 1 - Motivation and Background

The remaining part of Chapter 1 completes the motivation with an in-depth survey of contemporary gas turbine control techniques. The survey's qualitative comparison, with respect to the desirable considerations specified later in this chapter, is used to support the choice of studying adaptive model predictive control.

Chapter 2 - Adaptive Model Predictive Control

Chapter 2 focuses the literature review on adaptive model predictive control, which forms the basis of the technical contributions in chapters 3, 4 and 5.

Chapter 3 - Learning Model Predictive Control

Chapter 3 describes the technical details of the basic learning model predictive control algorithm.

Chapter 4 - Preview Information and Inherent Robustness

Chapter 4 describes the extension of the learning MPC framework, which incorporates the benefits of utilising preview disturbance information.

Chapter 5 - Adaptation in MPC under GT Degradation

Chapter 5 completes the technical contribution of this thesis by proposing conditions under which the learning model predictive control algorithm can perform robust and safe adaptation. A case-study gas turbine is used to demonstrate the algorithm in simulation.

Chapter 6 - Conclusion and Future Work

Chapter 6 concludes the work with a discussion of the achieved result: the complete description of an inherently robust adaptive model predictive control framework. Performance improvements of the algorithm are proposed as future topics of research.

1.3 Publications

- K. Grzędziński & P. Trodden. Learning MPC: System stability and convergent identification under bounded modelling error. In 2018 Australian & New Zealand Control Conference (ANZCC), pages 125-130. dec 2018. DOI: 10.1109/ANZCC.2018.8606603
- K. Grzędziński, P. R. Baldivieso Monasterios, & P. Trodden. Learning MPC: enlarging the region of attraction using preview perturbation information. Peer-reviewed abstract and poster presented at Control 2018 (UKACC). sep 2018.

1.4 Background: Desirables for Controller Comparison

In terms of desirable algorithm properties, the previously discussed design considerations (performance, economics and safety) can be further narrowed down to three sub-categories: performance and optimality; robustness and reconfigurability and finally, the resultant design and implementation. These categories have been defined in the following paragraphs such that the proposed control methods can be compared using common themes.

1.4.1 Performance and Optimality

Optimality is concerned with the maximal obtainable performance and efficiency of the gas turbine. The ideal engine controller should make sure that the engine produces the demanded thrust, using the minimal amount of fuel, whilst minimising engine degradation and respecting engine design constraints. Even though GTs are predominantly designed to operate at several distinct operating points such as cruise, hot day take-off and top of climb conditions, control during and between these operating points should be optimal with respect to a chosen metric.

Multivariable Control

With an increase in control variables on newer engine designs, utilising multivariable dynamics can be used to improve performance [118]. Historically, fuel flow into the combustion chamber was the only controllable input of the GT, with shaft speed or engine pressure ratio (EPR) being a proxy for the output engine thrust, which is not a directly measurable quantity. Fuel flow retains the position of dominant control variable as it dictates the energy content within the GT's gas path. The remaining control parameters determine how much energy can be extracted from the constrained aero-thermodynamic environment produced by the engine geometry and thermodynamic state of the incoming flow; to produce the desired thrust. Hence, it can be presumed that there is some optimal combination of all control variables, for a given operating condition, such that the energy extracted from the fuel is maximised. These complex interactions can only be captured in a multivariable framework.

1.4.2 Robustness and Reconfigurability

Uncertainty from environmental disturbances, sensor noise and even human error is omnipresent; it must be acknowledged when designing the GT controller. As

a safety critical component, the controller must prove its robustness by demonstrating the ability to handle operating conditions associated with the airframe's mission profile. This requires certification via means of mathematical proof, simulation and physical experiments demonstrating performance within the expected flight envelope, whilst complying with legislation under civil aviation authority codes. These same guarantees must also be ensured for all reconfigured instances of the controlled system.

1.4.3 Design and Implementation

The control algorithm must be implementable and certifiable, with straight-forward tuning rules that can be easily understood by industrialists that do not need detailed understanding of the underlying mathematical theory. This requires simple, tractable and repeatable control algorithms that guarantee performance margins, when implemented on hardware with modern processing capability.

1.5 Background: Gas Turbine Control Frameworks Survey

1.5.1 Introduction

The aim of this section is to outline the main control frameworks that are in active use and under research for GTs. Underlying theory and related applications will be briefly discussed to reveal some application specific subtleties that are relevant to GT control. The section will therefore highlight the advantages and disadvantages of each respective control algorithm, leading to an informed decision to further study MPC.

1.5.2 Classical Control

Proportional-integral-derivative (PID) controllers are employed on multivariable systems using two methods. Firstly, multivariable system dynamics can be decoupled to produce multiple single-input single-output (SISO) systems that can be handled directly with traditional loop shaping techniques. However, the designer must be aware that this may introduce right-half plane zeros into the model which limits the allowable controller gains for system stability [8]. On GTs, these SISO controllers have been historically designed for bandwidths up to 10 Hz which accounted for the slower inertia and pressure dynamics of the GT [63]. Secondly, a cascaded series of control loops have been proposed for use with a supervisory multivariable controller; [74] implemented a supervisory model predictive controller to generate output set-points as references to the low-level actuator PID

controllers. This preserves existing system architectures whilst introducing coordination between control loops. However, the low-level controllers still require additional tuning to avoid saturation.

Whilst PID control can be applied to non-linear systems, classical design techniques used to compute gains cannot be directly applied to non-linear systems, therefore, generating a controller based on a linearisation of the system becomes a necessity. PID gain-scheduling can be used for a set of linearised system dynamics at pre-defined equilibrium points where the sequence in which controller gains are activated depends on a scheduled parameter (chosen by the designer). The scheduled parameter is used to describe the current state of the system and its associated linear neighbourhood. As soon as a threshold on the scheduling parameter is exceeded, the controller associated with the newly entered valid linearisation region is activated [75]. An applied example of this type of control is the Rolls-Royce inverse model (RRIM) controller and is currently used on the Trent 1000 engine series [88].

The RRIM controller splits the engine's non-linear dynamics into steady-state and transient dynamics; look-up tables generated from high-fidelity propriety engine models are used to model speed and acceleration demands to fuel flow (as opposed to modelling fuel flow to speed/acceleration) [33]. Proportional gain is scheduled on the rotational speed of the high pressure shaft. RRIM control shows that it is able to provide equivalent transient performance and robustness as the classical PI engine controller, but with a significant reduction of tuning parameters; a reduction from a previous 54 to 5 [89]. This was achieved by using the inverse model and by using non-dimensional thermodynamic parameters, removing the control loop's dependence on flight altitude. However, no remarks have been made as to how RRIM can be integrated with an *active* health monitoring system. There is clear limitation on the use of stored look-up tables; engine parameters will change over time, which will degrade the controller performance if the models are not updated.

It is important to note that the main reason why this method has been used in industry is that it is based on well-understood classical control theory that is readily certifiable. PID control retains its popularity in the aerospace industry due to its mature analysis tools for performance and stability. However, this body of literature primarily focuses on the SISO applications, which contradicts its use for inherently multivariable GT architectures; the decoupling required to apply the SISO techniques may cause performance loss as it may be difficult to fully diagonalize the plant at all frequencies [118].

Additionally, PID control is prone to controller saturation which leads to the

phenomenon of integrator wind-up. A range of anti-wind-up schemes have been developed as an ad-hoc treatment of the issue [125], including more sophisticated methods using linear matrix inequalities and optimisation [122].

System safety constraints have been handled using minimum-maximum selection logic [110]. In this technique, there are regulators associated with system output limits. These regulators produce allowable control signals which are then compared against each other using min-max logic. A selection of the most appropriate control signal is made which prohibits a violation of the nearest to violation limit. Whilst this technique is shown to work in GTs, there is little academic proof that guarantees stability between limit switches [8]. Similarly to the treatment of integrator wind-up, min-max selection logic is an ad-hoc addition for a classical controller that is unaware of the system's inherent constraints.

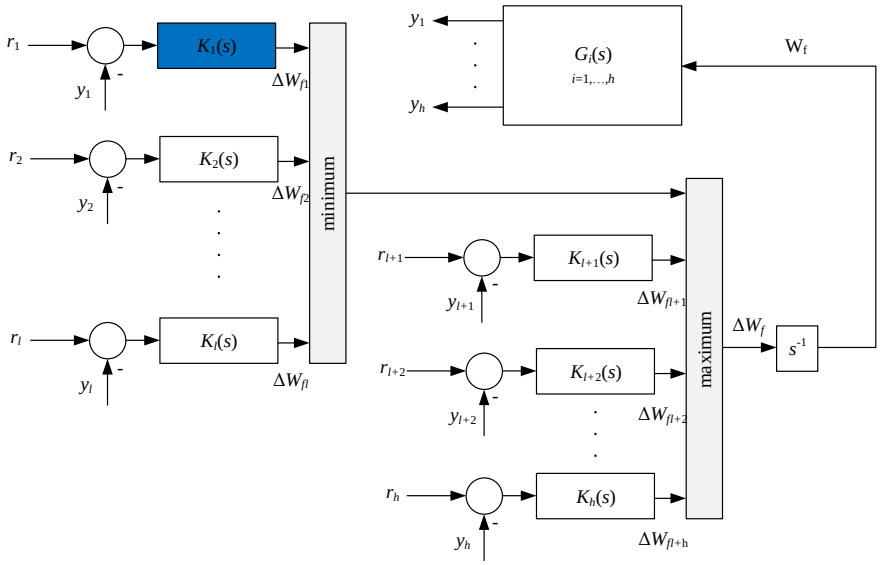

Figure 1.2: The min-max limiting logic used for constraint handling in contemporary gas turbine control. The compensators $K_i(s)$ are used to regulate set-points r_i for the chosen measured outputs y_i, including controlled and limited outputs. The blue block denotes the main compensator used for thrust control.

Fault tolerant control of an industrial GT has been suggested within the framework of classical PI control [115]. The suggested formulation incorporates open-loop control of variable guide vanes (VGVs); these are only used during transient demands to prevent surge. This implementation attempts to predict a set of possible faults that can occur in a GT; an appropriate family of PID controllers was designed using frequency response analysis [66]. The method suggests that a

measured off-line identified frequency response can be used to design a three-parameter PID low-order controller. However, the paper itself concludes that the design method is sensitive to the system identification method used, as some of the results showed that the identified plants were not PID stabilizable. Numerical round-off errors had an effect on predicting the poles of the identified system which raises certification issues of the proposed identification scheme.

The work described in [115] showed that the number of controllers did not have to equal the number of faulty plants. Some controllers were stabilising for multiple fault scenarios, which would reduce the storage requirements on-board a GT. The drawback of this method is that it is still fundamentally SISO and it assumes a priori knowledge of possible faults; this leaves the control method vulnerable to faults that cannot be predicted.

An unusual method of GT control system design been proposed using multiple objective genetic algorithms (MOGA) [27]. This design strategy involves a meta-heuristic process of selecting Pareto-optimal controller parameters to converge on an optimal design. In essence, instead of using a traditional compensator design and tuning methods, a selection of initially randomly generated compensator parameters is selected, mated and tested; a "survival of the fittest" approach is used to converge on an optimal controller which conforms to pre-defined performance requirements. In this approach, the GT architecture was also pre-defined and the controller was based on classical phase lead and lag compensator principles. Whilst the paper claims that MOGA is a superior multi-objective optimisation method to others mentioned in the paper, no comparison of the resultant MOGA controller was made with another control framework to show whether it actually gives better performance. Controller robustness and reconfigurability were not explicitly stated within the paper. Due to computational burden of the MOGA, on-line re-tuning of compensators by re-computing the MOGA could be deemed impractical even on modern hardware.

1.5.3 Dynamic Controllers

A multivariable gain-scheduled approach with dynamic linear controllers has been proposed with online stability verification [101]. This method was shown to work on a laboratory twin-spool turboprop with variable pitch angle blading. The GT dynamics were linearised at a chosen number of points along the equilibrium manifold. At each linearisation point, a dynamic controller was designed with appropriate gains. A stability preserving interpolation algorithm was used to move between the controllers at each one of equilibrium points. The scheduling parameter was unusually chosen to be the euclidean norm of the two spool

speeds, rather than the typical choice of fan/prop shaft speed or engine pressure ratio [109].

Rather than attempting to compute the eigenvalues of each instance of the closed loop system, proving stability was posed as an optimisation problem to show the existence of positive definite matrix which satisfies the Lyapunov equation. The optimisation was used to approximate the dynamics matrices as a linear combination of all linearised closed-loop dynamics in the operating envelope. Hence, a single positive definite matrix was computed to show stability of the entire envelope.

1.5.4 H-infinity

The H-infinity approach takes into account the worst case modelling uncertainty, disturbance and noise from the onset of controller synthesis. The general aim is therefore to minimise the infinity norm of the closed loop transfer function of a system, subject to constraints posed by the right half plane zero content of the original open loop plant. Minimising the H-infinity norm is equivalent to finding the approximate inverse of the plant, but through the use of feedback [51]. This control technique can be implemented in both transfer function and state-space formulations and it has been recommended that state-space is used for multivariable systems.

A scheduled multi-input multi-output (MIMO) approach has been described using a reduced state linear GT model [114]. However, only the surge margin of the low pressure compressor was considered as a constraint; high pressure compressor surge margin and turbine temperatures were ignored. The paper discusses appropriate input/output pairings based on eliminating right half plane zero content of each transfer function. Additionally, a relative gain array (RGA) matrix is constructed to identify the magnitude of cross-coupling between each input/output pairing. Whilst the paper describes a standard application of an H-infinity loop-shaping algorithm, emphasis is made on the selection of the best input/output pairings that leads to simplified controller design.

Engine control has been proposed using a combination of a set-point governor and a H-infinity controller [71]. The focus of this paper was on disturbances caused by variable inlet flow conditions and ice build up on the low-temperature sections of the GT. These were interpreted as linear interpolations between a nominal engine and an engine with 20% flow blockage using the C-MAPSS40k software model [46]. Numerical evidence showed that sensitivity to disturbances was dependent on the choice of output variable, which had disparate effects on the surge margins (SMs) of the low pressure (LP) and high pressure (HP) compressors. EPR,

as opposed to shaft speeds, was chosen as the measured output since it had moderate sensitivity on both LP and HP compressor SMs; surge must be avoided regardless of the component location in the engine. The paper also highlighted that at a particular steady-state condition, regulating EPR using fuel flow at the worst case disturbance, violated the surge margin. The paper suggests that either changing the EPR set-point or concurrent use variable bleed valves (VBVs) and VGVs would alleviate the detriment of the disturbances on the SMs; reinforcing the need for using MIMO control. The resulting 17th order H-infinity controller was simplified down to a 4th order controller with satisfactory performance, to reduce computational complexity. The set-point governor was used to pre-filter thrust such that SMs were not violated from large thrust demands. Trade-off between closed-loop bandwidth and surge margin was identified as a design choice.

Additionally, multivariable and fault tolerant control was developed for military twin-spool turbofans using the H-infinity framework [112]. The control design procedure used standard plant linearisation at several operating points. However, the novelty was introduced through design of four H-infinity controllers for each operating point, allowing for redundancy during sensor failure. Each controller was designed to handle pairs of measurable parameters, namely the shaft speeds, turbine exit pressure and temperature. The authors claimed that the proposed control method was computationally feasible on the full authority digital engine controller (FADEC) installed on the case study engine at the time. Additionally, the problem of constraint handling was dealt with in a traditional manner using saturation functions and anti-wind-up. However, bump-less transitions between controllers was not guaranteed.

1.5.5 Non-linear Control

A linear control law design, which utilises a non-linear approximation of the gas turbine, is proposed in [128]. The state-space model used for control design is:

$$\dot{x} = Ax + G(x) + Bu \tag{1.1}$$

with the function $G(x)$ assumed to satisfy Lebesgue integrability conditions. It is claimed, without proof or validation against real data, that the gas turbine's dynamics can be modelled in such a way. The method permits the design of a simple feedback $u = Kx$, providing that the initial states are bounded to a subset of the engine's operating space, using the Gronwall-Bellman lemma. This result allows the use of *any* standard linear control technique to choose the feedback gain e.g. linear quadratic regulator (LQR) or H-infinity etc. However, it is not obvious how

the choice of feedback gain can be utilised to maximise the controller's domain of attraction with the Gronwall-Bellman inequality. The performance of this regime is compared against the open-loop response of the non-linear case-study engine model, which given any K that provides more aggressive eigenvalues, it is obvious that the trajectories will converge quicker. Moreover, it is unclear how an increased number of input variables may complicate the tuning of such a controller.

The same research group also proposed a generalised minimum variance control (GMV) control which relies on the minimisation of the error and input covariance in a stochastic optimisation framework [129]. Under causality assumptions, the solution of a Diophantine equation leads to a non-linear compensator that does not require the use of an explicit model of the plant's non-linear dynamics; only linear disturbance and reference governor models are required. Moreover, the characterization of the non-linear controller requires an initial PID design, which is iterated upon, to obtain the non-linear compensator. However, the non-linear open-loop plant must be assumed stable for this control method to work. Such an implementation has been simulated on a gas turbine in [129] but the details are difficult to access due to the language barrier. Whilst this method does not explicitly consider constraints in the formulation, it is suggested that constraints could be included as barrier functions within the unknown non-linear term of the plant dynamics.

1.5.6 Sliding Mode Control

sliding mode control (SMC) involves the concept of states which "slide" across a hyperplane defined in the system's state-space. This hyperplane is used to denote the desired state-trajectory. The control input polarity is dependent on which side of the hyperplane the system states are in as it progresses through time. Further details on SMC can be found in the following survey [2].

Engine control has been proposed using a combination of SMC and H-infinity control [110]. Whilst SMC does support MIMO control, this controller explicitly assumes a SISO system with fuel flow and fan shaft speed being the input and output respectively. This implementation of SMC describes a stability proof, the first for SMC in gas turbine control. Faulty engine conditions were only modelled as additional exogenous inputs but modelling uncertainties were not discussed.

An inherent problem with SMC is known as chattering (noise in the state trajectory) which is caused by the mismatch between its discrete time implementation and the continuous nature of physical systems [73]. Significant theoretical efforts have been made to suppress this problem in control design. The paper by [110] does not explicitly state how this issue was mitigated even though the simulations

show smooth trajectories.

An application of SMC was applied with a novel modelling method where the rotational energy content of the shafts are the regulated parameters, rather than the typical fan shaft speed approach [130]. The paper suggests that overshoot free response can be guaranteed, something that until now, has been omitted from controller design. The paper presents a theoretical guarantee of this approach, but only in the case of fuel flow as the single input. However, the energy modelling grossly oversimplifies the gas dynamics at each stage of the GT by simply using the steady flow energy equation; normally these stages are considered using component maps.

1.5.7 LQR and LQG

The linear quadratic regulator (LQR) utilises a mathematical optimisation over a future horizon (typically infinite), which generates an analytically derived feedback control law [118]. The linear quadratic Gaussian (LQG) paradigm is an extension of LQR but with the addition of a state estimator that provides an estimated full state vector from measurable outputs. This is especially useful in turbofan control as some parameters cannot be directly measured within the engine; for example thrust, turbine inlet temperature and the compressor surge margins. The Kalman filter estimator includes Gaussian probabilities of noise and disturbances entering the states and outputs; this type of observer based control was shown to have no stability margins as opposed it its pure LQR counterpart [38].

LQG control has been applied to a 2-spool turbofan engine with 9 control parameters, including fuel flow, various guide vanes and variable duct geometries [34]. Singular value decomposition was used to reduce the linearised engine model to a 3 input 3 output system; only the influence of one set of guide vanes was assumed to be scheduled independently from the LQG controller. The paper claims a 20% reduction in settling time compared to gain-scheduled controller designed for the same engine but makes no remarks as to how close the controller is approaching the physical constraints of the turbofan.

Adaptive LQG was proposed using loop transfer recovery tuning (LTR) [119]. An RLS parameter estimator was used to correctly identify the gas turbine's linear model after switching occurred (unknown to the controller) between two different operating modes. After convergence of the RLS algorithm, the parameters of the newly identified dynamics were used to design an updated LQG controller with LTR. Whilst the intuition of this approach is straightforward, there are clear shortcomings; the convergence of the RLS to the true dynamics must be guaranteed. This can be ensured with an exponential forgetting factor [16]. The proposed

controller only considers switching between two neighbouring operating points; switching may induce instability if more than two, and perhaps non-adjacent, modes are entered if dwell-time arguments are not accounted for. Additionally, the excitation for this self-tuning regime is solely obtained through the switch of the operating mode. Whilst this is attractive in terms of an efficiency argument (because no additional energy is added to perturb the system inputs), only certain mode switches may guarantee frequency rich data for convergent RLS estimates.

1.5.8 Model Predictive Control

MPC originated as an industrial control technique from the chemical process industry, with its theoretical roots in optimal control theory [87]. MPC derivatives are composed of both transfer function and state-space representations, however, multivariable systems are more naturally handled in the latter form. The plant's model is used to predict its trajectories over a finite time prediction horizon. An optimisation problem is solved over this prediction horizon, using a chosen cost function with representative system constraints. This produces a sequence of optimal control inputs for regulating the system to a desired set-point. In basic formulations, only the first input of this sequence is ever delivered to the plant whilst the remainder of the sequence is discarded.

The process of prediction and optimisation is repeated over the same horizon length, at each sampling instant, yielding its alternate name of receding horizon control. The difference between MPC and other control frameworks that minimise a chosen cost function (LQR and H-infinity control) is that MPC incorporates the system's physical constraints within the resultant control law; constraint handling is by far the most powerful feature of MPC. Predictive control results in a *non-linear* control law when input and state constraints are active. This prohibits the use of standard linear stability theory for closed-loop stability analysis with MPC. Instead, more mathematically rigorous Lyapunov stability must be used [18].

The transfer function form of MPC, known as generalised predictive control (GPC), has been proposed for controlling the HP shaft speed with fuel-flow [98]. The prediction models for this controller were obtained by performing open-loop identification of a neural-network that was subsequently linearised across the engine's operating space. GPC controllers with the non-linear neural network (NN) model and linearised prediction models were compared against a gain-scheduled PID controller designed with the Ziegler-Nichols tuning criterion. The controllers were implemented on a Non-linear Auto-Regressive Moving Average with eXogenous input (NARMAX) representation of the gas turbine, obtained from real-data of a Spey engine test. Both GPC controller showed superior transient performance

compared to gain-scheduled controller, however, only the GPC with linear prediction models showed potential for real-time performance. However, there are no guarantees that the models identified from initial test-bed experiments will ever be representative of a real engine after degradation.

As observed in [98] and in many other time critical applications, MPC is a computationally demanding control method. In the so-called implicit formulation, MPC control requires optimisations to be solved on-line and faster than the chosen sampling interval. The optimisation problem must be carefully constructed for it to produce desired control performance; minimising an arbitrary cost function will produce a mathematically optimal input, but in reality, the result may not be physically meaningful. Common practise is to use a quadratic function (as an analogue of system energy), which can be conveniently used as a Lyapunov function for guaranteeing theoretical stability [68]. In other instances, the 1-norm cost function has been used in spacecraft orbit MPC to penalise absolute values of actuation effort [54]; this provides superior propellant saving but lacks the attractive properties of continuity that are essential for proving closed-loop Lyapunov stability.

Under a quadratic cost, an alternative formulation of the online optimisation problem has been proposed to ten-fold reduce the computational complexity associated with solving a quadratic program [72]. The paper proposes an alternative (but equivalent) to the standard linear quadratic online MPC optimisation. By optimising over perturbations to a known feedback law (obtained by solving infinite horizon LQ problem) the computational complexity of an online quadratic program is avoided. These perturbations are only used to handle constraints i.e. enforce feasibility of the control law $u = Kx$ that would otherwise violate constraints, had there been no perturbation of the control law (e.g near system constraints). The constraint imposed on this "closed-loop" paradigm ensures that membership of a terminal set is invoked at the consecutive time-step, rather than the terminal state in the prediction horizon. The inherent conservativeness of this approach is relaxed by solving an additional optimisation, online, to compute a sufficiently scaled-down version of the perturbations which still guarantees feasibility, but is less detrimental to the optimality of the approach. The simple approach of scaling the perturbation leads to solving a quadratic and a set of linear equations in one variable. A Monte Carlo study shows the comparison of standard LQ-MPC against the proposed extended Newton-Raphson (ENR) MPC showing marginal (if any, depending on system parameters) cost differences to the standard MPC. However, the conservativeness resulting from the use of ellipsoidal sets is not fully disclosed; the regions of attraction of the LQ-MPC and ENR-MPC

are not investigated. Moreover, the cost similarity between the LQ-MPC and ENR was observed to occur for systems where the invariant ellipsoids were “better fits” to the maximal positive invariant sets.

A major drawback of MPC is that the optimisations can become infeasible in the presence of large uncertainty. This translates to an MPC controller which either does not produce a new input or produces an input which violates constraints. Obviously, both cases can lead to undesirable if not catastrophic system performance. One suggested solution is to use "soft" constraints; these are constraints which can be varied in time to force optimisation feasibility [106]. However, softening constraints will always be limited by the maximum that is physically possible by the plant. Other suggestions include explicitly accounting for bounded disturbances, which adds conservatism by tightening original plant constraints by a factor related to the maximal expected disturbances [94] [42]. In this case, the occurrence of a worst case disturbance will still lead to satisfying the plant’s constraints. Ensuring recursive feasibility for a nominal MPC controller has been thoroughly demonstrated in MPC literature [95], however model-mismatch and disturbances can lead a nominal MPC controller out of its feasible operating regions. An attempt at creating "feasibility certificates" was proposed by using an algorithm which can test initial states for whether the MPC controller will lead to infeasibility [81]. Unfortunately, the algorithm lacks practicality as it requires the solution of an additional mixed-integer optimisation which only leads to further computational demand.

Despite the computational demand, there has been a growth in interest for applying MPC to GT engines due to continual computational improvements in hardware. Whilst it can be speculated as to when the first MPC controller was introduced on a GT, due to the attractiveness for military GT applications, MPC was experimentally shown to work on an industrial GT in real-time with a sampling interval of 1.5 seconds in the late 90s [40]. However, some constraint violation issues were reported during aggressive control demands. Additionally, this implementation of MPC only handled a small portion of the GT architecture; VGVs were scheduled independently.

A non-linear MPC formulation was simulated on a low by-pass turbofan with a sampling interval of 10 miliseconds, state dimension of 8 and only two control input variable: the fuel flow rate and the exhaust nozzle geometry [23]. Constraints on surge, temperature, shaft speed and acceleration were all shown to be satisfied. However, the authors highlighted that their simulation times were not possible in real-time and that "the average execution time for this algorithm running in MATLAB and Simulink in an uncompiled state on a Pentium 3 500 MHz with 512 MB

of RAM was 700 ms". Considering this work was conducted over a decade ago, it is reasonable that MPC should soon become computationally feasible even on systems with fast dynamics.

A method for combining MPC with fault tolerant control was first described by [69]. Furthermore, a theoretical example has been described for aircraft longitudinal control with unanticipated multiple actuator failures [136]. The paper suggests that Gaussian processes can be used as a framework to detect a fault from which the model, constraints or even cost function can be changed appropriately to account for the physical change to the system. Even though MPC supports reconfigurability for fault tolerant control, it may not always be obvious as to how the system formulation should be changed to account for a fault. Moreover, it can be difficult to show a-priori that an added fault constraint will still admit a feasible solution to the online optimisation. The main issue from a certification standpoint is that there is no academic proof that such a system will work in practise; the author argues that the problem is inherently dealing with an event that cannot be reasonably expected and would be the last line of defence to prevent a catastrophic failure. Whilst it was shown that the example used was not possible to execute in real-time, the fault tolerant MPC was able to adapt to a faulty actuator and consequently track a set-point after 25 seconds of simulated data, sampled at 0.2 seconds. Hence, the practicality of this method is still limited by computational requirements.

Active FTC using MPC on the C-MAPSS40k engine model, which was previously described in this literature review with H-infinity control, has been simulated showing superior performance to the traditional gain-scheduling approach [113]. This model allows MIMO control using fuel flow, VBVs and VGVs however, the MPC controller implemented only regulated fan shaft speed using fuel flow rate. A sampling interval of 15 ms with control and prediction horizons of 1 and 50 respectively was used with linearised models. This is computationally significant when compared with the work by [40], which used 1.5 seconds sampling interval and horizons of 13 time-steps. Consequently, a 60 second simulation took 143 seconds to compute using a constrained quadratic cost function, as opposed to the classical gain-scheduled controller which only took 3.3 seconds. This paper shows the first instance of MPC being directly compared against a PID gain-scheduled controller using the same engine model and fault scenarios. The paper presents simulations with improved performance, in terms of response time and fuel consumption reduction compared to the gain-scheduled engine controller after the fault occurred. However no tuning methodology was presented. The MPC formulation presented is limited in formal proofs of recursive feasibility for changing

fault scenarios. The simulated results showed that even without re-tuning the MPC controller, desired thrust demands were still met under faulty conditions, unlike with the gain-scheduled approach.

Explicit MPC on GTs has been discussed as a potential route for reducing computational on-line burden [121]. This method partitions the optimisation parameter space into several regions in which a dynamic programming solution of the constrained optimisation problem is computed off-line. The optimal control law is selected based on the portion of the state-space that the GT is operating in, leading to a simple look-up table. This method was shown to provide a significant time advantage over a conventional approach of solving a quadratic program on-line. The drawback of depending on off-line computed controls is that the off-line control laws are computed based on an offline model; the effects of degradation will not be accounted over time as the explicit MPC control laws are not re-computed.

The state-of-the-art application of non-linear MPC on a gas turbine can be seen in a recent paper by [133]. The paper demonstrates a fully functioning laboratory based GT air compressor which works in real-time to regulate compressor exit pressure whilst providing demand bleed air mass flow, in a complete multivariable fashion (all inputs used by MPC controller). Even though this is not an aero-engine application, to the author's knowledge, this is the highest technology readiness level (TRL) example of functioning MPC on a GT in the public domain. Fuel flow and bleed valve position are the control inputs. The paper emphasizes that real-time operation was achieved by formulating a non-dimensionalized high fidelity model which was subsequently order reduced through singular perturbation theory [62]. Whilst this reduced computational burden, an implication on the controller design can be observed. The non-dimensionalization requires, a-priori, constants that are obtained through system identification. This required some form of experimental data of the GT prior to the MPC control which may be costly to generate if scaled up to a full turbofan.

Learning-based Model Predictive Control (LBMPC) has been proposed as a method of reducing conservativeness of robust formulations of MPC [10]. Tube-MPC was chosen as the underlying robust framework of LBMPC [106]. Tube-MPC sacrifices optimality by assuming that control policies [†] are linear, to control the spread of trajectories during predictions over the horizon, caused by the unmeasured disturbances. A consequence of this is that tube-MPC enforces tighter constraint sets compared to the nominal formulation, which is the source of conservatism. LBMPC alleviates this conservatism by re-identifying more accurate

[†]A control policy is a sequence of *control laws* $\{\mu_0(x), \mu_1(x), \mu_2(x), ...\}$ where $\mu = Kx + v$ is assumed in tube-MPC. In standard MPC, a sequence of *inputs* $\{u_0(x), u_1(x), u_2(x), ...\}$ is just a degenerate example of a control policy, where the control laws are constants.

linear models that are intended to converge to the true dynamics of the system. A major benefit of such a controller on a turbofan is that the controller could learn the dynamics caused by the use of secondary control inputs such as the VGVs and VBVs; these dynamics cannot be accurately determined a-priori without complex engine simulations. Furthermore, the dynamics will change over time with common degradation modes such as fouling, erosion and corrosion. An example of this learning behaviour was claimed by the same authors but on a quadcopter experimental flight; LBMPC was able to overcome the ground effect through the controller's learning process [9]. One of the drawbacks of this method is that full-state measurements have been assumed to be available; this is not typically the case on a turbofan. However, the authors speculate that LBMPC is possible without full-state information through set theoretic estimation methods. This will ultimately lead to further computational demand, which is already high in nominal MPC formulations.

Table 1.1: Table showing historical computational performance of MPC.

Year Citation	Description	GT Type	Sampling Period (s)	Horizons (N_c,N_p)	Real-time?
1997 [40]	Exp - MPC	Industrial	1.5	13,13	Yes
2002 [23]	Sim - NMPC	Turbofan	0.01	?	No (0.7s)
2008 [111]	Sim - MMPC	Turbofan	?	20,20	No (0.015s)
2012 [113]	Sim - MPC	Turbofan	0.015	1,50	No (0.036s)
2015 [133]	Exp - NMPC	Industrial	0.2	2,2	Yes

The problem of simultaneously controlling and identifying a system has been coined the dual-control problem, which is what [9] tries to tackle. A persistently exciting tube-MPC method has been suggested, which partitions the control input into a regulatory and exciting part; the former is the signal that is trying to regulate the system to the origin whilst the latter is used to excite the system to identify it [59]. This persistently exciting disturbance is treated as an additional disturbance in the tube-MPC framework. The key research question is to identify the sufficient level of excitation which makes a trade-off between exciting the system enough to get a valuable model and avoiding perturbing in too much from its desired control tracking trajectory. In this paper, a non-convex constraint was enforced to ensure sufficient excitation of the system which leads to a non-convex optimisation to be solved, which is undesirable for performance, especially if a sub-optimal local optimum is found. From a practical point of view, such persistent excitation would require additional control inputs which could imply using more fuel. The fuel savings generated from obtaining a more accurate model for the controller would have to be more than the fuel used to identify it, to be of any

practical use. Additionally, with some insight into the system, it may be possible to understand for how long and how often such a persistently exciting controller should be turned on.

State-of-the-art research that is pushing computational feasibility utilises field programmable gate arrays (FPGAs) as the hardware platform [55]. Contrary to standard sequential micro-processors, FPGAs permit parallel-processing which can potentially reduce computation time of optimisations. A drawback of FPGAs is that they cannot handle floating point representations of words. However, [55] argue that most MPC formulations do not require such high levels of word accuracy provided that the optimisation problem is constructed in a numerically stable manner.

Time delay and quantization effects can become more prominent in MPC with a distributed architecture [45]. Up until this point in the literature review, strictly centralised controllers were considered. Distributed controllers are used to break down an overall control problem into a set of local control problems handled by physically separate controllers. Therefore it is critical for each controller to receive immediate information from the network of controllers as otherwise, instability may occur. A recent survey on the use of predictive control in highly complicated networks has suggested that MPC contains the necessary framework to consider these issues. Not only the constraints of the individual subsystems can be handled, but techniques for acknowledging data transmission delays and loss have been demonstrated whilst guaranteeing robust performance [85].

1.6 Conclusion

1.6.1 Summary of Control Frameworks

A range of control methodologies have been discussed to identify suitable candidate algorithms for an advanced GT controller. The summary will briefly discuss the findings of the literature survey with respect to the themes defined at the beginning of this thesis.

Performance and Optimality

MPC, MOGA, LQG and H-infinity use mathematical optimisation to compute the control law. However, MPC is the *only* framework which takes into account constraints when computing the control law. All the other methods require additional limiting logic operations, pre-filtering or tuning to avoid reaching critical design limitations, potentially reducing performance.

Robustness and Reconfigurability

Reconfigurability is supported on all control frameworks, to respectively varying degrees; by updating the transfer functions or state-space matrices with appropriate representations of the fault. However, not all faults can be treated in this way as fundamental requirements on stabilizability and detectability must be maintained. Representing the fault with a new exogenous input has also been suggested. Distinctively to MPC, reconfigurability can be explicitly handled by changing the prediction model, constraints or cost function to represent the fault; but it may be unclear as to how these must be changed.

Design and Implementation

Gain-scheduled controllers retain their dominance in the field of GT control due to their proven and certified behaviour over many years of operation. The remaining control methods are yet to be widely implemented on GTs due to fairly recent proofs of concept. More work has to be done to demonstrate stability under changing fault scenarios. Implicit MPC has the problem of requiring an optimisation to be solved on-line which requires a sensibly constructed optimisation problem that must be reliably solved in finite time, by the available computational power on the aircraft or engine itself.

1.6.2 Concluding Statement

There is clear incentive for integrating control and automated health monitoring for performance, economical and safety reasons. A multivariable approach would allow a more systematic method for including dynamic cross-coupling within the GT and would reduce design complications of traditional hierarchical loops, at the cost of requiring more computational power on-board. The performance improvements are expected from running engines closer to safety margins which traditionally, have been made conservative due to the limitations of classical control techniques. Active-health monitoring systems will therefore make the estimates of these safety margins more readily available and unique to each engine. MPC has been chosen as the candidate algorithm for further study. The main issue observed in literature is the need for guarantees under degrading gas turbine performance which motivates an adaptive MPC solution.

Chapter 2

Adaptive Model Predictive Control Literature Review

2.1 Introduction

Following the review of GT control methodologies, this chapter focuses on MPC and its adaptive derivatives. The prime motivation of studying MPC is that it allows safety constraints, numerous to the GT control problem, to be explicitly programmed into the control law without additional ad-hoc techniques (e.g. anti-windup and compensator selection logic). The main weakness: sensitivity to prediction model fidelity, motivates either a robust or adaptive approach when attempting control of a degrading engine. In this chapter, adaptive approaches are studied to avoid conservativeness that is typically observed with explicitly robust approaches.

It is therefore constructive to explore the adaptive mechanisms that enable prediction model updates. The review encompasses a wide range of adaptive MPC controllers, from active methods (using additional exploration/excitation signals) and passive methods (using existing trajectories) that enable convergent prediction model estimates to more accurate plant descriptions.

The *linear* prediction formulation is studied as most practical implementations use this form. Unfortunately, MPC is a notational disaster with many notational conventions. The chapter begins with a nominal formulation of MPC, to highlight the ingredients and technical assumptions that are required for retaining nominally stable performance, during slowly varying changes in the plant.

2.1.1 State-space Regulation

The concern of this thesis is the regulation of an uncertain plant's state $x_k \in \mathbb{R}^n$ to the origin using control input $u \in \mathbb{R}^m$ *in the presence of state and input constraints and plant uncertainty* (see Figure 2.1). Note that a change of variables can be performed to solve any tracking problem ($z_k = x_k - x_{\text{target}}$ and $v_k = u_k - u_{\text{target}}$) as an equivalent state error regulation problem.

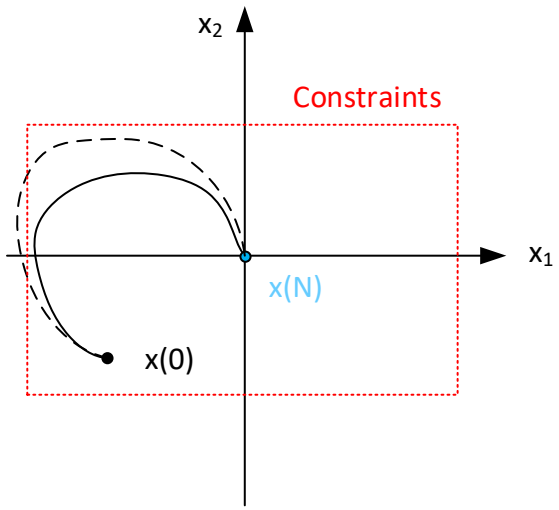

Figure 2.1: Phase portrait of a constrained regulation problem. Dashed closed-loop trajectories denote violation of plant constraints under an inaccurate model-based controller.

2.1.2 Plant Dynamics

In discrete-time, the uncertain plant's consecutive state is defined as:

Definition 1. *(True State Evolution) State-space definition of uncertain plant dynamics.*

$$x^+ \triangleq x_{k+1} = f(x_k, u_k, w_k) \tag{2.1}$$

with $w \in \mathbb{R}^p$ denoting the unknown disturbance, applied at the current time-step k and has the property of a Lipschitz continuous function.

Assumption 1 (Lipschitz continuity of the true state evolution)**.** *The true state evolution x_{k+1} satisfies $||x_{k_2} - x_{k_1}|| \leq L_f||k_2 - k_1||$ over $k_i \in \mathbb{Z}_{\geq 0}$, with Lipschitz constant L_f.*

Note that k may be omitted for notational clarity and will be re-introduced when reinforcement of time dependency is crucial.

In reality, the true (and possibly non-linear) plant dynamics (2.1) are not necessarily known exactly at k. To simplify analysis of (2.1), the following assumptions are introduced:

Assumption 2. *Reachable Linear Plant Dynamics*

$$x^+ \triangleq Ax + Bu + Gw \tag{2.2}$$

The system matrices A,B and G are compatible with their respectively multiplied vectors and describe the dynamic response of (2.1)*. Moreover, the matrices* (A, B) *are reachable but with unknown parameters; the parameters of G may also be unknown.*

Operation near a single operating condition *e.g.*, an equilibrium point of system (2.1) is, in practice, sufficient for Assumption 2 to hold. However, aerospace systems are non-linear over their operating envelopes and therefore it is common to approximate the non-linear dynamics as a linear parameter varying (LPV) model [39]. With an appropriate LPV modelling exercise and design of the scheduling parameter, switching between models within the LPV model set improves fidelity whilst retaining simplicity for computation and control synthesis.

The linearity assumption enables the results in this thesis; a proof-of-concept adaptive controller for a single operating point, which in theory, could be extended for an entire LPV representation. For now, this assumption is adequate when using MPC around a single known equilibrium point of the system [87].

Traditionally, the parameters of the system matrices around a chosen equilibrium would be obtained from standard linearisation methods; using Taylor series expansion or classical system identification performed on informative experimental data.

Assumption 3. *System Structure, State Availability and Constraints*

- *The state, input and disturbance dimensions n,m,p respectively, are known and remain constant. Moreover, the state x is measurable at every sampling instant.*
- *The state constraint* $x \in \mathbb{X}$ *is polyhedral; the input and disturbance constraint sets* $u \in \mathbb{U}$ *and* $w \in \mathbb{W}$ *are polytopic; each set contains the origin in its interior.*

The goal of MPC regulation is to find a control law $u = \kappa(x)$ online, using an explicit model of (2.2), such that the closed-loop $x^+ = f(x, w)$ is driven to the origin, or at least, the smallest neighbourhood of it (see Figure 2.1) whilst satisfying state and input constraints.

2.1.3 Prediction

An explicit prediction model is used to predict the evolution of (2.2) from the current state measurement x_k, under an optimal (depending on the metric) choice

of inputs, over a finite number of time-steps into the future $\{x_{k+1}, \ldots, x_{k+N}\}$. The prediction model inherits the same structure as (2.2) and is used to define the *predicted state*:

$$\bar{x}^+ \triangleq \bar{A}x + \bar{B}u + \bar{G}w. \tag{2.3}$$

Notice that a prediction model (2.3) can differ in parameters to the true plant dynamics (2.2), at a particular equilibrium point. Nevertheless, the parameters are time-invariant with respect to each equilibrium point. This implies that the degradation considered is time-invariant with respect to the plant's transients, which is reasonable when studying long-term wear-and-tear.

Assumption 4. *(Nominal Model Predictive Control) The external disturbance acting on* (2.2) *is zero.*

$$\bar{x}^+ \triangleq \bar{A}x + \bar{B}u \tag{2.4}$$

Nominal MPC considers the problem of regulation without explicitly accounting for *external* disturbances within the prediction model, as in (2.4). In summary, the dynamics (2.2) denote the true plant behaviour whereas the predictions from (2.4) are used to evaluate the cost of candidate control policies that steer (2.1) to the origin; the optimal sequence of policies found through an online optimisation.

2.1.4 The Online Optimisation

The nominal MPC regulator $u = \kappa_N(x)$ is an implicit control law that is obtained from an online optimisation problem that is solved at each time-step. At a state x, the problem is

$$\mathbb{P}_1(x)\colon V_N^0(x) = \min_{\mathbf{u}} \{V_N(x, \mathbf{u}) : \mathbf{u} \in \mathcal{U}_N(x)\} \tag{2.5}$$

where the decision variable is the sequence of controls over the prediction horizon of length N, *i.e.*, $\mathbf{u} = \{u_0, u_1, \ldots, u_{N-1}\}$. The cost function is

$$V_N(x, \mathbf{u}) = V_f(x_N) + \sum_{i=0}^{N-1} l(x_i, u_i)$$

where $V_f(x_N)$ is the terminal ($i = N$) cost function and $l(x_i, u_i)$ denotes the stage cost i.e. the cost at each point along the prediction horizon. The infinite-horizon linear quadratic cost function is equivalently denoted using a finite-horizon representation

$$V_N(x, \mathbf{u}) = x_N^\top P x_N + \sum_{i=0}^{N-1} x_i^\top Q x_i + u_i^\top R u_i \tag{2.6}$$

with $l(x_i, u_i) = ||\bar{x}_i||^2_{2,Q} + ||u_i||^2_{2,R}$ for $i = 0, \ldots, N-1$ and weighting matrix $P \succ 0$ for the terminal stage. The cost function is The admissible input set $\mathcal{U}_N(x)$ employs the nominal prediction model (2.4) as part of the following constraints:

$$x_0 = x \tag{2.7a}$$
$$\bar{x}_{i+1} = \bar{A}x_i + \bar{B}u_i \tag{2.7b}$$
$$\bar{x}_i \in \mathbb{X} \tag{2.7c}$$
$$u_i \in \mathbb{U} \tag{2.7d}$$
$$\bar{x}_N \in \mathbb{X}_f \tag{2.7e}$$

The optimal control problem $\mathbb{P}_1(x)$ is therefore a conventional linear quadratic MPC one, albeit with terminal conditions. The terminal ingredients P and $\mathbb{X}_f$ are designed in the well-established way that guarantees stability and recursive feasibility *in the absence of any uncertainty* (*i.e.*, $\bar{A} = A$ and $\bar{B} = B \implies \bar{x}_i = x_i$) [106]:

Assumption 5. *(Terminal Dynamics and Terminal Control Law Design - Mode 2) $\mathbb{X}_f$ is a PI set for $\bar{x}^+ = \bar{A}x + \bar{B}u$ under some stabilizing control law $u = Kx$; that is, $(\bar{A} + \bar{B}K)\mathbb{X}_f \subseteq \mathbb{X}_f$. Moreover, $\mathbb{X}_f$ is a polytope, contains the origin in its interior, and is admissible with respect to the constraints ($\mathbb{X}_f \subseteq \mathbb{X}$ and $K\mathbb{X}_f \subseteq \mathbb{U}$). The matrix P is the solution of the Lyapunov equation under the same $u = Kx$:*

$$(\bar{A} + \bar{B}K)^\top P(\bar{A} + \bar{B}K) - P = -(Q + K^\top RK)$$

Employing Assumptions 2-5, the optimal control problem for regulation $\mathbb{P}_1(x)$ becomes a quadratic program (QP) that is readily solvable using efficient convex solvers [20]. The solving of $\mathbb{P}_1(x)$ yields the optimal control sequence $\mathbf{u}^0(x) = \{u^0_0(x), u^0_1(x), \ldots, u^0_{N-1}(x)\}$; taking the first control input of the sequence optimized at each sampling instant and applying it to the system then defines the implicit feedback control law

$$\bar{\kappa}_N(x) = u^0_0(x). \tag{2.8}$$

The region of attraction of this control law (and value function $V^0_N(x)$) is, by definition, $\mathcal{X}_N = \{x : \mathcal{U}_N(x) \neq \emptyset\}$. This formulation of MPC, also known as dual-mode MPC, provides exponential stability guarantees of the origin under no uncertainty, in the linear framework [106]. Furthermore, dual-mode MPC possesses a degree of inherent robustness to small uncertainties [106]; the key property exercised in this thesis. Given that an adaptive controller has to deal with degrading plant

dynamics, it is constructive to study the implications of plant uncertainty on this nominal MPC recipe.

Changing Model

The initial prediction model of the GT dynamics will be out-dated by changing parameters within (2.4) caused by degradation. The GT monitoring system must be capable of tracking these changes. A basic requirement is that the parameters of the newly identified models lead to reachable systems. Given that degradation is restricted to "small" changes in parameters, such an assumption is not restrictive and may only be of concern when considering faults that lead to lost control degrees of freedom. These types of faults are out of the scope for this work.

It is important to consider the extreme case of the true system becoming open-loop unstable, despite a small change of model parameters. An unstable prediction model will cause open-loop predictions to diverge for a large prediction horizon. This in turn may lead to numerical conditioning issues when attempting to solve (2.5) as a QP. However, this can be alleviated by pre-stabilising the system and therefore predicting the pre-stabilised system trajectories. This pre-stabilising feedback law computation still requires an accurate model.

Changing Constraints

Not only do the constraints define the physical boundaries of the controlled system, but they are implemented as a means to guarantee its theoretical stability. Terminal state constraints ensure that the mode-2 stabilising control law does not violate constraints [95]. The terminal constraint set $\mathbb{X}_f$ defines the maximal positively invariant set of states in which the chosen terminal feedback control law produces feasible control inputs. This set would be classically determined a-priori by iteratively solving linear programs till convergence [48].

Supposing that degradation can be represented as a change in the linear model assumption 5 may no longer hold as $\mathbb{X}_f$ is not guaranteed to be PI for the newly degraded system. This motivates on-line re-computation of the terminal set, which is even more computationally demanding if the system is high-dimensional. In the context of robust control, rather than using an iterative approach proposed in [48], recent work has proposed finding the terminal robust positive invariant set (RPI) set by solving a single linear program [126]. Whilst the proposed method offers an attractive simplification of computing minimal robust control invariant (RCI) sets, the method assumes an a-priori number of inequalities to represent the outer approximation of the minimal RCI set. Without the algorithm in [48], it is currently

not possible to find an appropriate number of inequalities a-priori. Despite these difficulties, new results show promise in reducing the computation time of RCI sets; regardless of system dimension [97]. Online learning of the controllability and invariant sets using machine learning are being explored as alternatives to this fundamental problem [25].

2.2 Adaptive Model Predictive Control Review

Adaptive MPC was first demonstrated to work within the process control industry [108]. The problem of controlling relatively slow industrial processes, unlike gas turbines, lends itself well to the computational complexity of MPC. The dual-control problem of simultaneous identification and control was then not deemed necessary due to safe operating conditions during open-loop model re-identification. Clearly, this type of adaptive MPC is not possible in the GT control problem; degradation must be accounted for in closed-loop and in real-time.

Adaptive MPC, for true dual-control, gives scope for improving MPC controllers without bringing the plant off-line for model maintenance. The largest time requirement in commissioning of all MPC controllers is from the modelling of a satisfactory prediction model [120]. Given that a plant change occurs, due to a newly installed feature or slowly time-varying change, the cost of re-modelling using ad-hoc identification experiments may be economically infeasible. Hence, adaptive MPC can be used to reduce the effect of uncertainty in plant-model mismatch by automatically learning the best model within a chosen model structure [44].

Whilst conventional MPC theory has been firmly established in academia, industrial applications are still limited to basic formulations [39]. Robust, stochastic and adaptive flavours of MPC have all emerged to tackle modelling uncertainty, despite industry still being hesitant in using basic stability ingredients such as terminal invariant sets, as described in the introduction of this chapter [17, 92]. The issue of computational complexity is of main concern for industry, with significant research efforts proposing new solution methods and hardware implementations [55, 64, 67, 111]. Given the observations in table 1.1, computational concerns are becoming less of a priority.

Despite attention being focused on the computational implementation of MPC, the basic formulations have a fundamental problem; limited stability guarantees and potential infeasibility of the online optimisation under prediction model changes [87, 135]. Extensive simulations of possible operating scenarios have to be conducted to ensure that the controller does not become infeasible [39, 81]. In robust

MPC, the computation of the RPI sets requires knowledge of an accurate model for the invariance properties to be retained [70]; large model mismatch leads to large disturbance sets which increases conservatism when larger RPIs are used for constraint tightening. Clearly, this motivates the need for a dual controller that can match the prediction model with the true dynamics. Once a "good" enough model is obtained, a robust design can be implemented without introducing unnecessary conservatism (i.e. robustness to a bounded disturbance rather that the sum of a bounded disturbance *and* model mismatch).

2.2.1 Dual-control

A dual controller needs to provide a control signal that simultaneously regulates and excites. Intuitively, the excitation should minimise the deterioration on control performance whilst maximising the richness of the signal for parameter identification. *Persistency of excitation (PE)* of the regressor is an established ingredient for discriminating between linear models within a chosen model structure [49].

However, the idea of dual-control is argued to have a common misconception [56]: the presumption of an intuitive trade-off between a system's control performance and its exploration. In the framework of receding horizon control, the optimal control sequence can be found using dynamic programming (DP). Deviations from the optimal trajectories (in the classical DP sense), caused by exploration, will seemingly induce sub-optimal trajectories. This perceived sub-optimality is argued to occur because the classical optimality principle does not account for potential reduction of uncertainty in future predictions through exploration in the present. A re-formulation of the optimal control problem to include optimisation over *future uncertainty* permits the following example: performing "sub-optimal" decisions in the present (in the classical sense), to make an optimal control decision in the future (accounting for uncertainty), which is better than the optimal decision had their been no effort to reduce future uncertainty. A key issue is one of ambiguity; that is, how to model the uncertainty within this newly defined dual-control approach? Given that it is well known that optimising over functions in an MPC framework to choose the optimal feedback policy is intractable (optimising over functions is an infinite dimensional problem) [106], optimising over potential model structures shares exactly the same problem. With this observation and the little-known literature on variable structure MPC controllers, the dual-control problem will be considered in its classical sense of [41].

2.2.2 Excitation: Persistency and Optimality

Promoting excitation for closed-loop identification has been suggested through minimisation of the parameter covariance and also metrics associated with the chosen identification algorithm's variables; for example, maximising the determinant, eigenvalues and/or traces of a so-called information matrix [57, 105]. In the recursive least squares setting, the information matrix is defined as (3.7b) and is intimately linked to the estimation algorithm's "gain". In [57], online experiment design is proposed through the augmentation of a nominal MPC cost function and its constraints. The augmentation of the cost function introduces a term which promotes input signals that both regulate and excite the system, for generating frequency rich data for parameter identification. Two cost terms are compared in sensitivity studies. Firstly, the trace of the predicted parameter covariance is minimised, which introduces non-convex constraints into the optimisation problem. However, the problem can be relaxed to a quadratically constrained quadratic program (QCQP) through an appropriate reformulation; the simulations of this paper did not take advantage of this result. Secondly, the exponents of the negated diagonal elements of the least squares information matrix (as defined by (3.7b)), are minimised. Non-convexity is also introduced through performing predictions of the information matrix over the excitation horizon (an analogous concept to the prediction horizon which defines the number of future time-steps to excite the system). Through simulation it is shown that maximising properties of the information matrix are less disruptive to the regulation objective. The sensitivity study of this augmentation showed that varying the excitation horizon for values larger than unity do not benefit the identification convergence rate. However, this observation is predicated on the fact that a discount factor was used to de-emphasize future information matrix predictions in the cost function. In fact, this is contrary to the observations of [105] who argue that it is indeed valuable to predict the information matrix when maximising the minimum eigenvalues of the information matrix. For the dual controller utilizing the parameter covariance minimization, simulations show that parameters do not necessarily converge to the true values of the plant. This is explained by the fact that the notion of PE is completely ignored in this cost framework.

In [90], a backwards looking horizon is employed to ensure that the regressor remains persistently exciting whereas in [105], a SISO dual control problem is considered within an MPC framework that utilises the notion of maximising the minimum eigenvalue of the information matrix. This maximisation is used to satisfy the persistence of excitation condition. The paper suggests that maximising the information matrix over several prediction steps *does* benefit the parameter

identification process, but is only shown through simulation.

The conclusion of [43] is that optimising over a single measure for "richness" of excitation does not ensure optimal excitations; information about model parameter dynamics is not utilised, as discussed in the re-definition of the dual control problem. It is therefore sensible to interpret that the choice of excitation metric is a tuning parameter/heuristic of the adaptive MPC algorithm in the classical optimal control setting.

In [28], the authors attempt to avoid PE by storing a buffer of regressors which maximise the minimum singular value of the buffer. In [10], a non-deterministic excitation signal is accounted for using robust tube MPC [106] to update the prediction model whilst providing recursive feasibility. However, the explicit and implicit robust tubes remain conservative under the initial model parameters. Another adaptive robust approach is utilized in [3] where predictions of parameter error bounds are used to reduce the conservativeness of a min-max MPC approach.

Similarly in [131], an adaptive estimator is used to estimate the disturbance caused by mismatch. Both approaches briefly discuss the computational complexity of their respective algorithms. Moreover, computation of the predicted parameter error bounds in [3] are performed using nominal parameters, rather than the true parameters. Indirect adaptive MPC uses a polytopic linear differential inclusion description of system uncertainty that enters through the dynamics matrix [37]. Under the assumed problem structure, robust feasibility and stability results are obtained using predictions of a vector describing the convex combination of possible dynamics matrices, which must be known a-priori. However, the problem only requires the solution of a standard quadratic program.

Tracking to a PE reference input signal computed from a constrained least squares problem is shown in [26]. The resultant PE reference signal is then tracked by a standard MPC controller whilst ensuring parameter convergence under state and input constraints. The convex relaxation of the exciting reference control input has to be solved iteratively and online; convergence of the relaxed optimization problem is only briefly discussed. Similarly, model re-identification is proposed using a controller that tracks a PE reference input within a control invariant set that is claimed to be safe for identification experiments [50]. However, the control invariant set is computed based on an uncertain model (assumed strictly stable), which does not guarantee robust stability and feasibility after only performing prediction model update.

In [90], PE is guaranteed at the current time-step by taking advantage of natural transients that may already be PE, without inducing additional excitation. Hence,

the PE condition over a backward-in-time window of control signals is checked. This is a non-convex constraint on the online optimal control problem (OCP); a feasible solution of this OCP is shown to be an input that was implemented at the end of the backward-in-time window, which implicitly induces periodicity seen in classical adaptive control formulations [8]. However, the solution is feasible only under input constraints and is not optimal with respect to any excitation measure.

PE is addressed under state constraints in [60] using a tube MPC approach; it is shown that, under mild conditions, PE of the additive exciting input is transferred to PE of the regressor. Another robust dual-controller is described in [132] which modifies the regulation cost with an excitation cost that diminishes as a function of the prediction model error. Robustness is shown to an a-priori known polytopic inclusion of the system with both parametric and additive uncertainty. However, both [60] and [132] suffer from conservatism induced by using robust (control) invariant sets for constant uncertainty sets, and moreover, propose non-convex optimal control problems that are computationally infeasible in practice.

On the other hand, [123] and [82] rely on an implicit satisfaction of PE in a recursive set membership identification framework [12]. Here, the idea is to restrict uncertainty around the true dynamics to within a contracting fixed-complexity polytope; rather than attempting to identify a specific set of parameters that fit a model structure. An extension to this set identification method, using a robust tube MPC parametrization with $u = Kx + v$ as the control law, has been proposed for both time invariant and time-varying systems [83]. The linear feedback gain K is designed (for example, using H-infinity) to robustify against a bounded disturbance that, in addition to additive noise, also encompasses the modelling uncertainty. The degree of freedom v is optimised over a shrinking set of model candidates and a homothetic tube for ensuring robust positive invariance of the terminal constraint [104]. Least mean-squares estimation is used to identify point estimates within the set of model candidates, such that certainty equivalence MPC is solved. However, this terminal cost is chosen conservatively to satisfy the Lyapunov equation for the initial large set of model candidates. Conservatism is also present in computation of the terminal constraint which is the trade-off for reduced computational complexity. The work in [84] provides respite to the conservatism of [82] by optimising over the cross-section of the robust-tube, whilst maintaining the same number of inequalities. PE conditions are explicitly satisfied with an additional constraint in [83], which under the single input example, is still convex (and requires the iteration of two quadratic programs). For general multivariable systems, convexity is lost.

The results of [123] have since been updated in [124] which now guarantee

recursive feasibility for linear time-varying systems. This result is obtained by making predictions of the parameter sets using assumed bounds on the rate of change of these parameter sets. However, the proposed methodology is still conservative since the terminal constraint uses the largest parameter uncertainty set to ensure recursive feasibility. Moreover, the complexity of the parameter uncertainty polytope has to be tuned for a reasonable trade-off between conservatism and computational complexity. Generation of informative data for performing updates of the uncertainty polytope is omitted; the simulation example relies on the excitation from the square wave reference signal for tracking.

2.3 Conclusion

This chapter summarises a nominal and conventional recipe for synthesising MPC controllers that provide exponential stability whilst guaranteeing recursive feasibility of the online optimisation; in the unrealistic setting of no uncertainty. The recipe is examined for the case where a new model becomes available as a result of learning long-term degradation in gas turbines. The key issues that must be satisfied before adaptation can be performed include:

- The choice of prediction horizon N must be stabilising for the updated and potentially unstable identified dynamics.
- The Mode-2 control law $u = Kx$ in Assumption 5 must be stabilising for the newly identified prediction model; a solution is to re-compute K online for the new prediction model.
- The terminal constraint set $\mathbb{X}_f$ must be PI for the newly identified model; a solution is to re-compute $\mathbb{X}_f$ online for the new prediction model.

A range of learning mechanisms have been observed in literature which vary in implementation, computational complexity and efficacy. It is clear that in the linear framework presented in this chapter, an implicit or explicit PE condition of the regressor vectors must be satisfied to ensure convergent parameter identification. The choice of excitation cost is a heuristic within a classical dual-control framework. However, there is little (if any) recognition of the stability properties of the nominal formulation of MPC under additional excitation. Motivated by this observation, the following chapter proposes a computationally efficient (convex) precursor to a fully adaptive MPC controller which produces excitations that respect nominal robustness margins; thereby guaranteeing robust behaviour within a sub-set of the controller's region of attraction.

Chapter 3

Learning Model Predictive Control

3.1 Introduction

In the previous chapter, the adaptation of a prediction model is studied around an assumed operating point. In this chapter, the same assumption is utilised to study a generic dual-control problem for linear time invariant (LTI) constrained systems, without loss of applicability to the GT problem at hand.

It is assumed that the true (degraded) dynamics are unknown, and that only a nominal (healthy) model is available for control. Using the nominal model, a MPC based controller is designed which employs two online optimisation problems: the first problem is a conventional MPC regulation problem, albeit with stabilizing terminal conditions. The second problem adds perturbations to the optimal regulating control law in order to promote excitation of the system dynamics and facilitates accurate learning by a RLS algorithm. Maximising a deterministic analogue of an "information objective" [5] in the linear framework is implemented [91].

The mechanism behind the prediction model update with the more accurate model learned online is deferred to Chapter 5; instead, in this chapter, the focus is on what can be achieved with respect to stability and feasibility in the nominal MPC framework. The main novel result—compared with lack of stability results in similar approaches [90, 105, 137]—uses the well-established exponential stability result of LQ-MPC. With the inherent robustness of this exponentially stabilising controller, explicit bounds on the perturbation signal magnitude (in terms of model error) are derived such that, if met, guarantee closed-loop stability despite model uncertainty and the exciting perturbations.

3.2 Problem Statement

In this section, a statement of the closed-loop plant identification problem under LQ-MPC is presented. Firstly, recall the regulation problem of controlling the unknown plant

$$x^+ \triangleq Ax + Bu \tag{3.1}$$

to the smallest neighbourhood of the origin, as in Figure 2.1. The plant (3.1) represents slowly time-varying linear dynamics that have changed since initial controller synthesis (*i.e.*, slow enough to be considered constant with respect to the controller horizons and closed-loop transients). Suppose that an initial LQ-MPC controller, $u = \bar{\kappa}_N(x)$, employs the best available prediction model that was developed through off-line first principles physics and data-driven methods:

$$\bar{x}^+ \triangleq \bar{A}x + \bar{B}u. \tag{3.2}$$

Since the model (3.2) was used to compute the stabilising ingredients of the regulation optimisation $\mathbb{P}_1(x)$ (2.5) for the non-degraded plant, the theoretical guarantees of the stabilising ingredients for the closed-loop system $x^+ = Ax + B\bar{\kappa}_N(x)$ will no longer hold [106]. Moreover, performance of the MPC regulator may not match an equivalent MPC controller with prediction model (3.1). Furthermore, since it may not be possible to operate the plant in open-loop after controller deployment, conventional open-loop identification schemes will no longer be available for estimating the true plant dynamics. Therefore, this chapter proposes a modified LQ-MPC control law

$$u = \bar{\kappa}_N(x) + v, \tag{3.3}$$

where the perturbation term v is used to overcome the technical issues of closed-loop system identification, whilst retaining the attractive properties of constraint satisfaction under the MPC approach. These modifications are derived from fundamental properties of the well-known recursive least squares estimation algorithm.

3.3 Recursive Least Squares

In the deterministic linear receding horizon framework, convergent model estimates can be obtained using RLS estimation algorithm under informative data, generated from PE signals [78]. Moreover, RLS facilitates the use of parameter dynamics models, which if available, can accelerate convergence. Such an estima-

tor can be shown to be equivalent to the Kalman filter and provides the optimal trade-off between parameter tracking ability and sensitivity to noise, in terms of parameter error covariance [78]. Parameter dynamics models are application specific and extremely sensitive intellectual property, particularly in the aerospace industry. Without loss of generality, and the worst-case of no available parameter dynamics, the following section formalises the basic RLS definitions.

3.3.1 State-space Formulation of RLS

For online identification, the RLS estimation of the true dynamics (3.1) uses the following definitions of the regressor, predictor and parameter estimate.

Definition 2 (Regressor)**.** *The regressor vector used to perform the linear regression:*

$$\psi^\top = \begin{bmatrix} x^\top & u^\top \end{bmatrix} \tag{3.4}$$

Definition 3 (Predictor)**.** *The predictor, based on* (3.2)*, used to define the current state estimate:*

$$\hat{x}^\top = (\psi^-)^\top \hat{\theta}^- \tag{3.5}$$

Definition 4 (Parameter Estimates)**.** *The estimated elements of the dynamics matrices.*

$$\hat{\theta} = \begin{bmatrix} \hat{A} & \hat{B} \end{bmatrix}^\top \tag{3.6}$$

The superscript '$-$' denotes a predecessor value (the previous time-step) such that the update law equations for the estimates are

$$\hat{\theta} = \hat{\theta}^- + \mathcal{R}^{-1}\psi^- \left[x^\top - (\psi^-)^\top \hat{\theta}^- \right] \tag{3.7a}$$

$$\mathcal{R} = \gamma \mathcal{R}^- + \psi^- (\psi^-)^\top \tag{3.7b}$$

where $\gamma \in (0, 1)$ is a forgetting factor that is used to discount old measurements, whilst $\mathcal{R}$ is known as the information matrix. The inversion of the information matrix, with its recursive definition, lends itself to being efficiently computed using the matrix inversion lemma.

Lemma 1 (Matrix Inversion Lemma)**.** *Let $\mathcal{P} = \mathcal{R}^{-1}$ for notational convenience. Then, $\mathcal{P}$, can be recursively computed using*

$$\mathcal{P} = \left(\gamma \mathcal{R}^- + \psi^- (\psi^-)^\top \right)^{-1} = \mathcal{P}^- - \mathcal{P}^- \psi^- \underbrace{\left(\gamma + (\psi^-)^\top \mathcal{P}^- \psi^- \right)^{-1}}_{\in \mathbb{R}} (\psi^-)^\top \mathcal{P}^- \tag{3.8}$$

if the inversion $\mathcal{P}^-$ exists at the preceding time-step.

The proof of this statement can be found in Appendix 5. Using this lemma, the parameter update law can then be further simplified to:

$$\mathcal{L} = \frac{\mathcal{P}^-\psi^-}{\gamma + (\psi^-)^\top \mathcal{P}^-\psi^-} \tag{3.9a}$$

$$\hat{\theta} = \hat{\theta}^- + \mathcal{L}\left[x^\top - (\psi^-)^\top \hat{\theta}^-\right] \tag{3.9b}$$

It is well known that exponential convergence of the estimates $\hat{\theta}$ to (A, B) is assured if the regressor ψ is *persistently exciting* [52].

Definition 5 (Persistency of Excitation)**.** *The regressor ψ is said to be persistently exciting if for all k, there exists some integer $M > 0$ and real constants $\alpha > 0$, $\beta > 0$ such that:*

$$\alpha I_{(n+m)\times(n+m)} \preceq \sum_{i=k}^{k+M} \psi_i \psi_i^\top \preceq \beta I_{(n+m)\times(n+m)}. \tag{3.10}$$

To satisfy Lemma 1, the initial knowledge of $\mathcal{R}^-$ is required to initialize the RLS scheme with the initial inversion $\mathcal{P}^-$.

Assumption 6. *At time $k = 0$, the prior information matrix $\mathcal{R}^-(0)$ is positive definite.*

The initial value of $\mathcal{R}^-$ may incorporate any prior knowledge of the estimates and/or an optimised excitation sequence; it can be chosen freely by the designer to improve the convergence rate of RLS. However, this aspect of the design is beyond the scope of this thesis. The choice of a forgetting factor $\gamma \in (0, 1)$ will ensure that after a sufficient length of time, the initialisation will be discounted [65].

3.3.2 Closed-loop Identification

It is important to highlight an important fallacy in closed-loop identification: *persistence of excitation is not a sufficient condition for convergent parameter estimates in closed-loop identification* [78]. Under linear feedback it can be easily shown, with a simple example, that the predictor's structure leads to ambiguity.

Remark 1 (Ambiguity in closed-loop identification)**.** *The system* (3.2) *under a linear state feedback control $u = Kx$ yields the closed-loop dynamics $x^+ = (A + BK)x$; using* (3.5)*, an equivalent "open-loop" system can be constructed:*

$$\hat{x} = \hat{A}^* x^- + \hat{B}^* u^-$$

with $\hat{A}^* = (A + \sigma BK)$ *and* $\hat{B}^* = (B - \sigma B)$ *for some arbitrary* σ*. Clearly, identification under linear feedback leads to identification of the closed-loop dynamics since there is no way to distinguish between the open-loop and closed-loop dynamics for any* σ*.*

Nominal LQ-MPC is indeed a linear control law when constraints are inactive [106]. However, the affine structure of the proposed control law (3.3) alleviates the linear correlation between the measured states and inputs, through the purposefully designed perturbation term (3.13).

3.3.3 Learning Model Predictive Control Overview

With the preceding definitions, it is now possible to describe an optimisation problem that addresses the technical issues of closed-loop identification. The learning MPC controller architecture follows that of [105], with two optimal control problems solved in series and at each time-step, in order to determine the control inputs to the system for regulation and excitation. Note that unlike the previous works, the dual-control problem in the nominal setting is considered in the state-space framework; a formulation that naturally fits the multivariable problem.

First, a conventional MPC regulation problem (2.5) is solved, which implicitly defines the state feedback law (2.8). Subsequently, but at the same sampling instant, a secondary optimisation problem is solved in order to determine *perturbations* to the optimal regulating control law, that excite the system in order to promote convergent parameter estimation. Although the feedback law $\bar{\kappa}_N(x)$ regulates the system (to the origin in the absence of uncertainty), it does not necessarily provide the excitation of the system dynamics needed for convergent identification [78]. For example, consider the closed-loop trajectories under a deadbeat controller design; they may not provide enough informative data samples for the estimation algorithm before reaching the origin. Moreover, the data gathered under feedback will induce estimation bias as discussed in Section 3.3.2. Hence, immediately following the solving of $\mathbb{P}_1(x)$ at state x, which generates the sequence $\mathbf{u}^0(x)$, the secondary optimisation

$$\mathbb{P}_2(x; \mathbf{u}^0(x)) : H_N^0(x; \mathbf{u}^0(x)) = \max_{\mathbf{v}} \{H_N(x, \mathbf{v}) : \mathbf{u}^0(x) + \mathbf{v} \in \tilde{\mathcal{U}}_N(x)\} \tag{3.11}$$

is solved; where $H_N(\cdot)$ is an objective function designed to promote PE of the regressor (3.4), the decision variable is the N-step sequence of control *perturbations*, $\mathbf{v} = \{v_0, \ldots, v_{N-1}\}$, and the feasible set $\tilde{\mathcal{U}}_N(x)$ is defined by the constraints (2.7) together with the additional constraint:

$$||v_i|| \leq E \quad \text{for} \quad i = 0, \ldots, N-1. \tag{3.12}$$

This constraint permits a perturbation of the input from that of the optimal regulating control sequence; $E \geq 0$ is therefore a tuning parameter to be set by the designer, and its role in stability is analysed in Section 3.7. Assuming that the optimal solution to (3.11) achieves PE of the regressor, application of the first element of $\mathbf{v}^0(x; \mathbf{u}^0(x))$, the analogous MPC control law for excitation,

$$v = \pi_N(x) = v_0^0\left(x; \mathbf{u}^0(x)\right) \tag{3.13}$$

will ensure that the norm of the parameter estimate's error decays exponentially to zero, thus yielding an estimate that corresponds to the true plant dynamics. In practice, convergent identification must be checked with simulations; the discussion of this observation is left to Proposition 2, after the excitation optimisation has been defined with an appropriate solution method.

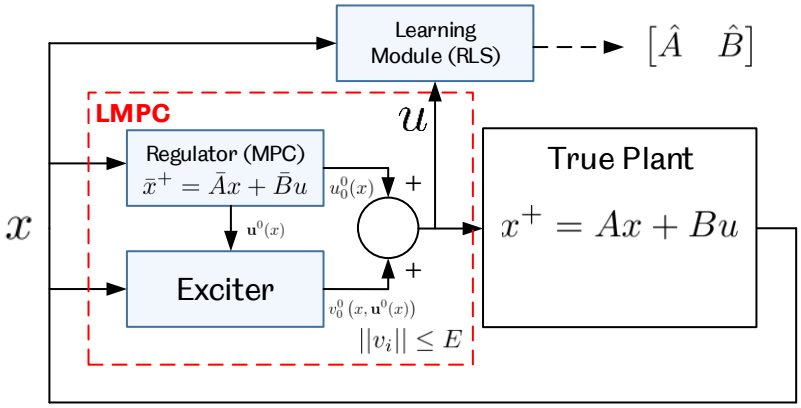

Figure 3.1: Block-diagram depicting the optimisation dependencies within the LMPC controller.

3.4 Excitation Optimisation Definition

It is easy to show that the summation of regressor outer products (3.10) and the information matrix recursion (3.7b), with $\gamma = 1$, is equal to the net change in the information matrix:

$$\begin{aligned}
\mathcal{R}_{k+1} &= \mathcal{R}_k + \psi_k \psi_k^\top \\
\mathcal{R}_{k+2} &= \mathcal{R}_k + \psi_k \psi_k^\top + \psi_{k+1} \psi_{k+1}^\top \\
&\vdots \\
\mathcal{R}_{k+M} - \mathcal{R}_k &= \sum_{i=k}^{k+M} \psi_i \psi_i^\top
\end{aligned}$$

and therefore, persistence of excitation can be ensured by achieving positive definiteness of $\mathcal{R}_{k+M} - \mathcal{R}_k$ *i.e.*, the positiveness of the smallest eigenvalue $\underline{\lambda}$ of $\mathcal{R}_{k+M} - \mathcal{R}_k$. Accordingly, with the current value of the information matrix equal to $\mathcal{R}$ (the output of (3.7b)) at the state x, the cost function of the secondary optimisation is designed as:

$$H_N(x, \mathbf{v}) = \underline{\lambda}(\mathcal{R}_N - \mathcal{R}_0) \tag{3.14a}$$

$$\mathcal{R}_0 = \mathcal{R} \tag{3.14b}$$

$$\mathcal{R}_{i+1} = \mathcal{R}_i + \psi_i \psi_i^\top \tag{3.14c}$$

$$\psi_i^\top = \begin{bmatrix} x_i^\top & (u_i^0(x) + v_i)^\top \end{bmatrix} \tag{3.14d}$$

Explicit predictions of future information matrices (3.14d) are used to provide additional degrees of freedom for the excitation optimisation, given that the regulator part will be acting against both the excitation and uncertainty.

3.4.1 Excitation Optimisation Tuning

Weighting constants σ_i can be introduced to emphasise perturbations on specific information matrices such that the cost term, (3.14a), can be modified to:

$$\sum_{i=k}^{k+M} \sigma_i \psi_i \psi_i^\top.$$

Choosing optimal weightings with respect to maximising the minimum eigenvalue, subject to a constraint on the norm of the weightings, in itself, is a semi-definite program (SDP) [4]. Such a tuning approach could be used to alleviate the problem of using open-loop predictions of the information matrix. Since the

information matrix is predicted using the nominal prediction model, there is no explicit guarantee that the future regressors of the true system (3.1) will be exciting. Hence, the weighting of information matrix predictions closer to the current time-step are more likely to be implemented in the receding horizon framework.

For more practical tuning, the designer may choose to truncate the "excitation" horizon of (3.14a) to both prioritise excitations earlier in the prediction horizon, and moreover, reduce computational demand of the excitation optimisation. For simplicity, the excitation horizon is chosen to equal the prediction horizon N.

3.4.2 Non-convexity of the Excitation Problem

Unfortunately, the product of the decision variables caused by the definition of the regressor (3.4), leads to an eigenvalue problem with linear matrix inequality (LMI) constraints and a *quadratic equality constraint* (from $\psi_i\psi_i^\top$ terms) [47]; which is a non-convex optimisation problem that does not fit any of the standard convex optimisation forms. The following section describes how this non-standard problem can be approximated and potentially solved online.

3.5 Excitation Optimisation Solution

The excitation optimisation problem $\mathbb{P}_2\left(x;\mathbf{u}^0(x)\right)$, as it stands, may be computationally expensive to implement. The following section addresses the non-linearity by utilising the optimal sequences from $\mathbb{P}_1(x)$ to form an appropriate approximation of the excitation problem, inspired by the results of [47] that also deal with a similar PE constraint. The novelty in this formulation is that, rather than enforcing a fixed magnitude of excitation as in [47], the optimiser of (3.11) is free to choose an excitation up to the bounded magnitude (3.12), providing that the state and input trajectories that develop from x satisfy their respective constraints.

3.5.1 Linearisation of the Eigenvalue Problem

The linearisation of problem $\mathbb{P}_2\left(x;\mathbf{u}^0(x)\right)$ now denoted $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$ makes use of the following regressor decomposition:

Definition 6 (Regressor Decomposition)**.** *Using the nominal optimal trajectories trajectories* $\mathbf{u}^0(x) = \{u_0^0(x), u_1^0(x), \ldots, u_{N-1}^0(x)\}$ *and* $\mathbf{x}^0(x) = \{x, \bar{x}_1^0(x), \ldots, \bar{x}_{N-1}^0(x)\}$,

the regressor ψ_i can be decomposed as

$$\psi_i = \psi_i^* + \delta\psi_i \tag{3.15a}$$

$$(\psi^*)_i^\top = [\bar{x}_i^\top, (u_i^0)^\top]^\top \tag{3.15b}$$

$$\delta\psi_i^\top = \begin{cases} [0^\top, v_i^\top]^\top, & \text{if } i = 0 \\ [\sum_{j=0}^{i-1} \bar{A}^{i-1-j}\bar{B}v_j^\top, v_i^\top]^\top, & \text{if } i > 0 \end{cases} \tag{3.15c}$$

where the subscript i denotes the index of the respective elements of the trajectories $\mathbf{u}^0(x)$ and $\mathbf{x}^0(x)$, and index $i = 0$ denotes the first element of each trajectory.

Note that the bar indicates trajectories from a nominal model MPC regulator, with the initial measured state $\bar{x}_0 = x$. The terms v_i belong to the candidate perturbation sequence $\mathbf{v} = \{v_0, \ldots, v_{N-1}\}$ *i.e.*, the decision variable of (3.11). With the decomposition defined, the approximation of $\mathbb{P}_2\left(x; \mathbf{u}^0(x)\right)$ is derived in the following proposition.

Proposition 1 (Linearisation of $\mathbb{P}_2\left(x; \mathbf{u}^0(x)\right)$)**.** *The problem* (3.11) *can be linearised to a new convex problem $\mathbb{P}_2^*\left(x; \mathbf{u}^0(x)\right)$, that lower bounds the original problem* (3.11)*, and is solvable using semi-definite programming.*

Proof. The cost function (3.14a) implies the excitation problem in LMI form [21]:

$$\max_{\lambda, \mathbf{v}} \lambda \quad \text{subject to:} \quad \mathcal{R}_N(\mathbf{v}) - \mathcal{R}_0 - \lambda I \succeq 0 \tag{3.16}$$

The constraint (3.16), using the regressor decomposition, is used to define the modified cost function

$$\sum_{i=0}^{N} \left(\psi_i^* + \delta\psi_i\right)\left(\psi_i^* + \delta\psi_i\right)^\top \succeq \lambda I$$

$$\underbrace{\sum_{i=0}^{N} \psi_i^*(\psi_i^*)^\top + \psi_i^*(\delta\psi_i)^\top + \delta\psi_i(\psi_i^*)^\top}_{H_N^*(x,\mathbf{v})} + \delta\psi_i(\delta\psi_i)^\top \succeq \lambda I$$

$$H_N^*(x, \mathbf{v}) + \delta\psi_i(\delta\psi_i)^\top \succeq H_N^*(x, \mathbf{v})$$

where $H_N^*(x, \mathbf{v})$ is a linearised version of the original excitation cost (3.14a). □

An explicit lower bound on (3.10) has been obtained by removing the new quadratic cross-terms $\delta\psi_i(\delta\psi_i)^\top$, yielding a LMI. The convex excitation problem $\mathbb{P}_2^*\left(x; \mathbf{u}^0(x)\right)$, subject to the same constraints of (2.7), can be readily solved using interior point methods [21].

Having defined both regulation (2.5) and excitation (3.11) problems, the following elementary result links their feasibility.

Proposition 2 (Connected feasibility)**.** *If $x \in \mathcal{X}_N$, then $\mathbb{P}_1(x)$ is feasible and, moreover, $\mathbb{P}_2\left(x; \mathbf{u}^*(x)\right)$ is feasible for any $\mathbf{u}^*(x) \in \mathcal{U}_N(x)$* (i.e., *not necessarily optimal for $\mathbb{P}_2$).*

Proof. This observation follows from the fact that the zero solution (*i.e.*, no perturbation) is always feasible for (3.11). □

This observation, whilst attractive from a feasibility standpoint, weakens the exponential result for the implemented RLS algorithm; the null perturbation solution may not necessarily satisfy the PE condition. Hence, convergence of the estimates to the true plant dynamics must be tested in simulation.

3.5.2 Successive Linearisation

By maximising the minimum eigenvalue of the lower bound of the information matrix increment, (i.e. the problem $\mathbb{P}_2^*\left(x; \mathbf{u}^*(x)\right)$), the real objective $\underline{\lambda}(\mathcal{R}_N - \mathcal{R}_0)$ will be only indirectly maximised. Moreover, the drawback of this approach is sub-optimality with respect to the regulation objective. However, because the original constraints are accounted for, the computed excitation sequence will still satisfy constraints for the nominal predicted trajectories.

The linearisation $\mathbb{P}_2^*\left(x; \mathbf{u}^0(x)\right)$ may, at most, lead to slower convergence of the RLS parameters. To combat this sub-optimality, as proposed in [47], re-linearisation of $\mathbb{P}_2^*\left(x; \mathbf{u}^0(x)\right)$ at an updated ψ^* can be implemented. Specifically, the updated linearisation point for the j-th iteration ($j = 0, \ldots, p$) is:

$$\psi^*_{\bar{i},j+1} = \psi^*_{\bar{i},j} + \delta\psi_{i,j} \tag{3.17}$$

where $\psi^*_{\bar{i},0} = \psi^*_i$. For $1 \leq j \leq p$, the linearisation iteration is

$$(\psi^*_{\bar{i},j})^\top = \left[(x^*_{\bar{i},j})^\top, (u^*_{\bar{i},j})^\top\right]. \tag{3.18}$$

The new predicted state trajectories are

$$x^*_{\bar{i},j} = \bar{A}^i x + \sum_{l=0}^{i-1} \bar{A}^l \bar{B} u^*_{\bar{i}-1-l,j} \tag{3.19}$$

with the perturbed inputs (and sub-optimal with respect to $\mathbb{P}_1(x)$):

$$u^*_{\bar{i},j} = u^0_i + \sum_{l=1}^{j} v_{i,l}. \tag{3.20}$$

The intuition of this recursion can be explained with the following example: suppose iteration $j = 0$ is the first linearised excitation problem $\mathbb{P}^*_{2,j=0}\left(x, \mathbf{u}^0(x)\right)$ that has already been solved to obtain perturbation sequence $\mathbf{v}_{j=0}$. By definition of the decomposition for $\mathbb{P}^*_2\left(x; \mathbf{u}^0(x)\right)$, the state and input trajectories $\mathbf{u}^0(x)$ and $\mathbf{x}^0(x)$ have already been used to define $\psi^*_{i,0}$. Hence, at the following iteration $j = 1$, to solve $\mathbb{P}^*_{2,1}\left(x, \mathbf{u}^0(x)\right)$, decomposition $\psi^*_{i,1}$ is composed of the adjusted state trajectories caused by the updated input sequence $\mathbf{u} = \mathbf{u}^0(x) + \mathbf{v}_{j=0}$. This process is repeated until satisfactory convergence criterion is met or until there is no time left for further optimisation during the sampling interval.

3.6 The LMPC Algorithm

Algorithm 3.1 provides a detailed description of the proposed learning model predictive controller for estimating the true model of the degraded system. The implementation and simulation results of this algorithm can be found in Section 3.9.

Algorithm 3.1 Learning MPC with RLS

Initialization

1: At $k = 0$, set $x = x(0)$, solve $\mathbb{P}_1(x)$ for $\mathbf{u}^0(x)$ and initialize $\mathcal{R}^-$, $\hat{\theta}^-$ for online steps

Online

2: Measure x and compute $\mathcal{R}$ from (3.9)
3: Solve $\mathbb{P}_1(x)$ for $\mathbf{u}^0(x)$
4: Solve $\mathbb{P}^*_{2,j=0}\left(x; \mathbf{u}^0(x)\right)$ for $\mathbf{v}^0_{j=0}\left(x; \mathbf{u}^0(x)\right)$
5: **while** remaining time in sampling instant **do**
6: Update linearisation $\psi^*_{i,j+1} = \psi^*_{i,j} + \delta\psi_{i,j}$
7: Solve $\mathbb{P}^*_{2,j+1}\left(x; \mathbf{u}^0(x)\right)$ for $\mathbf{v}^0_{j+1}\left(x; \mathbf{u}^0(x)\right)$
8: Update iteration index to $j + 1$
9: **end while**
10: Apply sum of first controls from each sequence to the system

$$u = \bar{\kappa}_N(x) + \pi_N(x) = u^0_0(x) + v^0_{0,j}\left(x; \mathbf{u}^0(x)\right)$$

11: Wait one time-step, then go to step 2

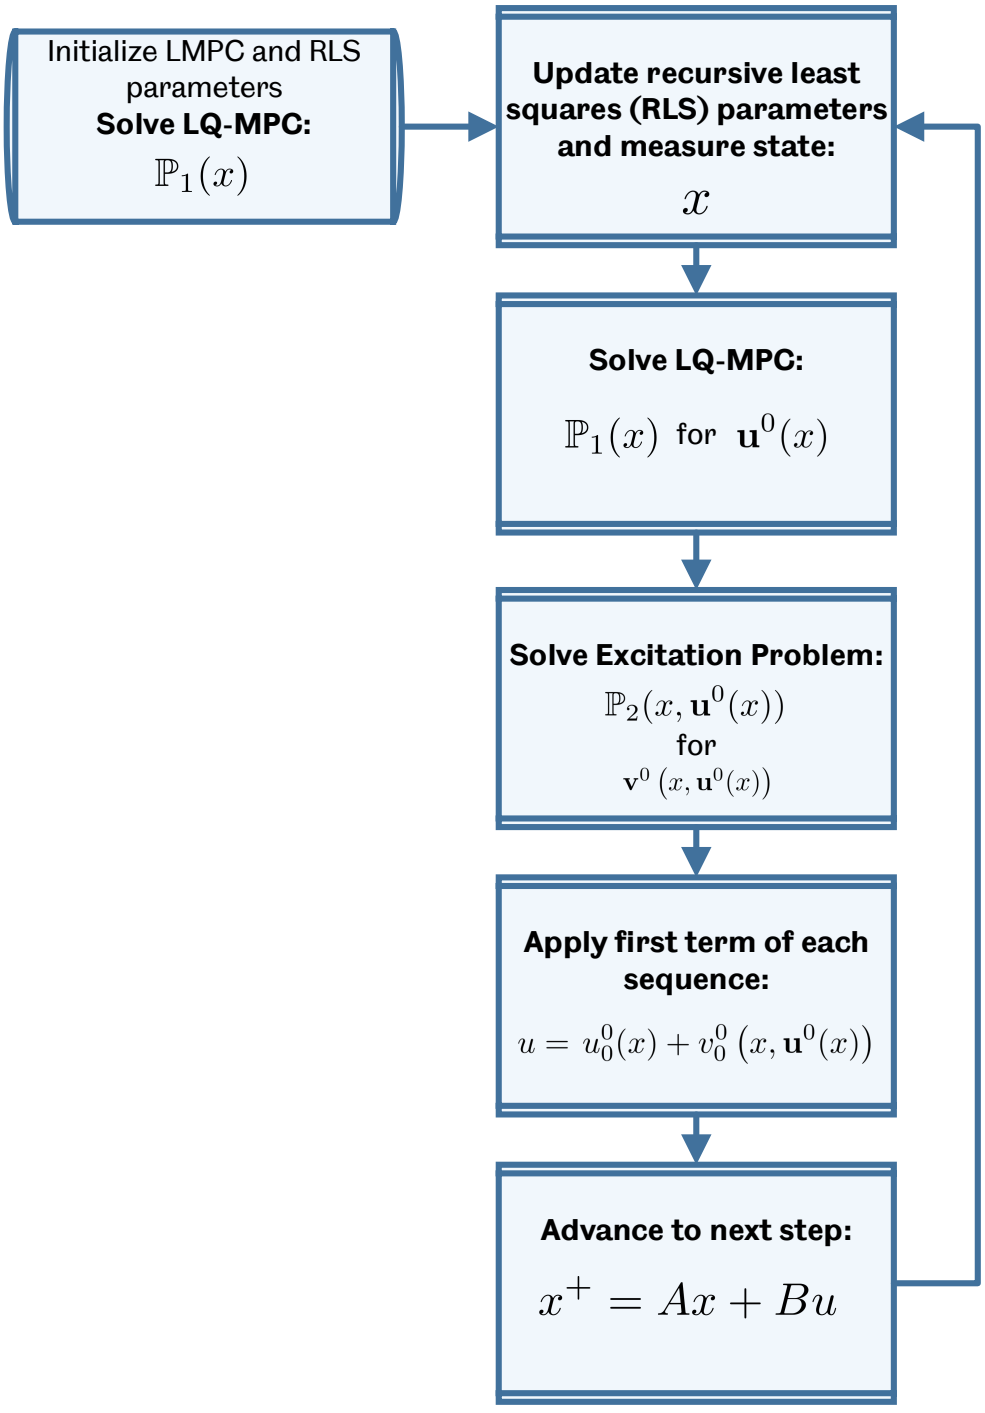

Figure 3.2: Flow chart describing the steps of a simplified LMPC algorithm (no iteration of the excitation optimisation).

3.7 Theoretical Stability Analysis

Having defined the control architecture and the optimisation problems involved, it is now essential to understand the impact of both model uncertainty and excitation on the closed-loop system under the LMPC controller:

$$x^+ = Ax + B(\bar{\kappa}_N(x) + \pi_N(x). \tag{3.21}$$

The main result establishes an upper bound on the model error (the mismatch between $(\bar{A}, \bar{B})$ and (A, B)) which, if met, guarantees stability of the closed-loop uncertain and perturbed system.

3.7.1 Main Result

To guarantee stability of a neighbourhood of the origin, which is the best that can be done given the excitation term, the well-known exponential stability results of nominal MPC with terminal conditions are employed [106]. Moreover, the inherent robustness of the regulating controller $\bar{\kappa}_N(x)$ is exploited to handle the state prediction error caused by modelling uncertainty and the exciting perturbation. This is a novel robust stability result that is explicitly in terms of the modelling uncertainty and available magnitude of excitation [53].

Proposition 3 (Exponential stability of nominal MPC)**.** *Suppose Assumptions 2,3,4,5 hold and the objective is quadratic and positive definite as in* (2.6)*, under a terminal control K that is LQ optimal. Then, the origin is exponentially stable for the nominal system $\bar{x}^+ = \bar{A}x + \bar{B}\bar{\kappa}_N(x)$, with region of attraction $\bar{\mathcal{X}}_N$. Moreover, there exist constants $c_2 > c_1 > 0$ such that the value function satisfies:*

$$c_1||x||^2 \leq V_N^0(x) \leq c_2||x||^2 \tag{3.22a}$$

$$V_N^0\left(\bar{A}x + \bar{B}\bar{\kappa}_N(x)\right) \leq \zeta V_N^0(x) \tag{3.22b}$$

for all $x \in \bar{\mathcal{X}}_N$, where $\zeta \triangleq (1 - c_1/c_2)$.

This result (textbook proof may be found in [106]) implies that in the absence of model uncertainty and the perturbing input for identification, the states of the regulated system decay exponentially fast to the origin. However, in reality, the model uncertainty and exciting perturbations are present. Moreover, the true successor state is *not* $\bar{A}x + \bar{B}\bar{\kappa}_N(x)$ but (3.21). To deal with this, the following property of the value function is utilised [106].

Lemma 2 (Lipschitz continuity of the value function)**.** *The value function $V_N^0(\cdot)$ satisfies $||V_N^0(x_1) - V_N^0(x_2)|| \leq L||x_1 - x_2||$ over $\mathcal{X}_N$, with Lipschitz constant L.*

Proof. This Lemma holds under the defined regulation formulation (2.5). The proof of Lemma 2 can be found in Theorem C.29 in Appendix C of [106]. □

Lemma 2 together with Proposition 3 are used to establish the main stability result:

Theorem 1 (Stability in terms of model mismatch)**.** *Suppose Assumptions 2-5 hold, and let* $x_0 \in \Omega_c \triangleq \{x : V_N^0(x) \leq c\} \subset \mathcal{X}_N$ *with positive constants* $b < c$*. If* $E \geq 0$ *for some* $\rho \in (\zeta, 1)$ *and* L *satisfying Lemma 2 such that,*

$$E \leq \frac{(\rho - \zeta)b - L||\delta A x + \delta B \bar{\kappa}_N(x)||}{L||\bar{B} + \delta B||} \tag{3.23}$$

for all $x \in \Omega_c$*, where* $\delta A \triangleq A - \bar{A}$ *and* $\delta B \triangleq B - \bar{B}$*, then (i) the set* Ω_b *is positively invariant for the system* (3.21)*, (ii) the set* Ω_c *is also positively invariant for* (3.21)*. Hence, starting from* $x_0 \in \Omega_c$*, the states of the system enter the set* Ω_b *in finite time and remain therein.*

Proof. (i) If $x \in \Omega_b$ then $V_N^0(x) \leq b$. Using (3.22) as an ISS Lyapunov function [106] for the true plant behaviour $x^+ = Ax + B(\bar{\kappa}_N(x) + \pi_N(x))$, the triangle inequality and induced matrix norm properties lead to

$$V_N^0(x^+) \leq \zeta b + L||\bar{B} + \delta B||E + L||\delta A x + \delta B \bar{\kappa}_N(x)||.$$

If there exists a $\rho \in (\zeta, 1]$ such that,

$$\zeta b + L||\bar{B} + \delta B||E + L||\delta A x + \delta B \bar{\kappa}_N(x)|| \leq \rho b \tag{3.24}$$

for all $x \in \Omega_b$, then the set Ω_b is PI for $x^+ = Ax + B(\bar{\kappa}_N(x) + \pi_N(x))$; the bound (3.23) is obtained by making E the subject of (3.24).

(ii) The claim for invariance of Ω_c and finite time convergence to Ω_b; consider some $x(0) \in \Omega_c \setminus \Omega_b$. Using the ISS property of the value function and the existence of $\rho \in (\zeta, 1]$ yields

$$\begin{aligned} V_N^0(x(1)) &\leq \zeta V_N^0(x(0)) + (\rho - \zeta) V_N^0(x(0)) \\ &\leq \rho V_N^0(x(0)). \end{aligned}$$

Thus, $V_N^0(x(k)) \leq \rho^k V_N^0(x(0))$, and so $V_N^0(x(k')) \leq b$ after some finite k', implying $x(k') \in \Omega_b$. □

Here, Ω_b is desired to be as small as possible to reduce regulation error, while Ω_c is the largest sublevel set contained in $\mathcal{X}_N$. The inequality (3.23) establishes

an upper bound on the value E which limits the perturbing input v_i in problem $\mathbb{P}_2$. Note that Ω_b can be chosen arbitrarily large so long as it is contained within the domain of attraction $\mathcal{X}_N$. The trade-off from a large choice of b is reduced regulation performance of the closed-loop due to larger permissible excitations.

3.7.2 Discussion

Though theoretically interesting, the bound (3.23) is of little practical use since several of the terms on the right-hand side are unknown or difficult to determine accurately:

1. The model errors δA, δB are unknown a-priori, but estimated bounds may be known.

2. The decay constant ζ can be estimated conservatively from theoretical considerations [106], or more accurately from numerical simulations as described later in Section 3.8.

3. The remaining constants b, c, ρ, and the Lipschitz constant L, have to be estimated from numerical simulations, in view of the value function being only implicitly defined (as the optimum of a parametric optimisation problem).

4. The use of the triangle inequality means that the developed upper bound on excitation is conservative; larger excitations may exist under which the closed-loop system remains stable and feasible.

On the other hand, some salient points that support intuition about the dual-control problem can be inferred:

- If model error is sufficiently small, the right-hand side of (3.23) is positive, and hence $E > 0$ is permitted in problem $\mathbb{P}_2$: that is, exciting perturbations to the regulating control signal are admissible while maintaining the closed-loop stability of the uncertain system.

- Conversely, for large model error the right-hand side of (3.23) may be negative, implying that no non-negative E exists. Hence, model uncertainty is too large to permit any exciting perturbations to the input, and if stability guarantees *are* to be achieved then a stronger regulatory action is required from the controller (*i.e.*, producing a small decay constant ζ).

Finally, it is noted that model error in both A and B means that it is difficult to establish a practical bound on model uncertainty that maintains the stability guarantee. If, however, error is exclusive to A then a bound that is more practical

(albeit still requiring numerical determination of the relevant constants) follows as a corollary of Theorem 1.

Corollary 1. *If $\bar{B} = B$, then the stability guarantee of Theorem 1 holds provided that:*

$$||\delta A|| \leq \frac{1}{\sqrt{b/c_1}}\left(\frac{(\rho-\zeta)b}{L} - ||\bar{B}||E\right) \tag{3.25}$$

Proof. With no uncertainty in the input matrix $\bar{B}$, then the inequality (3.23) simplifies to:

$$E \leq \frac{(\rho-\zeta)b - L||\delta A x||}{L||\bar{B}||}$$

Algebraic manipulation and using triangle inequality for a second time,

$$\frac{(\rho-\zeta)b}{L} - ||\bar{B}||E \geq ||\delta A||.||x||$$

According to Theorem 1, the state converges to $x \in \Omega_b$ and the lower bound of (3.22) implies the $c_1||x||^2 \leq b$. Re-arranging to make $||\delta A||$ the subject leads to Corollary 1. □

Note that, confirming intuition, large input-to-state gain implies the need for small E (and hence small exciting perturbations) in order to maintain stability of the regulator.

3.8 Computing Value Function Constants

This section describes how to compute the constants required to evaluate the bounds for Corollary 1.

3.8.1 $V_N^0(x)$ Lower Bound

Since the value function is, by definition, the optimal cost

$$V_N^0(x) = x_N^{0\top} P x_N^0 + \sum_{i=0}^{N-1} x_i^{0\top} Q x_i^0 + u_i^{0\top} R u_i^0 \tag{3.26}$$

with the superscripts denoting the optimal trajectories, it follows that $V_N^0(x) \geq x_0^{0\top} Q x_0^0 \geq \underline{\lambda}(Q)|x|^2$ with x being the measured state, hence $c_1 = \underline{\lambda}(Q)$.

3.8.2 $V_N^0(x)$ Upper Bound

From dynamic programming recursions of Section 2.4.4 in [106], the terminal cost upper bounds the value function such that

$$V_N^0(x) \leq x_N^\top P x_N, \quad \forall x \in \mathbb{X}_f \tag{3.27}$$

and hence, $V_N^0(x) \leq \bar{\lambda}(P)|x|^2 = \alpha(||x||)$ provides an upper bounding κ_∞ function within $\mathbb{X}_f$. However, this definition requires extension to the entire region of attraction $\mathcal{X}_N$. Utilising Proposition 2.18 of [106] guarantees the existence of $c > \eta > 0$ such that a new κ_∞ function $\beta(||x||) = \frac{c}{\eta}\alpha(||x||) = \frac{c}{\eta}\bar{\lambda}(P)||x||^2$ bounds the value function within $\mathcal{X}_N$. Therefore, $c_2 = \frac{c}{\eta}\bar{\lambda}(P)$. The remaining task is to find the least conservative pair c and η such that c_2 is the tightest upper bounding constant.

Computing η

Proposition 2.18 of [106] defines η as

$$\eta = \max_{x \in \mathbb{B}_r} \alpha(||x||) \tag{3.28}$$

where $\mathbb{B}_r = \{x \in \mathbb{R}^n : ||x|| \leq r\}$, which is equivalent to fitting the largest Euclidean norm ball within the terminal constraint $\mathbb{X}_f$.

Computing c

Proposition 2.18 of [106] defines c as the constant which implicitly defines the largest level set such that $V_N^0(x) \leq c$ for all $x \in \mathcal{X}_N$. By definition of the value function (3.27), the upper bound

$$V_N^0(x) \leq \underbrace{(N-1)\max_{x_c \in \mathbb{X}} x_c^\top Q x_c}_{l^c(.,u)} + \underbrace{(N-1)\max_{u_c \in \mathbb{U}} u_c^\top R u_c}_{l^c(x,.)} + \underbrace{\max_{x_c \in \mathbb{X}_f} x_c^\top P x_c}_{V_f^c(.)}, \quad \forall x \in \mathcal{X}_N \tag{3.29}$$

is obtained by selecting the worst-case stage cost at each step along the prediction horizon. With abuse of notation, the constant is defined $c = l^c(.,u) + l^c(x,.) + V_f^c(.)$. For a least conservative approximation of c_2, c has to be minimal whereas η maximal. By numerically testing near the boundaries of the controllability sets (4.2), rather than the boundary of $\mathbb{X}$, at each stage thus reducing the contribution from $l(.,u)$.

3.8.3 $V_N^0(x)$ Decay Rate

The decay rate ζ is a result of the stabilising ingredients utilised in Theorem 1. As the value function is used as a Lyapunov function for the closed-loop, the following conditions are satisfied:

$$c_1||x||^2 \leq V_N^0(x) \leq c_2||x||^2 \tag{3.30a}$$

$$V_N^0\left(\bar{A}x + \bar{B}\bar{\kappa}_N(x)\right) \leq V_N^0(x) - c_1||x||^2 \tag{3.30b}$$

for all $x \in \mathcal{X}_N$. The decay constant in $V_N^0\left(\bar{x}^+\right) \leq \zeta V_N^0(x)$ is found from a straightforward re-arrangement using the upper bound of the value function, yielding $\zeta = (1 - c_1/c_2)$.

3.8.4 $V_N^0(x)$ Lipschitz Constant

The Lipschitz constant is challenging to evaluate since the value function does not have an analytical expression and requires the solution of the MPC problem, hence, characterisation of the implicit control law $\bar{\kappa}_N(x)$. The best that can be hoped for is a tight upper bound [134]. In the context of (1), the smallest Lipschitz constant will lead to the most desirable bound on model mismatch. Methods for finding the best approximations of the Lipschitz constant are still an active field of research, even for quadratic programs and as such, the reader is referred to [30]. In the context of MPC, upper estimations have been shown in non-convex value function examples [117]. A conjecture for simplifying the computation of the Lipschitz constant has been proposed which uses tools from explicit MPC [106].

3.9 Illustrative Example

The learning MPC algorithm is demonstrated using the convex relaxation described in 3.5.1, on a system with the unstable dynamics:

$$x^+ = \underbrace{\begin{bmatrix} 1 & 1 \\ 0 & 1.2 \end{bmatrix}}_{A} x + \underbrace{\begin{bmatrix} 0 & 1 \\ 0.5 & 0.2 \end{bmatrix}}_{B} u$$

Contrary to the instability of the true dynamics, the prediction model used by

the controller is stable:

$$x^+ = \underbrace{\begin{bmatrix} 0.7 & 0.7 \\ 0 & 0.84 \end{bmatrix}}_{\bar{A}} x + \underbrace{\begin{bmatrix} 0 & 1 \\ 0.5 & 0.2 \end{bmatrix}}_{\bar{B}} u.$$

Table 3.1 shows the parameters selected for the controller and RLS algorithm, chosen to comply with Assumptions 2-5. The following constraints were considered:

$$\mathbb{X} = \{x \in \mathbb{R}^2 : |[x]_h| \leq 4, h = 1, 2\},$$
$$\mathbb{U} = \{u \in \mathbb{R}^2 : |[u]_h| \leq 1, h = 1, 2\}.$$

The set $\mathbb{X}_f$ is designed to satisfy Assumption 4 with the computed P and the unconstrained LQ-optimal terminal control law K.

Figure 3.3 shows the state and input trajectories of the system under the composite regulating and exciting control law (3.3). The controller achieves regulation of the system to a neighbourhood of the origin, despite the input perturbations and the model error.

Figure 3.4 shows, over the same time period, the learning of the true system dynamics by the RLS module. In this case, the perturbations applied to the regulating control law are sufficient to achieve accurate convergence of the model parameters.

Table 3.1: Controller parameters

Parameter	Value	Parameter	Value
State x_0	$[3, 0]^\top$	γ	0.7
R	$I_{2\times 2}$	E	0.7
Q	$5I_{2\times 2}$	N	10
P	$\left[\begin{smallmatrix} 5.2485 & 0.1314 \\ 0.1314 & 5.7667 \end{smallmatrix}\right]$	$\mathcal{R}^-$	$I_{4\times 4}$

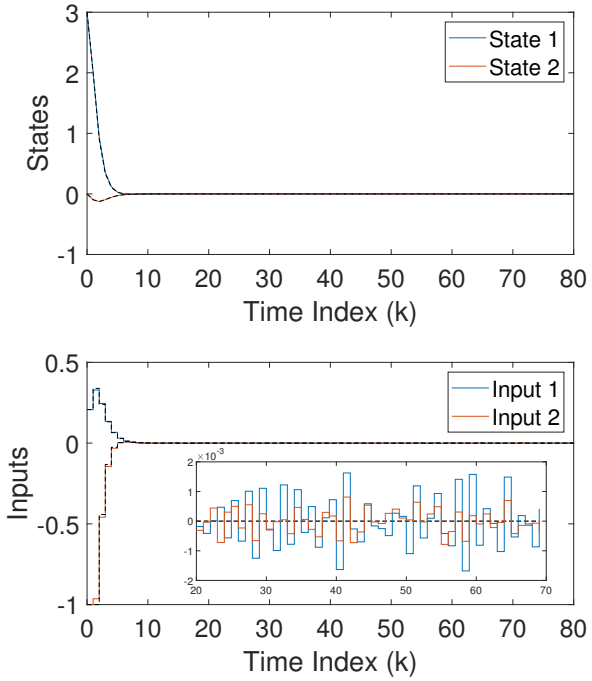

Figure 3.3: State and input trajectories under the composite control law, including a magnified view. The dashed lines denote conventional MPC trajectories (*i.e.*, without the exciting perturbations).

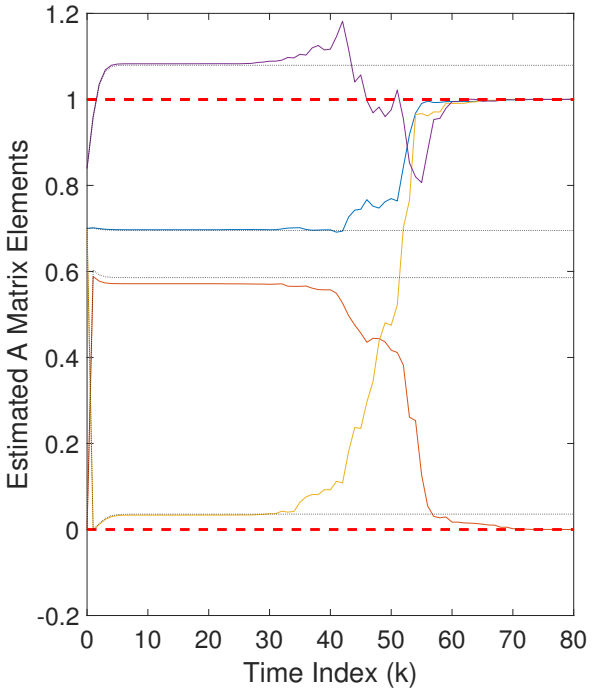

Figure 3.4: Estimated A-matrix parameters. Dashed red lines denote the true parameters. Dotted black lines denote the parameters estimated under only conventional LQ-MPC, which do not converge to the true parameters.

3.10 Summary and Conclusion

An approach to the dual-control problem has been presented in a linear MPC framework, where persistency of excitation is promoted through an additional optimisation problem solved online. Closed-loop stability of the true uncertain system, despite the model error and exciting perturbations applied to the regulating control law, was established by using the well-known exponential stability result of LQ-MPC. The main result gives an upper bound on the permitted magnitude of input perturbations determined by the secondary optimisation. When specialized to the case of uncertainty only in the A-matrix, the result provides a more practical, albeit still difficult to determine upper bound on model error.

Since an excitation sequence over a future horizon is made available from the optimiser $\mathbb{P}_2\left(x; \mathbf{u}^0(x)\right)$, the next logical step is to explore how the original regulating controller $\bar{\kappa}_N(x)$ can utilise the knowledge of a receded perturbation sequence to further robustify the regulating controller. The next chapter precisely explores this idea and shows several novel results that improve the inherent robustness of the control method without adding any additional computational complexity to the MPC scheme approach.

Chapter 4

Preview Information and Inherent Robustness

4.1 Introduction

Most industrial MPC controllers work despite their formulations not using the standard ingredients that provide stability and recursive feasibility guarantees [87]. These controllers utilise the property of inherent robustness that is present in feedback control [6]. The linear quadratic formulation of MPC presented in Chapter 2 does indeed provide a *quantifiable* degree of inherent robustness, related to the chosen tuning parameters of the controller. Thus far, the inherent robustness property has been exercised to deal with model mismatch *and* excitation for convergent identification.

The novelty of this chapter introduces feed-forward into the algorithm proposed in Chapter 3 by informing the controller of the predicted excitation sequence. With appropriate modifications of the nominal ingredients in LQ-MPC, the controller's robustness margins become independent of the magnitude of the excitations; instead, the change from the expected excitation sequence affects robustness in an analogous bound to (3.23). Moreover, the proposed approach enlarges the region of attraction when compared to a nominal LQ-MPC with no preview of the excitations.

This simple extension of learning MPC improves its robustness without any additional on-line computational complexity; simple arithmetic operations and storage of the perturbation sequence are the only additional computational requirements. Similar theoretical analysis to Chapter 3 and numerical arguments support the claim under appropriate bounds on modelling error.

4.2 Preview Information

In robust MPC, the common design practice is to assume that the disturbance exists within some worst-case bounded set of potential disturbances. Such a design prepares the MPC controller to handle any realisation of the disturbance at the cost of increased conservatism. Given that MPC performs predictions over a future horizon, it is sensible to investigate how the knowledge of future disturbances–if available and of sufficient accuracy–can be utilised to robustify LQ-MPC; rather than resorting to the classical robust approach. In the context of gas turbines, such a situation could exist if prior information of the flight schedule and future inlet conditions is available; either through accurate weather predictions or sensors that can detect the characteristics of the oncoming flow.

It is constructive to begin further discussion of this idea with a definition of preview information in the linear prediction model framework.

Definition 7 (Preview Disturbance Sequence)**.** *The finite preview disturbance sequence* $\mathbf{w} = \{w_0, w_1, \ldots, w_{N-1}\}$ *defines the knowledge of state disturbances over the prediction horizon* N*, with current disturbance* $w_k = w_0$ i.e.*,the first term of the sequence.*

Such an available sequence can be easily introduced into the prediction model (2.4) yielding

$$\bar{x}^+ = \bar{A}x + \bar{B}u + w, \tag{4.1}$$

and becomes implemented as an analogous constraint (2.7b) in a modified optimisation problem. This system description is equivalent to (2.3) with the assumption that the disturbance is entirely measurable. However, as will be illuminated in the coming sections, naively updating the prediction model with future disturbances is not sufficient for retaining the stability guarantees of LQ-MPC [13].

4.2.1 Issues with including Preview Information

Several theoretic considerations must be addressed when including predictions of future disturbances in the prediction horizon. Firstly, it is useful to recall and discuss how some of the assertions from nominal LQ-MPC are affected:

Assumption 2: Linearity and reachability of the dynamics.

Assumption 3: Fixed plant dimensionality and constraints.

Assumption 4: Lack of external disturbances affecting the plant.

Assumption 5: Positive invariance of the terminal constraint set $\mathbb{X}_f$.

Proposition 3: Exponential stability using the value function $V_N^0(x)$.

Assumption 2 holds for the system without the preview information. However, under deterministic knowledge of the preview information, the plant dynamics become affine. These modified dynamics are the cause of the following issues that deserve the most attention.

Loss of Invariance of $\mathbb{X}_f$

Assumption 5 is clearly broken with the addition of preview information, as the set $\mathbb{X}_f$ may no longer be PI for the system $\bar{x}^+ = \bar{A}x + \bar{B}u + w$ under the same terminal control law $u = \bar{K}x$ (even under a constant w). Moreover, simply shifting the original set $\mathbb{X}_f$ by w will result in violation of the original constraints defined under Assumption 3.

Ill-posed Objective $V_N(x, \mathbf{u})$

Since the origin is no longer an equilibrium point of the system $\bar{x}^+ = \bar{A}x + \bar{B}u + w$, the best that can be done is to steer the states into a *neighbourhood* of the origin. The value function $V_N^0(x)$ is unable to satisfy the conditions of a Lyapunov function under the new prediction model because the components $V_f(x)$ and $l(x, u)$ will not be zero for affine dynamics (4.1).

Loss of nesting of the Controllability Sets

The controllability sets, derived in Appendix 6.6, are useful for analysing feasibility of constrained control. It is well known that the system (2.4) under LQ-MPC, satisfying Assumptions 2-5, has nested controllability sets if $\mathbb{X}_f$ is PI for the system $\bar{x}^+ = \bar{A}x + \bar{B}u$ [93]. This property is attractive as it allows for straightforward arguments that can prove recursive feasibility. However, with a disturbed prediction model, the controllability sets are defined as:

$$\mathcal{X}_{i+1} = \mathbb{X} \cap A^{-1}(\mathcal{X}_i \oplus -B\mathbb{U} \oplus -\{w(N-i)\}) \quad , \quad \text{for} \quad i = 0, \ldots, N-1 \tag{4.2}$$

where $\mathcal{X}_0 = \mathbb{X}_f$. Notice that the final term in the set recursion (4.2) induces translation of each controllability set along the prediction horizon. Regardless of whether $\mathbb{X}_f$ is made PI for (4.1), with large enough $w(N-i)$, loss of guaranteed nesting makes it difficult to prove recursive feasibility of the MPC controller (but not impossible [13]).

4.2.2 Preview Information Definitions and Assumptions

This chapter does not suggest the use of additional sensing capability or perfect weather predictions. Rather, the goal is to use a specialised form of the disturbance sequence that is already available: the perturbation sequence $\mathbf{v}$ as a result of solving an ancillary optimisation for excitation.

Definition 8 (Preview Perturbation Sequence)**.** *The preview perturbation sequence* $\mathbf{v} = \{v_0, v_1, \dots, v_{N-1}\}$ *defines the knowledge of input perturbations, over the prediction horizon N, at some generic time-step k (which may be omitted for clarity).*

To begin analysis, some additional assumptions on the nature of this preview information are required. The following mild assumptions on the preview perturbation sequence are made:

Assumption 7 (Current knowledge)**.** *The current perturbation* v_k *is known exactly at time-step k (*$i = 0$ *in the prediction horizon index); estimates of future perturbation* v_i *for* $i \in \mathbb{Z}_{i\geq 1}$ *are available and satisfy* $Bv_i \in \mathbb{W}$*.*

Clearly the preview information definitions 7 and 8 are equivalent when each term of the latter definition is multiplied through the uncertain B matrix *i.e.*, $w = Bv$.

Since most of the issues of including preview information relate to the terminal conditions of the regulation problem, it is important to define what happens after the end of the sequence of available perturbation predictions. To simplify analysis of these terminal conditions, the following definitions define the behaviour of the perturbation sequence as part of the receding horizon framework.

Definition 9 (Receding perturbations)**.** *The tail-end of the perturbation sequence is defined as* $\tilde{\mathbf{v}}(\mathbf{v}) = \{v_1, \dots, v_{N-1}, v_{N-1}\}$ *where the first term of* $\mathbf{v}$ *is removed whilst the last term is repeated and appended, maintaining the original dimensions of the sequence such that* $B \odot \tilde{\mathbf{v}}(\mathbf{v}) \in \mathbb{W} \times \cdots \times \mathbb{W} \times \mathbb{W}_f \times \mathbb{W}_f$ *and* $\mathbb{W}_f \subseteq \mathbb{W}$*.*

Definition 10 (Extension beyond N)**.** *The infinite-length perturbation sequence* $\mathbf{v}_N(k)$ *is constructed by concatenating* $\mathbf{v}(k)$ *with* $\mathbf{v}_f(k) = \{v_i\}_{i\geq N}$*, where* $Bv_i \in \mathbb{W}_f \subseteq \mathbb{W}$ *for* $v_i \triangleq v_f(\mathbf{v})$ *when* $i \geq N$*. Hence, the infinite sequence is* $\mathbf{v}_N(k) = \{v_0, \dots, v_{N-1}, v_{N-1}, \dots\}$.

In other words, $\mathbf{v}_f(k)$ contains the final term of the finite sequence, v_{N-1}, and is repeated infinitely. Finally, the infinite sequence $\mathbf{v}_i(k)$ for $i = 0, \dots, N$ is a version of $\mathbf{v}_N(k)$ but without the first $N - i$ terms. The function $\tilde{\mathbf{v}}$ applied to an infinite sequence, $\tilde{\mathbf{v}}(\mathbf{v}_i)$, removes the first term of $\mathbf{v}_i$ yielding $\tilde{\mathbf{v}}(\mathbf{v}_i) = \mathbf{v}_{i-1}$.

With these definitions and assumptions, the results of [13] can be applied and more importantly *strengthened*; by utilising the properties that follow from considering input perturbations rather than state disturbances. The following section describes the necessary modifications of the optimal control problem that retain the nominal stability guarantees, whilst accounting for preview perturbation information.

4.3 LQ-MPC with Perturbation Preview

Knowledge of the perturbation sequence requires a change in notation of the original MPC problem. The perturbation sequence $\mathbf{v}$ becomes an additional parameter such that the new LQ-MPC optimisation problem is defined as:

$$\mathbb{P}_1(x;\mathbf{v})\colon V_N^0(x;\mathbf{v}) = \min_{\mathbf{u}}\{V_N(x,\mathbf{u};\mathbf{v}) : \mathbf{u} \in \mathcal{U}_N(x;\mathbf{v})\} \tag{4.3}$$

where the feasible region $\mathcal{U}_N(x;\mathbf{v})$ is defined for $i = 0,\dots,N-1$ using:

$$x_0 = x \tag{4.4a}$$
$$\bar{x}_{i+1} = \bar{A}x_i + \bar{B}(u_i + v_i) \tag{4.4b}$$
$$\bar{x}_i \in \mathbb{X} \tag{4.4c}$$
$$u_i \in \mathbb{U}_i(\mathbf{v}) \tag{4.4d}$$
$$\bar{x}_N \in \mathbb{X}_f(\mathbf{v}) \tag{4.4e}$$

with the cost function

$$V_N(x,\mathbf{u};\mathbf{v}) = V_f(x_N;\mathbf{v}) + \sum_{i=0}^{N-1} l(x_i,u_i;\mathbf{v})$$

and constraints that are modified appropriately to retain the nominal guarantees of LQ-MPC. Derivations of the critical components $V_N(x,\mathbf{u};\mathbf{v})$, $\mathbb{U}(\mathbf{v})$ and $\mathbb{X}_f(\mathbf{v})$ are developed in the following sections.

4.3.1 Modifications of the Stabilising Ingredients

The required modifications of $V_f(x;\mathbf{v})$, $l(x,u;\mathbf{v})$ and $\mathbb{X}_f(\mathbf{v})$, first proposed in [13], address the loss of invariance and ill-posedness of the regulation problem. The key assumptions from section 3.B of [13] are specialised to the preview perturbation definitions.

Assumption 8 (Nominal Terminal Invariant Set)**.** *The control invariant set $\bar{\mathbb{X}}_f$ is known and functions $\bar{l}(x,u)$, $\bar{V}_f(x)$ satisfy Assumptions 4 and 5 for the reachable system $\bar{x}^+ = \bar{A}x + \bar{B}u$ but with scaled constraint sets $\beta_x\mathbb{X}$, $\beta_u\mathbb{U}$ where β_x,$\beta_u \in [0,1)$.*

Assumption 8 implies the existence of a stabilising control law $\bar{\kappa}_f(x)$ in a set of stabilising control laws such that:

$$\begin{aligned} u_f(x) = &\{u \in \beta_u\mathbb{U} : Ax + Bu \in \beta_x\mathbb{X}, \\ &\bar{V}_f(\bar{A}x + \bar{B}u) + \bar{l}(x,u) \leq \bar{V}_f(x)\} \qquad \forall x \in \bar{\mathbb{X}}_f. \end{aligned}$$

Assumption 9. *The control law $\bar{\kappa}_f(x)$ is continuous over its domain of attraction $\bar{\mathbb{X}}_f$, $\bar{\kappa}_f(0) = 0$ and is continuously differentiable around the origin.*

Invoking Assumptions 8 and 9 lead to the nominal terminal dynamics $\bar{x}^+ = \bar{A}x + \bar{B}\bar{\kappa}_f(x) = f_{\bar{\kappa}_f}(x)$, with a domain of attraction $\bar{\mathbb{X}}_f$. However, the perturbed terminal system will in fact be

$$\bar{x}^+ = f_{\bar{\kappa}_f}(x) + w_f$$

with the terminal disturbance stated in terms of the terminal input perturbation $w_f = \bar{B}v_f \in \mathbb{W}_f$. Moreover, Assumption 9 permits the linearisation of $f_{\bar{\kappa}_f}(x)$ around the origin such that

$$\Pi \triangleq \frac{\delta\bar{\kappa}_f(x)}{\delta x} \quad , \quad \Phi \triangleq \frac{\delta f_{\bar{\kappa}_f}(x)}{\delta x} = A + B\Pi. \tag{4.5}$$

Note that any linear control law K can be used in place of Π, including a LQ regulator from the dual-mode approach described in Chapter 2. Given a constant terminal disturbance, and the technical assumptions thus far, it is now possible to obtain a new equilibrium point for the perturbed terminal dynamics.

Perturbed System Equilibrium

Steady-state analysis of the perturbed dynamics under the stabilising control law $\Pi = K$ yields the new equilibrium point:

$$x_f(v_f) = \Psi\bar{B}v_f \quad , \quad u_f(v_f) = K\Psi\bar{B}v_f \tag{4.6}$$

where $\Psi = (I - \Phi)^{-1}$. With Ψ being well defined and Schur, the following Lemma is proved:

Lemma 3. *Suppose Assumptions 8 and 9 hold. Then, the equilibrium point $\big(x_f(v_f), u_f(v_f)\big)$ of the terminal dynamics $x^+ = \Phi x + \bar{B}v_f$, computed using* (4.6)*, exists and is unique under a stabilising feedback control law $\Pi = K$ and constant perturbation v_f.*

Modified LQ-MPC Ingredients

The explicit equilibrium point defined in Lemma 3 can now be used to translate the original stabilising ingredients:

$$l(x,u;\mathbf{v}) \triangleq x_e^\top(v_f)Qx_e(v_f) + u_e^\top(v_f)Ru_e(v_f) \tag{4.7}$$
$$V_f(x;\mathbf{v}) \triangleq x_e^\top(v_f)Px_e(v_f) \tag{4.8}$$
$$\mathbb{X}_f(\mathbf{v}) \triangleq \bar{\mathbb{X}}_f \oplus \{x_f(v_f)\} \tag{4.9}$$
$$\bar{\kappa}_f(x;\mathbf{v}) \triangleq Kx_e(v_f) + K\Psi\bar{B}v_f = Kx \tag{4.10}$$

where the state and input error terms are $x_e(v_f) = x - x_f(v_f)$ and $u_e(v_f) = u - u_f(v_f)$ respectively.

Remark 2. *Notice that the modified terminal control law* (4.10) *can be simplified to the linear control law $u_f = Kx$ of the unmodified terminal control problem, when the definition of the equilibrium point $x_e(v_f)$ is substituted in. The significance of this is that the terminal cost matrix P remains unchanged, since the solution of the discrete time (DT) Lyapunov equation that encapsulates the cost-to-go for an infinite prediction horizon, can be computed with the unmodified terminal control law.*

Assumption 10 (Scaling of the terminal perturbation)**.** *There exist scalars $\alpha_x, \alpha_u \in [0,1)$ such that*

$$\Psi\mathbb{W}_f \subseteq \alpha_x\mathbb{X} \quad , \quad K\Psi\mathbb{W}_f \subseteq \alpha_u\mathbb{U}.$$

Scaling the terminal perturbation is required to satisfy Assumption 8; the constants α_x, α_u limit the size of the admissible terminal perturbation given the scaled down constraint sets $\beta_x\mathbb{X}, \beta_u\mathbb{U}$.

The following proposition establishes the stability result for the terminal dynamics, under the proposed set scaling and translations of the nominal ingredients.

Proposition 4 (Invariance of $\mathbb{X}_f(\mathbf{v})$)**.** *Suppose Assumptions 2,3,4 and 7,8,9,10 hold. Then, for any constant $\bar{B}v_f \in \mathbb{W}_f$, (i) the set*

$$\mathbb{X}_f(\mathbf{v}) = \bar{\mathbb{X}}_f \oplus \{x_f(v_f)\} \tag{4.11}$$

is PI for $\bar{x}^+ = \bar{A}x + \bar{B}\bar{\kappa}_f(x;\mathbf{v}) + \bar{B}v_f$, and (ii) the functions $l(x,u;\mathbf{v})$ and $V_f(x;\mathbf{v})$ satisfy

$$V_f\big(\underbrace{\bar{A}x + \bar{B}(\bar{\kappa}_f(x;\mathbf{v}) + v_f)}_{\bar{x}^+};\mathbf{v}\big) + l(x,\bar{\kappa}_f(x;\mathbf{v});\mathbf{v}) \leq V_f(x;\mathbf{v}) \quad \forall x \in \mathbb{X}_f(\mathbf{v}).$$

Furthermore, (iii) $\mathbb{X}_f(\mathbf{v}) \subseteq \mathbb{X}$ *and* $\bar{\kappa}_f\big(\mathbb{X}_f(\mathbf{v}); \mathbf{v}\big) \subseteq \mathbb{U}$ *if* $\alpha_x + \beta_x \leq 1$ *and* $\alpha_u + \beta_u \leq 1$. *Finally, (iv) the point* $\big(x_f(v_f), u_f(v_f)\big)$ *is an asymptotically stable equilibrium point for the system* $\bar{x}^+ = \bar{A}x + \bar{B}(\bar{\kappa}_f(x; \mathbf{v}) + v_f)$ *with a region of attraction* $\mathbb{X}_f(\mathbf{v})$.

Proof. The proof of Proposition 4 follows exactly the proof of Proposition 1 in [13], but only replacing w_f with $w_f = \bar{B}v_f$. □

Remark 3. *An elegant consequence of Proposition 4 is that, in the case of constant* v_f *and no modelling uncertainty, stability of a point is achieved rather than a neighbourhood of the origin for the terminal prediction dynamics. This observation is used to infer the same behaviour for the closed-loop system in Proposition 6.*

The online solving of the modified problem $\mathbb{P}_1(x; \mathbf{v})$ at each time-step, yields the analogous stabilising control law $\bar{\kappa}_N(x; \mathbf{v}) = u_0^0(x; \mathbf{v})$ that returns the first element of the optimal sequence, as in nominal LQ-MPC; but now accounting for preview perturbations.

4.4 Nesting of the Controllability Sets

This section addresses the final issue caused by using a perturbed prediction model: loss of nesting of the controllability sets. Unlike the results of [13], it is shown that nesting of the controllability sets can be restored under an additional translation of the input constraint set, yielding $\mathbb{U}_i(\mathbf{v})$. The trade-off, from performing the translation, is a slight reduction in the enlargement of the robust region of attraction and is demonstrated in the numerical study of Section 4.8. However, the extremely attractive property of recursive feasibility is recovered under assumptions on the evolution of the preview information sequence and zero modelling uncertainty [93].

4.4.1 Unchanging Perturbation

Following the MPC Assumptions 2, 3 and 5 with a generic state disturbance sequence, the controllability sets are recalled to be

$$\mathcal{X}_{i+1}(\mathbf{w}_{i+1}) = \mathbb{X} \cap A^{-1}(\mathcal{X}_i(\mathbf{w}_i) \oplus -B\mathbb{U} \oplus -\{w(N-i)\}) \tag{4.12}$$

with $\mathcal{X}_0 = \mathbb{X}_f$. The finite disturbance sequence $\mathbf{w} = \{w(0), \ldots, w(N-1)\}$ follows all of the same definitions as for $\mathbf{v}$ but in the space $\mathbf{w} \in \mathbb{R}^{Nn}$. Note that the subscripted sequence $\mathbf{w}_i$ is an infinite sequence, following the definitions of 4.2.2. As discussed, the computation of (4.12) in the generic disturbance case will lead

to translation of the controllability sets. Moreover, if A is unstable, the translation will be further exacerbated and can lead to a $\mathcal{X}_N$ that is not control invariant. Instead, by using input perturbations as the disturbance sequence and a suitable translation of the input constraint, nesting of the controllability sets over the prediction horizon is preserved.

Theorem 2 (Nesting of Controllability Sets)**.** *Suppose that the assumptions of Proposition 4 hold for control of $\bar{x}^+ = \bar{A}x + \bar{B}(u+v)$ with available preview sequence $\mathbf{v}$. If the input constraint set $\mathbb{U}$ is translated according to $\mathbb{U}_i(\mathbf{v}) = \mathbb{U} \ominus \{v(i)\}$ for $i = 0, \dots, N-1$, the controllability sets are nested* i.e., *$\mathcal{X}_N(\mathbf{v}_N) \supseteq \mathcal{X}_{N-1}(\mathbf{v}_{N-1}) \supseteq \cdots \supseteq \mathcal{X}_0(\mathbf{v}_0)$.*

Proof. The input constraint is translated according to $\mathbb{U} \ominus \{v(i)\}$ at each step along the prediction horizon. For example, at $i = 0$, $\mathbb{U} \ominus \{v(0)\}$ whereas at the consecutive prediction steps, $\mathbb{U} \ominus \{v(1)\}$ etc. Specialising equation (4.12) such that $w(i) = Bv(i)$ yields

$$\mathcal{X}_{i+1}(\mathbf{v}_{i+1}) = \mathbb{X} \cap A^{-1}(\mathcal{X}_i(\mathbf{v}_i) \oplus -B(\mathbb{U} - v(N-i)) \oplus -\{Bv(N-i)\}). \tag{4.13}$$

The Minkowski sum of a set with the negation of a singleton set is simply subtraction (see proof in Appendix 6.5). Straightforward algebra results in the controllability set recursion:

$$\mathcal{X}_{i+1}(\mathbf{v}_{i+1}) = \mathbb{X} \cap A^{-1}(\mathcal{X}_i(\mathbf{v}_i) \oplus -B\mathbb{U}). \tag{4.14}$$

Finally, nesting follows from two additional facts: (i) $\mathbb{X}$, $\mathbb{U}$ being polyhedral and containing the origin (ii) $\mathcal{X}_0 = \mathbb{X}_f(\mathbf{v}) \subset \mathbb{X}$ which by construction, satisfies Proposition 4 meaning $\mathbb{X}_f(\mathbf{v})$ is control invariant. □

The preceding proof specialises the arguments of [93].

Remark 4. *Theorem 2 proves nesting of the controllability sets over the prediction horizon but does not imply anything about the nesting between time-steps* i.e., *$x \in \mathcal{X}_N(\mathbf{v}_N) \implies x^+ \in \mathcal{X}_{N-1}(\mathbf{v}^+)$ if the perturbation sequence changes arbitrarily to $\mathbf{v}^+$. However, if the input perturbation sequence $\mathbf{v}$ recedes along the MPC horizon according to $\mathbf{v}^+ = \tilde{\mathbf{v}}(\mathbf{v})$, nesting of the controllability sets between time-steps is guaranteed and nominal recursive feasibility is recovered, since $x \in \mathcal{X}_N(\mathbf{v}_N) \implies x^+ \in \mathcal{X}_{N-1}(\mathbf{v}_{N-1})$.*

In other words, Corollary 1 of [93] is satisfied and moreover, forward propagation of $\mathbb{X}_f(\mathbf{v})$ via the recursion (4.14) implies each $\mathcal{X}_i(\mathbf{v}_i)$ is control invariant if the perturbation sequence $\mathbf{v}$ develops according to $\mathbf{v}^+ = \tilde{\mathbf{v}}(\mathbf{v})$. It is important to note that simply receding the known sequence using $\tilde{\mathbf{v}}(\mathbf{v})$ may not provide

sufficient excitation at the next time step, since the regulator acts against the excitation. Additionally, the sequence $\mathbf{v}_N(k)$ will converge to a constant after N time-steps. However, the modifications derived thus far provide the necessary baseline to study the more useful and challenging case of changing the perturbation sequence between time-steps.

4.5 Modification of LMPC

Hitherto, the discussion of the modifications to account for preview information has been limited to the newly defined regulation problem $\mathbb{P}_1(x;\mathbf{v})$. This section specifies how the ancillary excitation problem $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$ has to be modified to satisfy the relevant assumptions when utilising a preview perturbation sequence.

4.5.1 Issues with including Preview Information

Since the goal of the excitation problem is to promote informative data for convergent RLS parameter identification, the feasibility of $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$ is the main concern of this section. Optimality is of lesser importance since $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$ is already an approximation of the true excitation problem (3.11). Fortunately, the only assumption that affects the excitation optimisation is:

> **Assumption 7**: Exact knowledge of the current time-step perturbation v_k such that $w_k = Bv_k$ is known.

Clearly, this assumption causes a catch-22 situation with the new definition of $\mathbb{P}_1(x;\mathbf{v})$ and original $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$. To solve $\mathbb{P}_2^*\left(x(k),\mathbf{u}^0(x(k))\right)$ for $\mathbf{v}(k)$ the optimal regulation sequence $\mathbf{u}^0(x(k);\mathbf{v}(k))$ is required, which itself is computed using $\mathbb{P}_1(x(k);\mathbf{v}(k))$ and an already known $\mathbf{v}(k)$! A simple remedy to break the algebraic loop is to constrain the current time-step perturbation to the previous time-step prediction, thereby also satisfying Assumption 7.

Assumption 11 (Causality constraint)**.** *Given the perturbation sequence notation* $\mathbf{v}(k) = \{v_0(k), v_1(k), \ldots, v_{N-1}(k)\}$*, consider the previous time-step sequence* $\mathbf{v}(k-1)$*. The assumption requires that the perturbation realization at current time k satisfies* $v_k = v_1(k-1)$*.*

Remark 5. *At first glance, Assumption 11 can lead to a situation where the perturbations settle to a constant, depending on the initialisation of* $\mathbf{v}(k-1)$*. However, since the optimisation* $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$ *is free to optimise over the remaining degrees of freedom in* $\mathbf{v}$*, receding the perturbation sequence according to* $\tilde{\mathbf{v}}(\mathbf{v})$ *will lead to new perturbations* v_k*.*

Assumption 11 is implemented in a modified version of the excitation problem $\mathbb{P}_2$, which limits the degrees of freedom of the decision variable. Moreover, Assumption 11 implies that some initialisation sequence $\mathbf{v}(k-1)$ (also denoted $\mathbf{v}^-$) is required to solve the first optimisation problem $\mathbb{P}_1(x(k);\mathbf{v}(k))$. This initialisation sequence can be readily designed off-line. A classic example of an excitation sequence that is known to have good excitation properties is a pseudo-random binary signal (PRBS) [79]; providing that such an initial excitation sequence is feasible and does not divert trajectories away from reaching the terminal constraint.

With Assumption 7 and the discussion of Remark 5, the definition and role of $\mathbb{P}_2^*\left(x;\mathbf{u}^0(x)\right)$ is subtly changed. Rather than simply optimising the magnitude of perturbation along the horizon at each time-step, deviations from the tail-end of the previous time-step perturbation sequence are used for promoting PE conditions. This interpretation motivates the theoretical analysis on the rate of change of the excitation sequence and its affect on stability; the result of the analysis is presented in Section 4.6. However, before any further analysis, a statement of the modified excitation problem is given.

Summary of modified $\mathbb{P}_2$ Excitation Problem

With the prior discussion on how to deal with introducing preview information into the learning MPC algorithm, this section defines the modified excitation optimisation.

Immediately following the solving of $\mathbb{P}_1(x;\mathbf{v})$ at state x, which generates the sequence $\mathbf{u}^0(x)$, the secondary optimisation is

$$\mathbb{P}_2^*(x;\mathbf{v}^-,\mathbf{u}^0) : H_N^{0,*}(x;\mathbf{v}^-,\mathbf{u}^0) = \max_{\mathbf{v}}\{H_N^*(x,\mathbf{v};\mathbf{v}^-,\mathbf{u}) : \mathbf{u}^0+\mathbf{v}\in\tilde{\mathcal{U}}_N(x;\mathbf{v}^-)\}$$

with a perturbation sequence $\mathbf{v}=\{v_0,v_1,\ldots,v_{N-1}\}$ as the decision variable. The

cost function is designed using the same linearisation technique of Proposition 1:

$$H_N^*(x, \mathbf{u}) = \underline{\lambda}(\mathcal{R}_N - \mathcal{R}_0) \tag{4.15a}$$
$$\mathcal{R}_0 = \mathcal{R} \tag{4.15b}$$
$$\mathcal{R}_{i+1} = \mathcal{R}_i + \psi_i^*(\delta\psi_i)^\top + \delta\psi_i(\psi_i^*)^\top \tag{4.15c}$$
$$\psi_i = \psi_i^* + \delta\psi_i \tag{4.15d}$$
$$(\psi^*)_i^\top = [\bar{x}_i^\top, (u_i^0)^\top]^\top \tag{4.15e}$$
$$\delta\psi_i^\top = \begin{cases} [0^\top, v_i^\top]^\top, & \text{if } i = 0 \\ [\sum_{j=0}^{i-1} \bar{A}^{i-1-j}\bar{B}v_j^\top, v_i^\top]^\top, & \text{if } i > 0 \end{cases} \tag{4.15f}$$

The modified admissible input set $\tilde{\mathcal{U}}_N(x; \mathbf{v}^-)$ employs the *nominal* prediction model as part of the following constraints:

$$x_0 = x \tag{4.16a}$$
$$\bar{x}_{i+1} = \bar{A}x_i + \bar{B}(u_i^0 + v_i) \tag{4.16b}$$
$$\bar{x}_i \in \mathbb{X} \tag{4.16c}$$
$$(u_i^0 + v_i) \in \mathbb{U} \tag{4.16d}$$
$$\bar{x}_N \in \bar{\mathbb{X}}_f \oplus \{x_f(v_f^*)\} \tag{4.16e}$$
$$||v_i|| \leq E \tag{4.16f}$$
$$v_0 = v_1(k-1) \tag{4.16g}$$
$$||\mathbf{v}(k) - \tilde{\mathbf{v}}(\mathbf{v}(k-1))|| \leq \epsilon \tag{4.16h}$$

with $i = 0, \ldots, N-1$ and $v_f^* = v_f(\tilde{\mathbf{v}}(\mathbf{v}^-))$. Note that the main differences from the original optimisation problem $\mathbb{P}_2$ are:

1. **Input constraint translation** (4.16d); the constraint is equivalent to the constraint $v_i \in \mathbb{U}(\mathbf{v}) = \mathbb{U} \ominus \{u_i^0\}$, where the controls u_i^0 are obtained from solving regulation problem $\mathbb{P}_1(x; \mathbf{v})$.

2. **Terminal constraint modification** (4.16e); which is the same terminal constraint $\bar{x}_N \in \mathbb{X}_f(\mathbf{v})$ as in regulation problem $\mathbb{P}_1(x; \mathbf{v})$. The additional notation has been added to distinguish between the final term of the decision variable $\mathbf{v}$ and the final term of $\tilde{\mathbf{v}}(\mathbf{v}(k-1))$, which has already been used in $\mathbb{P}_1(x; \mathbf{v})$.

3. **Causality constraint** (4.16g); used to satisfy Assumption 7 and to remove the algebraic loop between $\mathbb{P}_1(x; \mathbf{v})$ and $\mathbb{P}_2^*(x; \mathbf{v}, \mathbf{u}^0(x))$.

4. **Rate-of-change constraint** (4.16h); the change from the tail-end of the perturbation sequence at the previous time-step is norm bounded for robustness reasons described in Section 4.6.

Note that constraints (4.16b)-(4.16h) are enforced on elements of $\mathbf{v}$, whilst the optimal sequence for regulation is a constant with respect to $\mathbb{P}_2^*(x; \mathbf{v}, \mathbf{u}^0(x))$, since it has already been computed for the current time-step. The control law that is analogous to (3.13) is therefore defined as the first term of the optimal perturbation sequence

$$\pi_N(x; \mathbf{v}) = v_0^0 \left(x; \mathbf{v}, \mathbf{u}^0(x)\right). \tag{4.17}$$

The input translation and modified terminal constraints ensure that the tail perturbation sequence $\tilde{\mathbf{v}}(\mathbf{v}(k-1))$ is always feasible (in the zero model mismatch scenario). This claim is explained in the following subsection.

Connected Feasibility

Assumption 11 requires careful consideration with respect to its effect on the feasibility of $\mathbb{P}_2^*\left(x; \mathbf{u}^0(x)\right)$. The following proposition is introduced to show that the tail-end perturbation sequence is feasible under the modified regulation problem $\mathbb{P}_1(x; \mathbf{v}(k))$.

Proposition 5 (Connected feasibility with preview)**.** *If $x \in \mathcal{X}_N(\mathbf{v}(k))$, then $\mathbb{P}_1(x; \mathbf{v}(k))$ is feasible and, moreover, $\mathbb{P}_2\left(x, \mathbf{u}^*(x); \mathbf{v}(k)\right)$ is feasible for any $\mathbf{u}^*(x; \mathbf{v}(k)) \in \mathcal{U}_N(x; \mathbf{v}(k))$ if $\mathbf{v}(k-1)$ is admissible at $k-1$ and $\mathbf{v}(k) = \tilde{\mathbf{v}}(\mathbf{v}(k-1))$ (but not necessarily optimal for $\mathbb{P}_2$).*

Proof. The proof follows directly from the nesting of the controllability sets in Theorem 2 and the definition of sequence $\mathbf{v}_i$, since $\tilde{\mathbf{v}}(\mathbf{v}_i) = \mathbf{v}_{i-1}$. □

The modified admissible input sequence set $\mathbf{v} \in \tilde{\mathcal{U}}_N(x; \mathbf{v}^-)$, depicted in Figure 4.1 shows how deviations from the the tail-end $\mathbf{v} = \tilde{\mathbf{v}}(\mathbf{v}^-)$ are feasible. Since the sequence $\mathbf{u}^*(x; \mathbf{v})^\dagger$ translates the original feasibility set $\mathcal{U}_N(x; \mathbf{v})$, the origin of set $\tilde{\mathcal{U}}_N(x; \mathbf{v}^-)$ defines an analogous "zero" solution to that of Proposition 2.

[†]which is $\mathbf{u}^0(x; \mathbf{v})$ if the excitation problem is only linearised once.

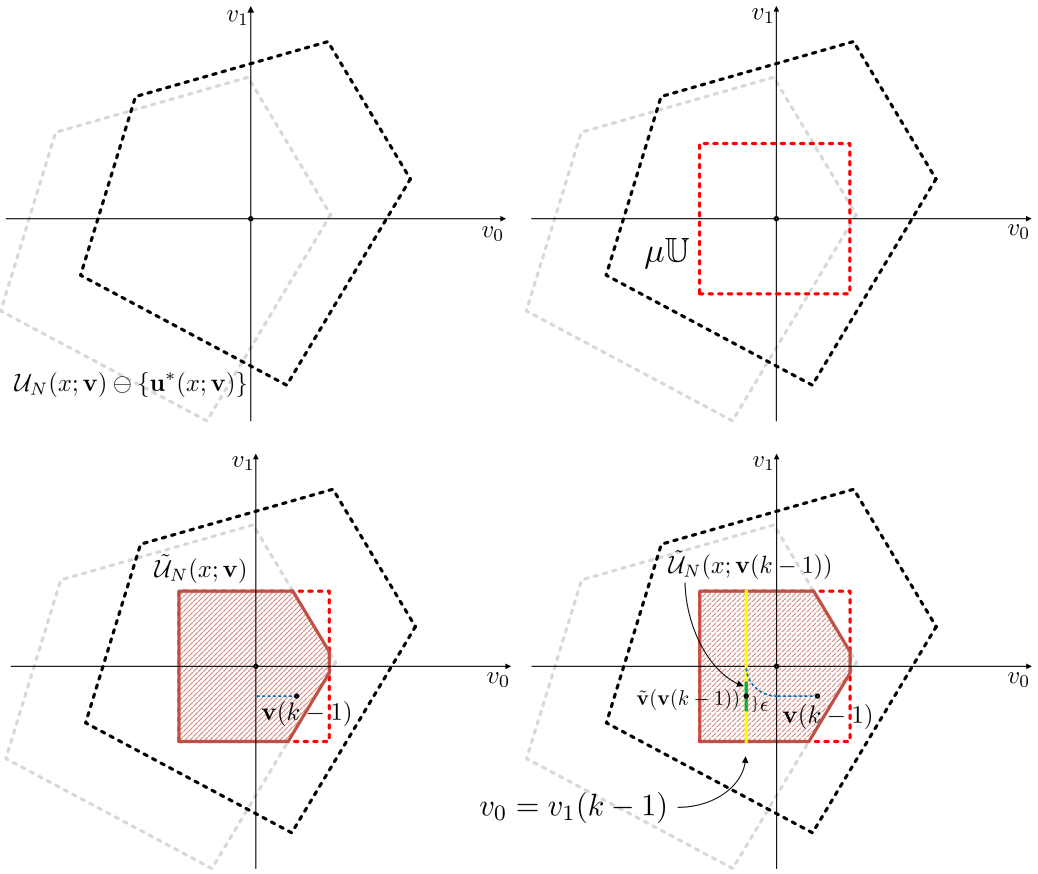

Figure 4.1: Construction of the feasible region $\tilde{\mathcal{U}}_N(x;\mathbf{v}(k-1))$ (denoted by the green line) of the optimisation problem $\mathbb{P}_2\left(x;\mathbf{u}^*(x),\mathbf{v}(k-1)\right)$ (left to right, top to bottom). The scaled constraint set $\mu\mathbb{U}$ is scaled by an arbitrary positive constant μ and is a polytopic approximation of (3.12).

Additionally, the same successive linearisation of the regressor, described in Section 3.5.2, can be directly applied to the modified $\mathbb{P}_2$ problem. In the case of modelling uncertainty, there may exist better excitation sequences (with respect to the $\mathbb{P}_2$ cost) that differ from the admissible tail-end $\tilde{\mathbf{v}}(\mathbf{v}(k-1))$ for the measured state x. Moreover, since the modified MPC controller acts in a receding horizon manner and attempts to reject disturbances (through the input term produced by $\mathbb{P}_1(x;\mathbf{v})$), PE may be lost at some point in the future. In an intuitive sense, the modified problem $\mathbb{P}_2^*(x;\mathbf{u}^0(x),\mathbf{v}^-)$ attempts to keep the excitation sequence PE, given some initial off-line excitation design. In the following section, it is shown how ϵ from (4.16h) is used as tuning parameter for trading-off between excitation and robustness.

4.6 Bounding the Rate of Change of the Perturbation Sequence

The main result of this section establishes an upper bound on the model error (the mismatch between $(\bar{A},\bar{B})$ and (A,B)) which, if met, guarantees stability of the closed-loop uncertain and perturbed system, under the preview information scheme. To achieve this, the analysis employs the exponential stability result of nominal MPC from Proposition 3 but with terminal conditions that invoke Proposition 4 (invariance of $\mathbb{X}_f(\mathbf{v})$). The inherent robustness of this newly modified controller is then exploited in order to handle the state prediction error caused by modelling uncertainty and a changing preview perturbation sequence.

Proposition 6 (Exponential stability of preview learning MPC)**.** *Suppose Assumptions 2-4 and 7-10 hold. Then, by Proposition 4, the equilibrium point for the terminal dynamics is exponentially stable for the nominal system* $\bar{x}^+ = \bar{A}x + \bar{B}(\bar{\kappa}_N(x;\mathbf{v}) + \pi_N(x;\mathbf{u}))$*, with region of attraction* $\mathcal{X}_N(\mathbf{v})$ *if* $\mathbf{v} = \tilde{\mathbf{v}}(\mathbf{v}(k-1))$ *according to Definition 9. Moreover, there exist constants* $c_2 > c_1 > 0$ *such that the value function satisfies:*

$$c_1||x||^2 \leq V_N^0(x;\mathbf{v}) \leq c_2||x||^2 \tag{4.18a}$$

$$V_N^0\left(\bar{x}^+;\tilde{\mathbf{v}}(\mathbf{v})\right) \leq \zeta V_N^0(x;\mathbf{v}) \tag{4.18b}$$

for all $x \in \mathcal{X}_N(\mathbf{v})$*, where* $\zeta \triangleq (1 - c_1/c_2)$.

This result follows exactly the same proof as Proposition 3 providing that the terminal ingredients have been modified as described in Section 4.3.1, the perturbation sequence is taken as the tail-end of the previously feasible sequence and no model uncertainty is present. Therefore, the states of the nominal regulated sy-

stem decay exponentially fast to the equilibrium point of the terminal dynamics. However, for the true system, model uncertainty yields:

$$\begin{aligned} x^+ &= Ax + Bu \\ x^+ &= (\bar{A} + \delta A)x + (\bar{B} + \delta B)u \\ x^+ &= \bar{A}x + \bar{B}u + \delta Ax + \delta Bu \end{aligned}$$

The successor state x^+, under closed-loop with the preview learning control law, is therefore:

$$x^+ = \bar{A}x + \bar{B}\underbrace{\left(\bar{\kappa}_N(x;\mathbf{v}) + \pi_N(x;\mathbf{u})\right)}_{\text{input}} + \underbrace{\delta Ax + \delta B\left(\bar{\kappa}_N(x;\mathbf{v}) + \pi_N(x;\mathbf{u})\right)}_{\text{disturbance}} \tag{4.19}$$

where $\delta A \triangleq A - \bar{A}$ and $\delta B \triangleq B - \bar{B}$. To simplify notation during analysis of the uncertain system, the following definition is made:

Definition 11 (Composite State)**.** *The composite state $z \triangleq (x, \mathbf{v})$, with the equivalent disturbance according to $\mathbf{w} = B \odot \mathbf{v} \in \mathbb{W} \times \cdots \times \mathbb{W} \times \mathbb{W}_f = \mathcal{W}$ leads to a new definition of the parametrised regions of feasibility:*

$$\mathcal{Z}_N \triangleq \{(x, \mathbf{v}) : x \in \mathcal{X}_N(\mathbf{v}), B \odot \mathbf{v} \in \mathcal{W}\}. \tag{4.20}$$

Under the properties of $\mathcal{X}_N(\mathbf{v})$ and $\mathcal{W}$ discussed in Sections 4.2.2 and 4.3, the set $\mathcal{Z}_N$ inherits convexity, boundedness, closedness and containment of the origin in its interior.

Assumption 12 (Bounded rate-of-change of excitation)**.** *The perturbation sequence evolves as $\mathbf{v}^+ = \tilde{\mathbf{v}}(\mathbf{v}) + \Delta\mathbf{v}$, where $\Delta\mathbf{v} = \mathbf{v}^+ - \tilde{\mathbf{v}}(\mathbf{v}) \in \Delta\mathcal{V} \subseteq \mathcal{V}$ such that $v^+ \in B\mathcal{V} \subseteq \mathcal{W}$.*

To deal with the uncertainty, the continuity property of the value function [106] is recalled:

Lemma 4 (Lipschitz continuity of the value function)**.** *The value function $V_N^0(\cdot)$ satisfies $||V_N^0(z_1) - V_N^0(z_2)|| \leq L||z_1 - z_2||$ over $\mathcal{Z}_N$, with Lipschitz constant L.*

Proof. Similarly to the properties of the value function with no preview information, this Lemma holds for the value function resulting from regulation problem (4.3). In the preview information formulation, Lipschitz continuity is defined with respect to the value function of the composite state z. The proof of Lemma 4 follows exactly the same argument of Theorem C.29 in Appendix C of [106] with the composite state z. □

Lipschitz continuity holds over bounded sets and affinity of the dynamics in this constrained control setting. Moreover, since the value function is the optimal cost of problem (4.3), the following analysis requires explicit characterisation of the value function. The guarantees based on Lipschitz continuity will no longer hold when the problem is extended to include non-linear dynamics. However, in this linear setting, an insightful robustness margin is developed in the following theorem.

Theorem 3 (Stability in terms of perturbation rate of change)**.** *Let Assumptions 2-4, 7-10 and 12 hold. Also, let* $z_0 = (x_0, \mathbf{v}(0)) \in \Omega_c^z \triangleq \{z : V_N^0(z) \leq c\} \subset \mathcal{Z}_N$ *for some positive constants* $c > b$*. If there exists a* $\rho \in (\zeta, 1)$*,* $E(\mu) \geq 0$ *(perturbation magnitude bound) and* $\epsilon \geq 0$ *(rate of change bound) such that,*

$$\epsilon \leq \frac{(\rho - \zeta)b}{L||\tilde{B}||} - \frac{(||\delta A x + \delta B \kappa_N(x; \mathbf{v})|| + ||\delta B|| E(\mu))}{||\tilde{B}||} \tag{4.21}$$

for all $z \in \Omega_c^z$*, then (i) the set* Ω_b^z *is positively invariant for the system* (4.19) *under the bounded rate-of-change of excitation; where* $\tilde{B} = diag[B, \ldots, B] \in \mathbb{R}^{Nn \times Nm}$*. Moreover, (ii) the set* Ω_c^z *is also positively invariant for* (4.19)*. Hence, starting from* $(x_0, \mathbf{v}(0)) \in \Omega_c^z$*, the states of the system enter the set* Ω_b^z *in finite time and remain therein.*

Proof. The proof utilises Lipschitz continuity of the true system's value function along with with the triangle inequality, is used to study the change between the value function under preview uncertainty and closed-loop state discrepancy. (i) Firstly, consider $z = (x, \mathbf{v}) \in \Omega_b^z \triangleq \{z : V_N^0(z) \leq b\}$ such that

$$V_N^0(x^+; \mathbf{v}^+) \leq V_N^0(x^+; \tilde{\mathbf{v}}(\mathbf{v})) + L||\tilde{B}||.||\mathbf{v}^+ - \tilde{\mathbf{v}}(\mathbf{v})|| \tag{4.22}$$

is formed from invoking Lipschitz continuity of the value function when *considering rate-limited deviation of the perturbation sequence* $\mathbf{v}^+$ *from the tail-end sequence* $\tilde{\mathbf{v}}(\mathbf{v})$. Additionally, in the nominal case, the descent property of the value function $V_N^0(\bar{x}^+; \tilde{\mathbf{v}})$ with preview information (Proposition 6) is

$$V_N^0(\bar{x}^+; \tilde{\mathbf{v}}(\mathbf{v})) \leq \zeta V_N^0(x; \mathbf{v}). \tag{4.23}$$

Invoking Lipschitz continuity for the value function *considering uncertainty in the dynamics caused by modelling error and excitation* yields

$$V_N^0(x^+; \tilde{\mathbf{v}}(\mathbf{v})) \leq V_N^0(\bar{x}^+; \tilde{\mathbf{v}}(\mathbf{v})) + L||w(x)||, \tag{4.24}$$

where $w(x) = \delta A x + \delta B \left(\bar{\kappa}_N(x; \mathbf{v}) + \pi_N(x; \mathbf{u})\right)$ is the disturbance due to the model

mismatch. Substituting (4.23) into (4.24) yields:

$$V_N^0(x^+;\tilde{\mathbf{v}}(\mathbf{v})) \leq \zeta V_N^0(x;\mathbf{v}) + L||w(x)|| \tag{4.25}$$

Suppose there exists a constant $\rho \in (\zeta, 1]$ that considers slower decay rates that still maintain the Lyapunov function property of the value function. Since $z = (x, \mathbf{v}) \in \Omega_b^z \triangleq \{z : V_N^0(z) \leq b\}$, substituting the bound (4.25) in to (4.22) yields:

$$V_N^0(x^+;\mathbf{v}^+) \leq \zeta b + L\Big(||w(x)|| + ||\tilde{B}||.\underbrace{||\mathbf{v}^+ - \tilde{\mathbf{v}}(\mathbf{v})||}_{\epsilon}\Big) \leq \rho b \tag{4.26}$$

Finally, the bound on the rate of change of the perturbation sequence ($||\mathbf{v}^+ - \tilde{\mathbf{v}}(\mathbf{v})|| \leq \epsilon$) is obtained when ϵ in (4.26) is made the subject.

$$\epsilon \leq \frac{(\rho - \zeta)b - L\left(||w(x)|| + ||\delta B||E(\mu)\right)}{L||\tilde{B}||}.$$

Therefore, if inequality (4.6) is satisfied, the trajectories remain in the smallest set $z_{k+1} \in \Omega_b^z$. (ii) The claim for finite time convergence to Ω_b^z; suppose Ω_c^z is the largest possible set within the parametrised region of feasibility $\mathcal{Z}_N$, and let $(x_0, \mathbf{v}(0)) \in \Omega_c^z \setminus \Omega_b^z$. If there exists a constant $\rho \in (\zeta, 1]$, then

$$\begin{aligned} V_N^0\left(x^+;\mathbf{v}^+\right) &\leq \zeta V_N^0\left(x_0;\mathbf{v}(0)\right) + (\rho - \zeta)V_N^0\left(x_0;\mathbf{v}(0)\right) \\ &\leq \rho V_N^0\left(x_0;\mathbf{v}(0)\right). \end{aligned}$$

Thus, $V_N^0\left(x_k;\mathbf{v}(k)\right) \leq \rho^k c$ for $x_0 \in \Omega_c^z \setminus \Omega_b^z$. Hence $V_N^0\left(x_{k'};\mathbf{v}(k')\right) \leq b$ after some finite k', implying $\left(x_{k'};\mathbf{v}(k')\right) \in \Omega_b^z$ and the perturbation rate-of-change is limited according to Assumption 4.6. □

Corollary 2. *If $||\delta B|| = 0$, then the stability guarantee of Theorem 3 holds provided that:*

$$||\delta A|| \leq \frac{1}{\sqrt{b/c_1}}\left(\frac{(\rho - \zeta)b}{L} - ||\tilde{B}||\epsilon\right) \tag{4.27}$$

Proof. The proof follows the similar reasoning as the proof for Corollary 1. With no uncertainty in the input matrix $\bar{B}$, then the inequality (4.21) simplifies to:

$$\epsilon \leq \frac{(\rho - \zeta)b - L||\delta A x||}{L||\tilde{B}||}. \tag{4.28}$$

Using the triangle inequality for a second time yields

$$\epsilon \leq \frac{(\rho - \zeta)b - L||\delta A||.||x||}{L||\tilde{B}||}. \tag{4.29}$$

Since the composite state converges to $(x_{k'}; \mathbf{v}(k')) \in \Omega_b^z$ by part (ii) of Theorem 3, (4.18a) bounds the state norm, providing the perturbation is rate-limited. Then, algebraic manipulation to make $||\delta A||$ the subject yields the inequality (4.27). □

Remark 6. *Under lack of input matrix uncertainty* ($||\delta B|| = 0$)*, the stability bound* (4.27) *is independent of the perturbation magnitude bound E. The consequence is that the control scheme employing the assumptions of Theorem 3 is inherently robust to size of perturbations, providing that perturbations are feasible.*

The permissible time-wise difference ϵ is therefore used as a tuning parameter and has to be chosen to satisfy (4.27) to guarantee robust stability for an assumed degree of uncertainty in the dynamics matrix A. The role of ϵ in the performance of the RLS algorithm is explored numerically in the following Section 4.8.

4.7 The PL-MPC Algorithm

Algorithm 4.1 details the procedure for implementing preview learning MPC. A simplified flow chart in Figure 4.3 has been added to outline the core steps of the algorithm. The flowchart highlights the difference between standard LMPC. The PL-MPC algorithm builds upon the results of Chapter 3 with the modifications outlined in sections 4.3-4.5, including the method to re-linearise problem $\mathbb{P}_2$.

Algorithm 4.1 Preview Learning MPC with RLS

Initialization

1: At $k = 0$, set $x = x(0)$, solve $\mathbb{P}_1(x; \mathbf{v})$ for $\mathbf{u}^0(x; \mathbf{v})$ and initialize $\mathcal{R}^-$, $\hat{\theta}^-$, $\mathbf{v}^-$ for online steps

Online

2: Measure x, compute $\mathcal{R}$ from (3.9) and recede to $\mathbf{v} = \tilde{\mathbf{v}}(\mathbf{v}^-)$

3: Solve $\mathbb{P}_1(x; \mathbf{v})$ for $\mathbf{u}^0(x; \mathbf{v})$

4: Solve $\mathbb{P}^*_{2,j=0}\left(x; \mathbf{u}^0(x; \mathbf{v}), \mathbf{v}^-\right)$ for $\mathbf{v}^0_{j=0}\left(x; \mathbf{u}^0(x)\right)$

5: **while** remaining time in sampling instant **do**

6: Update linearisation $\psi^*_{i,j+1} = \psi^*_{i,j} + \delta\psi_{i,j}$

7: Solve $\mathbb{P}^*_{2,j+1}\left(x; \mathbf{u}^0(x; \mathbf{v}), \mathbf{v}^-\right)$ for $\mathbf{v}^0_{j+1}\left(x, \mathbf{u}^0(x)\right)$

8: Update iteration index to $j + 1$

9: **end while**

10: Apply sum of first controls from each sequence to the system

$$u = \bar{\kappa}_N(x) + \pi_N(x) = u^0_0(x; \mathbf{v}) + v^0_{0,j}\left(x; \mathbf{u}^0(x; \mathbf{v})\right)$$

11: Wait one time-step, then go to step 2

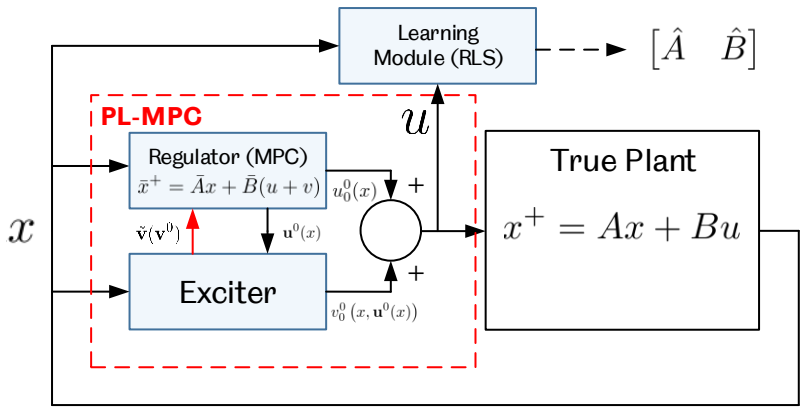

Figure 4.2: Block-diagram depicting the optimisations within PL-MPC control. The excitation module now informs the regulator of the perturbation sequence.

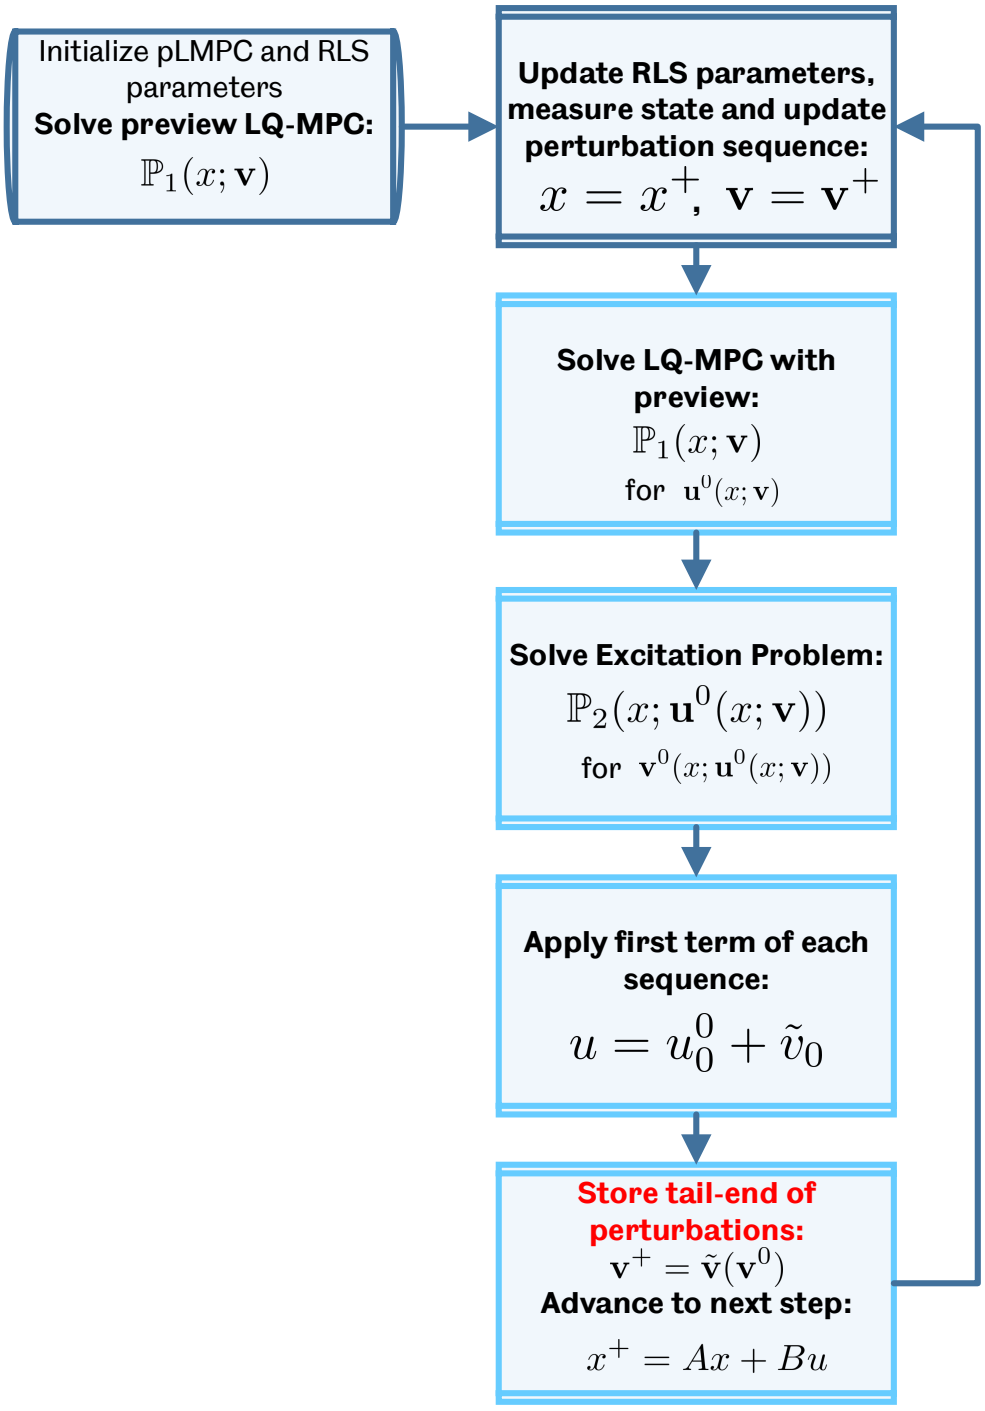

Figure 4.3: Flow chart describing the steps of a simplified PL-MPC algorithm (no iteration of the excitation optimisation).

4.7.1 Summary of Key Algorithmic Steps

Following the order of the flowchart in Figure 4.2, the key steps in algorithm include:

1. **Initialisation**: For a measured state x, the algorithm first solves the regulation problem $\mathbb{P}_1(x;\mathbf{v})$ using an off-line computed perturbation sequence $\mathbf{v}$. This control law is applied to advance the consecutive time-step.

2. **Measurement**: The tail-end of the perturbation sequence is receded to $\mathbf{v} = \tilde{\mathbf{v}}(\mathbf{v}(k-1))$. The state x is measured. The RLS equations are updated; including definitions of the regressor, parameter estimate and the information matrix.

3. **Regulation**: Using the measured state x, the regulation problem $\mathbb{P}_1(x;\mathbf{v})$ is solved with the pair $(x;\mathbf{v})$, yielding the control law $\bar{\kappa}_N(x;\mathbf{v})$.

4. **Excitation**: Using the measured state x, the excitation problem $\mathbb{P}_2(x;\mathbf{u})$ is solved with the pair $(x;\mathbf{u})$, yielding the control law $\pi_N(x;\mathbf{u})$. If there is remaining time left before the next sampling instant, then the excitation problem is re-linearised around an updated control sequence (as described in Section 3.5.2) and solved again.

5. **Control Assembly**: The first elements of the regulation and excitation optimisation solutions are summed to yield the PL-MPC control law $u = u_0^0 + v_0^0$. The notation in Figure 4.3 which uses $\tilde{v}_0$ emphasises the fact that the first term of the perturbation sequence is the perturbation that was computed at the previous time-step; v_0^0 and $\tilde{v}_0$ are numerically the same variable.

6. **Update**: The PL-MPC control law $u = u_0^0 + v_0^0$ is applied to the true system dynamics $x^+ = Ax + Bu$ to advance the state to the consecutive time-step. The perturbation sequence is stored such that $\mathbf{v} = \tilde{\mathbf{v}}(\mathbf{v}(k-1))$. The algorithm now returns to measurement step 2.

Note that steps 2-5 always occur at the same time-step; the time-index only advances between steps 1 to 2 and 6 to 2.

4.8 Numerical Study

The aim of this section is to numerically demonstrate the advantage of employing preview information of the input perturbation sequence, within a modified LQ-MPC setup. The practical implication of this modification–the enlargement of the closed-loop region of attraction– is a consequence of a corollary of Theorem 3:

Corollary 3 (Enlarged Region of Attraction)**.**

$$\bigcup_{k=0}^{\infty} \Omega_b(\mathbf{v}(k)) \quad \textit{with} \quad ||\mathbf{v}(k) - \tilde{\mathbf{v}}(\mathbf{v}(k-1))|| \leq \epsilon$$

where $\Omega_b(\mathbf{v}(k)) = \{x : (x, \mathbf{v}(k)) \in \Omega_b^z\} \subset \mathcal{X}_N(\mathbf{v}(k))$. Notice that for nominal MPC without preview information, the largest sub-level set $\Omega_b(\mathbf{v}(0))$ is the RoA when $\mathbf{v}(0) = \{0, \ldots, 0\}$. On the other hand, for the PL-MPC case, the union of all sub-level sets with every disturbance realisation produces the largest effective RoA [13]. This is a non-standard method of defining a region of attraction of the closed-loop uncertain system. The enlarged region arises from the translation of the nominal terminal ingredients according to the realisation of the preview disturbance. The individual sets for each disturbance realisation potentially yield smaller *instantaneous* regions of attraction $\Omega_b(\mathbf{v}(k))$ due to the required scaling; however, each instantaneous region may cover new portions of the state-space as a consequence of translating to new equilibria. Therefore, the union describes the effective robust region of attraction, given that the disturbance sequences are realized exactly as predicted. In the following example it is numerically shown that $\bigcup_{k=0}^{\infty} \Omega_b(\mathbf{v}(k))$ is made larger than the nominal RoA.

4.8.1 Simulating Worst-case Rates of Change

Since it is impractical to pre-compute the possible perturbation sequences resulting from $\mathbb{P}_2$ at each state $x \in \mathcal{X}_N(\mathbf{v})$ (to evaluate the value function $V_N^0(x; \mathbf{v})$), the sequences which deviate from the null perturbation, whilst satisfying the rate of change bound (4.21), are tested. Hence, the enlarged RoA is numerically evaluated at the vertices of a polytope that defines perturbation sequences with the largest rates of change.

$$\mathbf{v}^*(k) = \{0, v_1^*, \ldots, v_{N-1}^*\}$$

with each $v_i^* = \pm E\mathbf{1}$ for $i = 1, \ldots, N-1$ and $\mathbf{1} \in \mathbb{R}^m$. Therefore, a total of 2^{N-1} permutations of the sequence are tested over a discrete set of states within each $\mathcal{X}_N(\mathbf{v}(k))$. This implicitly assumes that $\mathbf{v}(k-1) = \{0, \ldots, 0\}$ was previously implemented perturbation sequence. Without limitation, these experimental conditions exercise the worst case scenario in terms of the feasibility of the problem.

4.8.2 Illustrative Example Parameters

The robust regions of attraction of the closed-loop system (4.19) under PL-MPC are constructed for the following system, with marginally stable dynamics:

$$x^+ = \underbrace{\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}}_{A=\bar{A}} x + \underbrace{\begin{bmatrix} 0 \\ 0.5 \end{bmatrix}}_{B=\bar{B}} u.$$

In this example, the prediction model is assumed to match the true dynamics; the point of this study is only to show and verify that the resultant robust region of attraction is enlarged. The sets $\mathbb{X}_f(\mathbf{v})$ are designed to satisfy the assumptions of Proposition 4 with the state and input constraints

$$\mathbb{X} = \{x \in \mathbb{R}^2 : |[x]_h| \leq 10, h = 1, 2\}$$
$$\mathbb{U} = \{u \in \mathbb{R} : |u| \leq 3\}$$

and the computed P from the linearised LQ-optimal terminal control law K. Table 4.1 shows the parameters chosen to comply with the PL-MPC description in this chapter.

Table 4.1: PL-MPC Controller parameters

Parameter	Value	Parameter	Value
R	1	E	0.3
Q	$5I_{2\times 2}$	N	3
P	$\left[\begin{smallmatrix} 8.9139 & -3.1250 \\ -3.1250 & 7.5615 \end{smallmatrix}\right]$	ϵ	1.56
b	800	α_x, β_u	0.5

4.8.3 Regions of Attraction

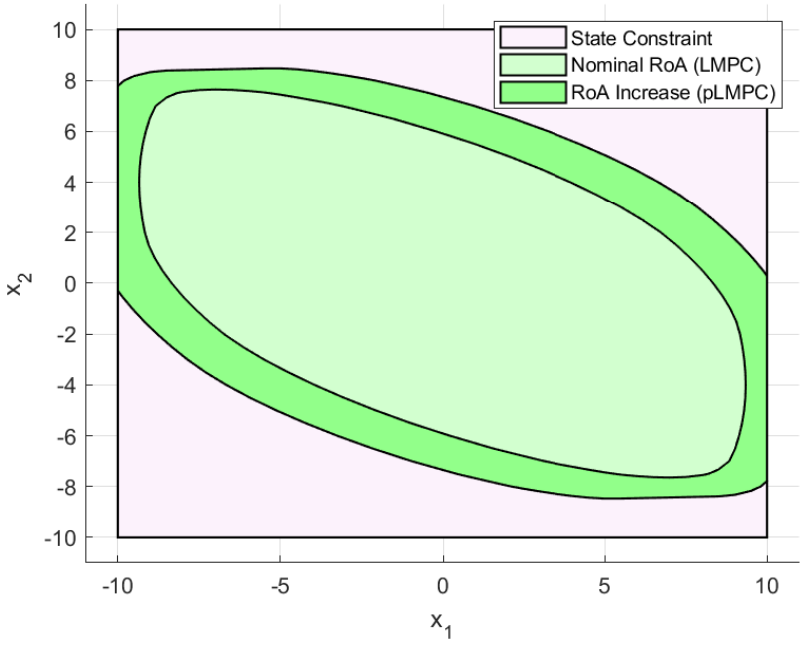

Figure 4.4: The RoA for the nominal LMPC and PL-MPC algorithms with nominal input constraint $\mathbb{U}$.

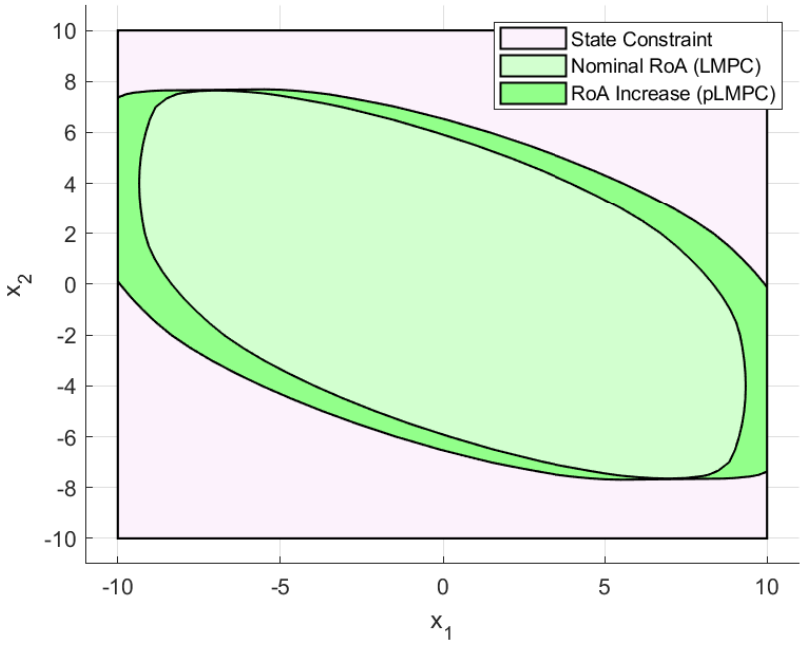

Figure 4.5: The RoA for the nominal LMPC and PL-MPC algorithms with input constraint translation $\mathbb{U}_i(\mathbf{v})$.

4.8.4 Discussion

Two cases in which satisfaction of Theorem 2 are explored; firstly, when the input constraint is simply $\mathbb{U}$; secondly, when the input constraint is modified to $\mathbb{U}_i(\mathbf{v})$. The simulation results show that the RoA is enlarged by 36% and 21% from $\Omega_b(\mathbf{v}(0))$, respectively. For both cases, the value function was evaluated over an evenly sampled grid contained within feasible regions $\mathbb{X}_f(\mathbf{v})$ and $\mathbb{U}_i(\mathbf{v})$, with approximately 8700 points.

State, input and parameter trajectories under PL-MPC converge as expected and have not been shown for brevity. For meaningful convergence plots, the reader is referred to the comparison in following section or the final numerical example where PL-MPC is applied to a linear GT model in Chapter 5.

The enlarged RoAs incorporate portions of the state constraints that would otherwise not be contained within the robust RoAs. The discussion of the robustness bound (4.27) follows analogously from the discussion from Section 3.7.2, but now in terms of the rate of change of the perturbation sequence between time-steps. The same numerical problems for evaluating the bound are inherited. However, the new bound (4.27) is now *independent of the excitation magnitude*, provided that the input constraint is modified to satisfy the new definition $\mathbb{U}_i(\mathbf{v})$; if the initialised excitation sequence $\mathbf{v}(k-1)$ is chosen such that the problem $\mathbb{P}_1$ is feasible, the controller's inherent robustness property only has to deal with deviations from the design perturbation sequence and modelling uncertainty. On the other hand, the standard LMPC formulation's inherent robustness has to deal with the entirety of the excitation sequence and modelling uncertainty. The modifications for preview information absorb a deterministic portion of the uncertainty such that more of the inherent robustness margin is reserved for modelling uncertainty and exogenous noise.

4.8.5 Comparison with LMPC

Figures 4.6 and 4.7 illuminate some of the differences between the behaviour of LMPC and PL-MPC; using the same numerical example as Chapter 3. These figures depict system trajectories under a PL-MPC control law that has been designed using the same parameters as in Table 3.1. The PL-MPC specific parameters α_x, α_u and ϵ are all selected as 0.5. A side-by-side comparison is complicated by the fact that standard LMPC does not allow the use of an a-priori designed excitation sequence. Hence, the simulation is initiated assuming a null perturbation sequence (to the disadvantage of PL-MPC).

Under a common choice of tuning parameters, the PL-MPC excites the system

more aggressively than LMPC. The resultant frequency-rich trajectories enable the RLS estimator to converge to the true dynamics roughly two-times faster than in the case of standard LMPC. Despite the small perturbations of standard LMPC, the true dynamics are identified since the PE condition 3.10 can be satisfied by any arbitrarily small constants. In reality, such small excitations may not even be implementable due to resolution of real actuators. Moreover, such small excitations will be obscured by real-sensor noise and disturbances. On the other hand, PL-MPC selects larger excitations despite the same magnitude constraints. It is speculated that controller's ability to predict the disturbing nature of the internally generated perturbation enables fuller use of the region of feasibility before reaching constraints. It was observed that the choice of forgetting factor γ had the greatest influence on the convergence time.

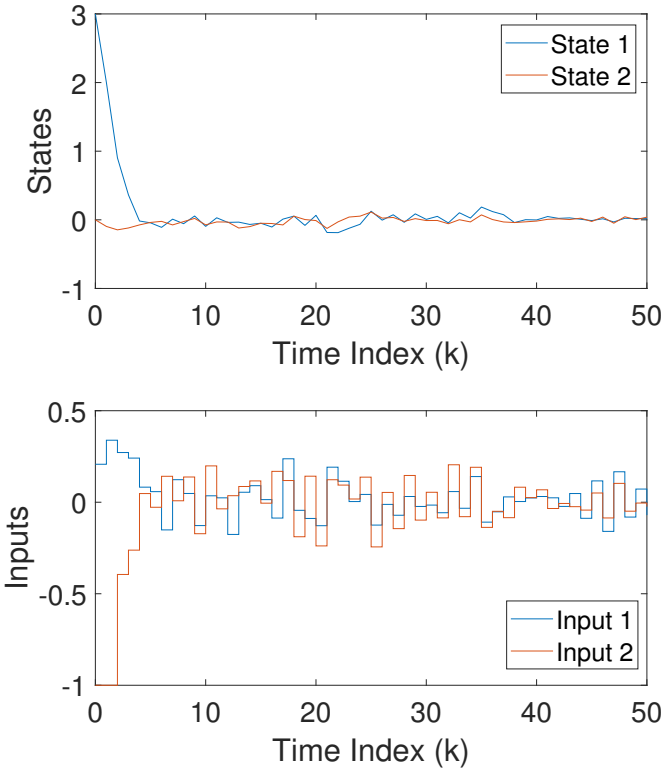

Figure 4.6: State and input trajectories under a PL-MPC control law, applied to the illustrative example of Section 3.9 with the same tuning parameters.

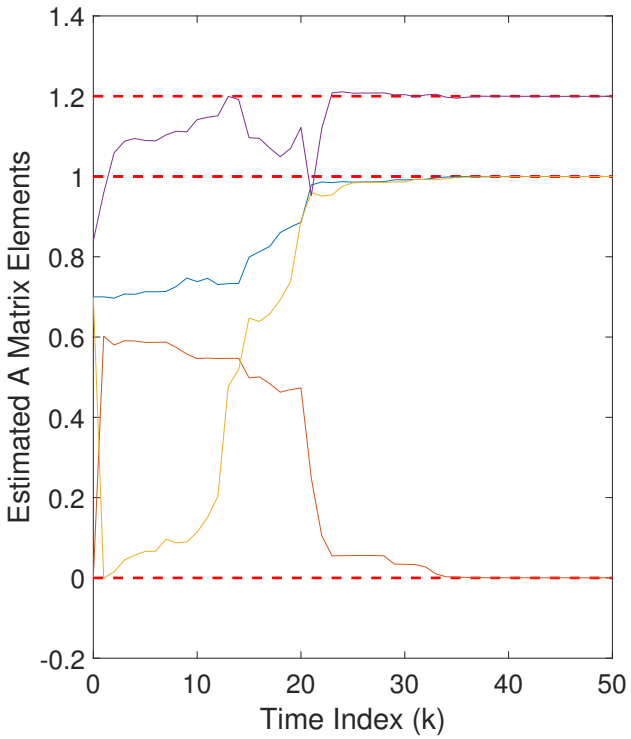

Figure 4.7: Estimated A-matrix parameters under PL-MPC for the illustrative example of Section 3.9. Dashed red lines denote the true parameters.

4.9 Summary and Conclusion

This chapter provides an extension of the LMPC controller, making it inherently robust to the perturbations that are required to excite the system for convergent closed-loop identification. Robust formulations of MPC have been avoided to mitigate unnecessary conservatism typically seen in robust control design. Provided that the modelling uncertainty is bounded and the rate change of the perturbation sequence is chosen to match this uncertainty through (4.27), robust performance of the closed-loop system is guaranteed under input and state constraints. Whilst the rate-of-change bound is conservative, the conservatism only affects the performance of the excitation optimisation - once parameter estimates converge, this excitation problem can be turned off without affecting the nominal regulating LQ-MPC controller.

It is emphasised that the additional computational complexity added is trivial. For completeness, these additional online complexities include:

- Storage of the perturbation sequence between time-steps.

- Online translations of the sets $\mathbb{X}_f$ and $\mathbb{U}$ to $\mathbb{X}_f(\mathbf{v})$ and $\mathbb{U}_i(\mathbf{v})$.
- Additional equality constraint due to Assumption 11 on excitation problem $\mathbb{P}_2$.

Only the tuning of this new controller is complicated by additional tightening factors β_x,β_u, excitation bound E and rate of change bound ϵ. However, these parameters are intuitive and simple to understand; the rationale of this proposed extension is well suited to the simplicity required in certifying safety-critical controllers. For a real engineering system, a-priori knowledge of good excitation signals can be utilised to design an offline excitation sequence for initialising the PL-MPC algorithm.

Given that the controller has been robustified with respect to its initial nominal modelling, the final logical step is to consider the conditions under which stable and feasible adaptation can be made to a newly identified model. The final technical chapter will demonstrate that a gas turbine model can be used within this inherently robust adaptive model predictive control approach.

Chapter 5

Adaptation in MPC under Gas Turbine Degradation

5.1 Introduction

The final contribution of this thesis considers the adaptation mechanism of the proposed adaptive MPC algorithm. The rationale of this mechanism is to switch to a LQ-MPC controller which utilises a refined prediction model and its associated terminal ingredients. Hence, after an update to a more accurate model, a larger proportion of the inherent robustness can be dedicated to uncertainty that cannot be modelled in practise; avoiding the conservatism of robust approaches. Whilst exhaustive simulations can show that serendipitously switching between controllers with small discrepancies in their prediction models may lead to stable and acceptable performance, there are no control-theoretic guarantees. Moreover, simulations for all possible pairs of nominal/identified prediction models under an assumed norm bounded uncertainty may be impractical.

Therefore, a survey of contemporary methods in MPC for conducting safe-updates, with an online identified prediction model, is conducted. The identified control-theoretic challenges include:

- Lyapunov stability of the closed-loop dynamics under switching.
- Feasibility of state and input trajectories between controller transitions.

Following the survey, mild technical assumptions are proposed to ensure that the MPC scheme can safely adapt the MPC controller's stabilising ingredients. Under these assumptions, the inherently robust adaptive control scheme is demonstrated in a simulation study of a hobby-scale gas turbine.

5.2 The Switched Control Problem

In this section, the problem of controlling a degrading plant is posed in a switched systems framework. This framework can be used to capture changes in the true system's dynamics. The control designer's aim is to synthesize a monitoring system that can correctly identify what the active mode of the system may be and therefore implement a control law that is most appropriate given the available knowledge.

Consider the true dynamics (2.1), where the exogenous vector $w(k)$ is now replaced with a switching signal $\sigma(k)$. Let these true dynamics be approximated by a linear system of the form:

$$x^+ \triangleq A_{\sigma(k)}x + B_{\sigma(k)}u \tag{5.1}$$

Definition 12 (Switched System)**.** *The signal $\sigma : \mathbb{Z}_{[0,\infty)} \to p \in \mathcal{P}$, where $\mathcal{P}$ is a compact set denoting the possible dynamics, and (A_p, B_p) defines the selected pair of plant matrices. Moreover, the set $\mathcal{P}$ contains both the nominal $(\bar{A}, \bar{B})$ and true (A, B) dynamics, with unique indexes $\bar{p}$ and p_t respectively.*

Definition 12 is used to represent the envelope of plant dynamics that can be safely controlled. The problem considered is one of autonomous switching; specifically, the sub-category of time-dependent switching with an *unknown* switching signal [76]. The objective is to update the LQ-MPC feedback law to a control design $u = \kappa_{N,p}(x)$ that corresponds to the active plant matrices. However, the only design that is available a-priori is the nominal control law design $u = \kappa_{N,\bar{p}}(x) = \bar{\kappa}_N(x)$.

Assumption 13. *The LQ-MPC regulator $\bar{\kappa}_N(x)$, designed for the nominal model* (3.2)*, asymptotically stabilises all plant dynamics (A_p, B_p).*

Assumption 13 is motivated by the the need to always have a controller that stabilises the currently active plant matrices. Without this assumption, it would be difficult to certify the control scheme on safety-critical aerospace system. The rationale for updating to $u = \kappa_{N,p}(x)$ is therefore to retain performance and robustness margins of nominal LQ-MPC designs. The closed-loop system under the this rationale is

$$x^+ = A_{\sigma(k)}x + B_{\sigma(k)}\kappa_{N,p}(x) \tag{5.2}$$

where the aim is to match the correctly designed control law $u = \kappa_{N,\sigma(k)}(x)$ to the currently active plant dynamics induced by $\sigma(k)$. Switched system stability and

feasibility issues naturally arise from this formulation. The following review of literature motivates further assumptions on the rate of the switching signal, such that the control laws proposed in the previous chapters can be readily applied in a quasi-time invariant setting.

5.2.1 Switching between Stable Sub-systems

In classical switched systems theory, where unconstrained control problems are considered, it is well-known that arbitrary switching between different globally stable dynamics can lead to overall instability of the switched system [76]. Conversely, arbitrary switching between different globally unstable dynamics can lead to overall stability of the switched system. This discussion focuses on the former and more relevant scenario, where dwell-time τ^s is employed to force a system to remain within a single stable operating mode, for a prescribed length of time [140].

Definition 13 (Stability Dwell-Time)**.** *An explicit computation of dwell-time is possible by comparing the Lyapunov functions of each mode before and after the switch such that*

$$\tau^s \geq ceil\Big[-\frac{\ln \mu}{\ln(1-\zeta)}\Big], \tag{5.3}$$

where the constant $\mu > 1$ denotes the allowable increase and $0 < \zeta < 1$ denotes entered mode's Lyapunov function's decay rate, respectively. The function ceil(n), returns the nearest integer that is greater than or equal to n due to the discrete time context.

In essence, stability dwell-time uses the decay rates of the value functions to dominate the sudden increases of the value between mode switches. Definition 13 holds in the discrete time domain framework. However, *average* dwell time is a more flexible approach since it permits more often switches when required, at the expense of prolonged single mode operation in the future [76]. Unfortunately, such considerations only apply to continuous time MPC systems [99].

Implication of dwell-time on MPC

There is a natural limitation for the choice of μ in the context of stability dwell-time for constrained MPC. The regions of attraction for the nominal and identified prediction model controllers are $\bar{\mathcal{X}}_N$ and $\hat{\mathcal{X}}_N$; both are compact sets under their respective stabilising MPC control laws (2.8). Hence, the maximum sub-level sets of the value functions $\bar{V}_N^0(x)$ for the nominal and $\hat{V}_N^0(x)$ for the identified mode, are upper bounded, thereby bounding μ. Moreover, despite a stability dwell-time being satisfied, the trajectories of a particular mode may not necessarily be

feasible for other modes [77], motivating use of a feasibility dwell-time τ^f. Prior-art tackling both feasibility and stability dwell-time in adaptive MPC is reviewed in the following section.

5.3 Switching for Adaptation in MPC

Switching between stabilizing MPC controllers in a DT setting was first considered in [29]. These early works considered switching systems with prescribed switching signals; including time uncertainty around each switching instant [96].

The predecessor paper of the adaptive MPC scheme in [37] describes how stability of the terminal dynamics under a changing prediction model can be preserved, without burdensome online recomputation of the terminal cost and control law [36]. This approach requires a-priori known modes that are parametrised by convex combinations of the A-matrix polytopic uncertainty set's vertices. Satisfaction of an appropriate LMI condition enforces cost decrease between two consecutive modes. A precomputed common robust PI set *for all modes* is used, introducing conservatism. Ultimately, these additional stability and feasibility technicalities depend on the true system being contained within the polytopic uncertainty set.

The LQ-MPC dual-mode paradigm for a DT switched linear system was first considered in [22]. In this paper, nominal stability and feasibility is guaranteed under modifications of the terminal constraints, if the switching signal is known exactly along the prediction horizon [22]. Since the dynamics of each mode are known a-priori, the off-line computed terminal ingredients that satisfy Assumption 5 (individually for each mode) are used to compute mode-wise common control invariant common control invariant (CCI) sets. Given an admissible switching signal, the terminal constraints are enlarged by using the dwell-time in a particular mode, to steer the system into a mode-wise PI set. The dwell-time effectively extends the MPC's control horizon, without adding additional decision variables to the online optimisation. The main assumption of this approach is that the modified terminal constraint sets of each mode must also be nested within the τ-step controllability set of every other mode (computed using equation (4.2)), which reduces the tolerable degree of collective heterogeneity in all of the modes. Whilst the computation of these sets can be performed offline, these sets may become restrictive; the numerical example only considers two modes. However, this approach is still less conservative than naively computing one single common invariant set for all modes. Ultimately, there is an increased memory requirement since the relaxation of the terminal constraints for each mode uses $\tau - 1$ number of relaxed terminal constraints. Moreover, the explicit characterisation of a mini-

mum dwell time for the system, with respect to the value function of each mode, is not explicitly stated. An extension of this work removed the requirement of a known switching signal over the horizon, but still requires a-priori knowledge of the minimum stability dwell-time for each mode and a directed graph of admissible mode transitions [32].

In [61], switching between a-priori known modes is considered in a robust tube MPC formulation [106], for an uncertain switched linear system with additive disturbances. The exponential decay properties of the value function are utilised to explicitly characterize conservative minimum mode-dependent dwell-times (MDTs) for stable and feasible mode transitions, despite the assumed heterogeneity of each mode. The proposed approach utilizes invariant multi-sets [11] in a transition controller to ensure feasibility of the admissible mode switches. However, to reduce the computational complexity of computing the bounds of the MDT, conservatism is introduced through using a 1-norm ball to conservatively bound the 2-norm of state trajectories. The exponential decay property characterizes the instances in which the state's trajectory enters the region of feasibility of an adjacent mode. This is the key difference from a similar approach proposed in [141], where the controllability sets are used to compute the MDTs. In general, computing the controllability sets for high dimensional systems can become computationally intractable, whereas the bounded 2-norm trajectories using a 1-norm proxy is relatively straightforward. Additionally, in [61], the use of multi-sets leads to larger approximations of the minimal robust positive invariant sets, leading to smaller regions of attraction than in a nominal MPC control implementation (*i.e.*, the formulation of Section 2.1.4). Due to the relative simplicity of [61] compared with [141], the characterisation of [61] is computationally more efficient, yet still difficult to be performed online.

Tackling feasibility issues has been extensively studied in both continuous and discrete time settings [35]. Feasible switching is shown in [100] using a min-max MPC approach, where under a known dwell-time restriction on the switching signal, all admissible permutations of the switching signal are considered over the MPC horizon. In this approach, only the switching signal at the current time-step is required. Feasibility between switches was enforced by implementing specialised input constraints, named consistency constraints. Whilst recursive feasibility is guaranteed, the paper shows that the min-max approach is more conservative than a nominal MPC implementation on a nominally feasible example.

A totally different perspective on safe transitioning between controllers has been studied in [127]. In this work, an a-priori safe set and safe controller are known, whilst a learning control law is used to probe a disturbed linear polyto-

pic difference inclusion. A tube-like MPC is employed to ensure that the learning control law, which is unaware of system constraints, evolves under feasible trajectories. This "safe" control law minimises the norm between a back-up control law and the probing control signal such that the the minimum objective is obtained when applying the probing control policy. If the predicted trajectories under the probing control law violate constraints, the back-up control law selects a control that satisfies the constraints whilst deviating as little as possible from the probing control signal. Under conservative modifications of the terminal safe-set, the scheme is shown to be recursively feasible. The conservatism is alleviated online by identifying and growing the initial safe-set with states that are verified to be feasible under the learning control law; resulting in a so-called model predictive safety certification scheme.

To the best of the author's knowledge, no work has been conducted on utilising the LQ-MPC's cost function's continuity to handle infrequent mode switching for small heterogeneity.

5.4 Performance Monitoring

The main idea of the inherently robust adaptive controller is to use an online identified prediction model, to avoid the conservatism and complexity of available algorithms in adaptive MPC literature. Whilst theoretical guarantees of convergence of the proposed RLS algorithm have been well-established, the measurements used to perform the identification will undoubtedly become corrupted by real world phenomena. Sophisticated feature extraction and signal processing may alleviate some degree of corruption from noise, but the glaring question remains: how can the engineer be certain that updating to an online obtained prediction model will lead to performance enhancement?

5.4.1 Explicit Performance Monitoring

An explicit answer can be found by exercising the model in a safe manner. Before such an exercise is proposed, it is useful to survey the juvenile field of controller performance monitoring. Much of the technical solutions to controller performance monitoring remain occluded by application specific and in-house developed solutions [138]. The survey [138] presents the most recent discussion of performance monitoring solutions in the context of PI, LQG and MPC. Classical performance metrics, including minimum variance based metrics, become complicated under multivariable and constrained systems. For MPC, there is no standardised method for performance supervision.

In [138] a key performance indicator (KPI) is derived using the linear quadratic Gaussian (LQG) controller's value function. The unconstrained setup of this problem enables the analytical computation of the KPI for both model mismatch and modelled disturbance scenarios. However, for an equivalent MPC controller using the same model of the dynamics and the disturbance, the KPI fails to converge to the analytical expected KPI; this is accredited to the constraints becoming active. Therefore, despite having an accurate model, the KPI may falsely infer inferior performance. Since the KPI converges to a constant value, it is suggested that Monte-Carlo simulations can be used to compute the constrained KPI equivalent, as a replacement to the analytical KPI.

Explicit Performance Assessment Metrics in MPC

In summary, two metrics for performance monitoring have been identified:

- Time-averaged value function $\tilde{V}_N^0$
  - Requires solution of the online optimisation that uses the identified model as the prediction model.
- Prediction model error $||x^\top - (\psi^-)^\top \hat{\theta}||$

5.5 Prediction Model Adaptation

Theoretical considerations for safe controller updates, in the presence of uncertainty, require careful consideration for preserving stability and feasibility of nominal LQ-MPC controllers. This section states the assumptions required to permit adaptation of the PL-MPC to a more accurate prediction model.

Once the estimates $(\hat{A}, \hat{B})$ converge to the true parameters (A, B), the solution of the excitation problem $\mathbb{P}_2$ is no longer required; the exciter in Figure 4.2 is effectively turned off. Clearly, with the knowledge of an improved prediction model, LQ-MPC can be updated for performance gains. However, the non-trivial question of when to safely update the LQ-MPC emerges. This problem is modelled by considering the switch between two closed-loop systems:

1. **Nominal Controller, True Dynamics**: $x^+ = Ax + B\bar{\kappa}(x)$, with the nominal LQ-MPC controller $u = \bar{\kappa}_N(x)$.
2. **Adapted Controller, True Dynamics**: $x^+ = Ax + B\kappa(x)$, with the updated LQ-MPC controller $u = \kappa_N(x)$.

Notice that only the regulator part of the adaptive PL-MPC is considered. Moreover, the regulator's dependency on $\mathbf{v}$ is removed, since the excitations have been turned off. The problem is now tantamount to that of a classical switched system [76], albeit with constraints and an unknown switching signal. The identified technical problems of constrained switching in MPC, discussed in Section 5.3, are recalled:

- Lyapunov stability of the closed-loop dynamics under switching.
- Feasibility of state and input trajectories between controller transitions.

Inherent robustness of LQ-MPC, without additional assumptions, does not provide explicit guarantees for dealing with these technical issues. Several assumptions on the nature of the model uncertainty are proposed to ensure safe updates in the LQ-MPC setting. Hence, the proposed adaptive PL-MPC controller avoids the use of robust/stochastic formulations to preserve the simplicity of a nominal LQ-MPC design.

5.5.1 Stability Considerations

In the LTI setting, only a single switch is required to adapt to the LQ-MPC for the true dynamics (3.1). Moreover, after convergence of the estimates to the true model, finite time is required to re-compute stabilising ingredients of Table 5.1.

Table 5.1: LQ-MPC Stabilising Ingredients

Ingredient	
Terminal control law	$u_f = Kx$
Terminal invariant set	$\mathbb{X}_f$
Terminal weighting matrix	P

Therefore, whilst the stability ingredients are being computed, dwell-time requirements can be potentially satisfied under normal operation; the time required to compute the LQR control law K, the PI set $\mathbb{X}_f$ and weighting matrix P for the system (3.1), is likely to be longer than the dwell-time for stability. In the case that there is excess computational power available, the following assumption is introduced.

Assumption 14. *Let the time required to compute the LQR control law K, the PI set $\mathbb{X}_f$ and weighting matrix P for the system* (3.1) *be τ^c. If τ^c is less than the minimum dwell-time τ^s for stability when switching between $x^+ = Ax + B\bar{\kappa}_N(x)$ to $x^+ = Ax + B\kappa_N(x)$, the system is held in mode $x^+ = Ax + B\bar{\kappa}_N(x)$ for at least $\tau^s - \tau^c$ additional steps.*

By performing the update only after the stabilising ingredients are computed, dwell-time for stability is satisfied. Whilst re-computing these ingredients online is non-ideal from a computational perspective, this only has to be performed when a new model is identified. For slowly-varying systems, this may only have to be done a few times. The work of [37] avoids the computational burden of re-computing these stability ingredients for vertices of a model uncertainty polytope, but at the cost of additional conservatism. In this work, it is deemed reasonable to re-compute the stabilising ingredients, since it is only done once for the identified LTI true dynamics.

5.5.2 Feasibility Considerations

The more difficult problem is that of guaranteeing feasibility, since the state trajectories under the nominal control law $u = \bar{\kappa}_N(x)$ may not lie within the region of feasibility $\mathcal{X}_N$ for the system $x^+ = Ax + B\kappa_N(x)$. By utilising the inherent robust properties of LQ-MPC, the following proposition is explored to derive conditions for a feasible update.

Proposition 7 (Feasible Update)**.** *Suppose Assumptions 2-5 hold such that the nominal closed-loop system $\bar{x}^+ = \bar{A}x + \bar{B}\bar{\kappa}_N(x)$ under LQ-MPC control law $u = \bar{\kappa}_N(x)$ is exponentially stable, in-line with Proposition 3. Consider Corollary 1, but with $E = 0$ implying excitations have been turned off, such that the system $x^+ = Ax + B\bar{\kappa}(x)$ converges to a smaller robust positive invariant set $\bar{\Omega}_{b'}$, where $b' < b$. Additionally, let $\mathcal{X}_N$ be the region of feasibility of $\mathbb{P}_1(x)$ for the true system* (3.1) *under the control law $u = \kappa_N(x)$.*

If $x \in \bar{\Omega}_{b'} \subset \mathcal{X}_N$ then the switch from $x^+ = Ax + B\bar{\kappa}(x)$ to $x^+ = Ax + B\kappa(x)$ is feasible.

Unfortunately, Proposition 7 is not helpful for a-priori design, since $\mathcal{X}_N$ can only be obtained after learning the true dynamics. Moreover, the size of the set $\bar{\Omega}_{b'}$ depends on the magnitude and structure of the uncertainty, and has to be obtained from simulating the closed-loop. However, the implicit assumptions on the model uncertainty (Assumption 13), allows the set $\bar{\Omega}_{b'}$ to be be collapsed to the origin.

Proposition 8. *If Assumption 13 holds, the smallest robust positive invariant for the system $x^+ = Ax + B\bar{\kappa}_N(x)$ is $\bar{\Omega}_{b'} = \{0\}$.*

The final assumption is introduced to allow the existence of control laws with the same tuning weightings as the nominally designed LQ-MPC.

Assumption 15. *The LQ-MPC regulator $\kappa_N(x)$ exists for the same weighting matrices Q, R and an updated P that satisfies Assumption 5 with the true system* (3.1)*.*

Since the the nominal control law exponentially stabilizes (5.5) to the smallest robust positive invariant set $\bar{\Omega}_{b'} = \{0\}$ and $\mathcal{X}_N$ contains the origin by construction, the state is guaranteed to enter $\mathcal{X}_N$ at some future time. The update from $x^+ = Ax + B\bar{\kappa}_N(x)$ to $x^+ = Ax + B\kappa_N(x)$ can therefore be performed after checking whether the measured state satisfies $x \in \mathcal{X}_N$. By simulating the proposed control algorithm on the SPT5 dynamics, Assumptions 13 and 15 are shown to hold for the gas turbine example in the following section.

5.6 JetCat SPT5 Control Implementation

In this section, a simulation study of the inherently robust adaptive control law is conducted on an example multivariable gas turbine model: a modified JetCat SPT5 turbo-prop [102]. This case-study two-shaft gas turbine is a 6kW hobby-scale engine that has been fitted with a variable pitch propeller, enabling variable pitch to be used as an additional input degree of freedom. The regulation of the engine's shaft speeds, to those at cruise, is demonstrated under pre-specified prediction model error; emulating a typical engine acceleration command under degraded plant conditions.

5.6.1 SPT5 Dynamics

The SPT5's dynamics at the cruise operating condition have been extracted from the LPV description of [101]. For brevity, the following discussion will use the term CRZ to refer to the linear model at cruise. The authors obtained linear models by linearising a high-fidelity non-linear physics based model, which was previously validated with data from an experimental set-up; depicted in Figure 5.1. For a complete description of the validation approach, the reader is directed to the work of [102]. In this chapter, the CRZ model is utilised to show a proof-of-concept of the adaptive PL-MPC algorithm around a single operating condition.

Modelling Degraded Engine Health

Since degradation is a complex phenomenon, that is unique to each engine, there is no easy method of predicting how each engine's dynamics will degrade. In this study, the convention is to use the CRZ model from [101] as the true *degraded* behaviour of the plant. On the other hand, a nominal model is used to represent dynamics that were deemed true at the beginning of service life. The nominal dynamics are artificially fabricated by scaling the true dynamics matrices by a constant Ξ. This inverted approach is used to demonstrate that the proposed algorithm is able to handle arbitrary model error; provided that the error is norm bounded in-line with the results of Theorem 3 and its related assumptions.

JetCat SPT5 Engine Specification

The data provided in tables 5.2 and 5.3 provide a specification of control related variables. It is highlighted that input rate of change constraints (slew rates) have been assumed using engineering judgment, as these were not disclosed by the authors of the experimental rig [101, 102].

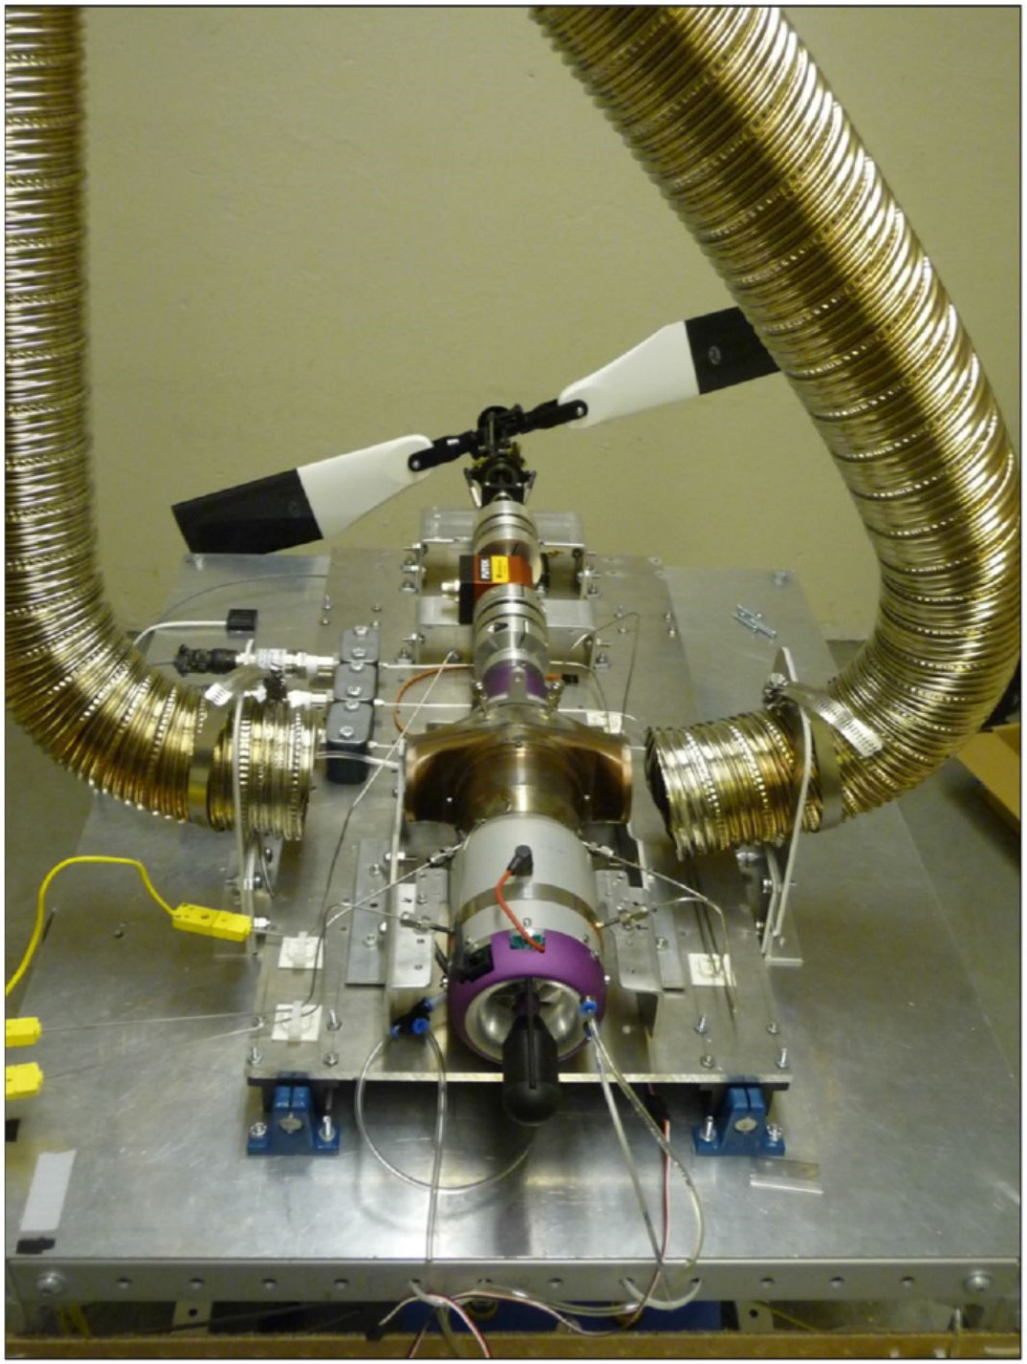

Figure 5.1: Head-on view of the modified JetCat SPT5. This image has been reproduced with the kind permission of the authors of [101].

Table 5.2: JetCat SPT5 Parameters

Description (Variable)	Normalisation (Units)	Constraints
HP Shaft Speed (N_2)	170,000 (RPM)	$50,000 \leq N_2 \leq 170,000$
Prop. Shaft Speed (N_1)	7000 (RPM)	$1500 \leq N_1 \leq 7000$
Fuel Flow (W_f)	3.532×10^{-3}(kg/s)	$(0.512 \leq W_f \leq 3.532) \times 10^{-3}$
Fuel Flow Rate ($\dot{W}_f$)	N/A (kg/s^2)	$? \leq \dot{W}_f \leq ?$
Prop. Pitch Angle ($[u]_2$)	N/A (deg)	$0 \leq [u]_2 \leq 35$
Pitch Angle Rate ($[\dot{u}]_2$)	N/A (deg)	$? \leq [\dot{u}]_2 \leq ?$

Table 5.2 describes the main engine variables; the states, inputs and their respective constraints used for control. These parameters have been extracted from the manufacturer's data-sheet [101].

Definition 14. *The input rate of change is defined as* $\Delta u = u - u^-$. *Moreover, it is bounded such that* $\Delta u \in \Delta\mathbb{U} \subseteq \mathbb{U} \subset \mathbb{R}^m$.

The JetCat SPT5's dynamics model has been provided in terms of normalised variables. Since the aim of this work is to study control around CRZ, the constraints on each variable are appropriately translated. These translated CRZ-specific constraints are presented in Table 5.3. The translation is made using the CRZ equilibrium point *i.e.,* the target steady-state and input from Table 5.4. These translations are imperative for satisfaction of the second part of Assumption 3. Interestingly, the propeller pitch angle was not normalised by the authors of [101].

Table 5.3: JetCat SPT5 constraints around CRZ

Normalised Variable	Translated Variable	Normalised Constraints
HP Shaft Speed (N_2)	$[x]_1$	$-0.432 \leq [x]_1 \leq 0.274$
Prop. Shaft Speed (N_1)	$[x]_2$	$-0.286 \leq [x]_2 \leq 0.500$
Fuel Flow (W_f)	$[u]_1$	$-0.324 \leq [u]_1 \leq 0.532$
Fuel Flow Rate (ΔW_f)	$\Delta[u]_1$	$-0.08 \leq \Delta[u]_1 \leq 0.08$
Prop. Pitch Angle ($[u]_2$)	$[u]_2$	$-16 \leq [u]_2 \leq 19$
Pitch Angle Rate ($\Delta[u]_2$)	$\Delta[u]_2$	$-0.25 \leq \Delta[u]_2 \leq 0.25$

The notation $[x]_h$,$[u]_p$ denotes the h-th state ($h = 1, 2$) and p-th input ($p = 1, 2$) for the vectors $x_k \in \mathbb{R}^n$, $u_k \in \mathbb{R}^m$ respectively. For the JetCat SPT5, $n = m = 2$. The Δ operator is used to denote the discrete-time rate of change $\Delta[u_h]_k = [u_h]_k - [u_h]_{k-1}$ between consecutive time-steps.

True Dynamics

The SPT5's continuous time dynamics at cruise (CT-CRZ) are:

$$\dot{x} = \underbrace{\begin{bmatrix} -1.700 & 0.100 \\ 0.600 & -1.100 \end{bmatrix}}_{A_c} x(t) + \underbrace{\begin{bmatrix} 1.200 & 0.000 \\ 0.300 & -0.023 \end{bmatrix}}_{B_c} u(t). \tag{5.4}$$

The CT-CRZ dynamics are discretised using zero-order hold sampling with a period of $t_s = 50$ms, yielding:

$$x_{k+1} = \underbrace{\begin{bmatrix} 0.919 & 0.005 \\ 0.028 & 0.947 \end{bmatrix}}_{A} x_k + \underbrace{\begin{bmatrix} 0.058 & 0.000 \\ 0.016 & -0.001 \end{bmatrix}}_{B} u_k. \tag{5.5}$$

The control problem is studied in the discrete-time framework because the proposed MPC controller requires finite time to compute the feedback signal; the inter-sample period limiting the available time for computing solutions to both $\mathbb{P}_1$ and $\mathbb{P}_2$. The zero-order-hold sampling rate $f_s = \frac{1}{t_s}$ is chosen to be at least 10 times faster than the fastest dynamic mode of the CT-CRZ model.

The resultant normalised discrete-time dynamics (DT-CRZ) (5.6) are controllable and open-loop stable. Moreover, it can be observed that the states are unidirectionally coupled *i.e.*, the affect of the propeller shaft dynamics on the HP shaft speed is more than five times the magnitude of the HP shaft's dynamics affect on the propeller speed. Additionally, propeller pitch input does not directly affect the HP shaft speed. This is a consequence of the engine's free turbine architecture [116].

Nominal Model Dynamics

The nominal discrete time dynamics (DT-CRZ) are

$$x_{k+1} = \underbrace{\Xi \begin{bmatrix} 0.919 & 0.005 \\ 0.028 & 0.947 \end{bmatrix}}_{\bar{A}} x_k + \underbrace{\begin{bmatrix} 0.058 & 0.000 \\ 0.016 & -0.001 \end{bmatrix}}_{\bar{B}} u_k. \tag{5.6}$$

This nominal model emulates a-priori control modelling of a healthy and undegraded engine. As the engine degrades, the performance and robustness margins of the nominal MPC design will deteriorate. By learning the true SPT5 dynamics around CRZ, the prediction model of the controller (and its stabilising ingredients) can be updated to recover robustness and optimality of an ideal LQ-MPC design.

5.6.2 Regulation to Cruise Equilibrium

The control problem is the regulation of the true SPT5 dynamics (5.5), from an initial state x_0, to the target equilibrium point x_e associated with CRZ. These initial and target states are defined in Table 5.4.

Table 5.4: Initial Condition and Target Equilibrium

Equilibrium	Normalised		Translated	
	$[N_1; N_2]^\top$	$[W_f; [u]_2]^\top$	$[[x]_1; [x]_2]^\top$	$[[u]_1; [u]_2]^\top$
Initial x_0	$[0.5327; 0.3678]^\top$	$[0.4685; 16]^\top$	$[-0.1937; -0.1322]^\top$	$[0; 0]^\top$
Target x_e	$[0.7264; 0.5]^\top$	$[0.4685; 16]^\top$	$[0; 0]^\top$	$[0; 0]^\top$

The initial state corresponds to operation near the E4 equilibrium point described in [101], which is a condition between CRZ and idle. The target state is the CRZ equilibrium. In the translated reference frame, the target state is the origin.

Assumption 16 (Equilibrium Points)**.** *Both the initial and target states are known a-priori and correspond to desired thrust levels.*

The objective of the PL-MPC controller is two-fold; firstly, the states must be regulated from the initial state to the target equilibrium whilst satisfying state and input constraints of Table 5.3. Secondly, the estimate of the open-loop dynamics must converge to the true dynamics at the CRZ condition. When the true dynamics (5.5) have been identified, the MPC ingredients are updated to improve closed-loop performance and robustness.

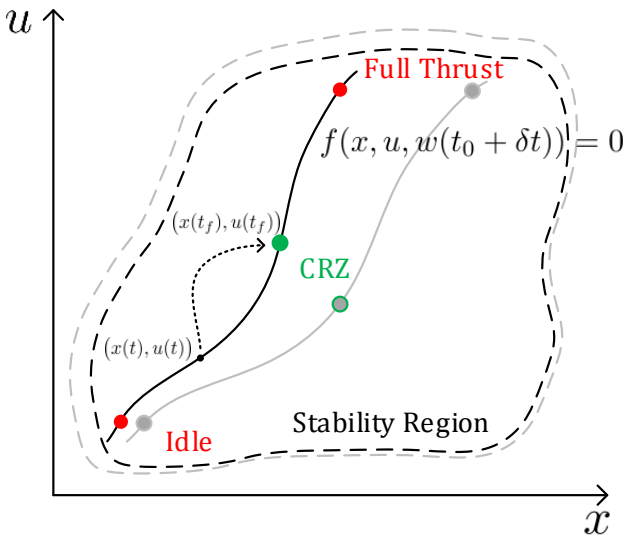

Figure 5.2: 2D schematic of regulation from an arbitrary lower power condition to the cruise equilibrium point. The PL-MPC scheme utilises the linear model associated with the CRZ point of the nominal non-linear plant (greyed out solid line). Note that this schematic is not to scale; it is an abstraction to help the reader understand the real high-dimensional control problem.

5.6.3 Simulation Results

The following simulations demonstrate the efficacy of the adaptive PL-MPC controller when applied to linear GT dynamics. These results are generated using the parameters from the problem setup of section 5.6. The plots are used to investigate the differences between a nominal LQ-MPC implementation and the updated LQ-MPC controller. The chosen tuning parameters are presented in Table 5.5. The reader is directed to Appendix 6.7.3 for the tuning method.

State and input trajectories of PL-MPC and nominal LQ-MPC

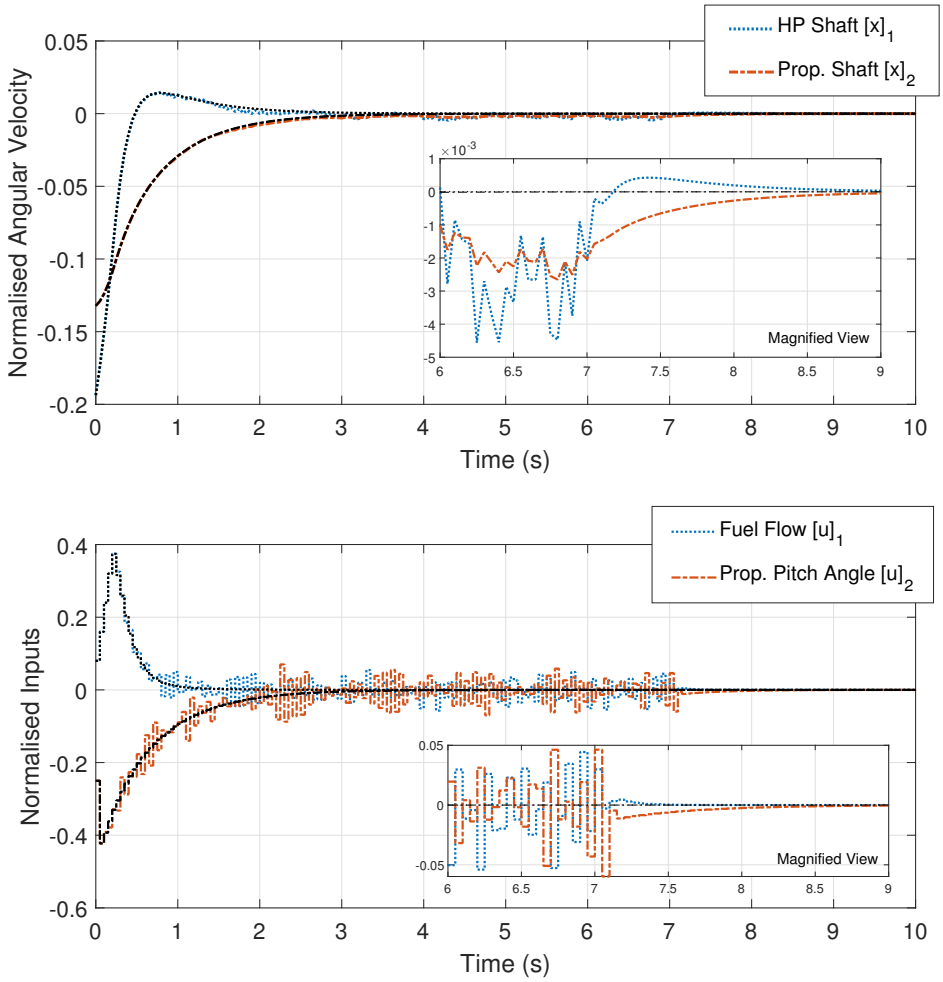

Figure 5.3: State and input trajectories under the PL-MPC (coloured) and nominal LQ-MPC (black). The magnified views demonstrate the trajectories near the switching event, where the closed-loop is switched to the updated LQ-MPC controller.

Figure 5.3 demonstrates the regulation of the SPT5 engine from the chosen initial condition to the origin *i.e.*, cruise. The trajectories illustrate the engine's acceleration response, caused by an initial rise in fuel flow and simultaneous reduction in propeller pitch. A small overshoot in HP shaft speed is observed. Figure 5.3 also depicts how the PL-MPC state and input trajectories deviate from the nominal LQ-MPC trajectories. Initially, the PL-MPC and nominal LQ-MPC trajectories coincide. As simulation time progresses, the excitations become more prominent, especially near the origin. The excitation optimisation is turned off once the model estimate convergence metric is satisfied, yielding the updated LQ-MPC trajectories as shown in the magnified views.

Input rates of change for under PL-MPC and nominal LQ-MPC

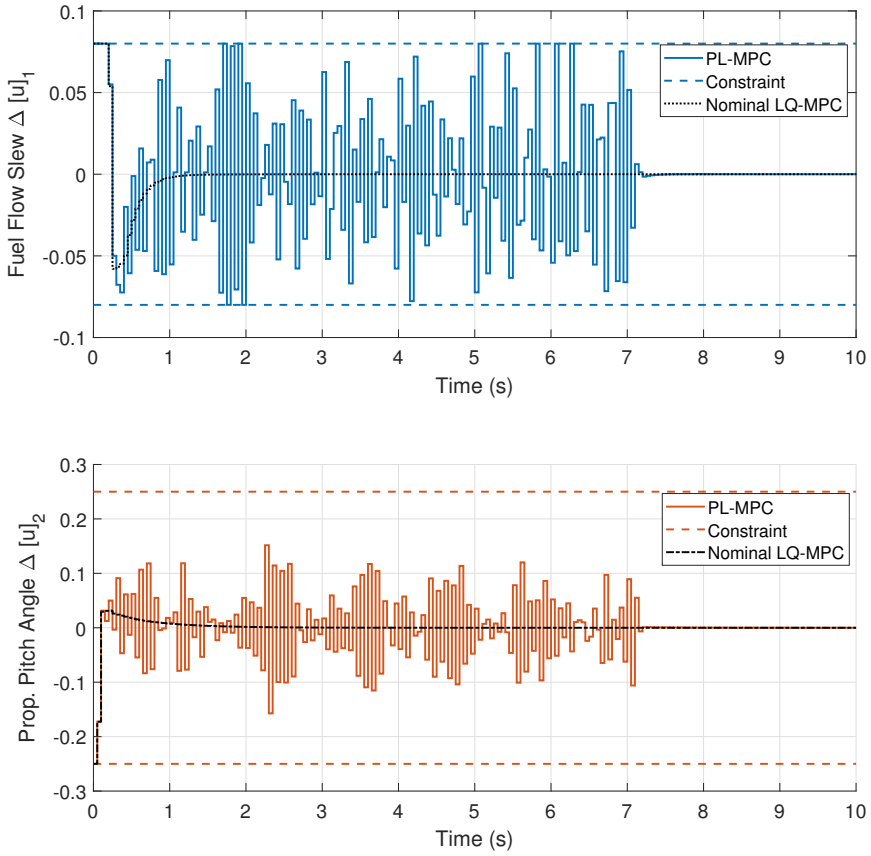

Figure 5.4: Input rate of change trajectories for the PL-MPC (coloured) and nominal LQ-MPC (black) controllers.

Figure 5.4 depicts the slew rates of fuel flow and propeller pitch. As in 5.3, the PL-MPC and the nominal LQ-MPC input slew trajectories coincide during the initial part of the transient behaviour. Both the fuel flow and the pitch angle rates of change satisfy the constraints specified in Table 5.3. Note that the deviations from the nominal input rates of change are caused by a deterministic perturbation, computed by the excitation optimisation that is solved online.

Parameter estimates under PL-MPC and nominal LQ-MPC

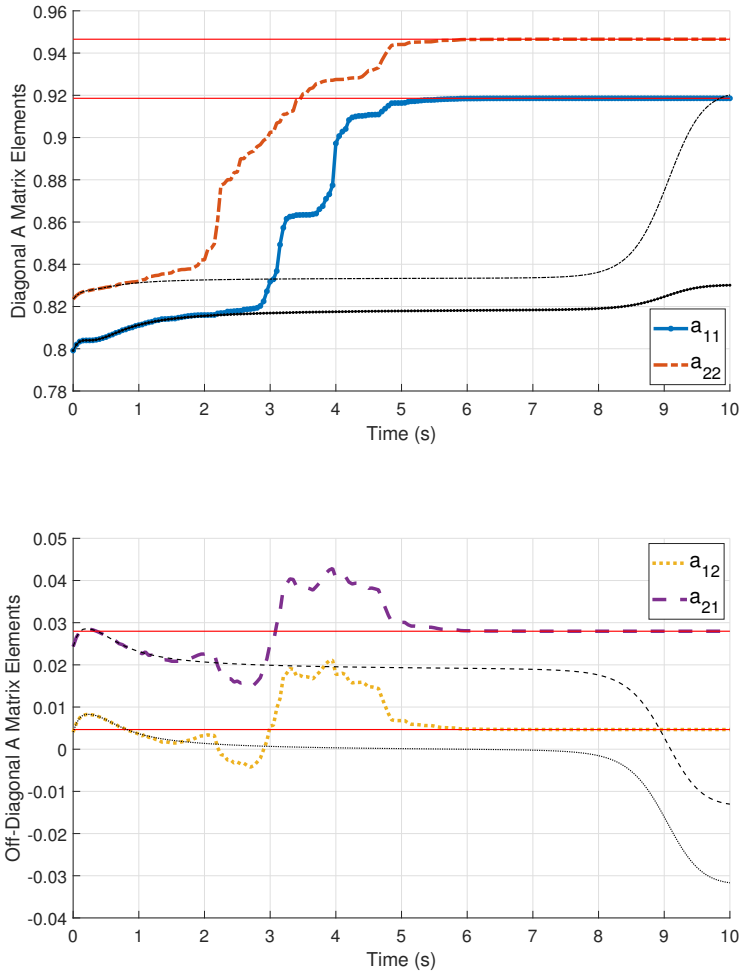

Figure 5.5: Estimated A-matrix parameters using RLS. The solid red lines denote the true elements. The lines described by the legends correspond to trajectories under the PL-MPC whereas the faint and diverging lines denote the trajectories under nominal LQ-MPC.

Figure 5.5 depicts the trajectories of the estimated state-transition matrix elements under two feedback controllers; the PL-MPC and the nominal LQ-MPC. The estimates for the PL-MPC case can be seen to converge to the respective elements of the engine's true open-loop dynamics. Convergence to the true state-transition matrix is obtained just above 7 seconds. On the other hand, after around 10 seconds, the estimate of the state-transition matrix under nominal LQ-MPC converges to an arbitrary value, that does not correspond to the true dynamics.

Minimum Eigenvalue Trajectory under PL-MPC and nominal LQ-MPC

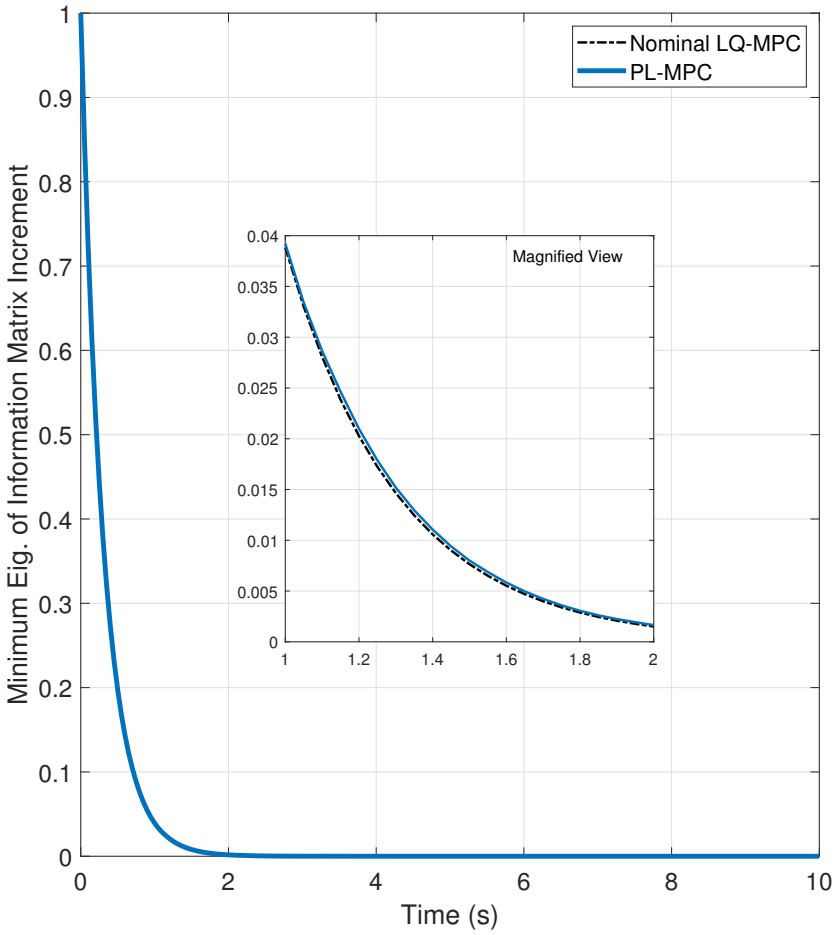

Figure 5.6: Trajectory of the information matrix $\mathcal{R}$ minimum eigenvalue. The magnified view highlights the PL-MPC affect on the minimum eigenvalue.

Figure 5.6 shows the minimum eigenvalue of the information matrix increment, plotted over the duration of the simulation for two feedback cases. Firstly,

the nominal LQ-MPC controller and secondly, the nominal PL-MPC. The magnified view reveals that the PL-MPC closed-loop trajectories facilitate the increase of the minimum value. For the LQ-MPC case, the eigenvalues monotonically decay to zero. On the other hand, for the PL-MPC case, the minimum eigenvalue decays to a neighbourhood of zero, until perturbations are turned off; once the state-transition matrix estimate satisfies the convergence criterion.

Cost function plots

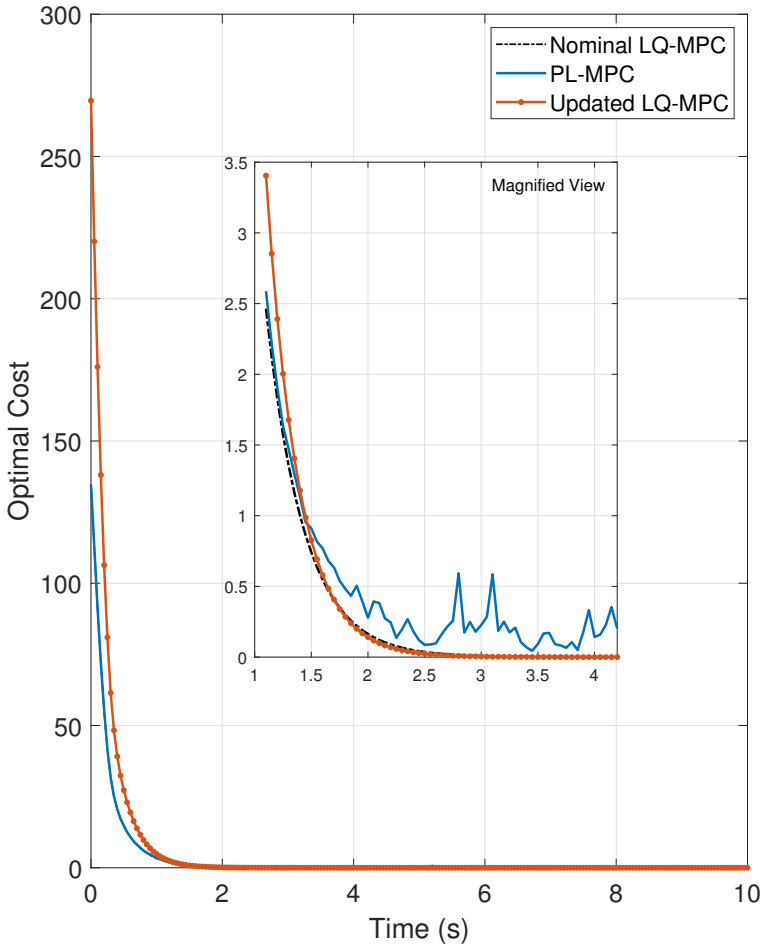

Figure 5.7: Cost function trajectories. The cost for each controller is computed using the quadratic formulation from (3.26).

Figure 5.7 illustrates the LQ cost of the closed-loop under three separate feedback control scenarios; the nominal LQ-MPC controller, the PL-MPC controller

under learning and finally, the updated LQ-MPC controller. The optimal cost for each controller scenario is computed using (3.26). For both instances of the LQ-MPC, the cost is monotonically decreasing. Due to model mismatch, there is a significant difference between the closed-loop costs for the nominal and updated LQ-MPC scenarios. For the PL-MPC scenario, the magnified view shows that the optimal cost is not monotonically decreasing due to the additional excitation inputs.

Comparison of nominal versus updated LQ-MPC

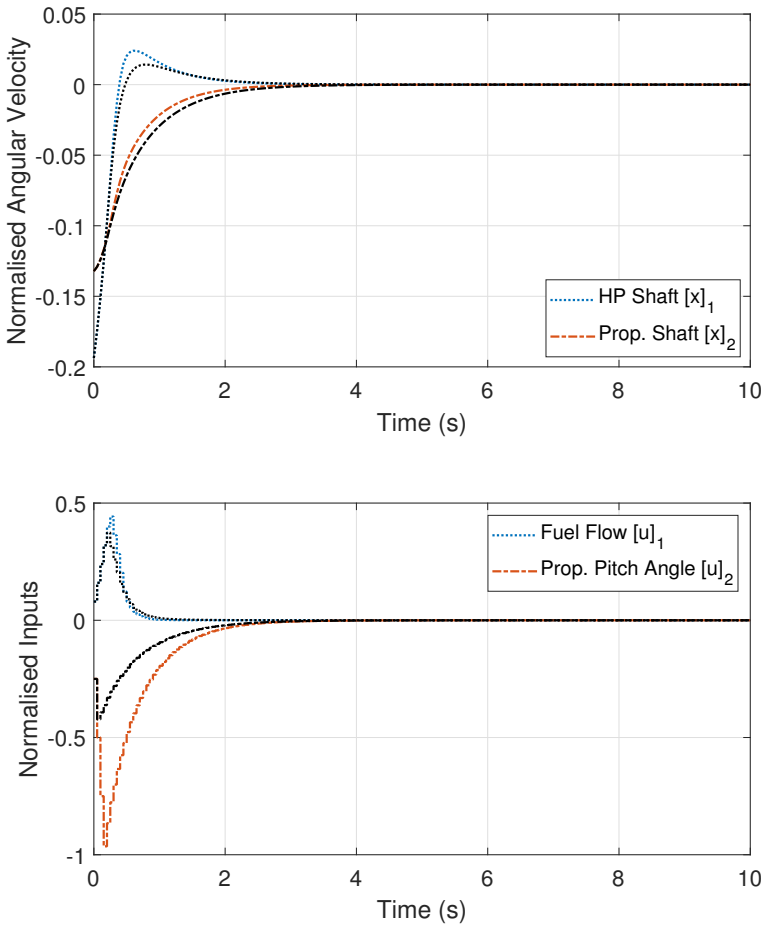

Figure 5.8: State and input trajectories under nominal and updated LQ-MPC controllers. The black lines, with matching legend line-styles, denote the *nominal* LQ-MPC trajectories whereas coloured lines denote updated LQ-MPC.

Figure 5.8 demonstrates the trajectory differences between a nominal and up-

dated LQ-MPC implementation that starts from the same initial condition as the regulation problem stated in Section 5.6.2. Note that only the prediction models and the stabilising terminal ingredients of Table 5.1 are updated; the remaining tuning constants, such as the Q and R, are kept constant between the adaptation. Consequently, a more aggressive closed-loop response is observed under the updated LQ-MPC control law. Significantly larger propeller pitch actuation effort is observed, whilst the fuel flow control input remains roughly the same. The state settling time is reduced with a small increase on the HP shaft speed overshoot.

Tuning Parameters

Table 5.5: Controller parameters

Parameter	Value	Parameter	Value	Parameter	Value
α_x, α_u	$0.5, 0.5$	γ	0.85	c_1	800
R	$\left[\begin{smallmatrix} 20 & 0 \\ 0 & 1 \end{smallmatrix}\right]$	Ξ	0.87	c_2	2426.6
Q	$\left[\begin{smallmatrix} 800 & 0 \\ 0 & 1200 \end{smallmatrix}\right]$	N	10	ζ	0.67
P	$\left[\begin{smallmatrix} 1465 & -486 \\ -486 & 3325 \end{smallmatrix}\right]$	$\mathcal{R}^-$	$\left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]$	L	46.5

Table 5.5 contains the tuning parameters used for the proposed PL-MPC controller. Some of the tuning parameters have been obtained by performing numerical simulations; these include b, c_2, ζ and L which have been obtained using the methods described in Section 3.8.

Optimisation Solution Times

Table 5.6: Simulation Computation Times

Optimisation	Minimum Time (sec)	Maximum Time (sec)	Average Time (sec)
$\mathbb{P}_1(x)$	0.1062	1.0572	0.1156
$\mathbb{P}_1(x; \mathbf{v})$	0.0719	0.2294	0.0901
$\mathbb{P}_2(x; \mathbf{u})$	0.1415	1.0837	0.1828

Table 5.6 denotes the solution time of the regulation and excitation optimsations under the PL-MPC controller designed using the tuning parameters of Table 5.5. The simulation times presented in Table 5.6 have been obtained by running MATLAB code on an average desktop computer running Windows 10 with a first generation i7-3GHz processor and 6GB of RAM.

Note that the times in Table 5.6 are extracted from running MATLAB code that uses generic optimisation solvers; the results do not reflect the smallest pos-

sible attainable time on an embedded controller. Interestingly, when solving the regulation optimisation $\mathbb{P}_1(x)$ without accounting for preview information or terminal constraint set tightening, all three timing metrics are greater than when accounting for preview information. More investigations are required to make any claims about reducing computational effort when using preview information. One possible conjecture is that accurate preview information can mitigate predicted trajectories from activating constraints, since run times were observed to be significantly longer when constraints were active.

The timing results only consider the time spent solving the respective optimisation problems and do not consider the time taken to compute the off-line MPC ingredients. In the worst case, the prediction, constraint and cost matrices only have to be re-computed once the decision is made to update to the new prediction model. The point of the table is to show the relative complexity between the regulation and excitation optimisation problems. As expected, the semi-definite program solution of the excitation problem takes the longest, on all three timing metrics.

5.7 Discussion

Despite the error in the dynamics matrix, the controller is able to simultaneously regulate and excite the system to the cruise equilibrium point whilst satisfying state, input and input slew rate constraints. Moreover, the LQ-MPC controller feasibly updates to the true plant model, whilst demonstrating a performance improvement through faster regulation of the states to the origin when compared to the original MPC design. The following subsections discuss the salient observations and limitations distilled from the simulation study.

Validation of the Robust Stability Inequality 4.27

Having synthesized the control law for the JetCat SPT5, the robust stability inequality (4.27) is verified. Firstly, the robust stability bound is recalled:

$$||\delta A|| \leq \frac{1}{\sqrt{b/c_1}}\left(\frac{(\rho-\zeta)b}{L} - ||\tilde{B}||\epsilon\right). \tag{5.7}$$

Substituting the values from the PL-MPC tuning constants and the numerically obtained parameters stated in Table 5.5 yields:

$$0.124 \leq \frac{1}{\sqrt{55.08/800}}\left(\frac{(1-0.6703)\times 55.08}{46.5409} - 0.0596\times 0.1\right) = 1.46 \quad \text{(to 3 s.f.)}$$

implying that the closed-loop system under PL-MPC, $x^+ = Ax + B(\bar{\kappa}_N(x;\mathbf{v}) + \pi_N(x;\mathbf{u}))$, is robustly stable within the sub-level parametrised by $b = 55.08$; despite the assumed prediction model mismatch and the additional input excitations. Notice that $\rho = 1$ since it is only required that the closed-loop value function of the uncertain system is monotonically decreasing.

Clearly, for this numerical example, the robustness margin is relatively large; the norm of the model mismatch could be increased by an order of magnitude and still be satisfied. However, it is important to note that this inequality is only useful if all the relevant assumptions hold; most importantly, Assumption 13. This assumption requires that the nominal LQ-MPC is stabilising for the true plant. Indeed, with such a large margin in this example, it is possible to find examples where a severely unstable prediction can be chosen such that $||\delta A||$ satisfies (5.7), but the resultant $\bar{\kappa}_N(x;\mathbf{v})$ is not stabilising. This highlights that the inequality (5.7) is not a sufficient condition for robust stability guarantees and requires careful analysis of all assumptions under specific numerical examples.

Similarity to standard LQ-MPC

Notice that when the simulation is begun, the PL-MPC trajectories behave similarly to those of the standard LQ-MPC solution. This can be explained by the fact that when the problem is initialised closer to constraint boundaries, the excitation problem $\mathbb{P}_2$ will favour smaller deviations from the tail-end of the excitation sequence to satisfy constraints; recall that by construction, the tail-end of the excitation sequence should be feasible for $\mathbb{P}_2$ (even with modelling uncertainty, provided that $\mathbb{P}_1$ is feasible). As the trajectories move deeper into the interior of $\mathcal{X}_N(\mathbf{v})$, the excitations become more pronounced; $\mathbb{P}_2$ is able to choose perturbations which would have previously been infeasible because of proximity to the constraints.

Convergent closed-loop identification

The online RLS identification scheme using closed-loop data under PL-MPC results in estimation of the true plant dynamics around CRZ. On the other hand, RLS under nominal LQ-MPC closed-loop data (with no excitation) does not converge to the true dynamics. Moreover, Figure 5.6 shows that the minimum eigenvalue of the current information matrix is increased in the PL-MPC compared to the nominal LQ-MPC. More importantly, the PL-MPC controller promotes excitations that satisfy the original constraints of the problem. On the other hand, the classical approach of conducting an equivalent closed-loop identification exercise would superimpose dither signals over the regulator input. In such an approach, ad-hoc

constraint handling methods would be required, leading to further suboptimality with respect to the regulation objective.

Computational complexity

As discussed in the previous chapters, the regulation and excitation optimisations are both convex. Solution time data has been gathered to give insight into the computational complexity of the PL-MPC algorithm. Note that these simulation times correspond to code executions in MATLAB; using YALMIP to interface with MATLAB's quadprog solver (to solve $\mathbb{P}_1$) and SeDuMi for semi-definite program solver (to solve $\mathbb{P}_2$) [80]. The code has not been optimised for solution time and is expected to be further reduced if implemented in hardware-specific programming languages, such as C/C++.

Excitation problem redundancy

The attractive feature of the PL-MPC algorithm is the redundancy of the excitation problem. Since its feasibility depends on the feasibility of the regulation problem, it can be turned on/off whenever the closed-loop is running under the LQ-MPC regulator. Moreover, the proposed re-linearisation method of Section 3.5.2 can be used to improve the solution of the excitation problem, if there is remaining time left before the next sampling instant.

Parameter Bursting

Adaptive control schemes can suffer from the phenemenon known as bursting [7]. This can occur when the regressor is not persistently exciting, leading to abrupt changes in the parameter estimates. It is a complex phenomenon which has been numerically explored in [7] on a second-order non-linear system; the phenomenon results from the interaction between the parameter dependent dynamics an estimator that uses data which is only PE over a finite interval. Since the LMPC scheme does not guarantee PE at all times, but rather promotes PE when the system constraints allow it, issues associated with bursting are avoided by not updating the prediction model with the estimates at arbitrary time-steps. Moreover, the RLS algorithm is turned off after satisfactory convergence is achieved.

Concluding Remarks

Updating the ingredients leads to a LQ-MPC design, described in Section 2.1.4 and [107], but with *the prediction model matching the true dynamics*. Thus, under lack

of any other exogenous uncertainty, recursive feasibility is recovered. Moreover, if exogenous disturbances are present, then the inherent robustness property of this updated LQ-MPC controller is fully dedicated to dealing with the exogenous uncertainty; rather than model mismatch.

5.7.1 Limitations of the Applied Algorithm

It is important to illuminate the limitations of this numerical study of PL-MPC. These limitations stem from the assumptions that have been made in order to make the theoretical arguments of Chapters 2-4.

- **Modelling and MPC Assumptions**: 2, 3, 4 and 5.

  These modelling assumptions simplify the control problem to that of regulating a linear system with fixed dimensions and no exogenous noise; actual practise is to use an output tracking formulation [87]. The only uncertainty comes from the the plant-model mismatch and thus, the controller's robustness is exercised to handle this single type of uncertainty. If exogenous noise were to be considered, it can still be included within derived robustness bounds under appropriate boundedness assumption.

  Without this linear framework, inherent robustness using the continuity property of the value function cannot be guaranteed [106]. The choice of using a quadratic cost function formulation of MPC is fundamental for showing the inherent robustness results; a 1-norm cost ($||.||_1$) would not provide the same a-priori stability guarantees.

  Whilst availability of the state vector is a standard assumption in academic literature, in reality, this assumption requires high accuracy of the on-board sensors or a robust state estimation design.

- **RLS Assumption**: 6.

  This assumption is necessary for the PE property to hold. However, because the RLS estimation algorithm utilises a forgetting factor $\gamma \in [0, 1)$, even an arbitrarily chosen initialisation of the information matrix will have an insignificant affect on the estimates after a long time period, and is therefore not considered an issue; at worst, slower convergence to the true parameters may be observed.

- **Preview Information Assumptions**: 7, 8, 9, 10, 11 and 12.

  The assumptions for including preview information rely on the nominal linear formulation of MPC [106] which has been well-established from the

discussion of the required modification in 4.2.1. The key limitation is the need to use a perfect $\bar{B}$ matrix; Assumption 7 implicitly requires that the B matrix is known exactly, such that the current perturbation is known, as required in the results of [13]. Evaluation of the inherent robustness bound in Corollary 1 and equation (4.21) also relies on perfect knowledge of the B matrix.

- **Knowledge of Initial and Target States**: 16

  For the specified regulation problem, this assumption is required to avoid steady-state offset. In reality, thrust control requires accurate mapping between the states and desired thrust level, since it is difficult to obtain a direct measurement of thrust whilst in-flight; shaft speeds are used as proxies for thrust [109]. Observer-based tracking formulations that introduce integrators can be used to relax this assumption [87].

- **Detecting True Parameter Convergence**

  In reality, the true system model is unknown and therefore the model error norm cannot be evaluated. This complicates the determination of convergence to the true dynamics. If a naive threshold on the derivative of each parameter estimate is used, then the controller may update prematurely without converging to the true dynamics. Under PL-MPC, PE conditions are promoted rather than guaranteed (as a consequence of Proposition 5), hence, some portions of the collected data may not have been PE. More sophisticated convergence detection methods must be employed in practise, to ensure that the prediction model is updated to a more accurate prediction model.

- **Online computation**: 14

  The computation of the stabilising ingredients online poses a challenge for an embedded control implementation. However, with the advent of cloud computing and the push for greater connectedness in gas turbine monitoring, the re-computations may not necessarily have to be performed on the engine itself.

- **Uncertainty Matching**: 13, 15.

  The assumptions that enable guaranteed feasible updates of the stabilising ingredients depend on the model uncertainty. This limits the degree of mismatch that can be tolerated. This assumption will always need extensive validation to check that the nominal and updated control laws regulate the states of the system to the origin.

- **Feasible Switch Check**

  To check if the switch to the updated controller is feasible, the controllability set recursion (4.2) is used to construct the domain of attraction for the updated LQ-MPC controller. The controllability sets under linear dynamics and polytopic constraints are convex. However, in practice, these convex sets can become difficult to compute, particularly for large horizons and system dimensionality, since numerous recursions of (4.2) can lead to polytopes with arbitrarily large number of facets.

Chapter 6

Conclusion and Future Work

Advanced control algorithms that utilise novel actuation have been identified as enablers for unprecedented engine capability. Apart from reducing specific fuel consumption, emissions and maintenance costs, novel control laws can increase the robustness and safety margins by adapting to engine-specific health states. To enable such capability, introducing further autonomy to the engine control laws is required; to avoid the costs of manual re-tuning. Therefore, automating the learning of engine-specific parameters has been identified as a potential route to alleviate over-design and conservatism, which is endemic in gas turbine control. This thesis derives a candidate control solution: the proposed preview learning MPC control law. The following sections conclude the observations, contributions and proposed improvements for future work.

6.1 Contributions

The chapter-wise contributions are stated in the following sub-sections.

Chapter 1 - Motivation and Background

Chapter 1 investigates and surveys a wide-range of control strategies applied to the gas turbine control problem. The survey distils features of the gas turbine control problem that translate to desirable features of a potential candidate control algorithm. The conclusions of this survey are used to support the choice of studying adaptive model predictive control.

Chapter 2 - Adaptive Model Predictive Control

Chapter 2 focuses the literature review on adaptive model predictive control, including frameworks that enforce PE conditions and those that do not. The nominal

LQ-MPC recipe is defined and investigated to consider the technical issues that arise under changes to the plant's prediction model. A summary of LQ-MPC's shortcomings is reported, motivating the need for the technical contributions of this thesis.

Chapter 3 - Learning Model Predictive Control

Chapter 3 derives stability results for the proposed LMPC algorithm. The LMPC algorithm is constructed to simultaneously regulate and excite the uncertain plant dynamics, subject to state and input constraints. Steering the plant's states to the origin whilst identifying its underlying open-loop dynamics, using RLS, is provisioned by the resultant "information rich" state and input trajectories. The inherent robustness property of LQ-MPC is employed to characterize robust stability of the closed-loop, in the form of a novel inequality. This inequality is expressed in terms of plant-model mismatch and excitation magnitude; it enables the selection of an appropriate magnitude of excitation such that robust stability is preserved. This inequality provides a tuning guideline that alleviates the trial-and-error selection of the excitation magnitude, which is frequently observed in prior-art. Finally, the theory is demonstrated with a numerical example.

Chapter 4 - Preview Information and Inherent Robustness

Chapter 4 describes the extension of the LMPC framework, by utilising existing results for including preview disturbance information within an MPC controller's prediction horizon. By specializing prior-art to the case of *known input excitation sequences*, several contributions are made as part of the proposed PL-MPC.

Firstly, in the uncertainty-free scenario, recursive feasibility is guaranteed under assumptions on the evolution of the preview information and appropriate translation of the input constraints. This is a consequence of controllability set nestedness, which is obtained through the proposed optimisation formulation. This nestedness property does not necessarily hold if a generic disturbance sequence is considered, as shown in prior-art. Hence, Chapter 4 strengthens the preview information framework when it is specialized to input disturbances.

Secondly, robust stability results are derived for the PL-MPC algorithm. The tuning rule of PL-MPC, analogous to the inequality derived in Chapter 3, is shown to depend on the rate-of-change of the excitation sequence between time-steps; rather than the magnitude of the excitation as shown previously. This subtly changes the role of the excitation optimisation that is used to promote PE conditions. Given an a-priori designed excitation sequence (which can even be the null perturbation

sequence), the PL-MPC seeks to promote PE condition of the trajectories, whilst handling constraints. To the author's best knowledge, there are no other methods that robustify against an internally generated excitation without inducing conservatism using known robust MPC methods.

Finally, the region of attraction under a PL-MPC controller is numerically shown to be significantly larger than in the nominal LQ-MPC case, confirming the proposition of prior art.

Chapter 5 - Adaptation in MPC under GT Degradation

Chapter 5 proposes conditions under which inherently robust LQ-MPC algorithms can perform stable and feasible adaptation. These conditions are identified by considering the LQ-MPC adaptation in the switched systems framework [76]. A set of assumptions on the permissible degradation is defined, to allow the LTI results of the previous chapters to hold for slowly time-varying systems. A JetCat SPT5 gas turbine control model is used to demonstrate the adaptive PL-MPC algorithm's efficacy on a validated linear gas turbine engine model. Input rate-of-change constraints are introduced to highlight the controller's ability to respect constraints, despite the additional input excitations required for learning the true dynamics. The numerical simulations demonstrate the controller's ability to perform regulation, identification and safe switching to an LQ-MPC control law that is matched to the true dynamics of the gas turbine.

6.2 Summary of the Proposed Algorithm

The proposed preview learning model predictive controller has known stability guarantees and inherent robustness margins to limited modelling uncertainty. The theoretical extension enables an initial MPC control law to learn an updated prediction model that corresponds to the true open-loop dynamics of the plant. Of course, no model is a perfect representation of reality; however, under the mild technical assumptions of this thesis, the PL-MPC is able to learn updated (*i.e.*, more accurate) plant parameters, without over-complicating off-line design effort. Since the controller automatically designs online and "optimal" closed-loop identification experiments, that also satisfy constraints, the classical recursive least squares estimator can converge to the true plant parameters; provided that the estimation data is sufficiently exciting. With a more accurate estimate of the true open-loop plant dynamics, the LQ-MPC prediction model and its associated stabilising ingredients are updated. By updating the LQ-MPC with the aforementioned stabilising

ingredients, the nominal MPC controller regains a portion of its robustness margin; since disturbances caused by model mismatch are eliminated. The inherent robustness of LQ-MPC feedback can therefore be fully committed to dealing with exogenous uncertainty, that exist in reality.

6.2.1 Desirable Features

The choice of using MPC as the control framework was justified by the following observations from gas turbine control literature:

- Explicit handling of constraints.

  Complicated gain-scheduled limiting-logic and anti-windup schemes are made redundant.

- Online reconfigurability for adaption to health condition.

  The implicit form of MPC solves an optimisation online that can be updated with online identified parameters when they become available.

- Optimal control policies.

  The closed-loop trajectories minimise a quadratic cost metric; the optimisation can be intuitively tuned to reduce specific fuel consumption through appropriate selection of weighting matrices. The excitation optimisation can be turned off once a better model is identified.

- Multivariable control strategy.

  The cross-coupling between states and inputs can be explicitly utilised for control.

6.2.2 Scope of Application

The scope of this research applies to fixed dimensionality LTI systems and slowly time-varying systems that satisfy continuity and dwell-time arguments. Polyhedral state and polytopic input constraints, with full state measurements available and no exogenous noise conditions are required. Whilst the proposed controller solves convex optimisations, the applicability is limited to low dimensional systems due to the computational difficulties when computing stabilising ingredients online. Moreover, obtaining the value function properties for high dimensional can become numerically burdensome and conservative if a-priori guarantees are to be verified.

6.3 Future Work

Several interesting avenues for extending the results of this thesis exist.

6.3.1 Linear Time-varying Plant

Future work should investigate the necessary modifications of the PL-MPC algorithm to cope with quickly time-varying true plant dynamics, that vary within the settling time of the closed-loop *e.g.*, a sudden actuator failure. Currently, the quasi-LTI assumption only allows to consider slowly varying dynamics that are orders of magnitude slower than the control transients. It would be interesting to see the results of this thesis united with the nominal time-varying MPC literature [106], to consider plant-model mismatch within the control horizon.

6.3.2 Linear Parameter-varying Plant

Since a single linear model cannot accurately represent the operation of a gas turbine over its entire envelope, a study of how PL-MPC can be systematically introduced into an LPV framework could yield interesting results. After all, switching between controllers can also provide excitation for identifying open-loop dynamics [78].

6.3.3 Non-linear Systems

Instead of focusing on linear approximations, non-linear systems can alleviate the problems of model mismatch. However, extending the results to non-linear systems may prove difficult since continuity, which is the basic property that the results of this thesis rely on, can be lost. Therefore, identifying classes of non-linear systems where continuity is not lost may be a useful exercise to confirm whether an equivalent non-linear learning MPC exists. Relatively easy NARMAX identification in an MPC framework may yield interesting results [15].

6.3.4 Reducing Excitation Suboptimality

The convex approximation of the excitation problem results in sub-optimality in the excitation signals, which further exacerbates the sub-optimality of the regulation objective. An investigation into alternative solution methods for solving the excitation problem, to provide better quality trajectories with respect to *both* conflicting regulation and excitation objectives, could reduce convergence time and reduce unnecessary actuation. Results in the field of multi-objective optimisation,

such as lexicographic programming, provide interesting approaches to reformulate the conflicting objectives in dual-control [1].

6.3.5 Alternative Identification Schemes

As for the identification scheme, RLS was chosen due to its simplicity and explicit relation between the information matrix and the parameter update law. Other recursive algorithms, such as least mean squares proposed in set-membership problem setting of [82], provide an alternative for robust adaptive MPC; but the resultant controllers, that utilise a set of possible prediction models, suffer from conservatism. Integration of the ideas from PL-MPC and set-membership could allow for the treatment of time-varying plant dynamics, whilst reducing conservatism. The issues of having an ill-posed cost function when excitations are present [84], could be avoided through equivalent translations that are proposed in section 4.3.1.

6.3.6 Degradation Modelling and Output Feedback

The engine degradation is assumed to enter through the state-transition matrix. Future work could consider alternative methods for modelling engine degradation; as has been done in [109], which considers an unknown health state vector that acts like an unknown exogenous disturbance. In such instances, tracking filters are being explored [19] to estimate the health state vector. However, careful consideration is required since introduction of a state estimator may eliminate the derived robustness margins; it it is well-known that LQG that not possess any, if the noise covariance is sufficiently large. Hence, for "large" noise models with unbounded supports (e.g. Gaussian white noise with large covariance), stochastic MPC methods may be more appropriate [58]. Dual-control in the stochastic framework is receiving recent attention with the works of [56].

6.3.7 Hardware-in-the-loop Validation

From a practical standpoint, the PL-MPC algorithm requires further validation with real hardware. Whilst the proposed theoretical algorithm solves convex optimisation problems, a real-time embedded implementation that copes with the realities of state-estimation error, sensor noise, exogenous disturbances and even delay and packet-loss, would be most valuable for demonstrating higher technology readiness levels.

Appendix

6.4 Matrix Inversion Lemma

The recursive definition of the information matrix (3.7b) is well suited for use with the matrix inversion lemma. In the sequel, it is shown that the problem of finding the inverse of the information matrix $\mathcal{R}^{-1}$ is simplified to the inversion of a scalar; given that the preceding inversion $(\mathcal{R}^{-1})^{-}$ is available. Therefore, the matrix inversion lemma significantly reduces the computational complexity of the RLS algorithm. The the matrix inversion lemma identity is verified in the sequel.

Lemma 5 (Matrix Inversion Lemma)**.** *Let A, C and $C^{-1} + DA^{-1}B$ be non-singular square matrices. Then the following equation holds*

$$\big(A + BCD\big)^{-1} = A^{-1} - A^{-1}B\big(C^{-1} + DA^{-1}B\big)^{-1}DA^{-1}.$$

Proof. The proof shows that the equation holds by multiplying both sides with $(A + BCD)$. This yields:

$$\begin{aligned} I &= I + BCDA^{-1} - B\big(C^{-1} + DA^{-1}B\big)^{-1}DA^{-1} - BCDA^{-1}B\big(C^{-1} + DA^{-1}B\big)^{-1}DA^{-1} \\ &= I + BCDA^{-1} - B(C^{-1} + DA^{-1}B\big)^{-1}\big(I + CDA^{-1}B\big)DA^{-1} \\ &= I + BCDA^{-1} - BCDA^{-1} \\ &= I \end{aligned}$$

□

For the RLS formulation, the lemma holds under the definitions of the matrices

Table 6.1: Matrix Inversion Lemma and Equivalent RLS Definitions

A	B	C	D
$\gamma\mathcal{R}^{-}$	ψ	I	$\psi^{\top}$

6.5 Singleton Set Manipulation in the Minkowski Sum

Let $\mathcal{A}, \mathcal{B} \subset \mathbb{R}^n$ be two sets in a common vector space. Consider some vector $b^* \in \mathbb{R}^n$ that is added:

$$\begin{aligned} \mathcal{A} + c &= \{a + b^* : a \in \mathcal{A}\}, \\ &= b^* + \{a : a \in \mathcal{A}\}. \end{aligned}$$

Now, recall the definition of the Minkowski sum

$$A \oplus B \triangleq \{a + b : a \in \mathcal{A}, b \in \mathcal{B}\}.$$

If the set $\mathcal{B}$ is constructed such that it is the singleton set $\mathcal{B} = \{b^*\}$, then the Minkowski sum is:

$$\begin{aligned} \mathcal{A} \oplus \mathcal{B} &= \{a + b : a \in \mathcal{A}, b \in \mathcal{B}\}, \\ &= \{a + b^* : a \in \mathcal{A}\}, \\ &= b^* + \{a : a \in \mathcal{A}\}. \end{aligned}$$

The Minkowski sum is equivalent to standard addition when one of the summands is a singleton set. With this fact, the right hand side term of equation 4.13 simplifies as follows:

$$-B(\mathbb{U} - v(N-i)) \oplus -\{Bv(N-i)\} = -B\mathbb{U}.$$

6.6 Derivation of the Controllability Sets

The controllability sets of a generic discrete-time non-linear system $x^+ = f(x, u)$, with constraint sets $x \in \mathbb{X}$ and $u \in \mathbb{U}$, are given by the recursive definition

$$\mathcal{X}_{i+1} \triangleq \{x \in \mathbb{X} : \exists u \in \mathbb{U}, f(x, u) \in \mathcal{X}_i\}.$$

The set $\mathcal{X}_{i+1}$ contains the states that can be steered, in one-step, to the controllability set at the i-th step, under the system dynamics $x^+ = f(x, u)$, whilst respecting state and input constraints. These sets are of importance in MPC because the recursion can be used to prove recursive feasibility. When the recursion is initialised with $\mathcal{X}_0 = \mathbb{X}_f$, the positively invariant terminal set for the system $x^+ = f(x, u)$,

the MPC controller's region of attraction can be explicitly obtained. However, for a generic non-linear system, these sets may be prohibitively difficult to compute because of the possible arbitrary complexity of the dynamics and the constraints. By specialising the recursion to the linear system with polytopic constraints, studied in this thesis, the recursion is simplified to

$$\begin{aligned}\mathcal{X}_{i+1} &= \{x \in \mathbb{X} : \exists u \in \mathbb{U}, Ax + Bu \in \mathcal{X}_i\} \\ &= \{x \in \mathbb{X} : \exists u \in \mathbb{U}, Ax \in \mathcal{X}_i \oplus \{-Bu\}\} \\ &= \{x \in \mathbb{X} : Ax \in \mathcal{X}_i \oplus \{-B\mathbb{U}\}\} \\ &= \{x \in \mathbb{X} : x \in A^{-1}(\mathcal{X}_i \oplus -B\mathbb{U})\} \\ &= \mathbb{X} \cap A^{-1}(\mathcal{X}_i \oplus -B\mathbb{U}).\end{aligned}$$

6.7 JetCat SPT5 Tuning Constants

6.7.1 Maximal Level Set of the Nominal Value Function

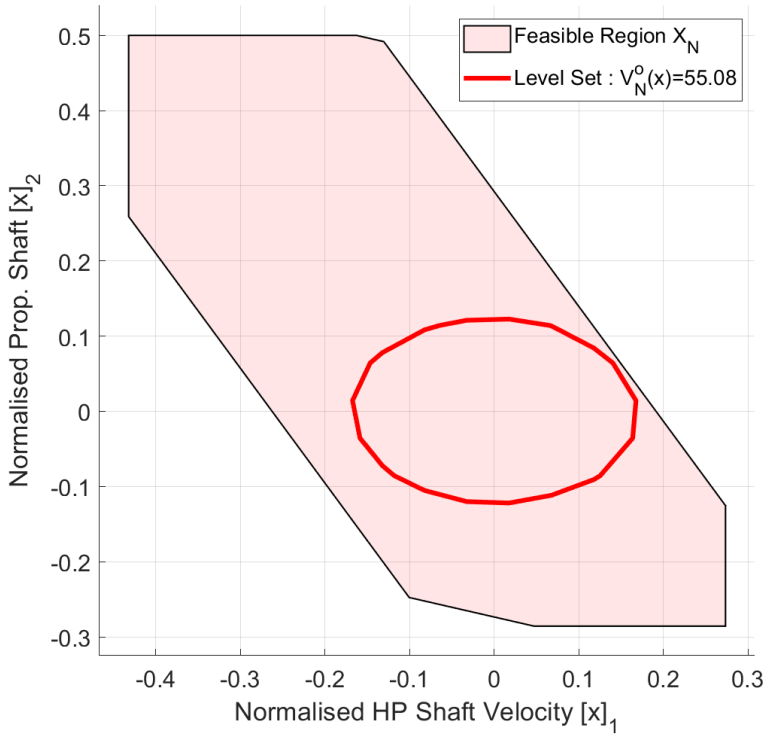

Figure 6.1: Numerical approximation of the largest level set of the value function $V_N^0(x)$ contained within the region of feasibility $\mathcal{X}_N$, for the nominal closed-loop system.

6.7.2 Nominal Explicit LQ-MPC

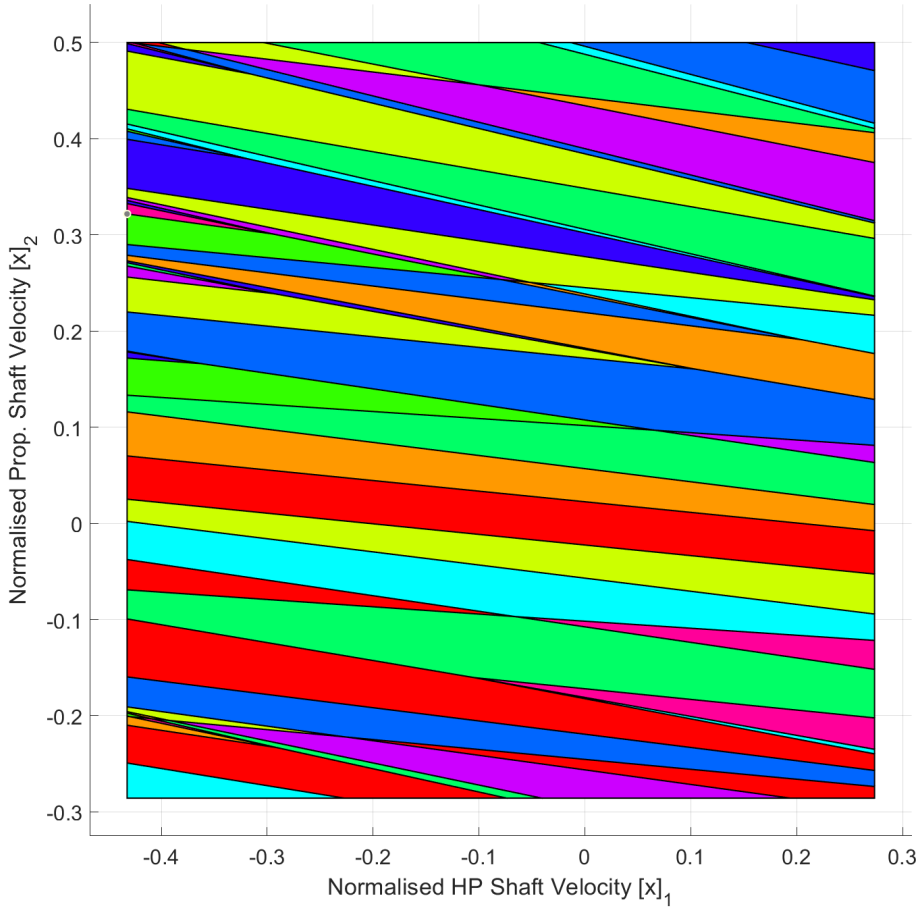

Figure 6.2: The nominal explicit LQ-MPC control law regions.

6.7.3 Tuning Method

1. Set $\rho = 1$.
2. Compute norm of Input matrix B
3. Choose a degradation constant Ξ E.g $\Xi = 0.87$.
4. Choose initial Q, R weighting matrices and control horizon N.
   (a) Set N to be as small as possible to be able to compute the explicit solution for finding the Lipschitz constant of the value function.

(b) Extract lower bound from min eig of Q.

5. Compute explicit LQ-MPC control law.

   (a) Obtain the Lipschitz constant L.

      i. Take the norm of each feedback matrix gain, select the largest norm.

6. Simulate closed-loop with nominal LQ-MPC controller, $\bar{\kappa}_N(x)$ for prediction model matching the plant as if the plant were to behave as predicted.

   - Obtain value function constants.
     - i. Max level set – from inspecting value contours.
     - ii. Decay rate – from fitting upper bounding exponential to closed-loop cost.

7. Compute Corollary inequality and check if satisfied. If not satisfied, reduce model error norm.

Bibliography

[1] B.A. Acikmese. Application of lexicographic goal programming with convex optimization in control systems. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, Reston, Virginia, 2013.

[2] J. Adamy and A. Flemming. Soft variable-structure controls: a survey. *Automatica*, 40(11):1821–1844, 2004.

[3] V. Adetola, D. DeHaan, and M. Guay. Adaptive model predictive control for constrained nonlinear systems. *Systems & Control Letters*, 58(5):320–326, 2009.

[4] J.C. Allwright. On maximizing the minimum eigenvalue of a linear combination of symmetric-matrices. *SIAM Journal on Matrix Analysis and Applications*, 10(3):347–382, 1989.

[5] T. Alpcan and I. Shames. An Information-Based Learning Approach to Dual Control. *IEEE Transactions on Neural Networks and Learning Systems*, 26(11): 2736–2748, 2015.

[6] K.J. Åström and R.M. Murray. *Feedback Systems An Introduction for Scientists and Engineers*. Number July. 2007.

[7] K.J. Åström and B. Wittenmark. *Adaptive Control*. Dover Books on Electrical Engineering. Dover Publications, 2 edition, 2008.

[8] K. Åström and T. Hägglund. The future of PID control. *Control Engineering Practice*, 9(11):1163–1175, 2001.

[9] A. Aswani, P. Bouffard, and C. Tomlin. Extensions of learning-based model predictive control for real-time application to a quadrotor helicopter. In *2012 American Control Conference (ACC)*, pages 4661–4666. 2012.

[10] A. Aswani, H. Gonzalez, S.S. Sastry, and C. Tomlin. Provably safe and robust learning-based model predictive control. *Automatica*, 49(5):1216–1226, 2013.

[11] N. Athanasopoulos, K. Smpoukis, and R.M. Jungers. Invariant sets analysis for constrained switching systems. *IEEE Control Systems Letters*, 1(2):256–261, 2017.

[12] E.W. Bai, H. Cho, and R. Tempo. Convergence Properties of the Membership Set. *Automatica*, 34(10):1245–1249, 1998.

[13] P.R. Baldivieso Monasterios and P.A. Trodden. Model predictive control of linear systems with preview information: feasibility, stability and inherent Robustness. *IEEE Transactions on Automatic Control*, 64(9):3831–3838, 2018.

[14] A.R. Behbahani, A. Von Moll, R. Zeller, and J. Ordo. Aircraft integration challenges and opportunities for distributed intelligent control, power, thermal management, and diagnostic and prognostic systems. 2014.

[15] S.A. Billings and L.A. Aguirre. Effects of the sampling time on the dynamics and identification of nonlinear models. *International Journal of Bifurcation and Chaos*, 05(06):1541–1556, 1995.

[16] R.R. Bitmead. Persistence of Excitation Conditions and the Convergence of Adaptive Schemes. *IEEE Transactions on Information Theory*, 30(2):183–191, 1984.

[17] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.

[18] F. Blanchini and S. Miani. *Set-Theoretic Methods in Control*. Birkhäuser Basel, 2008.

[19] S. Borguet and O. Loénard. Comparison of adaptive filters for gas turbine performance monitoring. *Journal of Computational and Applied Mathematics*, 234(7):2202–2212, 2010.

[20] S. Boyd and L. Vandenberghe. *Convex Optimization*. 2004.

[21] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994.

[22] L.J. Bridgeman, C. Danielson, and S. Di Cairano. Stability and feasibility of MPC for switched linear systems with dwell-time constraints. In *2016 American Control Conference (ACC)*, pages 2681–2686. 2016.

[23] B. Brunell, R. Bitmead, and A. Connolly. Nonlinear model predictive control of an aircraft gas turbine engine. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, volume 4, pages 4649–4651. 2002.

[24] S. Can Gülen. Beyond Brayton Cycle: It is time to change the paradigm. *Journal of Engineering for Gas Turbines and Power*, 140(11):111703, 2018.

[25] A. Chakrabarty, A. Raghunathan, S. Di Cairano, and C. Danielson. Data-Driven Estimation of Backward Reachable and Invariant Sets for Unmodeled Systems via Active Learning. In *2018 IEEE Conference on Decision and Control (CDC)*, number CDC, pages 372–377. 2018.

[26] Y. Cheng, S. Haghighat, and S. Di Cairano. Robust dual control MPC with application to soft-landing control. In *2015 American Control Conference (ACC)*, pages 3862–3867. 2015.

[27] A. Chipperfield and P. Fleming. Multiobjective gas turbine engine controller design using genetic algorithms. *IEEE Transactions on Industrial Electronics*, 43(5):583–587, 1996.

[28] G. Chowdhary, M. Mühlegg, and E. Johnson. Exponential parameter and tracking error convergence guarantees for adaptive controllers without persistency of excitation. *International Journal of Control*, 87(8):1583–1603, 2014.

[29] P. Colaneri and R. Scattolini. Robust model predictive control of discrete-time switched systems. *IFAC Proceedings Volumes*, 40(14):208–212, 2007.

[30] L. Coroianu. Best Lipschitz constants of solutions of quadratic programs. *Journal of Optimization Theory and Applications*, 170(3):853–875, 2016.

[31] D. Culley, S. Garg, S.J. Hiller, W. Horn, A. Kumar, H.K. Mathews, H. Moustapha, H. Pfoertner, T. Rosenfeld, and P. Rybarik. More Intelligent Gas Turbine Engines. Technical Report April, NATO, 2009.

[32] C. Danielson, L.J. Bridgeman, and S. Di Cairano. Necessary and sufficient conditions for constraint satisfaction in switched systems using switch-robust control invariant sets. *International Journal of Robust and Nonlinear Control*, 29(9):2589–2602, 2019.

[33] C. Davies, J.E. Holt, and I.A. Griffin. Benefits of inverse model control of Rolls-Royce civil gas turbines. Technical report, Rolls-Royce, University of Sheffield, 2006.

[34] W. Dellinger and A. Alouani. Application of linear multivariable control design for the GE-21 turbofan jet engine. In *[1989] Proceedings. The Twenty-First Southeastern Symposium on System Theory*, pages 167–171. 1989.

[35] E. DeSantis, M. DiBenedetto, and L. Berardi. Computation of Maximal Safe Sets for Switching Systems. *IEEE Transactions on Automatic Control*, 49(2): 184–195, 2004.

[36] S. Di Cairano. Model adjustable predictive control with stability guarantees. In *2015 American Control Conference (ACC)*, pages 226–231. 2015.

[37] S. Di Cairano. Indirect adaptive model predictive control for linear systems with polytopic uncertainty. In *2016 American Control Conference (ACC)*, pages 3570–3575. 2016.

[38] J. Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on Automatic Control*, 23(4):756–757, 1978.

[39] U. Eren, A. Prach, B.B. Koçer, S.V. Raković, E. Kayacan, and B. Açıkmeşe. Model predictive control in aerospace systems: current state and opportunities. *Journal of Guidance, Control, and Dynamics*, pages 1–25, 2017.

[40] H.V. Essen. *Modelling and model based control of turbomachinery*. PhD thesis, Technische Universiteit Eindhoven, 1998.

[41] A. Feldbaum. Dual Control Theory: Part 1. *Avtomat. i Telemekh*, 21(9):1240–1249, 1960.

[42] J. Fleming, B. Kouvaritakis, and M. Cannon. Regions of attraction and recursive feasibility in robust MPC. In *21st Mediterranean Conference on Control and Automation*, number 5, pages 801–806. 2013.

[43] G. Franceschini and S. Macchietto. Model-based design of experiments for parameter precision: state of the art. *Chemical Engineering Science*, 63(19): 4846–4872, 2008.

[44] H. Fukushima, T.H. Kim, and T. Sugie. Adaptive model predictive control for a class of constrained linear systems based on the comparison model. *Automatica*, 43(2):301–308, 2007.

[45] S. Garg. Overview of propulsion controls and diagnostics research at NASA Glenn. In *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, Reston, Virigina, 2012.

[46] S. Garg. Modular Aero-Propulsion System Simulations C-MAPSS40k, 2012. URL `http://www.grc.nasa.gov/WWW/cdtb/software/mapss.html`.

[47] H. Genceli and M. Nikolaou. New approach to constrained predictive control with simultaneous model identification. *AIChE Journal*, 42(10):2857–2868, 1996.

[48] E. Gilbert and K. Tan. Linear systems with state and control constraints: the theory and application of maximal output admissible sets. *IEEE Transactions on Automatic Control*, 36(9):1008–1020, 1991.

[49] S. Glad and L. Ljung. Model structure identifiability and persistence of excitation. In *29th IEEE Conference on Decision and Control*, pages 3236–3240 vol.6. 1990.

[50] A. González, A. Ferramosca, G. Bustos, J. Marchetti, M. Fiacchini, and D. Odloak. Model predictive control suitable for closed-loop re-identification. *Systems & Control Letters*, 69(1):23–33, 2014.

[51] M. Green and D.J. Limebeer. *Linear Robust Control*. 1995.

[52] M. Green and J.B. Moore. Persistence of excitation in linear systems. *Systems & Control Letters*, 7(5):351–360, 1986.

[53] K. Grzędziński and P.A. Trodden. Learning MPC: system stability and convergent identification under bounded modelling error. In *2018 Australian & New Zealand Control Conference (ANZCC)*, pages 125–130. 2018.

[54] E.N. Hartley, P.A. Trodden, A.G. Richards, and J.M. Maciejowski. Model predictive control system design and implementation for spacecraft rendezvous. *Control Engineering Practice*, 20(7):695–713, 2012.

[55] E.N. Hartley, J.L. Jerez, A. Suardi, J.M. Maciejowski, E.C. Kerrigan, and G.A. Constantinides. Predictive control using an FPGA with application to aircraft control. *IEEE Transactions on Control Systems Technology*, 22(3):1006–1017, 2014.

[56] T.A.N. Heirung. *Dual control: Optimal, adaptive decision-making under uncertainty*. PhD thesis, Norwegian University of Science and Technology, 2016.

[57] T.A.N. Heirung, B. Foss, and B.E. Ydstie. MPC-based dual control with online experiment design. *Journal of Process Control*, 32:64–76, 2015.

[58] T.A.N. Heirung, J.A. Paulson, J. O'Leary, and A. Mesbah. Stochastic model predictive control — how does it work? *Computers & Chemical Engineering*, 114:158–170, 2018.

[59] B.A. Hernandez and P.A. Trodden. Persistently exciting tube MPC. In *2016 American Control Conference (ACC)*, pages 948–953, Boston, MA, 2016.

[60] B.A. Hernandez Vicente and P.A. Trodden. Stabilizing predictive control with persistence of excitation for constrained linear systems. *Systems & Control Letters*, 126:58–66, 2019.

[61] B.A. Hernandez Vicente and P.A. Trodden. Switching tube-based MPC: characterization of minimum dwell-time for feasible and robustly stable switching. *IEEE Transactions on Automatic Control*, 64(10):4345–4352, 2019.

[62] M.H. Holmes. *Introduction to Perturbation Methods*, volume 20 of *Texts in Applied Mathematics*. Springer New York, New York, NY, 2013.

[63] L. Jaw and J. Mattingly. *Aircraft Engine Controls*. American Institute of Aeronautics and Astronautics, Inc., Washington, DC, 2009.

[64] J.L. Jerez, P.J. Goulart, S. Richter, G.A. Constantinides, E.C. Kerrigan, and M. Morari. Embedded online optimization for model predictive control at megahertz rates. pages 1–14, 2013.

[65] R.M. Johnstone, C. Richard Johnson, R.R. Bitmead, and B.D. Anderson. Exponential convergence of recursive least squares with exponential forgetting factor. *Systems and Control Letters*, 2(2):77–82, 1982.

[66] L.H. Keel and S.P. Bhattacharyya. Controller synthesis free of analytical models: three term controllers. *IEEE Transactions on Automatic Control*, 53(6): 1353–1369, 2008.

[67] B. Khusainov, E.C. Kerrigan, and G.A. Constantinides. Automatic software and computing hardware codesign for predictive control. *IEEE Transactions on Control Systems Technology*, 27(5):2295–2304, 2019.

[68] H. Kiendl, J. Adamy, and P. Stelzner. Vector norms as Lyapunov functions for linear systems. *IEEE Transactions on Automatic Control*, 37(6):839–842, 1992.

[69] J. Kocijan, R. Murray-Smith, C. Rasmussen, and B. Likar. Predictive control with Gaussian process models. In *The IEEE Region 8 EUROCON 2003. Computer as a Tool.*, volume 1, pages 352–356. 2003.

[70] I. Kolmanovsky and E.G. Gilbert. Theory and computation of disturbance invariant sets for discrete-time linear systems. *Mathematical Problems in Engineering*, 4(4):317–367, 1998.

[71] I.V. Kolmanovsky, L.C. Jaw, W. Merrill, and H. Tran Van. Robust control and limit protection in aircraft gas turbine engines. In *2012 IEEE International Conference on Control Applications*, pages 812–819. 2012.

[72] B. Kouvaritakis, M. Cannon, and J.A. Rossiter. Who needs QP for linear MPC anyway? *Automatica*, 38(5):879–884, 2002.

[73] H. Lee and V.I. Utkin. Chattering suppression methods in sliding mode control systems. *Annual Reviews in Control*, 31(2):179–188, 2007.

[74] Y. Lee, S. Park, and J.H. Lee. On interfacing model predictive controllers with low-level loops. *Industrial & Engineering Chemistry Research*, 39(1):92–102, 2000.

[75] D.J. Leith and W.E. Leithead. Survey of gain-scheduling analysis and design. *International Journal of Control*, 73(11):1001–1025, 2000.

[76] D. Liberzon. *Switching in Systems and Control*. Systems & Control: Foundations & Applications. Birkhäuser Boston, Boston, MA, 2003.

[77] Lixian Zhang and R.D. Braatz. On switched MPC of a class of switched linear systems with modal dwell time. In *52nd IEEE Conference on Decision and Control*, pages 91–96. 2013.

[78] L. Ljung. *System Identification: Theory for the User*. Prentice Hall, 2 edition, 1997.

[79] L. Ljung and T. McKelvey. A least squares interpretation of sub-space methods for system identification. In *Proceedings of 35th IEEE Conference on Decision and Control*, volume 1, pages 335–342. 1996.

[80] J. Löfberg. Yalmip : A toolbox for modeling and optimization in matlab. In *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.

[81] J. Löfberg. Oops! I cannot do it again: Testing for recursive feasibility in MPC. *Automatica*, 48(3):550–555, 2012.

[82] M. Lorenzen, F. Allgöwer, and M. Cannon. Adaptive model predictive control with robust constraint satisfaction. In *IFAC-PapersOnLine*, volume 50, pages 3313–3318, 2017.

[83] M. Lorenzen, M. Cannon, and F. Allgöwer. Robust MPC with recursive model update. *Automatica*, 103:461–471, 2019.

[84] X. Lu and M. Cannon. Robust adaptive tube model predictive control. In *2019 American Control Conference (ACC)*. 2019. Available at `https://ora.ox.ac.uk/objects/uuid:02d18ffb-d4ff-4c7f-b793-94ed8c040255`.

[85] S. Lucia, M. Kögel, P. Zometa, D.E. Quevedo, and R. Findeisen. Predictive control, embedded cyberphysical systems and systems of systems – A perspective. *Annual Reviews in Control*, 41:193–207, 2016.

[86] O. Lyantsev, T. Breikin, G. Kulikov, and V. Arkov. On-line performance optimisation of aero engine control system. *Automatica*, 39(12):2115–2121, 2003.

[87] J.M. Maciejowski. *Predictive Control with Constraints*. Pearson, 2002.

[88] S. Mahmood. Control system, 2012. URL `http://www.google.co.uk/patents/US8321104`.

[89] S. Mahmood. *Application of the Rolls-Royce Inverse Model to Trent Aero Engines*. PhD thesis, University of Sheffield, 2001.

[90] G. Marafioti, R.R. Bitmead, and M. Hovd. Persistently exciting model predictive control. *International Journal of Adaptive Control and Signal Processing*, 28(6):536–552, 2014.

[91] I. Mareels, R. Bitmead, M. Gevers, C. Johnson, R. Kosut, and M. Poubelle. How exciting can a signal really be? *Systems & Control Letters*, 8(3):197–204, 1987.

[92] D.Q. Mayne. Robust and stochastic MPC: are we going in the right direction? *IFAC-PapersOnLine*, 48(23):1–8, 2015.

[93] D.Q. Mayne. An apologia for stabilising terminal conditions in model predictive control. *International Journal of Control*, 86(11):2090–2095, 2013.

[94] D.Q. Mayne, E.C. Kerrigan, E.J. van Wyk, and P. Falugi. Tube-based robust nonlinear model predictive control. *International Journal of Robust and Nonlinear Control*, 21(11):1341–1353, 2011.

[95] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[96] P. Mhaskar, N.H. El-Farra, and P.D. Christofides. Robust predictive control of switched systems: Satisfying uncertain schedules subject to state and control constraints. *International Journal of Adaptive Control and Signal Processing*, 22(2):161–179, 2008.

[97] F. Mirko and A. Mazen. Computing control invariant sets is easy. 2017. Available at `http://arxiv.org/abs/1708.04797`.

[98] J. Mu, D. Rees, and G. Liu. Advanced controller design for aircraft gas turbine engines. *Control Engineering Practice*, 13(8):1001–1015, 2005.

[99] M.A. Müller, P. Martius, and F. Allgöwer. Model predictive control of switched nonlinear systems under average dwell-time. *Journal of Process Control*, 22(9):1702–1710, 2012.

[100] C.J. Ong, Z. Wang, and M. Dehghan. Model predictive control for switching systems with dwell-time restriction. *IEEE Transactions on Automatic Control*, 61(12):4189–4195, 2016.

[101] M. Pakmehr, N. Fitzgerald, E.M. Feron, J.S. Shamma, and A. Behbahani. Gain scheduled control of gas turbine engines: stability and verification. *Journal of Engineering for Gas Turbines and Power*, 136(3):031201, 2013.

[102] M. Pakmehr, N. Fitzgerald, E. Feron, J. Paduano, and A. Behbahani. Physics-based dynamic modeling of a turboshaft engine driving a variable pitch propeller. *Journal of Propulsion and Power*, 32(3):646–658, 2016.

[103] A.A. Perpignan, A. Gangoli Rao, and D.J. Roekaerts. Flameless combustion and its potential towards gas turbines. *Progress in Energy and Combustion Science*, 69:28–62, 2018.

[104] S.V. Raković, B. Kouvaritakis, R. Findeisen, and M. Cannon. Homothetic tube model predictive control. *Automatica*, 48(8):1631–1638, 2012.

[105] J. Rathouský and V. Havlena. MPC-based approximate dual controller by information matrix maximization. *International Journal of Adaptive Control and Signal Processing*, 27(11):974–999, 2013.

[106] J.B. Rawlings and D.Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill, 2009.

[107] J.B. Rawlings, D.Q. Mayne, and M.M. Diehl. *Model Predictive Control: Theory, Computation and Design*. Nob Hill, 2 edition, 2018.

[108] J. Richalet, A. Rault, J. Testud, and J. Papon. Model predictive heuristic control. *Automatica*, 14(5):413–428, 1978.

[109] H. Richter. *Advanced Control of Turbofan Engines*. Springer New York, New York, NY, 2012.

[110] H. Richter and J. Litt. A novel controller for gas turbine engines with aggressive Limit management. In *47th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, number August, pages 1–17, Reston, Virigina, 2011.

[111] H. Richter, A.V. Singaraju, and J.S. Litt. Multiplexed predictive control of a large commercial turbofan engine. *Journal of Guidance, Control, and Dynamics*, 31(2):273–281, 2008.

[112] D. Ring, A.K. Christiansson, and M. Härefors. Fault tolerant multivariable control of a military turbofan engine. *IFAC Proceedings Volumes*, 38(1):538–543, 2005.

[113] D.C. Saluru, R.K. Yedavalli, and R.K. Belapurkar. Active Fault Tolerant Model Predictive Control of a Turbofan Engine Using C-MAPSS40k. In *ASME 2012 5th Annual Dynamic Systems and Control Conference joint with the JSME 2012 11th Motion and Vibration Conference*, pages 349–358. 2012.

[114] R. Samar and I. Postlethwaite. Multivariable controller design for a high performance aero-engine. In *International Conference on Control '94*, number 389, pages 1312–1317. 1994.

[115] M. Sanchez-Parra, D.A. Suarez, and C. Verde. Fault tolerant control for gas turbines. In *2011 16th International Conference on Intelligent System Applications to Power Systems*, pages 1–6. 2011.

[116] H.I.H. Saravanamuttoo, G. Rogers, H. Cohen, and P. Straznicky. *Gas Turbine Theory*. Pearson, 6 edition, 2009.

[117] L. Schwenkel, M. Gharbi, S. Trimpe, and C. Ebenbauer. Online learning with stability guarantees: A memory-based real-time model predictive controller, 2018. URL `https://arxiv.org/abs/1812.09582`.

[118] S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control*. John Wiley & Sons, 2 edition, 2005.

[119] Q. Song, J. Wilkie, and M.J. Grimble. Robust controller for gas turbines based upon LQG/LTR design with self-tuning features. *Journal of Dynamic Systems, Measurement, and Control*, 115(3):569, 1993.

[120] Z. Sun, S.J. Qin, A. Singhal, and L. Megan. Performance monitoring of model-predictive controllers via model residual assessment. *Journal of Process Control*, 23(4):473–482, 2013.

[121] S. Surendran, R. Chandrawanshi, S. Kulkarni, S. Bhartiya, P.S.V. Nataraj, and S. Sampath. Multi-parametric model predictive control strategy on laboratory SR-30 gas turbine. In *ASME 2015 Gas Turbine India Conference*, page V001T06A001. 2015.

[122] D. Tan, A. He, X. Wang, and Y. Liu. Multivariable Aeroengine PID Control With Amplitude Saturation: An LMI Solution. In *Volume 3: Controls, Diagnostics and Instrumentation; Cycle Innovations; Marine*, pages 277–285. 2010.

[123] M. Tanaskovic, L. Fagiano, R. Smith, and M. Morari. Adaptive receding horizon control for constrained MIMO systems. *Automatica*, 50(12):3019–3029, 2014.

[124] M. Tanaskovic, L. Fagiano, and V. Gligorovski. Adaptive model predictive control for linear time varying MIMO systems. *Automatica*, 105(11):237–245, 2019.

[125] S. Tarbouriech and M. Turner. Anti-windup design: an overview of some recent advances and open problems. *IET Control Theory & Applications*, 3(1): 1–19, 2009.

[126] P. Trodden. A one-step approach to computing a polytopic robust positively invariant set. *IEEE Transactions on Automatic Control*, 61(12):4100–4105, 2016.

[127] K.P. Wabersich and M.N. Zeilinger. Linear model predictive safety certification for learning-based control. 2018.

[128] J. Wang, N. Min, Z. Ye, and Z. Hu. Nonlinear control of aircraft engines using a generalized minimum variance based approach. In *Proceedings of the 32nd Chinese Control Conference*, pages 5368–5372, 2013.

[129] J. Wang, Z. Ye, and Z. Hu. Nonlinear control of aircraft engines using a generalized minimum variance based approach. In *Proceedings of the 32nd Chinese Control Conference*, volume 134, page 094502. 2013.

[130] X. Wang, J. Zhao, and X.M. Sun. Overshoot-free acceleration of aero-engines: An energy-based switching control method. *Control Engineering Practice*, 47: 28–36, 2016.

[131] X. Wang, Y. Sun, and K. Deng. Adaptive model predictive control of uncertain constrained systems. In *2014 American Control Conference*, pages 2857–2862. 2014.

[132] A. Weiss and S. Di Cairano. Robust dual control MPC with guaranteed constraint satisfaction. *Proceedings of the IEEE Conference on Decision and Control*, pages 6713–6718, 2014.

[133] A. Wiese, M. Blom, C. Manzie, M. Brear, and A. Kitchener. Model reduction and MIMO model predictive control of gas turbine systems. *Control Engineering Practice*, 45:194–206, 2015.

[134] G. Wood and B. Zhang. Estimation of the Lipschitz constant of a function. *Journal of Global Optimization*, 8(1):91–103, 1996.

[135] Y.G. Xi, D.W. Li, and S. Lin. Model predictive control — status and challenges. *Acta Automatica Sinica*, 39(3):222–236, 2013.

[136] X. Yang and J.M. Maciejowski. Fault tolerant control using Gaussian processes and model predictive control. *International Journal of Applied Mathematics and Computer Science*, 25(1):1–12, 2015.

[137] E. Žáčeková, S. Prívara, and M. Pčolka. Persistent excitation condition within the dual control framework. *Journal of Process Control*, 23(9):1270–1280, 2013.

[138] M. Zagrobelny, L. Ji, and J.B. Rawlings. Quis custodiet ipsos custodes? *Annual Reviews in Control*, 37(2):260–270, 2013.

[139] M.N. Zeilinger, M. Morari, and C.N. Jones. Soft constrained model predictive control with robust stability guarantees. *IEEE Transactions on Automatic Control*, 59(5):1190–1202, 2014.

[140] L. Zhang, E.k. Boukas, and P. Shi. Exponential H filtering for uncertain discrete-time switched linear systems with average dwell time: A - dependent approach. *International Journal of Robust and Nonlinear Control*, 18(11):1188–1207, 2008.

[141] L. Zhang, S. Zhuang, and R.D. Braatz. Switched model predictive control of switched linear systems: Feasibility, stability and robustness. *Automatica*, 67:8–21, 2016.