

FORMALLY MODELLING AND VERIFYING THE FREERTOS REAL-TIME OPERATING SYSTEM

Shu Cheng

Doctor of Philosophy
University of York
Computer Science

September, 2014

Abstract

Formal methods is an alternative way to develop software, which applies mathematical techniques to software design and verification. It ensures logical consistency between the requirements and the behaviour of the software, because each step in the development process, i.e., abstracting the requirements, design, refinement and implementation, is verified by mathematical techniques. However, in ordinary software development, the correctness of the software is tested at the end of the development process, which means it is limited and incomplete. Thus formal methods provides higher quality software than ordinary software development. At the same time, real-time operating systems are playing increasingly more important roles in embedded applications. Formal verification of this kind of software is therefore of strong interest.

FreeRTOS has a wide community of users: it is regarded by many as the *de facto* standard for micro-controllers in embedded applications. This project formally specifies the behaviour of FreeRTOS in Z, and its consistency is verified using the Z/Eves theorem prover. This includes a precise statement of the preconditions for all API commands. Based on this model, (a) code-level annotations for verifying task related API are produced with Microsoft's Verifying C Compiler (VCC); and (b) an abstract model for extension of FreeRTOS to multi-core architectures is specified with the Z notation.

This work forms the basis of future work that is refinement of the models to code to produce a verified implementation for both single and multi-core platforms.

Contents

Abstract	iii
Contents	v
List of Figures	xi
List of Tables	xiii
List of Schemas	xv
Acknowledgements	xxix
Declaration	xxxix
1 Introduction	1
1.1 Formal Methods	1
1.2 FreeRTOS	3
1.2.1 Task Management	4
1.2.2 Communication and Synchronisation	8
1.2.3 Other API functions	11
1.3 VCC	12
1.4 Multi-core Processor	13
1.5 Objectives and Challenges	15
1.5.1 Objectives	15
1.5.2 Structure View	15
1.5.3 Challenges	16

1.6	Font and Name Styles	17
1.7	Structure of Thesis	19
2	Literature Review	21
2.1	Related Work	21
2.2	Z Notation	25
2.2.1	Tools for Z	27
2.3	FreeRTOS	28
2.4	Summary	29
3	Abstract FreeRTOS	31
3.1	Model Overview	31
3.2	Goal and Scope	32
3.3	Requirements	33
3.3.1	Functional Requirements	33
3.3.2	Non-functional Requirements	35
3.3.3	Environment Requirements	36
3.4	Summary	36
4	Modelling FreeRTOS	37
4.1	Iteration Process	37
4.2	Task Model	40
4.2.1	Basic Statements	40
4.2.2	Additional Schema for Reschedule	47
4.2.3	Creating and Deleting Tasks	49
4.2.4	Executing Tasks	60
4.2.5	Suspending/Resuming Tasks	60
4.2.6	Changing Priority of Tasks	62
4.3	Queue Model	63
4.3.1	Basic Statements	63
4.3.2	Extension	67
4.3.3	Creating and Deleting Queues	69

4.3.4	Sending and Receiving Items	69
4.4	Time Model	72
4.4.1	Basic Statements	72
4.4.2	Extension	75
4.4.3	Delaying Tasks	76
4.4.4	Checking Delayed Tasks	76
4.4.5	Time-Sharing	78
4.5	Mutex Model	78
4.5.1	Basic Statement	79
4.5.2	Extension	81
4.5.3	Creating and Deleting Semaphores and Mutexes	83
4.5.4	Taking Mutexes	84
4.5.5	Giving Mutexes	85
4.6	Summary of Interface	88
4.7	Some Properties	89
4.8	Summary	91
5	VCC Verification of FreeRTOS	93
5.1	Overview	93
5.2	Statement Definition	96
5.2.1	Translation from Z to VCC	100
5.3	Creating Tasks	101
5.4	Deleting Tasks	106
5.5	Getting and Setting Priority	110
5.6	Suspending and Resuming Tasks	112
5.7	Summary	113
6	Extension for Multi-Core	115
6.1	Overview	115
6.2	Task Model	116
6.2.1	Basic Statements	116
6.2.2	Additional Schemas	119

6.2.3	Creating and Deleting Tasks	121
6.2.4	Migrating Task	122
6.2.5	Other Operations	123
6.3	Queue Model	123
6.4	Time and Mutex Model	125
6.4.1	Time Slicing	126
6.4.2	Taking Mutexes	127
6.4.3	Giving Mutexes	128
6.5	Summary	131
7	Evaluation and Case Studies	133
7.1	Project Summary	133
7.2	Case Studies	134
7.2.1	Case 1	135
7.2.2	Case 2	138
7.2.3	Case 3	141
7.3	Issues of FreeRTOS	143
7.4	Summary	146
8	Conclusion and Future Work	147
8.1	Conclusion	147
8.2	Future Work	149
8.3	Task Model with Promotion	150
8.4	Summary	154
	Appendix A Introductory Appendix	155
	Appendix B Summary of Z/Eves Proof Commands	157
	Appendix C Summary of Interface	159
	Appendix D Specification for Task Model	171
	Appendix E Specification for Queue Model	185

Appendix F	Specification for Time Model	205
Appendix G	Specification for Mutex Model	225
Appendix H	Specification for Multi-core Task Model	281
Appendix I	Specification for Multi-core Queue Model	293
Appendix J	Specification for Multi-core Time Model	303
Appendix K	Specification for Multi-core Mutex Model	311
Appendix L	Spec for Multi-core Task with Promotion	331
Appendix M	VCC Annotated Source Code	343
Bibliography		361

List of Figures

1.1	State chart for Tasks	5
1.2	An example application that uses RTOS (Task related).	6
1.3	An example application that uses RTOS (Communication related).	10
1.4	Overview of project	17
5.1	Constraints of tskTCB	96
5.2	State and transition in VCC	97
5.3	FreeRTOS structure	98
5.4	Contract for creating tasks	103
5.5	Creating tasks pre-verification	104
5.6	Creating tasks verification part-1	105
5.7	Creating tasks verification part-2	107
5.8	Contract for deleting tasks	108
5.9	Postconditions for priority setting	111
7.1	API function execution history and result for Case 1	136
7.2	API function execution history (above) and result (bottom) for Case 2	140
7.3	API function execution history (above) and result (bottom) for Case 3	142
7.4	Scenario for priority inversion issue	145

List of Tables

1.1	Prefix of variable and function names used in FreeRTOS	18
1.2	Suffixes used in schema names	18
4.1	The constraints for giving mutexes (no waiting tasks)	87
4.2	The constraints for giving mutexes (with waiting tasks)	87
6.1	Conditions for giving mutex cases (have waiting tasks and the mutex holder inherits the priority)	130
C.1	API mappings & preconditions for operations	160
C.2	API mappings & preconditions for operations(continue)	161
C.3	API mappings & preconditions for operations(continue)	162
C.4	API mappings & preconditions for operations(continue)	163
C.5	API mappings & preconditions for operations(continue)	164
C.6	API mappings & preconditions for operations(continue)	165
C.7	API mappings & preconditions for operations(continue)	166
C.8	API mappings & preconditions for operations(continue)	167
C.9	API mappings & preconditions for operations(continue)	168
C.10	API mappings & preconditions for operations(continue)	169

List of Schemas

2.1	<i>Task</i>	26
2.2	<i>setPrio</i>	26
4.1	<i>FreeRTOS</i>	38
4.2	<i>TaskData</i>	42
4.3	<i>StateData</i>	43
4.4	<i>ContextData</i>	43
4.5	<i>PrioData</i>	43
4.6	<i>Task</i>	44
4.7	<i>Init_TaskData</i>	44
4.8	<i>Init_StateData</i>	45
4.9	<i>Init_ContextData</i>	45
4.10	<i>Init_PrioData</i>	45
4.11	<i>Init_Task</i>	45
4.12	Δ <i>Task</i>	47
4.13	<i>Reschedule</i>	48
4.14	<i>CreateTaskN_T</i>	49
4.15	<i>CreateTaskN_TFSBSig</i>	50
4.16	<i>CreateTaskS_T</i>	54
4.17	<i>CreateTaskS_TFSBSig</i>	55

4.18	<i>DeleteTaskS_T</i>	56
4.19	<i>DeleteTaskS_TFSBSig</i>	56
4.20	<i>QueueData</i>	64
4.21	<i>WaitingData</i>	64
4.22	<i>QReleasingData</i>	65
4.23	<i>Queue</i>	65
4.24	<i>TaskQueue</i>	65
4.25	<i>Init_TaskQueue</i>	66
4.26	<i>ExtendTaskXi</i>	67
4.27	<i>ExtTaskFSBSig</i>	67
4.28	<i>DeleteTaskN_TQ</i>	68
4.29	<i>QueueSendN_TQ</i>	70
4.30	<i>Time</i>	73
4.31	<i>TaskQueueTime</i>	73
4.32	<i>Init_Time</i>	75
4.33	<i>QueueSendF_TQT</i>	76
4.34	<i>CheckDelayedTaskN_TQT</i>	77
4.35	<i>MutexData</i>	79
4.36	<i>OriginalPrioData</i>	79
4.37	<i>MReleasingData</i>	80
4.38	<i>Mutex</i>	80
4.39	<i>TaskQueueTimeMutex</i>	80
4.40	<i>basePriorityMan</i>	86
6.1	<i>TaskData</i>	117
6.2	<i>Task</i>	118
6.3	<i>Reschedule</i>	119
6.4	<i>findTopReady</i>	119
6.5	<i>findACore_T</i>	120
6.6	<i>QueueData</i>	123
6.7	<i>ChangeQueueLevel_TQ</i>	125
6.8	<i>TimeSlicing_TQT</i>	126

8.1	<i>TaskData</i>	151
8.2	<i>Multi_Task</i>	151
8.3	<i>Init</i>	152
8.4	<i>PromoteC</i>	152
D.1	<i>TaskData</i>	171
D.2	<i>Init_TaskData</i>	172
D.3	<i>StateData</i>	172
D.4	<i>Init_StateData</i>	172
D.5	<i>ContextData</i>	172
D.6	<i>Init_ContextData</i>	172
D.7	<i>PrioData</i>	173
D.8	<i>Init_PrioData</i>	173
D.9	<i>Task</i>	173
D.10	Δ <i>Task</i>	173
D.11	<i>Init_Task</i>	174
D.12	<i>Reschedule</i>	174
D.13	<i>CreateTaskN_T</i>	174
D.14	<i>CreateTaskN_TFSBSig</i>	175
D.15	<i>CreateTaskS_T</i>	175
D.16	<i>CreateTaskS_TFSBSig</i>	176
D.17	<i>DeleteTaskN_T</i>	176
D.18	<i>DeleteTaskN_TFSBSig</i>	176
D.19	<i>DeleteTaskS_T</i>	177
D.20	<i>DeleteTaskS_TFSBSig</i>	177
D.21	<i>ExecuteRunningTask_T</i>	178
D.22	<i>ExecuteRunningTask_TFSBSig</i>	178
D.23	<i>SuspendTaskN_T</i>	178
D.24	<i>SuspendTaskN_TFSBSig</i>	178
D.25	<i>SuspendTaskS_T</i>	179
D.26	<i>SuspendTaskS_TFSBSig</i>	179
D.27	<i>SuspendTaskO_T</i>	180

D.28	<i>SuspendTaskO_TFSBSig</i>	180
D.29	<i>ResumeTaskN_T</i>	180
D.30	<i>ResumeTaskN_TFSBSig</i>	181
D.31	<i>ResumeTaskS_T</i>	181
D.32	<i>ResumeTaskS_TFSBSig</i>	181
D.33	<i>ChangeTaskPriorityN_T</i>	181
D.34	<i>ChangeTaskPriorityN_TFSBSig</i>	182
D.35	<i>ChangeTaskPriorityS_T</i>	182
D.36	<i>ChangeTaskPriorityS_TFSBSig</i>	182
D.37	<i>ChangeTaskPriorityD_T</i>	183
D.38	<i>ChangeTaskPriorityD_TFSBSig</i>	183
E.1	<i>QueueData</i>	185
E.2	<i>Init_QueueData</i>	185
E.3	<i>WaitingData</i>	185
E.4	<i>Init_WaitingData</i>	185
E.5	<i>QReleasingData</i>	186
E.6	<i>Init_QReleasingData</i>	186
E.7	<i>Queue</i>	186
E.8	<i>Init_Queue</i>	187
E.9	<i>TaskQueue</i>	187
E.10	<i>Init_TaskQueue</i>	187
E.11	<i>ExtendTaskXi</i>	187
E.12	<i>ExtTaskFSBSig</i>	187
E.13	<i>DeleteTaskN_TQ</i>	188
E.14	<i>SuspendTaskN_TQ</i>	189
E.15	<i>CreateQueue_TQ</i>	192
E.16	<i>CreateQueue_TQFSBSig</i>	193
E.17	<i>DeleteQueue_TQ</i>	193
E.18	<i>DeleteQueue_TQFSBSig</i>	193
E.19	<i>QueueSendN_TQ</i>	193
E.20	<i>QueueSendN_TQFSBSig</i>	194

E.21	<i>QueueSendF_TQ</i>	194
E.22	<i>QueueSendF_TQFSBSig</i>	195
E.23	<i>QueueSendW_TQ</i>	196
E.24	<i>QueueSendW_TQFSBSig</i>	196
E.25	<i>QueueSendWS_TQ</i>	197
E.26	<i>QueueSendWS_TQFSBSig</i>	197
E.27	<i>QueueReceiveN_TQ</i>	199
E.28	<i>QueueReceiveN_TQFSBSig</i>	199
E.29	<i>QueueReceiveE_TQ</i>	199
E.30	<i>QueueReceiveE_TQFSBSig</i>	200
E.31	<i>QueueReceiveW_TQ</i>	201
E.32	<i>QueueReceiveW_TQFSBSig</i>	201
E.33	<i>QueueReceiveWS_TQ</i>	202
E.34	<i>QueueReceiveWS_TQFSBSig</i>	203
F.1	<i>Time</i>	205
F.2	<i>Init_Time</i>	205
F.3	<i>TaskQueueTime</i>	205
F.4	<i>Init_TaskQueueTime</i>	206
F.5	<i>ExtendTaskQueueXi</i>	206
F.6	<i>DeleteTaskN_TQT</i>	207
F.7	<i>SuspendTaskN_TQT</i>	208
F.8	<i>QueueSendF_TQT</i>	212
F.9	<i>QueueSendF_TQTFSBSig</i>	212
F.10	<i>QueueSendW_TQT</i>	213
F.11	<i>QueueSendWS_TQT</i>	214
F.12	<i>QueueReceiveE_TQT</i>	215
F.13	<i>QueueReceiveE_TQTFSBSig</i>	215
F.14	<i>QueueReceiveW_TQT</i>	216
F.15	<i>QueueReceiveWS_TQT</i>	217
F.16	<i>DelayUntil_TQT</i>	218
F.17	<i>DelayUntil_TQTFSBSig</i>	219

F.18	<i>CheckDelayedTaskN_TQT</i>	220
F.19	<i>CheckDelayedTaskN_TQTFSSBSig</i>	220
F.20	<i>CheckDelayedTaskS_TQT</i>	221
F.21	<i>CheckDelayedTaskS_TQTFSSBSig</i>	222
F.22	<i>TimeSlicing_TQT</i>	223
F.23	<i>TimeSlicing_TQTFSSBSig</i>	223
F.24	<i>NoSlicing_TQT</i>	224
F.25	<i>NoSlicing_TQTFSSBSig</i>	224
G.1	<i>MutexData</i>	225
G.2	<i>Init_MutexData</i>	225
G.3	<i>OriginalPrioData</i>	225
G.4	<i>Init_OriginalPrioData</i>	226
G.5	<i>MReleasingData</i>	226
G.6	<i>Init_MReleasingData</i>	226
G.7	<i>Mutex</i>	226
G.8	<i>Init_Mutex</i>	226
G.9	<i>TaskQueueTimeMutex</i>	226
G.10	<i>Init_TaskQueueTimeMutex</i>	227
G.11	<i>ExtendTQTXi</i>	227
G.12	<i>DeleteTaskN_TQTM</i>	228
G.13	<i>DeleteTaskN_TQTMFSBSig</i>	228
G.14	<i>DeleteTaskS_TQTM</i>	228
G.15	<i>DeleteTaskS_TQTMFSBSig</i>	228
G.16	<i>ChangeTaskPriorityNNotHolder_TQTM</i>	231
G.17	<i>ChangeTaskPriorityNNotHolder_TQTMFSBSig</i>	231
G.18	<i>ChangeTaskPrioritySNotHolder_TQTM</i>	231
G.19	<i>ChangeTaskPrioritySNotHolder_TQTMFSBSig</i>	232
G.20	<i>ChangeTaskPriorityDNotHolder_TQTM</i>	232
G.21	<i>ChangeTaskPriorityDNotHolder_TQTMFSBSig</i>	232
G.22	<i>ChangeTaskPriorityNNotInherited_TQTM</i>	233
G.23	<i>ChangeTaskPriorityNNotInherited_TQTMFSBSig</i>	233

G.24	<i>ChangeTaskPrioritySNotInherited_TQTM</i>	233
G.25	<i>ChangeTaskPrioritySNotInherited_TQTMFSBSig</i>	233
G.26	<i>ChangeTaskPriorityDNotInherited_TQTM</i>	234
G.27	<i>ChangeTaskPriorityDNotInherited_TQTMFSBSig</i>	234
G.28	<i>ChangeTaskPriorityInheritedN_TQTM</i>	235
G.29	<i>ChangeTaskPriorityInheritedN_TQTMFSBSig</i>	235
G.30	<i>ChangeTaskPriorityInheritedU_TQTM</i>	236
G.31	<i>ChangeTaskPriorityInheritedU_TQTMFSBSig</i>	236
G.32	<i>ChangeTaskPriorityInheritedS_TQTM</i>	236
G.33	<i>ChangeTaskPriorityInheritedS_TQTMFSBSig</i>	236
G.34	<i>DeleteQueue_TQTM</i>	237
G.35	<i>DeleteQueue_TQTMFSBSig</i>	237
G.36	<i>QueueSendN_TQTM</i>	238
G.37	<i>QueueSendN_TQTMFSBSig</i>	238
G.38	<i>QueueSendF_TQTM</i>	238
G.39	<i>QueueSendF_TQTMFSBSig</i>	238
G.40	<i>QueueSendW_TQTM</i>	239
G.41	<i>QueueSendW_TQTMFSBSig</i>	239
G.42	<i>QueueSendWS_TQTM</i>	240
G.43	<i>QueueSendWS_TQTMFSBSig</i>	240
G.44	<i>QueueReceiveN_TQTM</i>	241
G.45	<i>QueueReceiveN_TQTMFSBSig</i>	242
G.46	<i>QueueReceiveE_TQTM</i>	242
G.47	<i>QueueReceiveE_TQTMFSBSig</i>	242
G.48	<i>QueueReceiveW_TQTM</i>	243
G.49	<i>QueueReceiveW_TQTMFSBSig</i>	243
G.50	<i>QueueReceiveWS_TQTM</i>	244
G.51	<i>QueueReceiveWS_TQTMFSBSig</i>	244
G.52	<i>CreateBinarySemaphore_TQTM</i>	249
G.53	<i>CreateBinarySemaphore_TQTMFSBSig</i>	249
G.54	<i>DeleteBinarySemaphore_TQTM</i>	250

G.55	<i>DeleteBinarySemaphore_TQTMFSBSig</i>	250
G.56	<i>CreateMutex_TQTM</i>	250
G.57	<i>CreateMutex_TQTMFSBSig</i>	251
G.58	<i>DeleteMutex_TQTM</i>	251
G.59	<i>DeleteMutex_TQTMFSBSig</i>	251
G.60	<i>MutexTakeNnonInh_TQTM</i>	252
G.61	<i>MutexTakeNnonInh_TQTMFSBSig</i>	252
G.62	<i>MutexTakeNInh_TQTM</i>	252
G.63	<i>MutexTakeNInh_TQTMFSBSig</i>	253
G.64	<i>MutexTakeEnonInh_TQTM</i>	253
G.65	<i>MutexTakeEnonInh_TQTMFSBSig</i>	253
G.66	<i>MutexTakeEInheritReady_TQTM</i>	254
G.67	<i>MutexTakeEInheritReady_TQTMFSBSig</i>	255
G.68	<i>MutexTakeEInheritHolder_TQTM</i>	256
G.69	<i>MutexTakeEInheritHolder_TQTMFSBSig</i>	257
G.70	<i>MutexTakeRecursive_TQTM</i>	257
G.71	<i>MutexTakeRecursive_TQTMFSBSig</i>	258
G.72	<i>basePriorityMan</i>	258
G.73	<i>MutexGiveNnonInh_TQTM</i>	258
G.74	<i>MutexGiveNnonInh_TQTMFSBSig</i>	259
G.75	<i>MutexGiveNInhN_TQTM</i>	259
G.76	<i>MutexGiveNInhN_TQTMFSBSig</i>	260
G.77	<i>MutexGiveNInhS_TQTM</i>	260
G.78	<i>MutexGiveNInhS_TQTMFSBSig</i>	261
G.79	<i>MutexGiveWnonInhN_TQTM</i>	263
G.80	<i>MutexGiveWnonInhN_TQTMFSBSig</i>	263
G.81	<i>MutexGiveWnonInhS_TQTM</i>	265
G.82	<i>MutexGiveWnonInhS_TQTMFSBSig</i>	266
G.83	<i>MutexGiveWInhN_TQTM</i>	268
G.84	<i>MutexGiveWInhN_TQTMFSBSig</i>	269
G.85	<i>MutexGiveWInhSR_TQTM</i>	271

G.86	<i>MutexGiveWinhSR_TQTMFSBSig</i>	272
G.87	<i>MutexGiveWinhSW_TQTM</i>	275
G.88	<i>MutexGiveWinhSW_TQTMFSBSig</i>	275
G.89	<i>MutexGiveNRecursive_TQTM</i>	278
G.90	<i>MutexGiveNRecursive_TQTMFSBSig</i>	278
H.1	<i>TaskData</i>	282
H.2	<i>Init_TaskData</i>	282
H.3	<i>StateData</i>	282
H.4	<i>Init_StateData</i>	282
H.5	<i>ContextData</i>	282
H.6	<i>Init_ContextData</i>	282
H.7	<i>PrioData</i>	282
H.8	<i>Init_PrioData</i>	282
H.9	<i>Task</i>	283
H.10	Δ <i>Task</i>	283
H.11	<i>Init_Task</i>	283
H.12	<i>createTaskSpeCoreN_T</i>	283
H.13	<i>findACore_T</i>	284
H.14	<i>Reschedule</i>	284
H.15	<i>createTaskSpeCoreS_T</i>	285
H.16	<i>DeleteTaskN_T</i>	285
H.17	<i>findTopReady</i>	285
H.18	<i>DeleteTaskS_T</i>	286
H.19	<i>ExecuteRunningTask_T</i>	286
H.20	<i>SuspendTaskN_T</i>	286
H.21	<i>SuspendTaskS_T</i>	287
H.22	<i>SuspendTaskO_T</i>	287
H.23	<i>ResumeTaskN_T</i>	287
H.24	<i>ResumeTaskS_T</i>	287
H.25	<i>ChangeTaskPriorityN_T</i>	288
H.26	<i>ChangeTaskPriorityS_T</i>	288

H.27	<i>ChangeTaskPriorityD_T</i>	288
H.28	<i>MigrationN_T</i>	289
H.29	<i>MigrationS_T</i>	289
H.30	<i>MigrationRuN_T</i>	290
H.31	<i>MigrationRuS_T</i>	290
I.1	<i>QueueData</i>	293
I.2	<i>Init_QueueData</i>	293
I.3	<i>WaitingData</i>	293
I.4	<i>Init_WaitingData</i>	294
I.5	<i>QReleasingData</i>	294
I.6	<i>Init_QReleasingData</i>	294
I.7	<i>Queue</i>	294
I.8	<i>Init_Queue</i>	294
I.9	<i>TaskQueue</i>	295
I.10	<i>Init_TaskQueue</i>	295
I.11	<i>ExtendTaskXi</i>	295
I.12	<i>DeleteTaskN_TQ</i>	295
I.13	<i>SuspendTaskN_TQ</i>	296
I.14	<i>CreateQueue_TQ</i>	296
I.15	<i>DeleteQueue_TQ</i>	297
I.16	<i>QueueSendN_TQ</i>	297
I.17	<i>QueueSendF_TQ</i>	297
I.18	<i>QueueSendW_TQ</i>	298
I.19	<i>QueueSendWS_TQ</i>	299
I.20	<i>QueueReceiveN_TQ</i>	299
I.21	<i>QueueReceiveE_TQ</i>	300
I.22	<i>QueueReceiveW_TQ</i>	300
I.23	<i>QueueReceiveWS_TQ</i>	301
I.24	<i>ChangeQueueLevel_TQ</i>	302
J.1	<i>Time</i>	303
J.2	<i>Init_Time</i>	303

J.3	<i>TaskQueueTime</i>	303
J.4	<i>Init_TaskQueueTime</i>	304
J.5	<i>ExtendTaskQueueXi</i>	304
J.6	<i>DeleteTaskN_TQT</i>	304
J.7	<i>SuspendTaskN_TQT</i>	304
J.8	<i>QueueSendF_TQT</i>	305
J.9	<i>QueueSendW_TQT</i>	305
J.10	<i>QueueSendWS_TQT</i>	305
J.11	<i>QueueReceiveE_TQT</i>	306
J.12	<i>QueueReceiveW_TQT</i>	306
J.13	<i>QueueReceiveWS_TQT</i>	306
J.14	<i>DelayUntil_TQT</i>	307
J.15	<i>CheckDelayedTaskN_TQT</i>	307
J.16	<i>CheckDelayedTaskS_TQT</i>	308
J.17	<i>TimeSlicing_TQT</i>	308
J.18	<i>NoSlicing_TQT</i>	309
K.1	<i>MutexData</i>	311
K.2	<i>Init_MutexData</i>	311
K.3	<i>OriginalPrioData</i>	311
K.4	<i>Init_OriginalPrioData</i>	311
K.5	<i>MReleasingData</i>	312
K.6	<i>Init_MReleasingData</i>	312
K.7	<i>Mutex</i>	312
K.8	<i>Init_Mutex</i>	312
K.9	<i>TaskQueueTimeMutex</i>	312
K.10	<i>Init</i>	312
K.11	<i>ExtendTQTXi</i>	313
K.12	<i>DeleteTask_TQTM</i>	313
K.13	<i>ChangeTaskPriorityNotHolder_TQTM</i>	313
K.14	<i>ChangeTaskPriorityNotInherited_TQTM</i>	313
K.15	<i>ChangeTaskPriorityInheritedN_TQTM</i>	314

K.16	<i>ChangeTaskPriorityInheritedU_TQTM</i>	314
K.17	<i>ChangeTaskPriorityInheritedS_TQTM</i>	314
K.18	<i>DeleteQueue_TQTM</i>	315
K.19	<i>QueueSend_TQTM</i>	315
K.20	<i>QueueReceive_TQTM</i>	315
K.21	<i>CreateBinarySemaphore_TQTM</i>	316
K.22	<i>DeleteBinarySemaphore_TQTM</i>	316
K.23	<i>CreateMutex_TQTM</i>	317
K.24	<i>DeleteMutex_TQTM</i>	317
K.25	<i>MutexTakeNnonInh_TQTM</i>	317
K.26	<i>MutexTakeNInh_TQTM</i>	318
K.27	<i>MutexTakeRecursive_TQTM</i>	318
K.28	<i>MutexTakeEnonInh_TQTM</i>	319
K.29	<i>MutexTakeEInheritSameCoreHolder_TQTM</i>	319
K.30	<i>MutexTakeEInheritSameCoreReady_TQTM</i>	320
K.31	<i>MutexTakeEInheritDiffCoreN_TQTM</i>	321
K.32	<i>MutexTakeEInheritDiffCoreS_TQTM</i>	321
K.33	<i>basePriorityMan</i>	322
K.34	<i>MutexGiveNRecursive_TQTM</i>	323
K.35	<i>MutexGiveNnonInh_TQTM</i>	323
K.36	<i>MutexGiveNInhN_TQTM</i>	324
K.37	<i>MutexGiveNInhS_TQTM</i>	324
K.38	<i>MutexGiveWnonInhN_TQTM</i>	325
K.39	<i>MutexGiveWnonInhS_TQTM</i>	325
K.40	<i>MutexGiveWInhN_TQTM</i>	326
K.41	<i>MutexGiveWInhSR_TQTM</i>	327
K.42	<i>MutexGiveWInhSW_TQTM</i>	328
K.43	<i>MutexGiveWInhSBoth_TQTM</i>	329
L.1	<i>TaskData</i>	331
L.2	<i>Init_TaskData</i>	332
L.3	<i>StateData</i>	332

L.4	<i>ContextData</i>	332
L.5	<i>Init_ContextData</i>	332
L.6	<i>PrioData</i>	332
L.7	<i>Init_PrioData</i>	332
L.8	<i>Task</i>	332
L.9	Δ <i>Task</i>	332
L.10	<i>Init_Task</i>	333
L.11	<i>Reschedule</i>	333
L.12	<i>CreateTaskN_T</i>	333
L.13	<i>CreateTaskS_T</i>	333
L.14	<i>DeleteTaskN_T</i>	334
L.15	<i>DeleteTaskS_T</i>	334
L.16	<i>SuspendTaskN_T</i>	335
L.17	<i>SuspendTaskS_T</i>	335
L.18	<i>SuspendTaskO_T</i>	335
L.19	<i>ResumeTaskN_T</i>	335
L.20	<i>ResumeTaskS_T</i>	336
L.21	<i>ChangeTaskPriorityN_T</i>	336
L.22	<i>ChangeTaskPriorityS_T</i>	336
L.23	<i>ChangeTaskPriorityD_T</i>	336
L.24	<i>Multi_Task</i>	337
L.25	<i>Init</i>	337
L.26	<i>PromoteC</i>	337
L.27	<i>findACore_MT</i>	338
L.28	<i>PromoteD</i>	338
L.29	<i>Promote</i>	338
L.30	<i>MigrationN_MT</i>	339
L.31	<i>MigrationS_MT</i>	340
L.32	<i>MigrationRuN_MT</i>	340
L.33	<i>MigrationRuS_MT</i>	341
L.34	<i>getRunningTask</i>	342

L.35 *getPriority* 342

Acknowledgements

First and foremost, I would like to thank my supervisor, Prof. Jim Woodcock, whose patient and continuous guidance, constructive suggestions and kind encouragement helped me throughout my PhD. It has been great pleasure to work with him.

I also want to express my great gratitude to my parents and family. Without their boundless love, teaching and guidance I would never be who I am.

And I also want to thank my colleagues in the Computer Science Department at the University of York, who have contributed to discussions and support about my work. Especially, I appreciate Steve King, Rob Alexander, Jeremy Jacob, Sam Simpson, Victor Bandur, Pakorn Waewsawangwong, Tasawer Khan, Chris Poskitt, James Williams, and Detlef Plump. Thanks are also due to members of the Indian Institute of Science in Bangalore for their helpful discussions and comments of this work: thanks to Deepak D'Souza, Sumesh Divakaran, Anuridh Kushwah, Virendra Singh, and Nigamanth Shridar.

I also want to thank my assessors, Steve King and Jonathan Bowen. Their precious feedback helps me to make the thesis much more better.

Last, but not least, I want to appreciate all my friends: especially to Lin Liu, Siyu Wang, Lei Chen, Hao Sun, YunFeng Ma, Hengyu Wu, Wen Luo and Keshu Xu. Their friendship, help and encouragement, always support me, even in the hardest times, and bring me the most blissful memories.

Declaration

The work in this thesis has been carried out by the author between September 2010 and September 2014 at the Department of Computer Science, University of York. This work has not previously been presented for an award at this, or any other, University. Apart from work whose authors are clearly acknowledged and referenced, all other works presented in this thesis were carried out by the author. First section of Chapter 4 of this thesis have appeared in previously published paper:

- S. Cheng, J. Woodcock, and D. D'Souza, Using formal reasoning on a model of tasks for FreeRTOS. *Formal Aspects of Computing*, pp. 1–26, 2014. [Online]. Available: DOI:10.1007/s00165-014-0308-9

All the technical works presented in this paper are carried out by the author of this thesis.

Chapter 1

INTRODUCTION

This chapter begins by introducing formal methods and the international Verified Software Initiative (VSI). It then introduces FreeRTOS, as formally verifying the correctness of FreeRTOS is one of the pilot projects of the VSI. Next, this chapter introduces VCC and shows that it is possible to combine formal specification and code level annotations together to verify the source code. It then discusses the limitations of single core processors and the benefits of multi-core processors. In addition the objectives and challenges of verifying FreeRTOS are clarified. Finally, the structure of the thesis is given.

1.1 Formal Methods

Formal methods apply mathematical techniques to software design and verification and are normally supported by tools [1]. The general development process for formal methods is:

- (a) Use mathematical expressions to specify the state and the behaviour of the software according to the documented requirements, which can generally be expressed by state transitions. An abstract specification will be produced in this step;
- (b) Apply mathematical theorems and lemmas to verify the specification or model;

- (c) Refine the specification from abstract level to concrete model;
- (d) Repeat steps b & c until executable code is generated. Note, the relation between each refinement also has to be verified.

Formal methods can provide higher quality software than ordinary software development, because all models produced in the development process can be verified and proved using mathematical logic. This can efficiently detect any faults in the software at the initial stage of development, which may later have lead to huge losses. For instance, an error in the Inertial Reference System (IRS) of Ariane 5 caused the explosion of the rocket in June 1996 [2], costing around half a billion US dollars. The US Department of Commerce also estimates that the losses caused by avoidable software errors is between 20 and 60 billion dollars every year [3, 4]. Without formal methods these kinds of faults are sometimes very hard to discover, and even if revealed, may be too expensive to correct at a later stage [5].

However, formal methods have not been widely applied in industry, although they have significant advantages. Hall [6] believes that there are seven common myths or misunderstandings about formal methods, some of which cause this situation. Firstly, it is thought that formal methods increase the cost of the development. Yet Hall [6] indicates that in his experience, applying formal methods in commercial projects decreases the development cost. Although there is a one-time cost for learning the non-user-friendly tools which often support formal methods, the developer gains more benefits from the reduction of cost in the amount of testing and maintenance. King *et al.* [7] claim that using formal methods is more efficient for detecting faults than the most efficient testing phase, which also increases the cost. The cost of verifying and testing software may occupy 30% to 50% of the total cost of a software project [8]. This can increase to 70% for hardware. Even with this huge investment, however, Dijkstra [9] believes testing can never guarantee that software is free of bugs. Secondly, “formal methods involve complex mathematics and are incomprehensible to clients” [6]. Although formal methods apply mathematical technology in documentation and design, it only needs knowledge related to logic and set theory, which is a fundamental part of mathematics. Meanwhile,

using formal methods, developers may experiment with the model and demonstrate it to clients using animation. Such animations can clearly show clients the behaviour of the system. Formal methods help developers to organise documentation much better and because of the mathematical rigour, the documentation is also more likely to be unambiguous and precise. This makes it easier for clients to use and understand the system [4]. Lastly, people believe that formal methods are only used in academic and research fields or in highly critical systems. However, it has been reported that formal methods are suitable for industrial-scale applications [6, 7, 10, 11]. Moreover, Berry [12], chief scientist at ESTEREL Technologies, shows that the control system of the Airbus A380, which has five million lines of code, was automatically generated by formal methods and all worked first time.

Fortunately, this situation is changing. In 2003, Hoare [13] suggested the international Grand Challenge for Computing Research to build a verifying compiler, which could automatically verify whether a program met its requirements [14]. Based on this idea, the international Verified Software Initiative (VSI) [4, 5, 13, 15], led by Hoare, was proposed. The main aim of the VSI is to work out a more approachable strategy for verifying software with the integrated support tools. Several pilot projects have been selected for VSI, such as the Mondex electronic purse [16], POSIX file store [17], etc.

1.2 FreeRTOS

As a widely used real-time operating system, the function of FreeRTOS can be divided into three large categories: (a) multitasking task management; (b) inter-task communication and synchronisation; and (c) memory management, interrupt management and other features. The three key elements of FreeRTOS are:

Tasks: user processes.

Queues: communication mechanisms between tasks and interrupts.

Semaphores and Mutexes: the facilities which are used for resource management, event counting, mutual exclusion locks, etc.

1.2.1 Task Management

Tasks in FreeRTOS can be regarded as occupying one of two top-level states, `running` or `notRunning`. The running task is recorded by the task control block handler `pxCurrentTCB` and simply indicates that the task is currently executing on the processor. The `notRunning` state can be decomposed into three sub-states: `ready`, `suspended`, and `blocked`. The following lists in FreeRTOS are used to manage this:

Ready Lists (`pxReadyTasksLists`) this is an array of the task lists, in which tasks are available to be scheduled to the running state.

Delay List (`xDelayedTaskList1`) & **Overflow Delay List** (`xDelayedTaskList2`) tasks in these lists are blocked by an event for a certain period. They are sorted by wake-up time. Because the time is expressed by ticks in FreeRTOS, if the wake-up time is later than the time represented by $maximum_delay_ticks - current_ticks$, the ticks for wake-up time could overflow. Therefore, an overflow-delay list is required.

Suspended List (`xSuspendedTaskList`) tasks in this list have been suspended, and wait until they are resumed by another task.

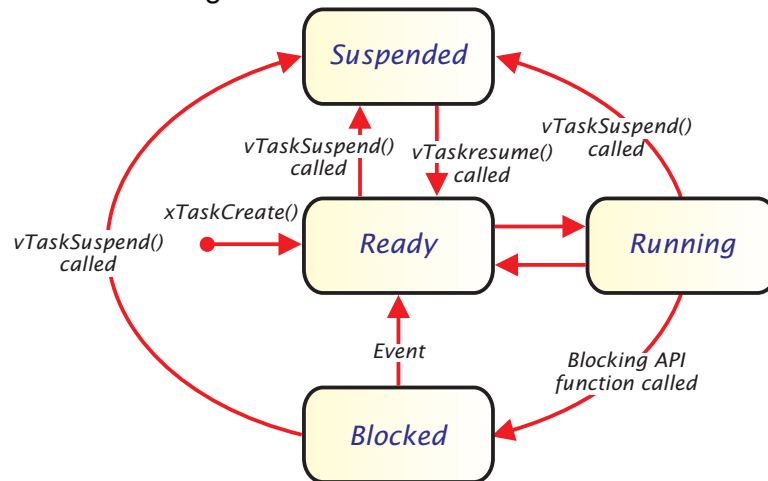
Pending Ready List (`xPendingReadyList`) tasks resumed from Interrupt Service Routines (ISRs) are kept in this list temporarily, while the scheduler is not running.

Waiting Termination List (`xTaskWaitingTermination`) deleted tasks stay here and wait to be removed by the `idle` task.

Tasks transit between these states as described in Fig. 1.1. For instance, a task cannot directly transit from `suspended` to `running`, because only `ready` tasks can be scheduled as `running` [18]

All tasks have their own priority, `uxPriority`, which is used by the scheduler. Tasks can have another priority, `uxBasePriority`, which records the original priority of

Figure 1.1: State chart for Tasks



tasks when priority inheritance occurs. The scheduler is responsible for counting the clock ticks, used to express time, and schedules the tasks. The scheduling policy adopted here is priority-based scheduling, which means that the task with the highest priority and in the ready state can be executed. As a result, it is impossible to use FreeRTOS in hard real-time environments. When a ready task has a higher priority than the running task, it will displace the running task from the CPU. The scheduler has two ways of switching tasks: pre-emptive and cooperative scheduling. In pre-emptive mode, the task with the highest priority will block the running task immediately and take the CPU. In cooperative mode, the running task can finish its CPU time before the task with the highest priority takes over. API functions are provided for task creation, deletion, and control. It is worth noting that the deletion API function does not actually delete a task from the system: it only adds the task to `xTasksWaitingTermination` and removes its reference from related task lists. The `idle` task, with permanent priority 0, the lowest priority, is used to do the deleting job and release the memory allocated by the kernel. However, it does not collect the memory allocated by the user, so tasks have to release used memory themselves, before being deleted.

In our specification, the function `state`, which is a total function from `TASK` to `STATE`, in schema `StateData` (see Sect. 4.2.1, Page. 43) is used to specify the states of tasks. Further, the reverse function of `state` can be used to cal-

Figure 1.2: An example application that uses RTOS (Task related).

```
1 xTaskHandle txh1 ;
2
3 void tx1(void * xPara){
4 xTaskCreate(tx3, (signed char *) "Task_3", 1000, NULL, 4, NULL);
5 for( ;; );
6 }
7
8 void tx2(void * xPara){
9 for( ;; ){
10 vTaskPrioritySet(txh1, 3);
11 }
12 }
13
14 void tx3(void * xPara){
15 for( ;; ){
16 vTaskDelete(NULL);
17 }
18 }
19
20 int main(void){
21 xTaskCreate(tx1, (signed char *) "Task_1", 1000, NULL, 1, & txh1);
22 xTaskCreate(tx2, (signed char *) "Task_2", 1000, NULL, 2, NULL);
23
24 vTaskStartScheduler();
25 return 0;
26 }
```

culate tasks in a specific state; for instance, `pxReadyTasksLists` can be represented by `state~(| {ready} |)`. This also works for running tasks: the result of `state~(| {running} |)` is a set with only one element—`running_task`, which represents the handler `pxCurrentTCB`. The function `priority` in schema `PrioData` and `old_priority` in schema `OriginalPrioData` of Mutex model represents tasks' `uxPriority` and `uxBasePriority` respectively. They are defined in Chap. 4.

We use a simple example application to illustrate the functionality provided by FreeRTOS. Fig. 1.2 shows the C code of an application that uses the FreeRTOS API function related to task management. Initially, the application creates two tasks: Task1 and Task2, with priority 1 and 2 respectively (a higher number indicates higher priority), and then starts the FreeRTOS scheduler. The scheduler then runs Task2, which immediately increases the priority of Task1 to 3. Task2 is now pre-empted by Task1, which gets to execute and creates a new task—Task3 with priority 4, which is the highest at the moment. Therefore, it pre-empts Task1 and can execute. Once Task3 is executing, it deletes itself, which triggers the scheduler to reschedule the system. As Task1 has the highest priority at this moment, it gets to execute again and will continue to execute.

We now describe in more detail what happens in the FreeRTOS implementation code. The application code for `main`, `tx1`, `tx2` and `tx3` is compiled along with the FreeRTOS code (for the scheduler and the API function calls including `xTaskCreate`) and loaded into memory. The scheduler code is loaded into the Interrupt Service Routine (ISR) code area so that it services software interrupts.

By analysing the source code of FreeRTOS, we see that execution begins with the first instruction in `main`, which is the call to the `xTaskCreate` API function. This code is provided by FreeRTOS. It allocates 1 kilobyte (defined in the parameters of `xTaskCreate`) of memory from the heap to the task stack, as well as space to store its Task Control Block (TCB) [19, 18]. From the source code, we see that the TCB contains all vital information about the task: where its code (`tx1` in this case) is located, where its stack begins, where its current top-of-stack pointer is, what its priority is, and so on. The API function call initialises the TCB entries for `Task1`. It then creates and initialises the various lists that the OS maintains, such as `pxReadyTasksLists`, `xSuspendedTaskList` and so on. It finally adds `Task1` to the ready list and returns. Next, `main` calls `xTaskCreate` for `Task2` and the API function call sets up the stack and TCB for `Task2` and adds it to the ready list, in a similar way. The next instruction in `main` is a call to the `vTaskStartScheduler` API function, which is also provided by FreeRTOS [19, 18]. This call creates the `idle` task with priority 0, and adds it to the ready list. It also sets the timer tick interrupt to occur at the required frequency. Finally, it does a context-switch to the highest priority ready task (i.e., it restores its execution state, namely the contents of its registers, from the task's stack where they were stored). The processor will next execute the instruction in the task that is resumed. In our example, this means that `Task2` will now begin execution.

When `Task2` begins execution it makes an API function call to `vTaskPrioritySet`. The code for this API function call compares the new priority and the current priority to decide whether scheduling is needed. If the API function increases the priority of a task or decreases the priority of the current running task, a reschedule will be requested. It then assigns the new priority to the target task, and moves the task to

the proper position in the ready list, if it is a ready task. In our case, the priority of Task1 is changed to 3, it is moved to the correct position in `pxReadyTasksLists`. The API function code then execute a `yield` (a kind of software interrupt) that is trapped by the scheduler. The scheduler picks the highest priority ready task, which in this case is Task1, and makes it the running task. Before this, the scheduler saves the registers of Task2 to its stack, and restores the register context of Task1 from its stack.

Task1 now creates the new task Task3. The process is similar to the `xTaskCreate` call to create Task1 and Task2. The difference is that here `xTaskCreate` triggers scheduling to make Task3 run.

When Task3 begins execution, it makes a call to the `vTaskDelete` API function. The code for this API function is simple. It removes the target task from the state list and related events list; in this case, Task3 is removed from `pxReadyTasksLists`. As it is the current running task, the API function code triggers scheduling again to make the highest priority ready task run, which is Task1. Task1 then executes its trivial `for`-loop, *ad infinitum*.

The animation and formal verification of our specification for this process will be illustrated in Chap. 7.

1.2.2 Communication and Synchronisation

In FreeRTOS, tasks and interrupts communicate and synchronise with each other through queues. When two tasks in FreeRTOS need to exchange information, they send and receive information to and from a queue. As items are exchanged between tasks and queues by being copied to or from a queue, the size of each item in the queue must be the same. Otherwise, when a task receives an item from a queue, it would be confused as to how many bytes needed to be received. Every queue fixes the size of all items it can receive using `uxItemSize`. All queues also have a capacity, `uxLength`, which indicates how many items can be held by the queue. The number of items currently stored in the queue needs to be

recorded as well. Tasks will be blocked while they attempt to send (receive) items to (from) a full (empty) queue. The following two sequences are used to manage this: (a) `xTasksWaitingToSend` records tasks blocked by sending operations; and (b) `xTasksWaitingToReceive` records tasks blocked by receiving operations. As well as these basic properties, a number of other fields are recorded also for a queue, such as, `pcHead` and `pcTail`, which represent where the queue starts and ends.

Semaphores and mutexes, which are used to manage resources, mutual exclusion locks and so on, are implemented by queues. They are considered to be special queues. Specifically, the item size for semaphores and mutexes is 0. This is because, when tasks take a semaphore or mutex, they do not copy items from semaphores and mutexes. What is of interest to the task which attempts to take the semaphore or mutex, is whether it is available or not. The initial state of semaphores and mutexes is full rather than of empty, which is the initial state for normal queues. The main difference between semaphores and mutexes is that the maximum length for mutexes is always 1; on the other hand, the size of semaphores can be any unsigned number. Moreover, mutexes support the priority inheritance mechanism when a higher priority task is waiting to take a mutex which is held by a lower priority task. Each mutex has its own mutex holder if it is taken by a task. The holder of a mutex can repeatedly take it at any time. Therefore, a mutex needs to know who is its holder. It overrides the field `pcTail` of normal queue to `pxMutexHolder` for this purpose.

In our specification, functions `q_max` and `q_size` in the `QueueData` schema of the `Queue` model (see Sect. 4.3, Page. 63) are used to represent the capacity and the current size of the queue respectively. Functions `wait_snd` and `wait_rcv` are used to indicate the blocked task for each queue in the system. Meanwhile, mutex related information is included in the `Mutex` model (see Sect. 4.5, Page. 79). For instance, the function `mutex_holder`, which represents the `pxMutexHolder`, is contained by `MutexData` schema of `Mutex` model. They are defined in Chap. 4.

Figure 1.3: An example application that uses RTOS (Communication related).

```
1 xTaskHandle tskH;
2 xSemaphoreHandle xMutex;
3
4 void tx1(void * xPara){
5 xSemaphoreTake(xMutex, portMAX_DELAY);
6 for( ;; );
7 }
8
9 void tx2(void * xPara){
10 vTaskDelay(10);
11
12 xSemaphoreTake(xMutex, portMAX_DELAY);
13 for( ;; );
14 }
15
16 int main(void){
17 xMutex = xSemaphoreCreateMutex();
18
19 xTaskCreate(tx1, (signed char *) "Task_1", 1000, NULL, 2, & tskH);
20 xTaskCreate(tx2, (signed char *) "Task_2", 1000, NULL, 3, NULL);
21
22 vTaskStartScheduler();
23 return 0;
24 }
```

Similar to Sect. 1.2.1, we use a simple example application (Fig. 1.3) to illustrate functionality related to communication in FreeRTOS. Initially, the application creates a mutex, `xMutex` and two tasks: Task1 and Task2, with priority 2 and 3 respectively and then starts the scheduler, which runs Task2. It requests to delay for 10ms, which blocks Task2 and lets Task1 execute. Once Task1 is executing, it takes the mutex `xMutex` then executes its infinite loop. After 10ms, Task2 wakes up. As it has higher priority than Task1, it preempts Task1 and starts to execute. Task2 also tries to take the mutex, `xMutex`. However, it has been held by Task1. Therefore, Task2 is blocked for `portMAX_DELAY` and Task1 can execute again.

In detail, `main` creates `xMutex` by calling `xSemaphoreCreateMutex` (Note, the operations, such as load code, create task, etc., which have been described in Sect. 1.2.1, are not repeated here). As declared in `semphr.h`, `xSemaphoreCreateMutex` actually executes the code of `xQueueCreateMutex`, which is defined in `queue.c` [19, 18]. This code allocates space to store the new queue structure (`xQueue`) and initialise the structure for `xMutex`. For instance, set the type to `queueQUEUE_IS_MUTEX`, set the holder of the mutex to `NULL`, set item size to 0, etc. Next, `main` creates Task1 and Task2 and starts the scheduler.

When Task2 begins execution, it makes a call to the `vTaskDelay` API function.

The code for this API function call will add `Task2` to an event list, which is a priority queue associated with the delayed tasks, with a value that corresponds to the current tick count plus 10. Then, `Task1` is scheduled as the running task.

Once `Task1` starts to execute, it calls `xSemaphoreTake` to take the mutex, `xMutex`. `xSemaphoreTake` is also declared in `semphr.h` and executes code in `queue.c`, which is `xQueueGenericReceive`. It checks whether there is an item available in the queue (i.e., `xMutex`). If there is, the calling task receives the item. Otherwise, it is blocked for a period which is specified by a parameter of the API function. As there is no task holding `xMutex` at the moment, `Task1` can successfully take the mutex. It then executes its infinite `for`-loop, until an interrupt for the next timer tick arrives from the hardware clock. This interrupt is again trapped by the scheduler and it increments its tick count. The scheduler then checks if any of the delayed tasks have a time-to-awake value that equals the current tick count. There is none and the scheduler hands back control to `Task1`. However, when the 10th timer interrupt takes place, the scheduler finds that `Task2`'s time-to-awake equals the current tick count, and moves it to the ready queue. Since there is now a higher priority ready task, `Task1` is swapped out and `Task2` is restored and made to execute. It then attempts to take the `xMutex`. As `Task1` holds the `xMutex` at the moment, `Task2` is blocked by the mutex. The event list item of `Task2` is added to the `xTaskWaitingToReceive` of `xMutex` as well. Therefore, `Task1` can execute again and stay in an infinite `for`-loop.

This process will also be animated and formally verified in Chap. 7.

1.2.3 Other API functions

FreeRTOS also provides an API function for other operations, such as memory management, interrupt management, etc. Memory management related API function calls can be used to allocate and free memory. When a task or queue needs memory, `pvPortMalloc` can be used to do this. It first locates one of the available memory blocks, and then returns its pointer to the task or the queue. To release memory, `vPortFree` can be used. Meanwhile, interrupt related API functions can

be used to serve interrupts, enter/exit critical sections and so on. Specifically, all the interrupts have a piece of server code, an Interrupt Service Routine (ISR). When the operating system services the interrupt, it cannot accept another interrupt. Furthermore, as a real-time operating system, some parts of the code may be critical, which means they are unable to be interrupted. When the program enters this section of code, the counter `uxCriticalNesting` would be increased and `portDISABLE_INTERRUPTS` is called to set a processor flag to refuse further interrupts. When it exits the critical section, the counter is decreased. At this time, the value of the counter will be checked. If it is greater than zero, the processor flag remains the same, refusing further interrupts. Otherwise, if and only if it decreases to zero, `portENABLE_INTERRUPTS` can be applied to reset the flag to enable interrupts.

1.3 VCC

The Verifying C Compiler (VCC) was developed by Microsoft for the Hypervisor Verification Project [20]. It verifies the correctness of annotated C programs. Annotations used for VCC include function specifications, state assertions, type invariants and so on [21]. As described above, normally, there are several steps of refinement and verification from the abstract model to the concrete specification and the executable code. Sometimes these are difficult and expensive to perform. Using the specifications from verified abstract models to directly verify the C code can be interesting and efficient.

Using Microsoft Visual Studio (MVS) for VCC is recommended. With macro definitions, the normal C compiler in Visual Studio can ignore the annotations used by VCC. On the other hand, the VCC verifier may use the C code and annotation together. They are translated into Boogie [22] files. These files are then used to generate *.sx files, which can be used by the Z3 prover [23]. VCC translates the C code and the annotation to mathematical formulas and verifies them using the Z3 prover, rather than analysing the code and looking for bugs. Once a piece of code has been verified by VCC, its correctness with respect to the preconditions

and post conditions can be guaranteed. Verification in VCC is modular. It does not go through every function call to verify the code of a function. Instead, it verifies a function with the information of a function contract for each called function. VCC assumes that the called functions are correct. In this case, it only needs to verify that when the function call happens the system state satisfies the preconditions of the called function. If it does, VCC knows that the post condition of the called function is satisfied as well. With this feature, developers can verify a function, even when its sub-functions are not finished or verified. In addition, with the benefits of the Z3 prover, VCC provides the Model Viewer which shows an example for each failure, when Z3 fails to verify the code. These examples contain a sequence of the system states which lead to the failure. This is helpful for the developer to understand why the code failed to verify.

1.4 Multi-core Processor

The Central Processing Unit (CPU) is the core component of the computer. Its performance determines the performance of the whole computer system. Therefore, the hardware industry has continued to try to improve the performance of the processor. There are two common ways to achieve this:

- (a) Increasing the number of transistors on the chip. The Intel 4004, the first micro-processor built in 1971, had 2,300 MOS transistors [24]. According to Moore's Law, the number of transistors on a single chip will double approximately every two years [25]. Thus, after around 50 years development, it is now possible to put more than 500 million transistors on a single chip, e.g. the Intel i7-680UM Processor [26]. Due to the large number of transistors, an increasing number of resources are now available on a chip and processors have become progressively more powerful.
- (b) Increasing the clock rate of the processor.

To use and control the resources on a single core processor efficiently, a large number of complex circuits have been designed and used. Because of this, design

and verification for a processor based on traditional single core architecture is increasingly difficult. Bose *et al.* have reported that verification activities can take up around 70% of the net development cost [27]. Meanwhile, increasing the clock rate is one of the direct ways to improve the performance. Nevertheless, this is limited by the physical features of the processors, power consumption and related thermal problems, which have also become ever more critical [27]. After every pulse of the clock, each transistor needs to take some time to transfer to a new state. If a clock pulse occurs before that, the data and the state of the processor will be incorrect, which is unacceptable. Furthermore, a higher clock rate means a higher power cost. Taylor *et al.* have indicated that it would increase power consumption by the increase in clock rate cubed [28]. In addition, the thermal issues also increase with energy expended. Due to these issues, speeding up the clock frequency to obtain higher performance has reached a bottleneck.

These issues are especially serious for embedded systems. This is because embedded systems are designed for a small number of dedicated functions [29] and they normally work in mobile systems and/or critical and real-time systems, such as sensor controllers and car control systems, etc. Due to the function of embedded systems, it is impossible to provide them with unlimited power.

Multi-core processors, which are composed of two or more independent cores on a chip, seem an alternative way to solve the problem. Multi-core processes are:

- (a) Easy to design – due to parallel computation, multi-core processors can use several smaller and more simple cores to achieve a higher performance than a huge, complex core. Therefore, the designer only needs to repeat the simple cores across the chip and focus on the design of the communication method between separate cores, such as on-chip networks, bus and so on. Intel Research [30] reports that to design a single core chip with 100 million transistors would take about twice as long with twice as many people than a multi-core processor with the same number of transistors.
- (b) Energy efficient – because multi-core processors can separate the task into

independent subtasks and share them with different cores on the chip, they gain high performance with a lower clock rate. As described above, the raising of the clock rate would lead to higher power-consumption and related thermal problems. With the reduction of the clock rate, these issues would be handled. For instance, the power of the Intel[®] Pentium[®] 4 Processor 531, which is a single core processor with a 3.00 GHz clock rate, is 84W [31]. However, the Intel Teraflops Research Chip, which is a multi-core processor with a 3.16 GHz clock frequency, consumes only 62W of power [30].

- (c) Scalable – because the multi-core processor repeats the simple core across the chip, it is possible to place as many cores as the limits of the technology. In the laboratory, processors with 1,024 cores on a chip are now available.

Due to these benefits, an increasing number of companies use multi-core architecture.

1.5 Objectives and Challenges

1.5.1 Objectives

Our aim is to carry out a systematic exercise towards the verification of FreeRTOS that will:

- (a) Produce a formal specification of its intended behaviour.
- (b) Produce an annotated version of the implementation for VCC to verify.
- (c) Identify aspects of its implementation that do not conform to this specification.
- (d) Produce a detailed model of the core scheduling-related functionality that can serve as a basis for fixing the current implementation to obtain a “verified” version of FreeRTOS, engineered as originally intended by the developers.
- (e) Produce an abstract model for a multi-core platform, which is an extension of FreeRTOS.

1.5.2 Structure View

Fig. 1.4 illustrates the whole structure of the project. To achieve these objectives defined above, three tools (Z/Eves, VCC and ProZ) are used in the project,

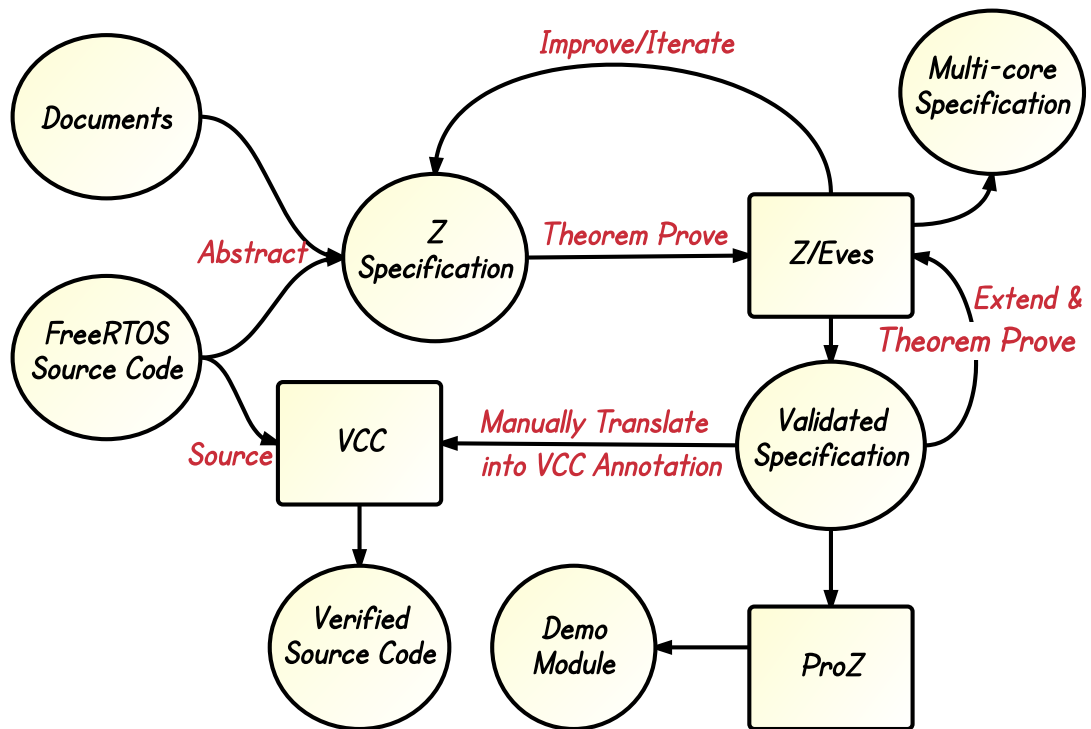
which are represented by rectangles in the figure. Specifically, VCC, introduced in Sect. 1.3, is used to verify the correctness of the annotated FreeRTOS source code, based on the function contracts derived from validated Z specification. Next, Z/Eves is a theorem prover [32], which we use in our project for analysing our Z specifications, checking syntax and proving theorems automatically with some help of human. Finally, ProZ [33] is used to animate the validated Z model. It demonstrates the behaviour of the software described by a Z specification. Sect. 2.2.1 describes Z/Eves and ProZ in detail. Furthermore, inputs and outputs of the tools are represented by circles, including documentations, FreeRTOS source code, Z specifications, etc. And the relations between the tools and their inputs and output are shown by arrows. In detail, as shown in Fig. 1.4, from the documentation [18, 19, 34] and source code of FreeRTOS, a basic version of the requirement is obtained by abstracting the documentation and reverse engineering the source code, which is described in detail in Chap. 3. Then, the first version of the Z specification is specified, based on the requirement, which is validated with the theorem prover, Z/Eves. After several rounds of iteration and improvements, the final version of validated specification is produced, which can be directly used by ProZ for animation. The validated specification and the iteration process is explained in Chap. 4. Afterwards, inspired by Multi-BSP model [35], the final specification is extended to multi-core platform and validated using Z/Eves (See Chap. 6) as well. Meanwhile, we manually translate the validated specification into VCC annotation, so, with FreeRTOS source code, the related source code can be verified (See Chap. 5).

1.5.3 Challenges

The main scientific difficulty with the verification of FreeRTOS is the low level of the code. The usual abstractions that make it easier to program systems software do not exist; it is the purpose of FreeRTOS to provide them. They include:

- (a) Communication and synchronisation.
- (b) Scheduling guarantees.
- (c) Interference freedom.

Figure 1.4: Overview of project



(d) Direct hardware interaction using clocks and interrupts.

These are provided through a number of complex pointer-based operations, which present yet another challenge: verifying pointer programs is a complex and difficult business.

1.6 Font and Name Styles

In this thesis, text from specification and implementation are distinguished by different font styles. In detail,

- File names of the source code are represented by italic font, e.g. *task.c*;
- Teletype font is used for text from FreeRTOS source code and VCC annotation, e.g. `pxCurrentTCB`;
- For text and formulas from specification, naturally, the mathematical font is used for them, $Task$.

Table 1.1: Prefix of variable and function names used in FreeRTOS

Prefix	Meaning
x	Non-standard integer, e.g. portBASE_TYPE, portTickType, xTaskHandle, etc.
v	void return type.
ux	Unsigned non-standard integer type, e.g. unsigned portBASE_TYPE
e	Enumerated type, e.g. eTaskState
prv	Private functions, e.g. prvDeleteTCB

Table 1.2: Suffixes used in schema names

Task Model	Queue Model	Time Model	Mutex Model
<i>_T</i>	<i>_TQ</i>	<i>_TQT</i>	<i>_TQTM</i>

In addition, as there are four levels of models in the specification, each operation may have more than one version of schema for different level of the model, for instance, the operation creating task has schemas for all four models. We use suffixes in Table. 1.2 to distinguish these schemas, e.g. *CreateTask_T* is the schema name for creating task in the task model and *CreateTask_TQ* is the schema name for creating task in the queue model. Furthermore, each operation may also be composed of different sub-operations. For example, depending on whether the operation requests rescheduling, creating task has two sub-operations. We append a suffix to the end of the operation name to indicate them, e.g. *N* and *S* used respectively in *CreatTaskN_T* and *CreatTaskS_T*.

Similarly, FreeRTOS uses prefixes of variable and function names to indicate the type of the variables and the return type of the functions. For instance, the prefix *x* shows that a variable, or the return type of a function, is a non-standard integer, e.g. portBASE_TYPE. There are two exceptions for the prefixes of the function names. Specifically, the prefix *v* indicates that the return type of a function is `void` and the prefix *prv* shows that a function is a private function. Table. 1.1 summarises some of the frequently used prefixes in FreeRTOS and the meaning of them. More details can be found from [36].

1.7 Structure of Thesis

The remainder of this thesis is divided into seven chapters:

Chapter 2 presents a review of the literature on formally specifying operating systems and related work in this field. It also reviews the tools used in this research.

Chapter 3 analyses the API functions of FreeRTOS and abstracts the requirements for each API function.

Chapter 4 describes the formalisation work for FreeRTOS and explains the specifications and theorems in detail. The task model shown in this chapter has been published in the journal *Formal Aspects of Computing* [37].

Chapter 5 illustrates the VCC verification for the API functions related to the task model of specifications.

Chapter 6 shows the formalisation work for the multi-core version of FreeRTOS, which is extended from the outcome of the previous chapter. The requirements are analysed and presented at the beginning. The extension to the specification for the multi-core version of FreeRTOS is then covered in the rest of the chapter.

Chapter 7 evaluates the research project in different ways. It firstly summarises the achievements of the research. It then also describes a carefully designed case study to show how the specification would work properly using the ProZ animator and the Z/Eves theorem prover.

Chapter 8 summarises the work done in this research project. Further, it reports the experience gained from undertaking the research process. Finally, suggestions for future work are presented.

Appendices introduces the structure of the following appendixes firstly. It, then, summarises the proof commands (See Page.157) frequently used in proofs of our model and the preconditions for the API functions of FreeRTOS (See Page. 159), which is followed by the specification of our FreeRTOS model and multi-core model, each of which is divided into four separate appendixes according to its sub-model structure. It also presents the specification for

multi-core task model with promotion. Finally, the annotated source code related to our VCC verification is listed.

Chapter 2

LITERATURE REVIEW

This chapter reviews the literature related to the research. It begins with a detailed discussion of FreeRTOS using examples to show how the API function is used to implement various applications. Z formal notation, which is adopted in this research, is then introduced, with related tools also described. Finally, the related research on formally verifying operating systems is examined.

2.1 Related Work

This section presents the existing work on verifying operating systems:

1. Craig describes the specification and refinement in Z notation of, Labrosse's μ C/OS operating system, a microkernel which is similar to FreeRTOS [38, 39, 40]. The refinement of the requirements targets mathematical data types at a level of abstraction well above program data types. The lowest level of refinement is also non-algorithmic and there are no real-time properties. Freitas & Woodcock [41] have continued Craig's refinement to target datatypes at the level of FreeRTOS, but without pointer implementation. Börger & Craig [42] also extend this work, modelling with pseudo-code descriptions as Abstract State Machines (ASMs), which produces an elegant restructuring of the model that makes it easier to understand and easier to refine into executable code.

2. Klein verifies seL4, a high-performance microkernel [43, 44, 45, 46, 47]. An abstract specification in Isabelle/HOL is refined into an executable specification in Haskell, which is then manually refined into a high performance implementation in the C programming language. The theoretical basis for the work is in separation logic. There is an almost complete handling of the features of seL4. The entire exercise involved 8,700 lines of C, 200,000+ lines of proof script, and 30 person-years of effort to establish the functional correctness of the operating system.
3. Buerki & Rueegsegger [48] introduce their design and implement a high assurance microkernel called Muen. They suggest that assuring the whole microkernel system is very difficult, but also unnecessary. However, the microkernel can be separated into different function blocks, some of which are critical. These are defined as the Trusted Computing Base (TCB) and it can lead to a fatal system error if the TCBs misbehave. Therefore, the set of TCBs can be treated as the smallest set of functions which are essential for verification.
4. Déharbe et al. have produced a specification in the B language of a restricted subset of FreeRTOS [49]. They provide a formalisation of a subset of the API function, verifying that all its expressions are well defined and demonstrate logical consistency. This model includes Task and Queue-related functions. The model contains seven basic B-machines, *FreeRTOSConfig*, *Type*, *Task*, *Queue*, *Scheduler*, *FreeRTOSBasic*, and *FreeRTOS*, with which the first model without priority is formalised. It is then refined to the second model, which takes priority into account. However, there are problems with this model; for instance, it prevents task creation while the scheduler is running, which is allowed by FreeRTOS. It also forbids tasks from sending and receiving messages to and from a queue when there is no task waiting to receive or send. Comparing this with our work, we introduce a model covering more functions of FreeRTOS. Due to a finer structure of definitions and abstractions, our specification has increased proof automation. Furthermore,

we correct the problems discovered in Déharbe et al.

5. Pronk looks at the verification problem for FreeRTOS [50]. He discusses and compares the advantages and disadvantages of theorem proving and refinement in this arena compared with model checking. He concentrates on the latter, using Promela and the SPIN model checker.
6. Lin, Freitas, & Woodcock produced a specification of FreeRTOS in Z covering the top-level functionality [51]. This was derived from Déharbe's B specification [49] (see above item 4), which was then extended to capture all the main FreeRTOS functionality. An attempt was made to verify the consistency using the Z/Eves theorem prover [32], although this could also be proved in ProofPowerZ [52], Isabelle/HOL [53], or PVS [54]. Originally, there were 30 unproven theorems out of 241. We have carried out further work to reduce the number of unproven theorems to around 10. During this process, we found the key reason for struggling with proofs is that the model is too concrete, which leads to proof complexity. For example, to represent the different states in FreeRTOS, the model uses seven different type variables, such as functions, sequences, finite sets, etc. Furthermore, even when Z/Eves can prove a theorem, it takes considerable time. Compared with this, our model is much more abstract and more tractable for proof.
7. Mühlberg & Freitas report on the application of the SOCA and VeriFast tools to FreeRTOS [55]. They focus on the verification of structural properties (e.g., pointer safety and arithmetic overflow) and liveness properties, but ultimately aim at demonstrating functional correctness. This includes the reconstruction of a formal specification of FreeRTOS in Z (mentioned above in item 6), bounded model-checking of the FreeRTOS code using the SOCAVerifier [56], as well as annotating the source code with assertions in separation logic to apply the VeriFast software verifier [57].
8. Ferreira uses separation logic to verify code-level pointer structures in FreeRTOS [58].

9. Abrial [59] has an unpublished specification of much of the functionality of FreeRTOS using the B method (excluding interrupts). The Z specification in this thesis is based on his work, although the verification is necessarily very different.
10. Mistry, Naylor, & Woodcock have developed a multi-core version of FreeRTOS on a Field Programmable Gate Array (FPGA), which is able to schedule tasks on multiple processors and support mutex in a concurrent environment [60, 61]. They present an adapted version of FreeRTOS that is able to schedule tasks on multiple processors, as well as provide full mutual exclusion support for use in concurrent applications, which is independent of the chosen platform, thus preserving one of FreeRTOS's most attractive features, portability.
11. In collaboration with the author of this thesis, an unpublished work from Kushwah, Divakaran, & D'Souza aims to give a proof of functional correctness by proving that the C implementation refines the abstract Z specification. The commonality with this work is that they also focus on the task-related functionality of FreeRTOS. The points of difference are that their specification is deterministic, more detailed, and closer to the implementation than ours. In addition, they do not check consistency or prove properties for their Z model.
12. Based on the previous work (item 11), Divakaran et al. [62] also attempt to use VCC to verify the implementation of FreeRTOS with abstract Z specification. The difference between these works is that they focus on checking refinement steps. They promote three approaches (*“Direct-Import”*, *“Combined”* and *“Two-Step”*) in VCC to check refinements between the abstract model and the implementation. By comparing and applying them to the case study, they claim that the *“Two-Step”* approach is much more efficient than the other two, which splits the verification process into two steps: (a) The behaviours of the function defined in the abstract model and the concrete function are verified; (b) Then, the outputs from both functions satisfying the

gluing invariants is checked. Furthermore, they also suggest how to translate a subset of the Z notation to VCC.

2.2 Z Notation

Z, developed by Sørensen [63] in 1982, is a formal notation and became an international standard in 2002 [64]. It is based on set theory and mathematical logic. Specifically, the Z notation uses set operators, set comprehensions, Cartesian products and power sets. The logic part uses first-order predicate calculus [65]. With these mathematical theories, Z can describe the state and properties of the system being specified. Z uses an abstract structure called a Schema to capture a number of concepts in one named block. Subsequent schemas that refer to the name of a previous schema can inherit all the concepts and constraints introduced in that schema [10, 63, 65, 66]. This provides reusability in the Z notation. In order to manage system complexity, schemas are vital to keep the specification flexible and manageable in a real, large-scale project [65]. Moreover, a schema can also be used to specify the behaviour of the system.

Hoare [67] introduced Hoare Triples in 1969, which describe the logical behaviour of a computer program. This triple can later be verified by related logic theories and lemmas. Specifically, the notation

$$P \{Q\} R$$

can be used to express that if the precondition P is true before the program Q is initialised, then the postcondition R will be true once Q terminates successfully. Therefore, a schema is composed of two parts, a precondition and a postcondition. The default relation between predicates in different lines of a schema is called logical conjunction \wedge .

For instance, the following shows part of a specification for a system. It is used to record the relationship between tasks and their priority. (Note, this is just a example to illustrate how schemas work; the definition here is not the same in our specification of FreeRTOS.) Firstly, $TASK$ is defined to represent tasks in the world using a given set, which is used to introduce uninterpreted domain-specific types in Z.

[*TASK*]

Subsequently, a schema, *Task*, is introduced to specify the basic abstract state of the system. It includes two component declarations: (a) *tasks*, which is a subset of *TASK*, indicates the tasks in our system; and (b) *priority*, which is a partial function from *TASK* to \mathbb{N} , illustrates that tasks may have a priority. As well as the declaration, the constraint for these properties is also defined, indicating that the domain of *priority* is *tasks*. This means tasks in set *tasks* have a priority and tasks not in *tasks* do not.

<i>Task</i>
<i>tasks</i> : $\mathbb{P} \text{TASK}$ <i>priority</i> : <i>TASK</i> $\leftrightarrow \mathbb{N}$
dom <i>priority</i> = <i>tasks</i>

Based on state schemas, operation schemas can be defined, such as *setPrio* below, which specifies the operation for setting the priority of a task. It refers to the *Task* schema, to obtain all the essential information about the system we defined. Therefore, the variables defined in *Task* can be directly used.

<i>setPrio</i>
ΔTask <i>t?</i> : <i>TASK</i> <i>prio?</i> : \mathbb{N} <i>out!</i> : \mathbb{N}
<i>t?</i> \in <i>tasks</i>
<i>tasks'</i> = <i>tasks</i> <i>priority'</i> = <i>priority</i> $\oplus \{(t? \mapsto \text{prio?})\}$ <i>out!</i> = <i>priority'</i> (<i>t?</i>)

The declaration part of the schema starts with ΔTask , which refers to the pre- and post-state of the *Task* schema and indicates that it is going to change the state *Task*. Following this, it introduces some inputs (e.g., *t?*) and outputs (e.g., *out!*). They are distinguished through a decoration convention. Variables in the pre-state (e.g., *tasks*) and post-state (e.g., *tasks'*) are similarly distinguished. This schema has only one precondition, *t?* \in *tasks*, indicating that before the operation

setPrio the target task $t?$ is in the system. Therefore, it is in the domain of *priority* as defined in *Task* schema. Consequently, after the operation has successfully finished, the new priority of $t?$ is updated in the function *priority'*. The override operator, \oplus , is used to achieve this, which is the most common way to update functions in the Z notation. If the first element of the pair specified in operation (e.g., $t?$ in this case) exists in the domain of the function (e.g., *priority*), it would update the result of function application (e.g., *priority*($t?$)) to the new value (e.g., *prio?*); otherwise it adds the pair into the function. Finally, the new priority of target task $t?$ is set as output, which should simply equal to *prio?*.

More explanation of the Z notation will be provided in Chap. 4 along with the explanation of our specification.

2.2.1 Tools for Z

A leading proof tool for the Z notation is Z/Eves, which can be used for analysing Z specifications, syntax checking and most importantly theorem proving. Using Z/Eves, specifications can be entered by importing \LaTeX source code [68] or by typing directly from the editing facility provided by Z/Eves [69]. Once the specification has been entered, Z/Eves can automatically perform syntax checking, type-checking, and some basic proving, by double clicking on paragraphs, such as schemas. This is very useful for users to avoid syntax errors. Furthermore, based on self-defined theorems or lemmas and the built-in theorems, which can be found in [70], Z/Eves can be used to prove specifications automatically. More details about how to use Z/Eves can be found in the user guide [69].

ProZ, which is extended from ProB [33], can be used to animate and check Z specifications. It is worth noting that each specification animated by ProZ should have one and only one schema named *Init*, which defines the initialisation state of the model. Once the model is initialised, based on the current state of the model, available operations will be shown and the user can apply them to the model by simply double clicking on them.

2.3 FreeRTOS

FreeRTOS is a widely used real-time operating system written by a team led by Richard Barry of Wittenstein High-Integrity Systems in the UK [19]. Introductions to FreeRTOS informally describe the application programming interface (API) for the real-time operating system kernel [18]. Verifying the correctness of FreeRTOS has also been proposed as a pilot project for the VSI. This verification experiment presents two distinct challenges: (a) Code-level verification to automatically analyse FreeRTOS for structural integrity properties; and (b) The creation of a rational reconstruction of the refinement of the FreeRTOS code starting from an abstract specification, discharging all verification conditions automatically. This project was chosen as a contribution to the VSI at a workshop held at Microsoft Research, Cambridge, in 2008, that was gathering difficult research problems from industry. Modelling and verifying operating system kernels is considered to be scientifically interesting, pushing the current capabilities of software verification research and technology. Klein is the first to formally verify an operating system kernel and describe the main scientific challenges [43, 71].

FreeRTOS has a large community of users programming embedded microcontrollers: it was downloaded 107,000 times in 2013, putting it high in the top 100 SourceForge codes (there are more than 200,000 available). Verification of FreeRTOS, which allowed the discovery of residual errors, would thus have a strong impact on the international embedded system community.

FreeRTOS is a lightweight, embeddable, multi-tasking, Real-Time Operating System (RTOS). It makes the key assumption that the target system has a single processing unit. It is really a library of types and functions that can be used to build microkernels using a combination of C and assembly language, and has been ported to most embedded systems architectures. It allows a very small kernel to be produced to target microcontrollers, somewhere between 4–9kB. In some special cases, it can be less than 4kB. For instance, it takes less than 4kB of RAM, when creating 13 tasks, 2 queues and 4 software timers for RL78 [34]. It provides

services for embedded programming tasks, communication and synchronisation, memory management, real-time events and I/O-device control.

Fourteen different compilers are used with FreeRTOS, giving complex configuration options and extensive parametrisation. A version of the software, SafeRTOS, has been certified to Safety Integrity Level 3 by the Technical University of Vienna for the following safety standards: IEC 61508, FDA 510(k), and DO-178B. These certificates are for the process of development, rather than for the correctness of the software against stated requirements.

The objective of formally verifying FreeRTOS would be to find any errors and make some guarantees about the code's behaviour. Since the requirements are distributed throughout the documentation, there is a clear need to produce a formal abstract specification. A broader aim of our work is to study the verification problem for an entire class of software, namely *real-time operating systems for embedded applications*, and we have chosen to focus on an exemplar of this class of system namely the FreeRTOS kernel. The techniques and methodology developed here can be expected to be applicable to other software in this class of system.

2.4 Summary

This chapter firstly reviewed previous work related to formalising operating systems, which is helpful for us to understand the background of the project. Secondly, it discussed the principle of formal verification and Z used in the project. In addition, it reviewed the target system of the project, FreeRTOS.

Based on this, the next chapter will define the goals and scope of the project and most importantly the requirements of the project can be abstracted.

Chapter 3

ANALYSIS AND ABSTRACT API FUNCTIONS OF FREERTOS

This chapter discusses the requirements for the research. Because the development process does not completely follow either a formal methods strategy or reverse engineering strategy, it begins by giving an overview of our model. It then discusses the goals and scope of the research, followed by the requirements analysis.

3.1 Model Overview

Normally, the development process in formal methods starts with the requirements, modelling the behaviour of the system, and refining through several steps to executable code. As FreeRTOS does not have explicitly articulated requirements and has been implemented in C, we consider the user manual and the practical guide to FreeRTOS [19, 18] the basis of the requirements. However, these sources are not detailed enough for us to build the model; they just provide a basic functional description of the API functions. Therefore, we also take the FreeRTOS source code into consideration for modelling. Thus, the model is mainly based on the API documentation to verify the functional correctness of FreeRTOS, with some of the details of the specification derived from the source code. For instance, the

`xTaskCreate` API function documentation states that if the API function returns `pdTRUE`, then the task has been created. (NB, to simplify the model, we just consider the successful case of API functions.) Nevertheless, it does not indicate how it was created. Therefore, we analysed the source code to find out how it works and formalised the behaviour of `xTaskCreate` based on that.

As described in the last chapter, the key elements of FreeRTOS can be divided into three categories: (a) task management; (b) communication and synchronisation; and (c) memory and interrupt management. However, functions related to memory and interrupt management are quite hardware-dependent, so we abandon them at this level of abstraction. Furthermore, according to whether they are time dependent, task API functions can be divided into pure task operations or time-related operations. Therefore, we focus on task, communication (i.e., queue), time and synchronisation (i.e., semaphore and mutex) related FreeRTOS API functions. Each of them is reasonably independent. Therefore, our modelling starts from the core part of the system, *Task* model, and then expands to cover other features. For each subsystem, we attempt to keep the model as simple as possible. This can significantly reduce the difficulty of modelling and verification. In detail, we first build the *Task* model. Based on this, the *Queue* model is added into the system, followed by the *Time* and *Mutex* models.

3.2 Goal and Scope

The goal of this research is to provide a verified high level abstract formal model for FreeRTOS, which can be used as the foundation for future research (i.e., refinement, extension, etc.). It describes the behaviour of FreeRTOS API functions. Based on FreeRTOS's manual [19], 61 API functions are provided to the user. However, some of them have similar functions, for instance, `xQueueReceive` and `xQueuePeek`. Both of them attempt to receive an item from a queue, but `xQueueReceive` actually receives data from the queue and `xQueuePeek` just checks if there are any items available in the queue. Moreover, due to the abstract level of the specification, some API functions perform the same function, for instance,

xQueueSend, *xQueueSendToBack*, *xQueueSendToFront*. All of these try to send an item to a queue. The difference is that with *xQueueSend* a user can send the item to either the front or the back of the queue and the latter two API functions can only send the item to the back and the front of the queue respectively. At this level of abstraction, we do not need to consider the detail of the order of items in a queue or how items are stored in a queue. Therefore, we only selected 15 of the API functions for our model (listed in Sect. 3.3.1). To simplify the modelling and verifying further, we assume that the system scheduler is continuously running and only focus on cases where the API functions succeed.

Due to the complexity of the traditional refinement process and time limitations, we have not refined our model into executable code. However, we attempt to use VCC and our model to verify task-related FreeRTOS API functions to illustrate that the VCC kind of verifier plus abstract specification can be an alternative approach to verification. In addition, we also extend this model for multi-core platforms.

3.3 Requirements

Generally, the use case diagram, which shows the relation between the actor and use cases, and the relation between different use cases, is the most common approach to describing requirements. It is also used as a fundamental document for further software development. However, in our project, FreeRTOS has been implemented already and does not have a proper set of requirements. Therefore, we summarise requirements for FreeRTOS according to its API function.

3.3.1 Functional Requirements

1. Task Related:

- 1.1 *xTaskCreate* Create a task and specify its priority, reschedule tasks when the priority of the new task is greater than the running task;
- 1.2 *vTaskDelete* Remove a non-idle task from the system, reschedule tasks if the deleted task is the running task;

- 1.3 *vTaskSuspend* Set the state of a task to suspended, reschedule tasks if the suspended task is the running task;
- 1.4 *vTaskResume* Set the state of a suspended task to ready, reschedule when the priority of the resumed task is greater than that of the running task;
- 1.5 *vTaskPrioritySet* Set new priority of a task (N.B. the priority of `idle` task is always 0), reschedule tasks:
 - if the priority of the running task is set lower than the priority of the highest-priority ready task;
 - if the priority of a ready task is set greater than that of the running task.

2. Queue Related:

- 2.1 *xQueueCreate* Create a queue and specify its size;
- 2.2 *vQueueDelete* Delete a queue;
- 2.3 *xQueueSend* Send an item to the queue;
 - block this task if the queue is full, add it to waiting send list;
 - if the queue is empty, wake up the highest-priority task, which is waiting to receive an item from the queue.
- 2.4 *xQueueReceive* Receive an item from the queue;
 - if the queue is empty block the task and add it to waiting receive list;
 - if the queue is full, wake up the highest-priority task, which is waiting to send an item to the queue.

3. Time Related:

- 3.1 *vTaskDelayUntil* Block current running task until the specified time, reschedule the highest-priority ready task as the new running task;

4. Semaphore & Mutex Related:

- 4.1 *vSemaphoreCreateBinary* Create a binary semaphore;

- 4.2 *vSemaphoreDelete* Delete a binary semaphore or a mutex;
- 4.3 *xSemaphoreCreateMutex* Create a mutex;
- 4.4 *xSemaphoreTake* Take a token from a semaphore or a mutex;
- 4.5 *xSemaphoreGive* Give a token back to a semaphore or mutex;

In addition, the following two time properties are related to task scheduling, which are also of interest to us:

1. The function *prvCheckDelayedTasks* checks expiry time for blocked tasks. When time increases, it checks if there are any blocked tasks that need to be woken up. If there are, it moves them to the ready state, which may cause rescheduling;
2. When the system increases the ticks counter, which represents the time, it also checks whether there are any ready tasks with the same priority as the running task to share the processor.
 - If there is more than one ready task sharing the highest priority, they need to share CPU time as well. Rescheduling is required in this case.
 - Otherwise, the current running task keeps running.

3.3.2 Non-functional Requirements

Non-functional requirements specify the constraints on the services or functions offered by the system. In our case, the non-functional requirements can be summarised as:

1. Well-definedness. The specification should be well-defined;
2. Animatable. The specification should be able to be animated by ProZ.
3. Feasibility. The specification should be feasible (i.e. initial state and precondition for each operation should be reachable);
4. Reproducible. The specification and verification should be easily reproduced by other users.
5. Reusable. The specification should be able to be reused and expanded easily with little or no modification;

3.3.3 Environment Requirements

Z/Eves and ProZ are used as the prover and animator during the project. To reproduce the experiment, they are essential. Additionally, Community Z Tools (CZT) [72] can be very helpful for modifying the specification source code. CZT also integrates an interface for Z/Eves, as the original graphic interface of Z/Eves is implemented in Python, which crashes easily in Windows 7. Further, FreeRTOS v7.1.1 is used for the project.

3.4 Summary

In this chapter, the first section showed how we abstract the requirement for FreeRTOS, following which the second section defined the goal and scope of the project. Finally, the last section discussed the requirements of the project in three parts: functional requirements, non-functional requirements and environment requirements. Specifically, the functional requirements described the requirements for FreeRTOS API functions we modelled in the project, which were divided into four categories.

In the next chapter, following the categories of the functional requirements, the abstract model of FreeRTOS, which is the core of the project, will be described in detail. In addition, the experience gained from the modelling process will be discussed, which is helpful for developers specifying large systems, like us.

Chapter 4

MODELLING FREERTOS

*This chapter describes the model in detail, following the structure described in the previous chapter. Firstly, it shows how we approached the model. The *Task* model is then described, being the simplest and most important part of the model. In this section, some auxiliary theorems, which are helpful during proving and modelling are also explained. Following this, the *Queue* model is illustrated, followed by the *Time* and *Mutex* models. For each of these models, we also briefly explain how the API functions of the previous model are expanded. In addition, we collect the preconditions for each API function and some properties of the system. We also give a summary of the proof commands in Appendix B, so that the reader can follow the general argument behind the formal proofs or even recreate the proof in Z/Eves.¹*

4.1 Iteration Process

Following the requirement and the refinement strategy of [59], we started the modelling process with the simplest and the most important part of the system, which includes creating tasks, deleting tasks and rescheduling. Basically, the idea of creating a task is adding a task, which does not belong to the system, to the system. To describe this, besides the given set (*TASK*) representing tasks, we need

¹Z/Eves project file and other related files can be found in supplementary material.

a set (*tasks*) to help us to distinguish the tasks known by the system and others. Similarly, to delete a task, we simply remove the task from the system, i.e. remove the target task from the set *tasks*. Finally, to define the behaviour of rescheduling, the simplest idea is setting the target task as the new running task, and performing a context switch. The variables, *running_task*, *log_context* and *phy_context* are defined for this purpose. Therefore, we specified the first and tiniest model of FreeRTOS, which contains only one base schema called *FreeRTOS*.

$ \begin{aligned} & \textit{FreeRTOS} \\ & \textit{phy_context} : \textit{CONTEXT} \\ & \textit{running_task} : \textit{TASK} \\ & \textit{log_context} : \textit{TASK} \rightarrow \textit{CONTEXT} \\ & \textit{tasks} : \mathbb{P} \textit{TASK} \\ \\ & \textit{running_task} \in \textit{tasks} \\ & \textit{idle} \in \textit{tasks} \end{aligned} $

The schema *FreeRTOS* contains all the necessary variables described above to describe creating, deleting and rescheduling tasks. Based on this, we defined the specification for these three operations. As this model is really tiny and simple, it is easy to understand and validate, but extremely incomplete. We then took some fundamental attributes of tasks into consideration. First, a task in FreeRTOS always has a state and it should be possible to change the state of a task from one to another. Second, the scheduling policy adopted by FreeRTOS is priority-based scheduling. A task in FreeRTOS must have a priority. The task then can be scheduled according to its priority. Therefore, the functions, *state* and *priority* are added into *FreeRTOS*. The specification for creating, deleting and rescheduling tasks are updated accordingly. Meanwhile, due to these two new functions, we can define the behaviour of suspending, resuming and changing priority of tasks. At this stage, we obtained a reasonable complete task model of FreeRTOS. As the model is still simple (compared to the expanded models) and we did not verify preconditions for each schema, there were no issues during the modelling and proving process. Identifying system attributes, encapsulating the attributes to a base schema and specifying the related behaviour of the system based on the schema can be considered as a good choice for a small system. Because of the

simplicity of the system, the size of base schema can be reasonably small. In this case, encapsulating all related attributes in the base schema does not raise the difficulty and complexity of proof for other schemas in the model.

After this, following the structure described in Chap. 3, we extended the model incrementally to contain queues, then added time, and finally added semaphores and mutexes. Similarly, each time, all base information, which will be introduced in detail in Sect. 4.3, 4.4 and 4.5, was added to the base schema *FreeRTOS* and the related operations are defined based on *FreeRTOS*. We found that during this process we produced an unacceptably large base schema. At the end, when semaphores and mutex related data were added to the model, the base schema *FreeRTOS* was longer than an A4 page. As a consequence, operation schemas also became unreasonably long, which increased the difficulty of validation and made the preconditions for operation schemas impossible to verify. The successful approach described above for the task model becomes unsuccessful, as the scale of the system increases the difficulty and complexity of the model dramatically.

To solve this problem, we broke the base schema, *FreeRTOS*, down into smaller pieces. As stated in the previous chapter, the base schema is split into four sub-schemas, *Task*, *Queue*, *Time* and *Mutex*. They are the four sub-models which are described in following sections. Using this strategy, the size and complexity of the schemas of the model are reduced dramatically. However, these sub-schemas still contain too much information to verify the preconditions for operation schemas, which produce too much unrelated information in the proof condition. This trivial information increases complexity and the difficulty of auto proving for the prover. Thus, the sub-base schemas were broken down further to obtain the current version of the model. Based on our experience, we can state that encapsulating all the attributes of the system to a single base schema is definitely an unsuccessful approach for a large system; on the other hand, hierarchically structured incremental base schema can be considered as a good choice, as it reduces the difficulty and complexity of the modelling and proving process by: (a) hiding as much unrelated information as possible, (b) proving theorems in the model which contains only the

information related to the theorem, (c) reusing the previous proved theorems for proving theorems in later complex model.

4.2 Task Model

As mentioned in Sect. 3.1, task related API functions are the core part of the model and fundamental to it. To define this model, it is essential to state some basic context that will be used in the specification.

4.2.1 Basic Statements

The given sets *CONTEXT* and *TASK* are provided as given sets to represent the environment of the processor and the tasks, respectively; in Z, given sets are basic, maximal types.

$$[CONTEXT, TASK]$$

Two constants, *bare_context* and *idle*, are introduced by an axiomatic definition, which contains a declaration and a constraint. Here, the constraint is trivially true and is omitted. The constant *bare_context* is an element of the set *CONTEXT*; it represents the initial state of the processor. The constant *idle* is of type *TASK*; it represents the system task that runs when no other task is scheduled.

$$\left| \begin{array}{l} \textit{bare_context} : \textit{CONTEXT} \\ \textit{idle} : \textit{TASK} \end{array} \right.$$

STATE is defined using a free type in its simplest form, enumerating exactly five distinct constants.

$$\textit{STATE} ::= \textit{nonexistent} \mid \textit{ready} \mid \textit{blocked} \mid \textit{suspended} \mid \textit{running}$$

The set of legal state transitions is described by an abbreviation: *transition* names the appropriate set that models the diagram in Fig. 1.1.

$$\begin{aligned} \textit{transition} == & (\{\textit{blocked}\} \times \{\textit{nonexistent}, \textit{ready}, \textit{running}, \textit{suspended}\}) \\ & \cup (\{\textit{nonexistent}\} \times \{\textit{ready}, \textit{running}\}) \\ & \cup (\{\textit{ready}\} \times \{\textit{nonexistent}, \textit{running}, \textit{suspended}\}) \\ & \cup (\{\textit{running}\} \times \{\textit{blocked}, \textit{nonexistent}, \textit{ready}, \textit{suspended}\}) \\ & \cup (\{\textit{suspended}\} \times \{\textit{nonexistent}, \textit{ready}, \textit{running}\}) \end{aligned}$$

In particular, transitions $(blocked, running)$ and $(suspended, running)$ are included because when a task is woken up from the `blocked` state or resumed from the `suspended` state, its state actually transits to `ready`. However, if it has a higher priority than the running task, it will be scheduled to `running`. At this level of abstraction, we consider these two steps as a single step, which makes state transitions $(blocked, running)$ and $(suspended, running)$ possible. The definition for *transition* turns out not to be very useful in automating proofs about transitions, because Z/Eves would expand *transition* into the set in all possible proof contexts. This greatly increases the load on the prover. Therefore, we `disable` the definition and add two theorems that are more helpful. The first is a typing lemma that states that *transition* is a set of pairs of *STATE*; its proof is a very simple consequence of the definition of *transition*. With the help of the proof command prefix “*with enabled (transition)*”, Z/Eves will take the disabled definition of *transition* into consideration during proof. The proof command “*prove by reduce*” requests Z/Eves to explore possible theorems and lemmas to prove the goal automatically. Therefore, the goal can be easily proved automatically by Z/Eves using the following command.

Theorem 1 (`gTransitionType`)

$$transition \in \mathbb{P}(STATE \times STATE)$$

proof [`gTransitionType`]

with enabled (transition) prove by reduce;

■

Next, we add the following lemma to tell Z/Eves about each individual pair in *transition*, which is helpful to Z/Eves for automatically proving. Similarly, the proof is very simple.

Theorem 2 (rule `lInTransition`)

$$\begin{aligned} \forall l, r : STATE \\ | (l, r) \in \{ & (nonexistent \mapsto ready), (running \mapsto ready), \\ & (blocked \mapsto ready), (suspended \mapsto ready), \\ & (ready \mapsto running), (blocked \mapsto running), \end{aligned}$$

- $(suspended \mapsto running), (nonexistent \mapsto running),$
 $(running \mapsto suspended), (ready \mapsto suspended),$
 $(blocked \mapsto suspended), (running \mapsto blocked),$
 $(running \mapsto nonexistent), (ready \mapsto nonexistent),$
 $(blocked \mapsto nonexistent), (suspended \mapsto nonexistent)\}$
- $(l, r) \in transition$

proof [*InTransition*]

with normalization with enabled (*transition*) prove by reduce;

■

Based on these definitions, the state schema of the model can be specified, describing basic system properties. For this stage of modelling, we focus only on task-related information in FreeRTOS. Further information will be introduced in a related model. To simplify the proof and the specification, we verify the system only when the scheduler is running. Therefore, we assume the scheduler is always running.

To describe the tasks in FreeRTOS, the following four kinds of data are needed, which are defined by a schema definition. In the Z notation, the schema is used to structure and compose descriptions. Once a schema is assigned a name, it is possible to use that name to reuse the schema in other expressions or schemas.

1. **Task data.** The variables recorded in this category are directly related to tasks. First, to simplify the description of the model and the following proofs, we need to distinguish tasks that are known to the system from others; therefore, a set *tasks* is defined as a finite subset of *TASK*. Second, in the FreeRTOS source code, *task.c* file, a pointer (*pxCurrentTCB*) is used to record the current running task, which is useful in several cases, such as scheduling. In the specification, a variable *running_task* of type *TASK* is used to represent this. Two constraints are specified: the *idle* task and the *running_task* have to be known to the system at all times.

$TaskData$ $tasks : \mathbb{F} TASK$ $running_task : TASK$

$running_task \in tasks$ $idle \in tasks$

2. **State data.** As described in Sect. 1.2.1, FreeRTOS uses different lists to manage the tasks known to the system. Abstractly, two tasks in different lists have different states. Therefore, the variable *state* is used to indicate the state of the tasks. Specifically, the *idle* task, which is a system task with responsibility for maintenance jobs for the system (such as garbage collection), can only be *ready* or *running*; it cannot be *blocked*, *suspended*, or *deleted* (*nonexistent*).

$stateData$ $state : TASK \rightarrow STATE$
$state(idle) \in \{ready, running\}$

3. **Context data.** The two variables *phys_context* and *log_context*, respectively, represent the physical system context (e.g., register values, some stacks, etc.) and the logical context for all the tasks that are not running (i.e., the system states of a task when it exits the running state).

$ContextData$ $phys_context : CONTEXT$ $log_context : TASK \rightarrow CONTEXT$

4. **Priority data.** FreeRTOS is a priority-based operating system: all the tasks in the system have their own priority, and a total function, *priority*, is introduced to record this. The priority of *idle* task must always be the lowest priority, which is 0.

$PrioData$ $priority : TASK \rightarrow \mathbb{N}$
$priority(idle) = 0$

Invariant Based on these definitions, we can describe the state schema for tasks that are maintained by this part of FreeRTOS.

$Task$ $TaskData$ $StateData$ $ContextData$ $PrioData$
$tasks = TASK \setminus (state \sim (\{nonexistent\}))$ $state \sim (\{running\}) = \{running_task\}$ $\forall pt : state \sim (\{ready\}) \bullet priority(running_task) \geq priority(pt)$

Apart from the four schemas describing the task, state, context, and priority data, three more constraints are added to this schema. They show that:

- All the tasks whose state is not *nonexistent* are known to the system. Here, as mentioned above, the *state* is a function, a special case of a relation. The operator, \sim , takes the inverse relation, so that $state \sim$ is a relation in $STATE \leftrightarrow TASK$. The operand, $(\)$ and $(\)$ calculates relational image. The result for this predicate is a set that contains all the *TASK*s whose states are *nonexistent*.
- Only one task can occupy the *running* state at any given time, which is *running_task*.
- The priority of the *running_task* is the highest of all the ready tasks.

Initialisation Based on the state definition and the assumptions mentioned above, we describe the initialisation of the *Task* state in a similar piecewise fashion: we separately initialise the four sub-states, and then combine them.

1. **Task data.** Initially, there are no user-defined tasks in the system; there is only one task in the system, *idle*, which is also the initial *running_task*.

$Init_TaskData$ $TaskData'$

$tasks' = \{idle\}$ $running_task' = idle$
--

2. **State data.** Furthermore, every other task is in the *nonexistent* state, except *idle* whose state is *running*.

$Init_StateData$ $StateData'$
$state' = (\lambda x : TASK \bullet nonexistent) \oplus \{(idle \mapsto running)\}$

3. **Context data.** Also, initially, the logical and physical contexts of all tasks is the *bare_context*.

$Init_ContextData$ $ContextData'$
$phys_context' = bare_context$ $log_context' = (\lambda x : TASK \bullet bare_context)$

4. **Priority data.** Finally, all tasks have the lowest priority, 0.

$Init_PrioData$ $PrioData'$
$priority' = (\lambda x : TASK \bullet 0)$

The initial state for *Task* can be defined using these four definitions.

$Init_Task$ $Task'$
$Init_TaskData$ $Init_StateData$ $Init_ContextData$ $Init_PrioData$

In order to prove that all the initial states are reachable, the following five theorems are introduced. They assert that there is at least one possible postcondition for initialising each sub-state schema and the overall schema. Due to the simplicity of these theorems, Z/Eves is able to fully prove them automatically.

Theorem 3 (TaskDataInit)

$\exists \text{TaskData}' \bullet \text{Init_TaskData}$

proof [*TaskDataInit*]
prove by reduce;

■

Theorem 4 (StateDataInit)

$\exists \text{StateData}' \bullet \text{Init_StateData}$

Theorem 5 (ContextDataInit)

$\exists \text{ContextData}' \bullet \text{Init_ContextData}$

Theorem 6 (PrioDataInit)

$\exists \text{PrioData}' \bullet \text{Init_PrioData}$

It is easy to prove these theorems with the proof command “*prove by reduce*”, except for *TaskInit*, because it has more constraints on its state variables.

Theorem 7 (TaskInit)

$\exists \text{Task}' \bullet \text{Init_Task}$

proof [*TaskInit*]
prove by reduce;
apply extensionality;
with enabled (applyOverride) prove;

■

After the automatic proving ordered by *prove by reduce*, Z/Eves is confused about the equivalence between sets defined in schema *Task*. The application of *override* also confuses the prover. Therefore, we need to guide the prover to apply theorems, *extensionality* and *applyOverride*, to discharge them. These theorems are provided by the Z/Eves toolkit [70].

We can check whether the state change respects the transition relation as a dynamic invariant that must be satisfied by all the operations on the $Task$ state by redefining $\Delta Task$:

$$\begin{array}{|l}
 \hline
 \Delta Task \\
 Task \\
 Task' \\
 \hline
 \forall st : TASK \mid state'(st) \neq state(st) \\
 \bullet state(st) \mapsto state'(st) \in transition \\
 \hline
 \end{array}$$

It is worth mentioning that in this schema we use $Task'$ to refer to the post state of the $Task$. Initially, the expression “ $\Delta Schema$ ” ($Schema$ refers to a state schema) has been defined to contain both the pre- and post-state of $Schema$. We redefine it here to add further constraints for $Task$.

Based on these fundamental definitions, operations related to tasks can be specified.

4.2.2 Additional Schema for Reschedule

In a multi-tasking real-time operating system, rescheduling tasks is essential and occurs frequently. Generally, depending on the purpose of the system, the operating system would follow some suitable algorithm to determine the task to be scheduled. Other system states can then be updated accordingly. Therefore, at this level of abstract specification, it is possible to define the rescheduling process nondeterministically. However, the model described in this chapter focuses on FreeRTOS. We will follow the algorithm used in FreeRTOS to specify rescheduling, which is based on task priority. Specifically, once a `ready` task obtains a higher priority than the running task, it will be scheduled as the new running task. Subsequently, the system will switch the context of the current running task out and swap in the context of the new running task. It is also necessary to manage related lists and system states properly, for instance, by setting the selected task as the running task and inserting the current running task in a suitable list.

In this specification, we introduce the schema $Reschedule$ to perform the swapping part of the rescheduling process, which can then be used by other schemas.

The priority-based scheduling algorithm is embedded in the operation schemas for different API functions that need rescheduling. The priority-based rescheduling behaviour depends on the destination to which the current running task is moved. For example, when suspending the running task, the destination of the running task is the suspended list; but when we create a task with a higher priority than the running task, the destination of the current running task is actually one of the ready lists. These lists are represented by the function *state*. Therefore, updating the *state* with the variable *st?* manages these lists. In the Z notation, variables marked with “?” and “!” indicate that they are I/O variables, respectively, for a schema. When other schemas reuse the *Reschedule* schema, *st?* will be introduced within these schemas with the value of the destination of the current running task. Because both schemas contain a variable with the same name, these two variables will be bound together. Consequently, the schema *Reschedule* can obtain the destination of the running task by accessing the value of *st?*. The operator, \oplus , is normally used to update functions in Z. If the first element of a pair exists in the domain of the function, it will update the second element of the pair in the function to the new value; otherwise it appends the pair to the function. Therefore, it is used here to update the *state* of *running_task* and *target?*. Similarly, for each case, the new running task, the final state of *tasks*, and the priority of tasks may also be different. We leave these decisions to the calling schemas. Therefore, variables—*target?*, *tasks?* and *pri?*—are introduced to represent these properties.

<i>Reschedule</i>
$\Delta Task$ <i>target?</i> : <i>TASK</i> <i>tasks?</i> : $\mathbb{P} TASK$ <i>st?</i> : <i>STATE</i> <i>pri?</i> : <i>TASK</i> $\rightarrow \mathbb{N}$
<i>tasks'</i> = <i>tasks?</i> <i>running_task'</i> = <i>target?</i> <i>state'</i> = <i>state</i> $\oplus \{(target? \mapsto running), (running_task \mapsto st?)\}$ <i>phys_context'</i> = <i>log_context</i> (<i>target?</i>) <i>log_context'</i> = <i>log_context</i> $\oplus \{(running_task \mapsto phys_context)\}$ <i>priority'</i> = <i>pri?</i>

The calling schema just needs to specify the correct values for these variables, the *Reschedule* schema then handles the rest of the work.

4.2.3 Creating and Deleting Tasks

After initialising the system, there is only one task (*idle*); in order to add more tasks to the system, the *Create* operation can be used. Once a task finishes, it should be *Deleted* (see Page. 55) to allow other tasks to use the resources held by it. *xTaskCreate* and *vTaskDelete* are also the first group of API functions provided by FreeRTOS. Generally, there are two cases for each of these two operations: one is to add or remove a task from the system; the other one leads to a re-scheduling of tasks.

First Case of Creating Tasks If the assigned priority is not greater than the priority of the current running task, it adds the new task that does not already exist. The input *target?* represents the task that will be created. The input *newpri?* contains the priority assigned to the new task. Therefore, the precondition is specified as: first, *target?* is not known by the system; second, the assigned priority, *newpri?* is no more than the priority of *running_task*. After the operation, the *target?* is known to the system, the task *target?* is added to *tasks* and updates the *state* function to record that the state of *target?* is *ready*. The input *newpri?* is assigned to the task *target?* by updating the function *priority*. Because this operation will not cause rescheduling, other properties of *Task* remain unchanged. The “ Ξ ” operation has been used here: it is defined in Z to show that the pre- and post-states are unchanged. The schema *CreateTaskN_T* can be introduced, which indicates that this schema is used to *Create Task operation* in the *normal case* for the *Task model*. Generally, we use postfix *N* for the *Normal case* of the operation, which does not lead to rescheduling; and *S* for the *Scheduling case*. The postfix after the underscore indicates which model it is specified for. For example, *T* in this case shows the schema is part of the task model.

$$\left\{ \begin{array}{l} \text{CreateTaskN_T} \\ \Delta \text{Task} \end{array} \right.$$

$target? : TASK$ $newpri? : \mathbb{N}$
$state(target?) = nonexistent$ $newpri? \leq priority(running_task)$ $tasks' = tasks \cup \{target?\}$ $running_task' = running_task$ $state' = state \oplus \{(target? \mapsto ready)\}$ $\exists ContextData$ $priority' = priority \oplus \{(target? \mapsto newpri?)\}$

Having defined this operation as a relation on *Task* states, we need to work out its precondition. We posit that the before-state, the inputs, and the first two predicates are the precondition, and collect these into the following schema, where the suffix *FSBSig* in the schema name stands for *Feasibility Signature*.

$CreateTaskN_TFBSig$
$Task$ $target? : TASK$ $newpri? : \mathbb{N}$
$state(target?) = nonexistent$ $newpri? \leq priority(running_task)$

These declarations and predicates are clearly necessary for the actual precondition as stated above. We show that they are also sufficient in the next theorem, which can be automatically generated. Specifically, for any “state” that satisfies the definition of *CreateTaskN_TFBSig*, the precondition of *CreateTaskN_T* is satisfied. The operator “pre *Schema*” is defined in Z to calculate the precondition schema of a schema [65, Chap. 14]. For instance, the predicate “pre *CreateTaskN_T*” in the following theorem obtains the precondition schema by calculating $\exists Task' \bullet CreateTaskN_T \setminus (outputs)$, where *outputs* refers to the list of output variables related to the operation, which will be hidden, and is empty in this case. The schema hiding operator, “\”, hides the variables listed in the *outputs* from the declaration of the operation by introducing them in the predicate part of the schema with an existential quantifier.

Theorem 8 (CreateTaskN_T_vc_ref)

$$\forall CreateTaskN_TFBSig \mid true \bullet \text{pre } CreateTaskN_T$$

It is interesting to understand the proof of this theorem. First of all, as mentioned above, the Z/Eves prover is used to verify our specification. All the proof scripts shown in this thesis are used to help Z/Eves to finish the proof work. Generally, there are two ways to finish a proof [69, Chap. 5]: (a) exploratory proof — directly prove the theorem without any previous plan and address any proof goals returned by the prover; (b) planned proof — carry out a detailed plan for the proof, which is enough to finish the proof by hand, then transfer the plan to a proof script for the prover. To maximise the benefit of proof automation, we adopt the exploratory proof approach in many cases. The general idea for this approach is:

1. Expand terms such as schema references and let Z/Eves prove the proof goal automatically.
2. When Z/Eves is stuck, stop at the proof goals, guide Z/Eves by using or applying related theorems or lemmas to rewrite the proof goals, provide more conditions, etc.
3. Let Z/Eves progress based on the new goal.
4. Repeat step 2 and step 3 until the proof is finished.

For efficiency, it is necessary to expand as few terms as possible in step 1. This can significantly reduce the proof time, especially when the system is complex. This is also one of the reasons for defining our system in parts.

Specifically, we first use the following proof command to expand all necessary terms and then let the prover automatically apply rules and theorems, which are included by Z/Eves, to prove the goal.

with disabled (ContextData) prove by reduce;

Meanwhile, because the *ContextData* is unchanged in this schema, we keep it unexpanded. The prefix *with disabled (ContextData)* can achieve this by making the prover ignore *ContextData*, when expanding the terms. Note that as some theorems are rarely used when proving and other theorems are time consuming, Z/Eves disables them by default. This is helpful for improving the efficiency of the proof process; however, it is also one of the reasons why Z/Eves may become stuck

in some cases. As a result, the original proof goal is transferred to the following five goals².

1. The tasks known by the system are finite.

$$TASK \setminus ((state \oplus \{(target?, ready)\}) \sim (\{nonexistent\} \setminus)) \in \mathbb{F} TASK$$

As defined in *CreateTaskN_T*, the expression $state \oplus \{(target?, ready)\}$ is equal to the post *state*. The left side of the expression indicates all the tasks known by the system after the operation. It should be a finite set as defined in *TaskData*.

2. The *running_task* remains the same before and after the operation:

$$(state \oplus \{(target?, ready)\}) \sim (\{running\} \setminus) = state \sim (\{running\} \setminus)$$

Similarly, the image of *running* under the inverse function $(state \oplus \{(target?, ready)\}) \sim$ represents the *running_task* after the operation.

3. The *target?* task is added into the system by the operation:

$$\begin{aligned} & TASK \setminus ((state \oplus \{(target?, ready)\}) \sim (\{nonexistent\} \setminus)) \\ & = \{target?\} \cup (TASK \setminus (state \sim (\{nonexistent\} \setminus))) \end{aligned}$$

After the operation, the tasks known by the system should be the same as the known tasks of the pre state of the system plus the created task, which is *target?*.

4. The priority of the *target?* task is less than or equal to the running task:

$$\begin{aligned} & ((state \oplus \{(target?, ready)\})(pt) = ready \wedge (pt = target? \vee pt \in TASK)) \\ & \Rightarrow \\ & priority(running_task) \geq (priority \oplus \{(target?, newpri?)\})(pt) \end{aligned}$$

Comparable to *state*, the post state of the *priority* function can also be written as: $priority \oplus \{(target?, newpri?)\}$. In this case, this expression is easy to understand.

²Because the proof goals are too long to present in this thesis, we only list the most important part here. Please download the Z/Eves project file from the supplementary material and open it with Z/Eves to find the full details.

5. Every state transition made by any task respects the *transition* relation.

$$\begin{aligned} & (st \in TASK \wedge \neg (state \oplus \{(target?, ready)\})(st) = state(st) \\ & \Rightarrow \\ & (state(st), (state \oplus \{(target?, ready)\})(st)) \in transition) \end{aligned}$$

It is easy to find that the key to proving both goals 1 & 3 is goal 3. As defined, $TASK \setminus (state \sim (\{nonexistent\}))$ which is *tasks*, is a finite set. If we can prove goal 3, we can easily show that the union of two finite sets is a finite set. Furthermore, for goal 3, the prover is actually confused by the complex set calculation on the left side of the equation. As these two goals are derived from the constraint of the state schema, we expect they will repeat frequently in the precondition proofs of other schemas. Therefore, we introduce a lemma, *setminUpdate*, to help Z/Eves to discharge this kind of goal automatically. In Z/Eves, it is possible to use the keyword, *rule*, to define an external lemma to help the proof. Z/Eves will use them automatically when the *prove* command is called.

Theorem 9 (rule setminUpdate)

$$\begin{aligned} & \forall f : TASK \rightarrow STATE; g : TASK \rightarrow STATE \bullet \\ & TASK \setminus ((f \oplus g) \sim (\{nonexistent\})) = \\ & TASK \setminus (f \sim (\{nonexistent\}) \setminus (g \sim (\{nonexistent\}))) \cup \\ & (\text{dom } g \setminus (g \sim (\{nonexistent\}))) \end{aligned}$$

Similarly, the proof goal 2 is also likely to repeat during the verification. We define another lemma, *runningUpdate*, to improve automation of proof.

Theorem 10 (rule runningUpdate)

$$\begin{aligned} & \forall f : TASK \rightarrow STATE; g : TASK \rightarrow STATE \mid \\ & running \notin \text{ran } g \wedge \\ & (f \sim (\{running\})) \cap \text{dom } g = \emptyset \bullet \\ & (f \oplus g) \sim (\{running\}) = f \sim (\{running\}) \end{aligned}$$

To prove these two lemmas, the *extensionality* rule in Z/Eves toolkit can be used. The detailed proof script can be found in the supplementary material. After adding these two lemmas before theorem 8 and restarting the proof, we can find that the first three proof goals are discharged automatically.

Proof goal 4 is then given by the constraint in *Task* schema. Tasks other than *target?* maintain the requirement that the priority of the running task is at least as great as that of all the ready tasks:

$$\forall pt : state \sim (\{ready\}) \bullet priority(running_task) \geq priority(pt)$$

A copy of this constraint is also in the assumption part of the goal, and to distinguish *pt* in these two, Z/Eves renames one from *pt* to *pt_0*. Therefore, to prove that tasks other than *target?* obey the constraint, we just need to indicate that *pt_0* and *pt* are the same. For *target?*, the priority is defined as *newpri?*, which is specified to be no higher than the priority of *running_task* as a precondition of this schema. The rule *applyOverride* is applied to analyse expressions that contain the operator \oplus . Finally, the command *with normalization prove;* is used to finish the proof³. Thus, the theorem *CreateTaskN_T_vc_ref* can be proved by following script in Z/Eves.

proof [*CreateTaskN_T_vc_ref*]
with disabled (ContextData) prove by reduce;
instantiate pt_0 == pt;
with enabled (applyOverride) prove;
apply applyOverride;
with normalization reduce;

■

Second Case of Creating Tasks If the priority assigned to the new task is greater than the priority of the running task, then rescheduling is required. This is achieved by calling the *Reschedule* schema. The current running task will be moved into the ready state; the new priority and initial context is allocated for the new task, which is then scheduled to be the running task. To reuse *Reschedule*, the variables *st?*, *pri?* and *tasks?* are declared and assigned appropriately. Note, the default logical context for the new tasks is *bare_context*. We do not need to set it separately. Therefore, the schema for the second case of the create task operation can be defined as follows:

$$\sqsubset \text{CreateTaskS_T} \text{_____}$$

³The details about the proof command, *with normalization prove;*, can be found in Appendix B.

$\Delta Task$ $target? : TASK$ $newpri? : \mathbb{N}$
$state(target?) = nonexistent$ $newpri? > priority(running_task)$ $\exists st? : STATE; tasks? : \mathbb{F} TASK; pri? : TASK \rightarrow \mathbb{N}$ $\quad st? = ready$ $\quad \wedge tasks? = tasks \cup \{target?\}$ $\quad \wedge pri? = priority \oplus \{(target? \mapsto newpri?)\} \bullet Reschedule$

Similarly to the previous case, the signature schema and the precondition theorem can be defined.

$CreateTaskS_TFSSig$
$Task$ $target? : TASK$ $newpri? : \mathbb{N}$
$state(target?) = nonexistent$ $newpri? > priority(running_task)$

Theorem 11 (CreateTaskS_T_vc_ref)

$$\forall CreateTaskS_TFSSig \mid true \bullet \text{pre } CreateTaskS_T$$

This indicates that the new task is unknown to the system before the operation and the priority of the new task is higher than the priority of the running task. This is sufficient and necessary for the precondition of schema $CreateTaskS_T$.

Deleting Tasks The first case for deleting a task is that it is not the running task: the state of this task—provided it is not the *idle* task—can be *ready*, *blocked*, or *suspended*, because, normally the handle of the *idle* task, `xIdleTaskHandle`, is private to the system and impossible for the user to obtain. After the operation, the deleted task will become unknown to the system by deleting it from *tasks*, setting its state to *nonexistent*, and setting its logical context to the *bare_context*.

It is worth mentioning that in the source code of `vTaskDelete` in FreeRTOS, the context of the deleted task is not actually deleted, but instead moved to the `xTasksWaitingTermination` list. It is the `idle` task that actually performs garbage collection to recover the resources allocated by the system. At this level of abstraction, we consider all this as part of the deletion operation, resetting the *log_context* of the deleted task to the *bare_context*. Note, due to space limitations, we only list the parts of our model which contain something of interest; the rest of the specifications, precondition theorems, and proof scripts can be found from the supplementary material.

Secondly, if the task to be deleted is the running task—but not the *idle* task—then we remove it from the system. This leaves a vacuum to be filled: we need to schedule another process to use the CPU. We will choose the task in a ready state with the highest priority. However, we cannot use *Reschedule* to achieve this because the logical context of the running task, which is requested by this operation but not supported by *Reschedule*, will be reset. The output variable *topReady!* is introduced. The universally quantified expression specifies that the *topReady!* holds the highest priority. It is worth mentioning here that if there are several solutions, then *topReady!* is chosen nondeterministically. Similarly, the *tasks*, *state*, *phys_context* and *log_context* are updated.

<i>DeleteTaskS_T</i>
$\Delta Task$ $target? : TASK$ $topReady! : TASK$
$target? \in tasks \setminus \{idle\}$ $state(target?) \in \{running\}$ $state(topReady!) = ready$ $\forall t : state \sim (\{ready\} \mid) \bullet priority(topReady!) \geq priority(t)$ $tasks' = tasks \setminus \{target?\}$ $running_task' = topReady!$ $state' = state \oplus \{(topReady! \mapsto running), (target? \mapsto nonexistent)\}$ $phys_context' = log_context(topReady!)$ $log_context' = log_context \oplus \{(target? \mapsto bare_context)\}$ $\exists PrioData$

The signature schema of this can be obtained as follows.

$DeleteTaskS_TFSSig$ $Task$ $target? : TASK$
$target? \in tasks \setminus \{idle\}$ $state(target?) \in \{running\}$

Theorem 12 (DeleteTaskS_T_vc_ref)

$$\forall DeleteTaskS_TFSSig \mid true \bullet \text{pre } DeleteTaskS_T$$

As mentioned above, the “pre” operator calculates the precondition schema for $DeleteTaskS_T$, which is the result of $\exists Task' \bullet DeleteTaskS_T \setminus (topReady!)$. When the prover automatically discharges this predicate, it attempts to eliminate existentially quantified variables. Because the post state of the system, $Task'$, has been defined in the operation, the one-point rule⁴ is applied to handle them. However, the variable $running_task'$ and $topReady!$ can only eliminate one of them, because the output variable $topReady!$ is assigned the value of $running_task'$ in this operation. Therefore, the proof goal will become:

$$\begin{aligned}
& \exists running_task' : TASK \bullet \\
& \quad Task[log_context := log_context \oplus \{(target?, bare_context)\}, \\
& \quad \quad phys_context := log_context(running_task'), \\
& \quad \quad running_task := running_task', \\
& \quad \quad state := state \oplus (\{(target?, nonexistent)\} \\
& \quad \quad \quad \cup \{(running_task', running)\}), \\
& \quad \quad tasks := tasks \setminus \{target?\}] \\
& \wedge (\forall st : TASK \mid \\
& \quad \neg (state \oplus (\{(target?, nonexistent)\} \cup \\
& \quad \quad \{(running_task', running)\}))(st) = state(st) \bullet \\
& \quad (state(st), (state \oplus (\{(target?, nonexistent)\} \cup \\
& \quad \quad \{(running_task', running)\}))(st)) \in transition) \\
& \wedge state(running_task') = ready \\
& \wedge (\forall t_0 : state \sim (\{ready\}) \bullet \\
& \quad priority(running_task') \geq priority(t_0)
\end{aligned}$$

⁴One-point rule: $\exists x : X \mid p \bullet q \wedge x = t \Leftrightarrow p[t/x] \wedge q[t/x] \wedge t \in X$, provided that x is not free in t .

Meanwhile, as defined in the specification, one of the highest priority ready tasks is nondeterministically assigned to the variable $running_task'$. In this case, with the existential-elimination rule, if we can find an instance of these tasks that satisfies this predicate, the proof goal can be verified. Therefore, we introduce the following function (f) to discover a member of a set of tasks that has the highest value of its g , which can be replaced by priority, among other tasks in that set. A label “findDelegate” is assigned to this lemma, which can be referred to during later proofs.

$$\begin{array}{|l}
 f : \mathbb{P} \text{ TASK} \rightarrow \text{ TASK} \\
 \hline
 \emptyset \notin \text{ dom } f \\
 \langle\langle \text{ findDelegate } \rangle\rangle \\
 \forall \text{ Task}; a : \mathbb{P} \text{ TASK}; g : \text{ TASK} \rightarrow \mathbb{Z} \\
 \bullet a \in \text{ dom } f \wedge f(a) \in a \wedge a \subseteq \text{ dom } g \wedge (\forall t : a \bullet g(f(a)) \geq g(t))
 \end{array}$$

It is then possible to use this function to find the highest priority task in a ready state and use it to instantiate the $running_task'$, when the function $priority$ is assigned to the variable g . If we let $p, \exists x : X \bullet q$ represent the conditions and the goal of proof above, the predicate to be proved can be considered as

$$p \Rightarrow \exists x : X \bullet q. \quad (4.1)$$

Further, let t represent $f(\text{state} \sim (\{ready\}))$. When we instantiate the $running_task'$ with the delegate, with one-point rule, we have $\exists x : X \bullet q \wedge x = t \Leftrightarrow t \in X \wedge q[t/x]$, which gives $\exists x : X \bullet q \Leftrightarrow \exists x : X \bullet q \vee (t \in X \wedge q[t/x])$. Therefore, equation (4.1) transfers into:

$$p \Rightarrow (\exists x : X \bullet q) \vee (t \in X \wedge q[t/x]). \quad (4.2)$$

Reorganising the equation, the relation

$$p \wedge \neg (t \in X \wedge q[t/x]) \Rightarrow \exists x : X \bullet q. \quad (4.3)$$

can be acquired.

Therefore, applying the proof command “ $instantiate \text{ running_task}' == f(\text{state} \sim (\{ready\}))$;”, a negative copy of this proof goal will be added to the condition part,

of which $running_task'$ will be replaced by $f(state \sim (\{ready\} \cup \emptyset))$. Analysing the negative copy of the goal, we find

$$\begin{aligned}
Task[log_context &:= log_context \oplus \{(target?, bare_context)\}, \\
phys_context &:= log_context(f(state \sim (\{ready\} \cup \emptyset))), \\
\dots &\Rightarrow \\
&t \in TASK \wedge \\
&state(t) = ready \wedge \\
&\neg priority(f(state \sim (\{ready\} \cup \emptyset))) \geq priority(t) \quad (4.4)
\end{aligned}$$

which conflicts with the definition of function f ; therefore, it is **not** *true*. However, according to the implication, if we can prove that the condition is *false*, the result of proof is *true*. As proving this condition is *false* presents difficulties due to the complexity, an auxiliary theorem, $lDeleteTaskS_T_Lemma$, is introduced and proved separately. When we use it to prove $DeleteTaskS_T_vc_ref$, the variable $topReady!$ can be substituted by $f(state \sim (\{ready\} \cup \emptyset))$.

Theorem 13 (lDeleteTaskS_T_Lemma)

$$\begin{aligned}
&\forall Task; topReady!, target? : TASK \\
&\quad | target? \in tasks \setminus \{idle\} \\
&\quad \wedge state(target?) \in \{running\} \\
&\quad \wedge state(topReady!) = ready \\
&\quad \wedge (\forall rtsk : state \sim (\{ready\} \cup \emptyset) \bullet priority(topReady!) \geq priority(rtsk)) \\
&\quad \bullet \neg (Task[log_context := log_context \oplus \{(target?, bare_context)\}, \\
&\quad\quad phys_context := log_context(topReady!), \\
&\quad\quad running_task := topReady!, \\
&\quad\quad state := state \oplus \\
&\quad\quad\quad (\{(target?, nonexistent)\} \cup \{(topReady!, running)\}), \\
&\quad\quad tasks := tasks \setminus \{target?\}] \\
&\quad \wedge (st \in TASK \\
&\quad\quad \wedge \neg (state \oplus (\{(target?, nonexistent)\} \cup \\
&\quad\quad\quad \{(topReady!, running)\}))(st) = state(st) \\
&\quad\quad \Rightarrow (state(st), (state \oplus (\{(target?, nonexistent)\} \cup \\
&\quad\quad\quad \{(topReady!, running)\}))(st) \in transition) \\
&\quad \Rightarrow t \in TASK \\
&\quad\quad \wedge state(t) = ready \\
&\quad\quad \wedge \neg priority(topReady!) \geq priority(t)
\end{aligned}$$

Generally, the purpose of Theorem 13 is to prove and notify Z/Eves that for all the states that satisfy the definition of *Task*, based on the precondition of schema *DeleteTaskS_T*, the proof goal (4.4) is *false*. With this information and the following script, Z/Eves can easily prove the result of Theorem 12 is *true*. *DeleteTaskS_T_vc_ref* can be continued.

```
proof [DeleteTaskS_T_vc_ref]
  use findDelegate[a := state ~ (| {ready} |), g := priority];
  with disabled (Task) prove by reduce;
  instantiate running_task' == f(state ~ (| {ready} |));
  prove;
  use lDeleteTaskS_T_Lemma[topReady! := f(state ~ (| {ready} |))];
  prove;
  instantiate t_0 == rtsk;
  prove;
```

■

4.2.4 Executing Tasks

In FreeRTOS, there is no API function for this: once the task is scheduled, it will be executed automatically. However, it is helpful for specifications to show the task being executed, especially when executing the specification with an animator (such as ProZ). In detail, when the processor executes a task, it updates registers, flags, memory location, and so on. We model this by updating the physical context of the processor. Here, we are not interested in the new value after the operation, but want to know that it has changed and the new value has some special property. Therefore, we use a nondeterministic definition again for updating *phys_context*. Because the schema *ExecuteRunningTask_T* describes executing the task, if the new value of *phys_context* is different from the original, it will be satisfied.

4.2.5 Suspending/Resuming Tasks

Like creating and deleting, suspending and resuming tasks also have two cases. When the system suspends a ready or blocked task, it does not lead to rescheduling. However, if the task to be suspended is the running task, then the system needs to find another task to take the processor. If a resumed task has a higher priority than the running task, it becomes the new running task, otherwise, it goes

to the ready state. As mentioned above, normally the handle of the *idle* task is not obtainable. Even though the user may extend the behaviour of the *idle* task by modifying the `vApplicationIdleHook` function, the *idle* task must never be suspended [18], and consequently can never be resumed. It is possible to suspend an already suspended task: the system keeps everything the same as before. So, the first case concerns suspending a task that is ready or blocked; the only change necessary is to update the task's state. The following script shows the precondition theorem and proof script of the schema of this case.

Theorem 14 (SuspendTaskN_T_vc_ref)

$\forall \text{ SuspendTaskN_TFSSig} \mid \text{true} \bullet \text{pre SuspendTaskN_T}$

proof [*SuspendTaskN_T_vc_ref*]

prove by reduce;

apply extensionality to predicate $\text{TASK} \setminus (\text{state} \sim (\downarrow \{ \text{nonexistent} \} \downarrow)) = \text{TASK} \setminus ((\text{state} \oplus \{(\text{target?}, \text{suspended} \})) \sim (\downarrow \{ \text{nonexistent} \} \downarrow))$;

apply extensionality to predicate $(\text{state} \oplus \{(\text{target?}, \text{suspended} \})) \sim (\downarrow \{ \text{running} \} \downarrow) = \text{state} \sim (\downarrow \{ \text{running} \} \downarrow)$;

instantiate $\text{pt_0} == \text{pt}$;

prove;

apply *applyOverride*;

with normalization prove;

■

Due to the complication of the proof goal, the final proof command “*with normalization prove;*” requires a significant amount of time to complete. However, if we use the “*cases, next*” commands to separate the proof goals into different cases and then apply “*with normalization prove;*”, it becomes much more efficient.

The second case of the suspend operation is when the suspended task is the running task. Clearly, this leads to rescheduling. This operation ensures that the running task is not the idle task. It selects a target that is ready and is one of the ready tasks with the highest priority (there may be many such tasks). The *Reschedule* schema is used to achieve the necessary rescheduling. Similar to *DeleteTaskS_T*, a nondeterministically chosen value is assigned

to *running_task'*. The prover is therefore confused about its value. An additional theorem, *lSuspendTaskS_T_Lemma*, is introduced to help the prover with the precondition. Finally, it is also possible to suspend a suspended task. According to the reference manual of FreeRTOS [19], nothing changes when a suspended task is suspended. A single call to `vTaskResume` can resume the task that has been suspended several times. For this reason, in schema *SuspendTaskO_T*, predicate $\exists Task$ is used to show that the pre- and post-value of all variables within *Task* schema are unchanged.

Similarly, the first case of resuming a task does not cause rescheduling. The priority of the resumed task must be no higher than the running task. The task is moved to the *ready* state and everything else is unchanged. In the second case, the resumed task has a higher priority than the running task, and rescheduling is required. Again, the schema *Reschedule* is used to approach this.

4.2.6 Changing Priority of Tasks

Because the priority of the *idle* task is permanently 0, if the target task is *idle*, the *newpri?* should equal 0. Specifically, to change the priority of tasks, there are three different cases that need to be considered. In the first case, there is no scheduling required, and this follows if one of the following conditions hold:

1. The target is the *running* task and the new priority is at least as high as every other *ready* task.
2. The target is *ready* and the new priority does not have a greater priority than the *running* task.
3. The target is the *idle* task and the new priority is 0.
4. The target is *blocked*.
5. The target is *suspended*.

Note that we cannot change the priority of *nonexistent* tasks. Further, as the set *TASK* is composed of *running*, *ready*, *blocked*, *suspend*, and *nonexistent*, tasks in these states are disjoint. Therefore, the predicate $state(target?) \neq nonexistent$

implies that the *target?* is in one of the other four states. That means for conditions related to the *blocked* and *suspended* states, we do not need other predicates. Finally, the effect of the operation is to only change the priority of the target, but nothing else. Then, we update the function *priority* by overriding the priority of the *target?* task with *newpri?*.

In the second case, the target is a ready task whose new priority is higher than that of the running task. The target displaces the running task as the tasks are rescheduled. Similarly, the *Reschedule* schema is used to achieve this.

In the third case, similar to the second, rescheduling is required. However, the target task, whose priority we wish to change, is the running task. Meanwhile, the new priority is not the highest of the ready tasks. The schema for this would firstly pick up the task with the highest priority in the ready tasks. It updates the value of the priority of the running task. Finally, it reschedules the system with the *Reschedule* schema. The variable “*topReady!*” here, similar to the schedule case of delete task and suspend task, is used to represent which ready task holds the highest priority among other ready tasks and would be scheduled as the new running task after the operation. Also, the schema *lChangeTaskPriorityD_T_Lemma* (See Page. 183) is introduced to handle the nondeterministically chosen value of “*running_task!*”.

4.3 Queue Model

Queue is the facility provided by FreeRTOS for communication between tasks. Similar to the task model, to define this model, we need to specify some basic states.

4.3.1 Basic Statements

Firstly, we define *QUEUE* to represent the queues.

[*QUEUE*]

The properties of queues in FreeRTOS, can generally be divided into the following three parts:

1. **Queue data.** In this schema, variables are used to record the basic properties of queues. First, set *queue* is given to distinguish the queues known to the system from others. Second, it is necessary to know the **maximum** size and the actual size of each queue in the system. Therefore, two functions, *q_max* and *q_size*, are respectively used for these, and the domain of these two functions should equal *queue*. Furthermore, for each queue in the system, its actual size cannot exceed its maximum size.

<i>QueueData</i>
$queue : \mathbb{P} \text{QUEUE}$ $q_max : \text{QUEUE} \rightarrow \mathbb{N}_1$ $q_size : \text{QUEUE} \rightarrow \mathbb{N}$
$\text{dom } q_max = \text{dom } q_size = queue$ $\forall q : \text{QUEUE} \mid q \in queue \bullet q_size(q) \leq q_max(q)$

2. **Waiting data.** As the maximum size of a queue is finite, queues can be full when a task attempts to send an item to them. In this case, FreeRTOS allows the task to wait until spaces become available in the queue. Similarly, when a task wants to receive items from a queue which is empty, the task may also wait for some resources to be available in the queue. Therefore, two functions *wait_snd* and *wait_rcv* are defined, respectively. Due to the definition of these two functions, a task cannot both be waiting to send and to receive items at the same time. A constraint is specified for this property, that the intersection of the domain of these two functions is the empty set.

<i>WaitingData</i>
$wait_snd : \text{TASK} \rightarrow \text{QUEUE}$ $wait_rcv : \text{TASK} \rightarrow \text{QUEUE}$
$\text{dom } wait_snd \cap \text{dom } wait_rcv = \emptyset$

3. **Releasing data.** According to the implementation of FreeRTOS, when a task is released from a waiting event, it should continue the operation which it was performing before the event. For instance, if the task is blocked by sending

an item to a queue, when it is released from *wait_snd* it should continue sending the item to the queue. To achieve this, two assistant functions are provided to indicate if a task has just been released from the waiting event and also record the queue it was dealing with.

$QReleasingData$ $release_snd : TASK \rightarrow QUEUE$ $release_rcv : TASK \rightarrow QUEUE$
$\text{dom } release_snd \cap \text{dom } release_rcv = \emptyset$

Gathering these three schemas, we can define the schema *Queue* for the properties across each sub-state.

$Queue$ $QueueData$ $WaitingData$ $QReleasingData$
$\text{ran } wait_snd \subseteq queue$ $\text{ran } wait_rcv \subseteq queue$ $\text{ran } release_snd \subseteq queue$ $\text{ran } release_rcv \subseteq queue$ $(\text{dom } wait_snd \cup \text{dom } wait_rcv) \cap (\text{dom } release_snd \cup \text{dom } release_rcv) = \emptyset$

First, the range of functions described in *WaitingData* and *QReleasingData* are *queue*, because a task cannot send or receive items from a queue which is unknown to the system. Additionally, the tasks of functions in *WaitingData* and *QReleasingData* are disjoint, because only when a task is removed from the functions in *WaitingData*, it can then be added to the related function in *QReleasingData*.

Finally, the state data for the queue level model can be defined using these definitions.

$TaskQueue$ $Task$ $Queue$

$$\begin{array}{l} \text{dom } wait_snd \subseteq state \sim (\{ blocked \}) \\ \text{dom } wait_rcv \subseteq state \sim (\{ blocked \}) \end{array}$$

There are two extra constraints for this schema, which indicate that the state of tasks in the domain of *wait_snd* and *wait_rcv* is blocked.

Initialisation We have defined the initialisation for the *Task* schema in the previous section. Now, we only need to specify the initialisation for *Queue*, then combine them to obtain the initialisation for the queue model (i.e. *TaskQueue* schema). Initially, no queues exist in the system, therefore, the initial state of the set *queue* and all other functions in *QueueData*, *WaitingData* and *QReleasingData* are the empty set. Finally, the initialisation of *TaskQueue* can be defined as

$$\begin{array}{l} Init_TaskQueue \\ TaskQueue' \\ Init_Task \\ Init_Queue \end{array}$$

To prove the reachability of these initial states, initialisation theorems are also introduced. As the definition of the *Queue* schema and its sub-state are quite simple, these theorems are easily proved using the command *prove by reduce*. With the definition of theorem *TaskInit*, the initialisation theorem of *TaskQueue* can be proved by the following script.

Theorem 15 (TaskQueueInit)

$$\exists TaskQueue' \bullet Init_TaskQueue$$

proof [*TaskQueueInit*]
use TaskInit;
prove by reduce;

■

4.3.2 Extension

Before specifying operation schemas for the queue model, it is necessary to extend the operations for the task model satisfying the state of the queue model. Firstly, as $\Delta Task$ is overridden in order to check the state transition, we need to override $\Delta TaskQueue$ for this purpose as well.

$$\Delta TaskQueue \cong TaskQueue \wedge TaskQueue' \wedge \Delta Task$$

Generally, most task related operations do not need to update information about queues. Schema $ExtendTaskXi$ is given to extend the base state of task related schemas from $\Delta Task$ to $\Delta TaskQueue$, to show that the running task is not released from sending or receiving events and to specify that queue related states are unchanged with $\Xi Queue$.

$\frac{ExtendTaskXi}{\Delta TaskQueue}$
$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
$\Xi Queue$

With this schema, a conjunction relation can be used to easily extend task related schemas to the queue model. For example, to extend $CreateTaskN_T$ to this level, the following script can be used.

$$CreateTaskN_TQ \cong ExtendTaskXi \wedge CreateTaskN_T$$

Similarly, the following scripts can be used to extend the signature schema of $CreateTaskN_T$ to the queue level.

$\frac{ExtTaskFSBSig}{TaskQueue}$
$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$$CreateTaskN_TQFSBSig \cong ExtTaskFSBSig \wedge CreateTaskN_TFSBSig$$

To show that the precondition of $CreateTaskN_TQ$ is sufficient, theorem 16 is introduced.

Theorem 16 (CreateTaskN_TQ_vc_ref)

$$\forall \text{CreateTaskN_TQFSBSig} \mid \text{true} \bullet \text{pre CreateTaskN_TQ}$$

As two key components of this theorem, $\text{CreateTaskN_TQFSBSig}$ and CreateTaskN_TQ , are extended from task level, theorem 8 is very helpful to simplify the proving. The command “use $\text{CreateTaskN_T_vc_ref}$ ” reuses theorem 8.

Exceptionally, there are some schemas that need to update information about queue related variables. If a task is blocked by a waiting event or it has just been released from a waiting event, it is recorded by waiting or releasing functions in WaitingData and QReleasingData . Deleting or suspending it requires updating related data. For instance, when expanding the schema to delete the task in the normal case to this level, information about the target task needs to be removed from wait_snd , wait_rcv , release_snd and release_rcv .

$$\begin{array}{l} \text{DeleteTaskN_TQ} \\ \text{DeleteTaskN_T} \\ \Delta \text{TaskQueue} \\ \hline \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\ \exists \text{QueueData} \\ \text{wait_snd}' = \{\text{target?}\} \triangleleft \text{wait_snd} \\ \text{wait_rcv}' = \{\text{target?}\} \triangleleft \text{wait_rcv} \\ \text{release_snd}' = \{\text{target?}\} \triangleleft \text{release_snd} \\ \text{release_rcv}' = \{\text{target?}\} \triangleleft \text{release_rcv} \end{array}$$

Similarly to the general case of extension, the task level schema, DeleteTaskN_T , is reused to introduce constraints and maintain task information. For queue information, variables in QueueData are unchanged, and target task related pairs are removed from functions in WaitingData and QReleasingData . Domain anti-restriction operator \triangleleft is used to approach this. It excludes the set on the left-hand side of an operator from the domain of a relation. Specifically, in our case, pairs whose first element is target? are removed from the functions. If target? does not exist in the domain of the functions, nothing happens to that function.

As only one constraint is added to this new schema, which is the same as *ExtTaskFSBSig*, the signature schema and precondition theorem of *DeleteTaskN_TQ* can be given with the same strategy as the creating task case.

4.3.3 Creating and Deleting Queues

The first API function related to queue is `xQueueCreate`, it is used to add new queues into the system. Once a queue is no longer needed, it is deleted by `xQueueDelete` to release the resources. The schema *CreateQueue_TQ* and *DeleteQueue_TQ* are defined for these operations respectively.

Specifically, to create a queue, the user needs to specify its capacity (*size?*) which should be greater than 0. A new queue (*que?*), which is unknown to the system before the operation, can then be added to the set *queue*. Meanwhile, its maximum and real size can be specified by *size?* and initial size, which is 0, respectively. In addition, other information should remain unchanged before and after the operation.

In contrast, to delete a queue, we need to remove the information related to the queue from *queue*, *q_max* and *q_size*. However, we can only delete a queue, which is known to the system when no task is using it. This means no task is waiting for it and no task has just been released from a waiting event related to the queue.

4.3.4 Sending and Receiving Items

Sending an item to a queue can be represented by increasing the current size of the queue. However, the exact behaviour depends on whether the queue is full, whether there is a task waiting to receive an item from the queue, and whether the waiting task has a higher priority than the running task. It can be divided into four cases, described below. In cases where the running task is released from a waiting-to-send event, it has to continue its attempt at sending and the target queue has to be the queue which the running task attempted to send to before the waiting event.

The first is the most general case. There is space left in the queue, which means the running task can successfully send an item to the queue. Meanwhile, there are no other tasks waiting to receive an item from the queue. Therefore, the running task can send the item to the queue normally. Consequentially, after the operation, the size of $que?$ should be increased by 1; the running task should be removed from the function $release_snd$ and all other data should remain unchanged.

$ \begin{array}{l} \text{QueueSendN_TQ} \\ \Delta \text{TaskQueue} \\ que? : \text{QUEUE} \\ topReady! : \text{TASK} \\ \\ running_task \notin \text{dom } release_rcv \\ running_task \in \text{dom } release_snd \\ \Rightarrow que? = release_snd(running_task) \\ que? \in queue \\ q_size(que?) < q_max(que?) \\ que? \notin \text{ran } wait_rcv \\ \exists \text{Task} \\ queue' = queue \\ q_max' = q_max \\ q_size' = q_size \oplus \{(que? \mapsto q_size(que?) + 1)\} \\ \exists \text{WaitingData} \\ release_snd' = \{running_task\} \triangleleft release_snd \\ release_rcv' = release_rcv \\ topReady! = running_task \end{array} $
--

The schema $QueueSendW_TQ$ is introduced for the second case. The postfix W used here represents that there are waiting tasks. In detail, there is also space in the queue. However, there are tasks waiting to receive an item from the queue and the priority of the highest priority task ($topReady!$) is lower than or equal to the current running task. As a result, $topReady!$ will be woken up and recorded in $release_rcv$, as it has just been released from a waiting-to-receive event. In other respects, it is the same as the normal case. As the $topReady!$ is chosen nondeterministically from the domain of $wait_rcv$, the assistant theorem $lQueueSendW_TQ_Lemma$ is used to help with the proof.

Third, similar to the previous case, there is space in the queue and tasks wait-

ing to receive, but the priority of the highest priority waiting task (*topReady!*) is greater than the running task. The schema *QueueSendWS_TQ* represents this case. When the waiting task is woken up, it is rescheduled as the new running task. Therefore, the schema for this case uses the schema *Reschedule* to maintain task related information and a similar script to the previous case is used to manage queue related data.

In the fourth case, when the target queue is full (the schema *QueueSendF_TQ* is defined for this case), the running task is blocked. Similar to the suspending running task, this leads to rescheduling, which is achieved by using the schema *Reschedule*. Furthermore, because the running task is blocked by a waiting-to-send event, the maplet (*running_task* \mapsto *que?*) is added to the function *wait_snd*. Finally, running task related information also needs to be removed from *release_snd*. Similarly, the new running task is nondeterministically selected. The theorem *lQueueSendF_TQ_Lemma* is introduced to help with the proof.

Just like sending an item to a queue, receiving an item from a queue can be specified as decreasing the size of the queue and there are four cases for the receiving operation: (a) Normal case (*QueueReceiveN_TQ*). There are items available in the queue and no other tasks are waiting to send an item to the queue. The size of the queue is decreased by one; we remove running task related information from *release_rcv* and keep everything else unchanged; (b) Waiting case (*QueueReceiveW_TQ*). There are items available in the queue and some tasks waiting to send an item to the queue. Moreover, the priority of the highest priority waiting task is no higher than the running task. The highest priority waiting task is moved to the *ready* state and recorded in *release_snd*. The schema maintains other variables as in the first case; (c) Waiting and scheduling case (*QueueReceiveWS_TQ*). There are items in the queue and some tasks waiting to be sent. The highest priority waiting task has a higher priority than the running task and rescheduling is required. The schema *Reschedule* is used for this. Other variables are maintained in a similar way to the previous case; (d) Queue is empty case (*QueueReceiveE*). The running task is blocked by the

waiting-to-receive event. The running task is rescheduled to the *blocked* state by the *Reschedule* schema and recorded in the function *wait_rcv*. Further, the operation removes running task related information from *release_rcv* and keeps other variables unchanged.

4.4 Time Model

At this level of modelling, we take time properties into consideration. As a real-time multi-tasking system, this is crucial. For instance, the processor is shared by tasks with the same priority based on time sharing. Different tasks may also need to cooperate with each other based on the same system clock, etc. As in previous models, to illustrate the time model, we start from the basic information.

4.4.1 Basic Statements

A constant, *slice_delay* is defined to represent the unit of time for each time slice, which is specified as 1 in this case. A label, “disabled slice_delay_def”, is given to this lemma. With the `disabled` mark, this lemma is automatically omitted by Z/Eves during proofs.

$$\left| \begin{array}{l} \textit{slice_delay} : \mathbb{N} \\ \hline \langle\langle \textit{disabled slice_delay_def} \rangle\rangle \\ \textit{slice_delay} = 1 \end{array} \right|$$

Compared with the task and queue models, the base state of the time model is very simple, as it only contains four variables. First, the system clock is represented by a variable, *clock*. Second, when tasks block themselves by calling delay API functions, the set *delayed_task* is used to mark them. The key difference between the `blocked` and the `suspended` state is blocking time. Tasks in the `suspended` state can only be resumed by certain types of events. However, tasks in a `blocked` state can be woken up by both events and time. Third, to record how long a `blocked` task will be blocked, a function *wait_time* is introduced. In a single-core multi-task operating system, CPU time is divided into individual time slices; tasks with the same priority can then take one slice in turn, to share CPU time equally. In FreeRTOS, a

similar strategy is used for CPU sharing. For this purpose, *time_slice* is defined to show the number of time slices that have passed.

<i>Time</i>
$clock : \mathbb{N}$ $delayed_task : \mathbb{P} TASK$ $wait_time : TASK \rightarrow \mathbb{N}$ $time_slice : \mathbb{N}$
$\forall t : \text{dom } wait_time \bullet wait_time(t) \geq clock$

In *wait_time*, each task is mapped to the time at which it should finish waiting, so each element of the range should be greater than (after) the system clock, which is *clock*.

Because the range of *wait_time* indicates the time tasks need to wait, no members should be less than the system clock, which is *clock*.

Similar to the schema *TaskQueue*, the state schema for this level of model can be specified by combining *TaskQueue* and *Time*.

<i>TaskQueueTime</i>
$TaskQueue$ $Time$
$\langle delayed_task, \text{dom } wait_snd, \text{dom } wait_rcv \rangle \text{ partition } \text{dom } wait_time$ $delayed_task \subseteq state \sim (\{ blocked \})$

As the function *wait_time* records blocking time for each blocked task, its domain has to contain all the tasks in blocked state (i.e., *delay_task*, $\text{dom } wait_snd$ and $\text{dom } wait_rcv$). Similar to the relation between $\text{dom } wait_snd$ and $\text{dom } wait_rcv$ described in Sect. 4.3.1, once a task is blocked by an event, it cannot continue its execution and be blocked by another event. The intersection of each two of these sets should be the empty set. Therefore, we can define that *delay_task*, $\text{dom } wait_snd$ and $\text{dom } wait_rcv$ partition $\text{dom } wait_time$. As mentioned above, the set *delayed_task* marks tasks which are blocked by themselves. All tasks in *delayed_task* are in the blocked state.

Additionally, it is clear that the state of tasks in the domain of *wait_time* have to be *blocked*. The prover, however, cannot get this assertion directly. Theorem 17 is provided to guide the prover. The proof is also quite simple; we need to indicate that: (a) The tasks in the domain of *wait_time* are also tasks in *delay_task*, *dom wait_rcv* or *dom wait_snd*; (b) The state of tasks in *delay_task*, *dom wait_rcv* and *dom wait_snd* is *blocked*. After this, Z/Eves can handle the rest of the work using the proof command “*with normalization rewrite*”. Specifically, the first three proof commands in the following script are used to expand *TaskQueueTime* and *TaskQueue* schemas, as they contain the necessary information, for instance, the relation between the domain of *wait_time* and the other three sets (*delay_task*, *dom wait_rcv* and *dom wait_snd*), and the state of the tasks in these sets. The fourth and fifth commands indicate to the prover that the union of *delay_task*, *dom wait_rcv* and *dom wait_snd* equals the domain of *wait_time* and *t* can be a member of one of these three sets. The final five proof commands request that the prover rewrites the proof goal with the theorem *inPower*⁵ and then proves that the state of the task *t* is *blocked*.

Theorem 17 (rule domTime)

$$\forall \text{TaskQueueTime}; t : \text{TASK} \mid t \in \text{dom wait_time} \bullet t \in \text{state} \sim (\{ \text{blocked} \})$$

proof [*domTime*]

invoke TaskQueueTime;

invoke TaskQueue;

prove;

apply extensionality to predicate delayed_task

$\cup (\text{dom wait_rcv} \cup \text{dom wait_snd}) = \text{dom wait_time}$;

instantiate y == t;

with enabled (inPower) prove;

instantiate e == t;

instantiate e_0 == t;

instantiate e_1 == t;

with normalization rewrite;

■

⁵ $X \in \mathbb{P} Y \Leftrightarrow (\forall e : X \bullet e \in Y)$

Initialisation Initially, the system clock is 0 and in the first CPU time slice. Moreover, there is no task blocked, so the set *delayed_task* and the function *wait_time* should be empty.

$Init_Time$ <hr/> $Time'$ <hr/> $clock' = 0$ $delayed_task' = \emptyset$ $wait_time' = \emptyset$ $time_slice' = slice_delay$

Combining the schema *Init_Time* and *Init_TaskQueue*, the initialisation schema for the time model can be generated. With the help of theorem 15, the initialisation theorem for this model is also easy to prove.

4.4.2 Extension

Similar to the queue model, we need to override the schema $\Delta TaskQueueTime$ to enable the state transition check for related schemas and define an extension schema to help upgrade the schemas for the queue model to this level, which expands the base state to $\Delta TaskQueueTime$ and keeps variables in *Time* unchanged. Most schemas can be easily upgraded by a conjunction between the queue level schema and the extension schema, like the extension for the general case described in Sect. 4.3.2. However, there are schemas that need to update variables in the *Time* schema as well.

- Similar to the queue model, deleting and suspending operations need to remove the target task from *delayed_task* and *wait_time*, if it is blocked;
- Schemas that unblock a task from the blocked state also need to remove the unblocked task from the function *wait_time*;
- Finally, schemas that block the running task need to add blocking time information into the function *wait_time*.

For instance, when a queue is full, sending items to that queue blocks the running task. Originally, the API function `xQueueSend` blocked the running task for a

period of time. To simplify our model, the running task will block until a specified time ($wtime?$ in the specification). Therefore, an extra precondition is given, that the waking time is later than the current clock. The maplet, $(running_task \mapsto wtime?)$, is appended to $wait_time$ to record this. (Signature schema and the precondition theorem proof of this schema can be found on Page. 212 and in the supplementary material as well.)

$ \begin{array}{l} \text{QueueSendF_TQT} \\ \Delta \text{TaskQueueTime} \\ \text{QueueSendF_TQ} \\ wtime? : \mathbb{N} \\ \hline wtime? > \text{clock} \\ \text{clock}' = \text{clock} \\ \text{delayed_task}' = \text{delayed_task} \\ \text{wait_time}' = \text{wait_time} \oplus \{(running_task \mapsto wtime?)\} \\ \text{time_slice}' = \text{time_slice} \end{array} $

4.4.3 Delaying Tasks

In FreeRTOS, there are two API functions for delaying tasks, `vTaskDelay` and `vTaskDelayUntil`. One delays tasks for a certain period and the other delays tasks until a specific time. To simplify the model, we only modelled one of them, specifying `vTaskDelayUntil`.

To delay a task, firstly, delaying time ($wtime?$) needs to be specified as later than the system clock. As the running task is blocked, rescheduling is requested. Like other schemas requesting rescheduling, the schema *Reschedule* is used to do this. Furthermore, *delayed_task* and *wait_time* have to be updated to record this information as well.

4.4.4 Checking Delayed Tasks

Once the blocking time of a blocked task has expired, it needs to be woken automatically by the system. This is performed by the function `prvCheckDelayedTasks` in FreeRTOS. When the scheduler increases the clock ticks, this function is called to check if there are any tasks that need to be woken up. If there are, these tasks

are moved to the ready list for scheduling. Two cases are used to model this. When the priority of a woken task is no higher than the running task, it is moved from blocked state to ready state. However, when its priority is higher than the running task, rescheduling is requested.

Specifically, in our model, when tasks' blocking time expires, we unblock them in the order of priority. Therefore, the next unblocked task (*topWaiting!*) should have the earliest waking time. If there is more than one task that holds the same wake up time, it should be the highest priority task first. Depending on whether its priority is greater than the running task or not, there are two separate schemas. For normal cases, it does not need to reschedule. For task related variables, we only need to override the *state* of *topWaiting!* to *ready* and keep everything else unchanged. Meanwhile, *topWaiting!* has to be removed from one of the block related lists (i.e., *delay_task*, *wait_snd*, *wait_rcv*) and *wait_time*. In addition, we also need to update the system clock to the wake up time of *topWaiting!* to show that time has passed. This is because at this level of abstraction, we are not interested in the behaviour of the system at each time tick. Our main objective is verifying the correctness of the API functions. In this way, we can ensure that tasks are unblocked in the correct order, while focusing on the behaviour of the operation and ignoring the trivial details (at this level of abstraction) of the clock ticks.

CheckDelayedTaskN_TQT

Δ *TaskQueueTime*

topWaiting! : *TASK*

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$topWaiting! \in \text{dom } wait_time$

$\forall wt : \text{dom } wait_time \bullet wait_time(topWaiting!) \leq wait_time(wt)$

$\forall wt : \text{dom } wait_time \mid wait_time(wt) = wait_time(topWaiting!)$

$\bullet priority(topWaiting!) \geq priority(wt)$

$priority(topWaiting!) \leq priority(running_task)$

$\exists TaskData$

$state' = state \oplus \{(topWaiting! \mapsto ready)\}$

$\exists ContextData$

$\exists PrioData$

$\exists QueueData$

$wait_snd' = \{topWaiting!\} \triangleleft wait_snd$

$wait_rcv' = \{topWaiting!\} \triangleleft wait_rcv$

$$\begin{array}{l} \exists QReleasingData \\ clock' = wait_time(top\ Waiting!) \\ delayed_task' = delayed_task \setminus \{top\ Waiting!\} \\ wait_time' = \{top\ Waiting!\} \triangleleft wait_time \\ time_slice' = time_slice \end{array}$$

In contrast, if the priority of *topWaiting!* is higher than the running task, rescheduling is required. The schema *Reschedule* is used to maintain task related variables. For other variables, it is the same as the normal case.

As *topWaiting!* is selected nondeterministically, assistant lemmas, *lCheckDelayedTaskN_TQT_Lemma* and *lCheckDelayedTaskS_TQT_Lemma*, are introduced to help prove the precondition theorems of these two cases respectively.

4.4.5 Time-Sharing

As a multi-task operating system, if there are any ready tasks that have the same priority as the running task, they will share the CPU time. In FreeRTOS, this is implemented by an interrupt service routine. When a time slice passes, it will trigger an internal interrupt to check whether it is necessary to perform rescheduling. If there are other tasks holding the same priority as the running task, the next task in the ready list for that priority will be rescheduled. In our specification, the schema *TimeSlicing_TQT* is specified for this case. It nondeterministically selects the next running task from ready tasks, which have the same priority as the current running task. The schema *Reschedule* can be used for rescheduling. Then *time_slice* is increased to indicate that one time slice has passed. However, if the running task is the only task that has the highest priority, it can continue to occupy CPU time. In this case, the schema *NoSlicing_TQT* only needs to increase the *time_slice* and keeps everything else unchanged.

4.5 Mutex Model

As described in Sect. 1.2.2, semaphores and mutexes are used to manage shared resources and are special queues. There are two types of semaphores

in FreeRTOS, counting semaphores and binary semaphores. The counting semaphores allow more than one task to hold the semaphore, which is specified by its size. However, the binary semaphores are similar to mutexes, which only allow one task to hold the semaphore. The difference between binary semaphores and mutexes is that binary semaphores do not support priority inheritance when competition happens. These two types of semaphores have similar properties. To simplify the specification, we only model the binary one.

4.5.1 Basic Statement

Basic information about semaphores and mutexes is gathered in the schema *MutexData*, which defines that semaphores and mutexes are members of *QUEUE*. As a task holds a mutex, it can take the mutex repeatedly. The function *mutex_holder* is introduced to record the mutex holder for each mutex. Meanwhile, it is also necessary to know how many times the mutex holder retakes the mutex; because, it has to return the mutex the same number of times to actually return the mutex. There are three constraints for *MutexData*: (a) The intersection of *semaphore* and *mutex* should be the empty set. Although, in FreeRTOS, the mutex is treated as a special semaphore, they have different properties and operations. Therefore, we separate them into two disjoint sets. (b) The domain of the function *mutex_recursive* has to be *mutex*, as it records how many time a mutex holder repeatedly takes the mutex. (c) If a mutex is not held by a task, its *mutex_recursive* has to be 0.

MutexData

$$\begin{aligned} \text{semaphore} &: \mathbb{P} \text{ QUEUE} \\ \text{mutex} &: \mathbb{P} \text{ QUEUE} \\ \text{mutex_holder} &: \text{ QUEUE} \rightarrow \text{ TASK} \\ \text{mutex_recursive} &: \text{ QUEUE} \rightarrow \mathbb{N} \end{aligned}$$

$$\begin{aligned} \text{mutex} \cap \text{ semaphore} &= \emptyset \\ \text{dom mutex_recursive} &= \text{ mutex} \\ \forall m : \text{ mutex} \bullet m \notin \text{ dom mutex_holder} &\Leftrightarrow \text{ mutex_recursive}(m) = 0 \end{aligned}$$

To enable the priority inheritance mechanism, the schema *OriginalPrioData* is provided, which has only one function *base_priority* recording the original priority of a mutex holder.

OriginalPrioData

$base_priority : TASK \rightarrow \mathbb{N}$

Similar to queues, when a task is released from waiting to take a mutex event, it has to continue its attempt. However, because only the mutex holder can give the mutex back, it is impossible to block a task from giving back a mutex. The schema *MReleasingData* is defined to record the task which has just been released from waiting to take a mutex.

MReleasingData

$release_mutex : TASK \rightarrow QUEUE$

Based on these definitions, the constraints between them can be defined in the schema *Mutex*.

Mutex

MutexData

OriginalPrioData

MReleasingData

$dom\ base_priority = ran\ mutex_holder$

$ran\ release_mutex \subseteq mutex$

Finally, the base state of the mutex model can be defined as follows. First, semaphores and mutexes are special cases of queues, which have a maximum size of 1. Second, if a mutex is held by a task, its size should be 0. Therefore, the domain of *mutex_holder* is all the mutexes with size 0. Third, only tasks known by the system can take and hold a mutex, which means the range of *mutex_holder* is a subset of *tasks*. Priority inheritance is used to avoid a higher priority task being blocked by a lower priority mutex holder. If a mutex holder inherits a priority from another task, the new priority has to be greater than its original priority. Furthermore, only the holder of a mutex or a semaphore can return the mutex or the semaphore, so it is impossible for a task to be blocked and released by a sending event. Finally, as mutexes are a subset of queues and taking a mutex is implemented by receiving from a queue, the function *release_mutex* should be a subset of the function *release_rcv* as well.

TaskQueueTimeMutex

TaskQueueTime

Mutex

$semaphore \subseteq queue$

$\forall s : semaphore \bullet q_max(s) = 1$

$mutex \subseteq queue$

$\forall m : mutex \bullet q_max(m) = 1$

$dom\ mutex_holder = \{m : mutex \mid q_size(m) = 0\}$

$ran\ mutex_holder \subseteq tasks$

$\forall mh : ran\ mutex_holder \bullet priority(mh) \geq base_priority(mh)$

$\forall ms : mutex \cup semaphore \bullet ms \notin ran\ wait_snd \cup ran\ release_snd$

$release_mutex \subseteq release_rcv$

Initially, no semaphore or mutex exists in the system. All variables in *Mutex* should be empty at this point.

4.5.2 Extension

A similar strategy of extension to that in Sect. 4.4.2 can be used to upgrade the specification of the time model to the mutex model. The schema *ExtendTQTXi* is introduced, which expands the base state and leaves variables in *Mutex* unchanged. However, there are some operations that need to update variables in *Mutex*.

- A mutex holder cannot be deleted by the system. The resource locked by the mutex would no longer be available for other tasks, if the mutex holder is deleted before it releases the mutex.
- Semaphores and mutexes are special queues, which means the queue operations may treat a semaphore or mutex as a normal queue. However, the operations for queues are not actually correct for semaphores and mutexes. For instance, when creating a queue, it should be empty by default. Then the user can use it for communication. However, for a semaphore or a mutex, the user wants it to be full initially as this indicates that a resource is available. It is important to prevent queue operations from dealing with the semaphores and mutexes. Therefore, the constraint " $que? \notin semaphore \cup mutex$ " is added to all schemas for queue operations.

- Similarly, as the function *release_mutex* is a sub-function of *release_rcv*, it is also vital to indicate that the running task is not released from taking a mutex for the queue receiving schemas.
- Finally, the changing priority operations are the most complex extensions for this level. In the previous model, we have three cases for changing task priorities. At this level, it depends on whether: (a) the target task is holding a mutex, or not; (b) the target task is inheriting a priority, or not; (c) the new priority is higher than the inherited priority and the priority of the running task, or not. Each case can be subdivided into three more cases. Specifically, for tasks not holding a mutex, we keep *Mutex* unchanged and append the precondition to indicate the target task is not a member of the domain of *base_priority*, which equals the domain of *mutex_holder*, for the three cases, *NNotHolder*⁶, *SNotHolder* and *DNotHolder*. For tasks holding a mutex but not inheriting a priority, the schemas for the three cases (*NNotInherited*, *SNotInherited* and *DNotInherited*) have to update the information of *base_priority* as well. When the target task is a member of the domain of the base priority and is equal to its current priority, it implies that the task is not inheriting a priority. The preconditions are appended to extend the three cases of changing task priority. Finally, in the case of the target inheriting a priority from another task, (a) if the new priority is not greater than its inheriting priority, the operation updates the base priority of the target task only; (b) meanwhile, if its new priority is greater than its inheriting priority but it is lower than or equal to the running task, it is based on the normal case of changing the priority of a task and updates the base priority of the target task; (c) however, if its new priority is greater than that of the running task, its base priority is updated and rescheduling is requested; the schema *ChangeTaskPriorityS_TQT* is used to simplify the specification. The schema *InheritedN*, *InheritedU* and *InheritedS* are defined for this case.

⁶Due to the length of these schema names, we present the postfix for each schema here only; i.e. the postfix *NNotHolder* represents the schema *ChangeTaskPriorityNNotHolder_TQTM*

4.5.3 Creating and Deleting Semaphores and Mutexes

As described in Sect. 4.5.1, initially there is no semaphore or mutex in the system. FreeRTOS provides the API functions `vSemaphoreCreateBinary` and `xSemaphoreCreateMutex` to introduce a new binary semaphore and mutex to the system. Later, when these structures are no longer needed, they can be removed from the system by the API function `vSemaphoreDelete`.

Semaphores and mutexes are special queues, which have a maximum capacity of 1. The behaviour of their creating and deleting operations is similar to that of queues. In FreeRTOS, creating a semaphore or mutex actually creates a new queue with capacity 1. It then sends an item to fill the semaphore or mutex, which makes the semaphore or mutex available. Meanwhile, the API function `vSemaphoreDelete` is directly defined by `vQueueDelete`.

We also try to follow this strategy to reuse existing specifications in new schemas. This not only simplifies the definition of new schema, but also dramatically reduces the complexity of the verification, because the precondition theorem for existing schemas can be used directly in the proof of the precondition theorem for the new schema. Therefore, when defining the schemas *DeleteBinarySemaphore_TQTM* and *DeleteMutex_TQTM*, the schema *DeleteQueue_TQT* is used to manage system information before the mutex model. The deleted semaphore and mutex are then removed from related variables, *semaphore*, *mutex* and *mutex_recursive* respectively. It is worth noting that when deleting a mutex from the system, it should not be held by any task.

However, when creating operations, as the initial size of the new semaphore and mutex is full instead of empty, *CreateQueue_TQT* cannot be used. The schema *CreateBinarySemaphore_TQTM* and *CreateMutex_TQTM* add a new queue to the *queue* and set its capacity and size to 1. The queue is also added to *semaphore*, *mutex* and *mutex_recursive*. All other variables remain unchanged.

4.5.4 Taking Mutexes

As a special queue, taking a mutex can be treated as receiving an item from the mutex. Basically, we attempt to reuse specifications for the queue receiving operation in our specification for mutex taking. In Sect. 4.3.4, four cases are defined for receiving an item from a queue. However, the second and third cases are impossible for mutex taking, as they are introduced for cases where there are tasks waiting to send an item to the queue and there is no wait to send events for semaphores and mutexes. Based on the normal and empty cases of the receiving operation, the specification for the mutex taking operation can be defined. Specifically, there are two sub-cases for the normal case, depending on whether the running task previously held a mutex.

- *MutexTakeNnonInh_TQTM* If the running task did not previously hold a mutex, which means its base priority was not initialised, we need to initialise it with the value of its current priority.
- *MutexTakeNInh_TQTM* Otherwise, the base priority of the running task should not be changed.

In addition, the remaining part of these two sub-cases are the same: (a) They reuse the schema *QueueReceiveN_TQT* to receive the item from the mutex, which shows that the running task takes the mutex; (b) Set the running task as the holder of the mutex; (c) Increase the value of *mutex_recursive*; (d) Finally, remove information about the running task from *release_mutex* to enable the running task to execute other operations.

When the mutex is not available, the running task will be blocked by waiting for a receiving event. Depending on the relationship between the priorities of the running task and the mutex holder, there are also three cases.

- *MutexTakeEnonInh_TQTM* If the priority of the running task is not higher than that of the mutex holder, the mutex holder keeps its priority and the running task is replaced by a ready task with the highest priority, which is

the same as the empty case of the queue receiving operation. The schema *QueueReceiveE_TQT* is used to achieve this and running task related information is also removed from *release_mutex*.

- On the other hand, if the priority of the running task is higher than the mutex holder, the mutex holder inherits the priority of the running task and the running task is blocked.
 - *MutexTakeEInheritReady_TQTM* If the state of the mutex holder is ready, it becomes the new running task.
 - *MutexTakeEInheritHolder_TQTM* Otherwise, a task with highest priority in the ready state is selected as the new running task.

In these cases, the schema *QueueReceiveE_TQT* cannot be used, because the priority of the mutex holder is updated. However, *Reschedule* is used to manage rescheduling. Similarly, it also needs to remove information about the running task from related release functions (i.e., *release_rcv* and *release_mutex*). As the running task is blocked, data about the blocking time needs to be appended to the *wait_time* function as well. Other variables remain unchanged.

Finally, the last case, *MutexTakeRecursive_TQTM*, is for recursively taking the mutex, which is the simplest case. It increases the value of the mutex in *mutex_recursive* and keeps everything else the same.

In summary, taking mutexes has six cases in total.

4.5.5 Giving Mutexes

Similar to taking mutexes, when a mutex holder gives the mutex back, it actually sends an item back to the mutex, even if the item size is zero. Therefore, the schemas for queue sending are reused here to develop new specifications for giving mutexes. As we define the domain of *base_priority* to be equal to the range of the *mutex_holder*, when a task gives all its mutexes back, its base priority has to be removed as well. This can happen to every case of giving mutexes. The schema *basePriorityMan* is introduced to perform base priority management. It checks

whether the running task is giving its last mutex. If it is, *basePriorityMan* removes the running task related pairs from *base_priority*; otherwise, *base_priority* remains unchanged.

$\begin{array}{l} \text{basePriorityMan} \\ \Delta \text{TaskQueueTimeMutex} \\ \text{mut?} : \text{QUEUE} \\ \\ \text{running_task} \in \text{ran}(\{\text{mut?}\} \triangleleft \text{mutex_holder}) \\ \quad \Rightarrow \exists \text{OriginalPrioData} \\ \text{running_task} \notin \text{ran}(\{\text{mut?}\} \triangleleft \text{mutex_holder}) \\ \quad \Rightarrow \text{base_priority}' = \{\text{running_task}\} \triangleleft \text{base_priority} \end{array}$

The first case of giving mutexes (*MutexGiveNRecursive_TQTM*) is the simplest one, which handles the recursive return. When the mutex holder has taken the mutex several times, the mutex holder has to return the mutex the same number of times to make the mutex available for other tasks. When the value of the mutex in *mutex_recursive* is greater than 1, the schema *MutexGiveNRecursive_TQTM* decreases *mutex_recursive* by 1 each time. In addition, there are eight cases for giving mutexes. The behaviour depends on: (a) Whether there are tasks waiting to take the mutex; (b) Whether the priority of the mutex holder is inherited from other task; (c) Which priority is the highest (the base priority of the running task; the priority of the running task, the priority of the top priority ready task or the priority of the top priority waiting task). Table 4.1 and 4.2 illustrate the relationship between each case⁷.

Table 4.1 shows that based on the normal case of the queue sending operation there are three different cases. Specifically, for the first case (*NnonInh* case), there are no tasks waiting to take the mutex and the mutex holder does not inherit priority from other tasks. As the mutex holder is the current running task, we know its priority is the highest of the ready tasks. Therefore, there is no rescheduling request. We just need to send an item to the mutex to indicate that the mutex holder returns the mutex. Meanwhile, information about the mu-

⁷Due to the length of schema names, only the postfixes for each schema are presented in the table; i.e., the postfix *NnonInh* stands for the schema *MutexGiveNnonInh_TQTM*

Table 4.1: The constraints for giving mutexes (no waiting tasks)

Inh	N	Y	
Highest Prio	-	Base	Ready
Postfix	<i>NnonInh</i>	<i>NInhN</i>	<i>NInhS</i>

Table 4.2: The constraints for giving mutexes (with waiting tasks)

Inh	N		Y		
	Run	Waiting	Base	Ready	Waiting
Postfix	<i>WnonInhN</i>	<i>WnonInhS</i>	<i>WInhN</i>	<i>WInhSR</i>	<i>WInhSW</i>

tex in *mutex_holder* and *mutex_recursive* has to be reset. Finally, the schema *basePriorityMan* is used to manage the base priority of the mutex holder. The schema *MutexGiveNnonInh_TQTM* uses *QueueSendN_TQT* to send an item to the mutex. If the mutex holder inherits a priority from another task, its priority has to be reset to its original priority when it gives the mutex back. The second and third cases in this group are defined for this. The difference between them is when the original priority of the mutex holder (i.e., the running task) is lower than the priority of a ready task, rescheduling is requested. Unfortunately, the schema *QueueSendN_TQT* cannot be used, as the priority of the mutex holder needs to be modified, which is not covered by the definition of *QueueSendN_TQT*.

Similar to the last group of cases, the cases in Table 4.2 are based on waiting, and waiting and rescheduling cases of queue sending operations. There are tasks waiting to take the mutex. It is necessary to wake these tasks when the mutex holder returns the mutex. When the mutex holder does not inherit a priority, these cases are extensions from *QueueSendW_TQT* and *QueueSendWS_TQT*. Otherwise, the priority of the mutex holder needs to be revised to its original priority. Moreover, depending on the relationship between the base priority of the running task, the priority of the top priority ready task and the priority of the top priority waiting task, each case needs to decide which task should be the new running task after the operation (i.e., the highest priority task can be executed).

4.6 Summary of Interface

The preconditions for the interface, and the API function mapping can be found in Appendix C. When we define the schemas for API functions, in order to simplify the specification, we use different schemas to define the different cases of API functions. We use disjunction to connect them together to get the schema that represents the API function in FreeRTOS. The precondition for these new schemas is also obtained from the preconditions of the old schemas, which are also disjoint. For instance, the FreeRTOS API function for creating a task, `xTaskCreate`, is represented by the schema *CreateTask_T*, which has two cases, *CreateTaskN_T* and *CreateTaskS_T*, as defined above. Therefore, it is defined by these two subschemas linked by “ \vee ”. In the first case, the precondition is that *target?* is not known to the system and the priority of the new task is lower than or equal to the priority of the running task. The precondition for the second case is that *target?* is not known to the system and the priority of the new task is greater than that of the running task. Therefore, the precondition for the new schema is only that *target?* is not known to the system before the operation.

It is worth noticing that the functions in *QReleasingData* and *MReleasingData* are auxiliary functions. They are used to help specify the behaviours of queue sending, receiving, and mutex taking operations. In the implementation of FreeRTOS, when a task is woken up after being blocked during the execution of these operations, (for instance, the running task is blocked by sending an item to a full queue) the task continues any unfinished work. However, in the specification, once a task is blocked in a schema, its state will be simply transferred to *blocked*. When it is rescheduled as the running task later, it does not have to continue its unfinished job. It can perform any schema whose precondition is satisfied. Therefore, the functions in *QReleasingData* and *MReleasingData* are necessary to distinguish continuing schemas from the others. These functions do not actually represent anything in FreeRTOS. Consequently, preconditions related to them in the specifications are omitted in the tables of Appendix C.

Based on the content of these tables, it is possible to produce code-level annotations for VCC. These preconditions can be used in VCC as the content of requires clauses, “_(requires ...)”. Further, the postconditions of the schemas can also be transferred into ensures clauses, “_(ensures ...)”, of the notation of VCC. We have also verified the task-related functions with VCC, which is presented in Chap. 5.

4.7 Some Properties

Finally, there are properties of the specifications that are important for the prover, which have been verified as well. Some of these help to ensure that the specifications have the correct behaviours, the properties of the system are consistent with the API document and source code, etc. Others are used to help Z/Eves prove our model correct. These theorems may seem trivial to the human eye; however, they are particularly helpful for the prover. Therefore, in this section, we present a few of these theorems as examples. Further details can be found in the supplementary material.

1. As described above, in some schemas we need to find the task with the highest priority of all ready tasks. In these cases, ensuring the running task is not a member of the ready tasks is important, otherwise the reschedule algorithm would be chaotic. Moreover, ensuring the running task does not belong to the ready tasks is also important for the prover to prove the related properties of the task. For instance it helps to prove the theorem *TaskProperty6*.

Theorem 18 (TaskProperty3)

$$\forall Task; t : state \sim (\{ready\}) \bullet t \in tasks \setminus \{running_task\}$$

2. The variable *tasks* is used to record tasks known to the system. In other words, if the task is not recorded in *tasks*, it is unknown to the system. This theorem is helpful for the prover to determine the state of this task in *nonexistent*.

Theorem 19 (TaskProperty4)

$$\forall Task; t : TASK \setminus tasks \bullet state(t) = nonexistent$$

3. As defined by FreeRTOS, the *idle* task always has the lowest priority. If the priority of a task is greater than 0, this task cannot be the *idle* task. The command *prove by reduce* can be used to prove this.

Theorem 20 (TaskPriority5)

$$\forall Task; t : tasks \mid priority(t) > 0 \bullet t \neq idle$$

4. It is also interesting to check that the behaviour of the operation schemas are properly described. To illustrate this, we select the schema *SuspendTaskS_T* to check. This theorem will check that for any proper case of *Task*, after the *SuspendTaskS_T* operation, the old running task will be suspended and the new *running_task* has the highest priority of the ready tasks. To prove this theorem, the theorem *TaskProperty3* will be used. Following that, we apply the one-point rule to the condition. The goal can then be proved.

Theorem 21 (TaskProperty6)

$$\begin{aligned} &\forall Task \mid SuspendTaskS_T \\ &\bullet state'(running_task) = suspended \\ &\wedge (\forall t : state \sim (\{ready\}) \\ &\bullet priority(running_task') \geq priority(t)) \end{aligned}$$

proof [*TaskProperty6*]

with disabled (Δ *Task*, *Task*) *prove by reduce*;
use TaskProperty3[*t := target!*];
instantiate *t_0 == t*;
prove;

■

5. As mentioned in Sect. 2.2, the operator \oplus is commonly used to update the values of a function. However, when the domain of two operands of the operation are disjoint, the effect of this operator is to comprise an union set with the first operand and the second one, which also means that the first function is a subset of the result of the operation.

Theorem 22 (overridelsAppend)

$$\forall f, g : X \rightarrow Y \mid \text{dom } f \cap \text{dom } g = \emptyset \bullet f \subseteq f \oplus g$$

6. For a normal function, it is possible that many different elements in the domain of the function can match to one element in the range of the function. In this case, if domain subtraction is applied to this kind of function and there are other elements that can match to the same result as the elements in the subtracted set, the range of the function should not be changed.

Theorem 23 (ranUnchanged)

$$\begin{aligned} \forall f : X \rightarrow Y; a : X \mid a \in \text{dom } f \wedge f(a) \in \text{ran}(\{a\} \triangleleft f) \\ \bullet \text{ran } f = \text{ran}(\{a\} \triangleleft f) \end{aligned}$$

4.8 Summary

In this chapter, the first section revealed how and why the model is structured in the way described in this chapter. The following four sections described the model in detail, also showing how the model can be extended layer by layer. Finally, the preconditions for the API modelled in the project and some properties of the system were summarised in the last two sections.

In the next chapter, the preconditions, which were calculated in this chapter for task related operations, will be used to develop the VCC annotations and the process of verifying the FreeRTOS implementation (task related API functions) with VCC is described.

Chapter 5

VCC VERIFICATION OF FREERTOS

As mentioned in Chap. 4, the specification defined in the model for FreeRTOS can be used to develop VCC annotations to verify the implementation of FreeRTOS. This chapter illustrates our work on this. Due to time limitations, only task-related functions are verified, to demonstrate the possibility of verifying source code with the abstract model and VCC.

This chapter begins by introducing an overview of VCC and our verification, describing how we organised the specification and the source code. A basic statement of the specification is then given. Finally, annotated functions explain how to use the specification to verify the implementation of FreeRTOS. The full code is not given in the thesis; however, full details can be found in the supplementary material. An explanation of VCC annotations is given, along with a description of the verification.

5.1 Overview

As introduced in Sect. 1.3, VCC is a verifier based on the Z3 prover. It uses annotations, which are ignored by a normal C compiler, to describe the virtual model and the properties of the source code. The following annotations are most frequently used in VCC:

- The most basic annotations in VCC are *assumptions* and *assertions*. Assumptions are used so that the prover considers predicates declared in the assumptions to be logically *true* and therefore does not attempt to prove them. In contrast, the prover does attempt to prove predicates specified in assertions. However, whether the result is *true* or *false*, they are also considered as logically *true* in later proofs.
- It is also possible to define virtual variables and code, which are hidden to an ordinary C compiler, with *ghost* code. This can be used to introduce a virtual model of the specification to the implementation of the C source code.
- Constraints can be defined for data structures (concrete and virtual), which are called *Object Invariants* in VCC.
- An instance of a structure in VCC, called an *object*, has two states, *open* and *closed*. When it is in the *open* state, it is *mutable* and its constraints can be broken. Moreover, if the object is owned by the executing thread, it is writeable for the thread. However, once it is closed, all constraints specified for the data structure have to hold for that instance.
- A function contract can be specified for each function, which is similar to pre- and post-conditions for a schema in Z. This defines the behaviours of a function. VCC assumes the predicates in *requires* annotations in the contract to be *true* at the beginning of the function and verifies the predicates in *ensures* at the end of the function.

If the function with the contract is called by another function, the prover does not attempt to prove the called function. It just checks whether the preconditions of the function hold when it is called. If they do, the prover assumes the function is verified, i.e., the postconditions are satisfied.

- In VCC, parameters and global variables are not writeable for a function by default. They have to be included in *writes* annotations to inform the prover

that the function is going to modify the value of the variable. When data is *closed* and owned by the thread, it is called *wrapped*. Even if it is in a *writes* annotation, we also need to *unwrap* it to update the data of the structure.

In addition, there are also other annotations, such as *claim* and *mutable*, etc., that are not as common as the annotations described above. Further details can be found in [21]. They will be briefly described below, where they are used in the verification. It is also worth noting that the concept of hierarchical ownership for variables in VCC exists. For instance, a thread can own several structures and each structure can have its own variables and structures. To update a variable, it has to be included in the writes annotations, opened and owned by the thread.

In theory, we can directly annotate the FreeRTOS source code to verify the implementations. However, as we are only verifying the implementation of task related API functions, due to time limitations, we need to minimise the FreeRTOS source code to only include task-related code. Specifically, we create two files, *vtask.h* and *vtask.c*. In the header file, *vtask.h*, we include some essential predefinitions for verification, such as type definitions for the given sets introduced in the Z specification. This file also contains the functions declarations and their contracts. The file *vtask.c* is an annotated and minimised version of the FreeRTOS source code *task.c*, which includes all definitions of the related API functions. We also modify the API functions to simplify the verification process. We first remove any code not related to the functions verified in the Z specification, such as memory management. Next, we use equivalent code to replace some of the function calls, for instance, `portYIELD_WITHIN_API()`. The most important reason for this decision is that during function calls, the prover may believe that global variables may be changed and it will “forget” information about them. For some cases, it is possible to use claims to inform the prover that this information is not changed during function calls. However, this makes the verification extremely complex. There are also other reasons for different function replacement, which will be explained in detail below, along with the verification.

Figure 5.1: Constraints of `tskTCB`

```
1  _(dynamic_owns) typedef struct tskTaskControlBlock
2  {
3      xListItem          xGenericListItem; /*< The list that the state list
        item of a task is reference from denotes the state of that task (
        Ready, Blocked, Suspended ). */
4      xListItem          xEventListItem; /*< Used to reference a task from an
        event list. */
5      unsigned portBASE_TYPE uxPriority; /*< The priority of the task.
        0 is the lowest priority. */
6
7      _(invariant uxPriority < configMAX_PRIORITIES)
8      _(invariant \mine(&xGenericListItem))
9      _(invariant \mine(&xEventListItem))
10 } tskTCB;
11
12 _(ghost typedef tskTCB * TASK;)
```

5.2 Statement Definition

To create the annotations to verify the FreeRTOS implementation with our abstract model, we first need to transform our specification to VCC style annotations. For the task model, the first definition of the specification is the given sets of *TASK* and *CONTEXT*. Similar to our Z model, we are not interested in the memory management and the detail of context switches, which are mainly implemented by assembly language. We use a pointer of void type (i.e., `void *`) to define *CONTEXT* in a ghost annotation. However, it is not actually used in the verification. As mentioned in Sect. 1.2.1, a structure called Task Control Block (`tskTCB`) is defined in FreeRTOS to record the properties of a task, such as the priority of the task, general list item (in state lists to represent the state of the task), event list item (in event lists), etc. We use a simplified version of the `tskTCB` structure to define *TASK* in the verification, which removes memory management-related declarations and some other unrelated fields from the `tskTCB` structure.

As shown in Fig. 5.1, firstly, the priority of a task has to be lower than the maximum priority defined in the configuration file. Then, the list items of each task have to be held by the task. The task cannot be modified by a thread without it being opened, otherwise, the constraints of the system might be broken. Therefore, we state that the pointers for these list items are owned by the `tskTCB` structure, which can later transfer to the thread when it unwraps the task structure. Because of these ownership declarations, it is necessary to mark the structure as “`_(dynamic_owns)`”.

Figure 5.2: State and transition in VCC

```

1  _(ghost typedef enum{
2  nonexistent = 0, ready, blocked, suspended, running
3  } STATE;)
4
5  _(ghost \bool transition[5][5];)
6  _(def \bool Trans()
7  {
8      // nonexistent, ready, blocked, suspended, running |
9      //-----+-----
10     //{{ \false,    \true, \false, \false,  \true } | nonexistent
11     // { \true,    \false, \false, \true,   \true } | ready
12     // { \true,    \true, \false, \true,   \true } | blocked
13     // { \true,    \true, \false, \true,   \true } | suspended
14     // { \true,    \true, \true,  \true,   \false }}| running
15     return
16     transition[0][0] == \false && transition[0][1] == \true &&
17     ... &&
18     transition[4][4] == \false; // running
19 }
20 )

```

This will request the user to manage the ownership of the object components manually (e.g., manage the ownership of the generic list item, `&xGenericListItem`).

Following this, in Fig. 5.2, an enumerated type, which contains five states of tasks defined in the Z model, is used to translate the free type definition for *STATE* into VCC. In Z, a set of mappings is used to represent the transition rules, because we know all the elements of this set. Similarly here, the following boolean type two-dimensional array is used to represent this set. Each mapping included in *transition* can be matched to a boolean *true* in the array to represent that the transition is valid. Otherwise, if an element of the array is *false*, the represented transition is invalid in this case. It is worth noticing that, in VCC, the curly brackets are overridden to declare object sets. We cannot directly use it to initialise the array. An assistant *logic function* is used to do this. The logic function in VCC is a virtual function, which contains a specification for a given activity. It is normally used to perform some logic checking. As a virtual function, it should not have any operations about memory writes. Specifically, the function `Trans()` returns logical *true*, once the value of each element of *transition* satisfies the table included in the comments.

Based on these definitions, the virtual structure, `FreeRTOS`, for the base model can be given as below, which organises all the essential elements of verification.

The relationship between these elements and the concrete variables in the original source code can be specified in this virtual structure. Their constraints are also defined here (see Fig. 5.3).

Figure 5.3: FreeRTOS structure

```

1  _(ghost _(dynamic_owns) typedef struct{
2  \bool tasks[TASK];
3  STATE state[TASK];
4  //CONTEXT phys_context;
5  //CONTEXT log_context[TASK];
6
7  \natural priority[TASK];
8
9  //READY
10  _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE)
    ==>
11      (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
        pxReadyTasksLists[((tskTCB *) t)->uxPriority] &&
12      ((tskTCB *) t) != pxCurrentTCB) <==> state[t] == ready))
13  //BLOCKED
14  _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE)
    ==>
15      (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
        xDelayedTaskList1 ||
16      ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
        xDelayedTaskList2) <==> state[t] == blocked))
17  //SUSPENDED
18  _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE)
    ==>
19      (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
        xSuspendedTaskList <==> state[t] == suspended))
20  //RUNNING
21  _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE)
    ==>
22      (t == (TASK) pxCurrentTCB <==> state[t] == running))
23  _(invariant ((xList *) pxCurrentTCB->xGenericListItem.pvContainer) ==
        &pxReadyTasksLists[pxCurrentTCB->uxPriority])
24  //NONEXISTENT
25  _(invariant \forall TASK t; (t->\closed && xSchedulerRunning !=
        pdFALSE) ==>
26      (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
        xTasksWaitingTermination ||
27      t == NULL) <==> state[t] == nonexistent))
28
29  _(invariant \forall TASK t; xSchedulerRunning != pdFALSE ==> (state[t]
        != nonexistent <==> tasks[t]))
30  _(invariant \forall TASK t; tasks[t] ==> \mine(t))
31  _(invariant \forall TASK t; tasks[t] ==> t->\closed)
32  _(invariant \forall TASK t; tasks[t] ==> priority[t] == t->uxPriority)
33  _(invariant \forall TASK t; tasks[t] ==> state[t] <= 4)
34  _(invariant \mine(\embedding(& xIdleTaskHandle)))
35  _(invariant \mine(\embedding(& pxCurrentTCB)))
36  _(invariant xSchedulerRunning != pdFALSE ==> (tasks[xIdleTaskHandle]
        && tasks[pxCurrentTCB]))
37  _(invariant xSchedulerRunning != pdFALSE && xIdleTaskHandle != NULL
        ==> priority[xIdleTaskHandle] == 0)
38  _(invariant \forall TASK t; xSchedulerRunning != pdFALSE && state[t]
        == ready ==> priority[pxCurrentTCB] >= priority[t])
39
40  } * FreeRTOS;)

```

This structure translates the schema *Task* into a VCC style base model. The map `tasks` from `TASK` to `\bool` in Ln. 2 is used to represent *tasks* in the Z model. In VCC, map type is a type similar to an array, which can only be used in ghost

code. It maps a value of the type in the square bracket to a value of the type before the variable name. In this case, `\bool tasks[TASK]` maps each task to a boolean type value to indicate whether the task is known by the system or not. Similarly, a map from TASK to STATE and a map from TASK to natural numbers are used to represent *state* and *priority*, respectively. As we mentioned above, the context related variables are not used in the verification. They are declared, but commented out, which has no effect on our model.

After the declarations of the virtual variables, the links between them and the real variables are specified. Specifically, lines 9 to 27 in Fig. 5.3 show the relationship between the virtual state in the specification to the real state lists in FreeRTOS. For instance, FreeRTOS declares an array of `tskTCB` lists to store ready tasks. An invariant (Ln. 9 to 12) is used to specify this, which defines that all the tasks in a list of `pxReadyTasksLists` are exactly those that are in a ready state, i.e., `state[t] == ready`.

Following that, line 29 translates a constraint from the *tasks* schema. It states that all tasks which are not in a `nonexistent` state are known by the system (i.e., they map to *true* in `tasks`). However, if a task is known to the system, we can also say it belongs to the structure `FreeRTOS` (Ln. 30), as this structure describes the base of the whole system and is the root of the hierarchy ownership tree. When a thread unwraps the `FreeRTOS` structure from its *wrapped* state, the thread automatically obtains ownership of all tasks, which makes these tasks in the *wrapped* state at the moment. Otherwise, to verify the properties of the tasks for a function, it would have to have all the tasks in the parameter list of the function, which is not possible. A binding from the virtual priority to the real priority of each task is stated in Ln. 32.

Next, the handlers of the `idle` task and `pxCurrent` task are declared as global variables in FreeRTOS. If we want to specify the properties of these two handlers, we need to transfer their ownership to the `FreeRTOS` structure. To do this, we first need to transfer ownership of their container to `FreeRTOS`. In VCC, all global fields are included in a virtual container. Due to their definition in Ln. 30, setting

their entry in `tasks` to `true` transfers ownership to the FreeRTOS structure. The constraint that the priority of `idle` task has to be 0 can then be specified. The last invariant shows that the priority of the running task is the highest of all ready tasks, which is also the last constraint of the *Task* schema in Z.

Finally, another logic function, `excList()` (details of which can be found from the supplementary material), is given to ensure that the state lists for tasks are exclusive (i.e., one list cannot be used as more than one state).

5.2.1 Translation from Z to VCC

In summary, to translate types from Z to VCC, the following rules can be considered as a guide.

- Given sets in Z can be translated to point types or structures defined in C, e.g. *TASK* is translated to `tskTCB *`.
- Generally, functions can be translated into two ways, (a) As functions are sets of mappings in Z, they can be naturally translated to maps in VCC, e.g. the function *priority* is translated to `\natural priority[TASK]`; (b) If a function is defined from a given set which has been translated as a structure, to another type, the value of this function can be defined as a ghost field in the structure, e.g. the function *priority* can also be translated as a ghost field in the structure `tskTCB`. We chose the former way to translate *priority*, as it is closer to our Z specification.
- Sets in Z can be understood as power sets of their own type. Therefore, in VCC a set can be defined by a boolean type map that maps each element of its type into a boolean type value to indicate whether the element is a member of the set e.g. *tasks* is translated to `\bool tasks[TASK]`.
- Simple free types can be translated as enumerated types e.g. *STATE* defined in Fig. 5.2.
- Finally, static functions can be defined as two-dimensional arrays of Booleans, e.g. *transition* is represented by `\bool transition[5][5]`.

VCC provides plenty of operators for the ghost fields and logical operations, including quantifiers (\forall and \exists), implication, etc. Once all the variables in our Z specification have been translated to corresponding VCC variables or C variables, we can use operators defined in VCC and C for these types of variables to represent predicates used in the Z specification. For example,

- For functions, which are translated as mappings, we can simply apply the function to an element in the domain to obtain the value of the function, e.g. `state[t]` used in Ln. 12 of Fig. 5.3.
- Similarly, for sets in Z,
 - we can calculate the result of a map applied to an element to check whether the element is in the set, e.g. `tasks[t] == \true;`
 - we can also set the value of a map for an element to `\true` or `\false` to add or remove the element to or from the set respectively.
- As VCC supports quantifiers, we can directly translate such predicates from Z specification to VCC annotations, e.g. we use plenty of quantifier `\forall` in Fig. 5.3.

More detailed explanation will be provided during the explanation of each API functions in following sections.

Following this, we can use the preconditions and postconditions verified in the Z model to verify the implementation of FreeRTOS task API functions.

5.3 Creating Tasks

To verify the API function for creating tasks in FreeRTOS, `xTaskGenericCreate`, we first append two ghost parameters, `FreeRTOS` and `newTask`, to the function parameter list. These represent the virtual model and the new task created by the function respectively. The contract of the API function can then be specified as in Fig. 5.4. The first annotation, `_(updates FreeRTOS)`, is used to tell the prover that the structure is *wrapped* before and after the operation. It may be modified by the operation (i.e., `FreeRTOS` is writeable and in a *writes* clause). It is

actually equivalent to `_(requires \wrapped(FreeRTOS)), _(writes FreeRTOS)` and `_(ensures \wrapped(FreeRTOS))`.

As our specification in Z assumes that the scheduler is always executing (see Sect. 4.2.1), it is stated for VCC verification as well. Following this, the logic function `excList()` is used to ensure that all items in the state list are exclusive when the API function starts. Because the value of `newTask` and `pxCreatedTask` will be updated to the handler of the created task, they have to be included in the writes clause of the VCC annotation. In addition, the postconditions can be specified with ensures clauses.

In the Z specification, we separate the creating operation into two cases, normal and rescheduling. The postconditions for these need to be mixed together in the same way as the postconditions for the API function. VCC provides `\old()` function to obtain the pre-state of a variable, which is used here to refer to the old value of `pxCurrentTCB`. Then the priority of the old running task can be accessed. Comparing the new priority of the new task and the priority of the old running task, the difference between the two cases can be specified separately. In the normal case of task creation (see Ln. 9 to 11), the priority of the new task is less or equal to the priority of the old running task. The created task is set to the ready state. Otherwise (see Ln. 12 to 14), it requests rescheduling and replaces the value of `pxCurrentTCB`. The state of the created task and the old running task should also be set to `running` and `ready`, respectively. For both cases, the priority in the parameter list should be assigned to the new task.

It is worth noting that in the Z model, we do not restrict the maximum value of the priority. In FreeRTOS, however, the maximum value of priority, `configMAX_PRIORITIES`, is defined in the configuration file. If the new priority is not less than the maximum priority, it is set to the maximum priority of the system (i.e., `configMAX_PRIORITIES-1`).

The only precondition for creating tasks in the Z model is that the created task is not in the system before the operation. As it is newly created during the operation,

Figure 5.4: Contract for creating tasks

```

1  signed portBASE_TYPE xTaskGenericCreate( pdTASK_CODE pxTaskCode, const
    signed char * const pcName, unsigned short usStackDepth, void *
    pvParameters, unsigned portBASE_TYPE uxPriority, xTaskHandle *
    pxCreatedTask, portSTACK_TYPE *puxStackBuffer, const xMemoryRegion
    * const xRegions _(ghost FRTOS FreeRTOS) _(ghost TASK *newTask) )
2
3  _(updates FreeRTOS)
4  _(requires xSchedulerRunning == pdTRUE)
5  _(requires excList())
6
7  _(writes newTask, pxCreatedTask)
8
9  _(ensures \result == pdPASS ==>
10     uxPriority <= \old(pxCurrentTCB)->uxPriority ==>
11     (FreeRTOS->state[*newTask]) == ready)
12  _(ensures \result == pdPASS ==>
13     uxPriority > \old(pxCurrentTCB)->uxPriority ==>
14     (FreeRTOS->state[(TASK)\old(pxCurrentTCB)]) == ready && (FreeRTOS->
        state[(TASK)pxCurrentTCB]) == running)
15  _(ensures \result == pdPASS ==> FreeRTOS->priority[*newTask] ==
16     (\natural)(uxPriority < configMAX_PRIORITIES ? uxPriority :
        configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U))
17  _(ensures \result == pdPASS ==> pxCurrentTCB->\closed)
18  _(ensures \result == pdPASS ==> \fresh(*newTask))
19  _(ensures \forall TASK t; (FreeRTOS->tasks[t] &&
20     \old(FreeRTOS->state[t]) != FreeRTOS->state[t]) ==>
21     transition [\old (FreeRTOS->state [ t ] )][FreeRTOS->state [ t ]])

```

it is not easy to describe with requires clauses; therefore, it is stated in the post-condition instead. The function, `\fresh()` is provided by VCC to indicate that the object is freshly allocated. Finally, the last predicate requests VCC to check that the state transitions obey transition.

To improve the efficiency of the prover, VCC does not make forward inferences from the precondition by default. For instance, in the precondition, we state that `FreeRTOS` is *wrapped*, which means all the constraints have to hold at the beginning of the function. In the constraints of the `FreeRTOS` structure, we state that `pxCurrentTCB` is owned by `FreeRTOS` and all the tasks owned by `FreeRTOS` have to be closed. From these definitions, we can easily conclude that `pxCurrentTCB` has to be in the *closed* state, when `FreeRTOS` is *wrapped*. However, VCC cannot obtain this result automatically. Therefore, the script in Fig. 5.5 is used to show that `pxCurrentTCB` is *closed* at the beginning along with some other properties which may be helpful for later verification. After the declaration of some local variables, the first statement calls the function `prvInitialiseTCBVariables`, which allocates memory for the new task and its stack. To simplify the verification, we replaced it with `malloc`. An annotation is added here to make sure that the prover

Figure 5.5: Creating tasks pre-verification

```
1  _(assert \wrapped(FreeRTOS))
2  _(assert \inv(FreeRTOS))
3  _(assert xSchedulerRunning == pdTRUE)
4  _(assert FreeRTOS->tasks[pxCurrentTCB])
5  _(assert pxCurrentTCB \in FreeRTOS->\owns)
6
7  _(assert \forall TASK t; FreeRTOS->tasks[t] ==> \inv(t))
8  _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->priority[t] <
    configMAX_PRIORITIES)
9
10 _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->state[t] <= 4)
11 _(assert pxCurrentTCB->\closed)
12
13 signed portBASE_TYPE xReturn;
14 tskTCB * pxNewTCB;
15
16 \pxNewTCB = prvInitialiseTCBVariables( usStackDepth, puxStackBuffer );
17 pxNewTCB = (tskTCB *) malloc(sizeof(tskTCB));
18 _(assert pxNewTCB != pxCurrentTCB)
```

ensures that the handler for the new task is not equal to `pxCurrentTCB`, otherwise, when it is updated, the prover may confuse these two tasks as they refer to the same task.

As shown in Fig. 5.6, when `pxNewTCB` is successfully assigned memory, two assertions (Ln. 2 and 3) are added to help the prover ensure that `pxNewTCB` is successfully and newly allocated by the function. It is then necessary to complete the details of the created task. As we only focus on the functions we specified in the Z model, we remove all unrelated code and only keep functions related to the priority and the generic list item of the new task, which sets the new priority for the task and places it in the proper position in the ready task lists. The task is then physically created. The rest of the creation function manages the system state for the new task. In VCC, before verifying system properties related to `pxNewTCB`, the structure has to be wrapped in advance. To do this, the ownership of all `pxNewTCB` components has to be taken by the structure. However, for the new created task, the thread keeps the ownership of `pxNewTCB` and all its components, i.e., the generic and event list items, which need to be transferred to `pxNewTCB`. Two ghost statements in Ln. 13 and 14 are used to perform this and then the statement `_(wrap pxNewTCB)` is used to wrap it. In the critical section, the original source code manages the system states according to the properties of the new tasks, for instance, the maximum priority of existing tasks, the total

number of the tasks in the system, etc. However, they are not included in our reduced version. The variable `xReturn` is then set to `pdPASS` to indicate that the new task has been successfully created and added to the system. In cases of failure to allocate memory to `pxNewTCB`, the function does nothing but set the value of `xReturn` to `errCOULD_NOT_ALLOCATE_REQUIRED_MEMORY`. In this case, an assertion is needed to remind the prover that the invariants of `FreeRTOS` still hold at that time.

Figure 5.6: Creating tasks verification part-1

```

1  if ( pxNewTCB != NULL ) {
2    _(assert pxNewTCB)
3    _(assert \fresh(pxNewTCB))
4
5    // prvInitialiseTCBVariables( pxNewTCB, pcName, uxPriority, xRegions,
6      usStackDepth );
7    pxNewTCB->uxPriority = (uxPriority < configMAX_PRIORITIES ? uxPriority
8      : configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U);
9    pxNewTCB->xGenericListItem.pvContainer = &pxReadyTasksLists[pxNewTCB->
10     uxPriority];
11   _(assert \writable(&(pxNewTCB->xGenericListItem)))
12   _(assert \writable(&(pxNewTCB->xEventListItem)))
13   _(wrap &(pxNewTCB->xGenericListItem))
14   _(wrap &(pxNewTCB->xEventListItem))
15   _((ghost pxNewTCB->\owns = (\objset) {&(pxNewTCB->xGenericListItem)}))
16   _((ghost pxNewTCB->\owns += &(pxNewTCB->xEventListItem)))
17   _(assert !(FreeRTOS \in pxNewTCB->\owns))
18   _(wrap pxNewTCB)
19
20   if ( ( void * ) pxCreatedTask != NULL ) {
21     *pxCreatedTask = ( xTaskHandle ) pxNewTCB;
22   }
23   taskENTER_CRITICAL();
24   { //...
25     xReturn = pdPASS;
26   }
27   taskEXIT_CRITICAL();
28 } else {
29   _(assert \inv(FreeRTOS))
30   xReturn = errCOULD_NOT_ALLOCATE_REQUIRED_MEMORY;
31 }

```

When `xReturn` is set to `pdPASS`, the system needs to check whether rescheduling is requested and the ghost code for managing system states is placed here, as shown in Fig. 5.7. Firstly, before the system checks for rescheduling, we unwrap the `FreeRTOS`, set the priority of `pxNewTCB`, temporarily set the state of `pxNewTCB` to `ready`, add it to tasks and transfer its ownership to `FreeRTOS`. Finally, we assign `pxNewTCB` to the virtual parameter `* newTask`. Some assertions are used here to ensure that the assignments of the virtual code work. The system then checks the relationship between the priority for the new task and the priority for the running

task to determine whether the system needs to reschedule or not.

As a global volatile variable, `pxCurrentTCB` can be accessed and updated by a thread in VCC which does not hold its ownership. To access `pxCurrentTCB`, an *atomic_read* annotation is applied to the container of `pxCurrentTCB`, which is `\embedding(&pxCurrentTCB)` in this case. When the priority of `pxCurrentTCB` is lower than the new priority, the system needs to reschedule. We also modify the state of `pxNewTCB` and `pxCurrentTCB` to `running` and `ready`, respectively. A system function call to `portYIELD_WITHIN_API` is used to perform rescheduling, which is replaced by directly assigning the handler of the new task to `pxCurrentTCB` in our reduced version.

Similar to accessing `pxCurrentTCB`, in order to modify it, we need to apply the *atomic* annotation to its container. It is worth noting that during the atomic operation, information about global variable `FreeRTOS` is lost. Ideally, this can be handled by the *claim* annotation, which claims that the state of `FreeRTOS` is kept unchanged. However, this complicates the verification and due to the time limits of the project, we simply assert that the invariant of `FreeRTOS` holds before the atomic operation and assume them after the operation. In addition to these, some assertions are used to ensure that the properties for `FreeRTOS` are well maintained. It can then be wrapped. Finally, the assumption of the logic function, `Trans()`, is used to set the value of the array `transition`.

5.4 Deleting Tasks

Similar to creating tasks, to prove the API function for deleting tasks, `vTaskDelete`, two extra ghost parameters are appended, `FreeRTOS` and `topReady`. As shown in Fig. 5.8, the first four annotations in the function contract are similar to used in creating tasks. They describe the basic properties of the system. In addition, the preconditions specific to deleting tasks can be defined following our Z specification, which states the target task has to be a member of *tasks* but not *idle* and the *topReady!* is the highest priority ready task in the system. In the source code, the target task is named `pxTaskToDelete`. We first state that at the beginning of

Figure 5.7: Creating tasks verification part-2

```

1  if( xReturn == pdPASS )
2  {
3      _(unwrapping FreeRTOS) {
4          _(assert \inv(FreeRTOS))
5
6          _(ghost {
7              FreeRTOS->priority[(TASK) pxNewTCB] = pxNewTCB->uxPriority;
8              FreeRTOS->state[(TASK) pxNewTCB] = ready;
9              FreeRTOS->tasks[(TASK) pxNewTCB] = \true;
10             FreeRTOS->\owns += pxNewTCB;
11             *newTask = pxNewTCB;
12         })
13
14         _(assert FreeRTOS->state[(TASK) pxNewTCB] == ready)
15         _(assert FreeRTOS->state[(TASK) pxCurrentTCB] == running)
16         _(assert FreeRTOS->priority[(TASK) pxNewTCB] == pxNewTCB->uxPriority
17             )
18         _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *) t !=
19             pxNewTCB) && xSchedulerRunning != pdFALSE) ==>
20             (((xList *) ((tskTCB *) t)->xGenericListItem.pvContainer) == &
21             pxReadyTasksLists[((tskTCB *) t)->uxPriority] &&
22             ((tskTCB *) t) != pxCurrentTCB) <=> FreeRTOS->state[t] ==
23             ready))
24         ...
25     }
26     if( _(atomic_read \embedding(&pxCurrentTCB))
27         pxCurrentTCB->uxPriority < uxPriority )
28     {
29         _(ghost {
30             FreeRTOS->state[(TASK) pxCurrentTCB] = ready;
31             FreeRTOS->state[(TASK) pxNewTCB] = running;
32         })
33         _(assert FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready)
34         //portYIELD_WITHIN_API();
35         _(atomic \embedding(&pxCurrentTCB)){
36             pxCurrentTCB = pxNewTCB;
37             _(bump_volatile_version \embedding(&pxCurrentTCB))
38         }
39     }
40     _(assert \old(pxCurrentTCB)->uxPriority >= uxPriority ==> (FreeRTOS
41         ->state[(TASK) pxNewTCB] == ready && FreeRTOS->state[(TASK)
42         pxCurrentTCB] == running))
43     _(assert \old(pxCurrentTCB)->uxPriority < uxPriority ==>
44         (FreeRTOS->state[(TASK) pxNewTCB] == running &&
45         pxNewTCB == pxCurrentTCB &&
46         FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready
47         )
48     )
49     _(assert FreeRTOS->priority[(TASK) pxNewTCB] == pxNewTCB->uxPriority
50         )
51     _(assert FreeRTOS->priority[(TASK) \old(pxCurrentTCB)] == \old(
52         pxCurrentTCB)->uxPriority)
53     _(assert FreeRTOS->tasks[(TASK) \old(pxCurrentTCB)])
54     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *) t !=
55         pxNewTCB) && xSchedulerRunning != pdFALSE) ==>
56         (((xList *) ((tskTCB *) t)->xGenericListItem.pvContainer) == &
57         pxReadyTasksLists[((tskTCB *) t)->uxPriority] &&
58         ((tskTCB *) t) != pxCurrentTCB) <=> FreeRTOS->state[t] ==
59         ready))
60     ...
61     ...
62     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> t->\closed)
63     } //wrapping FreeRTOS
64 }
65 _(assume Trans())
66 return xReturn;

```

Figure 5.8: Contract for deleting tasks

```

1  _(updates FreeRTOS)
2  _(requires \mutable(&xSchedulerRunning))
3  _(requires xSchedulerRunning == pdTRUE)
4  _(requires excList())
5
6  _(requires FreeRTOS->tasks[(TASK) pxTaskToDelete])
7  _(requires (tskTCB *)pxTaskToDelete != (tskTCB *)xIdleTaskHandle)
8
9  _(requires FreeRTOS->tasks[topReady])
10 _(requires FreeRTOS->state[topReady] == ready)
11 _(requires \forall TASK rts;
12     (FreeRTOS->tasks[rts] && FreeRTOS->state[rts] == ready)
13     ==> FreeRTOS->priority[topReady] >= FreeRTOS->priority[rts])
14 _(requires topReady != (TASK) pxTaskToDelete)
15 _(requires topReady != (TASK) pxCurrentTCB)
16
17 _(ensures ((TASK) pxTaskToDelete)->\closed)
18 _(ensures pxTaskToDelete != NULL ==> ! FreeRTOS->tasks[(TASK)
19     pxTaskToDelete])
19 _(ensures pxTaskToDelete == NULL ==> ! FreeRTOS->tasks[(TASK) \old(
20     pxCurrentTCB)])
20 _(ensures pxTaskToDelete == NULL ==> (TASK) pxCurrentTCB == topReady)
21 _(ensures \forall TASK t; (FreeRTOS->tasks[t] &&
22     \old(FreeRTOS->state[t]) != FreeRTOS->state[t]) ==>
23     transition[\old(FreeRTOS->state[t])][FreeRTOS->state[t]])

```

the function, `pxTaskToDelete` has to be in `tasks` and not equal to the handler of the idle task, `xIdleTaskHandle`. Then, `topReady` has to be in `tasks` with a state of `ready` and have the highest priority of the ready tasks. Moreover, it is also important to state that the `topReady` must not be the same as the running task or the task to be deleted. Otherwise, the deleted task can be set to the new running task after the operation, which is not sensible.

Compared to creating tasks, the postconditions for deleting tasks are quite simple. They ensure that the target task is removed from the system by checking the value of the `tasks`. When the deleted task is the old running task, `topReady` is used to replace the running task. It needs to be noted that in FreeRTOS, if the handler of the deleted task in parameter list is `NULL`, the running task will be deleted. Like creating tasks, the last postcondition ensures that the state transfer for all tasks follows the restriction defined by `transition`.

Again, similar to creating tasks, we firstly need to ensure the prover retains the necessary information about the system. Before entering the critical section, we unwrap `FreeRTOS` (see Ln. 347 on Page. 348), as the deleted task is already owned by `FreeRTOS`. To delete it, `FreeRTOS` has to be opened. As mentioned above,

FreeRTOS defines a `NULL` value for the target task, which indicates that the target task is the running task. This is because the macro `prvGetTCBFromHandle` is used to obtain the control block for a task from its handle.

Meanwhile, the macro returns the control block of the running task, while the handle is `NULL`. To simplify the code, FreeRTOS sets the value of the parameter `pxTaskToDelete` to `NULL`, if it is equal to `pxCurrentTCB`. It then assigns the task obtained from `pxTaskToDelete` to `pxTCB`, which represents the task to be deleted afterwards. These two steps require access to the global volatile variable, `pxCurrentTCB`. The atomic read declaration is needed for this. As the virtual structure `FreeRTOS` has been unwrapped, the ownership of tasks in the system is now transferred to the thread and hence they are now wrapped. Assertion (in Ln. 370 on Page. 349) is used to ensure that `pxTCB` is wrapped at that time, and can then be unwrapped.

The function, `uxListRemove`, is used to remove `pxTCB`'s generic list item from the system. Furthermore, it can be added to the list `xTasksWaitingTermination` to mark that it needs to be removed. Then the idle task will make the deletion when it is executed. It is also necessary to remove the event list of the target task, if it is in any event list. Three assertions are used to ensure these operations perform properly. Following this, a piece of ghost virtual code (between Ln. 403 and Ln. 406 on Page. 349) is used to maintain the system state. The state of `pxTCB` is set to `nonexistent`, and the value of `pxTCB` in `tasks` is set to `\false` as well. Exiting the critical session, the API function checks whether the running task has been deleted. If it has been, rescheduling is required. Specifications and code similar to creating task are used here to perform rescheduling. Again, due to the atomic operation, we assert and assume the necessary properties of `FreeRTOS` before and after atomic access. It can then be wrapped. Finally, assumption of `Trans` is used to initialise the array `transition`, which can be used for the postcondition check, at a later point.

5.5 Getting and Setting Priority

The API function, `uxTaskPriorityGet`, is the simplest to verify (see Ln. 508-539 on Page. 351). It keeps everything unchanged and returns the priority of the target task. The precondition states that the target task (i.e., `pxTask`) is known to the system and other general system preconditions, such as the one to create tasks. The postcondition ensures that the returned value, `\result`, equals `FreeRTOS->priority[(TASK) pxTask]`.

To set the priority of a task, the API function `vTaskPrioritySet` can be used. To verify this, similar to deleting a task, the extra virtual parameters `FreeRTOS` and `topReady` are appended. However, in this case, the details of `topReady` remain undefined at the beginning. The key reason for this is that `vTaskPrioritySet` requests a system reschedule, when the priority of the running task is reduced. In this case, the running task may still be scheduled as the new running task again. Therefore, the definitions for the details of `topReady` are specified when rescheduling is required.

As well as the four general preconditions, there are three extra annotations for the precondition (see Ln. 551-560 on Page. 351). These state that `topReady` and `pxTask` are in the system and that when the target task is the idle task, the priority has to be 0. The postconditions for the `vTaskPrioritySet` (see Fig. 5.9) first state that the target task has to be closed and known by the system. The key feature for this API function is then to update the priority of the target task. An annotation is used to ensure that this has been done successfully. In our specification, there are two cases where the system needs to be rescheduled, (a) when the target task is in the ready state and obtains a higher priority than the running task; and (b) when the priority of the running task is reduced to less than one of the ready tasks. The following two predicates are specified to check that in these two cases the system is rescheduled properly. Finally, it is verified that the state transitions are valid for this operation.

Similarly, some assertions are used to show basic properties. After the dec-

Figure 5.9: Postconditions for priority setting

```

1  _(ensures ((TASK) pxTask)->\closed)
2  _(ensures FreeRTOS->tasks[(TASK) pxTask])
3  _(ensures FreeRTOS->priority[(TASK) pxTask] ==
4    (\natural)(uxNewPriority < configMAX_PRIORITIES ? uxNewPriority :
      configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U))
5  _(ensures (FreeRTOS->state[(TASK) pxTask] == ready && FreeRTOS->priority
6    [(TASK) pxTask] > \old(pxCurrentTCB)->uxPriority) ==>
7    (FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready && FreeRTOS->
8    state[(TASK) pxTask] == running)
9  )
10  _(ensures ((pxTask == NULL || (tskTCB *) pxTask == \old(pxCurrentTCB))
11    && !(
12      \forall TASK t; FreeRTOS->state[t] == ready ==> FreeRTOS->
13      priority[(TASK) pxTask] >= FreeRTOS->priority[t]
14    )) ==> (FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready &&
15      FreeRTOS->state[(TASK) pxCurrentTCB] == running &&
16      \old(FreeRTOS->state[(TASK) pxCurrentTCB]) == ready)
17  )
18  _(ensures \forall TASK t; (FreeRTOS->tasks[t] && \old(FreeRTOS->state[t]
19    ) != FreeRTOS->state[t]) ==> transition[\old(FreeRTOS->state[t])][
20    FreeRTOS->state[t]])

```

laration, the API function first checks whether the new priority is less than `configMAX_PRIORITIES` (see Ln. 594 on Page. 352). If it is not, it is set to the maximum priority. An assertion is added to ensure this. It then enters the critical session to update the priority of the target task and manage the system states. In this case, the `FreeRTOS` structure needs to be unwrapped. It transfers the ownership of `pxTCB` to the threads. As in the case of deleting tasks, if `pxTask` is equal to `pxCurrentTCB`, it is set to `NULL`. With the macro `prvGetTCBFromHandle` the task control block for the target task is obtained and assigned to `pxTCB`. Atomic read annotations are also used for reading from `pxCurrentTCB`. The functions use a local variable, `uxCurrentPriority` to record the old priority of the target task. After obtaining the control block for the target task, its priority is also accessible and assigned to `uxCurrentPriority`.

Afterwards, the local variable `xYieldRequired` is set to `pdTRUE` to indicate that rescheduling is requested, according to the relation between the old and new priority of the target task (see Ln. 641-663 on Page. 353). Specifically, if it is increased and the target task is not the running task, or the target task is the running task and its priority is decreased, it indicates that there may be a ready task that has a higher priority than the running task. `xYieldRequired` is set. Following this, the function starts to update the priority of `pxTCB` and manage the related system variables. At

this point, some assertions are inserted into the source code to verify that `pxTCB` is wrapped. We then unwrap it to allow the thread to update the priority of `pxTCB`. The priority of the target task is also updated to the mapping `priority` of the virtual model. Furthermore, if a task is in the ready or running state, its generic list item should be placed in one of the ready lists, according to its priority. Therefore, the target task has to be replaced in the new list based on its new priority, which helps the system perform rescheduling correctly.

During this process, we first need to inform the prover that while `pxTCB`'s generic list item is placed in one of `pxReadyTasksLists`, it is in the ready or running state (see Ln. 702 on Page. 354). Moreover, to replace its generic list item to the new position of the `pxReadyTasksLists`, `&(pxTCB->xGenericListItem)` is unwrapped.

Finally, assertions are inserted to verify that after this process, the target task is still in the mapping `tasks` and its generic list item is replaced in the correct position. The last part of the function is used to reschedule the system, if `xYieldRequired` is set to `pdTRUE` (see Ln. 726 on Page. 354). Similar to the previous API functions, `portYIELD_WITHIN_API` is used to perform this. It is replaced with ghost code as well. The assertions and assumptions are used to ensure the operation works. The only thing different here is that we need to specify the details of `topReady` above the ghost code. Because the API function requests rescheduling when the priority of the running task is reduced, it can still have the highest priority compared to the other ready tasks. Therefore, `topReady` can be `pxCurrentTCB` for this API function. Further, it has to have the highest priority compared to the other ready tasks. If the running task still has the highest priority, it is set to `topReady`. Like the previous API functions, the atomic access has lost some of the information about `FreeRTOS`, so we simply assume that they hold. It can then be wrapped before exiting the critical session. Finally, the assumption for `Trans()` is inserted.

5.6 Suspending and Resuming Tasks

In the FreeRTOS task model in Z, we also verify the API functions for suspending (see Page. 355-358) and resuming (see Page. 358-359) tasks, which are sim-

ilar to deleting and creating tasks. The key difference between suspending and deleting tasks is the destiny of the target task, which are the suspended state and the non-existent state respectively. Due to space limitations they are not presented here, but can be found from the supplementary material.

5.7 Summary

This chapter firstly introduced VCC and the overview of our verification with VCC. The rest of the chapter showed our verification for task related functions in the FreeRTOS implementation. The work described in this chapter shows the possibility of verifying executable code with an abstract specification and code verifier.

In the next chapter, the FreeRTOS model described in Chap. 4 will be extended to suit a multi-core platform.

Chapter 6

EXTENSION FOR MULTI-CORE

This chapter extends the model described in Chap. 4 to a multi-core platform. It also follows the structure adapted in Chap. 4, which develops the Task, Queue, Time and Mutex models. Due to time limitations, we were only able to extend the FreeRTOS specification to a multi-core platform; the proof of the consistency of the model could not be completed and is part of our future work. In addition, it was not possible to show all details of the model in the thesis. Full details can be found from the supplementary material.

6.1 Overview

To migrate FreeRTOS to a multi-core platform, it is important to find a new scheduling policy for tasks. Similar to the model for FreeRTOS, this is an abstract model. Therefore, one option is to leave the scheduling algorithm as nondeterministic. It could then be refined as required at the refinement stage. However, because the system described is an extension of FreeRTOS, we decided to extend the scheduling algorithm for FreeRTOS to a multi-core algorithm. This provides benefits for the development in later stages. For instance, if we apply the promotion technique to develop the new specification of the multi-core model¹, the specification for FreeRTOS can be reused. As mentioned previously, in FreeRTOS, tasks are

¹We produced another version of the task model with promotion, which is shown in Sect. 8.2.

scheduled based on their priority. On a multi-core platform, tasks execute in parallel. Therefore, rather than time-slicing a single processor, a simple scheduling policy is used to extend the notion of a highest-priority ready task to a set of highest-priority tasks sufficient for the number of cores available. This is called *global scheduling*. However, it is known to be less efficient than *partitioned scheduling*, where each task is bound to a specified core and scheduling occurs within the core [73]. The main reason for this is the cost of migrating tasks from one core to another. We adopted a priority-based partitioned scheduling policy.

As this was the first attempt, we aimed to keep everything as simple as possible. Therefore, we avoided specifying details about the architecture of the multi-core platform, memory and interrupt management. Following the structure of the FreeRTOS model, we divided our specification for the multi-core platform into four major parts: *Task*, *Queue*, *Time* and *Mutex*. The well-definedness of these models was checked by Z/Eves. Furthermore, they can also be animated by ProZ.

6.2 Task Model

This is the core part of the specification. It includes operations for task management, such as creating tasks and deleting tasks, *etc.* Similar to Sect. 4.2.1, we also need to define some basic statements that are used in the specification.

6.2.1 Basic Statements

As well as the given sets defined in Sect. 4.2.1, we introduce another given set, *CORE*, to represent the cores.

$$[CONTEXT, TASK, CORE]$$

The constant, *bare_context* represents the initial state of the processors, similar to the FreeRTOS model. The multi-core operating system may have multiple *idle* tasks, one for each core; therefore, the constant *idle* need to be updated to *idles*, a finite set of tasks. The last constant *cores* represents all the cores available for the system. As each core can have only one *idle* task, the number for cores and *idle* tasks should be the same.

$$\text{bare_context} : \text{CONTEXT}$$
$$\text{idles} : \mathbb{F} \text{ TASK}$$
$$\text{cores} : \mathbb{F} \text{ CORE}$$
$$\# \text{cores} = \# \text{idles}$$
$$\text{cores} \neq \emptyset$$

As in FreeRTOS, we defined five states for tasks in our system. Their state transitions should also obey *transition* defined in Sect. 4.2.1.

Similar to the FreeRTOS model, we also have four sub-state schemas to describe all essential properties of tasks in the system: *TaskData*, *StateData*, *ContextData* and *PrioData*. *TaskData* includes the most basic properties of the system. It first introduces a finite set *tasks* to represent all the tasks in the system. As a multi-core system, there is more than one running task in the system: each core must always have a running task and the running task is uniquely run on one core. Therefore, the definition of the running task is updated to an injective total function, *running_tasks*, which shows the relation between cores and their running tasks. Finally, a partial surjective function, *executable*, is given to record the relation between each task and its core. As we adopted a partitioned scheduling policy in our system, this information is essential for scheduling. To guarantee the properties mentioned above, four constraints are specified.

$$\text{TaskData}$$
$$\text{tasks} : \mathbb{F} \text{ TASK}$$
$$\text{running_tasks} : \text{cores} \rightarrow \text{TASK}$$
$$\text{executable} : \text{TASK} \rightarrow \text{cores}$$
$$\text{ran running_tasks} \subseteq \text{tasks}$$
$$\text{idles} \subseteq \text{tasks}$$
$$\text{dom executable} = \text{tasks}$$
$$\forall t : \text{ran running_tasks} \bullet \text{running_tasks}^{-1}(t) = \text{executable}(t)$$

The definition for *StateData* and *PrioData* are similar to the FreeRTOS model. The only difference is that rather than checking only one idle task, we need to ensure that for all idle tasks, their states are ready or running and their priorities are 0. Originally, there was only one core in the system, which meant that only one physical context needed to be recorded. Here, in the new *ContextData*, a function

from *cores* to *CONTEXT* is used to record the physical context for each core. With these schema definitions, the schema for the base system, *Task*, can be defined as follows:

$ \begin{array}{l} \textit{Task} \\ \textit{TaskData} \\ \textit{StateData} \\ \textit{ContextData} \\ \textit{PrioData} \\ \\ \textit{tasks} = \textit{TASK} \setminus (\textit{state} \sim (\{ \textit{nonexistent} \} \cup)) \\ \textit{state} \sim (\{ \textit{running} \} \cup) = \text{ran } \textit{running_tasks} \\ \forall \textit{pt} : \textit{state} \sim (\{ \textit{ready} \} \cup); r : \text{ran } \textit{running_tasks} \\ \quad \textit{executable}(\textit{pt}) = \textit{executable}(r) \bullet \textit{priority}(r) \geq \textit{priority}(\textit{pt}) \end{array} $

It appends three extra constraints to the system. Comparing it to the *Task* schema in the FreeRTOS model, we can see that the meaning of these constraints are the same. However, as the number of running tasks changes from one to many, the expression has to be updated accordingly.

It is essential to provide the schema to initialise *Task*. When the system starts, (a) There are only *idle* tasks in the system; (b) The state for all tasks is *nonexistent*, except *idle* tasks, which are the *running* tasks for their cores; (c) The physical and logical context is empty at the moment, i.e., *bare_context*; (d) The priority is 0 for all tasks. Because each core only has an *idle* task in the initial state, the initial value for *running_tasks* and *executable* should be pairs of *idles* and *cores*. However, although it is important to know there is an *idle* task from *idles* executing on each core from *cores*, we do not care which *idle* task maps to which core. Therefore, we leave the initial value for *running_tasks* and *executable* as nondeterministic.

Similar to the previous model, the Δ schema for *Task* has been overridden to insist that the state transfer for each task should follow the rule defined in *transition* as well.

6.2.2 Additional Schemas

First, the schema *Reschedule* is modified to satisfy priority-based partitioned scheduling for the multi-core environment. Specifically, all the tasks are bound to a core, which is available in the system; they will be scheduled later within the core, based on their priority. When there is a task whose priority is higher than the running task of its core, it will be scheduled as the new running task.

<p><i>Reschedule</i></p> <p>$\Delta Task$</p> <p>$target? : TASK$</p> <p>$tasks? : \mathbb{F} TASK$</p> <p>$executable? : TASK \rightsquigarrow cores$</p> <p>$st? : STATE$</p> <p>$pri? : TASK \rightarrow \mathbb{N}$</p> <hr/> <p>$target? \in tasks?$</p> <p>$dom\ executable? = tasks?$</p> <p>$tasks' = tasks?$</p> <p>$running_tasks' = running_tasks$</p> <p>$\oplus \{(executable?(target?) \mapsto target?)\}$</p> <p>$executable' = executable?$</p> <p>$state' = state \oplus \{(target? \mapsto running),$</p> <p>$(running_tasks(executable?(target?)) \mapsto st?)\}$</p> <p>$phys_context' = phys_context$</p> <p>$\oplus \{(executable?(target?) \mapsto log_context(target?))\}$</p> <p>$log_context' = log_context$</p> <p>$\oplus \{(running_tasks(executable?(target?))$</p> <p>$\mapsto phys_context(executable?(target?))\}$</p> <p>$priority' = pri?$</p>

As well as the interface variables introduced in the FreeRTOS model, *target?*, *tasks?*, *st?*, *pri?*, another input variable is provided, *executable?*, which updates the executable information for tasks.

Second, there is also a frequently reused operation to search for the ready task that is bound to the same core as the given task and holds the highest priority.

<p><i>findTopReady</i></p> <p><i>Task</i></p> <p>$target? : TASK$</p> <p>$topReady! : TASK$</p>

$$\begin{array}{l}
target? \in tasks \\
state(topReady!) = ready \\
executable(topReady!) = executable(target?) \\
\forall rt : state \sim (\{ready\}) \mid executable(rt) = executable(topReady!) \\
\bullet priority(topReady!) \geq priority(rt)
\end{array}$$

This schema is reasonably simple compared to the previous one. It takes the given task in the input $target?$ and returns the highest-priority ready task with $topReady!$, which is nondeterministically selected because we are not concerned about the details of any particular scheduling algorithm, only that this task holds the highest priority among the other ready tasks within the same core.

Third, due to partitioned scheduling, when creating a task it is necessary to allocate it to one specific core, which may have access to a particular resource or have the shortest distance property. If, however, the location of the task is not of interest to the user, then the system will allocate a suitable core. The algorithm used here is inspired by Best-Fit Algorithm [74] for memory allocation. The schema $findACore_T$ is designed to be used by the task creation schemas to find a proper core for new task.

$$\begin{array}{l}
\textit{findACore_T} \\
\textit{Task} \\
newpri? : \mathbb{N} \\
executeCore? : CORE \\
executeCore : CORE \\
\\
executeCore? \notin cores \\
executeCore \in cores \\
\exists tcs, cs : \mathbb{F} cores \mid \\
tcs = \{ pc : cores \mid newpri? > priority(running_tasks(pc)) \} \\
\bullet (tcs = \emptyset \Rightarrow cs = cores) \wedge (tcs \neq \emptyset \Rightarrow cs = tcs) \\
\wedge (\forall oc : cs \bullet executeCore \in cs \\
\wedge \#(executable \sim (\{executeCore\})) \leq \\
\#(executable \sim (\{oc\})))
\end{array}$$

The key aim of this algorithm is to find the best core for the task, so that it can be executed as soon as possible. It also attempts to minimise the maximum loads of all the cores, thus helping other tasks to meet their deadlines. Specifically, this

schema takes two input variables from the task creation schemas, *newpri?* and *executeCore?*, which represent the new priority of the target task and the target core specified by the user. When *executeCore?* is specified by the developer, this schema will have no effect. Otherwise, it compares the priority of the new task with the priority of all running tasks to find out whether it is possible to schedule the new task immediately. If it is possible, then the set *tcs* includes all possible cores. Otherwise, it is set to *cores*, which indicates all the cores available in the system. Subsequently, it examines which core has the minimum load and sets it to *executeCore*, which can be used later by related schemas.

6.2.3 Creating and Deleting Tasks

Similar to creating tasks for FreeRTOS, the relation between the priority of the new task and the priority of the running task, both of which are executed in the same core, splits task creation into two separate cases, normal and rescheduling. Depending on whether the associated core for the new task is specified, each of these cases can be further divided into two sub-cases. Therefore, the definition for task creation can be divided into four cases. When the user specifies the executing core, the new task *target?* is added to the system with the new priority *newpri?* and the bound core information recorded in related functions. If the priority of the new task is not higher than the priority of the running task, the system does not need rescheduling; the state of the new task is set to *ready*. Otherwise, *Reschedule* is used to reschedule the system. When the executing core is not provided by the developer, the schema *findACore_T* is used to determine the core to which the new task will be bound. Therefore these cases, where the user does not provide the executing core, are covered by the cases where the executing core is specified.

The delete operation is simpler than the create operation, because it considers only two cases, rescheduling or not. These cases are similar to the FreeRTOS model. If the deleted task is not the running task, it removes the target task from the system and updates related functions; otherwise, it leads to rescheduling. As tasks are bound to different cores, we cannot simply get the task with the highest priority from the *ready* state, like the rescheduling case for task deletion in the FreeRTOS model.

The schema *findTopReady* is used to find which task is going to be rescheduled as the new running task for that core after the operation. Otherwise, it does a similar job to the FreeRTOS model: removes *target?*, which is a running task, from *tasks*, and *executable*; updates *topReady!*, which is defined by *findTopReady*, to be the running task in its core; updates the physical context of the executing core of *topReady!* to its logical context; and finally, sets the logical context of the previous running task to *bare_context*.

6.2.4 Migrating Task

As a multi-core system, it is also necessary to provide facilities for the user to move a task from one core to another. The schema, *Migration_T*, is introduced for this purpose. There are four cases for migrating tasks:

1. *MigrationN_T* The task that is going to be migrated is a non-running task with a priority that is not greater than the priority of the running task of the target core, and therefore does not cause rescheduling, either in the original core or in the target core. All that needs to be done is to update the information about the target task in the new core for *executable*.
2. *MigrationS_T* The migrating task is a non-running task with priority higher than the running task of the new core. The original core does not perform rescheduling, but the migrating task causes a reschedule in the target core. The schema *Reschedule* is used to do this.
3. *MigrationRuN_T* The migrating task is running in its core, but has a priority lower than or a priority equal to the running task of the target core. As the migrating task is moved out, rescheduling is requested to find a suitable ready task to fill the core. In the target core, however, the migrant is simply added to the ready list.
4. *MigrationRuS_T* The migrant task is the running task in its core and also has a higher priority than the running task in the target core. Both source and target cores need rescheduling (and *Reschedule* cannot do this).

6.2.5 Other Operations

In addition to the operations described above, other operations are also for task management, such as suspending tasks, resuming tasks and changing the priority of tasks. However, their effects are isolated to a single core, the one in which the target task is executing. Their definitions are almost the same as in the FreeRTOS model. Therefore, they are not described in detail here. Their specification can be found in the supplementary material. It is worth noticing that as the definition of the running tasks is different, the related expressions are different and we need to keep the function *executable* unchanged for these operations.

6.3 Queue Model

Queue is also defined as a communication facility for tasks in our system. As tasks which need to communicate with each other may be resident in the different cores, queues should be accessible for different cores. At the same time, the base statement definition for the queue model in the FreeRTOS model includes all the essential information about a queue. However, each queue is limited to a single core only. Therefore, the function *q_ava* is appended to the schema *QueueData* to enable multiple accessibility for the new model. This function records a set of cores for each queue, which can access the queue. Similar to the functions, *q_size* and *q_max*, the domain of this function should also be the queues known by the system, i.e., *queue*. The range also has to be a set of cores known by the system, i.e., *cores*.

QueueData

$$\begin{aligned} &queue : \mathbb{P} \textit{QUEUE} \\ &q_max : \textit{QUEUE} \rightarrow \mathbb{N}_1 \\ &q_size : \textit{QUEUE} \rightarrow \mathbb{N} \\ &q_ava : \textit{QUEUE} \rightarrow \mathbb{F} \textit{CORE} \end{aligned}$$
$$\begin{aligned} &\text{dom } q_max = \text{dom } q_size \\ &\text{dom } q_size = \text{dom } q_ava \\ &\text{dom } q_ava = \textit{queue} \\ &\text{ran } q_ava \subseteq \mathbb{F} \textit{cores} \\ &\forall q : \textit{QUEUE} \mid q \in \textit{queue} \bullet q_size(q) \leq q_max(q) \end{aligned}$$

To extend operations from the task model to the queue level, the same strategy in Sect. 4.3.2 can be used. In addition, when deleting or suspending a task, it is necessary to remove related data from functions in *WaitingData* and *QReleasingData*. It is worth noticing that although it is easy to use the variable *running_task* to refer to the task executing the operation in the FreeRTOS model, we have multiple running tasks in the new model. A new input variable, *self?*, has to be introduced to each schema which needs to refer to the task executing the operation. For instance, to distinguish whether a task has just been released from the waiting event, we check whether the running task belongs to the domain of releasing functions in the FreeRTOS model. Here, we need to verify whether *self?* belongs to the releasing functions.

To create a queue, as well as the constraints and the behaviours described in Sect. 4.3.3, we also need to know the set of cores, *cset?*, which can access the queue. This information is recorded by appending the ordered pair, (*que?*, *cset?*), to the function *q_ava*.

When deleting a queue, we first need to check whether the calling task belongs to a core which can access the queue. Because there is more than one running task in the system, it is impossible to refer to the calling task by using the running task. The input variable *self?* is introduced to indicate the calling task. Then, with the function *executable*, its bound core can be identified. If this core is one of the cores which can access the queue, deleting can be performed. The information related to the queue should be removed from *queue*, *q_max*, *q_size* and *q_ava*.

For sending and receiving items to and from a queue, we also use the variable *self?* to identify the calling task, and the schema *findTopReady* is used to find the correct highest-priority ready task when rescheduling is required. The behaviours of these operations are the same as the FreeRTOS model (see Sect. 4.3.4).

Because a queue can only be accessed by a particular set of cores in the system, if a task which is using the queue migrates to another core which is not authorised to access the queue, the set of available cores for the queue has to be

updated to include the core, or the task loses its ability to access the queue. The last operation for the queue model is introduced to handle this, which is called *ChangeQueueLevel_TQ*.

<p><i>ChangeQueueLevel_TQ</i></p> <p>$\Delta TaskQueue$</p> <p><i>que?</i> : <i>QUEUE</i></p> <p><i>self?</i> : <i>TASK</i></p> <p><i>cset?</i> : $\mathbb{F} cores$</p> <hr/> <p><i>self?</i> $\notin \text{dom } release_snd \cup \text{dom } release_rcv$</p> <p><i>que?</i> $\in queue$</p> <p>$state(self?) = running$</p> <p>$executable(self?) \in q_ava(que?)$</p> <p>$cset? \neq q_ava(que?)$</p> <p>$cset? \neq \emptyset$</p> <p>$\forall t : wait_rcv \sim (\{ que? \}) \cup wait_snd \sim (\{ que? \})$ $\cup release_rcv \sim (\{ que? \}) \cup release_snd \sim (\{ que? \})$</p> <ul style="list-style-type: none"> • $executable(t) \in cset?$ <p>$\exists Task$</p> <p>$queue' = queue$</p> <p>$q_max' = q_max$</p> <p>$q_size' = q_size$</p> <p>$q_ava' = q_ava \oplus \{(que? \mapsto cset?)\}$</p> <p>$\exists WaitingData$</p> <p>$\exists QReleasingData$</p>
--

As we can see from the specification, the behaviour of the operation is really simple. It updates the value of $q_ava(que?)$ and keeps everything else unchanged. However, it can be performed only when the new set of cores is not empty and not equal to the original set. In addition, it also has to include all the cores in which there are some tasks using the queue. Otherwise some of these tasks may lose access to the queue.

6.4 Time and Mutex Model

The semaphores and mutexes in this multi-core model are also defined as special cases of queues. The properties of time facilities and semaphores and mutexes are the same as the FreeRTOS model. The operation schema of the lower level model can be extended to the higher level by the same strategy. Moreover, for the

operation schemas for the time and the mutex model except *TimeSlicing_TQT*, schemas for mutex tasking and mutex given, they focus on the behaviour of one task in one core, so their definitions are close to the FreeRTOS model. Due to length limitations, they are not repeated here (see Sect. 4.4 & 4.5 and Page. 205 and 225 for the details).

6.4.1 Time Slicing

The most interesting schema for this model is *TimeSlicing_TQT*. When time-slicing happens, the running tasks need to be replaced in all the cores in which there are some ready tasks that have the same priority as their running task. To achieve this, we first define a set of ready tasks with the same priority as their running task, called *topReadys!*. With this set and the inverse function of *executable*, we can find all the cores which need to be rescheduled. Therefore, the following schema is given.

$$\begin{array}{l}
 \text{--- } \underline{\text{TimeSlicing_TQT}} \text{ ---} \\
 \Delta \text{TaskQueueTime} \\
 \text{topReadys!} : \mathbb{F} \text{ TASK} \\
 \\
 \# \text{topReadys!} \leq \# \text{cores} \\
 \forall t : \text{topReadys!} \bullet \text{state}(t) = \text{ready} \\
 \quad \wedge \text{priority}(t) = \text{priority}(\text{running_tasks}(\text{executable}(t))) \\
 \forall t1, t2 : \text{topReadys!} \mid \text{executable}(t1) = \text{executable}(t2) \bullet t1 = t2 \\
 \forall c : \text{cores} \mid (\forall t : \text{topReadys!} \bullet \text{executable}(t) \neq c) \\
 \quad \bullet (\forall t : \text{executable}^{-1}(\{c\}) \mid \text{state}(t) = \text{ready} \\
 \quad \quad \bullet \text{priority}(t) < \text{priority}(\text{running_tasks}(c))) \\
 \text{topReadys!} \neq \emptyset \\
 \forall t : \text{dom } \text{time} \bullet \text{time_slice} \leq \text{time}(t) \\
 \text{tasks}' = \text{tasks} \\
 \text{executable}(\text{topReadys!}) \triangleleft \text{running_tasks}' \\
 \quad = \text{executable}(\text{topReadys!}) \triangleleft \text{running_tasks} \\
 \text{executable}' = \text{executable} \\
 (\text{running_tasks}(\text{executable}(\text{topReadys!})) \cup \text{topReadys!}) \triangleleft \text{state}' \\
 \quad = (\text{running_tasks}(\text{executable}(\text{topReadys!})) \cup \text{topReadys!}) \triangleleft \text{state} \\
 \text{executable}(\text{topReadys!}) \triangleleft \text{phys_context}' \\
 \quad = \text{executable}(\text{topReadys!}) \triangleleft \text{phys_context} \\
 \text{running_tasks}(\text{executable}(\text{topReadys!})) \triangleleft \text{log_context}' \\
 \quad = \text{running_tasks}(\text{executable}(\text{topReadys!})) \triangleleft \text{log_context} \\
 \text{priority}' = \text{priority} \\
 \forall trt : \text{topReadys!}
 \end{array}$$

- $running_tasks'(executable(trt)) = trt$
 $\wedge state'(trt) = running$
 $\wedge state'(running_tasks(executable(trt))) = ready$
 $\wedge phys_context'(executable(trt)) = log_context(trt)$
 $\wedge log_context'(running_tasks(executable(trt)))$
 $= phys_context(executable(trt))$

$\exists Queue$

$clock' = clock$

$delayed_task' = delayed_task$

$time' = time$

$time_slice' = time_slice + slice_delay$

For this operation, the post condition for *running_tasks* has to be discussed in two parts: (a) for the cores which do not need to be rescheduled, the value has to be equal to its original value; (b) on the other hand, the value has to be updated with tasks in *topReadys!* according to its executable core. Similarly, the post conditions for variables *state*, *phys_context* and *log_context* also need to be considered in two parts. Moreover, the value of *time_slice* has to be increased and the rest of the variables should be the same as before.

6.4.2 Taking Mutexes

As mentioned in Sect. 4.5.4, for a single core system, there are six cases for taking a mutex. When this operation migrates to a multi-core platform, there are two extra cases. First, if the task which is executing is taking the mutex and holds the mutex already, the specification simply increases the value of *mutex_recursive* for the mutex. Then, if the mutex is available, there are two cases for taking a mutex depending on whether the calling task has held a mutex already. Similar to the FreeRTOS model, the schema *QueueReceiveN_TQT* is used to simplify the specification. When the mutex is not available and the priority of the calling task is not greater than the priority of the mutex holder, the calling task will be blocked by the operation and no priority inheritance requested, *QueueReceiveE_TQT* is used for this case. In addition, once the priority of the calling task is higher than the mutex holder, the mutex holder needs to inherit the priority of the calling task. As a multi-core system, the mutex holder and the calling task may be executing in different cores. If they are executing in the same core, the operations are close

to the FreeRTOS model. The mutex holder or the top priority ready task replaces the calling task, which is blocked by the waiting event. However, when the mutex holder belongs to a different core of the calling task, there are two additional cases. When the inherited priority of the mutex holder is lower than or equal to its running task or the mutex holder is not in the ready state, for that core, no rescheduling is requested. We just need to reschedule the top priority ready task of the calling task as the new running task. On the other hand, once the mutex holder is in the ready state and the inherited priority is higher than its running task, rescheduling needs to be performed in both cores. For these two schemas, *Reschedule* cannot be used for rescheduling. This is because of the need to update the state of three and four tasks respectively. *Reschedule* cannot handle this.

6.4.3 Giving Mutexes

Like the FreeRTOS model, the mutex giving operation is the most complex of the whole model. The new version of *basePriorityMan* is also introduced to help manage the base priority of the calling task. The first case for this operation is also for a recursively returning mutex. It decreases the *mutex_recursive* of the mutex and keeps everything else unchanged if the mutex holder takes the mutex several times. When there is no task waiting to take the mutex, we have three cases: (a) The mutex holder, i.e., the calling task, did not inherit priority from another task. It just uses the schema *QueueSendN_TQT* to return the mutex, removes the mutex holder from the mutex, sets the *mutex_recursive* of the mutex to 0, uses the schema *basePriorityMan* to manage the base priority of the calling task and keeps the rest of the variables unchanged; (b) When the priority of the mutex holder is inherited from another task, it needs to reset its priority to the original. Further, if its original priority is still the highest of the ready tasks in the same core, it can keep executing. As the priority of the mutex holder has to be updated, the schema *QueueSendN_TQT* cannot be used. The behaviour has to be defined from scratch. It first updates the priority of the mutex holder, sets the size of the mutex to 1 make the mutex available for other tasks and also updates *mutex_holder*, *mutex_recursive*, etc. as in the previous case; (c) Finally, when the

original priority of the mutex holder is not the highest; rescheduling is requested and the schema *Reschedule* is used. Furthermore, if there are tasks waiting to take the mutex, when the mutex holder returns the mutex, the highest priority waiting task is woken up. Meanwhile, if the mutex holder did not inherit the priority, the priority of the mutex holder does not need to be updated after the operation. What we need to do is return the mutex, and unblock the highest priority waiting task. However, depending on the priority of the woken-up task, rescheduling needs to be considered. If its priority is not higher than the running task of its core, the woken task is placed in the ready state. Otherwise, the running task is replaced in its core. The behaviour of these two cases is the same as that of the queue sending operation except variables related to mutexes. Therefore, the specification for these two schemas can be defined with the schema of *QueueSendW_TQT* and *QueueSendWS_TQT* respectively, and an update of the functions *mutex_holder* and *mutex_recursive*.

Finally, different from last two cases, the mutex holder inherits priority from other tasks in the remaining cases. After the returning operation, the priority of the mutex holder has to be reset and the highest priority waiting task has to be woken up. In detail, the following conditions have to be considered to decide the behaviour of these cases:

1. Do the woken task and the calling task belong to the same core?
2. The relationship between the original priority of the mutex holder, the top priority ready task which executes in the same core as the mutex holder, the woken task and the running task in the same core as the woken task.

Table 6.1 illustrates the cases and their conditions in detail. In this table, *Running Task* refers to the running task which executes on the same core as the woken task, if the executing core for the woken up task is different from the executing core for the calling task; *Top Waiting* refers to the highest priority waiting task of the mutex; and *Top Ready* represents the highest priority ready task of the core of the calling task.

Table 6.1: Conditions for giving mutex cases (have waiting tasks and the mutex holder inherits the priority)

	Running Task	Top Waiting	Calling Task	Top Ready
1	c_2 P_H Running	c_1 - Ready	c_1 - Ready	c_1 P_H Running
2	c_2 P_H Running	c_1 P_H Running	c_1 - Ready	c_1 - Ready
3	c_2 P_H Running	c_1 - Ready	c_1 P_H Running	c_1 - Ready
4	c_2 - Ready	c_2 P_H Running	c_1 P_H Running	c_1 - Ready
5	c_2 - Ready	c_2 P_H Running	c_1 - Ready	c_1 P_H Running
6	c_2 P_H Running	c_2 - Ready	c_1 P_H Running	c_1 - Ready
7	c_2 P_H Running	c_2 - Ready	c_1 - Ready	c_1 P_H Running

From this table, firstly we can see that the effect of cases 3 & 6 are the same. They do not need to reschedule the tasks in any core. The schema *MutexGiveWinhN_TQTM* is defined for this case, which is the normal case for mutex returning. Secondly, cases 1 & 7 have the same effect as well. Although the relation between the executing cores for the top priority waiting task and the calling task are different in these two cases, they both request to use the top priority ready task replacing calling task as the new running task and do not reschedule for any other cores. When the woken task belongs to the core of the calling task, the priority of the top priority ready task has to be the highest compared with the original priority of the calling task and the priority of the woken task. When the operation resets the priority of the calling task, the system needs to be rescheduled. Similarly, cases 2 & 4 can also be recognised as the same case. Even through the post

states for these four tasks are different in these two cases, it can be considered as the top priority waiting task being rescheduled as the new running task in its core.

Specifically, if it belongs to the core of the calling task and its priority is the highest compared with the original priority of the calling task and the priority of the top priority ready task, it is scheduled as the new running task in its core (c_1). There is no rescheduling for other cores (c_2). On the other hand, when the woken task belongs to a different core from the calling task, and its priority is also higher than the priority of the running task in the core and the original priority of the calling task is no less than the priority of the top priority ready task in that core, the rescheduling also occurs in the core of the woken task (c_2) and no rescheduling is requested for other cores (c_1). Finally, case 5 is the last case for giving mutexes. It requests rescheduling for both cores of the calling task and the woken task. As shown in Table 6.1, the top waiting task and the calling task belong to different cores. The priority of the top waiting task and the top ready task are the highest in their cores. Rescheduling happens in both cores. As well as these differences for task related variables, the new values for the other variables are the same as the other cases of mutex giving. The specification for these schemas can be found from Page. 258.

6.5 Summary

This chapter presented the extended FreeRTOS model, which is used for a multi-core platform. The first section highlighted the differences between the original version and the multi-core version of the model. Following the structure used in Chap. 4, the model for the multi-core platform was shown, which focused on the differences between the original model and the new model. This piece of work shows that the model is easy to reuse and extend.

The next chapter will evaluate our project. Case studies will also be shown in the next chapter.

Chapter 7

EVALUATION AND CASE STUDIES

This chapter evaluates our research project. It compares the project with the objectives listed in Sect. 1.5, it also summarises the achievements of the project, followed by animations and proofs for the case studies discussed in Chap. 2. Finally, some issues about the implementation of FreeRTOS are discussed.

7.1 Project Summary

We will now show that, except objective d in Sect. 1.5 (which is included in Sect. 8.2 as a future work), our project meets the objectives (Sect. 1.5) and requirements (Sect. 3.3) proposed at the beginning of the project. During the project, we produced an abstract formal model for FreeRTOS (Obj. a). This includes 514 Z paragraphs, including basic definitions for the model, schemas for the operations, theorems for the consistency verifications and some assistant lemmas. From these paragraphs, 598 theorems were derived, all of which have been proved with Z/Eves (**Well-definedness** and **Feasibility**). In particular, around half of them (240) were proved automatically. The specified behaviour of the FreeRTOS API functions meets the requirements list in Sect. 3.3.1. The summary of the relationship between the API interfaces and the schemas in the specification is included in Appendix C. In addition, the model can also be animated by

ProZ [33] (**Animatable**). Based on our specification, the function contracts for task-related API functions were developed and the functions verified with VCC (Obj. b). However, due to time limitations, we only verified the task-related API functions and reduced the code complexity to focus on the functions modelled with the Z notation, we believe it is enough to demonstrate the possibility of directly verifying the implementation with the abstract formal specification. During the modelling and verification of FreeRTOS, we detected some issues with the FreeRTOS implementation, which are discussed in detail in Sect. 7.3 (Obj. c). Finally, we extended the model of FreeRTOS to a multi-core platform, as described in Chap. 6. The well-definedness of our multi-core model was also verified with Z/Eves, which includes the basic syntax and domain check for all the definitions. Further, the multi-core model can also be animated by ProZ (Obj. e). These extensions show that our specification for FreeRTOS is **reusable**. We also provide the source code and project file of our work in the supplementary material. This makes it easy to reproduce (**Reproducible**). Therefore, our model satisfies non-functional requirement stated in Sect. 3.3.2.

In summary, except objective d, this project meets the objectives and requirements, both functional and non-functional, proposed at the beginning of the project.

7.2 Case Studies

As discussed in Chap. 1, we have shown how our specification illustrates the execution of FreeRTOS code. Due to the different format requirements between ProZ, CZT and Z/Eves, we modified the source files to fit them. However, to animate our model with ProZ, some extra modifications were still needed, as follows:

- Remove the definition for the *findDelegate* function, which is defined to help in proof and has no effect for animation. Most importantly, ProZ struggles with it and runs out of time;
- Remove the disabled mark for *transition*, i.e., use `\begin{zed}` to replace `\begin[disabled]{zed}`. Otherwise, ProZ cannot use this definition, as `[disabled]` is not defined in ProZ;

- Use “Init” to replace the name of the top-level initialisation schema of the model. For instance, to animate the Task model, the name of the schema *Init_Task*, should be replaced. This is the entrance for ProZ to detect the model state and initialise the model;
- Remove the label, $\langle\langle \textit{disabled slice_delay_def} \rangle\rangle$, which is not defined in ProZ;
- For the mutex model, use “XXX?/que?” to replace all “que? := XXX?” in the schema reuse, as CZT cannot recognise the second format; in addition, ProZ cannot handle the first one.
- For the mutex model, we use the schema *basePriorityMan* (See Page. 258) to simplify the definitions for returning mutex, which is helpful for modelling. However, it dramatically increases the load for the animator. Therefore, it has to be removed and its contents have to be used to replace all calls to the schema. For example, the schema *MutexGiveNnonInh_TQTM* (See Page. 258) uses the schema *basePriorityMan* to manage the base priority of the running task, which should be replaced by:

$$\begin{aligned}
& \textit{running_task} \in \text{ran}(\{\textit{mut?}\} \triangleleft \textit{mutex_holder}) \\
& \Rightarrow \exists \textit{OriginalPriorityData} \\
& \textit{running_task} \notin \text{ran}(\{\textit{mut?}\} \triangleleft \textit{mutex_holder}) \\
& \Rightarrow \textit{base_priority}' = \{\textit{running_task}\} \triangleleft \textit{base_priority}
\end{aligned}$$

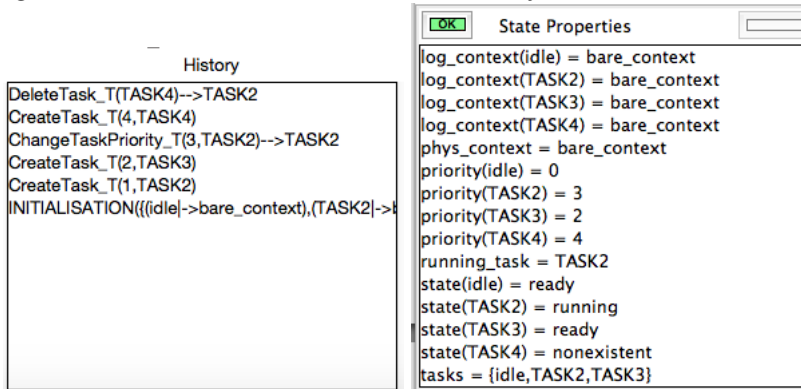
7.2.1 Case 1

The application for the first case is shown in Fig. 1.2 Page. 6. Firstly, ProZ is used to animate the model. Before animation, we set the size of our given sets to 4; the maximum integer needs to be set to 4 as well. This is because we have four tasks in the application, *idle*, *Task1*, *Task2* and *Task3*, and the maximum number used is the priority of *Task3*, namely 4. Afterwards, the *.tex* file is loaded in ProZ. Although we have three tasks in the application, these tasks execute in a single-core processor. Therefore, it is possible to predict the sequence of API function calling, which is¹:

xTaskCreate Create *Task2* with priority of 1;

¹As idle task is created when the animator is initialising, it is actually *Task1* for the animator. Therefore, the first task we created in the application (i.e. *Task1*) is *Task2* in ProZ.

Figure 7.1: API function execution history and result for Case 1



- xTaskCreate** Create Task3 with priority of 2;
- vTaskPrioritySet** Change the priority of Task2 to 3;
- xTaskCreate** Create Task4 with priority of 4;
- vTaskDelete** Delete Task4;

After initialising the machine in ProZ, we call the API function in this order (see Fig. 7.1). As we analysed in Sect. 1.2.1, the expected final state of execution should be:

1. There are three tasks left in the system, `idle`, `Task2`, and `Task3`;
2. `Task2` is the running task, as it has priority of 3.
3. `Task3` is in ready state with priority 2.
4. `Task4` is unknown to the system, therefore its state is nonexistent.

As we can see from the screen shot of state properties, the result generated from our model matches our expectations.

In addition to this, we also let Z/Eves verify our result. We use a theorem, similar to Theorem: 21 (see Page. 90), to show that the behaviour of the API function matches our expectation. The API function call `xTaskCreate` is repeated three times in the application, but we only show the theorem for one of these calls. Therefore, we have the following theorems to show our model works for the application.

1. Create `Task1`. When we execute this, there is only the `idle` task in the system, which is the running task; therefore, we have to indicate this situation

to the prover. The input variables need to be introduced and we need to specify the value of $newpri?$ as 1. After this operation, we expect that the new task is created, which means it is in the set $tasks'$. It should be the running task with priority 1, which is higher than the priority of the `idle` task.

Theorem 24 (CaseStudyStep1)

$$\begin{aligned} & \forall Task; target? : TASK; newpri? : \mathbb{N} \\ & | tasks = \{idle\} \wedge running_task = idle \\ & \quad \wedge newpri? = 1 \wedge CreateTask_T \\ & \bullet target? \in tasks' \wedge state'(target?) = running \\ & \quad \wedge priority'(target?) = 1 \end{aligned}$$

To prove this, we know that the system needs to be scheduled. Therefore, we try to eliminate the non-schedule part of the specification of $CreateTask_T$. The key condition to distinguish these two cases is whether the priority of the new task is greater than the running task. Thus, we expand the necessary schemas of the proof goal and then let the prover discharge the proof goal automatically by the `prove_by_reduce;` command.

proof [*CaseStudyStep1*]
with disabled (CreateTaskS_T, StateData,
TaskData, ContextData) reduce;
prove by reduce;

■

2. Change the priority of Task2 to 3. Similarly, we need to inform the prover about the pre-state of the system. The key element of the expected result of this API function call is that Task2 is scheduled as running task with priority of 3. By eliminating the unrelated case of $ChangeTaskPriority_T$, like the previous case for $CreateTask_T$, it is easy to prove this theorem.

Theorem 25 (CaseStudyStep3)

$$\begin{aligned} & \forall Task; Task2, Task3, target? : TASK; newpri? : \mathbb{N} \\ & | tasks = \{idle, Task2, Task3\} \\ & \quad \wedge priority(Task2) = 1 \wedge priority(Task3) = 2 \\ & \quad \wedge state(Task2) = ready \wedge running_task = Task3 \end{aligned}$$

$$\begin{aligned}
& \wedge \text{target?} = \text{Task2} \wedge \text{newpri?} = 3 \\
& \wedge \text{ChangeTaskPriority_T} \\
\bullet & \text{priority}'(\text{Task2}) = 3 \wedge \text{running_task}' = \text{Task2}
\end{aligned}$$

3. Finally, we verify the properties related to the last step of the API function call, delete Task4. Following the strategy introduced before, the theorem and proof can be obtained. The only difficulty in proving this theorem is in the nondeterministic definition for *topReady!*. In order to solve this, it is necessary to inform Z/Eves that (a) the possible value of *topReady!* is one of the elements of *tasks*; and (b) the priority of *topReady!* is the greatest amongst all *ready* tasks, i.e., the priority of *topReady!* has to be greater than or equal to the priority of Task2.

Theorem 26 (CaseStudyStep5)

$$\begin{aligned}
& \forall \text{Task}; \text{Task2}, \text{Task3}, \text{Task4}, \text{target?} : \text{TASK}; \text{newpri?} : \mathbb{N} \\
& | \text{tasks} = \{\text{idle}, \text{Task2}, \text{Task3}, \text{Task4}\} \\
& \quad \wedge \text{priority}(\text{Task2}) = 3 \wedge \text{priority}(\text{Task3}) = 2 \\
& \quad \wedge \text{priority}(\text{Task4}) = 4 \wedge \text{state}(\text{Task2}) = \text{ready} \\
& \quad \wedge \text{state}(\text{Task3}) = \text{ready} \wedge \text{state}(\text{Task4}) = \text{running} \\
& \quad \wedge \text{target?} = \text{Task4} \wedge \text{DeleteTask_T} \\
\bullet & \text{state}'(\text{Task4}) = \text{nonexistent} \wedge \text{running_task}' = \text{Task2}
\end{aligned}$$

7.2.2 Case 2

In Fig. 1.3, we provided example code demonstrating communication and synchronisation related API functions (see Page. 10). As described, it is also animated and verified with ProZ and Z/Eves, respectively. Similar to the previous case, we first animate it with the ProZ animator. In this case, we have two tasks in the system, together with the *idle* task. Therefore, it is necessary to set the size of the given set to 3 for this case. Similarly, 3 would be sufficient for the maximum integer, as the highest priority in the system is 3, the priority for Task3. After loading the model to the animator, the following sequence of API functions is called.

xSemaphoreCreateMutex Create a mutex *xMutex*;

xTaskCreate Create Task2 with priority of 2;

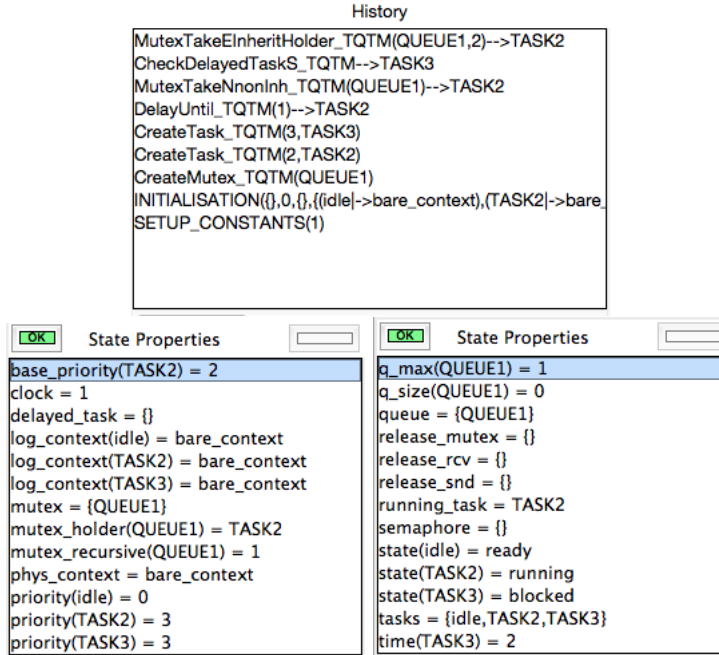
xTaskCreate Create Task3 with priority of 3;
vTaskDelay Delay Task3 for 1 time units;
xSemaphoreTask Task2 attempts to take the mutex xMutex;
xSemaphoreTask Task3 attempts to take the mutex xMutex.

After initialising the machine in ProZ, the sequence of API functions are called as shown in Fig. 7.2(Above). It should be noted that when we invoke the schema *DelayUntil_TQTM*, we selected 1 time unit instead of 10 which was originally defined on Page. 10, because the maximum number of the integer is 3. If we increase it to 10 the load of the animator would be dramatically increased. At the same time, the purpose of this function call is to block the high priority task. This replacement has no effect on the result. Moreover, the schema *CheckDelayedTaskS_TQTM* is called after Task2 has taken the mutex to release Task3 from the blocked state. In addition, we use the sub operation schemas for mutex take instead of *MutexTake_TQTM*. The key reason is that due to the complexity of *MutexTake_TQTM*, if we use it directly, it considerably increases the load of the animator. The animator will therefore take a long time to calculate and may fail to respond at all.

The main behaviour of this piece of code is that two tasks compete for a single mutex. The lower priority task holds the mutex and blocks the higher priority task. As the higher priority task is blocked by the mutex taking operation, priority inheritance happens. The lower priority task inherits the priority of the higher priority task. Therefore, the final priority of Task2 should be 3 and it should be in the running state, while Task3 is in the blocked state with its own priority. As shown on the bottom side of Fig. 7.2-Bottom, the actual animation result matches our expectation.

Because we have shown the verification of the theorem of xTaskCreate in the previous section, it is not repeated here. The following theorems show two cases of xSemaphoreTake. The first one illustrates the case when Task2 attempts to take the mutex. As the mutex is free initially, it successfully takes the mutex with nor-

Figure 7.2: API function execution history (above) and result (bottom) for Case 2



mal case in the specification. Meanwhile, the second represents the case when Task3 attempts to take `xMutex`, which matches the case when the mutex is unavailable, priority inheritance is requested and the mutex holder is scheduled as the new running task. Both of them can be easily proved by the Z/Eves command “*prove by reduce*”.

Theorem 27 (caseStudyTask1Take)

- $$\forall TaskQueueTimeMutex; mut? : QUEUE; topReady! : TASK$$
- $$| tasks = \{idle, Task2, Task3\} \wedge queue = \{QUEUE1\}$$
- $$\wedge priority(Task2) = 2 \wedge priority(Task3) = 3$$
- $$\wedge state(Task3) = blocked \wedge running_task = Task2$$
- $$\wedge mut? = QUEUE1 \wedge QUEUE1 \notin \text{dom } mutex_holder$$
- $$\wedge QUEUE1 \in mutex \wedge release_rcv = \emptyset$$
- $$\wedge release_snd = \emptyset \wedge base_priority = \emptyset$$
- $$\wedge MutexTakeNnonInh_TQTM$$
- $mutex_holder'(QUEUE1) = Task2$

$$\wedge priority'(Task2) = 2 \wedge priority'(Task3) = 3$$

$$\wedge running_task' = Task2$$

Theorem 28 (caseStudyTask2Take)

- $$\forall TaskQueueTimeMutex; mut? : QUEUE; topReady! : TASK$$
- $$| tasks = \{idle, Task2, Task3\} \wedge queue = \{QUEUE1\}$$

$$\begin{aligned}
& \wedge \text{priority}(\text{Task2}) = 2 \wedge \text{priority}(\text{Task3}) = 3 \\
& \wedge \text{state}(\text{Task2}) = \text{ready} \wedge \text{running_task} = \text{Task3} \\
& \wedge \text{mut?} = \text{QUEUE1} \wedge \text{QUEUE1} \in \text{mutex} \\
& \wedge \text{QUEUE1} \in \text{dom mutex_holder} \\
& \wedge \text{mutex_holder}(\text{QUEUE1}) = \text{Task2} \\
& \wedge \text{release_rcv} = \emptyset \wedge \text{release_snd} = \emptyset \\
& \wedge \text{base_priority} = \emptyset \wedge \text{clock} = 1 \wedge n? = 2 \\
& \wedge \text{MutexTakeEInheritHolder_TQTM} \\
\bullet & \text{mutex_holder}'(\text{QUEUE1}) = \text{Task2} \wedge \text{state}'(\text{Task3}) = \text{blocked} \\
& \wedge \text{priority}'(\text{Task2}) = 3 \wedge \text{priority}'(\text{Task3}) = 3 \\
& \wedge \text{running_task}' = \text{Task2}
\end{aligned}$$

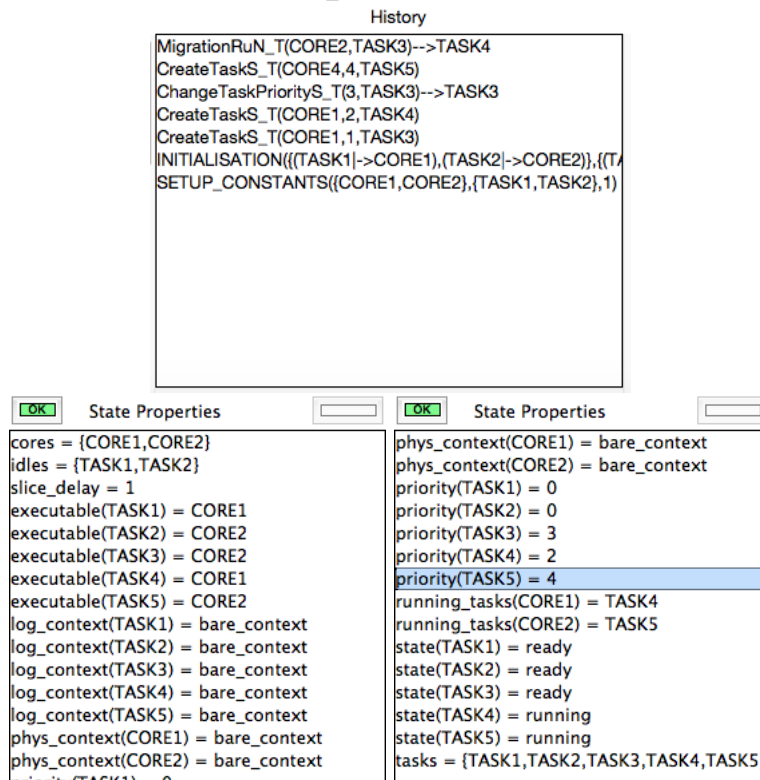
7.2.3 Case 3

Our extension model for the multi-core platform can also be animated with ProZ. Similarly, we formulate some sequences of operation calls that illustrate the behaviour of the model. Consider the following sequence, assuming there are two cores available to the system.

1. Initially, there are two tasks, Task3 and Task4, created on Core1 with priority of 1 and 2, respectively.
2. Change the priority of Task3 to 3.
3. Create Task5 with priority of 4.
4. Move Task3 to Core2.

When the system is initialised, Task4 occupies Core1, as it has the highest priority within Core1. Meanwhile, on Core2, the idle task is executing. When the priority of Task3 is changed to 3, it preempts Task4 and can execute. Thereupon, Task5 is created with a priority of 4. As it is not specified on which core the new task is created, the system selects one to accommodate the new task, based on the algorithm described in Sect. 6.2.3. In this case, Task5 will be created on Core2. Finally, Task3 is moved to Core2. Because Task3 is the current running task on Core1, moving it to Core2 causes Task4 to be scheduled as the next running task. Therefore, the final state of the system should be as follows: (a) there are five tasks in the system; (b) Task4 and Task5 are executing on Core1 and Core2, respectively; (c) Task3 is ready in Core2. To animate this

Figure 7.3: API function execution history (above) and result (bottom) for Case 3



process, we need at least five individual tasks and two cores available from the animator with a maximum natural number of 4. We initialise the system with the `setup_constants` and `initialisation` commands. For the `setup_constants` command, we need sets $\{CORE1, CORE2\}$ & $\{TASK1, TASK2\}$ as the parameters. Next, the `initialisation` command is used to initialise the system with maplets $TASK1 \mapsto CORE1$, $TASK2 \mapsto CORE2$ and so on, as parameters. This initialises the system with some definitions, such as, $TASK1$ and $TASK2$ as the `idle` task for $CORE1$ and $CORE2$, respectively. At this moment, we can animate the system following the sequence described above (see Fig. 7.3-Above). After these operations, we find the system status has been changed to Fig. 7.3(Bottom), which matches our expectation.

Again, we can verify the API functions with the Z/Eves theorem prover. For instance, consider the scenario where we initialise the state as shown above and then call `CreateTask_T` to create `Task3` on `Core1` with priority 1.

Theorem 29 (createTaskOverTwoCores)

$$\begin{aligned} & \forall Task; target? : TASK; newpri? : \mathbb{N}; c? : CORE \\ & | cores = \{c1, c2\} \wedge running_tasks = \{(c1 \mapsto i1), (c2 \mapsto i2)\} \\ & \quad \wedge newpri? = 1 \wedge tasks = \{i1, i2\} \wedge c? = c1 \\ & \quad \wedge executable = \{(i1 \mapsto c1), (i2 \mapsto c2)\} \wedge CreateTask_T \\ & \bullet target? \in tasks' \wedge state'(target?) = running \\ & \quad \wedge priority'(target?) = 1 \wedge executable'(target?) = c1 \end{aligned}$$

, where $i1, i2$ represent idle tasks for two cores (i.e. $c1$ and $c2$), respectively.

7.3 Issues of FreeRTOS

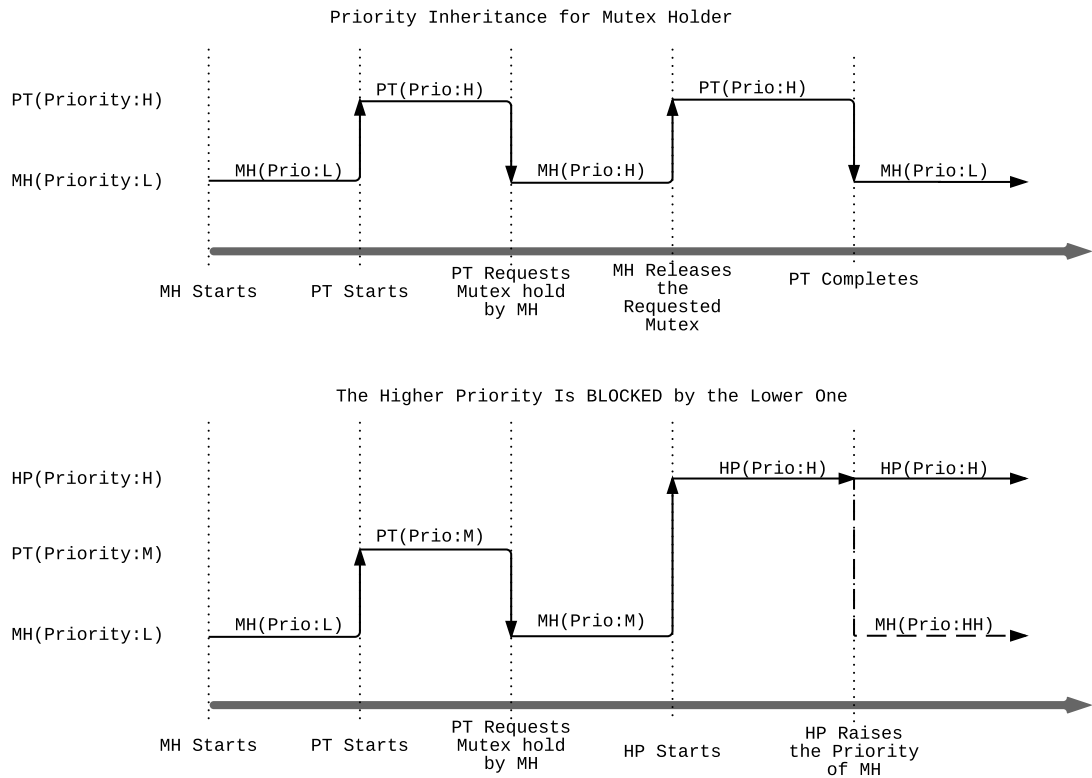
During the modelling and verifying process of our project, we detected some issues with FreeRTOS. These issues were revealed at different stages of the project. For instance, when we analysed the implementation to complete the requirements for modelling, we found an issue with changing the priority of a mutex holder (see, item 5). When we proved the precondition theorems, we found that deleting the `idle` task violated the system constraints (see, item 1) and so on. Due to the limitation of our VCC model, only item 1 can be detected during VCC verification. But it shows the possibility of verifying the implementation with the VCC and the abstract specification.

1. As shown in the precondition for `vTaskDelete` (Sect. 5.4), there is one predicate, which states the target task does not equal the `idle` task. If this predicate is removed, the verification would fail. The reason is that the operation may delete the `idle` task, which breaks the constraints of the system. Generally, the handler of the `idle` task is hidden from the user. However, it is possible to obtain it from `vApplicationIdleHook`, which is provided by FreeRTOS and used to define extra behaviours of the `idle` task. Therefore, we need an extra precondition to ensure that the `idle` task cannot be deleted by the operation.
2. Similar to the previous issue, the delete API function does not check if the target is the holder of a mutex. Because our VCC model does not include the

context related to mutex, this issue has no effect on our verification. However, if a mutex holder is deleted without returning its mutex, then the resource is locked permanently.

3. According to [19], if `vTaskSuspendAll` is called, `xTaskResumeAll` has to be called as many times as the suspend operation to resume the scheduler. FreeRTOS uses an unsigned counter, `uxSchedulerSuspended`, to record this. However, the API function `vTaskSuspendAll` does not check the overflow of the counter. In other words, when it overflows, only one call to `xTaskResumeAll` can resume the scheduler from millions of calls of `vTaskSuspendAll`. Although calling `vTaskSuspendAll` millions of times might not happen in real life, this could still be a weak point in the system.
4. Similar to the issues related to `vTaskDelete`, the API function `vQueueDelete` also does not verify that the target queue is not in use. If a task is blocked by waiting to send/receive an event to/from the queue, when it is woken up before the expiration time, it attempts to continue its operation. At this moment, the target queue could be an invalid pointer, if the queue has been deleted. Furthermore, in FreeRTOS, the API function `vSemaphoreDelete` is used to delete a mutex, which actually uses `vQueueDelete` directly to perform the deleting operation. It is then possible to delete a mutex, which is held by a task. This can even be performed by a task which is not the holder of the mutex. When the holder returns the mutex or recursively takes it, the handler of the mutex can be an invalid pointer again.
5. Finally, we reveal an issue, which is not about the implementation, but related to the design. When the running task fails to take a mutex of which the holder has a lower priority, then the higher priority task is blocked and a lower priority task can be executed prior to a high priority task. This is called priority inversion. FreeRTOS adopts priority inheritance to solve this. To implement this, a base priority is introduced for each task. This records the original priority of each task. Afterwards, when a user calls the API function

Figure 7.4: Scenario for priority inversion issue



to change the priority of the mutex holder, which has inherited priority from another task, FreeRTOS only updates its base priority. Generally, this works fine. However, if the new priority of the mutex holder is even higher than its inherited priority, this can cause priority inversion again.

Consider the following scenario (see Fig.7.4). The upper chart shows the normal case for priority inheritance. When a higher priority task requests a mutex held by a lower priority task, the lower priority task inherits the priority of the higher priority task. When it is scheduled as running, it finishes the job, releases the mutex and recovers its own priority. Then the higher priority task can take the mutex and continue its task. This works fine in this case. However, the problem happens at the end of the scenario shown in the lower chart. There are three tasks in this scenario, MH, PT, HP, which represent the mutex holder, the preemptive task, and the higher priority task respectively. Firstly, MH starts running and takes the mutex. When the preemptive task, which has a higher priority than the mutex holder, joins the system, it pre-

empties the system and starts to execute. It then attempts to take the mutex and is blocked by the failure to take the unavailable mutex. MH resumes its execution and inherits priority from the PT. Before it finishes its job and releases the mutex, HP preempts the system and updates the priority of MH to a even higher priority than HP. Theoretically, the mutex holder has the highest priority in the system at the moment, since it should be running. However, because FreeRTOS only updates the base priority of the mutex holder, which is not used for scheduling, HP would continues its execution.

We have produced sample code for each of these cases, which are included in the supplementary material. They can be directly used with the FreeRTOS v7.3.0 simulator.

7.4 Summary

This chapter has evaluated our project in terms of achievement and the case study. It first described the achievement of this project against the objectives, which were stated in Sect. 1.5, and requirements, which were abstracted in Chap. 3. Then, it illustrated the animation and theorem proof for case studies in ProZ and Z/Eves respectively to show the correctness of the model.

It is worth noting that the theorems for the case studies are composed manually. Due to the complexity of the system, it is possible that there are mistakes in these theorems. Especially, when a mistake happens in the condition part of a theorem, the theorem is still able to be proved by Z/Eves. Because, according to the definition of implication, the false condition can imply anything. Therefore, if there is contradiction in the condition part, it is extremely difficult to locate. To avoid this kind of problem, we wrote some simple theorems to check the condition for these theorems, which simply state the condition implies to false. If there is a contradiction in the condition, which makes the condition false, the theorem can be proved easily by Z/Eves. Otherwise, the condition is correct.

Chapter 8

CONCLUSION AND FUTURE WORK

This chapter concludes the project, including experience gained during the project. It discusses possible future work for the project and reports on some attempts related to future work.

8.1 Conclusion

We have produced the first complete abstract specification of FreeRTOS. The model can be animated by the ProZ tool to show how FreeRTOS works. We have shown that the model is internally consistent by discharging all the verification conditions for well-definedness of the specification and by calculating the exact preconditions for the successful operation of each part of the FreeRTOS API function. Experiments were performed based on this model.

Firstly, we translated the base state of the model, and pre and postconditions of task-related schemas into the virtual model structure and function contracts of VCC, respectively, to build an annotated version of the FreeRTOS implementation, which was then verified by VCC.

However, due to time limitations, we were only able to focus on task API functions and a simplified implementation. This still shows the possibility of verifying software implementation directly with a high-level abstract specification and a code

verifier. The typical way to obtain the implementation in formal methods is refinement, which is normally complex and time consuming. Although it guarantees the correctness of the design and the implementation, it also increases the difficulty and expense of the development. However, not every piece of software requires this high level of correctness, and for such pieces of software, we expect the way we have demonstrated to be an easier approach to formalisation.

Secondly, we have developed an extended model for a multi-core platform. The requirements of the specification were inspired (but not limited) by the Multi-BSP model [35]. As a high-level abstract model, we described the general behaviours for each operation, which can then be refined for a specific architecture. We also validated the model with the theorem prover, Z/Eves, by performing the syntax and domain checking for all the definitions. This work demonstrates the reusability of our verified specification. It also builds the foundations for developing a verified RTOS for the multi-core platform.

During the project, there were plenty of difficulties in both modelling and verifying, for instance:

1. Some operations (e.g., sending items to a queue) may have intermediate states and it was necessary to find the correct way to describe them.
2. The Z/Eves prover provided reasonable proof automation. However, it also had some problems. For example, in some cases, the order of predicates in a schema had an effect on the result of a proof, which should not be the case.
3. We also found that although there were many similar proof goals, we needed to guide the prover to prove them repeatedly. This increased the work load dramatically. A well-designed proof structure would be helpful for simplifying the proof process.
4. We spent considerable time on VCC experiments. It was not as easy as we imagined at the beginning. For simple cases, it was very clear and easy.

However, when the hierarchical ownership tree, which is one of the most important concepts of VCC, became involved, it was easy to become confused. Ideally, VCC should notice and highlight all false assertions or predicates that lead to conflict. However, if conflicts happen in the code, VCC may consider *false* to be *true*, which makes later proofs meaningless. Furthermore, the error model report provided by VCC was not detailed enough for the user to understand why the proof failed. This increased the difficulty of working with VCC.

We also learnt that the Z notation is sufficient for verifying the correctness of a single function, which is similar to the unit test in normal software engineering. However, for some of the system properties of the RTOS (such as time-related issues), other techniques have to be used (like CSP).

We provide all the definitions and proofs in the supplementary material, so that the entire verification can be replayed to check its authenticity. This means that crucially our experiment is repeatable.

8.2 Future Work

For future work,

1. As mentioned in Chap. 6, promotion is an alternative way to extend our specification to a multi-core platform. This would be the first project we expect to finish as future work and provides several benefits for the specification. The most important is the possibility for code reuse during development. We have finished the first part of the model, the task model, with the Z promotion approach, which is described below.
2. Currently, we translated Z specification to VCC manually without any verification about data refinement. As Z notation provides richer data types and operations than VCC, we would expect that the data types of variables in our FreeRTOS model can be refined to data types, which can be directly used

with VCC, in Z notation and verified with Z/Eves automatically. Then directly use these refined variables in VCC. This will increase the reliability of our verification during translation process. Furthermore, when the data refinement is achieved and the variables in the Z specification can be directly used in VCC, the translation process will become more straightforward. We would also expect the translation process can be performed automatically by some software.

3. We expect to verify the whole implementation of FreeRTOS with VCC and our abstract specification. This will not only provide more solid evidence to support our expectations in the previous section, but also assure the correctness of the implementation of FreeRTOS.
4. The models, both for FreeRTOS and multi-core RTOS, can be refined, level by level, to executable source code in order to obtain a fully verified implementation of RTOS. This will provide a highly assured RTOS for industrial and research use.

We are keen to encourage others to use our specifications and proofs as benchmarks for comparing other notation and tools.

8.3 Task Model with Promotion

We attempted to apply the Z promotion technique to improve our multi-core model. The model¹ for task API functions is developed as an example to illustrate how that can be achieved, and how it can also be animated with ProZ [33]. To promote the task model for FreeRTOS to a multi-core platform, we consider that there is a sub system, which is an instance of FreeRTOS, executing in each core. Therefore, we can use the task model for FreeRTOS, except the definition for the `idle` task, directly as a part of the new model. In the FreeRTOS model, we define the `idle`

¹Although it can be very slow, this model can be animated with ProZ with some modification, which is stated in the comments of the source file. We also provide example animation in the supplementary material.

task as a global constant of $TASK$. However, for the multi-core model, we need one `idle` task for each core. We modify the definition of the `idle` task for a core as a “local” variable of the schema $TaskData$ and introduce constants, $idles$ and $cores$, globally with constraints similar to the model shown in Chap. 6.

$TaskData$
$tasks : \mathbb{F} TASK$ $running_task : TASK$ $idle : TASK$
$running_task \in tasks$ $idle \in tasks$ $idle \in idles$

As the `idle` task should never change for all operations, we appended an additional constraint to the schema $\Delta Task$, which states that the `idle` task is equal *before* and *after* the operation. The rest of the task model can be used directly. The base state schema for the multi-core model can then be defined as follows:

$Multi_Task$
$subTask : cores \rightarrow Task$ $exeCore : TASK \rightarrow cores$
$\forall c1, c2 : cores \mid c1 \neq c2 \bullet$ $(subTask(c1)).tasks \cap (subTask(c2)).tasks = \emptyset$ $dom\ exeCore \in \mathbb{F} TASK$ $dom\ exeCore = \bigcup \{c : cores \bullet (subTask(c)).tasks\}$

Firstly, the total function, $subTask$, is used to match cores and their local FreeRTOS. In Z, a state schema (e.g., $Task$) can be used as a data type. Its characteristic binding can be obtained by the operator “ θ ”, which binds the values of variables of an instance of the schema to the name of the variables of the schema, and this can then be used to assign the instance of the schema to a variable. Simply, it can be understood as a handler of an instance of the schema. Similar to the model described in Chap. 6, the partial surjective function, $exeCore$, is given to record the relationship between the tasks and their executable core. Three constraints are introduced to describe the properties of the system. First, we state that the $tasks$ set for each core are disjoint; then, the domain of $exeCore$ is defined as the union of

all the *tasks* for each core; based on the previous constraints, we can easily prove that the domain of *exeCore* has to be a finite set, because *tasks* for each core and the set of all the cores in the system, *cores*, are finite sets. However, it is hard to prove this with Z/Eves. We chose to append another constraint to allow Z/Eves to recognise this. Initially, there were only the *idle* tasks for each core in the system. Therefore, we defined the domain of the initial state of *exeCore* as equal to *idles*. In addition to this, the initialise schema for the task model, *Init_Task*, was used to help us to initialise the system for the multi-core model by initialising each subsystem for each core.

$$\begin{array}{l}
 \text{Init} \\
 \hline
 \text{Multi_Task}' \\
 \hline
 \text{dom } \text{exeCore}' = \text{idles} \\
 \forall c : \text{cores} \\
 \quad \bullet \exists \text{Task}' \mid \text{Init_Task} \\
 \quad \quad \bullet \text{subTask}'(c) = \theta \text{Task}' \wedge \text{exeCore}'((\text{subTask}'(c)).\text{idle}) = c
 \end{array}$$

Based on these definitions, the promotion schema was defined. This describes the link between the global and local operations. In this case, three promotion schemas are specified, because the creating, deleting, and remaining operations have different behaviours. Specifically:

1. The promotion schema for the creating operations can be defined as the schema, *PromoteC*, which takes ΔTask , the input variable (i.e., *target?*) and an extra temporary variable (i.e., *executeCore*) as parameters. In particular, ΔTask is actually provided by the local operation of FreeRTOS. As a new task, the value of *target?* should not be in the system. The first precondition is used to restrict the target task from being included in the domain of *exeCore*. At the same time, we need to identify which instance of *Task* is executing in the core, *executeCore*. The second predicate of the preconditions is used to achieve this. Finally, for the post-state of operations, the maplet of *executeCore* and the characteristic binding of the post state of *Task* is updated to *subTask* and the relation between *target?* and *executeCore* is added to *exeCore*.

$PromoteC$ $\Delta Multi_Task$ $\Delta Task$ $target? : TASK$ $executeCore : cores$
$target? \notin \text{dom } executeCore$ $subTask(executeCore) = \theta Task$ $subTask' = subTask \oplus \{executeCore \mapsto \theta Task'\}$ $exeCore' = exeCore \oplus \{target? \mapsto executeCore\}$

2. Unlike the promotion operation for creating tasks, we need to remove the target task from the system. The promotion schema for deleting tasks, $PromoteD$, needs to remove $target?$ from $exeCore$, instead of adding it to $exeCore$. The other constraints should be the same, in order to identify the local $Task$ and update it with the post state of $Task$ to $subTask$.
3. For the remaining operations, which can be described with promotion, there is no effect on the function executable core, whether suspending tasks, resuming tasks or changing the priority of tasks. Similar to the previous case, the difference between the promotion schemas, $Promote$ and $PromoteC$, is the function $exeCore$, which stays the same.

Furthermore, with the promotion schemas, the sub-definitions for creating tasks can be defined as:

$$createTaskN_MT \cong \exists \Delta Task \bullet CreateTaskN_T \wedge PromoteC$$

$$createTaskS_MT \cong \exists \Delta Task \bullet CreateTaskS_T \wedge PromoteC$$

Like the multi-core model described in Chap. 6, the schema, $findACore_MT$, is introduced to locate the best position for a new task if its executable core is not specified by a user. The creating operation can be defined similarly. The same strategy can be applied to specify the rest of the operations.

We also specify the behaviours of migrating a task from one core to another, which also contains four cases, as in Sect. 6.2.4. However, the definitions are slightly

more complex than those for the old multi-core model. The reason for this is that in the old model, the function *state*, *priority*, *etc.*, are global total functions, which are easy to access by the operation schema. Instead, in the promotion model, they are private for each subsystem and can only be accessed by the function *subTask*. Due to the complexity of the expressions, we introduce two assistant variables, *srcSys* and *tarSys*, to represent the subsystems for the original core and the target core respectively. The behaviour of this operation can then be defined similarly to the one for the old model.

8.4 Summary

This chapter has summarised the project and suggested some future works after the project. It firstly described the whole process of the project, and identified the difficulties we experienced during the project. It then shows the potential future works based on the results of the project. Finally, as an example, we presented our attempt on specifying multi-core task model with promotion technique.

In general, in this project, we produced the first complete abstract specification in Z of FreeRTOS, together with proofs of consistency (well-definedness, initialisation, precondition, and a few properties). Then, the model is extended for multicore platform, with basic proofs of well-definedness (including syntax checking and domain validating). The abstract characterisation of both models is a first step towards a verified implementation of FreeRTOS on multicore. We were the first to promote FreeRTOS as a pilot project in VSI, and the work presented continues this by establishing a benchmark for others to follow. We believe that this is an important contribution to both the verification community and also the embedded systems community. We also demonstrated the possibility to verify a software system by combining the formal modelling (Z model) and code verifier (VCC), which can be an easier approach for improving the quality of the software.

Appendix A

INTRODUCTORY APPENDIX

There are twelve appendixes provided after this introductory appendix. They contain auxiliary information for the main body of the thesis. Specifically,

- (a) The summary of frequently used proof commands is shown in Appendix B. It is helpful for understanding the proof script provided in supplementary material.
- (b) As stated in Sect. 4.6, the mapping between API functions and the precondition for schema interface from FreeRTOS model can be found in Appendix C.
- (c) The specification for our FreeRTOS model can be found in Appendix D-G. They include all the definitions, schemas and theorems (i.e. precondition theorems for operation schemas, auxiliary theorems and theorems for some system properties). Unfortunately, the proof scripts for the model cannot be included, due to the length of the script itself. However, they can be found in supplementary material and can be used directly in the theorem prover, Z/Eves.
- (d) Similarly, the specification for multi-core model can also be found in Appendix H-K. And the the specification for the multi-core task model with promotion technique is shown in Appendix L.

(e) Finally, the VCC annotated source code of task related API functions (i.e. creating, deleting, suspending, resuming tasks and changing the priority of tasks) is listed in Appendix M.

Appendix B

SUMMARY OF Z/EVES PROOF

COMMANDS

We summarise the proof commands used in proving of the model. For full instruction of the proof commands, please see Chap. 5 of [32].

prove The prover automatically applies sequences of proof commands. For example, *simplify*, *rewrite*, *rearrange*. Besides this, the mathematical rules included in Z/Eves' mathematical toolkit [70] are applied, if possible.

prove by reduce The prover repeatedly reduces the current proof goal. In addition to what **prove** does, the prover expands all names.

with enabled (theorem) This is a prefix that is applied to the *prove*, *prove by reduce*, or an already prefixed command. Many inefficient rules are disabled by default, and this prefix enables them for the current command. For example, *with enabled (applyOverride) prove* allows the prover to use the disabled theorem *applyOverride* within the scope of the *prove* command.

with disabled (theorem) This is similar to the previous command, except that it disables the theorem rather than enabling it.

with normalization This is also a prefix for prove commands. It allows the prover to use “if-then-else” normal-form to represent all logical connectives [70].

instantiate This command allows the prover to instantiate quantified variables (universal in the assumptions, existential in the goal).

apply theorem As mentioned above, there are plenty of disabled rules in Z/Eves’ mathematical toolkit. This command applies the specified theorem to rewrite the goal.

use theorem This command allows a specified theorem to be used to deduce additional assumptions.

extensionality A theorem included in Z/Eves Mathematical Toolkits [70], which defined as:

$$X = Y \Leftrightarrow (\forall x : X \bullet x \in Y) \wedge (\forall y : Y \bullet y \in X)$$

Appendix C

SUMMARY OF INTERFACE

Table C.1: API mappings & preconditions for operations

API	Operation	Precondition
xTaskCreate	$\text{CreateTask}_T \triangleq \text{CreateTaskN}_T$ $\vee \text{CreateTasks}_T$	$\text{state}(\text{target}?) = \text{nonexistent}$
	CreateTaskN_T	$\text{state}(\text{target}?) = \text{nonexistent}$
	CreateTasks_T	$\text{newprv}^? \leq \text{priority}(\text{running_task})$
		$\text{state}(\text{target}?) = \text{nonexistent}$
		$\text{newprv}^? > \text{priority}(\text{running_task})$
vTaskDelete	$\text{DeleteTask}_T \triangleq \text{DeleteTaskN}_T$ $\vee \text{DeleteTasks}_T$	$\text{target}^? \in \text{tasks} \setminus \{\text{idle}\}$ $\text{state}(\text{target}?) = \text{running} \Rightarrow$ $(\exists \text{topReady}! : \text{state}^~(\{\{\text{ready}\}\})$ $\bullet (\forall t : \text{state}^~(\{\{\text{ready}\}\}) \mid$ $\bullet \text{priority}(\text{topReady}!) \geq \text{priority}(t)))$
	DeleteTaskN_T	$\text{target}^? \in \text{tasks} \setminus \{\text{idle}\}$ $\text{state}(\text{target}?) \in \{\text{ready}, \text{blocked}, \text{suspended}\}$
	DeleteTasks_T	$\text{target}^? \in \text{tasks} \setminus \{\text{idle}\}$ $\text{state}(\text{target}?) \in \{\text{running}\}$ $\exists \text{topReady}! : \text{state}^~(\{\{\text{ready}\}\}) \mid$ $\bullet (\forall t : \text{state}^~(\{\{\text{ready}\}\}) \mid$ $\bullet \text{priority}(\text{topReady}!) \geq \text{priority}(t))$
-	$\text{ExecuteRunningTask}_T$	$\exists \text{phys_context}' : \text{CONTEXT}$ $\bullet \text{phys_context}' \neq \text{phys_context}$

Table C.2: API mappings & preconditions for operations(continue)

VTaskSuspend	$SuspendTask_T \triangleq SuspendTaskN_T$ $\vee SuspendTasks_T$ $\vee SuspendTaskO_T$	$target? \in tasks \setminus \{idle\}$ $state(target?) = running \Rightarrow$ $(\exists topReady! : state \sim \{\{ready\}\} \Downarrow$ $\bullet (\forall t : state \sim \{\{ready\}\} \Downarrow$ $\bullet priority(topReady!) \geq priority(t))$
	$SuspendTaskN_T$	$target? \in tasks \setminus \{idle\}$ $state(target?) \in \{ready, blocked\}$
	$SuspendTasks_T$	$target? \in tasks \setminus \{idle\}$ $state(target?) \in \{running\}$ $\exists topReady! : state \sim \{\{ready\}\} \Downarrow$ $\bullet (\forall t : state \sim \{\{ready\}\} \Downarrow$ $\bullet priority(topReady!) \geq priority(t))$
	$SuspendTaskO_T$	$state(target?) = suspended$
VTaskResume	$ResumeTask_T \triangleq ResumeTaskN_T$ $\vee ResumeTasks_T$	$state(target?) = suspended$
	$ResumeTaskN_T$	$state(target?) = suspended$ $priority(target?) \leq priority(running_task)$
	$ResumeTasks_T$	$state(target?) = suspended$ $priority(target?) > priority(running_task)$

Table C.3: API mappings & preconditions for operations(continue)

TaskPrioritySet	$\text{ChangeTaskPriority}_T$	$state(target?) \neq \text{nonexistent}$ $target? = \text{idle} \Rightarrow \text{newpri?} = 0$ $\exists \text{topReady!} : state \sim (\{\text{ready}\}) \mid$ $\vee \text{ChangeTaskPriorityS}_T$ <ul style="list-style-type: none"> • $(state(target?) \in \{\text{running}\})$ $\wedge \neg (\forall \text{rtask} : \text{TASK} \mid state(\text{rtask}) = \text{ready}$ <ul style="list-style-type: none"> • $\text{newpri?} \geq \text{priority}(\text{rtask}))$ • $\text{newpri?} < \text{priority}(\text{topReady!})$ $\wedge (\forall t : state \sim (\{\text{ready}\}) \mid$ <ul style="list-style-type: none"> • $\text{priority}(\text{topReady!}) \geq \text{priority}(t))$
	$\text{ChangeTaskPriorityN}_T$	$state(target?) = \text{ready} \Rightarrow$ $\text{newpri} \leq \text{priority}(\text{running_task})$ $state(target?) = \text{running}$ $\Rightarrow (\forall t : state \sim (\{\text{ready}\}) \mid$ <ul style="list-style-type: none"> • $\text{newpri?} \geq \text{priority}(t)$ $state(target?) \neq \text{nonexistent}$ $target? = \text{idle} \Rightarrow \text{newpri?} = 0$
	$\text{ChangeTaskPriorityS}_T$	$state(target?) = \text{ready}$ $\text{newpri?} > \text{priority}(\text{running_task})$ $target? = \text{idle} \Rightarrow \text{newpri?} = 0$
	$\text{ChangeTaskPriorityD}_T$	$state(target?) \in \{\text{running}\}$ $target? = \text{idle} \Rightarrow \text{newpri?} = 0$ $\exists \text{topReady!} : state \sim (\{\text{ready}\}) \mid$ <ul style="list-style-type: none"> • $\text{newpri} < \text{priority}(\text{topReady!})$ $\wedge (\forall t : state \sim (\{\text{ready}\}) \mid$ <ul style="list-style-type: none"> • $\text{priority}(\text{topReady!}) \geq \text{priority}(t)$

Table C.4: API mappings & preconditions for operations(continue)

xQueueCreate	<i>CreateQueue_TQ</i>	$que? \notin queue$ $size? > 0$
xQueueDelete	<i>DeleteQueue_TQ</i>	$que? \in queue$ $que? \notin ran\ wait_snd \cup ran\ wait_rcv$ $que? \notin ran\ release_snd \cup ran\ release_rcv$
xQueueSend	<i>QueueSend_TQ</i> $\equiv QueueSendN_TQ$ $\vee QueueSendW_TQ$ $\vee QueueSendWS_TQ$ $\vee QueueSendF_TQ$	$que? \in queue$ $q_size(que?) = q_max(que?) \Rightarrow running_task \neq idle$
	<i>QueueSendN_TQ</i>	$que? \in queue$ $q_size(que?) < q_max(que?)$ $que? \notin ran\ wait_rcv$
	<i>QueueSendW_TQ</i>	$que? \in queue$ $q_size(que?) < q_max(que?)$ $\forall wr : wait_rcv \sim \{\{que?\}\} \triangleright \bullet priority(running_task) \geq priority(wr)$
	<i>QueueSendWS_TQ</i>	$que? \in queue$ $q_size(que?) < q_max(que?)$ $\exists topReady! : wait_rcv \sim \{\{que?\}\} \triangleright$ $\quad (\forall wr : wait_rcv \sim \{\{que?\}\} \triangleright \bullet priority(topReady!) \geq priority(wr))$ $\quad \bullet priority(topReady!) > priority(running_task)$
	<i>QueueSendF_TQ</i>	$que? \in queue$ $q_size(que?) = q_max(que?)$ $running_task \neq idle$

Table C.5: API mappings & preconditions for operations(continue)

xQueueReceive	$QueueReceive_TQ$	$que? \in queue$
	$\cong QueueReceiveN_TQ$	$q_size(que?) = 0 \Rightarrow running_task \neq idle$
	$\vee QueueReceiveW_TQ$	
	$\vee QueueReceiveWS_TQ$	
	$\vee QueueReceiveF_TQ$	
	$QueueReceiveN_TQ$	$que? \in queue$
		$q_size(que?) > 0$
		$que? \notin ran\ wait_snd$
	$QueueReceiveW_TQ$	$que? \in queue$
		$q_size(que?) > 0$
		$\forall us : wait_snd \sim \{\{que?\}\} \bullet \bullet priority(running_task) \geq priority(us)$
	$QueueReceiveWS_TQ$	$que? \in queue$
		$q_size(que?) > 0$
		$\exists topReady! : wait_snd \sim \{\{que?\}\} \bullet$
		$\quad (\forall us : wait_snd \sim \{\{que?\}\} \bullet \bullet priority(topReady!) \geq priority(us))$
		$\quad \bullet priority(topReady!) > priority(running_task)$
	$QueueReceiveE_TQ$	$que? \in queue$
		$q_size(que?) = 0$
		$running_task \neq idle$
vTaskDelayUntil	$DelayUntil_TQT$	$n? > clock$

Table C.6: API mappings & preconditions for operations(continue)

vSemaphoreCreateBinary	<i>CreateBinarySemaphore_TQTM</i>	$sem? \notin queue$
vSemaphoreDelete	<i>DeleteBinarySemaphore_TQTM</i>	$sem? \in semaphore$
xSemaphoreCreateMutex	<i>CreateMutex_TQTM</i>	$mut? \notin queue$
vSemaphoreDelete	<i>DeleteMutex_TQTM</i>	$mut? \in mutex \setminus dom\ mutex_holder$
xSemaphoreTake	<i>MutexTake_TQTM</i>	$mut? \in mutex$
	$\equiv MutexTakeNonInh_TQTM$	$q_size(mut?) = 0 \Rightarrow$
	$\vee MutexTakeNinh_TQTM$	$(running_task \neq idle$
	$\vee MutexTakeEnonInh_TQTM$	$\wedge n? > clock$
	$\vee MutexTakeElnheritReady_TQTM$	$\wedge mutex_holder(mut?) \neq idle)$
	$\vee MutexTakeElnheritHolder_TQTM$	
	$\vee MutexTakeRecursive_TQTM$	
	<i>MutexTakeNonInh_TQTM</i>	$mut? \in mutex$
		$q_size(mut?) > 0$
		$running_task \notin dom\ base_priority$
	<i>MutexTakeNinh_TQTM</i>	$mut? \in mutex$
		$q_size(mut?) > 0$
		$running_task \in dom\ base_priority$
	<i>MutexTakeEnonInh_TQTM</i>	$mut? \in dom\ mutex_holder$
		$priority(running_task) \leq priority(mutex_holder(mut?))$
		$n? > clock$
		$running_task \neq idle$
		$running_task \neq mutex_holder(mut?)$

Table C.7: API mappings & preconditions for operations(continue)

<i>MutexTakeEInheritReady_TQTM</i>	<i>mut?</i> ∈ dom <i>mutex_holder</i>
	<i>priority</i> (<i>running_task</i>) > <i>priority</i> (<i>mutex_holder</i> (<i>mut?</i>))
	<i>n?</i> > <i>clock</i>
	<i>mutex_holder</i> (<i>mut?</i>) ≠ <i>idle</i>
	<i>state</i> (<i>mutex_holder</i> (<i>mut?</i>)) ≠ <i>ready</i>
<i>MutexTakeEInheritHolder_TQTM</i>	<i>mut?</i> ∈ dom <i>mutex_holder</i>
	<i>priority</i> (<i>running_task</i>) > <i>priority</i> (<i>mutex_holder</i> (<i>mut?</i>))
	<i>n?</i> > <i>clock</i>
	<i>mutex_holder</i> (<i>mut?</i>) ≠ <i>idle</i>
	<i>state</i> (<i>mutex_holder</i> (<i>mut?</i>)) = <i>ready</i>
<i>MutexTakeRecursive_TQTM</i>	<i>mut?</i> ∈ dom <i>mutex_holder</i>
	<i>running_task</i> = <i>mutex_holder</i> (<i>mut?</i>)
xSemaphoreGive <i>MutexGive_TQTM</i>	<i>mut?</i> ∈ dom <i>mutex_holder</i>
	<i>running_task</i> = <i>mutex_holder</i> (<i>mut?</i>)
	≡ <i>MutexGiveNRecursive_TQTM</i>
	∨ <i>MutexGiveNnonInh_TQTM</i>
	∨ <i>MutexGiveNInh_TQTM</i>
	∨ <i>MutexGiveNInhS_TQTM</i>
	∨ <i>MutexGiveWnonInhN_TQTM</i>
	∨ <i>MutexGiveWnonInhS_TQTM</i>
	∨ <i>MutexGiveWinhN_TQTM</i>
	∨ <i>MutexGiveWinhSR_TQTM</i>
	∨ <i>MutexGiveWinhSW_TQTM</i>

Table C.8: API mappings & preconditions for operations(continue)

<i>MutexGiveNRecursive_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) > 1$
<i>MutexGiveNnonInh_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $mut? \notin \text{ran } wait_rcv$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $base_priority(running_task) = priority(running_task)$
<i>MutexGiveNinh_N_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $mut? \notin \text{ran } wait_rcv$ $base_priority(running_task) \neq priority(running_task)$ $\forall rt : state^{\sim}(\{ready\}) \mid \bullet base_priority(running_task) \geq priority(rt)$
<i>MutexGiveNinh_S_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $mut? \notin \text{ran } wait_rcv$ $base_priority(running_task) \neq priority(running_task)$ $\exists topReady! : state^{\sim}(\{ready\}) \mid \forall rt : state^{\sim}(\{ready\}) \mid$ <ul style="list-style-type: none"> • $priority(topReady!) \geq priority(rt)$ • $base_priority(running_task) < priority(topReady!)$

Table C.9: API mappings & preconditions for operations(continue)

<i>MutexGiveWhenInhN_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $base_priority(running_task) = priority(running_task)$ $\forall wr : wait_rcv \sim (\{mut?\}) \mid \bullet priority(running_task) \geq priority(wr)$
<i>MutexGiveWhenInhS_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $base_priority(running_task) = priority(running_task)$ $\exists topReady! : wait_rcv \sim (\{mut?\}) \mid$ $\quad \mid (\forall wr : wait_rcv \sim (\{mut?\}) \mid \bullet priority(topReady!) \geq priority(wr))$ $\quad \bullet priority(topReady!) > priority(running_task)$
<i>MutexGiveWhenN_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $\forall wr : wait_rcv \sim (\{mut?\}) \mid \bullet base_priority(running_task) \geq priority(wr)$ $base_priority(running_task) \neq priority(running_task)$ $\forall rt : state \sim (\{ready\}) \mid \bullet base_priority(running_task) \geq priority(rt)$

Table C.10: API mappings & preconditions for operations(continue)

<i>MutexGiveWInhSR_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $base_priority(running_task) \neq priority(running_task)$ $\exists topReady! : state \sim \{\{ready\}\} \perp$ <ul style="list-style-type: none"> • $(\forall rt : state \sim \{\{ready\}\} \perp) \bullet priority(topReady!) \geq priority(rt)$ $\wedge (\forall wr : wait_rcv \sim \{\{mut?\}\} \perp)$ <ul style="list-style-type: none"> • $priority(topReady!) > priority(wr)$ • $priority(topReady!) > base_priority(running_task)$
<i>MutexGiveWInhSW_TQTM</i>	$mut? \in \text{dom } mutex_holder$ $running_task = mutex_holder(mut?)$ $mutex_recursive(mut?) = 1$ $base_priority(running_task) \neq priority(running_task)$ $\exists topWaiting! : wait_rcv \sim \{\{mut?\}\} \perp$ <ul style="list-style-type: none"> • $(\forall wr : wait_rcv \sim \{\{mut?\}\} \perp) \bullet priority(topWaiting!) \geq priority(wr)$ $\wedge (\forall rt : state \sim \{\{ready\}\} \perp)$ <ul style="list-style-type: none"> • $priority(topWaiting!) \geq priority(rt)$ • $priority(topWaiting!) > base_priority(running_task)$

Appendix D

SPECIFICATION FOR TASK MODEL

[*CONTEXT*, *TASK*]

| *bare_context* : *CONTEXT*
| *idle* : *TASK*

STATE ::= *nonexistent* | *ready* | *blocked* | *suspended* | *running*

transition == (*blocked* × {*nonexistent*, *ready*, *running*, *suspended*})
∪ (*nonexistent* × {*ready*, *running*})
∪ (*ready* × {*nonexistent*, *running*, *suspended*})
∪ (*running* × {*blocked*, *nonexistent*, *ready*, *suspended*})
∪ (*suspended* × {*nonexistent*, *ready*, *running*})

theorem grule gTransitionType
transition ∈ ℙ(*STATE* × *STATE*)

theorem rule lInTransition
∀ *l*, *r* : *STATE* | (*l*, *r*) ∈ {*nonexistent* ↦ *ready*, *running* ↦ *ready*,
blocked ↦ *ready*, *suspended* ↦ *ready*, *ready* ↦ *running*,
blocked ↦ *running*, *suspended* ↦ *running*,
nonexistent ↦ *running*, *running* ↦ *suspended*,
ready ↦ *suspended*, *blocked* ↦ *suspended*,
running ↦ *blocked*, *running* ↦ *nonexistent*,
ready ↦ *nonexistent*, *blocked* ↦ *nonexistent*,
suspended ↦ *nonexistent*}

- (*l*, *r*) ∈ *transition*

TaskData

$tasks : \mathbb{F} \text{ TASK}$

$running_task : \text{TASK}$

$running_task \in tasks$

$idle \in tasks$

Init_TaskData

TaskData'

$tasks' = \{idle\}$

$running_task' = idle$

theorem TaskDataInit

$\exists \text{ TaskData}' \bullet \text{Init_TaskData}$

StateData

$state : \text{TASK} \rightarrow \text{STATE}$

$state(idle) \in \{ready, running\}$

Init_StateData

StateData'

$state' = (\lambda x : \text{TASK} \bullet nonexistent) \oplus \{(idle \mapsto running)\}$

theorem StateDataInit

$\exists \text{ StateData}' \bullet \text{Init_StateData}$

ContextData

$phys_context : \text{CONTEXT}$

$log_context : \text{TASK} \rightarrow \text{CONTEXT}$

Init_ContextData

ContextData'

$phys_context' = bare_context$

$log_context' = (\lambda x : \text{TASK} \bullet bare_context)$

theorem ContextDataInit
$$\exists \text{ContextData}' \bullet \text{Init_ContextData}$$

$$\text{PrioData}$$

$$\text{priority} : \text{TASK} \rightarrow \mathbb{N}$$

$$\text{priority}(\text{idle}) = 0$$

$$\text{Init_PrioData}$$

$$\text{PrioData}'$$

$$\text{priority}' = (\lambda x : \text{TASK} \bullet 0)$$
theorem PrioDataInit
$$\exists \text{PrioData}' \bullet \text{Init_PrioData}$$

$$\text{Task}$$

$$\text{TaskData}$$

$$\text{StateData}$$

$$\text{ContextData}$$

$$\text{PrioData}$$

$$\text{tasks} = \text{TASK} \setminus (\text{state} \sim (\{ \text{nonexistent} \} \cup))$$

$$\text{state} \sim (\{ \text{running} \} \cup) = \{ \text{running_task} \}$$

$$\forall pt : \text{state} \sim (\{ \text{ready} \} \cup) \bullet \text{priority}(\text{running_task}) \geq \text{priority}(pt)$$

$$\Delta \text{Task}$$

$$\text{Task}$$

$$\text{Task}'$$

$$\forall st : \text{TASK} \mid \text{state}'(st) \neq \text{state}(st)$$

$$\bullet \text{state}(st) \mapsto \text{state}'(st) \in \text{transition}$$

$$f : \mathbb{P} \text{TASK} \rightarrow \text{TASK}$$

$$\langle\langle \text{findDelegate} \rangle\rangle$$

$$\forall \text{Task}; a : \mathbb{P} \text{TASK}; g : \text{TASK} \leftrightarrow \mathbb{Z} \bullet$$

$$f(a) \in a \wedge a \subseteq \text{dom } g \wedge$$

$$(\forall t : a \bullet g(f(a)) \geq g(t))$$
theorem TaskProperty1
$$\forall \text{Task} \bullet \text{state}(\text{running_task}) = \text{running}$$

theorem TaskProperty2
$$\forall Task \bullet \forall t : TASK \mid t \in state \sim (\{blocked\}) \bullet t \in tasks$$
theorem TaskProperty3
$$\forall Task \bullet \forall t : state \sim (\{ready\}) \bullet t \in tasks \setminus \{running_task\}$$
theorem TaskProperty6
$$\forall Task; t : TASK \mid 0 < priority(t) \bullet idle \neq t$$
*Init_Task**Task'**Init_TaskData**Init_StateData**Init_ContextData**Init_PrioData***theorem** TaskInit
$$\exists Task' \bullet Init_Task$$
Reschedule $\Delta Task$ *target?* : TASK*tasks?* : $\mathbb{F} TASK$ *st?* : STATE*pri?* : TASK $\rightarrow \mathbb{N}$ $tasks' = tasks?$ $running_task' = target?$ $state' = state \oplus \{(target? \mapsto running), (running_task \mapsto st?)\}$ $phys_context' = log_context(target?)$ $log_context' = log_context \oplus \{(running_task \mapsto phys_context)\}$ $priority' = pri?$

$$disableReschedule \hat{=} [Task \mid false] \wedge Reschedule$$
CreateTaskN_T $\Delta Task$ *target?* : TASK*newpri?* : \mathbb{N} $state(target?) = nonexistent$

$$\begin{aligned}
& newpri? \leq priority(running_task) \\
& tasks' = tasks \cup \{target?\} \\
& running_task' = running_task \\
& state' = state \oplus \{(target? \mapsto ready)\} \\
& \exists ContextData \\
& priority' = priority \oplus \{(target? \mapsto newpri?)\}
\end{aligned}$$

CreateTaskN_TFSBSig

Task
target? : TASK
newpri? : N

state(target?) = nonexistent
newpri? ≤ priority(running_task)

theorem rule runningUpdate

$$\begin{aligned}
& \forall f : TASK \rightarrow STATE; g : TASK \rightarrow STATE \mid running \notin \text{ran } g \\
& \quad \wedge (f \sim \{\{running\}\}) \cap \text{dom } g = \emptyset \\
& \quad \bullet (f \oplus g) \sim \{\{running\}\} = f \sim \{\{running\}\}
\end{aligned}$$

theorem rule setminUpdate

$$\begin{aligned}
& \forall f : TASK \rightarrow STATE; g : TASK \rightarrow STATE \\
& \quad \bullet TASK \setminus ((f \oplus g) \sim \{\{nonexistent\}\}) \\
& \quad = TASK \setminus (f \sim \{\{nonexistent\}\}) \setminus (g \sim \{\{nonexistent\}\}) \\
& \quad \cup (\text{dom } g \setminus (g \sim \{\{nonexistent\}\}))
\end{aligned}$$

theorem CreateTaskN_T_vc_ref

$$\forall CreateTaskN_TFSBSig \mid true \bullet \text{pre } CreateTaskN_T$$

CreateTaskS_T

$\Delta Task$
target? : TASK
newpri? : N

state(target?) = nonexistent
newpri? > priority(running_task)
 $\exists st? : STATE; tasks? : \mathbb{F} TASK; pri? : TASK \rightarrow \mathbb{N}$
 $\mid st? = ready \wedge tasks? = tasks \cup \{target?\}$
 $\quad \wedge pri? = priority \oplus \{(target? \mapsto newpri?)\}$
 $\bullet Reschedule$

CreateTaskS_TFSBSig

Task

target? : *TASK*

newpri? : \mathbb{N}

$state(target?) = nonexistent$

$newpri? > priority(running_task)$

theorem *CreateTaskS_T_vc_ref*

$\forall CreateTaskS_TFSBSig \mid true \bullet \text{pre } CreateTaskS_T$

$CreateTask_T \cong CreateTaskN_T \vee CreateTaskS_T$

DeleteTaskN_T

$\Delta Task$

target? : *TASK*

topReady! : *TASK*

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{ready, blocked, suspended\}$

$tasks' = tasks \setminus \{target?\}$

$running_task' = running_task$

$state' = state \oplus \{(target? \mapsto nonexistent)\}$

$phys_context' = phys_context$

$log_context' = log_context \oplus \{(target? \mapsto bare_context)\}$

$\exists PrioData$

$topReady! = running_task$

DeleteTaskN_TFSBSig

Task

target? : *TASK*

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{ready, blocked, suspended\}$

theorem *finsetIsFinset*

$\forall X : \mathbb{F} TASK; x : TASK \bullet X \setminus \{x\} \in \mathbb{F} TASK$

theorem *DeleteTaskN_T_vc_ref*

$\forall DeleteTaskN_TFSBSig \mid true \bullet \text{pre } DeleteTaskN_T$

DeleteTaskS_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{running\}$

$state(topReady!) = ready$

$\forall t : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(t)$

$tasks' = tasks \setminus \{target?\}$

$running_task' = topReady!$

$state' = state \oplus \{(topReady! \mapsto running), (target? \mapsto nonexistent)\}$

$phys_context' = log_context(topReady!)$

$log_context' = log_context \oplus \{(target? \mapsto bare_context)\}$

$\exists PrioData$

DeleteTaskS_TFSBSig

Task

$target? : TASK$

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{running\}$

theorem *lDeleteTaskS_T_Lemma*

$\forall Task; topReady!, target? : TASK$

$| target? \in tasks \setminus \{idle\}$

$\wedge state(target?) \in \{running\}$

$\wedge state(topReady!) = ready$

$\wedge (\forall rtsk : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(rtsk))$

$\bullet \neg (Task[log_context := log_context \oplus \{(target?, bare_context)\},$

$phys_context := log_context(topReady!),$

$running_task := topReady!,$

$state := state \oplus$

$(\{(target?, nonexistent)\} \cup \{(topReady!, running)\}),$

$tasks := tasks \setminus \{target?\}]$

$\wedge (st \in TASK$

$\wedge \neg (state \oplus (\{(target?, nonexistent)\} \cup$

$\{(topReady!, running)\}))(st) = state(st)$

$\Rightarrow (state(st), (state \oplus (\{(target?, nonexistent)\} \cup$

$\{(topReady!, running)\}))(st) \in transition)$

$\Rightarrow t \in TASK$

$\wedge state(t) = ready$

$\wedge \neg priority(topReady!) \geq priority(t)$

theorem DeleteTaskS_T_vc_ref
$$\forall \text{DeleteTaskS_TFSSig} \mid \text{true} \bullet \text{pre DeleteTaskS_T}$$

$$\text{DeleteTask_T} \cong \text{DeleteTaskN_T} \vee \text{DeleteTaskS_T}$$

$$\text{ExecuteRunningTask_T}$$

$$\Delta \text{Task}$$

$$\text{target!} : \text{TASK}$$

$$\exists \text{TaskData}$$

$$\exists \text{StateData}$$

$$\text{phys_context}' \neq \text{phys_context}$$

$$\text{log_context}' = \text{log_context}$$

$$\exists \text{PrioData}$$

$$\text{target!} = \text{running_task}$$

$$\text{ExecuteRunningTask_TFSSig}$$

$$\text{Task}$$

$$\exists \text{phys_context}' : \text{CONTEXT} \bullet \text{phys_context}' \neq \text{phys_context}$$
theorem ExecuteRunningTask_T_vc_ref
$$\forall \text{ExecuteRunningTask_TFSSig} \mid \text{true} \bullet \text{pre ExecuteRunningTask_T}$$

$$\text{SuspendTaskN_T}$$

$$\Delta \text{Task}$$

$$\text{target?} : \text{TASK}$$

$$\text{topReady!} : \text{TASK}$$

$$\text{target?} \in \text{tasks} \setminus \{\text{idle}\}$$

$$\text{state}(\text{target?}) \in \{\text{ready}, \text{blocked}\}$$

$$\exists \text{TaskData}$$

$$\text{state}' = \text{state} \oplus \{(\text{target?} \mapsto \text{suspended})\}$$

$$\exists \text{ContextData}$$

$$\exists \text{PrioData}$$

$$\text{topReady!} = \text{running_task}$$

$$\text{SuspendTaskN_TFSSig}$$

$$\text{Task}$$

$$\text{target?} : \text{TASK}$$

$$\text{target?} \in \text{tasks} \setminus \{\text{idle}\}$$

$$state(target?) \in \{ready, blocked\}$$

theorem SuspendTaskN_T_vc_ref

$$\forall SuspendTaskN_TFSSig \mid true \bullet \text{pre } SuspendTaskN_T$$

$$SuspendTaskS_T$$

$$\Delta Task$$

$$target? : TASK$$

$$topReady! : TASK$$

$$target? \in tasks \setminus \{idle\}$$

$$state(target?) \in \{running\}$$

$$state(topReady!) = ready$$

$$\forall t : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(t)$$

$$\exists st? : STATE \mid st? = suspended$$

- Reschedule[*tasks/tasks?, priority/pri?, topReady!/target?*]

theorem TaskProperty4

$$\forall Task \mid SuspendTaskS_T$$

- $state'(running_task) = suspended$

$$\wedge (\forall t : state \sim (\{ready\}) \bullet priority(running_task') \geq priority(t))$$

theorem TaskProperty5

$$\forall Task \bullet \forall t : TASK \mid t \notin tasks \bullet state(t) = nonexistent$$

$$SuspendTaskS_TFSSig$$

$$Task$$

$$target? : TASK$$

$$target? \in tasks \setminus \{idle\}$$

$$state(target?) \in \{running\}$$

theorem lSuspendTaskS_T_Lemma

$$\forall Task; target?, topReady! : TASK$$

$$\mid target? \in tasks \setminus \{idle\}$$

$$\wedge state(target?) \in \{running\}$$

$$\wedge state(topReady!) = ready$$

$$\wedge (\forall rtsk : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(rtsk))$$

- $\neg (Task[log_context := log_context \oplus$

$$\{(running_task, phys_context)\},$$

$$phys_context := log_context(topReady!),$$

$$running_task := topReady!,$$

$$\begin{aligned}
& \text{state} := \text{state} \oplus \\
& \quad (\{(running_task, suspended)\} \cup \{(topReady!, running)\}) \\
& \wedge (st \in TASK \\
& \quad \wedge \neg (\text{state} \oplus (\{(running_task, suspended)\} \\
& \quad \cup \{(topReady!, running)\}))st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(running_task, suspended)\} \\
& \quad \cup \{(topReady!, running)\}))st) \in transition) \\
& \Rightarrow t \in TASK \\
& \quad \wedge \text{state}(t) = ready \\
& \quad \wedge \neg \text{priority}(topReady!) \geq \text{priority}(t)
\end{aligned}$$

theorem SuspendTaskS_T_vc_ref

$$\forall \text{SuspendTaskS_TFSSig} \mid true \bullet \text{pre SuspendTaskS_T}$$

SuspendTaskO_T
$\exists Task$
$target? : TASK$
$topReady! : TASK$
<hr/>
$\text{state}(target?) \in \{suspended\}$
$topReady! = running_task$

$\text{SuspendTaskO_TFSSig}$
$Task$
$target? : TASK$
$topReady! : TASK$
<hr/>
$\text{state}(target?) \in \{suspended\}$

theorem SuspendTaskO_T_vc_ref

$$\forall \text{SuspendTaskO_TFSSig} \mid true \bullet \text{pre SuspendTaskO_T}$$

$$\begin{aligned}
\text{SuspendTask_T} & \hat{=} \text{SuspendTaskN_T} \\
& \vee \text{SuspendTaskS_T} \\
& \vee \text{SuspendTaskO_T}
\end{aligned}$$

ResumeTaskN_T
$\Delta Task$
$target? : TASK$
<hr/>
$\text{state}(target?) = suspended$
$\text{priority}(target?) \leq \text{priority}(running_task)$
$\exists TaskData$

$$state' = state \oplus \{(target? \mapsto ready)\}$$

$$\exists ContextData$$

$$\exists PrioData$$

$$ResumeTaskN_TFSSig$$

$$Task$$

$$target? : TASK$$

$$state(target?) = suspended$$

$$priority(target?) \leq priority(running_task)$$

theorem ResumeTaskN_T_vc_ref

$$\forall ResumeTaskN_TFSSig \mid true \bullet \text{pre } ResumeTaskN_T$$

$$ResumeTaskS_T$$

$$\Delta Task$$

$$target? : TASK$$

$$state(target?) = suspended$$

$$priority(target?) > priority(running_task)$$

$$\exists st? : STATE \mid st? = ready$$

- Reschedule[*tasks/tasks?*, *priority/pri?*]

$$ResumeTaskS_TFSSig$$

$$Task$$

$$target? : TASK$$

$$state(target?) = suspended$$

$$priority(target?) > priority(running_task)$$

theorem ResumeTaskS_T_vc_ref

$$\forall ResumeTaskS_TFSSig \mid true \bullet \text{pre } ResumeTaskS_T$$

$$ResumeTask_T \hat{=} ResumeTaskN_T \vee ResumeTaskS_T$$

$$ChangeTaskPriorityN_T$$

$$\Delta Task$$

$$newpri? : \mathbb{N}$$

$$target? : TASK$$

$$topReady! : TASK$$

$state(target?) = ready \Rightarrow newpri? \leq priority(running_task)$
 $state(target?) = running \Rightarrow$
 $(\forall t : state \sim (\{ready\}) \bullet newpri? \geq priority(t))$
 $state(target?) \neq nonexistent$
 $target? = idle \Rightarrow newpri? = 0$
 $\exists TaskData$
 $\exists StateData$
 $\exists ContextData$
 $priority' = priority \oplus \{(target? \mapsto newpri?)\}$
 $topReady! = running_task$

ChangeTaskPriorityN_TFSBSig

Task
 $newpri? : \mathbb{N}$
 $target? : TASK$

$state(target?) = ready \Rightarrow newpri? \leq priority(running_task)$
 $state(target?) = running \Rightarrow$
 $(\forall t : state \sim (\{ready\}) \bullet newpri? \geq priority(t))$
 $state(target?) \neq nonexistent$
 $target? = idle \Rightarrow newpri? = 0$

theorem *ChangeTaskPriorityN_T_vc_ref*

$\forall ChangeTaskPriorityN_TFSBSig \mid true \bullet pre\ ChangeTaskPriorityN_T$

ChangeTaskPriorityS_T

$\Delta Task$
 $target? : TASK$
 $newpri? : \mathbb{N}$
 $topReady! : TASK$

$state(target?) = ready$
 $newpri? > priority(running_task)$
 $target? = idle \Rightarrow newpri? = 0$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $\mid st? = ready$
 $\quad \wedge pri? = priority \oplus \{(target? \mapsto newpri?)\}$
 $\quad \bullet Reschedule[tasks/tasks?]$
 $topReady! = target?$

ChangeTaskPriorityS_TFSBSig

Task
 $newpri? : \mathbb{N}$

$target? : TASK$

$state(target?) = ready$
 $newpri? > priority(running_task)$
 $target? = idle \Rightarrow newpri? = 0$

theorem ChangeTaskPriorityS_T_vc_ref

$\forall ChangeTaskPriorityS_TFSBSig \mid true \bullet \text{pre } ChangeTaskPriorityS_T$

$ChangeTaskPriorityD_T$

$\Delta Task$

$target? : TASK$
 $topReady! : TASK$
 $newpri? : \mathbb{N}$

$state(target?) = running$
 $target? = idle \Rightarrow newpri? = 0$
 $state(topReady!) = ready$
 $\forall t : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(t)$
 $newpri? < priority(topReady!)$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $\mid st? = ready$
 $\wedge pri? = priority \oplus \{(target? \mapsto newpri?)\}$
 $\bullet Reschedule[tasks/tasks?, topReady!/target?]$

$ChangeTaskPriorityD_TFSBSig$

$Task$

$newpri? : \mathbb{N}$
 $target? : TASK$

$state(target?) = running$
 $target? = idle \Rightarrow newpri? = 0$
 $\exists readyTask : state \sim (\{ready\}) \bullet newpri? < priority(readyTask)$

theorem lChangeTaskPriorityD_T_Lemma

$\forall Task; target?, topReady! : TASK; newpri? : \mathbb{N}$

$\mid state(target?) = running$
 $\wedge (target? = idle \Rightarrow newpri? = 0)$
 $\wedge state(topReady!) = ready$
 $\wedge (\forall rtsk : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(rtsk))$
 $\wedge newpri? < priority(topReady!)$
 $\bullet \neg (Task[log_context := log_context \oplus$
 $\{(running_task, phys_context)\},$

$$\begin{aligned}
& \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \text{priority} := \text{priority} \oplus \{(target?, \text{newpri?})\}, \\
& \text{running_task} := \text{topReady!}, \\
& \text{state} := \text{state} \oplus \\
& \quad (\{(running_task, \text{ready})\} \cup \{(topReady!, \text{running})\})] \\
\wedge & (st \in TASK \\
& \quad \wedge \neg (\text{state} \oplus (\{(running_task, \text{ready})\} \cup \\
& \quad \quad \{(topReady!, \text{running})\}))st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(running_task, \text{ready})\} \cup \\
& \quad \quad \{(topReady!, \text{running})\})))st \in \text{transition}) \\
\Rightarrow & t \in TASK \\
& \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t)
\end{aligned}$$

theorem ChangeTaskPriorityD_T_vc_ref

$\forall \text{ChangeTaskPriorityD_TFSSig} \mid \text{true} \bullet \text{pre ChangeTaskPriorityD_T}$

$$\begin{aligned}
\text{ChangeTaskPriority_T} & \hat{=} \text{ChangeTaskPriorityN_T} \\
& \vee \text{ChangeTaskPriorityS_T} \\
& \vee \text{ChangeTaskPriorityD_T}
\end{aligned}$$

Appendix E

SPECIFICATION FOR QUEUE MODEL

[*QUEUE*]

QueueData

$queue : \mathbb{P} \text{ QUEUE}$
 $q_max : \text{ QUEUE} \rightarrow \mathbb{N}_1$
 $q_size : \text{ QUEUE} \rightarrow \mathbb{N}$

$\text{dom } q_max = \text{dom } q_size = queue$
 $\forall q : \text{ QUEUE} \mid q \in queue \bullet q_size(q) \leq q_max(q)$

Init_QueueData

QueueData'

$queue' = \emptyset$
 $q_max' = \emptyset$
 $q_size' = \emptyset$

theorem *QueueDataInit*

$\exists \text{ QueueData}' \bullet \text{ Init_QueueData}$

WaitingData

$wait_snd : \text{ TASK} \rightarrow \text{ QUEUE}$
 $wait_rcv : \text{ TASK} \rightarrow \text{ QUEUE}$

$\text{dom } wait_snd \cap \text{dom } wait_rcv = \emptyset$

Init_WaitingData

WaitingData'

$wait_snd' = \emptyset$

$wait_rcv' = \emptyset$

theorem *WaitingDataInit*

$\exists \textit{WaitingData}' \bullet \textit{Init_WaitingData}$

QReleasingData

$release_snd : \textit{TASK} \rightarrow \textit{QUEUE}$

$release_rcv : \textit{TASK} \rightarrow \textit{QUEUE}$

$\text{dom } release_snd \cap \text{dom } release_rcv = \emptyset$

Init_QReleasingData

QReleasingData'

$release_snd' = \emptyset$

$release_rcv' = \emptyset$

theorem *QReleasingDataInit*

$\exists \textit{QReleasingData}' \bullet \textit{Init_QReleasingData}$

Queue

QueueData

WaitingData

QReleasingData

$\text{ran } wait_snd \subseteq \textit{queue}$

$\text{ran } wait_rcv \subseteq \textit{queue}$

$\text{ran } release_snd \subseteq \textit{queue}$

$\text{ran } release_rcv \subseteq \textit{queue}$

$(\text{dom } wait_snd \cup \text{dom } wait_rcv)$

$\cap (\text{dom } release_snd \cup \text{dom } release_rcv) = \emptyset$

theorem *QueueProperty1*

$\forall \textit{Queue} \bullet \forall q : \textit{queue} \bullet q_max(q) > 0$

theorem *ImageProperty1*

$\forall f : \textit{TASK} \rightarrow \textit{QUEUE} \bullet \forall y : \textit{QUEUE} \bullet f^{-1}(\{y\}) \neq \emptyset \Leftrightarrow y \in \text{ran } f$

Init_Queue

Queue'

Init_QueueData

Init_WaitingData

Init_QReleasingData

theorem QueueInit

$\exists \text{Queue}' \bullet \text{Init_Queue}$

TaskQueue

Task

Queue

$\text{dom } \text{wait_snd} \subseteq \text{state} \sim (\{ \text{blocked} \})$

$\text{dom } \text{wait_rcv} \subseteq \text{state} \sim (\{ \text{blocked} \})$

Init_TaskQueue

TaskQueue'

Init_Task

Init_Queue

theorem TaskQueueInit

$\exists \text{TaskQueue}' \bullet \text{Init_TaskQueue}$

$\Delta \text{TaskQueue} \cong \text{TaskQueue} \wedge \text{TaskQueue}' \wedge \Delta \text{Task}$

ExtendTaskXi

$\Delta \text{TaskQueue}$

$\text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

$\exists \text{Queue}$

ExtTaskFSBSig

TaskQueue

$\text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

$\text{CreateTaskN_TQ} \cong \text{ExtendTaskXi} \wedge \text{CreateTaskN_T}$

$$\text{CreateTaskN_TQFSBSig} \cong \text{ExtTaskFSBSig} \wedge \text{CreateTaskN_TFSBSig}$$

theorem CreateTaskN_TQ_vc_ref

$$\forall \text{CreateTaskN_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{CreateTaskN_TQ}$$

$$\text{CreateTaskS_TQ} \cong \text{ExtendTaskXi} \wedge \text{CreateTaskS_T}$$

$$\text{CreateTaskS_TQFSBSig} \cong \text{ExtTaskFSBSig} \wedge \text{CreateTaskS_TFSBSig}$$

theorem CreateTaskS_TQ_vc_ref

$$\forall \text{CreateTaskS_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{CreateTaskS_TQ}$$

$$\text{CreateTask_TQ} \cong \text{CreateTaskN_TQ} \vee \text{CreateTaskS_TQ}$$

$$\text{DeleteTaskN_TQ}$$

$$\text{DeleteTaskN_T}$$

$$\Delta \text{TaskQueue}$$

$$\text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$$

$$\exists \text{QueueData}$$

$$\text{wait_snd}' = \{\text{target?}\} \triangleleft \text{wait_snd}$$

$$\text{wait_rcv}' = \{\text{target?}\} \triangleleft \text{wait_rcv}$$

$$\text{release_snd}' = \{\text{target?}\} \triangleleft \text{release_snd}$$

$$\text{release_rcv}' = \{\text{target?}\} \triangleleft \text{release_rcv}$$

$$\text{DeleteTaskN_TQFSBSig} \cong \text{ExtTaskFSBSig} \wedge \text{DeleteTaskN_TFSBSig}$$

theorem DeleteTaskN_TQ_vc_ref

$$\forall \text{DeleteTaskN_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{DeleteTaskN_TQ}$$

$$\text{DeleteTaskS_TQ} \cong \text{ExtendTaskXi} \wedge \text{DeleteTaskS_T}$$

$$\text{DeleteTaskS_TQFSBSig} \cong \text{ExtTaskFSBSig} \wedge \text{DeleteTaskS_TFSBSig}$$

theorem lDeleteTaskS_TQ_Lemma

$$\forall \text{TaskQueue}; \text{topReady!}, \text{target?} : \text{TASK}$$

$$\mid \text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$$

$$\wedge \text{target?} \in \text{tasks} \setminus \{\text{idle}\}$$

$$\wedge \text{state}(\text{target?}) \in \{\text{running}\}$$

$$\wedge \text{state}(\text{topReady!}) = \text{ready}$$

$$\begin{aligned}
& \wedge (\forall rtsk : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(rtsk)) \\
& \bullet \neg (TaskQueue[log_context := log_context \oplus \\
& \quad \{(target?, bare_context)\}, \\
& \quad phys_context := log_context(topReady!), \\
& \quad running_task := topReady!, \\
& \quad state := state \oplus \\
& \quad \quad (\{(target?, nonexistent)\} \cup \{(topReady!, running)\}), \\
& \quad tasks := tasks \setminus \{target?\}] \\
& \wedge (st \in TASK \\
& \quad \wedge \neg (state \oplus (\{(target?, nonexistent)\} \cup \\
& \quad \quad \{(topReady!, running)\})) st = state(st) \\
& \quad \Rightarrow (state(st), (state \oplus (\{(target?, nonexistent)\} \cup \\
& \quad \quad \{(topReady!, running)\})) st) \in transition) \\
& \Rightarrow t \in TASK \\
& \quad \wedge state(t) = ready \\
& \quad \wedge \neg priority(topReady!) \geq priority(t)
\end{aligned}$$

theorem DeleteTaskS_TQ_vc_ref

$$\forall DeleteTaskS_TQFSBSig \mid true \bullet \text{pre } DeleteTaskS_TQ$$

$$DeleteTask_TQ \cong DeleteTaskN_TQ \vee DeleteTaskS_TQ$$

$$ExecuteRunningTask_TQ \cong ExtendTaskXi \wedge ExecuteRunningTask_T$$

$$\begin{aligned}
ExecuteRunningTask_TQFSBSig & \cong ExtTaskFSBSig \\
& \wedge ExecuteRunningTask_TFBSBSig
\end{aligned}$$

theorem ExecuteRunningTask_TQ_vc_ref

$$\forall ExecuteRunningTask_TQFSBSig \mid true \bullet \text{pre } ExecuteRunningTask_TQ$$

SuspendTaskN_TQ

SuspendTaskN_T

$\Delta TaskQueue$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$\exists QueueData$

$wait_snd' = \{target?\} \triangleleft wait_snd$

$wait_rcv' = \{target?\} \triangleleft wait_rcv$

$\exists QReleasingData$

$$SuspendTaskN_TQFSBSig \cong ExtTaskFSBSig \wedge SuspendTaskN_TFBSBSig$$

theorem SuspendTaskN_TQ_vc_ref

$\forall \text{SuspendTaskN_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{SuspendTaskN_TQ}$

$\text{SuspendTaskS_TQ} \cong \text{ExtendTaskXi} \wedge \text{SuspendTaskS_T}$

$\text{SuspendTaskS_TQFSBSig} \cong \text{ExtTaskFSBSig} \wedge \text{SuspendTaskS_TFSBSig}$

theorem lSuspendTaskS_TQ_Lemma

$\forall \text{TaskQueue}; \text{target?}, \text{topReady!} : \text{TASK}$

$\mid \text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

$\wedge \text{target?} \in \text{tasks} \setminus \{\text{idle}\}$

$\wedge \text{state}(\text{target?}) \in \{\text{running}\}$

$\wedge \text{state}(\text{topReady!}) = \text{ready}$

$\wedge (\forall \text{rtsk} : \text{state} \sim (\{\text{ready}\}) \mid \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}))$

$\bullet \neg (\text{TaskQueue}[\text{log_context} := \text{log_context} \oplus$
 $\quad \{(\text{running_task}, \text{phys_context})\},$
 $\quad \text{phys_context} := \text{log_context}(\text{topReady!}),$
 $\quad \text{running_task} := \text{topReady!},$
 $\quad \text{state} := \text{state} \oplus (\{(\text{running_task}, \text{suspended})\} \cup$
 $\quad \{(\text{topReady!}, \text{running})\})]$

$\wedge (st \in \text{TASK}$

$\wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{suspended})\} \cup$

$\{(\text{topReady!}, \text{running})\}))st = \text{state}(st)$

$\Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{suspended})\} \cup$

$\{(\text{topReady!}, \text{running})\}))st) \in \text{transition})$

$\Rightarrow t \in \text{TASK}$

$\wedge \text{state}(t) = \text{ready}$

$\wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t)$

theorem SuspendTaskS_TQ_vc_ref

$\forall \text{SuspendTaskS_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{SuspendTaskS_TQ}$

$\text{SuspendTaskO_TQ} \cong \text{ExtendTaskXi} \wedge \text{SuspendTaskO_T}$

$\text{SuspendTaskO_TQFSBSig} \cong \text{ExtTaskFSBSig} \wedge \text{SuspendTaskO_TFSBSig}$

theorem SuspendTaskO_TQ_vc_ref

$\forall \text{SuspendTaskO_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{SuspendTaskO_TQ}$

$\text{SuspendTask_TQ} \cong \text{SuspendTaskN_TQ}$

$\vee \text{SuspendTaskS_TQ}$

$\vee \text{SuspendTaskO_TQ}$

$ResumeTaskN_TQ \cong ExtendTaskXi \wedge ResumeTaskN_T$

$ResumeTaskN_TQFSBSig \cong ExtTaskFSBSig \wedge ResumeTaskN_TFSBSig$

theorem ResumeTaskN_TQ_vc_ref

$\forall ResumeTaskN_TQFSBSig \mid true \bullet pre ResumeTaskN_TQ$

$ResumeTaskS_TQ \cong ExtendTaskXi \wedge ResumeTaskS_T$

$ResumeTaskS_TQFSBSig \cong ExtTaskFSBSig \wedge ResumeTaskS_TFSBSig$

theorem ResumeTaskS_TQ_vc_ref

$\forall ResumeTaskS_TQFSBSig \mid true \bullet pre ResumeTaskS_TQ$

$ResumeTask_TQ \cong ResumeTaskN_TQ \vee ResumeTaskS_TQ$

$ChangeTaskPriorityN_TQ \cong ExtendTaskXi \wedge ChangeTaskPriorityN_T$

$ChangeTaskPriorityN_TQFSBSig \cong ExtTaskFSBSig$
 $\wedge ChangeTaskPriorityN_TFSBSig$

theorem ChangeTaskPriorityN_TQ_vc_ref

$\forall ChangeTaskPriorityN_TQFSBSig \mid true$
 $\bullet pre ChangeTaskPriorityN_TQ$

$ChangeTaskPriorityS_TQ \cong ExtendTaskXi \wedge ChangeTaskPriorityS_T$

$ChangeTaskPriorityS_TQFSBSig \cong ExtTaskFSBSig$
 $\wedge ChangeTaskPriorityS_TFSBSig$

theorem ChangeTaskPriorityS_TQ_vc_ref

$\forall ChangeTaskPriorityS_TQFSBSig \mid true$
 $\bullet pre ChangeTaskPriorityS_TQ$

$ChangeTaskPriorityD_TQ \cong ExtendTaskXi \wedge ChangeTaskPriorityD_T$

$ChangeTaskPriorityD_TQFSBSig \cong ExtTaskFSBSig$
 $\wedge ChangeTaskPriorityD_TFSBSig$

theorem lChangeTaskPriorityD_TQ_Lemma
$$\begin{aligned}
& \forall \text{TaskQueue}; \text{target?}, \text{topReady!} : \text{TASK}; \text{newpri?} : \mathbb{N} \\
& \quad | \text{state}(\text{target?}) = \text{running} \\
& \quad \wedge (\text{target?} = \text{idle} \Rightarrow \text{newpri?} = 0) \\
& \quad \wedge \text{state}(\text{topReady!}) = \text{ready} \\
& \quad \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} |) \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\
& \quad \wedge \text{newpri?} < \text{priority}(\text{topReady!}) \\
& \quad \bullet \neg (\text{TaskQueue}[\text{log_context} := \text{log_context} \oplus \\
& \quad \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \quad \text{priority} := \text{priority} \oplus \{(\text{target?}, \text{newpri?})\}, \\
& \quad \quad \text{running_task} := \text{topReady!}, \\
& \quad \quad \text{state} := \text{state} \oplus \\
& \quad \quad \quad \{(\text{running_task}, \text{ready})\} \cup \{(\text{topReady!}, \text{running})\}]) \\
& \quad \wedge \text{newpri?} < \text{priority}(\text{topReady!}) \\
& \quad \wedge (\text{st} \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\})) \text{st} = \text{state}(\text{st}) \\
& \quad \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\})) \text{st}) \in \text{transition}) \\
& \quad \Rightarrow t \in \text{TASK} \\
& \quad \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t)
\end{aligned}$$
theorem ChangeTaskPriorityD_TQ_vc_ref
$$\begin{aligned}
& \forall \text{ChangeTaskPriorityD_TQFSBSig} \mid \text{true} \\
& \quad \bullet \text{pre } \text{ChangeTaskPriorityD_TQ}
\end{aligned}$$

CreateQueue_TQ

$$\begin{aligned}
& \Delta \text{TaskQueue} \\
& \text{que?} : \text{QUEUE} \\
& \text{size?} : \mathbb{N}
\end{aligned}$$

$$\begin{aligned}
& \text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv} \\
& \text{que?} \notin \text{queue} \\
& \text{size?} > 0 \\
& \exists \text{Task} \\
& \text{queue}' = \text{queue} \cup \{\text{que?}\} \\
& \text{q_max}' = \text{q_max} \oplus \{(\text{que?} \mapsto \text{size?})\} \\
& \text{q_size}' = \text{q_size} \oplus \{(\text{que?} \mapsto 0)\} \\
& \exists \text{WaitingData} \\
& \exists \text{QReleasingData}
\end{aligned}$$

CreateQueue_TQFSBSig

TaskQueue

que? : *QUEUE*

size? : \mathbb{N}

running_task \notin dom *release_snd* \cup dom *release_rcv*

que? \notin *queue*

size? > 0

theorem *CreateQueue_TQ_vc_ref*

\forall *CreateQueue_TQFSBSig* | *true* • pre *CreateQueue_TQ*

DeleteQueue_TQ

Δ *TaskQueue*

que? : *QUEUE*

running_task \notin dom *release_snd* \cup dom *release_rcv*

que? \in *queue*

que? \notin ran *wait_snd* \cup ran *wait_rcv*

que? \notin ran *release_snd* \cup ran *release_rcv*

\exists *Task*

queue' = *queue* \setminus {*que?*}

q_max' = {*que?*} \triangleleft *q_max*

q_size' = {*que?*} \triangleleft *q_size*

\exists *WaitingData*

\exists *QReleasingData*

DeleteQueue_TQFSBSig

TaskQueue

que? : *QUEUE*

running_task \notin dom *release_snd* \cup dom *release_rcv*

que? \in *queue*

que? \notin ran *wait_snd* \cup ran *wait_rcv*

que? \notin ran *release_snd* \cup ran *release_rcv*

theorem *DeleteQueue_TQ_vc_ref*

\forall *DeleteQueue_TQFSBSig* | *true* • pre *DeleteQueue_TQ*

QueueSendN_TQ

Δ *TaskQueue*

que? : *QUEUE*

topReady! : *TASK*

$running_task \notin \text{dom } release_rcv$
 $running_task \in \text{dom } release_snd \Rightarrow que? = release_snd(running_task)$
 $que? \in queue$
 $q_size(que?) < q_max(que?)$
 $que? \notin \text{ran } wait_rcv$
 $\exists Task$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto q_size(que?) + 1)\}$
 $\exists WaitingData$
 $release_snd' = \{running_task\} \triangleleft release_snd$
 $release_rcv' = release_rcv$
 $topReady! = running_task$

QueueSendN_TQFSBSig

TaskQueue
 $que? : QUEUE$

$running_task \notin \text{dom } release_rcv$
 $running_task \in \text{dom } release_snd \Rightarrow que? = release_snd(running_task)$
 $que? \in queue$
 $q_size(que?) < q_max(que?)$
 $que? \notin \text{ran } wait_rcv$

theorem *QueueSendN_TQ_vc_ref*

$\forall QueueSendN_TQFSBSig \mid true \bullet \text{pre } QueueSendN_TQ$

QueueSendF_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $topReady! : TASK$

$running_task \notin \text{dom } release_rcv$
 $running_task \in \text{dom } release_snd \Rightarrow que? = release_snd(running_task)$
 $que? \in queue$
 $q_size(que?) = q_max(que?)$
 $running_task \neq idle$
 $state(topReady!) = ready$
 $\forall t : state \sim (\{ready\} \mid) \bullet priority(topReady!) \geq priority(t)$
 $\exists st? : STATE \mid st? = blocked$

- *Reschedule*[$topReady! / target?, tasks / tasks?, priority / pri?$]

 $\exists QueueData$
 $wait_snd' = wait_snd \oplus \{(running_task \mapsto que?)\}$

$$\begin{aligned}
& \text{wait_rcv}' = \text{wait_rcv} \\
& \text{release_snd}' = \{\text{running_task}\} \triangleleft \text{release_snd} \\
& \text{release_rcv}' = \text{release_rcv}
\end{aligned}$$

QueueSendF_TQFSBSig

TaskQueue
que? : *QUEUE*

$$\begin{aligned}
& \text{running_task} \notin \text{dom release_rcv} \\
& \text{running_task} \in \text{dom release_snd} \Rightarrow \text{que?} = \text{release_snd}(\text{running_task}) \\
& \text{que?} \in \text{queue} \\
& q_size(\text{que?}) = q_max(\text{que?}) \\
& \text{running_task} \neq \text{idle}
\end{aligned}$$

theorem *lQueueSendF_TQ_Lemma*

$$\begin{aligned}
& \forall \text{TaskQueue}; \text{topReady!} : \text{TASK}; \text{que?} : \text{QUEUE} \\
& \quad | \text{running_task} \notin \text{dom release_rcv} \\
& \quad \wedge (\text{running_task} \in \text{dom release_snd} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_snd}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge q_size(\text{que?}) = q_max(\text{que?}) \\
& \quad \wedge \text{running_task} \neq \text{idle} \\
& \quad \wedge \text{state}(\text{topReady!}) = \text{ready} \\
& \quad \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \mid) \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\
& \quad \bullet \neg (\text{TaskQueue}[\text{log_context} := \text{log_context} \oplus \\
& \quad \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \quad \text{release_snd} := \{\text{running_task}\} \triangleleft \text{release_snd}, \\
& \quad \quad \text{running_task} := \text{topReady!}, \\
& \quad \quad \text{state} := \text{state} \oplus \\
& \quad \quad \quad \{(\text{running_task}, \text{blocked})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \quad \text{wait_snd} := \text{wait_snd} \oplus \{(\text{running_task}, \text{que?})\}] \\
& \quad \wedge (t \in \text{TASK} \wedge \text{state}(t) = \text{ready} \\
& \quad \quad \Rightarrow \text{priority}(\text{topReady!}) \geq \text{priority}(t)) \\
& \quad \Rightarrow st \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\}))st = \text{state}(st) \\
& \quad \quad \wedge \neg (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\})))st \in \text{transition})
\end{aligned}$$

theorem *QueueSendF_TQ_vc_ref*

$$\forall \text{QueueSendF_TQFSBSig} \mid \text{true} \bullet \text{pre } \text{QueueSendF_TQ}$$

QueueSendW_TQ

Δ *TaskQueue*

que? : *QUEUE*

topReady! : *TASK*

$running_task \notin \text{dom } release_rcv$

$running_task \in \text{dom } release_snd \Rightarrow que? = release_snd(running_task)$

$que? \in queue$

$q_size(que?) < q_max(que?)$

$topReady! \in wait_rcv \sim (\{ que? \} \mid)$

$\forall wr : wait_rcv \sim (\{ que? \} \mid) \bullet priority(topReady!) \geq priority(wr)$

$priority(topReady!) \leq priority(running_task)$

\exists *TaskData*

$state' = state \oplus \{(topReady! \mapsto ready)\}$

\exists *ContextData*

\exists *PrioData*

$queue' = queue$

$q_max' = q_max$

$q_size' = q_size \oplus \{(que? \mapsto q_size(que?) + 1)\}$

$wait_snd' = wait_snd$

$wait_rcv' = \{topReady!\} \triangleleft wait_rcv$

$release_snd' = \{running_task\} \triangleleft release_snd$

$release_rcv' = release_rcv \oplus \{(topReady! \mapsto que?)\}$

QueueSendW_TQFSBSig

TaskQueue

que? : *QUEUE*

$running_task \notin \text{dom } release_rcv$

$running_task \in \text{dom } release_snd \Rightarrow que? = release_snd(running_task)$

$que? \in queue$

$q_size(que?) < q_max(que?)$

$\forall wr : wait_rcv \sim (\{ que? \} \mid) \bullet priority(running_task) \geq priority(wr)$

theorem *lQueueSendW_TQ_Lemma*

\forall *TaskQueue*; *que?* : *QUEUE*; *topReady!* : *TASK*

$\mid running_task \notin \text{dom } release_rcv$

$\wedge (running_task \in \text{dom } release_snd$

$\Rightarrow que? = release_snd(running_task))$

$\wedge que? \in queue$

$\wedge q_size(que?) < q_max(que?)$

$\wedge topReady! \in wait_rcv \sim (\{ que? \} \mid)$

$\wedge (\forall wrct : wait_rcv \sim (\{ que? \} \mid)$

$\bullet priority(topReady!) \geq priority(wrct))$

$\wedge priority(running_task) \geq priority(topReady!)$

- $\neg (TaskQueue[q_size := q_size \oplus \{(que?, (1 + q_size(que?))\}],$
 $release_rcv := release_rcv \oplus \{(topReady!, que?)\},$
 $release_snd := \{running_task\} \triangleleft release_snd,$
 $state := state \oplus \{(topReady!, ready)\},$
 $wait_rcv := \{topReady!\} \triangleleft wait_rcv]$
 $\wedge priority(topReady!) \leq priority(running_task)$
 $\wedge (st \in TASK \wedge \neg (state \oplus \{(topReady!, ready)\})st = state(st)$
 $\Rightarrow (state(st), (state \oplus \{(topReady!, ready)\})st)$
 $\in transition)$
 $\Rightarrow wr \in \text{dom } wait_rcv$
 $\wedge wait_rcv(wr) = que?$
 $\wedge \neg priority(topReady!) \geq priority(wr))$

theorem QueueSendW_TQ_vc_ref

$\forall QueueSendW_TQFSBSig \mid true \bullet \text{pre } QueueSendW_TQ$

theorem TaskQueueProperty1

$\forall TaskQueue \bullet \forall t : TASK \mid t \in \text{dom } wait_rcv \bullet state(t) = blocked$

QueueSendWS_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $topReady! : TASK$

$running_task \notin \text{dom } release_rcv$
 $running_task \in \text{dom } release_snd \Rightarrow que? = release_snd(running_task)$
 $que? \in queue$
 $q_size(que?) < q_max(que?)$
 $topReady! \in wait_rcv \sim (\{que?\} \mid)$
 $\forall wr : wait_rcv \sim (\{que?\} \mid) \bullet priority(topReady!) \geq priority(wr)$
 $priority(topReady!) > priority(running_task)$
 $\exists st? : STATE \mid st? = ready$

- $Reschedule[topReady!/target?, tasks/tasks?, priority/pri?]$

 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto q_size(que?) + 1)\}$
 $wait_snd' = wait_snd$
 $wait_rcv' = \{topReady!\} \triangleleft wait_rcv$
 $release_snd' = \{running_task\} \triangleleft release_snd$
 $release_rcv' = release_rcv \oplus \{(topReady! \mapsto que?)\}$

QueueSendWS_TQFSBSig

$TaskQueue$
 $que? : QUEUE$

$$\begin{array}{l}
\text{running_task} \notin \text{dom release_rcv} \\
\text{running_task} \in \text{dom release_snd} \Rightarrow \text{que?} = \text{release_snd}(\text{running_task}) \\
\text{que?} \in \text{queue} \\
q_size(\text{que?}) < q_max(\text{que?}) \\
\exists \text{topReady!} : \text{wait_rcv} \sim (\{ \text{que?} \}) \\
\quad | \forall \text{wr} : \text{wait_rcv} \sim (\{ \text{que?} \}) \\
\quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wr}) \\
\quad \bullet \text{priority}(\text{topReady!}) > \text{priority}(\text{running_task})
\end{array}$$

theorem lQueueSendWS_TQ_Lemma

$$\begin{array}{l}
\forall \text{TaskQueue}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK} \\
\quad | \text{running_task} \notin \text{dom release_rcv} \\
\quad \wedge (\text{running_task} \in \text{dom release_snd} \\
\quad \quad \Rightarrow \text{que?} = \text{release_snd}(\text{running_task})) \\
\quad \wedge \text{que?} \in \text{queue} \\
\quad \wedge q_size(\text{que?}) < q_max(\text{que?}) \\
\quad \wedge \text{topReady!} \in \text{wait_rcv} \sim (\{ \text{que?} \}) \\
\quad \wedge (\forall \text{wrct} : \text{wait_rcv} \sim (\{ \text{que?} \}) \\
\quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wrct})) \\
\quad \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{running_task}) \\
\quad \bullet \neg (\text{TaskQueue}[\text{log_context} := \text{log_context} \oplus \\
\quad \quad \{(\text{running_task}, \text{phys_context})\}, \\
\quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
\quad \quad q_size := q_size \oplus \{(\text{wait_rcv}(\text{topReady!}), (1 + \\
\quad \quad \quad q_size(\text{wait_rcv}(\text{topReady!})))\}, \\
\quad \quad \text{release_rcv} := \text{release_rcv} \oplus \\
\quad \quad \quad \{(\text{topReady!}, \text{wait_rcv}(\text{topReady!}))\}, \\
\quad \quad \text{release_snd} := \{\text{running_task}\} \triangleleft \text{release_snd}, \\
\quad \quad \text{running_task} := \text{topReady!}, \\
\quad \quad \text{state} := \text{state} \oplus \\
\quad \quad \quad (\{(\text{running_task}, \text{ready})\} \cup \{(\text{topReady!}, \text{running})\}), \\
\quad \quad \text{wait_rcv} := \{\text{topReady!}\} \triangleleft \text{wait_rcv}] \\
\quad \wedge (\text{wr} \in \text{dom wait_rcv} \wedge \text{wait_rcv}(\text{wr}) = \text{wait_rcv}(\text{topReady!}) \\
\quad \quad \Rightarrow \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wr})) \\
\quad \Rightarrow \text{st} \in \text{TASK} \\
\quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
\quad \quad \quad \{(\text{topReady!}, \text{running})\})) \text{st} = \text{state}(\text{st}) \\
\quad \quad \wedge \neg (\text{state}(\text{st}), (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
\quad \quad \quad \{(\text{topReady!}, \text{running})\}))) \text{st} \in \text{transition})
\end{array}$$

theorem QueueSendWS_TQ_vc_ref

$$\forall \text{QueueSendWS_TQFSBSig} \mid \text{true} \bullet \text{pre QueueSendWS_TQ}$$

$$\begin{aligned} \text{QueueSend_TQ} &\triangleq \text{QueueSendN_TQ} \\ &\vee \text{QueueSendF_TQ} \\ &\vee \text{QueueSendW_TQ} \\ &\vee \text{QueueSendWS_TQ} \end{aligned}$$

QueueReceiveN_TQ

$\Delta \text{TaskQueue}$
 $que? : \text{QUEUE}$
 $topReady! : \text{TASK}$

$running_task \notin \text{dom } release_snd$
 $running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task)$
 $que? \in queue$
 $q_size(que?) > 0$
 $que? \notin \text{ran } wait_snd$
 $\exists \text{Task}$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto q_size(que?) - 1)\}$
 $\exists \text{WaitingData}$
 $release_snd' = release_snd$
 $release_rcv' = \{running_task\} \triangleleft release_rcv$
 $topReady! = running_task$

QueueReceiveN_TQFSBSig

TaskQueue
 $que? : \text{QUEUE}$

$running_task \notin \text{dom } release_snd$
 $running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task)$
 $que? \in queue$
 $q_size(que?) > 0$
 $que? \notin \text{ran } wait_snd$

theorem *QueueReceiveN_TQ_vc_ref*

$\forall \text{QueueReceiveN_TQFSBSig} \mid true \bullet \text{pre } \text{QueueReceiveN_TQ}$

QueueReceiveE_TQ

$\Delta \text{TaskQueue}$
 $que? : \text{QUEUE}$
 $topReady! : \text{TASK}$

$running_task \notin \text{dom } release_snd$
 $running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task)$

$$\begin{aligned}
& que? \in queue \\
& q_size(que?) = 0 \\
& running_task \neq idle \\
& state(topReady!) = ready \\
& \forall t : state^{\sim}(\{ready\}) \bullet priority(topReady!) \geq priority(t) \\
& \exists st? : STATE \mid st? = blocked \\
& \quad \bullet Reschedule[topReady!/target?, tasks/tasks?, priority/pri?] \\
& \exists QueueData \\
& wait_snd' = wait_snd \\
& wait_rcv' = wait_rcv \oplus \{(running_task \mapsto que?)\} \\
& release_snd' = release_snd \\
& release_rcv' = \{running_task\} \triangleleft release_rcv
\end{aligned}$$

QueueReceiveE_TQFSBSig

TaskQueue
que? : QUEUE

$$\begin{aligned}
& running_task \notin \text{dom } release_snd \\
& running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task) \\
& que? \in queue \\
& q_size(que?) = 0 \\
& running_task \neq idle
\end{aligned}$$

theorem *lQueueReceiveE_TQ_Lemma*

$$\begin{aligned}
& \forall TaskQueue; que? : QUEUE; topReady! : TASK \\
& \quad \mid running_task \notin \text{dom } release_snd \\
& \quad \wedge (running_task \in \text{dom } release_rcv \\
& \quad \quad \Rightarrow que? = release_rcv(running_task)) \\
& \quad \wedge que? \in queue \\
& \quad \wedge q_size(que?) = 0 \\
& \quad \wedge running_task \neq idle \\
& \quad \wedge topReady! \in state^{\sim}(\{ready\}) \\
& \quad \wedge (\forall rtsk : state^{\sim}(\{ready\}) \\
& \quad \quad \bullet priority(topReady!) \geq priority(rtsk)) \\
& \quad \bullet \neg (TaskQueue[log_context := log_context \oplus \\
& \quad \quad \quad \{(running_task, phys_context)\}, \\
& \quad \quad \quad phys_context := log_context(topReady!), \\
& \quad \quad \quad release_rcv := \{running_task\} \triangleleft release_rcv, \\
& \quad \quad \quad running_task := topReady!, \\
& \quad \quad \quad state := state \oplus \\
& \quad \quad \quad \quad \{(running_task, blocked)\} \cup \{(topReady!, running)\}], \\
& \quad \quad \quad wait_rcv := wait_rcv \oplus \{(running_task, que?)\}) \\
& \quad \wedge (st \in TASK \\
& \quad \quad \wedge \neg (state \oplus (\{(running_task, blocked)\} \cup \\
& \quad \quad \quad \{(topReady!, running)\}))) st = state(st)
\end{aligned}$$

$$\begin{aligned}
&\Rightarrow (state(st), (state \oplus (\{(running_task, blocked)\} \cup \\
&\quad \{(topReady!, running)\}))st) \in transition) \\
&\Rightarrow t \in TASK \\
&\quad \wedge state(t) = ready \\
&\quad \wedge \neg priority(topReady!) \geq priority(t)
\end{aligned}$$

theorem QueueReceiveE_TQ_vc_ref

$$\forall QueueReceiveE_TQFSBSig \mid true \bullet \text{pre } QueueReceiveE_TQ$$

QueueReceiveW_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $topReady! : TASK$

$running_task \notin \text{dom } release_snd$
 $running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task)$
 $que? \in queue$
 $q_size(que?) > 0$
 $topReady! \in wait_snd \sim (\{que?\})$
 $\forall ws : wait_snd \sim (\{que?\}) \bullet priority(topReady!) \geq priority(ws)$
 $priority(topReady!) \leq priority(running_task)$
 $\exists TaskData$
 $state' = state \oplus \{(topReady! \mapsto ready)\}$
 $\exists ContextData$
 $\exists PrioData$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto q_size(que?) - 1)\}$
 $wait_snd' = \{topReady!\} \triangleleft wait_snd$
 $wait_rcv' = wait_rcv$
 $release_snd' = release_snd \oplus \{(topReady! \mapsto que?)\}$
 $release_rcv' = \{running_task\} \triangleleft release_rcv$

QueueReceiveW_TQFSBSig

$TaskQueue$
 $que? : QUEUE$

$running_task \notin \text{dom } release_snd$
 $running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task)$
 $que? \in queue$
 $q_size(que?) > 0$
 $\forall ws : wait_snd \sim (\{que?\}) \bullet priority(running_task) \geq priority(ws)$

theorem lQueueReceiveW_TQ_Lemma
$$\begin{aligned}
& \forall \text{TaskQueue}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \\
& \quad \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge \text{q_size}(\text{que?}) > 0 \\
& \quad \wedge \text{topReady!} \in \text{wait_snd} \sim (\{ \text{que?} \} \downarrow) \\
& \quad \wedge (\forall \text{wsnt} : \text{wait_snd} \sim (\{ \text{que?} \} \downarrow) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wsnt})) \\
& \quad \wedge \text{priority}(\text{running_task}) \geq \text{priority}(\text{topReady!}) \\
& \quad \bullet \neg (\text{TaskQueue}[\text{q_size} := \text{q_size} \oplus \{(\text{que?}, (\text{q_size}(\text{que?}) - 1))\}, \\
& \quad \quad \text{release_rcv} := \{ \text{running_task} \} \triangleleft \text{release_rcv}, \\
& \quad \quad \text{release_snd} := \text{release_snd} \oplus \{(\text{topReady!}, \text{que?})\}, \\
& \quad \quad \text{state} := \text{state} \oplus \{(\text{topReady!}, \text{ready})\}, \\
& \quad \quad \text{wait_snd} := \{ \text{topReady!} \} \triangleleft \text{wait_snd}] \\
& \quad \wedge \text{priority}(\text{topReady!}) \leq \text{priority}(\text{running_task}) \\
& \quad \wedge (\text{st} \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus \{(\text{topReady!}, \text{ready})\}) \text{st} = \text{state}(\text{st}) \\
& \quad \quad \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus \{(\text{topReady!}, \text{ready})\}) \text{st}) \\
& \quad \quad \quad \in \text{transition}) \\
& \quad \Rightarrow \text{ws} \in \text{dom wait_snd} \\
& \quad \quad \wedge \text{wait_snd}(\text{ws}) = \text{que?} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{ws})
\end{aligned}$$
theorem QueueReceiveW_TQ_vc_ref
$$\forall \text{QueueReceiveW_TQFSBSig} \mid \text{true} \bullet \text{pre QueueReceiveW_TQ}$$

$$\text{QueueReceiveWS_TQ}$$

$$\Delta \text{TaskQueue}$$

$$\text{que?} : \text{QUEUE}$$

$$\text{topReady!} : \text{TASK}$$

$$\text{running_task} \notin \text{dom release_snd}$$

$$\text{running_task} \in \text{dom release_rcv} \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})$$

$$\text{que?} \in \text{queue}$$

$$\text{q_size}(\text{que?}) > 0$$

$$\text{topReady!} \in \text{wait_snd} \sim (\{ \text{que?} \} \downarrow)$$

$$\forall \text{ws} : \text{wait_snd} \sim (\{ \text{que?} \} \downarrow) \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{ws})$$

$$\text{priority}(\text{topReady!}) > \text{priority}(\text{running_task})$$

$$\exists \text{st?} : \text{STATE} \mid \text{st?} = \text{ready}$$

$$\bullet \text{Reschedule}[\text{topReady!}/\text{target?}, \text{tasks}/\text{tasks?}, \text{priority}/\text{pri?}]$$

$$\text{queue}' = \text{queue}$$

$$\text{q_max}' = \text{q_max}$$

$$\text{q_size}' = \text{q_size} \oplus \{(\text{que?} \mapsto \text{q_size}(\text{que?}) - 1)\}$$

$$\begin{aligned}
wait_snd' &= \{topReady!\} \triangleleft wait_snd \\
wait_rcv' &= wait_rcv \\
release_snd' &= release_snd \oplus \{(topReady! \mapsto que?)\} \\
release_rcv' &= \{running_task\} \triangleleft release_rcv
\end{aligned}$$

QueueReceiveWS_TQFSBSig

TaskQueue
 $que? : QUEUE$

$$\begin{aligned}
&running_task \notin \text{dom } release_snd \\
&running_task \in \text{dom } release_rcv \Rightarrow que? = release_rcv(running_task) \\
&que? \in queue \\
&q_size(que?) > 0 \\
&\exists topReady! : wait_snd \sim (\{que?\}) \\
&\quad | \forall ws : wait_snd \sim (\{que?\}) \\
&\quad \bullet priority(topReady!) \geq priority(ws) \\
&\quad \bullet priority(topReady!) > priority(running_task)
\end{aligned}$$

theorem *lQueueReceiveWS_TQ_Lemma*

$$\begin{aligned}
&\forall TaskQueue; que? : QUEUE; topReady! : TASK \\
&\quad | running_task \notin \text{dom } release_snd \\
&\quad \wedge (running_task \in \text{dom } release_rcv \\
&\quad \quad \Rightarrow que? = release_rcv(running_task)) \\
&\quad \wedge que? \in queue \\
&\quad \wedge q_size(que?) > 0 \\
&\quad \wedge topReady! \in wait_snd \sim (\{que?\}) \\
&\quad \wedge (\forall wsnt : wait_snd \sim (\{que?\}) \\
&\quad \quad \bullet priority(topReady!) \geq priority(wsnt)) \\
&\quad \wedge priority(topReady!) > priority(running_task) \\
&\quad \bullet \neg (TaskQueue[log_context := log_context \oplus \\
&\quad \quad \{(running_task, phys_context)\}, \\
&\quad \quad phys_context := log_context(topReady!), \\
&\quad \quad q_size := q_size \oplus \{(wait_snd(topReady!), \\
&\quad \quad (q_size(wait_snd(topReady!)) - 1)\}, \\
&\quad \quad release_rcv := \{running_task\} \triangleleft release_rcv, \\
&\quad \quad release_snd := release_snd \oplus \\
&\quad \quad \{(topReady!, wait_snd(topReady!))\}, \\
&\quad \quad running_task := topReady!, \\
&\quad \quad state := state \oplus \\
&\quad \quad \{(running_task, ready)\} \cup \{(topReady!, running)\}), \\
&\quad \quad wait_snd := \{topReady!\} \triangleleft wait_snd] \\
&\quad \wedge (st \in TASK \\
&\quad \quad \wedge \neg (state \oplus (\{(running_task, ready)\} \cup \\
&\quad \quad \{(topReady!, running)\})) st = state(st) \\
&\quad \quad \Rightarrow (state(st), (state \oplus (\{(running_task, ready)\} \cup
\end{aligned}$$

$$\begin{aligned}
& \{(topReady!, running)\})st) \in transition) \\
\Rightarrow & ws \in \text{dom } wait_snd \\
& \wedge wait_snd(ws) = wait_snd(topReady!) \\
& \wedge \neg priority(topReady!) \geq priority(ws)
\end{aligned}$$

theorem QueueReceiveWS_TQ_vc_ref

$$\forall QueueReceiveWS_TQFSBSig \mid true \bullet \text{pre } QueueReceiveWS_TQ$$

$$\begin{aligned}
QueueReceive_TQ & \cong QueueReceiveN_TQ \\
& \vee QueueReceiveE_TQ \\
& \vee QueueReceiveW_TQ \\
& \vee QueueReceiveWS_TQ
\end{aligned}$$

Appendix F

SPECIFICATION FOR TIME MODEL

$slice_delay : \mathbb{N}$

$\langle\langle \text{disabled } slice_delay_def \rangle\rangle$

$slice_delay = 1$

$Time$

$clock : \mathbb{N}$

$delayed_task : \mathbb{P} \text{ TASK}$

$wait_time : \text{ TASK} \rightarrow \mathbb{N}$

$time_slice : \mathbb{N}$

$\forall t : \text{ dom } wait_time \bullet wait_time(t) \geq clock$

$Init_Time$

$Time'$

$clock' = 0$

$delayed_task' = \emptyset$

$wait_time' = \emptyset$

$time_slice' = slice_delay$

theorem TimeInit

$\exists Time' \bullet Init_Time$

$TaskQueueTime$

$TaskQueue$

$Time$

$$\langle \text{delayed_task}, \text{dom wait_snd}, \text{dom wait_rcv} \rangle \text{ partition dom wait_time}$$

$$\text{delayed_task} \subseteq \text{state}^{\sim}(\{ \text{blocked} \})$$

theorem rule domTime

$$\forall \text{TaskQueueTime}; t : \text{TASK} \mid t \in \text{dom wait_time}$$

- $t \in \text{state}^{\sim}(\{ \text{blocked} \})$

$$\text{Init_TaskQueueTime}$$

$$\text{TaskQueueTime}'$$

$$\text{Init_TaskQueue}$$

$$\text{Init_Time}$$

theorem TaskQueueTimeInit

$$\exists \text{TaskQueueTime}' \bullet \text{Init_TaskQueueTime}$$

$$\Delta \text{TaskQueueTime} \cong \text{TaskQueueTime} \wedge \text{TaskQueueTime}' \wedge \Delta \text{Task}$$

$$\text{ExtendTaskQueueXi}$$

$$\Delta \text{TaskQueueTime}$$

$$\exists \text{Time}$$

$$\text{CreateTaskN_TQT} \cong \text{ExtendTaskQueueXi} \wedge \text{CreateTaskN_TQ}$$

$$\text{CreateTaskN_TQTFSSig} \cong \text{TaskQueueTime} \wedge \text{CreateTaskN_TQTFSSig}$$

theorem CreateTaskN_TQT_vc_ref

$$\forall \text{CreateTaskN_TQTFSSig} \mid \text{true} \bullet \text{pre CreateTaskN_TQT}$$

$$\text{CreateTaskS_TQT} \cong \text{ExtendTaskQueueXi} \wedge \text{CreateTaskS_TQ}$$

$$\text{CreateTaskS_TQTFSSig} \cong \text{TaskQueueTime} \wedge \text{CreateTaskS_TQTFSSig}$$

theorem CreateTaskS_TQT_vc_ref

$$\forall \text{CreateTaskS_TQTFSSig} \mid \text{true} \bullet \text{pre CreateTaskS_TQT}$$

$$\text{CreateTask_TQT} \cong \text{CreateTaskN_TQT} \vee \text{CreateTaskS_TQT}$$

DeleteTaskN_TQT

DeleteTaskN_TQ

Δ *TaskQueueTime*

$clock' = clock$

$delayed_task' = delayed_task \setminus \{target?\}$

$wait_time' = \{target?\} \triangleleft wait_time$

$time_slice' = time_slice$

$DeleteTaskN_TQTFSSig \cong TaskQueueTime \wedge DeleteTaskN_TQFSBSig$

theorem *DeleteTaskN_TQT_vc_ref*

$\forall DeleteTaskN_TQTFSSig \mid true \bullet \text{pre } DeleteTaskN_TQT$

$DeleteTaskS_TQT \cong ExtendTaskQueueXi \wedge DeleteTaskS_TQ$

$DeleteTaskS_TQTFSSig \cong TaskQueueTime \wedge DeleteTaskS_TQFSBSig$

theorem *IDeleteTaskS_TQT_Lemma*

$\forall TaskQueueTime; topReady!, target? : TASK$

$\mid running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$\wedge target? \in tasks \setminus \{idle\}$

$\wedge state(target?) \in \{running\}$

$\wedge state(topReady!) = ready$

$\wedge (\forall rtsk : state \sim (\{ready\})$

$\bullet \text{priority}(topReady!) \geq \text{priority}(rtsk))$

$\bullet \neg (TaskQueueTime[log_context := log_context \oplus$

$\{(target?, bare_context)\},$

$phys_context := log_context(topReady!),$

$running_task := topReady!,$

$state := state \oplus$

$(\{(target?, nonexistent)\} \cup \{(topReady!, running)\}),$

$tasks := tasks \setminus \{target?\}]$

$\wedge (st \in TASK$

$\wedge \neg (state \oplus (\{(target?, nonexistent)\} \cup$

$\{(topReady!, running)\}))st = state(st)$

$\Rightarrow (state(st), (state \oplus (\{(target?, nonexistent)\} \cup$

$\{(topReady!, running)\}))st \in transition)$

$\Rightarrow t \in TASK$

$\wedge state(t) = ready$

$\wedge \neg \text{priority}(topReady!) \geq \text{priority}(t))$

theorem DeleteTaskS_TQT_vc_ref

\forall DeleteTaskS_TQTFSSig | true • pre DeleteTaskS_TQT

$DeleteTask_TQT \cong DeleteTaskN_TQT \vee DeleteTaskS_TQT$

$ExecuteRunningTask_TQT \cong ExtendTaskQueueXi$
 $\wedge ExecuteRunningTask_TQ$

$ExecuteRunningTask_TQTFSSig \cong TaskQueueTime$
 $\wedge ExecuteRunningTask_TQFSSig$

theorem ExecuteRunningTask_TQT_vc_ref

$\forall ExecuteRunningTask_TQTFSSig$ | true
• pre ExecuteRunningTask_TQT

$SuspendTaskN_TQT$

$SuspendTaskN_TQ$

$\Delta TaskQueueTime$

$clock' = clock$

$delayed_task' = delayed_task \setminus \{target?\}$

$wait_time' = \{target?\} \triangleleft wait_time$

$time_slice' = time_slice$

$SuspendTaskN_TQTFSSig \cong TaskQueueTime$
 $\wedge SuspendTaskN_TQFSSig$

theorem SuspendTaskN_TQT_vc_ref

$\forall SuspendTaskN_TQTFSSig$ | true • pre SuspendTaskN_TQT

$SuspendTaskS_TQT \cong ExtendTaskQueueXi \wedge SuspendTaskS_TQ$

$SuspendTaskS_TQTFSSig \cong TaskQueueTime$
 $\wedge SuspendTaskS_TQFSSig$

theorem lSuspendTaskS_TQT_Lemma

$\forall TaskQueueTime; target?, topReady! : TASK$
| $running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $\wedge target? \in tasks \setminus \{idle\}$
 $\wedge state(target?) \in \{running\}$
 $\wedge state(topReady!) = ready$

$$\begin{aligned}
& \wedge (\forall rtsk : state \sim (\{ready\} \mid)) \\
& \bullet \text{priority}(topReady!) \geq \text{priority}(rtsk) \\
& \bullet \neg (TaskQueueTime[log_context := log_context \oplus \\
& \quad \{(running_task, phys_context)\}, \\
& \quad phys_context := log_context(topReady!), \\
& \quad running_task := topReady!, \\
& \quad state := state \oplus (\{(running_task, suspended)\} \cup \\
& \quad \{(topReady!, running)\})]) \\
& \wedge (st \in TASK \\
& \quad \wedge \neg (state \oplus (\{(running_task, suspended)\} \cup \\
& \quad \{(topReady!, running)\}))st = state(st) \\
& \quad \Rightarrow (state(st), (state \oplus (\{(running_task, suspended)\} \cup \\
& \quad \{(topReady!, running)\}))st) \in transition) \\
& \Rightarrow t \in TASK \\
& \quad \wedge state(t) = ready \\
& \quad \wedge \neg \text{priority}(topReady!) \geq \text{priority}(t)
\end{aligned}$$

theorem SuspendTaskS_TQT_vc_ref

$$\forall SuspendTaskS_TQTFSSig \mid true \bullet \text{pre } SuspendTaskS_TQT$$

$$SuspendTaskO_TQT \cong ExtendTaskQueueXi \wedge SuspendTaskO_TQ$$

$$\begin{aligned}
SuspendTaskO_TQTFSSig & \cong TaskQueueTime \\
& \wedge SuspendTaskO_TQFSBSig
\end{aligned}$$

theorem SuspendTaskO_TQT_vc_ref

$$\forall SuspendTaskO_TQTFSSig \mid true \bullet \text{pre } SuspendTaskO_TQT$$

$$\begin{aligned}
SuspendTask_TQT & \cong SuspendTaskN_TQT \\
& \vee SuspendTaskS_TQT \\
& \vee SuspendTaskO_TQT
\end{aligned}$$

$$ResumeTaskN_TQT \cong ExtendTaskQueueXi \wedge ResumeTaskN_TQ$$

$$\begin{aligned}
ResumeTaskN_TQTFSSig & \cong TaskQueueTime \\
& \wedge ResumeTaskN_TQFSBSig
\end{aligned}$$

theorem ResumeTaskN_TQT_vc_ref

$$\forall ResumeTaskN_TQTFSSig \mid true \bullet \text{pre } ResumeTaskN_TQT$$

$$ResumeTaskS_TQT \cong ExtendTaskQueueXi \wedge ResumeTaskS_TQ$$

$$\begin{aligned} \text{ResumeTaskS_TQTFSSig} &\hat{=} \text{TaskQueueTime} \\ &\wedge \text{ResumeTaskS_TQFSBSig} \end{aligned}$$

theorem ResumeTaskS_TQT_vc_ref

$$\forall \text{ResumeTaskS_TQTFSSig} \mid \text{true} \bullet \text{pre ResumeTaskS_TQT}$$

$$\text{ResumeTask_TQT} \hat{=} \text{ResumeTaskN_TQT} \vee \text{ResumeTaskS_TQT}$$

$$\begin{aligned} \text{ChangeTaskPriorityN_TQT} &\hat{=} \text{ExtendTaskQueueXi} \\ &\wedge \text{ChangeTaskPriorityN_TQ} \end{aligned}$$

$$\begin{aligned} \text{ChangeTaskPriorityN_TQTFSSig} &\hat{=} \text{TaskQueueTime} \\ &\wedge \text{ChangeTaskPriorityN_TQFSBSig} \end{aligned}$$

theorem ChangeTaskPriorityN_TQT_vc_ref

$$\begin{aligned} &\forall \text{ChangeTaskPriorityN_TQTFSSig} \mid \text{true} \\ &\bullet \text{pre ChangeTaskPriorityN_TQT} \end{aligned}$$

$$\begin{aligned} \text{ChangeTaskPriorityS_TQT} &\hat{=} \text{ExtendTaskQueueXi} \\ &\wedge \text{ChangeTaskPriorityS_TQ} \end{aligned}$$

$$\begin{aligned} \text{ChangeTaskPriorityS_TQTFSSig} &\hat{=} \text{TaskQueueTime} \\ &\wedge \text{ChangeTaskPriorityS_TQFSBSig} \end{aligned}$$

theorem ChangeTaskPriorityS_TQT_vc_ref

$$\begin{aligned} &\forall \text{ChangeTaskPriorityS_TQTFSSig} \mid \text{true} \\ &\bullet \text{pre ChangeTaskPriorityS_TQT} \end{aligned}$$

$$\begin{aligned} \text{ChangeTaskPriorityD_TQT} &\hat{=} \text{ExtendTaskQueueXi} \\ &\wedge \text{ChangeTaskPriorityD_TQ} \end{aligned}$$

$$\begin{aligned} \text{ChangeTaskPriorityD_TQTFSSig} &\hat{=} \text{TaskQueueTime} \\ &\wedge \text{ChangeTaskPriorityD_TQFSBSig} \end{aligned}$$

theorem lChangeTaskPriorityD_TQT_Lemma

$$\begin{aligned} &\forall \text{TaskQueueTime}; \text{target?}, \text{topReady!} : \text{TASK}; \text{newpri?} : \mathbb{N} \\ &\mid \text{state}(\text{target?}) = \text{running} \\ &\wedge (\text{target?} = \text{idle} \Rightarrow \text{newpri?} = 0) \\ &\wedge \text{state}(\text{topReady!}) = \text{ready} \\ &\wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \})) \end{aligned}$$

- $priority(topReady!) \geq priority(rtsk)$

$$\wedge newpri? < priority(topReady!)$$

- $\neg (TaskQueueTime[log_context := log_context \oplus$
 $\{(running_task, phys_context)\},$
 $phys_context := log_context(topReady!),$
 $priority := priority \oplus \{(target?, newpri?)\},$
 $running_task := topReady!,$
 $state := state \oplus$
 $\{(running_task, ready)\} \cup \{(topReady!, running)\}])$
 $\wedge newpri? < priority(topReady!)$
 $\wedge (st \in TASK$
 $\wedge \neg (state \oplus (\{(running_task, ready)\} \cup$
 $\{(topReady!, running)\}))st = state(st)$
 $\Rightarrow (state(st), (state \oplus (\{(running_task, ready)\} \cup$
 $\{(topReady!, running)\}))st) \in transition)$
 $\Rightarrow t \in TASK$
 $\wedge state(t) = ready$
 $\wedge \neg priority(topReady!) \geq priority(t))$

theorem ChangeTaskPriorityD_TQT_vc_ref

$\forall ChangeTaskPriorityD_TQTFSSig \mid true$

- pre $ChangeTaskPriorityD_TQT$

$CreateQueue_TQT \hat{=} ExtendTaskQueueXi \wedge CreateQueue_TQ$

$CreateQueue_TQTFSSig \hat{=} TaskQueueTime \wedge CreateQueue_TQTFSSig$

theorem CreateQueue_TQT_vc_ref

$\forall CreateQueue_TQTFSSig \mid true$ • pre $CreateQueue_TQT$

$DeleteQueue_TQT \hat{=} ExtendTaskQueueXi \wedge DeleteQueue_TQ$

$DeleteQueue_TQTFSSig \hat{=} TaskQueueTime \wedge DeleteQueue_TQTFSSig$

theorem DeleteQueue_TQT_vc_ref

$\forall DeleteQueue_TQTFSSig \mid true$ • pre $DeleteQueue_TQT$

$QueueSendN_TQT \hat{=} ExtendTaskQueueXi \wedge QueueSendN_TQ$

$QueueSendN_TQTFSSig \hat{=} TaskQueueTime \wedge QueueSendN_TQTFSSig$

theorem QueueSendN_TQT_vc_ref
$$\forall \text{QueueSendN_TQTFSSig} \mid \text{true} \bullet \text{pre } \text{QueueSendN_TQT}$$

QueueSendF_TQT
$\Delta \text{TaskQueueTime}$
QueueSendF_TQ
$wtime? : \mathbb{N}$
<hr/>
$wtime? > \text{clock}$
$\text{clock}' = \text{clock}$
$\text{delayed_task}' = \text{delayed_task}$
$\text{wait_time}' = \text{wait_time} \oplus \{(running_task \mapsto wtime?)\}$
$\text{time_slice}' = \text{time_slice}$

$\text{QueueSendF_TQTFSSig}$
TaskQueueTime
$\text{QueueSendF_TQTFSSig}$
$wtime? : \mathbb{N}$
<hr/>
$wtime? > \text{clock}$

theorem lQueueSendF_TQT_Lemma
$$\forall \text{TaskQueueTime}; \text{topReady!} : \text{TASK}; \text{que?} : \text{QUEUE}; wtime? : \mathbb{N}$$

$$\mid \text{running_task} \notin \text{dom } \text{release_rcv}$$

$$\wedge (\text{running_task} \in \text{dom } \text{release_snd} \\ \Rightarrow \text{que?} = \text{release_snd}(\text{running_task}))$$

$$\wedge \text{que?} \in \text{queue}$$

$$\wedge q_size(\text{que?}) = q_max(\text{que?})$$

$$\wedge \text{running_task} \neq \text{idle}$$

$$\wedge \text{state}(\text{topReady!}) = \text{ready}$$

$$\wedge (\forall \text{rtsk} : \text{state} \sim \{ \text{ready} \} \mid$$

$$\bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}))$$

$$\wedge wtime? > \text{clock}$$

$$\bullet \neg (\text{TaskQueueTime}[\text{log_context} := \text{log_context} \oplus \\ \{(running_task, \text{phys_context})\},$$

$$\text{phys_context} := \text{log_context}(\text{topReady!}),$$

$$\text{release_snd} := \{ \text{running_task} \} \triangleleft \text{release_snd},$$

$$\text{running_task} := \text{topReady!},$$

$$\text{state} := \text{state} \oplus$$

$$\{(\text{running_task}, \text{blocked} \} \cup \{ (\text{topReady!}, \text{running}) \}),$$

$$\text{wait_time} := \text{wait_time} \oplus \{(\text{running_task}, wtime? \}),$$

$$\text{wait_snd} := \text{wait_snd} \oplus \{(\text{running_task}, \text{que?} \})]$$

$$\wedge (st \in \text{TASK}$$

$$\wedge \neg (\text{state} \oplus \{(\text{running_task}, \text{blocked} \}) \cup$$

$$\begin{aligned}
& \{(topReady!, running)\})st = state(st) \\
\Rightarrow & (state(st), (state \oplus (\{running_task, blocked\}) \cup \\
& \{(topReady!, running)\})st) \in transition) \\
\Rightarrow & t \in TASK \\
& \wedge state(t) = ready \\
& \wedge \neg priority(topReady!) \geq priority(t)
\end{aligned}$$

theorem QueueSendF_TQT_vc_ref

$\forall QueueSendF_TQTFSSig \mid true \bullet \text{pre } QueueSendF_TQT$

QueueSendW_TQT

$\Delta TaskQueueTime$

QueueSendW_TQ

clock' = *clock*
delayed_task' = *delayed_task*
wait_time' = {*topReady!*} \triangleleft *wait_time*
time_slice' = *time_slice*

$$\begin{aligned}
QueueSendW_TQTFSSig & \hat{=} TaskQueueTime \\
& \wedge QueueSendW_TQTFSSig
\end{aligned}$$

theorem lQueueSendW_TQT_Lemma

$\forall TaskQueueTime; topReady! : TASK; que? : QUEUE$
 $\mid running_task \notin \text{dom } release_rcv$
 $\wedge (running_task \in \text{dom } release_snd$
 $\Rightarrow que? = release_snd(running_task))$
 $\wedge que? \in queue$
 $\wedge q_size(que?) < q_max(que?)$
 $\wedge topReady! \in wait_rcv \sim (\{que?\})$
 $\wedge (\forall wrct : wait_rcv \sim (\{que?\})$
 $\bullet priority(topReady!) \geq priority(wrct))$
 $\wedge priority(running_task) \geq priority(topReady!)$
 $\bullet \neg (TaskQueueTime[q_size := q_size \oplus \{(que?, (1+$
 $q_size(que?))\},$
 $release_rcv := release_rcv \oplus \{(topReady!, que?\}),$
 $release_snd := \{running_task\} \triangleleft release_snd,$
 $state := state \oplus \{(topReady!, ready)\},$
 $wait_time := \{topReady!\} \triangleleft wait_time,$
 $wait_rcv := \{topReady!\} \triangleleft wait_rcv]$
 $\wedge priority(topReady!) \leq priority(running_task)$
 $\wedge (st \in TASK$
 $\wedge \neg (state \oplus \{(topReady!, ready)\})st = state(st)$
 $\Rightarrow (state(st), (state \oplus$

$$\begin{aligned}
& \{(topReady!, ready)\}st) \in transition) \\
\Rightarrow & wr \in \text{dom } wait_rcv \\
& \wedge wait_rcv(wr) = que? \\
& \wedge \neg priority(topReady!) \geq priority(wr)
\end{aligned}$$

theorem QueueSendW_TQT_vc_ref

$$\forall QueueSendW_TQTFSSig \mid true \bullet \text{pre } QueueSendW_TQT$$

$QueueSendWS_TQT$
$\Delta TaskQueueTime$
$QueueSendWS_TQ$
$clock' = clock$
$delayed_task' = delayed_task$
$wait_time' = \{topReady!\} \triangleleft wait_time$
$time_slice' = time_slice$

$$\begin{aligned}
QueueSendWS_TQTFSSig & \hat{=} TaskQueueTime \\
& \wedge QueueSendWS_TQTFSSig
\end{aligned}$$

theorem lQueueSendWS_TQT_Lemma

$$\begin{aligned}
& \forall TaskQueueTime; topReady! : TASK; que? : QUEUE \\
& \mid running_task \notin \text{dom } release_rcv \\
& \wedge (running_task \in \text{dom } release_snd \\
& \quad \Rightarrow que? = release_snd(running_task)) \\
& \wedge que? \in queue \\
& \wedge q_size(que?) < q_max(que?) \\
& \wedge topReady! \in wait_rcv \sim (\{que?\}) \\
& \wedge (\forall wrct : wait_rcv \sim (\{que?\}) \\
& \quad \bullet priority(topReady!) \geq priority(wrct)) \\
& \wedge priority(topReady!) > priority(running_task) \\
& \bullet \neg (TaskQueueTime[log_context := log_context \oplus \\
& \quad \{(running_task, phys_context)\}, \\
& \quad phys_context := log_context(topReady!), \\
& \quad q_size := q_size \oplus \{(wait_rcv(topReady!), (1+ \\
& \quad q_size(wait_rcv(topReady!)))\}, \\
& \quad release_rcv := release_rcv \oplus \\
& \quad \{(topReady!, wait_rcv(topReady!))\}, \\
& \quad release_snd := \{running_task\} \triangleleft release_snd, \\
& \quad running_task := topReady!, \\
& \quad state := state \oplus \\
& \quad \{(running_task, ready)\} \cup \{(topReady!, running)\}), \\
& \quad wait_time := \{topReady!\} \triangleleft wait_time, \\
& \quad wait_rcv := \{topReady!\} \triangleleft wait_rcv]
\end{aligned}$$

$$\begin{aligned}
& \wedge (st \in TASK \\
& \quad \wedge \neg (state \oplus (\{(running_task, ready)\} \cup \\
& \quad \quad \{(topReady!, running)\}))st = state(st) \\
& \quad \Rightarrow (state(st), (state \oplus (\{(running_task, ready)\} \cup \\
& \quad \quad \{(topReady!, running)\}))st) \in transition) \\
& \Rightarrow wr \in \text{dom } wait_rcv \\
& \quad \wedge wait_rcv(wr) = wait_rcv(topReady!) \\
& \quad \wedge \neg priority(topReady!) \geq priority(wr))
\end{aligned}$$

theorem QueueSendWS_TQT_vc_ref

$$\forall QueueSendWS_TQTFSSig \mid true \bullet \text{pre } QueueSendWS_TQT$$

$$\begin{aligned}
QueueSend_TQT & \cong QueueSendN_TQT \\
& \vee QueueSendF_TQT \\
& \vee QueueSendW_TQT \\
& \vee QueueSendWS_TQT
\end{aligned}$$

$$QueueReceiveN_TQT \cong ExtendTaskQueueXi \wedge QueueReceiveN_TQ$$

$$\begin{aligned}
QueueReceiveN_TQTFSSig & \cong TaskQueueTime \\
& \wedge QueueReceiveN_TQFSSig
\end{aligned}$$

theorem QueueReceiveN_TQT_vc_ref

$$\forall QueueReceiveN_TQTFSSig \mid true \bullet \text{pre } QueueReceiveN_TQT$$

$ \begin{aligned} & QueueReceiveE_TQT \\ & \Delta TaskQueueTime \\ & QueueReceiveE_TQ \\ & wtime? : \mathbb{N} \end{aligned} $
$ \begin{aligned} & wtime? > clock \\ & clock' = clock \\ & delayed_task' = delayed_task \\ & wait_time' = wait_time \oplus \{(running_task \mapsto wtime?)\} \\ & time_slice' = time_slice \end{aligned} $

$ \begin{aligned} & QueueReceiveE_TQTFSSig \\ & TaskQueueTime \\ & QueueReceiveE_TQFSSig \\ & wtime? : \mathbb{N} \end{aligned} $
$wtime? > clock$

theorem *lQueueReceiveE_TQT_Lemma*

$$\begin{aligned}
& \forall \text{TaskQueueTime}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK}; \text{wtime?} : \mathbb{N} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \\
& \quad \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge \text{q_size}(\text{que?}) = 0 \\
& \quad \wedge \text{running_task} \neq \text{idle} \\
& \quad \wedge \text{topReady!} \in \text{state} \sim (\{ \text{ready} \}) \\
& \quad \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \}) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\
& \quad \wedge \text{wtime?} > \text{clock} \\
& \quad \bullet \neg (\text{TaskQueueTime}[\text{log_context} := \text{log_context} \oplus \\
& \quad \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \quad \text{release_rcv} := \{ \text{running_task} \} \triangleleft \text{release_rcv}, \\
& \quad \quad \text{running_task} := \text{topReady!}, \\
& \quad \quad \text{state} := \text{state} \oplus \\
& \quad \quad \quad \{(\text{running_task}, \text{blocked})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \quad \text{wait_time} := \text{wait_time} \oplus \{(\text{running_task}, \text{wtime?})\}, \\
& \quad \quad \text{wait_rcv} := \text{wait_rcv} \oplus \{(\text{running_task}, \text{que?})\}] \\
& \quad \wedge (st \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\})) st = \text{state}(st) \\
& \quad \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\})) st) \in \text{transition}) \\
& \quad \Rightarrow t \in \text{TASK} \\
& \quad \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t)
\end{aligned}$$
theorem *QueueReceiveE_TQT_vc_ref*

$$\forall \text{QueueReceiveE_TQTFSSBSig} \mid \text{true} \bullet \text{pre } \text{QueueReceiveE_TQT}$$

$$\text{QueueReceiveW_TQT}$$

$$\Delta \text{TaskQueueTime}$$

$$\text{QueueReceiveW_TQ}$$

$$\text{clock}' = \text{clock}$$

$$\text{delayed_task}' = \text{delayed_task}$$

$$\text{wait_time}' = \{ \text{topReady!} \} \triangleleft \text{wait_time}$$

$$\text{time_slice}' = \text{time_slice}$$

$$\text{QueueReceiveW_TQTFSSBSig} \cong \text{TaskQueueTime}$$

$$\wedge \text{QueueReceiveW_TQTFSSBSig}$$

theorem lQueueReceiveW_TQT_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTime}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \\
& \quad \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge \text{q_size}(\text{que?}) > 0 \\
& \quad \wedge \text{topReady!} \in \text{wait_snd} \sim (\{ \text{que?} \} \setminus \emptyset) \\
& \quad \wedge (\forall \text{wsnt} : \text{wait_snd} \sim (\{ \text{que?} \} \setminus \emptyset) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wsnt})) \\
& \quad \wedge \text{priority}(\text{running_task}) \geq \text{priority}(\text{topReady!}) \\
& \quad \bullet \neg (\text{TaskQueueTime}[\text{q_size} := \text{q_size} \oplus \\
& \quad \quad \{(\text{que?}, (\text{q_size}(\text{que?}) - 1))\}, \\
& \quad \quad \text{release_rcv} := \{ \text{running_task} \} \triangleleft \text{release_rcv}, \\
& \quad \quad \text{release_snd} := \text{release_snd} \oplus \{(\text{topReady!}, \text{que?})\}, \\
& \quad \quad \text{state} := \text{state} \oplus \{(\text{topReady!}, \text{ready})\}, \\
& \quad \quad \text{wait_time} := \{ \text{topReady!} \} \triangleleft \text{wait_time}, \\
& \quad \quad \text{wait_snd} := \{ \text{topReady!} \} \triangleleft \text{wait_snd}] \\
& \quad \wedge \text{priority}(\text{topReady!}) \leq \text{priority}(\text{running_task}) \\
& \quad \wedge (\text{st} \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus \{(\text{topReady!}, \text{ready})\}) \text{st} = \text{state}(\text{st}) \\
& \quad \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus \\
& \quad \quad \quad \{(\text{topReady!}, \text{ready})\}) \text{st}) \in \text{transition}) \\
& \quad \Rightarrow \text{ws} \in \text{dom wait_snd} \\
& \quad \quad \wedge \text{wait_snd}(\text{ws}) = \text{que?} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{ws})
\end{aligned}$$
theorem QueueReceiveW_TQT_vc_ref
$$\forall \text{QueueReceiveW_TQTFSBSig} \mid \text{true} \bullet \text{pre QueueReceiveW_TQT}$$

$$\text{QueueReceiveWS_TQT} \text{-----}$$

$$\Delta \text{TaskQueueTime}$$

$$\text{QueueReceiveWS_TQ}$$

$$\text{clock}' = \text{clock}$$

$$\text{delayed_task}' = \text{delayed_task}$$

$$\text{wait_time}' = \{ \text{topReady!} \} \triangleleft \text{wait_time}$$

$$\text{time_slice}' = \text{time_slice}$$

$$\begin{aligned}
\text{QueueReceiveWS_TQTFSBSig} & \hat{=} \text{TaskQueueTime} \\
& \wedge \text{QueueReceiveWS_TQFSBSig}
\end{aligned}$$
theorem lQueueReceiveWS_TQT_Lemma
$$\forall \text{TaskQueueTime}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK}$$

$$\begin{aligned}
& | \text{running_task} \notin \text{dom release_snd} \\
& \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})) \\
& \wedge \text{que?} \in \text{queue} \\
& \wedge q_size(\text{que?}) > 0 \\
& \wedge \text{topReady!} \in \text{wait_snd} \sim (\{ \text{que?} \}) \\
& \wedge (\forall \text{wsnt} : \text{wait_snd} \sim (\{ \text{que?} \}) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wsnt})) \\
& \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{running_task}) \\
& \bullet \neg (\text{TaskQueueTime}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad q_size := q_size \oplus \{(\text{wait_snd}(\text{topReady!}), \\
& \quad (q_size(\text{wait_snd}(\text{topReady!})) - 1))\}, \\
& \quad \text{release_rcv} := \{\text{running_task}\} \triangleleft \text{release_rcv}, \\
& \quad \text{release_snd} := \text{release_snd} \oplus \\
& \quad \{(\text{topReady!}, \text{wait_snd}(\text{topReady!}))\}, \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus \\
& \quad \{(\text{running_task}, \text{ready})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \text{wait_time} := \{\text{topReady!}\} \triangleleft \text{wait_time}, \\
& \quad \text{wait_snd} := \{\text{topReady!}\} \triangleleft \text{wait_snd}] \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\})) st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\}))) st \in \text{transition}) \\
& \Rightarrow ws \in \text{dom wait_snd} \\
& \quad \wedge \text{wait_snd}(ws) = \text{wait_snd}(\text{topReady!}) \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(ws))
\end{aligned}$$

theorem QueueReceiveWS_TQT_vc_ref

$\forall \text{QueueReceiveWS_TQTFSSBSig} \mid \text{true} \bullet \text{pre QueueReceiveWS_TQT}$

$$\begin{aligned}
\text{QueueReceive_TQT} & \triangleq \text{QueueReceiveN_TQT} \\
& \vee \text{QueueReceiveE_TQT} \\
& \vee \text{QueueReceiveW_TQT} \\
& \vee \text{QueueReceiveWS_TQT}
\end{aligned}$$

DelayUntil_TQT

$\Delta \text{TaskQueueTime}$

$wtime? : \mathbb{N}$

$topReady! : \text{TASK}$

$\text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$

$$\begin{aligned}
& state(topReady!) = ready \\
& running_task \neq idle \\
& \forall t : TASK \mid state(t) = ready \bullet priority(topReady!) \geq priority(t) \\
& wtime? > clock \\
& \exists st? : STATE \mid st? = blocked \\
& \quad \bullet Reschedule[topReady!/target?, tasks/tasks?, priority/pri?] \\
& \exists Queue \\
& clock' = clock \\
& delayed_task' = delayed_task \cup \{running_task\} \\
& wait_time' = wait_time \oplus \{(running_task \mapsto wtime?)\} \\
& time_slice' = time_slice
\end{aligned}$$

DelayUntil_TQTFSSig

TaskQueueTime

$wtime? : \mathbb{N}$

$$\begin{aligned}
& running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv \\
& running_task \neq idle \\
& wtime? > clock
\end{aligned}$$

theorem |DelayUntil_TQT_Lemma

$$\begin{aligned}
& \forall TaskQueueTime; wtime? : \mathbb{N}; topReady! : TASK \\
& \quad \mid running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv \\
& \quad \wedge running_task \neq idle \\
& \quad \wedge state(topReady!) = ready \\
& \quad \wedge wtime? > clock \\
& \quad \wedge (\forall rtsk : state \sim (\{ready\} \mid) \\
& \quad \quad \bullet priority(topReady!) \geq priority(rtsk)) \\
& \quad \bullet \neg (TaskQueueTime[delayed_task := delayed_task \cup \\
& \quad \quad \{running_task\}, \\
& \quad \quad log_context := log_context \oplus \\
& \quad \quad \quad \{(running_task, phys_context)\}, \\
& \quad \quad phys_context := log_context(topReady!), \\
& \quad \quad running_task := topReady!, \\
& \quad \quad state := state \oplus \\
& \quad \quad \quad \{(running_task, blocked)\} \cup \{(topReady!, running)\}), \\
& \quad \quad wait_time := wait_time \oplus \{(running_task, wtime?)\}] \\
& \quad \wedge (st \in TASK \\
& \quad \quad \wedge \neg (state \oplus (\{(running_task, blocked)\} \cup \\
& \quad \quad \quad \{(topReady!, running)\})) st = state(st) \\
& \quad \quad \Rightarrow (state(st), (state \oplus (\{(running_task, blocked)\} \cup \\
& \quad \quad \quad \{(topReady!, running)\})) st) \in transition) \\
& \quad \Rightarrow t \in TASK \\
& \quad \quad \wedge state(t) = ready \\
& \quad \quad \wedge \neg priority(topReady!) \geq priority(t)
\end{aligned}$$

theorem DelayUntil_TQT_vc_ref
$$\forall \text{DelayUntil_TQTFSSBSig} \mid \text{true} \bullet \text{pre DelayUntil_TQT}$$

CheckDelayedTaskN_TQT

 $\Delta \text{TaskQueueTime}$ $\text{topWaiting!} : \text{TASK}$

$$\begin{aligned} & \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\ & \text{topWaiting!} \in \text{dom wait_time} \\ & \forall wt : \text{dom wait_time} \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(wt) \\ & \forall wt : \text{dom wait_time} \mid \text{wait_time}(wt) = \text{wait_time}(\text{topWaiting!}) \\ & \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(wt) \\ & \text{priority}(\text{topWaiting!}) \leq \text{priority}(\text{running_task}) \\ & \exists \text{TaskData} \\ & \text{state}' = \text{state} \oplus \{(\text{topWaiting!} \mapsto \text{ready})\} \\ & \exists \text{ContextData} \\ & \exists \text{PrioData} \\ & \exists \text{QueueData} \\ & \text{wait_snd}' = \{\text{topWaiting!}\} \triangleleft \text{wait_snd} \\ & \text{wait_rcv}' = \{\text{topWaiting!}\} \triangleleft \text{wait_rcv} \\ & \exists \text{QReleasingData} \\ & \text{clock}' = \text{wait_time}(\text{topWaiting!}) \\ & \text{delayed_task}' = \text{delayed_task} \setminus \{\text{topWaiting!}\} \\ & \text{wait_time}' = \{\text{topWaiting!}\} \triangleleft \text{wait_time} \\ & \text{time_slice}' = \text{time_slice} \end{aligned}$$

CheckDelayedTaskN_TQTFSSBSig

TaskQueueTime

$$\begin{aligned} & \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\ & \exists \text{topWaiting!} : \text{dom wait_time} \\ & \quad \bullet (\forall wt : \text{dom wait_time} \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(wt)) \\ & \quad \wedge (\forall wt : \text{dom wait_time} \mid \text{wait_time}(wt) = \text{wait_time}(\text{topWaiting!}) \\ & \quad \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(wt)) \\ & \quad \wedge \text{priority}(\text{topWaiting!}) \leq \text{priority}(\text{running_task}) \end{aligned}$$

theorem lCheckDelayedTaskN_TQT_Lemma
$$\forall \text{TaskQueueTime}; \text{topWaiting!} : \text{TASK}$$

$$\mid \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$$

$$\wedge \text{topWaiting!} \in \text{dom wait_time}$$

$$\wedge (\forall \text{dtk} : \text{dom wait_time}$$

$$\quad \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(\text{dtk}))$$

$$\wedge (\forall \text{detk} : \text{dom wait_time}$$

$$\quad \mid \text{wait_time}(\text{detk}) = \text{wait_time}(\text{topWaiting!}))$$

- $priority(topWaiting!) \geq priority(detk)$

$$\wedge priority(topWaiting!) \leq priority(running_task)$$

- $\neg (TaskQueueTime[*clock* := $wait_time(topWaiting!)$,
 $delayed_task := delayed_task \setminus \{topWaiting!\}$,
 $state := state \oplus \{(topWaiting!, ready)\}$,
 $wait_time := \{topWaiting!\} \triangleleft wait_time$,
 $wait_rcv := \{topWaiting!\} \triangleleft wait_rcv$,
 $wait_snd := \{topWaiting!\} \triangleleft wait_snd$]$

$$\wedge (st \in TASK$$

$$\wedge \neg (state \oplus \{(topWaiting!, ready)\})st = state(st)$$

$$\Rightarrow (state(st), (state \oplus$$

$$\{(topWaiting!, ready)\})st) \in transition)$$

$$\wedge (wt \in \text{dom } wait_time$$

$$\Rightarrow wait_time(topWaiting!) \leq wait_time(wt))$$

$$\Rightarrow wt_0 \in \text{dom } wait_time$$

$$\wedge wait_time(wt_0) = wait_time(topWaiting!)$$

$$\wedge \neg priority(topWaiting!) \geq priority(wt_0))$$

theorem CheckDelayedTaskN_TQT_vc_ref

$\forall CheckDelayedTaskN_TQTFSSBSig \mid true$

- pre $CheckDelayedTaskN_TQT$

CheckDelayedTaskS_TQT

$\Delta TaskQueueTime$

$topWaiting! : TASK$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$topWaiting! \in \text{dom } wait_time$

$\forall wt : \text{dom } wait_time \bullet wait_time(topWaiting!) \leq wait_time(wt)$

$\forall wt : \text{dom } wait_time \mid wait_time(wt) = wait_time(topWaiting!)$

- $priority(topWaiting!) \geq priority(wt)$

$priority(topWaiting!) > priority(running_task)$

$\exists st? : STATE \mid st? = ready$

- $Reschedule[topWaiting!/target?, tasks/tasks?, priority/pri?]$

$\exists QueueData$

$wait_snd' = \{topWaiting!\} \triangleleft wait_snd$

$wait_rcv' = \{topWaiting!\} \triangleleft wait_rcv$

$\exists QReleasingData$

$clock' = wait_time(topWaiting!)$

$delayed_task' = delayed_task \setminus \{topWaiting!\}$

$wait_time' = \{topWaiting!\} \triangleleft wait_time$

$time_slice' = time_slice$

$$\begin{aligned}
 & \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
 & \exists \text{topWaiting!} : \text{dom wait_time} \\
 & \bullet (\forall wt : \text{dom wait_time} \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(wt)) \\
 & \wedge (\forall wt : \text{dom wait_time} \mid \text{wait_time}(wt) = \text{wait_time}(\text{topWaiting!}) \\
 & \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(wt)) \\
 & \wedge \text{priority}(\text{topWaiting!}) > \text{priority}(\text{running_task})
 \end{aligned}$$

theorem lCheckDelayedTaskS_TQT_Lemma

$$\begin{aligned}
 & \forall \text{TaskQueueTime}; \text{topWaiting!} : \text{TASK} \\
 & \mid \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
 & \wedge \text{topWaiting!} \in \text{dom wait_time} \\
 & \wedge (\forall dtk : \text{dom wait_time} \\
 & \quad \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(dtk)) \\
 & \wedge (\forall detk : \text{dom wait_time} \\
 & \quad \mid \text{wait_time}(detk) = \text{wait_time}(\text{topWaiting!}) \\
 & \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(detk)) \\
 & \wedge \text{priority}(\text{topWaiting!}) > \text{priority}(\text{running_task}) \\
 & \bullet \neg (\text{TaskQueueTime}[\text{clock} := \text{wait_time}(\text{topWaiting!}), \\
 & \quad \text{delayed_task} := \text{delayed_task} \setminus \{\text{topWaiting!}\}, \\
 & \quad \text{log_context} := \text{log_context} \oplus \\
 & \quad \quad \{(\text{running_task}, \text{phys_context} \}, \\
 & \quad \text{phys_context} := \text{log_context}(\text{topWaiting!}), \\
 & \quad \text{running_task} := \text{topWaiting!}, \\
 & \quad \text{state} := \text{state} \oplus \\
 & \quad \quad \{(\text{running_task}, \text{ready} \}) \cup \{(\text{topWaiting!}, \text{running} \}) \}, \\
 & \quad \text{wait_time} := \{ \text{topWaiting!} \} \triangleleft \text{wait_time}, \\
 & \quad \text{wait_rcv} := \{ \text{topWaiting!} \} \triangleleft \text{wait_rcv}, \\
 & \quad \text{wait_snd} := \{ \text{topWaiting!} \} \triangleleft \text{wait_snd}] \\
 & \wedge (st \in \text{TASK} \\
 & \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready} \}) \cup \\
 & \quad \quad \{(\text{topWaiting!}, \text{running} \} \}))st = \text{state}(st) \\
 & \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{ready} \}) \cup \\
 & \quad \quad \{(\text{topWaiting!}, \text{running} \} \}))st) \in \text{transition}) \\
 & \wedge (wt \in \text{dom wait_time} \\
 & \quad \Rightarrow \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(wt)) \\
 & \Rightarrow wt_0 \in \text{dom wait_time} \\
 & \quad \wedge \text{wait_time}(wt_0) = \text{wait_time}(\text{topWaiting!}) \\
 & \quad \wedge \neg \text{priority}(\text{topWaiting!}) \geq \text{priority}(wt_0))
 \end{aligned}$$

theorem CheckDelayedTaskS_TQT_vc_ref

$$\begin{aligned}
 & \forall \text{CheckDelayedTaskS_TQTFSSBSig} \mid \text{true} \\
 & \bullet \text{pre CheckDelayedTaskS_TQT}
 \end{aligned}$$

TimeSlicing_TQT

Δ *TaskQueueTime*

topReady! : *TASK*

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$state(topReady!) = ready$

$priority(topReady!) = priority(running_task)$

$\forall t : \text{dom } wait_time \bullet time_slice \leq wait_time(t)$

$\exists st? : STATE \mid st? = ready$

• *Reschedule*[*topReady!*/*target?*, *tasks/tasks?*, *priority/pri?*]

$\exists Queue$

$clock' = clock$

$delayed_task' = delayed_task$

$wait_time' = wait_time$

$time_slice' = time_slice + slice_delay$

TimeSlicing_TQTFSSBSig

TaskQueueTime

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$\forall t : \text{dom } wait_time \bullet time_slice \leq wait_time(t)$

$\exists topReady! : state \sim (\{ready\})$

• $priority(topReady!) = priority(running_task)$

theorem *!TimeSlicing_TQT_Lemma*

$\forall TaskQueueTime; topReady! : TASK$

$\mid running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$\wedge state(topReady!) = ready$

$\wedge priority(topReady!) = priority(running_task)$

$\wedge (\forall ts : \text{dom } wait_time \bullet time_slice \leq wait_time(ts))$

• $\neg (TaskQueueTime[log_context := log_context \oplus$

$\{(running_task, phys_context)\},$

$phys_context := log_context(topReady!),$

$running_task := topReady!,$

$state := state \oplus$

$\{(running_task, ready)\} \cup \{(topReady!, running)\}),$

$time_slice := time_slice + slice_delay]$

$\Rightarrow st \in TASK$

$\wedge \neg (state \oplus (\{(running_task, ready)\} \cup$

$\{(topReady!, running)\}))st = state(st)$

$\wedge \neg (state(st), (state \oplus (\{(running_task, ready)\} \cup$

$\{(topReady!, running)\}))st \in transition)$

theorem TimeSlicing_TQT_vc_ref

$\forall \text{TimeSlicing_TQTFSSig} \mid \text{true} \bullet \text{pre } \text{TimeSlicing_TQT}$

NoSlicing_TQT

$\Delta \text{TaskQueueTime}$

$\text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

$\forall t : \text{dom } \text{wait_time} \bullet \text{time_slice} \leq \text{wait_time}(t)$

$\forall t : \text{state} \sim (\{ \text{ready} \} \mid) \bullet \text{priority}(t) < \text{priority}(\text{running_task})$

$\exists \text{Task}$

$\exists \text{Queue}$

$\text{clock}' = \text{clock}$

$\text{delayed_task}' = \text{delayed_task}$

$\text{wait_time}' = \text{wait_time}$

$\text{time_slice}' = \text{time_slice} + \text{slice_delay}$

$\text{NoSlicing_TQTFSSig}$

TaskQueueTime

$\text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

$\forall t : \text{dom } \text{wait_time} \bullet \text{time_slice} \leq \text{wait_time}(t)$

$\forall t : \text{state} \sim (\{ \text{ready} \} \mid) \bullet \text{priority}(t) < \text{priority}(\text{running_task})$

theorem NoSlicing_TQT_vc_ref

$\forall \text{NoSlicing_TQTFSSig} \mid \text{true} \bullet \text{pre } \text{NoSlicing_TQT}$

Appendix G

SPECIFICATION FOR MUTEX MODEL

MutexData

$semaphore : \mathbb{P} \text{ QUEUE}$
 $mutex : \mathbb{P} \text{ QUEUE}$
 $mutex_holder : \text{ QUEUE} \rightarrow \text{ TASK}$
 $mutex_recursive : \text{ QUEUE} \rightarrow \mathbb{N}$

$mutex \cap semaphore = \emptyset$
 $\text{dom } mutex_recursive = mutex$
 $\forall m : mutex \bullet m \notin \text{dom } mutex_holder \Leftrightarrow mutex_recursive(m) = 0$

Init_MutexData

MutexData'

$semaphore' = \emptyset$
 $mutex' = \emptyset$
 $mutex_holder' = \emptyset$
 $mutex_recursive' = \emptyset$

theorem *MutexDataInit*

$\exists \text{ MutexData}' \bullet \text{ Init_MutexData}$

OriginalPrioData

$base_priority : \text{ TASK} \rightarrow \mathbb{N}$

Init_OriginalPrioData

OriginalPrioData'

$base_priority' = \emptyset$

theorem OriginalPrioDataInit

$\exists OriginalPrioData' \bullet Init_OriginalPrioData$

MReleasingData

$release_mutex : TASK \rightarrow QUEUE$

Init_MReleasingData

MReleasingData'

$release_mutex' = \emptyset$

theorem MReleasingDataInit

$\exists MReleasingData' \bullet Init_MReleasingData$

Mutex

MutexData

OriginalPrioData

MReleasingData

$dom\ base_priority = ran\ mutex_holder$

$ran\ release_mutex \subseteq mutex$

Init_Mutex

Mutex'

Init_MutexData

Init_OriginalPrioData

Init_MReleasingData

theorem MutexInit

$\exists Mutex' \bullet Init_Mutex$

TaskQueueTimeMutex

TaskQueueTime

Mutex

$$\begin{aligned}
& semaphore \subseteq queue \\
& \forall s : semaphore \bullet q_max(s) = 1 \\
& mutex \subseteq queue \\
& \forall m : mutex \bullet q_max(m) = 1 \\
& dom\ mutex_holder = \{m : mutex \mid q_size(m) = 0\} \\
& ran\ mutex_holder \subseteq tasks \\
& \forall mh : ran\ mutex_holder \bullet priority(mh) \geq base_priority(mh) \\
& \forall ms : mutex \cup semaphore \bullet ms \notin ran\ wait_snd \cup ran\ release_snd \\
& release_mutex \subseteq release_rcv
\end{aligned}$$

$$\begin{aligned}
\Delta TaskQueueTimeMutex & \hat{=} TaskQueueTimeMutex \\
& \wedge TaskQueueTimeMutex' \\
& \wedge \Delta Task
\end{aligned}$$

$$Init_TaskQueueTimeMutex$$

$$TaskQueueTimeMutex'$$

$$Init_TaskQueueTime$$

$$Init_Mutex$$

theorem TaskQueueTimeMutexInit

$$\exists TaskQueueTimeMutex' \bullet Init_TaskQueueTimeMutex$$

$$ExtendTQTXi$$

$$\Delta TaskQueueTimeMutex$$

$$\exists Mutex$$

$$CreateTaskN_TQTM \hat{=} ExtendTQTXi \wedge CreateTaskN_TQT$$

$$\begin{aligned}
CreateTaskN_TQTMFSBSig & \hat{=} TaskQueueTimeMutex \\
& \wedge CreateTaskN_TQTFBSBSig
\end{aligned}$$

theorem CreateTaskN_TQTM_vc_ref

$$\forall CreateTaskN_TQTMFSBSig \mid true \bullet pre\ CreateTaskN_TQTM$$

$$CreateTaskS_TQTM \hat{=} ExtendTQTXi \wedge CreateTaskS_TQT$$

$$\begin{aligned}
CreateTaskS_TQTMFSBSig & \hat{=} TaskQueueTimeMutex \\
& \wedge CreateTaskS_TQTFBSBSig
\end{aligned}$$

theorem CreateTaskS_TQTM_vc_ref

\forall CreateTaskS_TQTMFSBSig | true • pre CreateTaskS_TQTM

$CreateTask_TQTM \cong CreateTaskN_TQTM \vee CreateTaskS_TQTM$

DeleteTaskN_TQTM

Δ TaskQueueTimeMutex

DeleteTaskN_TQT

$target? \notin \text{ran } mutex_holder$

\exists MutexData

\exists OriginalPrioData

$release_mutex' = \{target?\} \triangleleft release_mutex$

DeleteTaskN_TQTMFSBSig

TaskQueueTimeMutex

DeleteTaskN_TQTFBSBSig

$target? \notin \text{ran } mutex_holder$

theorem DeleteTaskN_TQTM_vc_ref

\forall DeleteTaskN_TQTMFSBSig | true • pre DeleteTaskN_TQTM

DeleteTaskS_TQTM

Δ TaskQueueTimeMutex

DeleteTaskS_TQT

$target? \notin \text{ran } mutex_holder$

\exists Mutex

DeleteTaskS_TQTMFSBSig

TaskQueueTimeMutex

DeleteTaskS_TQTFBSBSig

$target? \notin \text{ran } mutex_holder$

theorem lDeleteTaskS_TQTM_Lemma

\forall TaskQueueTimeMutex; topReady!, target? : TASK

| running_task \notin dom release_snd \cup dom release_rcv

\wedge target? \in tasks \setminus {idle}

\wedge state(target?) \in {running}

$$\begin{aligned}
& \wedge \text{state}(\text{topReady!}) = \text{ready} \\
& \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \mid)) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}) \\
& \wedge \text{target?} \notin \text{ran } \text{mutex_holder} \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{target?}, \text{bare_context})\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus \\
& \quad \{(\text{target?}, \text{nonexistent})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \text{tasks} := \text{tasks} \setminus \{\text{target?}\}] \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{target?}, \text{nonexistent})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\}))st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{target?}, \text{nonexistent})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\})))st \in \text{transition}) \\
& \Rightarrow t \in \text{TASK} \\
& \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t))
\end{aligned}$$

theorem DeleteTaskS_TQTM_vc_ref

$\forall \text{DeleteTaskS_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{DeleteTaskS_TQTM}$

$\text{DeleteTask_TQTM} \cong \text{DeleteTaskN_TQTM} \vee \text{DeleteTaskS_TQTM}$

$\text{ExecuteRunningTask_TQTM} \cong \text{ExtendTQTXi}$
 $\wedge \text{ExecuteRunningTask_TQT}$

$\text{ExecuteRunningTask_TQTMFSBSig} \cong \text{TaskQueueTimeMutex}$
 $\wedge \text{ExecuteRunningTask_TQTFFSBSig}$

theorem ExecuteRunningTask_TQTM_vc_ref

$\forall \text{ExecuteRunningTask_TQTMFSBSig} \mid \text{true}$
 $\bullet \text{pre } \text{ExecuteRunningTask_TQTM}$

$\text{SuspendTaskN_TQTM} \cong \text{ExtendTQTXi} \wedge \text{SuspendTaskN_TQT}$

$\text{SuspendTaskN_TQTMFSBSig} \cong \text{TaskQueueTimeMutex}$
 $\wedge \text{SuspendTaskN_TQTFFSBSig}$

theorem SuspendTaskN_TQTM_vc_ref

$\forall \text{SuspendTaskN_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{SuspendTaskN_TQTM}$

$$\text{SuspendTaskS_TQTM} \cong \text{ExtendTQTXi} \wedge \text{SuspendTaskS_TQT}$$

$$\begin{aligned} \text{SuspendTaskS_TQTMFSBSig} &\cong \text{TaskQueueTimeMutex} \\ &\wedge \text{SuspendTaskS_TQTFBSBSig} \end{aligned}$$

theorem lSuspendTaskS_TQTM_Lemma

$$\begin{aligned} &\forall \text{TaskQueueTimeMutex}; \text{target?}, \text{topReady!} : \text{TASK} \\ &| \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\ &\wedge \text{target?} \in \text{tasks} \setminus \{\text{idle}\} \\ &\wedge \text{state}(\text{target?}) \in \{\text{running}\} \\ &\wedge \text{state}(\text{topReady!}) = \text{ready} \\ &\wedge (\forall \text{rtsk} : \text{state} \sim (\{\text{ready}\}) | \\ &\quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\ &\bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\ &\quad \{(\text{running_task}, \text{phys_context} \}), \\ &\quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\ &\quad \text{running_task} := \text{topReady!}, \\ &\quad \text{state} := \text{state} \oplus (\{(\text{running_task}, \text{suspended} \}) \cup \\ &\quad \{(\text{topReady!}, \text{running} \})]) \\ &\wedge (st \in \text{TASK} \\ &\quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{suspended} \}) \cup \\ &\quad \{(\text{topReady!}, \text{running} \})})) st = \text{state}(st) \\ &\quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{suspended} \}) \cup \\ &\quad \{(\text{topReady!}, \text{running} \})})) st \in \text{transition}) \\ &\Rightarrow t \in \text{TASK} \\ &\quad \wedge \text{state}(t) = \text{ready} \\ &\quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t)) \end{aligned}$$

theorem SuspendTaskS_TQTM_vc_ref

$$\forall \text{SuspendTaskS_TQTMFSBSig} | \text{true} \bullet \text{pre SuspendTaskS_TQTM}$$

$$\text{SuspendTaskO_TQTM} \cong \text{ExtendTQTXi} \wedge \text{SuspendTaskO_TQT}$$

$$\begin{aligned} \text{SuspendTaskO_TQTMFSBSig} &\cong \text{TaskQueueTimeMutex} \\ &\wedge \text{SuspendTaskO_TQTFBSBSig} \end{aligned}$$

theorem SuspendTaskO_TQTM_vc_ref

$$\forall \text{SuspendTaskO_TQTMFSBSig} | \text{true} \bullet \text{pre SuspendTaskO_TQTM}$$

$$\begin{aligned} \text{SuspendTask_TQTM} &\cong \text{SuspendTaskN_TQTM} \\ &\vee \text{SuspendTaskS_TQTM} \\ &\vee \text{SuspendTaskO_TQTM} \end{aligned}$$

$$\text{ResumeTaskN_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{ResumeTaskN_TQT}$$
$$\begin{aligned} \text{ResumeTaskN_TQTMFSBSig} \hat{=} & \text{TaskQueueTimeMutex} \\ & \wedge \text{ResumeTaskN_TQTFBSBSig} \end{aligned}$$

theorem ResumeTaskN_TQTM_vc_ref

$$\forall \text{ResumeTaskN_TQTMFSBSig} \mid \text{true} \bullet \text{pre ResumeTaskN_TQTM}$$
$$\text{ResumeTaskS_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{ResumeTaskS_TQT}$$
$$\begin{aligned} \text{ResumeTaskS_TQTMFSBSig} \hat{=} & \text{TaskQueueTimeMutex} \\ & \wedge \text{ResumeTaskS_TQTFBSBSig} \end{aligned}$$

theorem ResumeTaskS_TQTM_vc_ref

$$\forall \text{ResumeTaskS_TQTMFSBSig} \mid \text{true} \bullet \text{pre ResumeTaskS_TQTM}$$
$$\text{ResumeTask_TQTM} \hat{=} \text{ResumeTaskN_TQTM} \vee \text{ResumeTaskS_TQTM}$$

ChangeTaskPriorityNNotHolder_TQTM _____

$\Delta \text{TaskQueueTimeMutex}$
 $\text{ChangeTaskPriorityN_TQT}$

$\text{target?} \notin \text{dom base_priority}$
 $\exists \text{Mutex}$

ChangeTaskPriorityNNotHolder_TQTMFSBSig _____

$\text{ChangeTaskPriorityN_TQTFBSBSig}$
 $\text{TaskQueueTimeMutex}$

$\text{target?} \notin \text{dom base_priority}$

theorem ChangeTaskPriorityNNotHolder_TQTM_vc_ref

$$\begin{aligned} & \forall \text{ChangeTaskPriorityNNotHolder_TQTMFSBSig} \mid \text{true} \\ & \bullet \text{pre ChangeTaskPriorityNNotHolder_TQTM} \end{aligned}$$

ChangeTaskPrioritySNotHolder_TQTM _____

$\text{ChangeTaskPriorityS_TQT}$
 $\Delta \text{TaskQueueTimeMutex}$

$\text{target?} \notin \text{dom base_priority}$
 $\exists \text{Mutex}$

ChangeTaskPrioritySNotHolder_TQTMFSBSig

ChangeTaskPriorityS_TQTFBSBSig

TaskQueueTimeMutex

$target? \notin \text{dom } base_priority$

theorem *ChangeTaskPrioritySNotHolder_TQTM_vc_ref*

$\forall \text{ChangeTaskPrioritySNotHolder_TQTMFSBSig} \mid \text{true}$

- $\text{pre } \text{ChangeTaskPrioritySNotHolder_TQTM}$

ChangeTaskPriorityDNotHolder_TQTM

$\Delta \text{TaskQueueTimeMutex}$

ChangeTaskPriorityD_TQT

$target? \notin \text{dom } base_priority$

$\exists \text{Mutex}$

ChangeTaskPriorityDNotHolder_TQTMFSBSig

TaskQueueTimeMutex

ChangeTaskPriorityD_TQTFBSBSig

$target? \notin \text{dom } base_priority$

theorem *lChangeTaskPriorityDNotHolder_TQTM_Lemma*

$\forall \text{TaskQueueTimeMutex}; target?, topReady! : \text{TASK}; newpri? : \mathbb{N}$

$\mid \text{state}(target?) = \text{running}$

$\wedge (target? = \text{idle} \Rightarrow newpri? = 0)$

$\wedge \text{state}(topReady!) = \text{ready}$

$\wedge (\forall rtsk : \text{state} \sim (\{ \text{ready} \} \mid))$

- $\text{priority}(topReady!) \geq \text{priority}(rtsk)$

$\wedge newpri? < \text{priority}(topReady!)$

$\wedge target? \notin \text{dom } base_priority$

- $\neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus$

$\{(running_task, \text{phys_context})\},$

$\text{phys_context} := \text{log_context}(topReady!),$

$\text{priority} := \text{priority} \oplus \{(target?, newpri?)\},$

$running_task := topReady!,$

$\text{state} := \text{state} \oplus$

$\{(running_task, \text{ready})\} \cup \{(topReady!, \text{running})\}])$

$\wedge newpri? < \text{priority}(topReady!)$

$\wedge (st \in \text{TASK}$

$$\begin{aligned}
& \wedge \neg (state \oplus (\{(running_task, ready)\} \cup \\
& \quad \{(topReady!, running)\})) st = state(st) \\
& \Rightarrow (state(st), (state \oplus (\{(running_task, ready)\} \cup \\
& \quad \{(topReady!, running)\}))) st \in transition) \\
& \Rightarrow t \in TASK \\
& \quad \wedge state(t) = ready \\
& \quad \wedge \neg priority(topReady!) \geq priority(t)
\end{aligned}$$

theorem ChangeTaskPriorityDNotHolder_TQTM_vc_ref
 \forall ChangeTaskPriorityDNotHolder_TQTMFSBSig | true
• pre ChangeTaskPriorityDNotHolder_TQTM

ChangeTaskPriorityNNotInherited_TQTM

ChangeTaskPriorityN_TQT
 Δ TaskQueueTimeMutex

target? \in dom base_priority
base_priority(target?) = priority(target?)
 \exists MutexData
base_priority' = base_priority \oplus {(target? \mapsto newpri?)}
 \exists MReleasingData

ChangeTaskPriorityNNotInherited_TQTMFSBSig

ChangeTaskPriorityN_TQTFBSBSig
TaskQueueTimeMutex

target? \in dom base_priority
base_priority(target?) = priority(target?)

theorem ChangeTaskPriorityNNotInherited_TQTM_vc_ref
 \forall ChangeTaskPriorityNNotInherited_TQTMFSBSig | true
• pre ChangeTaskPriorityNNotInherited_TQTM

ChangeTaskPrioritySNotInherited_TQTM

ChangeTaskPriorityS_TQT
 Δ TaskQueueTimeMutex

target? \in dom base_priority
base_priority(target?) = priority(target?)
 \exists MutexData
base_priority' = base_priority \oplus {(target? \mapsto newpri?)}
 \exists MReleasingData

ChangeTaskPrioritySNotInherited_TQTMFSBSig _____

ChangeTaskPriorityS_TQTFFSBSig

TaskQueueTimeMutex

$target? \in \text{dom } base_priority$

$base_priority(target?) = priority(target?)$

theorem *ChangeTaskPrioritySNotInherited_TQTM_vc_ref*

$\forall \text{ChangeTaskPrioritySNotInherited_TQTMFSBSig} \mid \text{true}$

• $\text{pre } \text{ChangeTaskPrioritySNotInherited_TQTM}$

ChangeTaskPriorityDNotInherited_TQTM _____

$\Delta \text{TaskQueueTimeMutex}$

ChangeTaskPriorityD_TQT

$target? \in \text{dom } base_priority$

$base_priority(target?) = priority(target?)$

$\exists \text{MutexData}$

$base_priority' = base_priority \oplus \{(target? \mapsto newpri?)\}$

$\exists \text{MReleasingData}$

ChangeTaskPriorityDNotInherited_TQTMFSBSig _____

TaskQueueTimeMutex

ChangeTaskPriorityD_TQTFFSBSig

$target? \in \text{dom } base_priority$

$base_priority(target?) = priority(target?)$

theorem *lChangeTaskPriorityDNotInherited_TQTM_Lemma*

$\forall \text{TaskQueueTimeMutex}; target?, topReady! : \text{TASK}; newpri? : \mathbb{N}$

$\mid \text{state}(target?) = \text{running}$

$\wedge (target? = \text{idle} \Rightarrow newpri? = 0)$

$\wedge \text{state}(topReady!) = \text{ready}$

$\wedge (\forall rtsk : \text{state} \sim (\{ \text{ready} \} \mid))$

• $priority(topReady!) \geq priority(rtsk)$

$\wedge newpri? < priority(topReady!)$

$\wedge target? \in \text{dom } base_priority$

$\wedge base_priority(target?) = priority(target?)$

• $\neg (\text{TaskQueueTimeMutex}[base_priority := base_priority \oplus$

$\{(target?, newpri?)\},$

$log_context := log_context \oplus$

$\{(running_task, phys_context)\},$

$phys_context := log_context(topReady!),$

$$\begin{aligned}
& \text{priority} := \text{priority} \oplus \{(target?, newpri?)\}, \\
& \text{running_task} := \text{topReady!}, \\
& \text{state} := \text{state} \oplus \\
& \quad \{(\text{running_task}, \text{ready})\} \cup \{(\text{topReady!}, \text{running})\}] \\
& \wedge \text{newpri?} < \text{priority}(\text{topReady!}) \\
& \wedge (st \in TASK \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \quad \{(\text{topReady!}, \text{running})\}))st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \quad \{(\text{topReady!}, \text{running})\}))st) \in \text{transition}) \\
& \Rightarrow t \in TASK \\
& \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t))
\end{aligned}$$

theorem *ChangeTaskPriorityDNotInherited_TQTM_vc_ref*
 $\forall \text{ChangeTaskPriorityDNotInherited_TQTMFSBSig} \mid \text{true}$
• pre *ChangeTaskPriorityDNotInherited_TQTM*

ChangeTaskPriorityInheritedN_TQTM _____

$\Delta \text{TaskQueueTimeMutex}$

$\text{newpri?} : \mathbb{N}$

$\text{target?} : TASK$

$\text{topReady!} : TASK$

$\text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$

$\text{target?} \in \text{dom base_priority}$

$\text{base_priority}(\text{target?}) \neq \text{priority}(\text{target?})$

$\text{newpri?} \leq \text{priority}(\text{target?})$

$\text{state}(\text{target?}) \neq \text{nonexistent}$

$\text{target?} = \text{idle} \Rightarrow \text{newpri?} = 0$

$\exists \text{TaskQueueTime}$

$\exists \text{MutexData}$

$\text{base_priority}' = \text{base_priority} \oplus \{(target? \mapsto \text{newpri?})\}$

$\exists M\text{ReleasingData}$

$\text{topReady!} = \text{running_task}$

ChangeTaskPriorityInheritedN_TQTMFSBSig _____

$\text{TaskQueueTimeMutex}$

$\text{newpri?} : \mathbb{N}$

$\text{target?} : TASK$

$\text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$

$\text{target?} \in \text{dom base_priority}$

$\text{base_priority}(\text{target?}) \neq \text{priority}(\text{target?})$

$$\begin{aligned} & \text{newpri?} \leq \text{priority}(\text{target?}) \\ & \text{state}(\text{target?}) \neq \text{nonexistent} \\ & \text{target?} = \text{idle} \Rightarrow \text{newpri?} = 0 \end{aligned}$$

theorem *ChangeTaskPriorityInheritedN_TQTM_vc_ref*
 $\forall \text{ChangeTaskPriorityInheritedN_TQTMFSBSig} \mid \text{true}$
 • pre *ChangeTaskPriorityInheritedN_TQTM*

ChangeTaskPriorityInheritedU_TQTM _____

Δ *TaskQueueTimeMutex*
ChangeTaskPriorityN_TQT

$$\begin{aligned} & \text{target?} \in \text{dom } \text{base_priority} \\ & \text{base_priority}(\text{target?}) \neq \text{priority}(\text{target?}) \\ & \text{newpri?} > \text{priority}(\text{target?}) \\ & \exists \text{MutexData} \\ & \text{base_priority}' = \text{base_priority} \oplus \{(\text{target?} \mapsto \text{newpri?})\} \\ & \exists \text{MReleasingData} \end{aligned}$$

ChangeTaskPriorityInheritedU_TQTMFSBSig _____

ChangeTaskPriorityN_TQTFBSBSig
TaskQueueTimeMutex

$$\begin{aligned} & \text{target?} \in \text{dom } \text{base_priority} \\ & \text{base_priority}(\text{target?}) \neq \text{priority}(\text{target?}) \\ & \text{newpri?} > \text{priority}(\text{target?}) \end{aligned}$$

theorem *ChangeTaskPriorityInheritedU_TQTM_vc_ref*
 $\forall \text{ChangeTaskPriorityInheritedU_TQTMFSBSig} \mid \text{true}$
 • pre *ChangeTaskPriorityInheritedU_TQTM*

ChangeTaskPriorityInheritedS_TQTM _____

Δ *TaskQueueTimeMutex*
ChangeTaskPriorityS_TQT

$$\begin{aligned} & \text{target?} \in \text{dom } \text{base_priority} \\ & \text{base_priority}(\text{target?}) \neq \text{priority}(\text{target?}) \\ & \text{newpri?} > \text{priority}(\text{target?}) \\ & \exists \text{MutexData} \\ & \text{base_priority}' = \text{base_priority} \oplus \{(\text{target?} \mapsto \text{newpri?})\} \\ & \exists \text{MReleasingData} \end{aligned}$$

ChangeTaskPriorityInheritedS_TQTMFSBSig

ChangeTaskPriorityS_TQTFBSBSig

TaskQueueTimeMutex

$target? \in \text{dom } base_priority$

$base_priority(target?) \neq priority(target?)$

$newpri? > priority(target?)$

theorem *ChangeTaskPriorityInheritedS_TQTM_vc_ref*

$\forall \text{ChangeTaskPriorityInheritedS_TQTMFSBSig} \mid \text{true}$

• $\text{pre } \text{ChangeTaskPriorityInheritedS_TQTM}$

$\text{ChangeTaskPriority_TQTM} \equiv \text{ChangeTaskPriorityNNotHolder_TQTM}$
 $\vee \text{ChangeTaskPrioritySNotHolder_TQTM}$
 $\vee \text{ChangeTaskPriorityDNotHolder_TQTM}$
 $\vee \text{ChangeTaskPriorityNNotInherited_TQTM}$
 $\vee \text{ChangeTaskPrioritySNotInherited_TQTM}$
 $\vee \text{ChangeTaskPriorityDNotInherited_TQTM}$
 $\vee \text{ChangeTaskPriorityInheritedN_TQTM}$
 $\vee \text{ChangeTaskPriorityInheritedU_TQTM}$
 $\vee \text{ChangeTaskPriorityInheritedS_TQTM}$

$\text{CreateQueue_TQTM} \equiv \text{ExtendTQTXi} \wedge \text{CreateQueue_TQT}$

$\text{CreateQueue_TQTMFSBSig} \equiv \text{TaskQueueTimeMutex}$

$\wedge \text{CreateQueue_TQTFBSBSig}$

theorem *CreateQueue_TQTM_vc_ref*

$\forall \text{CreateQueue_TQTMFSBSig} \mid \text{true}$ • $\text{pre } \text{CreateQueue_TQTM}$

DeleteQueue_TQTM

DeleteQueue_TQT

$\Delta \text{TaskQueueTimeMutex}$

$que? \notin \text{semaphore} \cup \text{mutex}$

$\exists \text{Mutex}$

DeleteQueue_TQTMFSBSig

TaskQueueTimeMutex

DeleteQueue_TQTFBSBSig

$que? \notin \text{semaphore} \cup \text{mutex}$

theorem DeleteQueue_TQTM_vc_ref

$\forall DeleteQueue_TQTMFSBSig \mid true \bullet \text{pre } DeleteQueue_TQTM$

QueueSendN_TQTM

QueueSendN_TQT

$\Delta TaskQueueTimeMutex$

$que? \notin mutex \cup semaphore$

$\exists Mutex$

QueueSendN_TQTMFSBSig

TaskQueueTimeMutex

QueueSendN_TQTFBSBSig

$que? \notin mutex \cup semaphore$

theorem QueueSendN_TQTM_vc_ref

$\forall QueueSendN_TQTMFSBSig \mid true \bullet \text{pre } QueueSendN_TQTM$

QueueSendF_TQTM

$\Delta TaskQueueTimeMutex$

QueueSendF_TQT

$que? \notin mutex \cup semaphore$

$\exists Mutex$

QueueSendF_TQTMFSBSig

TaskQueueTimeMutex

QueueSendF_TQTFBSBSig

$que? \notin mutex \cup semaphore$

theorem lQueueSendF_TQTM_Lemma

$\forall TaskQueueTimeMutex; topReady! : TASK; que? : QUEUE; wtime? : \mathbb{N}$

$\mid running_task \notin \text{dom } release_rcv$

$\wedge (running_task \in \text{dom } release_snd$

$\Rightarrow que? = release_snd(running_task))$

$\wedge que? \in queue$

$\wedge q_size(que?) = q_max(que?)$

$\wedge running_task \neq idle$

$$\begin{aligned}
& \wedge \text{state}(\text{topReady!}) = \text{ready} \\
& \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \mid)) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}) \\
& \wedge \text{wtime?} > \text{clock} \\
& \wedge \text{que?} \notin \text{mutex} \cup \text{semaphore} \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{release_snd} := \{ \text{running_task} \} \triangleleft \text{release_snd}, \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus \\
& \quad \{(\text{running_task}, \text{blocked})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \text{wait_time} := \text{wait_time} \oplus \{(\text{running_task}, \text{wtime?})\}, \\
& \quad \text{wait_snd} := \text{wait_snd} \oplus \{(\text{running_task}, \text{que?})\}] \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\}))st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\})))st \in \text{transition}) \\
& \Rightarrow t \in \text{TASK} \\
& \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t))
\end{aligned}$$

theorem QueueSendF_TQTM_vc_ref

$$\forall \text{QueueSendF_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{QueueSendF_TQTM}$$

$ \begin{aligned} & \text{QueueSendW_TQTM} \\ & \Delta \text{TaskQueueTimeMutex} \\ & \text{QueueSendW_TQT} \end{aligned} $
$ \begin{aligned} & \text{que?} \notin \text{mutex} \cup \text{semaphore} \\ & \exists \text{Mutex} \end{aligned} $

$ \begin{aligned} & \text{QueueSendW_TQTMFSBSig} \\ & \text{TaskQueueTimeMutex} \\ & \text{QueueSendW_TQTFBSBSig} \end{aligned} $
$ \text{que?} \notin \text{mutex} \cup \text{semaphore} $

theorem overrideIsDisjointUnion [X, Y]

$$\forall f, g : X \rightarrow Y \mid \text{dom } f \cap \text{dom } g = \emptyset \bullet f \subseteq f \oplus g$$

theorem lQueueSendW_TQTM_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{topReady!} : \text{TASK}; \text{que?} : \text{QUEUE} \\
& \quad | \text{running_task} \notin \text{dom release_rcv} \\
& \quad \wedge (\text{running_task} \in \text{dom release_snd} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_snd}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge \text{q_size}(\text{que?}) < \text{q_max}(\text{que?}) \\
& \quad \wedge \text{topReady!} \in \text{wait_rcv} \sim (\{ \text{que?} \} \mid) \\
& \quad \wedge (\forall \text{wrct} : \text{wait_rcv} \sim (\{ \text{que?} \} \mid) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wrct})) \\
& \quad \wedge \text{priority}(\text{running_task}) \geq \text{priority}(\text{topReady!}) \\
& \quad \wedge \text{que?} \notin \text{mutex} \cup \text{semaphore} \\
& \quad \bullet \neg (\text{TaskQueueTimeMutex}[\text{q_size} := \text{q_size} \oplus \{(\text{que?}, (1+ \\
& \quad \quad \text{q_size}(\text{que?})))\}, \\
& \quad \quad \text{release_rcv} := \text{release_rcv} \oplus \{(\text{topReady!}, \text{que?})\}, \\
& \quad \quad \text{release_snd} := \{ \text{running_task} \} \triangleleft \text{release_snd}, \\
& \quad \quad \text{state} := \text{state} \oplus \{(\text{topReady!}, \text{ready})\}, \\
& \quad \quad \text{wait_time} := \{ \text{topReady!} \} \triangleleft \text{wait_time}, \\
& \quad \quad \text{wait_rcv} := \{ \text{topReady!} \} \triangleleft \text{wait_rcv}] \\
& \quad \wedge \text{priority}(\text{topReady!}) \leq \text{priority}(\text{running_task}) \\
& \quad \wedge (\text{st} \in \text{TASK} \wedge \neg (\text{state} \oplus \{(\text{topReady!}, \text{ready})\})\text{st} = \text{state}(\text{st}) \\
& \quad \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus \{(\text{topReady!}, \text{ready})\})\text{st}) \\
& \quad \quad \quad \in \text{transition}) \\
& \quad \Rightarrow \text{wr} \in \text{dom wait_rcv} \\
& \quad \quad \wedge \text{wait_rcv}(\text{wr}) = \text{que?} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wr}))
\end{aligned}$$
theorem QueueSendW_TQTM_vc_ref
$$\forall \text{QueueSendW_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{QueueSendW_TQTM}$$

$ \begin{aligned} & \text{QueueSendWS_TQTM} \\ & \Delta \text{TaskQueueTimeMutex} \\ & \text{QueueSendWS_TQT} \end{aligned} $
$ \begin{aligned} & \text{que?} \notin \text{mutex} \cup \text{semaphore} \\ & \exists \text{Mutex} \end{aligned} $

$ \begin{aligned} & \text{QueueSendWS_TQTMFSBSig} \\ & \text{TaskQueueTimeMutex} \\ & \text{QueueSendWS_TQTFSBSig} \end{aligned} $
$ \text{que?} \notin \text{mutex} \cup \text{semaphore} $

theorem lQueueSendWS_TQTM_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{topReady!} : \text{TASK}; \text{que?} : \text{QUEUE} \\
& \quad | \text{running_task} \notin \text{dom release_rcv} \\
& \quad \wedge (\text{running_task} \in \text{dom release_snd} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_snd}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge q_size(\text{que?}) < q_max(\text{que?}) \\
& \quad \wedge \text{topReady!} \in \text{wait_rcv} \sim (\{ \text{que?} \}) \\
& \quad \wedge (\forall \text{wrct} : \text{wait_rcv} \sim (\{ \text{que?} \}) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wrct})) \\
& \quad \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{running_task}) \\
& \quad \wedge \text{que?} \notin \text{mutex} \cup \text{semaphore} \\
& \quad \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \quad q_size := q_size \oplus \{(\text{wait_rcv}(\text{topReady!}), (1 + \\
& \quad \quad \quad q_size(\text{wait_rcv}(\text{topReady!})))\}, \\
& \quad \quad \text{release_rcv} := \text{release_rcv} \oplus \\
& \quad \quad \{(\text{topReady!}, \text{wait_rcv}(\text{topReady!}))\}, \\
& \quad \quad \text{release_snd} := \{\text{running_task}\} \triangleleft \text{release_snd}, \\
& \quad \quad \text{running_task} := \text{topReady!}, \\
& \quad \quad \text{state} := \text{state} \oplus \\
& \quad \quad \{(\text{running_task}, \text{ready})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \quad \text{wait_time} := \{\text{topReady!}\} \triangleleft \text{wait_time}, \\
& \quad \quad \text{wait_rcv} := \{\text{topReady!}\} \triangleleft \text{wait_rcv}] \\
& \quad \wedge (st \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\}))st = \text{state}(st) \\
& \quad \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{ready})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\}))st) \in \text{transition}) \\
& \quad \Rightarrow wr \in \text{dom wait_rcv} \\
& \quad \quad \wedge \text{wait_rcv}(wr) = \text{wait_rcv}(\text{topReady!}) \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(wr))
\end{aligned}$$
theorem QueueSendWS_TQTM_vc_ref
$$\forall \text{QueueSendWS_TQTMFSBSig} \mid \text{true} \bullet \text{pre QueueSendWS_TQTM}$$

$$\begin{aligned}
\text{QueueSend_TQTM} & \cong \text{QueueSendN_TQTM} \\
& \vee \text{QueueSendF_TQTM} \\
& \vee \text{QueueSendW_TQTM} \\
& \vee \text{QueueSendWS_TQTM}
\end{aligned}$$

$$\text{QueueReceiveN_TQTM}$$

$$\text{QueueReceiveN_TQT}$$

$$\Delta \text{TaskQueueTimeMutex}$$

$running_task \notin \text{dom } release_mutex$
 $que? \notin mutex \cup semaphore$
 $\exists Mutex$

$QueueReceiveN_TQTMFSBSig$
 $TaskQueueTimeMutex$
 $QueueReceiveN_TQTFBSBSig$

$running_task \notin \text{dom } release_mutex$
 $que? \notin mutex \cup semaphore$

theorem mutexDiffQue

$\forall TaskQueueTimeMutex \mid que? \notin mutex \cup semaphore$
 $\bullet \text{ dom } mutex_holder = \{f_1 : mutex$
 $\mid (q_size \oplus \{(que?, (q_size(que?) - 1))\})f_1 = 0\}$

theorem QueueReceiveN_TQTM_vc_ref

$\forall QueueReceiveN_TQTMFSBSig \mid true \bullet \text{ pre } QueueReceiveN_TQTM$

$QueueReceiveE_TQTM$
 $\Delta TaskQueueTimeMutex$
 $QueueReceiveE_TQT$

$running_task \notin \text{dom } release_mutex$
 $que? \notin mutex \cup semaphore$
 $\exists Mutex$

$QueueReceiveE_TQTMFSBSig$
 $TaskQueueTimeMutex$
 $QueueReceiveE_TQTFBSBSig$

$running_task \notin \text{dom } release_mutex$
 $que? \notin mutex \cup semaphore$

theorem lQueueReceiveE_TQTM_Lemma

$\forall TaskQueueTimeMutex; que? : QUEUE; topReady! : TASK; wtime? : \mathbb{N}$
 $\mid running_task \notin \text{dom } release_snd$
 $\wedge (running_task \in \text{dom } release_rcv$
 $\Rightarrow que? = release_rcv(running_task))$
 $\wedge que? \in queue$
 $\wedge q_size(que?) = 0$

$$\begin{aligned}
& \wedge \text{running_task} \neq \text{idle} \\
& \wedge \text{topReady!} \in \text{state} \sim (\{ \text{ready} \} \downarrow) \\
& \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \downarrow) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\
& \wedge \text{wtime?} > \text{clock} \\
& \wedge \text{running_task} \notin \text{dom release_mutex} \\
& \wedge \text{que?} \notin \text{mutex} \cup \text{semaphore} \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{release_rcv} := \{\text{running_task}\} \triangleleft \text{release_rcv}, \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus \\
& \quad \quad \{(\text{running_task}, \text{blocked})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \text{wait_time} := \text{wait_time} \oplus \{(\text{running_task}, \text{wtime?})\}, \\
& \quad \text{wait_rcv} := \text{wait_rcv} \oplus \{(\text{running_task}, \text{que?})\}] \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \{(\text{topReady!}, \text{running})\})) st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \{(\text{topReady!}, \text{running})\})) st) \in \text{transition}) \\
& \Rightarrow t \in \text{TASK} \\
& \quad \wedge \text{state}(t) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(t))
\end{aligned}$$

theorem QueueReceiveE_TQTM_vc_ref

$\forall \text{QueueReceiveE_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{QueueReceiveE_TQTM}$

$\text{QueueReceiveW_TQTM}$

$\Delta \text{TaskQueueTimeMutex}$

QueueReceiveW_TQT

$\text{running_task} \notin \text{dom release_mutex}$

$\text{que?} \notin \text{mutex} \cup \text{semaphore}$

$\exists \text{Mutex}$

$\text{QueueReceiveW_TQTMFSBSig}$

$\text{TaskQueueTimeMutex}$

$\text{QueueReceiveW_TQTFBSBSig}$

$\text{running_task} \notin \text{dom release_mutex}$

$\text{que?} \notin \text{mutex} \cup \text{semaphore}$

theorem lQueueReceiveW_TQTM_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \\
& \quad \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \quad \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})) \\
& \quad \wedge \text{que?} \in \text{queue} \\
& \quad \wedge \text{q_size}(\text{que?}) > 0 \\
& \quad \wedge \text{topReady!} \in \text{wait_snd} \sim (\{ \text{que?} \}) \\
& \quad \wedge (\forall \text{wsnt} : \text{wait_snd} \sim (\{ \text{que?} \}) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wsnt})) \\
& \quad \wedge \text{priority}(\text{running_task}) \geq \text{priority}(\text{topReady!}) \\
& \quad \wedge \text{running_task} \notin \text{dom release_mutex} \\
& \quad \wedge \text{que?} \notin \text{mutex} \cup \text{semaphore} \\
& \quad \bullet \neg (\text{TaskQueueTimeMutex}[q_size := q_size \oplus \\
& \quad \quad \{(\text{que?}, (\text{q_size}(\text{que?}) - 1))\}, \\
& \quad \quad \text{release_rcv} := \{ \text{running_task} \} \triangleleft \text{release_rcv}, \\
& \quad \quad \text{release_snd} := \text{release_snd} \oplus \{(\text{topReady!}, \text{que?})\}, \\
& \quad \quad \text{state} := \text{state} \oplus \{(\text{topReady!}, \text{ready})\}, \\
& \quad \quad \text{wait_time} := \{ \text{topReady!} \} \triangleleft \text{wait_time}, \\
& \quad \quad \text{wait_snd} := \{ \text{topReady!} \} \triangleleft \text{wait_snd}] \\
& \quad \wedge \text{priority}(\text{topReady!}) \leq \text{priority}(\text{running_task}) \\
& \quad \wedge (st \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus \{(\text{topReady!}, \text{ready})\})st = \text{state}(st) \\
& \quad \quad \quad \Rightarrow (\text{state}(st), (\text{state} \oplus \{(\text{topReady!}, \text{ready})\})st) \\
& \quad \quad \quad \in \text{transition}) \\
& \quad \Rightarrow ws \in \text{dom wait_snd} \\
& \quad \quad \wedge \text{wait_snd}(ws) = \text{que?} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(ws))
\end{aligned}$$
theorem QueueReceiveW_TQTM_vc_ref
$$\forall \text{QueueReceiveW_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{QueueReceiveW_TQTM}$$

$ \begin{aligned} & \text{QueueReceiveWS_TQTM} \\ & \Delta \text{TaskQueueTimeMutex} \\ & \text{QueueReceiveWS_TQT} \end{aligned} $
$ \begin{aligned} & \text{running_task} \notin \text{dom release_mutex} \\ & \text{que?} \notin \text{mutex} \cup \text{semaphore} \\ & \exists \text{Mutex} \end{aligned} $

$ \begin{aligned} & \text{QueueReceiveWS_TQTMFSBSig} \\ & \text{TaskQueueTimeMutex} \\ & \text{QueueReceiveWS_TQTFBSBSig} \end{aligned} $

$$\begin{array}{l} \text{running_task} \notin \text{dom release_mutex} \\ \text{que?} \notin \text{mutex} \cup \text{semaphore} \end{array}$$

theorem lQueueReceiveWS_TQTM_Lemma

$$\begin{array}{l} \forall \text{TaskQueueTimeMutex}; \text{que?} : \text{QUEUE}; \text{topReady!} : \text{TASK} \\ | \text{running_task} \notin \text{dom release_snd} \\ \wedge (\text{running_task} \in \text{dom release_rcv} \\ \Rightarrow \text{que?} = \text{release_rcv}(\text{running_task})) \\ \wedge \text{que?} \in \text{queue} \\ \wedge \text{q_size}(\text{que?}) > 0 \\ \wedge \text{topReady!} \in \text{wait_snd} \sim (\{ \text{que?} \} \downarrow) \\ \wedge (\forall \text{wsnt} : \text{wait_snd} \sim (\{ \text{que?} \} \downarrow) \\ \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wsnt})) \\ \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{running_task}) \\ \wedge \text{running_task} \notin \text{dom release_mutex} \\ \wedge \text{que?} \notin \text{mutex} \cup \text{semaphore} \\ \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\ \{(\text{running_task}, \text{phys_context} \}), \\ \text{phys_context} := \text{log_context}(\text{topReady!}), \\ \text{q_size} := \text{q_size} \oplus \{ (\text{wait_snd}(\text{topReady!}), \\ (\text{q_size}(\text{wait_snd}(\text{topReady!})) - 1) \}), \\ \text{release_rcv} := \{ \text{running_task} \} \triangleleft \text{release_rcv}, \\ \text{release_snd} := \text{release_snd} \oplus \\ \{ (\text{topReady!}, \text{wait_snd}(\text{topReady!})) \}, \\ \text{running_task} := \text{topReady!}, \\ \text{state} := \text{state} \oplus \\ (\{ (\text{running_task}, \text{ready}) \} \cup \{ (\text{topReady!}, \text{running}) \}), \\ \text{wait_time} := \{ \text{topReady!} \} \triangleleft \text{wait_time}, \\ \text{wait_snd} := \{ \text{topReady!} \} \triangleleft \text{wait_snd}] \\ \wedge (\text{st} \in \text{TASK} \\ \wedge \neg (\text{state} \oplus (\{ (\text{running_task}, \text{ready}) \} \cup \\ \{ (\text{topReady!}, \text{running}) \})) \text{st} = \text{state}(\text{st}) \\ \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus (\{ (\text{running_task}, \text{ready}) \} \cup \\ \{ (\text{topReady!}, \text{running}) \}))) \text{st} \in \text{transition}) \\ \Rightarrow \text{ws} \in \text{dom wait_snd} \\ \wedge \text{wait_snd}(\text{ws}) = \text{wait_snd}(\text{topReady!}) \\ \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{ws})) \end{array}$$

theorem QueueReceiveWS_TQTM_vc_ref

$$\forall \text{QueueReceiveWS_TQTMFSBSig} \mid \text{true} \bullet \text{pre } \text{QueueReceiveWS_TQTM}$$

$$\begin{array}{l} \text{QueueReceive_TQTM} \cong \text{QueueReceiveN_TQTM} \\ \vee \text{QueueReceiveE_TQTM} \\ \vee \text{QueueReceiveW_TQTM} \end{array}$$

$$\vee \text{QueueReceiveWS_TQTM}$$

$$\text{DelayUntil_TQTM} \cong \text{ExtendTQTXi} \wedge \text{DelayUntil_TQT}$$

$$\begin{aligned} \text{DelayUntil_TQTMFSBSig} &\cong \text{TaskQueueTimeMutex} \\ &\wedge \text{DelayUntil_TQTF SBSig} \end{aligned}$$

theorem lDelayUntil_TQTM_Lemma

$$\begin{aligned} &\forall \text{TaskQueueTimeMutex}; \text{wtime?} : \mathbb{N}; \text{topReady!} : \text{TASK} \\ &| \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\ &\wedge \text{running_task} \neq \text{idle} \\ &\wedge \text{state}(\text{topReady!}) = \text{ready} \\ &\wedge \text{wtime?} > \text{clock} \\ &\wedge (\forall \text{rtsk} : \text{state} \sim \{ \text{ready} \}) \\ &\quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}) \\ &\quad \bullet \neg (\text{TaskQueueTimeMutex}[\text{delayed_task} := \text{delayed_task} \cup \\ &\quad \quad \{ \text{running_task} \}, \\ &\quad \quad \text{log_context} := \text{log_context} \oplus \\ &\quad \quad \{ (\text{running_task}, \text{phys_context}) \}, \\ &\quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\ &\quad \quad \text{running_task} := \text{topReady!}, \\ &\quad \quad \text{state} := \text{state} \oplus \\ &\quad \quad \{ (\text{running_task}, \text{blocked}) \} \cup \{ (\text{topReady!}, \text{running}) \}), \\ &\quad \quad \text{wait_time} := \text{wait_time} \oplus \{ (\text{running_task}, \text{wtime?}) \}] \\ &\wedge (\text{st} \in \text{TASK} \\ &\quad \wedge \neg (\text{state} \oplus (\{ (\text{running_task}, \text{blocked}) \} \cup \\ &\quad \quad \{ (\text{topReady!}, \text{running}) \})) \text{st} = \text{state}(\text{st}) \\ &\quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus (\{ (\text{running_task}, \text{blocked}) \} \cup \\ &\quad \quad \{ (\text{topReady!}, \text{running}) \}))) \text{st} \in \text{transition}) \\ &\Rightarrow \text{t} \in \text{TASK} \\ &\quad \wedge \text{state}(\text{t}) = \text{ready} \\ &\quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{t}) \end{aligned}$$

theorem DelayUntil_TQTM_vc_ref

$$\forall \text{DelayUntil_TQTMFSBSig} | \text{true} \bullet \text{pre DelayUntil_TQTM}$$

$$\text{CheckDelayedTaskN_TQTM} \cong \text{ExtendTQTXi} \wedge \text{CheckDelayedTaskN_TQT}$$

$$\begin{aligned} \text{CheckDelayedTaskN_TQTMFSBSig} &\cong \text{TaskQueueTimeMutex} \\ &\wedge \text{CheckDelayedTaskN_TQTF SBSig} \end{aligned}$$

theorem lCheckDelayedTaskN_TQTM_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{topWaiting!} : \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \quad \wedge \text{topWaiting!} \in \text{dom wait_time} \\
& \quad \wedge (\forall \text{dtk} : \text{dom wait_time} \\
& \quad \quad \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(\text{dtk})) \\
& \quad \wedge (\forall \text{detk} : \text{dom wait_time} \\
& \quad \quad | \text{wait_time}(\text{detk}) = \text{wait_time}(\text{topWaiting!}) \\
& \quad \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(\text{detk})) \\
& \quad \wedge \text{priority}(\text{topWaiting!}) \leq \text{priority}(\text{running_task}) \\
& \quad \bullet \neg (\text{TaskQueueTimeMutex}[\text{clock} := \text{wait_time}(\text{topWaiting!}), \\
& \quad \quad \text{delayed_task} := \text{delayed_task} \setminus \{\text{topWaiting!}\}, \\
& \quad \quad \text{state} := \text{state} \oplus \{(\text{topWaiting!}, \text{ready})\}, \\
& \quad \quad \text{wait_time} := \{\text{topWaiting!}\} \triangleleft \text{wait_time}, \\
& \quad \quad \text{wait_rcv} := \{\text{topWaiting!}\} \triangleleft \text{wait_rcv}, \\
& \quad \quad \text{wait_snd} := \{\text{topWaiting!}\} \triangleleft \text{wait_snd}] \\
& \quad \wedge (st \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus \{(\text{topWaiting!}, \text{ready})\})st = \text{state}(st) \\
& \quad \quad \quad \Rightarrow (\text{state}(st), (\text{state} \oplus \{(\text{topWaiting!}, \text{ready})\})st) \\
& \quad \quad \quad \in \text{transition}) \\
& \quad \wedge (wt \in \text{dom wait_time} \\
& \quad \quad \Rightarrow \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(wt)) \\
& \quad \Rightarrow \text{wt_0} \in \text{dom wait_time} \\
& \quad \quad \wedge \text{wait_time}(\text{wt_0}) = \text{wait_time}(\text{topWaiting!}) \\
& \quad \quad \wedge \neg \text{priority}(\text{topWaiting!}) \geq \text{priority}(\text{wt_0}))
\end{aligned}$$
theorem CheckDelayedTaskN_TQTM_vc_ref
$$\begin{aligned}
& \forall \text{CheckDelayedTaskN_TQTMFSBSig} \mid \text{true} \\
& \quad \bullet \text{pre CheckDelayedTaskN_TQTM}
\end{aligned}$$

$$\text{CheckDelayedTaskS_TQTM} \cong \text{ExtendTQTXi} \wedge \text{CheckDelayedTaskS_TQT}$$

$$\begin{aligned}
\text{CheckDelayedTaskS_TQTMFSBSig} & \cong \text{TaskQueueTimeMutex} \\
& \wedge \text{CheckDelayedTaskS_TQTFBSBSig}
\end{aligned}$$
theorem lCheckDelayedTaskS_TQTM_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{topWaiting!} : \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \quad \wedge \text{topWaiting!} \in \text{dom wait_time} \\
& \quad \wedge (\forall \text{dtk} : \text{dom wait_time} \\
& \quad \quad \bullet \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(\text{dtk})) \\
& \quad \wedge (\forall \text{detk} : \text{dom wait_time} \\
& \quad \quad | \text{wait_time}(\text{detk}) = \text{wait_time}(\text{topWaiting!}) \\
& \quad \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(\text{detk}))
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{priority}(\text{topWaiting!}) > \text{priority}(\text{running_task}) \\
& \bullet \neg (\text{wait_time}(\text{topWaiting!}) \in \mathbb{Z}) \\
& \quad \wedge \text{TaskQueueTimeMutex}[\text{clock} := \text{wait_time}(\text{topWaiting!}), \\
& \quad \quad \text{delayed_task} := \text{delayed_task} \setminus \{\text{topWaiting!}\}, \\
& \quad \quad \text{log_context} := \text{log_context} \oplus \\
& \quad \quad \quad \{(\text{running_task}, \text{phys_context} \}, \\
& \quad \quad \text{phys_context} := \text{log_context}(\text{topWaiting!}), \\
& \quad \quad \text{running_task} := \text{topWaiting!}, \\
& \quad \quad \text{state} := \text{state} \oplus \\
& \quad \quad \quad \{(\text{running_task}, \text{ready} \} \cup \{(\text{topWaiting!}, \text{running} \} \}), \\
& \quad \quad \text{wait_time} := \{ \text{topWaiting!} \} \triangleleft \text{wait_time}, \\
& \quad \quad \text{wait_rcv} := \{ \text{topWaiting!} \} \triangleleft \text{wait_rcv}, \\
& \quad \quad \text{wait_snd} := \{ \text{topWaiting!} \} \triangleleft \text{wait_snd}] \\
& \quad \wedge (st \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{ready} \} \cup \\
& \quad \quad \quad \{(\text{topWaiting!}, \text{running} \} \}))st = \text{state}(st) \\
& \quad \quad \Rightarrow (\text{state}(st), (\text{state} \oplus (\{(\text{running_task}, \text{ready} \} \cup \\
& \quad \quad \quad \{(\text{topWaiting!}, \text{running} \} \})))st \in \text{transition}) \\
& \quad \wedge (wt \in \text{dom wait_time} \\
& \quad \quad \Rightarrow \text{wait_time}(\text{topWaiting!}) \leq \text{wait_time}(wt)) \\
& \Rightarrow wt_0 \in \text{dom wait_time} \\
& \quad \wedge \text{wait_time}(wt_0) = \text{wait_time}(\text{topWaiting!}) \\
& \quad \wedge \neg \text{priority}(\text{topWaiting!}) \geq \text{priority}(wt_0))
\end{aligned}$$

theorem CheckDelayedTaskS_TQTM_vc_ref

$\forall \text{CheckDelayedTaskS_TQTMFSBSig} \mid \text{true}$

$\bullet \text{pre CheckDelayedTaskS_TQTM}$

$\text{TimeSlicing_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{TimeSlicing_TQT}$

$\text{TimeSlicing_TQTMFSBSig} \hat{=} \text{TaskQueueTimeMutex}$
 $\quad \wedge \text{TimeSlicing_TQTFBSBSig}$

theorem lTimeSlicing_TQTM_Lemma

$\forall \text{TaskQueueTimeMutex}; \text{topReady!} : \text{TASK}$

$\mid \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$

$\wedge \text{state}(\text{topReady!}) = \text{ready}$

$\wedge \text{priority}(\text{topReady!}) = \text{priority}(\text{running_task})$

$\wedge (\forall ts : \text{dom wait_time} \bullet \text{time_slice} \leq \text{wait_time}(ts))$

$\bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus$

$\quad \{(\text{running_task}, \text{phys_context} \},$

$\quad \text{phys_context} := \text{log_context}(\text{topReady!}),$

$\quad \text{running_task} := \text{topReady!},$

$\quad \text{state} := \text{state} \oplus$

$$\begin{aligned}
& (\{(running_task, ready)\} \cup \{(topReady!, running)\}), \\
& time_slice := time_slice + slice_delay] \\
\Rightarrow st \in TASK \\
& \wedge \neg (state \oplus (\{(running_task, ready)\} \cup \\
& \quad \{(topReady!, running)\})) st = state(st) \\
& \wedge \neg (state(st), (state \oplus (\{(running_task, ready)\} \cup \\
& \quad \{(topReady!, running)\}))) st \in transition)
\end{aligned}$$

theorem TimeSlicing_TQTM_vc_ref

$$\forall TimeSlicing_TQTMFSBSig \mid true \bullet \text{pre } TimeSlicing_TQTM$$

$$NoSlicing_TQTM \hat{=} ExtendTQTXi \wedge NoSlicing_TQT$$

$$\begin{aligned}
NoSlicing_TQTMFSBSig \hat{=} TaskQueueTimeMutex \\
\wedge NoSlicing_TQTFBSBSig
\end{aligned}$$

theorem NoSlicing_TQTM_vc_ref

$$\forall NoSlicing_TQTMFSBSig \mid true \bullet \text{pre } NoSlicing_TQTM$$

CreateBinarySemaphore_TQTM _____

$\Delta TaskQueueTimeMutex$
 $sem? : QUEUE$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $sem? \notin queue$
 $\exists Task$
 $queue' = queue \cup \{sem?\}$
 $q_max' = q_max \oplus \{(sem? \mapsto 1)\}$
 $q_size' = q_size \oplus \{(sem? \mapsto 1)\}$
 $\exists WaitingData$
 $\exists QReleasingData$
 $\exists Time$
 $semaphore' = semaphore \cup \{sem?\}$
 $mutex' = mutex$
 $mutex_holder' = mutex_holder$
 $mutex_recursive' = mutex_recursive$
 $\exists OriginalPrioData$
 $\exists MReleasingData$

CreateBinarySemaphore_TQTMFSBSig _____

$TaskQueueTimeMutex$
 $sem? : QUEUE$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $sem? \notin queue$

theorem CreateBinarySemaphore_TQTM_vc_ref

$\forall CreateBinarySemaphore_TQTMFSBSig \mid true$

- pre $CreateBinarySemaphore_TQTM$

$DeleteBinarySemaphore_TQTM$

$\Delta TaskQueueTimeMutex$

$sem? : QUEUE$

$sem? \in semaphore$

$DeleteQueue_TQT[sem?/que?]$

$semaphore' = semaphore \setminus \{sem?\}$

$mutex' = mutex$

$mutex_holder' = mutex_holder$

$mutex_recursive' = mutex_recursive$

$\exists OriginalPrioData$

$\exists MReleasingData$

$DeleteBinarySemaphore_TQTMFSBSig$

$DeleteQueue_TQTFBSBSig[sem?/que?]$

$TaskQueueTimeMutex$

$sem? : QUEUE$

$sem? \in semaphore$

theorem DeleteBinarySemaphore_TQTM_vc_ref

$\forall DeleteBinarySemaphore_TQTMFSBSig \mid true$

- pre $DeleteBinarySemaphore_TQTM$

$CreateMutex_TQTM$

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$mut? \notin queue$

$\exists Task$

$queue' = queue \cup \{mut?\}$

$q_max' = q_max \oplus \{(mut? \mapsto 1)\}$

$q_size' = q_size \oplus \{(mut? \mapsto 1)\}$

$\exists WaitingData$

$\exists QReleasingData$
 $\exists Time$
 $semaphore' = semaphore$
 $mutex' = mutex \cup \{mut?\}$
 $mutex_holder' = mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $\exists OriginalPrioData$
 $\exists MReleasingData$

CreateMutex_TQTMFSBSig

TaskQueueTimeMutex
 $mut? : QUEUE$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \notin queue$

theorem CreateMutex_TQTM_vc_ref

$\forall CreateMutex_TQTMFSBSig \mid true \bullet \text{pre } CreateMutex_TQTM$

DeleteMutex_TQTM

$\Delta TaskQueueTimeMutex$
 $mut? : QUEUE$

$mut? \in mutex \setminus \text{dom } mutex_holder$
 $DeleteQueue_TQT[que? := mut?]$
 $semaphore' = semaphore$
 $mutex' = mutex \setminus \{mut?\}$
 $mutex_holder' = mutex_holder$
 $mutex_recursive' = \{mut?\} \triangleleft mutex_recursive$
 $\exists OriginalPrioData$
 $\exists MReleasingData$

DeleteMutex_TQTMFSBSig

TaskQueueTimeMutex
 $mut? : QUEUE$

$DeleteQueue_TQTFSSBSig[que? := mut?]$
 $mut? \in mutex \setminus \text{dom } mutex_holder$

theorem subPfun [X, Y]

$\forall f, g : X \mapsto Y; y : Y \mid g \subseteq f \bullet y \notin \text{ran } f \Rightarrow y \notin \text{ran } g$

theorem DeleteMutex_TQTM_vc_ref

$\forall DeleteMutex_TQTMFSBSig \mid true \bullet \text{pre } DeleteMutex_TQTM$

MutexTakeNnonInh_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$topReady! : TASK$

$running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$

$mut? \in mutex$

$running_task \notin \text{dom } base_priority$

$QueueReceiveN_TQT[que? := mut?]$

$semaphore' = semaphore$

$mutex' = mutex$

$mutex_holder' = mutex_holder \oplus \{(mut? \mapsto running_task)\}$

$mutex_recursive' = mutex_recursive \oplus$

$\{(mut? \mapsto mutex_recursive(mut?) + 1)\}$

$base_priority' = base_priority \oplus$

$\{(running_task \mapsto priority(running_task))\}$

$release_mutex' = \{running_task\} \triangleleft release_mutex$

MutexTakeNnonInh_TQTMFSBSig

TaskQueueTimeMutex

$mut? : QUEUE$

$QueueReceiveN_TQTFSSBSig[que? := mut?]$

$running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$

$mut? \in mutex$

$running_task \notin \text{dom } base_priority$

theorem *MutexTakeNnonInh_TQTM_vc_ref*

$\forall MutexTakeNnonInh_TQTMFSBSig \mid true$

$\bullet \text{pre } MutexTakeNnonInh_TQTM$

MutexTakeNinh_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$topReady! : TASK$

$running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$

$mut? \in mutex$

$running_task \in \text{dom } base_priority$

$QueueReceiveN_TQT[que? := mut?]$

$semaphore' = semaphore$

$mutex' = mutex$

$$\begin{aligned} \text{mutex_holder}' &= \text{mutex_holder} \oplus \{(mut? \mapsto \text{running_task})\} \\ \text{mutex_recursive}' &= \text{mutex_recursive} \oplus \\ &\quad \{(mut? \mapsto \text{mutex_recursive}(mut?) + 1)\} \\ &\exists \text{OriginalPrioData} \\ \text{release_mutex}' &= \{\text{running_task}\} \triangleleft \text{release_mutex} \end{aligned}$$

MutexTakeNInh_TQTMFSBSig

TaskQueueTimeMutex
mut? : QUEUE

QueueReceiveN_TQTFBSBSig[que? := mut?]
running_task \in dom *release_rcv* \Rightarrow *running_task* \in dom *release_mutex*
mut? \in *mutex*
running_task \in dom *base_priority*

theorem *MutexTakeNInh_TQTM_vc_ref*

\forall *MutexTakeNInh_TQTMFSBSig* | *true* • pre *MutexTakeNInh_TQTM*

MutexTakeEnonInh_TQTM

Δ *TaskQueueTimeMutex*
mut? : QUEUE
topReady! : TASK
wtime? : N

running_task \in dom *release_rcv* \Rightarrow *running_task* \in dom *release_mutex*
mut? \in dom *mutex_holder*
priority(running_task) \leq *priority(mutex_holder(mut?))*
running_task \neq *mutex_holder(mut?)*
QueueReceiveE_TQT[que? := mut?]
 \exists *MutexData*
 \exists *OriginalPrioData*
release_mutex' = {*running_task*} \triangleleft *release_mutex*

MutexTakeEnonInh_TQTMFSBSig

TaskQueueTimeMutex
mut? : QUEUE
wtime? : N

QueueReceiveE_TQTFBSBSig[que? := mut?]
running_task \in dom *release_rcv* \Rightarrow *running_task* \in dom *release_mutex*
mut? \in dom *mutex_holder*
priority(running_task) \leq *priority(mutex_holder(mut?))*
running_task \neq *mutex_holder(mut?)*

theorem lMutexTakeEnonInh_TQTM_Lemma
$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{mut?} : \text{QUEUE}; \text{topReady!} : \text{TASK}; \\
& \quad \text{wtime?} : \mathbb{N} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \\
& \quad \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \quad \Rightarrow \text{mut?} = \text{release_rcv}(\text{running_task})) \\
& \quad \wedge \text{mut?} \in \text{queue} \\
& \quad \wedge \text{q_size}(\text{mut?}) = 0 \\
& \quad \wedge \text{running_task} \neq \text{idle} \\
& \quad \wedge \text{topReady!} \in \text{state} \sim (\{ \text{ready} \} \mid) \\
& \quad \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \mid) \\
& \quad \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\
& \quad \wedge \text{wtime?} > \text{clock} \\
& \quad \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \quad \Rightarrow \text{running_task} \in \text{dom release_mutex}) \\
& \quad \wedge \text{mut?} \in \text{dom mutex_holder} \\
& \quad \wedge \text{q_size}(\text{mut?}) = 0 \\
& \quad \wedge \text{priority}(\text{running_task}) \leq \text{priority}(\text{mutex_holder}(\text{mut?})) \\
& \quad \wedge \text{running_task} \neq \text{mutex_holder}(\text{mut?}) \\
& \quad \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \quad \text{release_mutex} := \{\text{running_task}\} \triangleleft \text{release_mutex}, \\
& \quad \quad \text{release_rcv} := \{\text{running_task}\} \triangleleft \text{release_rcv}, \\
& \quad \quad \text{running_task} := \text{topReady!}, \\
& \quad \quad \text{state} := \text{state} \oplus \\
& \quad \quad \quad (\{(\text{running_task}, \text{blocked})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \quad \text{wait_time} := \text{wait_time} \oplus \{(\text{running_task}, \text{wtime?})\}, \\
& \quad \quad \text{wait_rcv} := \text{wait_rcv} \oplus \{(\text{running_task}, \text{mut?})\}] \\
& \quad \wedge (\text{st} \in \text{TASK} \\
& \quad \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\})) \text{st} = \text{state}(\text{st}) \\
& \quad \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \quad \quad \{(\text{topReady!}, \text{running})\}))) \text{st} \in \text{transition}) \\
& \quad \Rightarrow \text{t} \in \text{TASK} \\
& \quad \quad \wedge \text{state}(\text{t}) = \text{ready} \\
& \quad \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{t})
\end{aligned}$$
theorem MutexTakeEnonInh_TQTM_vc_ref
$$\begin{aligned}
& \forall \text{MutexTakeEnonInh_TQTMFSBSig} \mid \text{true} \\
& \quad \bullet \text{pre MutexTakeEnonInh_TQTM}
\end{aligned}$$

$$\text{MutexTakeEInheritReady_TQTM}$$

$$\Delta \text{TaskQueueTimeMutex}$$

$$\text{mut?} : \text{QUEUE}$$

$topReady! : TASK$

$wtime? : \mathbb{N}$

$running_task \notin \text{dom } release_snd$

$running_task \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(running_task)$

$running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$

$mut? \in \text{dom } mutex_holder$

$priority(running_task) > priority(mutex_holder(mut?))$

$wtime? > clock$

$mutex_holder(mut?) \notin state \sim (\{ready\})$

$state(topReady!) = ready$

$\forall rt : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(rt)$

$\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$

$| st? = blocked \wedge pri? = priority \oplus$

$\{(mutex_holder(mut?) \mapsto priority(running_task))\}$

$\bullet Reschedule[topReady!/target?, tasks/tasks?]$

$\exists QueueData$

$wait_snd' = wait_snd$

$wait_rcv' = wait_rcv \oplus \{(running_task \mapsto mut?)\}$

$release_snd' = release_snd$

$release_rcv' = \{running_task\} \triangleleft release_rcv$

$clock' = clock$

$delayed_task' = delayed_task$

$wait_time' = wait_time \oplus \{(running_task \mapsto wtime?)\}$

$time_slice' = time_slice$

$\exists MutexData$

$\exists OriginalPrioData$

$release_mutex' = \{running_task\} \triangleleft release_mutex$

$MutexTakeEInheritReady_TQTMFSBSig$

$TaskQueueTimeMutex$

$mut? : QUEUE$

$wtime? : \mathbb{N}$

$running_task \notin \text{dom } release_snd$

$running_task \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(running_task)$

$running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$

$mut? \in \text{dom } mutex_holder$

$priority(running_task) > priority(mutex_holder(mut?))$

$wtime? > clock$

$mutex_holder(mut?) \notin state \sim (\{ready\})$

theorem $!MutexTakeEInheritReady_TQTM_Lemma$

$\forall TaskQueueTimeMutex; mut? : QUEUE; topReady! : TASK;$

$wtime? : \mathbb{N}$

$$\begin{aligned}
& | \text{running_task} \notin \text{dom release_snd} \\
& \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \Rightarrow \text{mut?} = \text{release_rcv}(\text{running_task})) \\
& \wedge (\text{running_task} \in \text{dom release_rcv} \\
& \quad \Rightarrow \text{running_task} \in \text{dom release_mutex}) \\
& \wedge \text{mut?} \in \text{dom mutex_holder} \\
& \wedge \text{priority}(\text{running_task}) > \text{priority}(\text{mutex_holder}(\text{mut?})) \\
& \wedge \text{wtime?} > \text{clock} \\
& \wedge \text{mutex_holder}(\text{mut?}) \notin \text{state} \sim (\{ \text{ready} \}) \\
& \wedge \text{state}(\text{topReady!}) = \text{ready} \\
& \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \}) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk})) \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{running_task}, \text{phys_context})\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{priority} := \text{priority} \oplus \\
& \quad \{(\text{mutex_holder}(\text{mut?}), \text{priority}(\text{running_task}))\}, \\
& \quad \text{release_mutex} := \{ \text{running_task} \} \triangleleft \text{release_mutex}, \\
& \quad \text{release_rcv} := \{ \text{running_task} \} \triangleleft \text{release_rcv}, \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus \\
& \quad (\{(\text{running_task}, \text{blocked})\} \cup \{(\text{topReady!}, \text{running})\}), \\
& \quad \text{wait_time} := \text{wait_time} \oplus \{(\text{running_task}, \text{wtime?})\}, \\
& \quad \text{wait_rcv} := \text{wait_rcv} \oplus \{(\text{running_task}, \text{mut?})\}] \\
& \wedge (\text{st} \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\})) \text{st} = \text{state}(\text{st}) \\
& \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus (\{(\text{running_task}, \text{blocked})\} \cup \\
& \quad \{(\text{topReady!}, \text{running})\}))) \text{st} \in \text{transition}) \\
& \Rightarrow \text{rt} \in \text{TASK} \\
& \quad \wedge \text{state}(\text{rt}) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rt}))
\end{aligned}$$

theorem *MutexTakeEInheritReady_TQTM_vc_ref*
 $\forall \text{MutexTakeEInheritReady_TQTMFSBSig} \mid \text{true}$
 $\bullet \text{pre } \text{MutexTakeEInheritReady_TQTM}$

MutexTakeEInheritHolder_TQTM _____

$\Delta \text{TaskQueueTimeMutex}$

mut? : *QUEUE*

topReady! : *TASK*

wtime? : \mathbb{N}

$\text{running_task} \notin \text{dom release_snd}$

$\text{running_task} \in \text{dom release_rcv} \Rightarrow \text{mut?} = \text{release_rcv}(\text{running_task})$

$running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$
 $mut? \in \text{dom } mutex_holder$
 $priority(running_task) > priority(mutex_holder(mut?))$
 $wtime? > clock$
 $mutex_holder(mut?) \in state \sim (\{ready\})$
 $topReady! = mutex_holder(mut?)$
 $topReady! \neq idle$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $\quad | st? = blocked \wedge pri? = priority \oplus$
 $\quad \quad \{ (topReady! \mapsto priority(running_task)) \}$
 $\quad \bullet Reschedule[topReady!/target?, tasks/tasks?]$
 $\exists QueueData$
 $wait_snd' = wait_snd$
 $wait_rcv' = wait_rcv \oplus \{ (running_task \mapsto mut?) \}$
 $release_snd' = release_snd$
 $release_rcv' = \{ running_task \} \triangleleft release_rcv$
 $clock' = clock$
 $delayed_task' = delayed_task$
 $wait_time' = wait_time \oplus \{ (running_task \mapsto wtime?) \}$
 $time_slice' = time_slice$
 $\exists MutexData$
 $\exists OriginalPrioData$
 $release_mutex' = \{ running_task \} \triangleleft release_mutex$

MutexTakeEInheritHolder_TQTMFSBSig

TaskQueueTimeMutex

$mut? : QUEUE$

$wtime? : \mathbb{N}$

$running_task \notin \text{dom } release_snd$
 $running_task \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(running_task)$
 $running_task \in \text{dom } release_rcv \Rightarrow running_task \in \text{dom } release_mutex$
 $mut? \in \text{dom } mutex_holder$
 $priority(running_task) > priority(mutex_holder(mut?))$
 $wtime? > clock$
 $mutex_holder(mut?) \neq idle$
 $mutex_holder(mut?) \in state \sim (\{ready\})$

theorem *MutexTakeEInheritHolder_TQTM_vc_ref*

$\forall MutexTakeEInheritHolder_TQTMFSBSig \mid true$

$\bullet \text{pre } MutexTakeEInheritHolder_TQTM$

MutexTakeRecursive_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$
 $topReady! : TASK$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $running_task = mutex_holder(mut?)$
 $\exists TaskQueueTime$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus$
 $\{(mut? \mapsto mutex_recursive(mut?) + 1)\}$
 $\exists OriginalPrioData$
 $\exists MReleasingData$
 $topReady! = running_task$

$MutexTakeRecursive_TQTMFSBSig$

$TaskQueueTimeMutex$
 $mut? : QUEUE$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $running_task = mutex_holder(mut?)$

theorem $MutexTakeRecursive_TQTM_vc_ref$

$\forall MutexTakeRecursive_TQTMFSBSig \mid true$

- pre $MutexTakeRecursive_TQTM$

$MutexTake_TQTM \cong MutexTakeNnonInh_TQTM$
 $\vee MutexTakeNInh_TQTM$
 $\vee MutexTakeEnonInh_TQTM$
 $\vee MutexTakeEInheritReady_TQTM$
 $\vee MutexTakeEInheritHolder_TQTM$
 $\vee MutexTakeRecursive_TQTM$

$basePriorityMan$

$\Delta TaskQueueTimeMutex$
 $mut? : QUEUE$

$running_task \in \text{ran}(\{mut?\} \triangleleft mutex_holder) \Rightarrow \exists OriginalPrioData$
 $running_task \notin \text{ran}(\{mut?\} \triangleleft mutex_holder)$
 $\Rightarrow base_priority' = \{running_task\} \triangleleft base_priority$

MutexGiveNnonInh_TQTM

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

topReady! : *TASK*

running_task \notin $\text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

mut? \in $\text{dom } \text{mutex_holder}$

running_task = *mutex_holder*(*mut?*)

mutex_recursive(*mut?*) = 1

base_priority(*running_task*) = *priority*(*running_task*)

QueueSendN_TQT[*que?* := *mut?*]

semaphore' = *semaphore*

mutex' = *mutex*

mutex_holder' = {*mut?*} \triangleleft *mutex_holder*

mutex_recursive' = *mutex_recursive* \oplus {(*mut?* \mapsto 0)}

basePriorityMan

\exists *MReleasingData*

MutexGiveNnonInh_TQTMFSBSig

TaskQueueTimeMutex

mut? : *QUEUE*

*QueueSendN_TQTF**FSBSig*[*que?* := *mut?*]

running_task \notin $\text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

mut? \in $\text{dom } \text{mutex_holder}$

running_task = *mutex_holder*(*mut?*)

mutex_recursive(*mut?*) = 1

base_priority(*running_task*) = *priority*(*running_task*)

theorem *ranUnchanged* [*X*, *Y*]

$\forall f : X \mapsto Y; a : X \mid a \in \text{dom } f \wedge f(a) \in \text{ran}(\{a\} \triangleleft f)$

- $\text{ran } f = \text{ran}(\{a\} \triangleleft f)$

theorem *MutexGiveNnonInh_TQTM_vc_ref*

\forall *MutexGiveNnonInh_TQTMFSBSig* | *true*

- *pre* *MutexGiveNnonInh_TQTM*

MutexGiveNinhN_TQTM

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

topReady! : *TASK*

running_task \notin $\text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

mut? \in $\text{dom } \text{mutex_holder}$

$running_task = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $mut? \notin \text{ran } wait_rcv$
 $base_priority(running_task) \neq priority(running_task)$
 $\forall rt : state \sim (\{ready\}) \bullet base_priority(running_task) \geq priority(rt)$
 $\exists TaskData$
 $\exists StateData$
 $\exists ContextData$
 $priority' = priority \oplus \{(running_task \mapsto base_priority(running_task))\}$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(mut? \mapsto 1)\}$
 $\exists WaitingData$
 $\exists QReleasingData$
 $\exists Time$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $basePriorityMan$
 $\exists MReleasingData$
 $topReady! = running_task$

MutexGiveNInhN_TQTMFSBSig

TaskQueueTimeMutex

$mut? : QUEUE$

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $running_task = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $mut? \notin \text{ran } wait_rcv$
 $base_priority(running_task) \neq priority(running_task)$
 $\forall rt : state \sim (\{ready\}) \bullet base_priority(running_task) \geq priority(rt)$

theorem *MutexGiveNInhN_TQTMF_vc_ref*

$\forall MutexGiveNInhN_TQTMFSBSig \mid true$

• pre *MutexGiveNInhN_TQTM*

MutexGiveNInhS_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$topReady! : TASK$

$$\begin{aligned}
& \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \text{mut?} \in \text{dom mutex_holder} \\
& \text{running_task} = \text{mutex_holder}(\text{mut?}) \\
& \text{mutex_recursive}(\text{mut?}) = 1 \\
& \text{mut?} \notin \text{ran wait_rcv} \\
& \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \text{state}(\text{topReady!}) = \text{ready} \\
& \forall \text{rt} : \text{state} \sim (\{ \text{ready} \}) \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rt}) \\
& \text{base_priority}(\text{running_task}) < \text{priority}(\text{topReady!}) \\
& \exists \text{st?} : \text{STATE}; \text{pri?} : \text{TASK} \rightarrow \mathbb{N} \\
& \quad | \text{st?} = \text{ready} \wedge \text{pri?} = \text{priority} \oplus \\
& \quad \quad \{ (\text{running_task} \mapsto \text{base_priority}(\text{running_task})) \} \\
& \quad \bullet \text{Reschedule}[\text{topReady!}/\text{target?}, \text{tasks}/\text{tasks?}] \\
& \text{queue}' = \text{queue} \\
& \text{q_max}' = \text{q_max} \\
& \text{q_size}' = \text{q_size} \oplus \{ (\text{mut?} \mapsto 1) \} \\
& \exists \text{WaitingData} \\
& \exists \text{QReleasingData} \\
& \exists \text{Time} \\
& \text{semaphore}' = \text{semaphore} \\
& \text{mutex}' = \text{mutex} \\
& \text{mutex_holder}' = \{ \text{mut?} \} \triangleleft \text{mutex_holder} \\
& \text{mutex_recursive}' = \text{mutex_recursive} \oplus \{ (\text{mut?} \mapsto 0) \} \\
& \text{basePriorityMan} \\
& \exists \text{MReleasingData}
\end{aligned}$$

MutexGiveNInhS_TQTMFSBSig _____

TaskQueueTimeMutex
mut? : QUEUE

$$\begin{aligned}
& \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \text{mut?} \in \text{dom mutex_holder} \\
& \text{running_task} = \text{mutex_holder}(\text{mut?}) \\
& \text{mutex_recursive}(\text{mut?}) = 1 \\
& \text{mut?} \notin \text{ran wait_rcv} \\
& \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \exists \text{topReady!} : \text{state} \sim (\{ \text{ready} \}) \\
& \quad | \forall \text{rt} : \text{state} \sim (\{ \text{ready} \}) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rt}) \\
& \quad \bullet \text{base_priority}(\text{running_task}) < \text{priority}(\text{topReady!})
\end{aligned}$$

theorem *IMutexGiveNInhS_TQTM_Lemma*

$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{mut?} : \text{QUEUE}; \text{topReady!} : \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \quad \wedge \text{mut?} \in \text{dom mutex_holder}
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{running_task} = \text{mutex_holder}(\text{mut}?) \\
& \wedge \text{mutex_recursive}(\text{mut}?) = 1 \\
& \wedge \text{mut}? \notin \text{ran wait_rcv} \\
& \wedge \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \wedge \text{state}(\text{topReady}!) = \text{ready} \\
& \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \mid)) \\
& \quad \bullet \text{priority}(\text{topReady}!) \geq \text{priority}(\text{rtsk}) \\
& \wedge \text{base_priority}(\text{running_task}) < \text{priority}(\text{topReady}!) \\
& \wedge \text{running_task} \in \text{ran}(\{ \text{mut}?\} \triangleleft \text{mutex_holder}) \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{mutex_holder}(\text{mut}?), \text{phys_context})\}, \\
& \quad \text{mutex_holder} := \{ \text{mut}?\} \triangleleft \text{mutex_holder}, \\
& \quad \text{mutex_recursive} := \text{mutex_recursive} \oplus \{(\text{mut}?, 0)\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady}!), \\
& \quad \text{priority} := \text{priority} \oplus \{(\text{mutex_holder}(\text{mut}?), \\
& \quad \quad \text{base_priority}(\text{mutex_holder}(\text{mut}?)\})\}, \\
& \quad \text{q_size} := \text{q_size} \oplus \{(\text{mut}?, 1)\}, \\
& \quad \text{running_task} := \text{topReady}!, \\
& \quad \text{state} := \text{state} \oplus (\{(\text{mutex_holder}(\text{mut}?), \text{ready})\} \cup \\
& \quad \quad \{(\text{topReady}!, \text{running})\}) \\
& \wedge (\text{st} \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{mutex_holder}(\text{mut}?), \text{ready})\} \cup \\
& \quad \quad \{(\text{topReady}!, \text{running})\})) \text{st} = \text{state}(\text{st}) \\
& \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus \\
& \quad \quad \{(\text{mutex_holder}(\text{mut}?), \text{ready})\} \cup \\
& \quad \quad \{(\text{topReady}!, \text{running})\})) \text{st} \\
& \quad \quad \in \text{transition}) \\
& \Rightarrow \text{rt} \in \text{TASK} \\
& \quad \wedge \text{state}(\text{rt}) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady}!) \geq \text{priority}(\text{rt}))
\end{aligned}$$

theorem lMutexGiveNInhS_TQTM_Lemma1

$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{mut}?: \text{QUEUE}; \text{topReady}!: \text{TASK} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \quad \wedge \text{mut}? \in \text{dom mutex_holder} \\
& \quad \wedge \text{running_task} = \text{mutex_holder}(\text{mut}?) \\
& \quad \wedge \text{mutex_recursive}(\text{mut}?) = 1 \\
& \quad \wedge \text{mut}? \notin \text{ran wait_rcv} \\
& \quad \wedge \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \quad \wedge \text{state}(\text{topReady}!) = \text{ready} \\
& \quad \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} \mid)) \\
& \quad \quad \bullet \text{priority}(\text{topReady}!) \geq \text{priority}(\text{rtsk}) \\
& \quad \wedge \text{base_priority}(\text{running_task}) < \text{priority}(\text{topReady}!) \\
& \quad \wedge \text{running_task} \notin \text{ran}(\{ \text{mut}?\} \triangleleft \text{mutex_holder}) \\
& \quad \bullet \neg (\text{TaskQueueTimeMutex}[\\
& \quad \quad \text{base_priority} := \{ \text{mutex_holder}(\text{mut}?)\} \triangleleft \text{base_priority},
\end{aligned}$$

$$\begin{aligned}
& \log_context := \log_context \oplus \\
& \quad \{(mutex_holder(mut?), phys_context)\}, \\
& mutex_holder := \{mut?\} \triangleleft mutex_holder, \\
& mutex_recursive := mutex_recursive \oplus \{(mut?, 0)\}, \\
& phys_context := \log_context(topReady!), \\
& priority := priority \oplus \{(mutex_holder(mut?), \\
& \quad base_priority(mutex_holder(mut?)))\}, \\
& q_size := q_size \oplus \{(mut?, 1)\}, \\
& running_task := topReady!, \\
& state := state \oplus (\{(mutex_holder(mut?), ready)\} \cup \\
& \quad \{(topReady!, running)\}) \\
& \wedge (st \in TASK \\
& \quad \wedge \neg (state \oplus (\{(mutex_holder(mut?), ready)\} \cup \\
& \quad \{(topReady!, running)\})) st = state(st) \\
& \quad \Rightarrow (state(st), (state \oplus \\
& \quad \quad (\{(mutex_holder(mut?), ready)\} \cup \\
& \quad \quad \{(topReady!, running)\})) st) \in transition) \\
& \Rightarrow rt \in TASK \\
& \quad \wedge state(rt) = ready \\
& \quad \wedge \neg priority(topReady!) \geq priority(rt)
\end{aligned}$$

theorem MutexGiveNInhS_TQTM_vc_ref

\forall MutexGiveNInhS_TQTMFSBSig | true • pre MutexGiveNInhS_TQTM

MutexGiveWnonInhN_TQTM _____

Δ TaskQueueTimeMutex

mut? : QUEUE

topReady! : TASK

running_task \notin dom *release_snd* \cup dom *release_rcv*

mut? \in dom *mutex_holder*

running_task = *mutex_holder*(*mut?*)

mutex_recursive(*mut?*) = 1

base_priority(*running_task*) = *priority*(*running_task*)

QueueSendW_TQT[*que?* := *mut?*]

semaphore' = *semaphore*

mutex' = *mutex*

mutex_holder' = $\{mut?\} \triangleleft mutex_holder$

mutex_recursive' = *mutex_recursive* \oplus $\{(mut? \mapsto 0)\}$

basePriorityMan

release_mutex' = *release_mutex* \oplus $\{(topReady! \mapsto mut?)\}$

_____ *MutexGiveWnonInhN_TQTMFSBSig* _____

TaskQueueTimeMutex

$$mut? : QUEUE$$

$$\begin{aligned}
& QueueSendW_TQTFSSig[que? := mut?] \\
& running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv \\
& mut? \in \text{dom } mutex_holder \\
& running_task = mutex_holder(mut?) \\
& mutex_recursive(mut?) = 1 \\
& base_priority(running_task) = priority(running_task)
\end{aligned}$$

theorem lMutexGiveWnonInhN_TQTM_Lemma

$$\begin{aligned}
& \forall TaskQueueTimeMutex; topReady! : TASK; mut? : QUEUE \\
& \quad | running_task \notin \text{dom } release_rcv \cup \text{dom } release_snd \\
& \quad \wedge mut? \in queue \\
& \quad \wedge q_size(mut?) < q_max(mut?) \\
& \quad \wedge topReady! \in wait_rcv \sim (\{mut?\}) \\
& \quad \wedge (\forall wrct : wait_rcv \sim (\{mut?\}) \\
& \quad \quad \bullet priority(topReady!) \geq priority(wrct)) \\
& \quad \wedge priority(running_task) \geq priority(topReady!) \\
& \quad \wedge mut? \in \text{dom } mutex_holder \\
& \quad \wedge running_task = mutex_holder(mut?) \\
& \quad \wedge mutex_recursive(mut?) = 1 \\
& \quad \wedge base_priority(running_task) = priority(running_task) \\
& \quad \wedge running_task \in \text{ran}(\{mut?\} \triangleleft mutex_holder) \\
& \quad \bullet \neg (TaskQueueTimeMutex[\\
& \quad \quad mutex_holder := \{mut?\} \triangleleft mutex_holder, \\
& \quad \quad mutex_recursive := mutex_recursive \oplus \{(mut?, 0)\}, \\
& \quad \quad q_size := q_size \oplus \{(mut?, (1 + q_size(mut?))\}, \\
& \quad \quad release_mutex := release_mutex \oplus \{(topReady!, mut?)\}, \\
& \quad \quad release_rcv := release_rcv \oplus \{(topReady!, mut?)\}, \\
& \quad \quad release_snd := \{mutex_holder(mut?)\} \triangleleft release_snd, \\
& \quad \quad running_task := mutex_holder(mut?), \\
& \quad \quad state := state \oplus \{(topReady!, ready)\}, \\
& \quad \quad wait_time := \{topReady!\} \triangleleft wait_time, \\
& \quad \quad wait_rcv := \{topReady!\} \triangleleft wait_rcv] \\
& \quad \wedge priority(topReady!) \leq priority(mutex_holder(mut?)) \\
& \quad \wedge (st \in TASK \\
& \quad \quad \wedge \neg (state \oplus \{(topReady!, ready)\})st = state(st) \\
& \quad \quad \Rightarrow (state(st), (state \oplus \{(topReady!, ready)\})st) \\
& \quad \quad \in transition) \\
& \quad \Rightarrow wr \in \text{dom } wait_rcv \\
& \quad \quad \wedge wait_rcv(wr) = mut? \\
& \quad \quad \wedge \neg priority(topReady!) \geq priority(wr))
\end{aligned}$$

theorem lMutexGiveWnonInhN_TQTM_Lemma1

$$\forall TaskQueueTimeMutex; topReady! : TASK; mut? : QUEUE$$

$$\begin{aligned}
& | \text{running_task} \notin \text{dom release_rcv} \cup \text{dom release_snd} \\
& \wedge \text{mut?} \in \text{queue} \\
& \wedge q_size(\text{mut?}) < q_max(\text{mut?}) \\
& \wedge \text{topReady!} \in \text{wait_rcv} \sim (\{ \text{mut?} \}) \\
& \wedge (\forall wrct : \text{wait_rcv} \sim (\{ \text{mut?} \}) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(wrct)) \\
& \wedge \text{priority}(\text{running_task}) \geq \text{priority}(\text{topReady!}) \\
& \wedge \text{mut?} \in \text{dom mutex_holder} \\
& \wedge \text{running_task} = \text{mutex_holder}(\text{mut?}) \\
& \wedge \text{mutex_recursive}(\text{mut?}) = 1 \\
& \wedge \text{base_priority}(\text{running_task}) = \text{priority}(\text{running_task}) \\
& \wedge \text{running_task} \notin \text{ran}(\{ \text{mut?} \}) \triangleleft \text{mutex_holder} \\
& \bullet \neg (\text{TaskQueueTimeMutex} [\\
& \quad \text{base_priority} := \{ \text{mutex_holder}(\text{mut?}) \} \triangleleft \text{base_priority}, \\
& \quad \text{mutex_holder} := \{ \text{mut?} \} \triangleleft \text{mutex_holder}, \\
& \quad \text{mutex_recursive} := \text{mutex_recursive} \oplus \{ (\text{mut?}, 0) \}, \\
& \quad \text{q_size} := \text{q_size} \oplus \{ (\text{mut?}, (1 + \text{q_size}(\text{mut?}))) \}, \\
& \quad \text{release_mutex} := \text{release_mutex} \oplus \{ (\text{topReady!}, \text{mut?}) \}, \\
& \quad \text{release_rcv} := \text{release_rcv} \oplus \{ (\text{topReady!}, \text{mut?}) \}, \\
& \quad \text{release_snd} := \{ \text{mutex_holder}(\text{mut?}) \} \triangleleft \text{release_snd}, \\
& \quad \text{running_task} := \text{mutex_holder}(\text{mut?}), \\
& \quad \text{state} := \text{state} \oplus \{ (\text{topReady!}, \text{ready}) \}, \\
& \quad \text{wait_time} := \{ \text{topReady!} \} \triangleleft \text{wait_time}, \\
& \quad \text{wait_rcv} := \{ \text{topReady!} \} \triangleleft \text{wait_rcv}] \\
& \wedge \text{priority}(\text{topReady!}) \leq \text{priority}(\text{mutex_holder}(\text{mut?})) \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus \{ (\text{topReady!}, \text{ready}) \}) st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus \{ (\text{topReady!}, \text{ready}) \}) st) \\
& \quad \in \text{transition}) \\
& \Rightarrow wr \in \text{dom wait_rcv} \\
& \quad \wedge \text{wait_rcv}(wr) = \text{mut?} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(wr))
\end{aligned}$$

theorem *MutexGiveWnonInhN_TQTM_vc_ref*
 $\forall \text{MutexGiveWnonInhN_TQTMFSBSig} \mid \text{true}$
 $\bullet \text{pre } \text{MutexGiveWnonInhN_TQTM}$

MutexGiveWnonInhS_TQTM

$\Delta \text{TaskQueueTimeMutex}$
 $\text{mut?} : \text{QUEUE}$
 $\text{topReady!} : \text{TASK}$

$\text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$
 $\text{mut?} \in \text{dom mutex_holder}$
 $\text{running_task} = \text{mutex_holder}(\text{mut?})$

$$\begin{aligned}
& \text{mutex_recursive}(\text{mut}?) = 1 \\
& \text{base_priority}(\text{running_task}) = \text{priority}(\text{running_task}) \\
& \text{QueueSendWS_TQT}[\text{que}? := \text{mut}?] \\
& \text{semaphore}' = \text{semaphore} \\
& \text{mutex}' = \text{mutex} \\
& \text{mutex_holder}' = \{\text{mut}?\} \triangleleft \text{mutex_holder} \\
& \text{mutex_recursive}' = \text{mutex_recursive} \oplus \{(\text{mut}? \mapsto 0)\} \\
& \text{basePriorityMan} \\
& \text{release_mutex}' = \text{release_mutex} \oplus \{(\text{topReady!} \mapsto \text{mut}?)\}
\end{aligned}$$

MutexGiveWnonInhS_TQTMFSBSig

TaskQueueTimeMutex
 $\text{mut}?: \text{QUEUE}$

$$\begin{aligned}
& \text{QueueSendWS_TQTFBSBSig}[\text{que}? := \text{mut}?] \\
& \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \text{mut}? \in \text{dom mutex_holder} \\
& \text{running_task} = \text{mutex_holder}(\text{mut}?) \\
& \text{mutex_recursive}(\text{mut}?) = 1 \\
& \text{base_priority}(\text{running_task}) = \text{priority}(\text{running_task})
\end{aligned}$$

theorem lMutexGiveWnonInhS_TQTM_Lemma

$\forall \text{TaskQueueTimeMutex}; \text{topReady!} : \text{TASK}; \text{mut}?: \text{QUEUE}$

$$\begin{aligned}
& | \text{running_task} \notin \text{dom release_rcv} \\
& \wedge (\text{running_task} \in \text{dom release_snd} \\
& \quad \Rightarrow \text{mut}? = \text{release_snd}(\text{running_task})) \\
& \wedge \text{mut}? \in \text{queue} \\
& \wedge \text{q_size}(\text{mut}?) < \text{q_max}(\text{mut}?) \\
& \wedge \text{topReady!} \in \text{wait_rcv} \sim (\{\text{mut}?\}) \\
& \wedge (\forall \text{wrct} : \text{wait_rcv} \sim (\{\text{mut}?\}) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{wrct})) \\
& \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{running_task}) \\
& \wedge \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \wedge \text{mut}? \in \text{dom mutex_holder} \\
& \wedge \text{running_task} = \text{mutex_holder}(\text{mut}?) \\
& \wedge \text{mutex_recursive}(\text{mut}?) = 1 \\
& \wedge \text{base_priority}(\text{running_task}) = \text{priority}(\text{running_task}) \\
& \wedge \text{running_task} \in \text{ran}(\{\text{mut}?\} \triangleleft \text{mutex_holder}) \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{mutex_holder}(\text{wait_rcv}(\text{topReady!})), \text{phys_context})\}, \\
& \quad \text{mutex_holder} := \{\text{wait_rcv}(\text{topReady!})\} \triangleleft \text{mutex_holder}, \\
& \quad \text{mutex_recursive} := \text{mutex_recursive} \oplus \\
& \quad \{(\text{wait_rcv}(\text{topReady!}), 0)\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{q_size} := \text{q_size} \oplus \{(\text{wait_rcv}(\text{topReady!}), (1+
\end{aligned}$$

$$\begin{aligned}
& q_size(wait_rcv(topReady!))), \\
& release_mutex := release_mutex \oplus \\
& \quad \{(topReady!, wait_rcv(topReady!))\}, \\
& release_rcv := release_rcv \oplus \\
& \quad \{(topReady!, wait_rcv(topReady!))\}, \\
& release_snd := \{mutex_holder(wait_rcv(topReady!))\} \\
& \quad \triangleleft release_snd, \\
& running_task := topReady!, \\
& state := state \oplus \\
& \quad (\{(mutex_holder(wait_rcv(topReady!)), ready)\} \cup \\
& \quad \quad \{(topReady!, running)\}), \\
& wait_time := \{topReady!\} \triangleleft wait_time, \\
& wait_rcv := \{topReady!\} \triangleleft wait_rcv \\
& \wedge (st \in TASK \\
& \quad \wedge \neg (state \oplus \\
& \quad \quad (\{(mutex_holder(wait_rcv(topReady!)), ready)\} \cup \\
& \quad \quad \quad \{(topReady!, running)\})) st = state(st) \\
& \quad \Rightarrow (state(st), (state \oplus \\
& \quad \quad (\{(mutex_holder(wait_rcv(topReady!)), ready)\} \cup \\
& \quad \quad \quad \{(topReady!, running)\})) st) \in transition) \\
& \Rightarrow wr \in \text{dom } wait_rcv \\
& \quad \wedge wait_rcv(wr) = wait_rcv(topReady!) \\
& \quad \wedge \neg priority(topReady!) \geq priority(wr))
\end{aligned}$$

theorem lMutexGiveWnonInhS_TQTM_Lemma1

$$\begin{aligned}
& \forall TaskQueueTimeMutex; topReady! : TASK; mut? : QUEUE \\
& \quad | running_task \notin \text{dom } release_rcv \\
& \quad \wedge (running_task \in \text{dom } release_snd \\
& \quad \quad \Rightarrow mut? = release_snd(running_task)) \\
& \quad \wedge mut? \in queue \\
& \quad \wedge q_size(mut?) < q_max(mut?) \\
& \quad \wedge topReady! \in wait_rcv \sim (\{mut?\}) \\
& \quad \wedge (\forall wrct : wait_rcv \sim (\{mut?\}) \\
& \quad \quad \bullet priority(topReady!) \geq priority(wrct)) \\
& \quad \wedge priority(topReady!) > priority(running_task) \\
& \quad \wedge running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv \\
& \quad \wedge mut? \in \text{dom } mutex_holder \\
& \quad \wedge running_task = mutex_holder(mut?) \\
& \quad \wedge mutex_recursive(mut?) = 1 \\
& \quad \wedge base_priority(running_task) = priority(running_task) \\
& \quad \wedge running_task \notin \text{ran}(\{mut?\} \triangleleft mutex_holder) \\
& \quad \bullet \neg (TaskQueueTimeMutex [\\
& \quad \quad base_priority := \{mutex_holder(wait_rcv(topReady!))\} \\
& \quad \quad \triangleleft base_priority, \\
& \quad \quad log_context := log_context \oplus \\
& \quad \quad \quad \{(mutex_holder(wait_rcv(topReady!)), phys_context)\},
\end{aligned}$$

$$\begin{aligned}
& \text{mutex_holder} := \{ \text{wait_rcv}(\text{topReady!}) \} \triangleleft \text{mutex_holder}, \\
& \text{mutex_recursive} := \text{mutex_recursive} \oplus \\
& \quad \{ (\text{wait_rcv}(\text{topReady!}), 0) \}, \\
& \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \text{q_size} := \text{q_size} \oplus \{ (\text{wait_rcv}(\text{topReady!}), (1 + \\
& \quad \text{q_size}(\text{wait_rcv}(\text{topReady!}))) \}, \\
& \text{release_mutex} := \text{release_mutex} \oplus \\
& \quad \{ (\text{topReady!}, \text{wait_rcv}(\text{topReady!})) \}, \\
& \text{release_rcv} := \text{release_rcv} \oplus \\
& \quad \{ (\text{topReady!}, \text{wait_rcv}(\text{topReady!})) \}, \\
& \text{release_snd} := \{ \text{mutex_holder}(\text{wait_rcv}(\text{topReady!})) \} \\
& \quad \triangleleft \text{release_snd}, \\
& \text{running_task} := \text{topReady!}, \\
& \text{state} := \text{state} \oplus \\
& \quad (\{ (\text{mutex_holder}(\text{wait_rcv}(\text{topReady!})), \text{ready}) \} \cup \\
& \quad \quad \{ (\text{topReady!}, \text{running}) \}), \\
& \text{wait_time} := \{ \text{topReady!} \} \triangleleft \text{wait_time}, \\
& \text{wait_rcv} := \{ \text{topReady!} \} \triangleleft \text{wait_rcv} \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus \\
& \quad \quad (\{ (\text{mutex_holder}(\text{wait_rcv}(\text{topReady!})), \text{ready}) \} \cup \\
& \quad \quad \quad \{ (\text{topReady!}, \text{running}) \})) st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus \\
& \quad \quad (\{ (\text{mutex_holder}(\text{wait_rcv}(\text{topReady!})), \text{ready}) \} \cup \\
& \quad \quad \quad \{ (\text{topReady!}, \text{running}) \})) st) \in \text{transition}) \\
& \Rightarrow wr \in \text{dom } \text{wait_rcv} \\
& \quad \wedge \text{wait_rcv}(wr) = \text{wait_rcv}(\text{topReady!}) \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(wr)
\end{aligned}$$

theorem *MutexGiveWnonInhS_TQTM_vc_ref*

$\forall \text{MutexGiveWnonInhS_TQTMFSBSig} \mid \text{true}$

- pre *MutexGiveWnonInhS_TQTM*

MutexGiveWinhN_TQTM

$\Delta \text{TaskQueueTimeMutex}$

$\text{mut?} : \text{QUEUE}$

$\text{topReady!} : \text{TASK}$

$\text{running_task} \notin \text{dom } \text{release_snd} \cup \text{dom } \text{release_rcv}$

$\text{mut?} \in \text{dom } \text{mutex_holder}$

$\text{running_task} = \text{mutex_holder}(\text{mut?})$

$\text{mutex_recursive}(\text{mut?}) = 1$

$\text{topReady!} \in \text{wait_rcv} \sim (\{ \text{mut?} \} \mid)$

$\forall wr : \text{wait_rcv} \sim (\{ \text{mut?} \} \mid) \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(wr)$

$\text{priority}(\text{topReady!}) \leq \text{base_priority}(\text{running_task})$

$$\begin{aligned}
& \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \forall rt : \text{state} \sim (\{ \text{ready} \} \downarrow) \bullet \text{base_priority}(\text{running_task}) \geq \text{priority}(rt) \\
& \exists \text{TaskData} \\
& \text{state}' = \text{state} \oplus \{ (\text{topReady!} \mapsto \text{ready}) \} \\
& \exists \text{ContextData} \\
& \text{priority}' = \text{priority} \oplus \{ (\text{running_task} \mapsto \text{base_priority}(\text{running_task})) \} \\
& \text{queue}' = \text{queue} \\
& \text{q_max}' = \text{q_max} \\
& \text{q_size}' = \text{q_size} \oplus \{ (\text{mut?} \mapsto 1) \} \\
& \text{wait_snd}' = \text{wait_snd} \\
& \text{wait_rcv}' = \{ \text{topReady!} \} \triangleleft \text{wait_rcv} \\
& \text{release_snd}' = \text{release_snd} \\
& \text{release_rcv}' = \text{release_rcv} \oplus \{ (\text{topReady!} \mapsto \text{mut?}) \} \\
& \text{clock}' = \text{clock} \\
& \text{delayed_task}' = \text{delayed_task} \\
& \text{wait_time}' = \{ \text{topReady!} \} \triangleleft \text{wait_time} \\
& \text{time_slice}' = \text{time_slice} \\
& \text{semaphore}' = \text{semaphore} \\
& \text{mutex}' = \text{mutex} \\
& \text{mutex_holder}' = \{ \text{mut?} \} \triangleleft \text{mutex_holder} \\
& \text{mutex_recursive}' = \text{mutex_recursive} \oplus \{ (\text{mut?} \mapsto 0) \} \\
& \text{basePriorityMan} \\
& \text{release_mutex}' = \text{release_mutex} \oplus \{ (\text{topReady!} \mapsto \text{mut?}) \}
\end{aligned}$$

MutexGiveWinhN_TQTMFSBSig

TaskQueueTimeMutex

mut? : QUEUE

$$\begin{aligned}
& \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \text{mut?} \in \text{dom mutex_holder} \\
& \text{running_task} = \text{mutex_holder}(\text{mut?}) \\
& \text{mutex_recursive}(\text{mut?}) = 1 \\
& \forall wr : \text{wait_rcv} \sim (\{ \text{mut?} \} \downarrow) \\
& \quad \bullet \text{base_priority}(\text{running_task}) \geq \text{priority}(wr) \\
& \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \forall rt : \text{state} \sim (\{ \text{ready} \} \downarrow) \bullet \text{base_priority}(\text{running_task}) \geq \text{priority}(rt)
\end{aligned}$$

theorem *lMutexGiveWinhN_TQTM_Lemma*

$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{topReady!} : \text{TASK}; \text{mut?} : \text{QUEUE} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \quad \wedge \text{mut?} \in \text{dom mutex_holder} \\
& \quad \wedge \text{running_task} = \text{mutex_holder}(\text{mut?}) \\
& \quad \wedge \text{mutex_recursive}(\text{mut?}) = 1 \\
& \quad \wedge \text{topReady!} \in \text{wait_rcv} \sim (\{ \text{mut?} \} \downarrow) \\
& \quad \wedge (\forall wrct : \text{wait_rcv} \sim (\{ \text{mut?} \} \downarrow)
\end{aligned}$$

- $priority(topReady!) \geq priority(wrct)$

$$\wedge priority(topReady!) \leq base_priority(running_task)$$

$$\wedge base_priority(running_task) \neq priority(running_task)$$

$$\wedge (\forall rtsk : state \sim (\{ready\}) \mid)$$

- $base_priority(running_task) \geq priority(rtsk)$

$$\wedge running_task \in ran(\{mut?\} \triangleleft mutex_holder)$$

- $\neg (TaskQueueTimeMutex[$
 - $mutex_holder := \{mut?\} \triangleleft mutex_holder,$
 - $mutex_recursive := mutex_recursive \oplus \{(mut?, 0)\},$
 - $priority := priority \oplus \{(mutex_holder(mut?),$
 - $base_priority(mutex_holder(mut?))\},$
 - $q_size := q_size \oplus \{(mut?, 1)\},$
 - $release_mutex := release_mutex \oplus \{(topReady!, mut?)\},$
 - $release_rcv := release_rcv \oplus \{(topReady!, mut?)\},$
 - $running_task := mutex_holder(mut?),$
 - $state := state \oplus \{(topReady!, ready)\},$
 - $wait_time := \{topReady!\} \triangleleft wait_time,$
 - $wait_rcv := \{topReady!\} \triangleleft wait_rcv]$
$$\wedge priority(topReady!) \leq base_priority(mutex_holder(mut?))$$

$$\wedge (st \in TASK$$
 - $\wedge \neg (state \oplus \{(topReady!, ready)\})st = state(st)$
 - $\Rightarrow (state(st), (state \oplus \{(topReady!, ready)\})st$
 - $\in transition)$
$$\Rightarrow wr \in dom wait_rcv$$
 - $\wedge wait_rcv(wr) = mut?$
 - $\wedge \neg priority(topReady!) \geq priority(wr))$

theorem lMutexGiveWinhN_TQTM_Lemma1

$$\forall TaskQueueTimeMutex; topReady! : TASK; mut? : QUEUE$$

$$\mid running_task \notin dom release_snd \cup dom release_rcv$$

$$\wedge mut? \in dom mutex_holder$$

$$\wedge running_task = mutex_holder(mut?)$$

$$\wedge mutex_recursive(mut?) = 1$$

$$\wedge topReady! \in wait_rcv \sim (\{mut?\} \mid)$$

$$\wedge (\forall wrct : wait_rcv \sim (\{mut?\} \mid)$$

- $priority(topReady!) \geq priority(wrct)$

$$\wedge priority(topReady!) \leq base_priority(running_task)$$

$$\wedge base_priority(running_task) \neq priority(running_task)$$

$$\wedge (\forall rtsk : state \sim (\{ready\}) \mid)$$

- $base_priority(running_task) \geq priority(rtsk)$

$$\wedge running_task \notin ran(\{mut?\} \triangleleft mutex_holder)$$

- $\neg (TaskQueueTimeMutex[$
 - $base_priority := \{mutex_holder(mut?)\} \triangleleft base_priority,$
 - $mutex_holder := \{mut?\} \triangleleft mutex_holder,$
 - $mutex_recursive := mutex_recursive \oplus \{(mut?, 0)\},$
 - $priority := priority \oplus \{(mutex_holder(mut?),$

$$\begin{aligned}
& \text{base_priority}(\text{mutex_holder}(\text{mut?}))), \\
& q_size := q_size \oplus \{(\text{mut?}, 1)\}, \\
& \text{release_mutex} := \text{release_mutex} \oplus \{(\text{topReady!}, \text{mut?})\}, \\
& \text{release_rcv} := \text{release_rcv} \oplus \{(\text{topReady!}, \text{mut?})\}, \\
& \text{running_task} := \text{mutex_holder}(\text{mut?}), \\
& \text{state} := \text{state} \oplus \{(\text{topReady!}, \text{ready})\}, \\
& \text{wait_time} := \{\text{topReady!}\} \triangleleft \text{wait_time}, \\
& \text{wait_rcv} := \{\text{topReady!}\} \triangleleft \text{wait_rcv} \\
& \wedge \text{priority}(\text{topReady!}) \leq \text{base_priority}(\text{mutex_holder}(\text{mut?})) \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus \{(\text{topReady!}, \text{ready})\})st = \text{state}(st) \\
& \quad \Rightarrow (\text{state}(st), (\text{state} \oplus \{(\text{topReady!}, \text{ready})\})st) \\
& \quad \in \text{transition}) \\
& \Rightarrow wr \in \text{dom wait_rcv} \\
& \quad \wedge \text{wait_rcv}(wr) = \text{mut?} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(wr))
\end{aligned}$$

theorem *MutexGiveWinhN_TQTM_vc_ref*
 $\forall \text{MutexGiveWinhN_TQTMFSBSig} \mid \text{true}$

- *pre MutexGiveWinhN_TQTM*

MutexGiveWinhSR_TQTM

$\Delta \text{TaskQueueTimeMutex}$

mut? : *QUEUE*

topWaiting! : *TASK*

topReady! : *TASK*

$\text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv}$
 $\text{mut?} \in \text{dom mutex_holder}$
 $\text{running_task} = \text{mutex_holder}(\text{mut?})$
 $\text{mutex_recursive}(\text{mut?}) = 1$
 $\text{topWaiting!} \in \text{wait_rcv} \sim (\mid \{\text{mut?}\} \mid)$
 $\forall wr : \text{wait_rcv} \sim (\mid \{\text{mut?}\} \mid) \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(wr)$
 $\text{state}(\text{topReady!}) = \text{ready}$
 $\forall rt : \text{state} \sim (\mid \{\text{ready}\} \mid) \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(rt)$
 $\text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task})$
 $\text{priority}(\text{topReady!}) > \text{priority}(\text{topWaiting!})$
 $\text{priority}(\text{topReady!}) > \text{base_priority}(\text{running_task})$
 $\text{tasks}' = \text{tasks}$
 $\text{running_task}' = \text{topReady!}$
 $\text{state}' = \text{state} \oplus \{(\text{running_task} \mapsto \text{ready}),$
 $\quad (\text{topReady!} \mapsto \text{running}), (\text{topWaiting!} \mapsto \text{ready})\}$
 $\text{phys_context}' = \text{log_context}(\text{topReady!})$
 $\text{log_context}' = \text{log_context} \oplus \{(\text{running_task} \mapsto \text{phys_context})\}$
 $\text{priority}' = \text{priority} \oplus \{(\text{running_task} \mapsto \text{base_priority}(\text{running_task}))\}$

$$\begin{aligned}
queue' &= queue \\
q_max' &= q_max \\
q_size' &= q_size \oplus \{(mut? \mapsto 1)\} \\
wait_snd' &= wait_snd \\
wait_rcv' &= \{topWaiting!\} \triangleleft wait_rcv \\
release_snd' &= release_snd \\
release_rcv' &= release_rcv \oplus \{(topWaiting! \mapsto mut?)\} \\
clock' &= clock \\
delayed_task' &= delayed_task \\
wait_time' &= \{topWaiting!\} \triangleleft wait_time \\
time_slice' &= time_slice \\
semaphore' &= semaphore \\
mutex' &= mutex \\
mutex_holder' &= \{mut?\} \triangleleft mutex_holder \\
mutex_recursive' &= mutex_recursive \oplus \{(mut? \mapsto 0)\} \\
basePriorityMan & \\
release_mutex' &= release_mutex \oplus \{(topWaiting! \mapsto mut?)\}
\end{aligned}$$

MutexGiveWinhSR_TQTMFSBSig

TaskQueueTimeMutex

mut? : QUEUE

$$\begin{aligned}
&running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv \\
&mut? \in \text{dom } mutex_holder \\
&running_task = mutex_holder(mut?) \\
&mutex_recursive(mut?) = 1 \\
&base_priority(running_task) \neq priority(running_task) \\
&\exists topReady! : state \sim (\{ready\}) \\
&\quad \bullet (\forall rt : state \sim (\{ready\}) \mid \bullet priority(topReady!) \geq priority(rt)) \\
&\quad \quad \wedge (\forall wr : wait_rcv \sim (\{mut?\}) \mid \\
&\quad \quad \quad \bullet priority(topReady!) > priority(wr)) \\
&\quad \quad \wedge priority(topReady!) > base_priority(running_task)
\end{aligned}$$

theorem lMutexGiveWinhSR_TQTM_Lemma

$\forall TaskQueueTimeMutex; topReady!, topWaiting! : TASK;$

$mut? : QUEUE$

$\mid running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$\wedge mut? \in \text{dom } mutex_holder$

$\wedge running_task = mutex_holder(mut?)$

$\wedge mutex_recursive(mut?) = 1$

$\wedge topWaiting! \in wait_rcv \sim (\{mut?\})$

$\wedge (\forall wrct : wait_rcv \sim (\{mut?\}) \mid$

$\quad \bullet priority(topWaiting!) \geq priority(wrct))$

$\wedge state(topReady!) = ready$

$\wedge (\forall rtsk : state \sim (\{ready\}) \mid$

$$\begin{aligned}
& \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}) \\
& \wedge \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{topWaiting!}) \\
& \wedge \text{priority}(\text{topReady!}) > \text{base_priority}(\text{running_task}) \\
& \wedge \text{running_task} \in \text{ran}(\{\text{mut?}\} \triangleleft \text{mutex_holder}) \\
& \bullet \neg (\text{TaskQueueTimeMutex}[\text{log_context} := \text{log_context} \oplus \\
& \quad \{(\text{mutex_holder}(\text{mut?}), \text{phys_context})\}, \\
& \quad \text{mutex_holder} := \{\text{mut?}\} \triangleleft \text{mutex_holder}, \\
& \quad \text{mutex_recursive} := \text{mutex_recursive} \oplus \{(\text{mut?}, 0)\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{priority} := \text{priority} \oplus \{(\text{mutex_holder}(\text{mut?}), \\
& \quad \quad \text{base_priority}(\text{mutex_holder}(\text{mut?}))\}, \\
& \quad \text{q_size} := \text{q_size} \oplus \{(\text{mut?}, 1)\}, \\
& \quad \text{release_mutex} := \text{release_mutex} \oplus \{(\text{topWaiting!}, \text{mut?})\}, \\
& \quad \text{release_rcv} := \text{release_rcv} \oplus \{(\text{topWaiting!}, \text{mut?})\}, \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus (\{(\text{mutex_holder}(\text{mut?}), \text{ready})\} \cup \\
& \quad \quad (\{(\text{topReady!}, \text{running})\} \cup \{(\text{topWaiting!}, \text{ready})\})), \\
& \quad \text{wait_time} := \{\text{topWaiting!}\} \triangleleft \text{wait_time}, \\
& \quad \text{wait_rcv} := \{\text{topWaiting!}\} \triangleleft \text{wait_rcv}] \\
& \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{topWaiting!}) \\
& \wedge (\text{st} \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{mutex_holder}(\text{mut?}), \text{ready})\} \cup \\
& \quad \quad (\{(\text{topReady!}, \text{running})\} \cup \\
& \quad \quad \quad \{(\text{topWaiting!}, \text{ready})\}))) \text{st} = \text{state}(\text{st}) \\
& \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus \\
& \quad \quad (\{(\text{mutex_holder}(\text{mut?}), \text{ready})\} \cup \\
& \quad \quad \quad (\{(\text{topReady!}, \text{running})\} \cup \\
& \quad \quad \quad \quad \{(\text{topWaiting!}, \text{ready})\}))) \text{st}) \\
& \quad \quad \in \text{transition}) \\
& \wedge (\text{wr} \in \text{dom wait_rcv} \wedge \text{wait_rcv}(\text{wr}) = \text{mut?} \\
& \quad \Rightarrow \text{priority}(\text{topWaiting!}) \geq \text{priority}(\text{wr})) \\
& \Rightarrow \text{rt} \in \text{TASK} \\
& \quad \wedge \text{state}(\text{rt}) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rt}))
\end{aligned}$$

theorem lMutexGiveWinhSR_TQTM_Lemma1

$$\begin{aligned}
& \forall \text{TaskQueueTimeMutex}; \text{topReady!}, \text{topWaiting!} : \text{TASK}; \\
& \quad \text{mut?} : \text{QUEUE} \\
& \quad | \text{running_task} \notin \text{dom release_snd} \cup \text{dom release_rcv} \\
& \quad \wedge \text{mut?} \in \text{dom mutex_holder} \\
& \quad \wedge \text{running_task} = \text{mutex_holder}(\text{mut?}) \\
& \quad \wedge \text{mutex_recursive}(\text{mut?}) = 1 \\
& \quad \wedge \text{topWaiting!} \in \text{wait_rcv} \sim (\{ \text{mut?} \} \mid) \\
& \quad \wedge (\forall \text{wrct} : \text{wait_rcv} \sim (\{ \text{mut?} \} \mid) \\
& \quad \quad \bullet \text{priority}(\text{topWaiting!}) \geq \text{priority}(\text{wrct}))
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{state}(\text{topReady!}) = \text{ready} \\
& \wedge (\forall \text{rtsk} : \text{state} \sim (\{ \text{ready} \} |)) \\
& \quad \bullet \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rtsk}) \\
& \wedge \text{base_priority}(\text{running_task}) \neq \text{priority}(\text{running_task}) \\
& \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{topWaiting!}) \\
& \wedge \text{priority}(\text{topReady!}) > \text{base_priority}(\text{running_task}) \\
& \wedge \text{running_task} \notin \text{ran}(\{\text{mut?}\} \triangleleft \text{mutex_holder}) \\
& \bullet \neg (\text{TaskQueueTimeMutex} [\\
& \quad \text{base_priority} := \{\text{mutex_holder}(\text{mut?})\} \triangleleft \text{base_priority}, \\
& \quad \text{log_context} := \text{log_context} \oplus \\
& \quad \quad \{(\text{mutex_holder}(\text{mut?}), \text{phys_context})\}, \\
& \quad \text{mutex_holder} := \{\text{mut?}\} \triangleleft \text{mutex_holder}, \\
& \quad \text{mutex_recursive} := \text{mutex_recursive} \oplus \{(\text{mut?}, 0)\}, \\
& \quad \text{phys_context} := \text{log_context}(\text{topReady!}), \\
& \quad \text{priority} := \text{priority} \oplus \{(\text{mutex_holder}(\text{mut?}), \\
& \quad \quad \text{base_priority}(\text{mutex_holder}(\text{mut?})))\}, \\
& \quad \text{q_size} := \text{q_size} \oplus \{(\text{mut?}, 1)\}, \\
& \quad \text{release_mutex} := \text{release_mutex} \oplus \{(\text{topWaiting!}, \text{mut?})\}, \\
& \quad \text{release_rcv} := \text{release_rcv} \oplus \{(\text{topWaiting!}, \text{mut?})\}, \\
& \quad \text{running_task} := \text{topReady!}, \\
& \quad \text{state} := \text{state} \oplus (\{(\text{mutex_holder}(\text{mut?}), \text{ready})\} \cup \\
& \quad \quad (\{(\text{topReady!}, \text{running})\} \cup \{(\text{topWaiting!}, \text{ready})\})), \\
& \quad \text{wait_time} := \{\text{topWaiting!}\} \triangleleft \text{wait_time}, \\
& \quad \text{wait_rcv} := \{\text{topWaiting!}\} \triangleleft \text{wait_rcv} \\
& \wedge \text{priority}(\text{topReady!}) > \text{priority}(\text{topWaiting!}) \\
& \wedge (\text{st} \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus (\{(\text{mutex_holder}(\text{mut?}), \text{ready})\} \cup \\
& \quad \quad (\{(\text{topReady!}, \text{running})\} \cup \\
& \quad \quad \quad \{(\text{topWaiting!}, \text{ready})\}))) \text{st} = \text{state}(\text{st}) \\
& \quad \Rightarrow (\text{state}(\text{st}), (\text{state} \oplus \\
& \quad \quad (\{(\text{mutex_holder}(\text{mut?}), \text{ready})\} \cup \\
& \quad \quad \quad (\{(\text{topReady!}, \text{running})\} \cup \\
& \quad \quad \quad \quad \{(\text{topWaiting!}, \text{ready})\}))) \text{st}) \in \text{transition}) \\
& \wedge (\text{wr} \in \text{dom } \text{wait_rcv} \wedge \text{wait_rcv}(\text{wr}) = \text{mut?} \\
& \quad \Rightarrow \text{priority}(\text{topWaiting!}) \geq \text{priority}(\text{wr})) \\
& \Rightarrow \text{rt} \in \text{TASK} \\
& \quad \wedge \text{state}(\text{rt}) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topReady!}) \geq \text{priority}(\text{rt}))
\end{aligned}$$

theorem MutexGiveWinhSR_TQTM_vc_ref

$\forall \text{MutexGiveWinhSR_TQTMFSBSig} \mid \text{true}$

\bullet pre *MutexGiveWinhSR_TQTM*

MutexGiveWinhSW_TQTM

*TaskQueueTimeMutex**mut?* : *QUEUE**topWaiting!* : *TASK*

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $running_task = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $topWaiting! \in wait_rcv \sim (\{mut?\})$
 $\forall wr : wait_rcv \sim (\{mut?\}) \bullet priority(topWaiting!) \geq priority(wr)$
 $\forall rt : state \sim (\{ready\}) \bullet priority(topWaiting!) \geq priority(rt)$
 $base_priority(running_task) \neq priority(running_task)$
 $priority(topWaiting!) > base_priority(running_task)$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $| st? = ready \wedge pri? = priority \oplus$
 $\{(running_task \mapsto base_priority(running_task))\}$
 $\bullet Reschedule[topWaiting!/target?, tasks/tasks?]$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(mut? \mapsto 1)\}$
 $wait_snd' = wait_snd$
 $wait_rcv' = \{topWaiting!\} \triangleleft wait_rcv$
 $release_snd' = release_snd$
 $release_rcv' = release_rcv \oplus \{(topWaiting! \mapsto mut?)\}$
 $clock' = clock$
 $delayed_task' = delayed_task$
 $wait_time' = \{topWaiting!\} \triangleleft wait_time$
 $time_slice' = time_slice$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
basePriorityMan
 $release_mutex' = release_mutex \oplus \{(topWaiting! \mapsto mut?)\}$

MutexGiveWinhSW_TQTMFSBSig

*TaskQueueTimeMutex**mut?* : *QUEUE*

$running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $running_task = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $base_priority(running_task) \neq priority(running_task)$
 $\exists topWaiting! : wait_rcv \sim (\{mut?\})$

- $(\forall wr : wait_rcv \sim (\{ mut? \}) \mid \bullet \text{priority}(topWaiting!) \geq \text{priority}(wr))$
- $(\forall rt : state \sim (\{ ready \}) \mid \bullet \text{priority}(topWaiting!) \geq \text{priority}(rt))$
- $\wedge \text{priority}(topWaiting!) > \text{base_priority}(running_task)$

theorem lMutexGiveWinhSW_TQTM_Lemma

$\forall TaskQueueTimeMutex; topWaiting! : TASK;$

$mut? : QUEUE$

$\mid running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$\wedge mut? \in \text{dom } mutex_holder$

$\wedge running_task = mutex_holder(mut?)$

$\wedge mutex_recursive(mut?) = 1$

$\wedge topWaiting! \in wait_rcv \sim (\{ mut? \}) \mid$

$\wedge (\forall wrct : wait_rcv \sim (\{ mut? \}) \mid$

• $\text{priority}(topWaiting!) \geq \text{priority}(wrct))$

$\wedge (\forall rtsk : state \sim (\{ ready \}) \mid$

• $\text{priority}(topWaiting!) \geq \text{priority}(rtsk))$

$\wedge \text{base_priority}(running_task) \neq \text{priority}(running_task)$

$\wedge \text{priority}(topWaiting!) > \text{base_priority}(running_task)$

$\wedge running_task \in \text{ran}(\{ mut? \}) \triangleleft mutex_holder$

• $\neg (TaskQueueTimeMutex[log_context := log_context \oplus$

$\{(mutex_holder(wait_rcv(topWaiting!)),$
 $phys_context)\},$

$mutex_holder := \{wait_rcv(topWaiting!)\} \triangleleft mutex_holder,$

$mutex_recursive := mutex_recursive \oplus$

$\{(wait_rcv(topWaiting!), 0)\},$

$phys_context := log_context(topWaiting!),$

$priority := priority \oplus$

$\{(mutex_holder(wait_rcv(topWaiting!)),$

$base_priority(mutex_holder($
 $wait_rcv(topWaiting!)))\},$

$q_size := q_size \oplus \{(wait_rcv(topWaiting!), 1)\},$

$release_mutex := release_mutex \oplus$

$\{(topWaiting!, wait_rcv(topWaiting!))\},$

$release_rcv := release_rcv \oplus$

$\{(topWaiting!, wait_rcv(topWaiting!))\},$

$running_task := topWaiting!,$

$state := state \oplus$

$\{(mutex_holder(wait_rcv(topWaiting!)), ready)\}$

$\cup \{(topWaiting!, running)\}\},$

$wait_time := \{topWaiting!\} \triangleleft wait_time,$

$wait_rcv := \{topWaiting!\} \triangleleft wait_rcv]$

$\wedge (st \in TASK$

$\wedge \neg (state \oplus$

$\{(mutex_holder(wait_rcv(topWaiting!)), ready)\} \cup$

$\{(topWaiting!, running)\}) st = state(st)$

$$\begin{aligned}
&\Rightarrow (state(st), (state \oplus \\
&\quad (\{(mutex_holder(wait_rcv(topWaiting!)), ready)\} \cup \\
&\quad \{(topWaiting!, running)\}))st) \in transition) \\
&\wedge (wr \in \text{dom } wait_rcv \wedge wait_rcv(wr) = wait_rcv(topWaiting!) \\
&\quad \Rightarrow priority(topWaiting!) \geq priority(wr)) \\
&\Rightarrow rt \in TASK \\
&\quad \wedge state(rt) = ready \\
&\quad \wedge \neg priority(topWaiting!) \geq priority(rt)
\end{aligned}$$

theorem lMutexGiveWinhSW_TQTM_Lemma1

$$\begin{aligned}
&\forall TaskQueueTimeMutex; topWaiting! : TASK; \\
&\quad mut? : QUEUE \\
&\quad | running_task \notin \text{dom } release_snd \cup \text{dom } release_rcv \\
&\quad \wedge mut? \in \text{dom } mutex_holder \\
&\quad \wedge running_task = mutex_holder(mut?) \\
&\quad \wedge mutex_recursive(mut?) = 1 \\
&\quad \wedge topWaiting! \in wait_rcv \sim (\{mut?\}) \\
&\quad \wedge (\forall wrct : wait_rcv \sim (\{mut?\}) \\
&\quad \quad \bullet priority(topWaiting!) \geq priority(wrct)) \\
&\quad \wedge (\forall rtsk : state \sim (\{ready\}) \\
&\quad \quad \bullet priority(topWaiting!) \geq priority(rtsk)) \\
&\quad \wedge base_priority(running_task) \neq priority(running_task) \\
&\quad \wedge priority(topWaiting!) > base_priority(running_task) \\
&\quad \wedge running_task \notin \text{ran}(\{mut?\}) \triangleleft mutex_holder \\
&\quad \bullet \neg (TaskQueueTimeMutex [\\
&\quad \quad base_priority := \{mutex_holder(wait_rcv(topWaiting!))\} \\
&\quad \quad \triangleleft base_priority, \\
&\quad \quad log_context := log_context \oplus \\
&\quad \quad \quad \{(mutex_holder(wait_rcv(topWaiting!)), phys_context)\}, \\
&\quad \quad mutex_holder := \{wait_rcv(topWaiting!)\} \triangleleft mutex_holder, \\
&\quad \quad mutex_recursive := mutex_recursive \oplus \\
&\quad \quad \quad \{(wait_rcv(topWaiting!), 0)\}, \\
&\quad \quad phys_context := log_context(topWaiting!), \\
&\quad \quad priority := priority \oplus \\
&\quad \quad \quad \{(mutex_holder(wait_rcv(topWaiting!)), \\
&\quad \quad \quad \quad base_priority(\\
&\quad \quad \quad \quad \quad mutex_holder(wait_rcv(topWaiting!)))\}, \\
&\quad \quad q_size := q_size \oplus \{(wait_rcv(topWaiting!), 1)\}, \\
&\quad \quad release_mutex := release_mutex \oplus \\
&\quad \quad \quad \{(topWaiting!, wait_rcv(topWaiting!))\}, \\
&\quad \quad release_rcv := release_rcv \oplus \\
&\quad \quad \quad \{(topWaiting!, wait_rcv(topWaiting!))\}, \\
&\quad \quad running_task := topWaiting!, \\
&\quad \quad state := state \oplus \\
&\quad \quad \quad (\{(mutex_holder(wait_rcv(topWaiting!)), ready)\} \cup \\
&\quad \quad \quad \{(topWaiting!, running)\}),
\end{aligned}$$

$$\begin{aligned}
& \text{wait_time} := \{\text{topWaiting!}\} \triangleleft \text{wait_time}, \\
& \text{wait_rcv} := \{\text{topWaiting!}\} \triangleleft \text{wait_rcv} \\
& \wedge (st \in \text{TASK} \\
& \quad \wedge \neg (\text{state} \oplus \\
& \quad \quad \{(\text{mutex_holder}(\text{wait_rcv}(\text{topWaiting!})), \text{ready})\} \cup \\
& \quad \quad \Rightarrow (\text{state}(st), (\text{state} \oplus \\
& \quad \quad \quad \{(\text{mutex_holder}(\text{wait_rcv}(\text{topWaiting!})), \text{ready})\} \cup \\
& \quad \quad \quad \{(topWaiting!, \text{running})\}))st) \in \text{transition}) \\
& \wedge (wr \in \text{dom wait_rcv} \wedge \text{wait_rcv}(wr) = \text{wait_rcv}(\text{topWaiting!}) \\
& \quad \Rightarrow \text{priority}(\text{topWaiting!}) \geq \text{priority}(wr)) \\
& \Rightarrow rt \in \text{TASK} \\
& \quad \wedge \text{state}(rt) = \text{ready} \\
& \quad \wedge \neg \text{priority}(\text{topWaiting!}) \geq \text{priority}(rt))
\end{aligned}$$

theorem *MutexGiveWinhSW_TQTM_vc_ref*

\forall *MutexGiveWinhSW_TQTMFSBSig* | *true*

- pre *MutexGiveWinhSW_TQTM*

MutexGiveNRecursive_TQTM _____

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

topReady! : *TASK*

running_task \notin *dom release_snd* \cup *dom release_rcv*

mut? \in *dom mutex_holder*

running_task = *mutex_holder*(*mut?*)

mutex_recursive(*mut?*) > 1

\exists *TaskQueueTime*

semaphore' = *semaphore*

mutex' = *mutex*

mutex_holder' = *mutex_holder*

mutex_recursive' = *mutex_recursive* \oplus

$\{(\text{mut?} \mapsto \text{mutex_recursive}(\text{mut?}) - 1)\}$

\exists *OriginalPrioData*

\exists *MReleasingData*

topReady! = *running_task*

MutexGiveNRecursive_TQTMFSBSig _____

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

running_task \notin *dom release_snd* \cup *dom release_rcv*

mut? \in *dom mutex_holder*

running_task = *mutex_holder*(*mut?*)

$mutex_recursive(mut?) > 1$

theorem MutexGiveNRecursive_TQTM_vc_ref
 \forall *MutexGiveNRecursive_TQTMFSBSig* | *true*
• pre *MutexGiveNRecursive_TQTM*

MutexGive_TQTM \cong *MutexGiveNnonInh_TQTM*
 \vee *MutexGiveNInhN_TQTM*
 \vee *MutexGiveNInhS_TQTM*
 \vee *MutexGiveWnonInhN_TQTM*
 \vee *MutexGiveWnonInhS_TQTM*
 \vee *MutexGiveWInhN_TQTM*
 \vee *MutexGiveWInhSR_TQTM*
 \vee *MutexGiveWInhSW_TQTM*
 \vee *MutexGiveNRecursive_TQTM*

Appendix H

SPECIFICATION FOR MULTI-CORE TASK MODEL

[*CONTEXT*, *TASK*, *CORE*]

$\text{bare_context} : \text{CONTEXT}$ $\text{idles} : \mathbb{F} \text{TASK}$ $\text{cores} : \mathbb{F} \text{CORE}$

$\# \text{cores} = \# \text{idle}$ $\text{cores} \neq \emptyset$

$\text{STATE} ::= \text{nonexistent} \mid \text{ready} \mid \text{blocked} \mid \text{suspended} \mid \text{running}$

$\text{transition} ::= (\{\text{blocked}\} \times \{\text{nonexistent}, \text{ready}, \text{running}, \text{suspended}\})$
 $\cup (\{\text{nonexistent}\} \times \{\text{ready}, \text{running}\})$
 $\cup (\{\text{ready}\} \times \{\text{nonexistent}, \text{running}, \text{suspended}\})$
 $\cup (\{\text{running}\} \times \{\text{blocked}, \text{nonexistent}, \text{ready}, \text{suspended}\})$
 $\cup (\{\text{suspended}\} \times \{\text{nonexistent}, \text{ready}, \text{running}\})$

$\text{slice_delay} : \mathbb{N}$

$\text{slice_delay} = 1$

$\text{BOOL} ::= \text{TRUE} \mid \text{FALSE}$

TaskData

$tasks : \mathbb{F} \text{ TASK}$
 $running_tasks : \text{cores} \multimap \text{TASK}$
 $executable : \text{TASK} \multimap \text{cores}$

$\text{ran } running_tasks \subseteq tasks$
 $idles \subseteq tasks$
 $\text{dom } executable = tasks$
 $\forall t : \text{ran } running_tasks \bullet running_tasks \sim t = executable(t)$

Init_TaskData

TaskData'

$tasks' = idles$

StateData

$state : \text{TASK} \rightarrow \text{STATE}$

$\forall i : idles \bullet state(i) \in \{ready, running\}$

Init_StateData

StateData'

$state' = (\lambda x : \text{TASK} \bullet nonexistent) \oplus (idles \times \{running\})$

ContextData

$phys_context : \text{cores} \rightarrow \text{CONTEXT}$
 $log_context : \text{TASK} \rightarrow \text{CONTEXT}$

Init_ContextData

ContextData'

$phys_context' = (\lambda c : \text{cores} \bullet bare_context)$
 $log_context' = (\lambda x : \text{TASK} \bullet bare_context)$

PrioData

$priority : \text{TASK} \rightarrow \mathbb{N}$

$\forall i : idles \bullet priority(i) = 0$

Init_PrioData

PrioData'

$priority' = (\lambda x : TASK \bullet 0)$

Task

TaskData

StateData

ContextData

PrioData

$tasks = TASK \setminus (state \sim (\{nonexistent\}))$

$state \sim (\{running\}) = \text{ran } running_tasks$

$\forall pt : state \sim (\{ready\}); r : \text{ran } running_tasks$

$| executable(pt) = executable(r) \bullet priority(r) \geq priority(pt)$

$\Delta Task$

Task

Task'

$\forall st : TASK | state'(st) \neq state(st) \bullet state(st) \mapsto state'(st) \in transition$

Init_Task

Task'

Init_TaskData

Init_StateData

Init_ContextData

Init_PrioData

createTaskSpeCoreN_T

$\Delta Task$

$target? : TASK$

$newpri? : \mathbb{N}$

$executeCore : CORE$

$executeCore \in cores$

$state(target?) = nonexistent$

$newpri? \leq priority(running_tasks(executeCore))$

$tasks' = tasks \cup \{target?\}$

$running_tasks' = running_tasks$

$executable' = executable \oplus \{(target? \mapsto executeCore)\}$

$state' = state \oplus \{(target? \mapsto ready)\}$

$$\exists \text{ContextData}$$

$$\text{priority}' = \text{priority} \oplus \{(target? \mapsto \text{newpri?})\}$$

$$\text{findACore}_T$$

$$\Delta \text{Task}$$

$$target? : \text{TASK}$$

$$\text{newpri?} : \mathbb{N}$$

$$\text{executeCore?} : \text{CORE}$$

$$\text{executeCore} : \text{CORE}$$

$$\text{executeCore?} \notin \text{cores}$$

$$\text{executeCore} \in \text{cores}$$

$$\exists tcs, cs : \mathbb{F} \text{cores} \mid$$

$$tcs = \{pc : \text{cores} \mid \text{newpri?} > \text{priority}(\text{running_tasks}(pc))\}$$

$$\bullet (tcs = \emptyset \Rightarrow cs = \text{cores}) \wedge (tcs \neq \emptyset \Rightarrow cs = tcs)$$

$$\wedge (\forall oc : cs \bullet \text{executeCore} \in cs$$

$$\wedge \#(\text{executable}^{\sim}(\{ \text{executeCore} \}))$$

$$\leq \#(\text{executable}^{\sim}(\{ oc \})))$$

$$\text{CreateTaskN}_T \hat{=} ([\text{executeCore?}, \text{executeCore} : \text{CORE}$$

$$\mid \text{executeCore?} \in \text{cores} \wedge \text{executeCore} = \text{executeCore?}] \vee \text{findACore}_T)$$

$$\wedge \text{createTaskSpeCoreN}_T$$

$$\text{Reschedule}$$

$$\Delta \text{Task}$$

$$target? : \text{TASK}$$

$$\text{tasks?} : \mathbb{F} \text{TASK}$$

$$\text{executable?} : \text{TASK} \rightsquigarrow \text{cores}$$

$$st? : \text{STATE}$$

$$pri? : \text{TASK} \rightarrow \mathbb{N}$$

$$\text{tasks}' = \text{tasks?}$$

$$\text{running_tasks}' = \text{running_tasks} \oplus \{(\text{executable?}(target?) \mapsto target?)\}$$

$$\text{executable}' = \text{executable?}$$

$$\text{state}' = \text{state} \oplus \{(target? \mapsto \text{running}),$$

$$(\text{running_tasks}(\text{executable?}(target?)) \mapsto st?)\}$$

$$\text{phys_context}' = \text{phys_context} \oplus$$

$$\{(\text{executable?}(target?) \mapsto \text{log_context}(target?))\}$$

$$\text{log_context}' = \text{log_context} \oplus$$

$$\{(\text{running_tasks}(\text{executable?}(target?))$$

$$\mapsto \text{phys_context}(\text{executable?}(target?)))\}$$

$$\text{priority}' = pri?$$

$$\text{disableReschedule} \hat{=} [\text{Task} \mid \text{false}] \wedge \text{Reschedule}$$

createTaskSpeCoreS_T

$\Delta Task$

$target? : TASK$

$newpri? : \mathbb{N}$

$executeCore : CORE$

$executeCore \in cores$

$state(target?) = nonexistent$

$newpri? > priority(running_tasks(executeCore))$

$\exists st? : STATE; tasks? : \mathbb{F} TASK; executable? : TASK \rightsquigarrow cores;$

$pri? : TASK \rightarrow \mathbb{N}$

$| st? = ready \wedge tasks? = tasks \cup \{target?\}$

$\wedge executable? = executable \oplus \{(target? \mapsto executeCore)\}$

$\wedge pri? = priority \oplus \{(target? \mapsto newpri?)\} \bullet Reschedule$

$CreateTaskS_T \triangleq ([executeCore?, executeCore : CORE$

$| executeCore? \in cores \wedge executeCore = executeCore?] \vee findACore_T)$

$\wedge createTaskSpeCoreS_T$

$CreateTask_T \triangleq CreateTaskN_T \vee CreateTaskS_T$

DeleteTaskN_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus idles$

$state(target?) \in \{ready, blocked, suspended\}$

$tasks' = tasks \setminus \{target?\}$

$running_tasks' = running_tasks$

$executable' = \{target?\} \triangleleft executable$

$state' = state \oplus \{(target? \mapsto nonexistent)\}$

$phys_context' = phys_context$

$log_context' = log_context \oplus \{(target? \mapsto bare_context)\}$

$\exists PrioData$

$topReady! = running_tasks(executable(target?))$

findTopReady

$Task$

$target? : TASK$

$topReady! : TASK$

$state(topReady!) = ready$

$executable(topReady!) = executable(target?)$

$$\forall rt : state \sim (\{ready\}) \mid executable(rt) = executable(topReady!) \\ \bullet \text{priority}(topReady!) \geq \text{priority}(rt)$$

DeleteTaskS_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus idles$

$state(target?) = running$

$findTopReady$

$tasks' = tasks \setminus \{target?\}$

$running_tasks' = running_tasks \oplus \{executable(target?) \mapsto topReady!\}$

$executable' = \{target?\} \triangleleft executable$

$state' = state \oplus \{(topReady! \mapsto running), (target? \mapsto nonexistent)\}$

$phys_context' = phys_context \oplus$

$\{(executable(target?) \mapsto log_context(topReady!))\}$

$log_context' = log_context \oplus \{(target? \mapsto bare_context)\}$

$\exists PrioData$

$DeleteTask_T \cong DeleteTaskN_T \vee DeleteTaskS_T$

ExecuteRunningTask_T

$\Delta Task$

$target! : \mathbb{F} TASK$

$\exists TaskData$

$\exists StateData$

$\forall c : cores \bullet phys_context'(c) \neq phys_context(c)$

$log_context' = log_context$

$\exists PrioData$

$target! = ran\ running_tasks$

SuspendTaskN_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$state(target?) \in \{ready, blocked\}$

$target? \notin idles$

$\exists TaskData$

$state' = state \oplus \{(target? \mapsto suspended)\}$

$\exists ContextData$

$\exists PrioData$
 $topReady! = running_tasks(executable(target?))$

$SuspendTaskS_T$

$\Delta Task$
 $target? : TASK$
 $topReady! : TASK$

$state(target?) = running$
 $target? \notin idles$
 $findTopReady$
 $\exists st? : STATE \mid st? = suspended$

- $Reschedule[tasks/tasks?, executable/executable?, priority/pri?, topReady!/target?]$

$SuspendTaskO_T$

$\exists Task$
 $target? : TASK$
 $topReady! : TASK$

$state(target?) = suspended$
 $topReady! = running_tasks(executable(target?))$

$SuspendTask_T \cong SuspendTaskN_T$
 $\vee SuspendTaskS_T$
 $\vee SuspendTaskO_T$

$ResumeTaskN_T$

$\Delta Task$
 $target? : TASK$

$state(target?) = suspended$
 $priority(target?) \leq priority(running_tasks(executable(target?)))$
 $\exists TaskData$
 $state' = state \oplus \{(target? \mapsto ready)\}$
 $\exists ContextData$
 $\exists PrioData$

$ResumeTaskS_T$

$\Delta Task$
 $target? : TASK$

$state(target?) = suspended$
 $priority(target?) > priority(running_tasks(executable(target?)))$
 $\exists st? : STATE \mid st? = ready$

- $Reschedule[tasks/tasks?, executable/executable?, priority/pri?]$

$ResumeTask_T \triangleq ResumeTaskN_T \vee ResumeTaskS_T$

$ChangeTaskPriorityN_T$

$\Delta Task$

$newpri? : \mathbb{N}$

$target? : TASK$

$topReady! : TASK$

$state(target?) = ready$

$\Rightarrow newpri? \leq priority(running_tasks(executable(target?)))$

$state(target?) = running \Rightarrow$

$(\forall rt : state \sim (\mid \{ready\} \mid) \mid executable(rt) = executable(target?)$

- $newpri? \geq priority(rt)$

$state(target?) \neq nonexistent$

$target? \in idles \Rightarrow newpri? = 0$

$\exists TaskData$

$\exists StateData$

$\exists ContextData$

$priority' = priority \oplus \{(target? \mapsto newpri?)\}$

$topReady! = running_tasks(executable(target?))$

$ChangeTaskPriorityS_T$

$\Delta Task$

$newpri? : \mathbb{N}$

$target? : TASK$

$topReady! : TASK$

$state(target?) = ready$

$newpri? > priority(running_tasks(executable(target?)))$

$target? \in idles \Rightarrow newpri? = 0$

$\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$

$\mid st? = ready \wedge pri? = priority \oplus \{(target? \mapsto newpri?)\}$

- $Reschedule[tasks/tasks?, executable/executable?]$

$topReady! = target?$

$ChangeTaskPriorityD_T$

$\Delta Task$

$newpri? : \mathbb{N}$

$target? : TASK$
 $topReady! : TASK$

$state(target?) = running$
 $target? \in idles \Rightarrow newpri? = 0$
 $findTopReady$
 $newpri? < priority(topReady!)$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $| st? = ready \wedge pri? = priority \oplus \{(target? \mapsto newpri?)\}$

- $Reschedule[tasks/tasks?, executable/executable?, topReady!/target?]$

$ChangeTaskPriority_T \hat{=} ChangeTaskPriorityN_T$
 $\vee ChangeTaskPriorityS_T$
 $\vee ChangeTaskPriorityD_T$

$MigrationN_T$

$\Delta Task$
 $target? : TASK$
 $topReady! : TASK$
 $newCore? : cores$

$state(target?) \in \{ready, blocked, suspended\}$
 $state(target?) = ready \Rightarrow$
 $priority(target?) \leq priority(running_tasks(newCore?))$
 $target? \notin idles$
 $newCore? \neq executable(target?)$
 $tasks' = tasks$
 $running_tasks' = running_tasks$
 $executable' = executable \oplus \{(target? \mapsto newCore?)\}$
 $\exists StateData$
 $\exists ContextData$
 $\exists PrioData$
 $topReady! = running_tasks(newCore?)$

$MigrationS_T$

$\Delta Task$
 $target? : TASK$
 $topReady! : TASK$
 $newCore? : cores$

$state(target?) = ready$
 $priority(target?) > priority(running_tasks(newCore?))$
 $target? \notin idles$
 $newCore? \neq executable(target?)$

$\exists st? : STATE; executable? : TASK \rightsquigarrow cores$
 $| st? = ready \wedge executable? = executable \oplus \{(target? \mapsto newCore?)\}$
 $\bullet Reschedule[tasks/tasks?, priority/pri?]$
 $topReady! = target?$

MigrationRuN_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$newCore? : cores$

$state(target?) = running$

$priority(target?) \leq priority(running_tasks(newCore?))$

$target? \notin idles$

$newCore? \neq executable(target?)$

$findTopReady$

$\exists st? : STATE; executable? : TASK \rightsquigarrow cores$

$| st? = ready \wedge executable? = executable \oplus \{(target? \mapsto newCore?)\}$

$\bullet Reschedule[tasks/tasks?, priority/pri?]$

MigrationRuS_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$newCore? : cores$

$state(target?) = running$

$priority(target?) > priority(running_tasks(newCore?))$

$target? \notin idles$

$newCore? \neq executable(target?)$

$findTopReady$

$tasks' = tasks$

$running_tasks' = running_tasks$

$\oplus \{(executable(target?) \mapsto topReady!), (newCore? \mapsto target?)\}$

$executable' = executable \oplus \{(target? \mapsto newCore?)\}$

$state' = state$

$\oplus \{(topReady! \mapsto running), (running_tasks(newCore?) \mapsto ready)\}$

$phys_context' = phys_context$

$\oplus \{(executable(topReady!) \mapsto log_context(topReady!)),$

$(newCore? \mapsto phys_context(executable(topReady!)))\}$

$log_context' = log_context$

$\oplus \{(running_tasks(newCore?) \mapsto phys_context(newCore?))\}$

$priority' = priority$

$$\begin{aligned} \text{Migration}_T &\cong \text{Migration}_{N_T} \\ &\vee \text{Migration}_{S_T} \\ &\vee \text{Migration}_{RuN_T} \\ &\vee \text{Migration}_{RuS_T} \end{aligned}$$

Appendix I

SPECIFICATION FOR MULTI-CORE QUEUE MODEL

[*QUEUE*]

QueueData

$queue : \mathbb{P} \textit{QUEUE}$
 $q_max : \textit{QUEUE} \rightarrow \mathbb{N}_1$
 $q_size : \textit{QUEUE} \rightarrow \mathbb{N}$
 $q_ava : \textit{QUEUE} \rightarrow \mathbb{F} \textit{CORE}$

$\text{dom } q_max = \text{dom } q_size$
 $\text{dom } q_size = \text{dom } q_ava$
 $\text{dom } q_ava = queue$
 $\text{ran } q_ava \subseteq \mathbb{F} \textit{cores}$
 $\forall q : \textit{QUEUE} \mid q \in queue \bullet q_size(q) \leq q_max(q)$

Init_QueueData

QueueData'

$queue' = \emptyset$
 $q_max' = \emptyset$
 $q_size' = \emptyset$
 $q_ava' = \emptyset$

WaitingData

$wait_snd : \textit{TASK} \rightarrow \textit{QUEUE}$

$wait_rcv : TASK \rightarrow QUEUE$
$dom\ wait_snd \cap dom\ wait_rcv = \emptyset$

$Init_WaitingData$
$WaitingData'$
$wait_snd' = \emptyset$
$wait_rcv' = \emptyset$

$QReleasingData$
$release_snd : TASK \rightarrow QUEUE$
$release_rcv : TASK \rightarrow QUEUE$
$dom\ release_snd \cap dom\ release_rcv = \emptyset$

$Init_QReleasingData$
$QReleasingData'$
$release_snd' = \emptyset$
$release_rcv' = \emptyset$

$Queue$
$QueueData$
$WaitingData$
$QReleasingData$
$ran\ wait_snd \subseteq queue$
$ran\ wait_rcv \subseteq queue$
$ran\ release_snd \subseteq queue$
$ran\ release_rcv \subseteq queue$
$(dom\ wait_snd \cup dom\ wait_rcv)$
$\quad \cap (dom\ release_snd \cap dom\ release_rcv) = \emptyset$

$Init_Queue$
$Queue'$
$Init_QueueData$
$Init_WaitingData$
$Init_QReleasingData$

<i>TaskQueue</i> <i>Task</i> <i>Queue</i>
$\text{dom } \textit{wait_snd} \subseteq \textit{state} \sim (\{ \textit{blocked} \})$ $\text{dom } \textit{wait_rcv} \subseteq \textit{state} \sim (\{ \textit{blocked} \})$

<i>Init_TaskQueue</i> <i>TaskQueue'</i>
<i>Init_Task</i> <i>Init_Queue</i>

$$\Delta \textit{TaskQueue} \hat{=} \textit{TaskQueue} \wedge \textit{TaskQueue}' \wedge \Delta \textit{Task}$$

<i>ExtendTaskXi</i> $\Delta \textit{TaskQueue}$ $\textit{self} ? : \textit{TASK}$
$\textit{state}(\textit{self}?) = \textit{running}$ $\textit{self} ? \notin \text{dom } \textit{release_snd} \cup \text{dom } \textit{release_rcv}$ $\exists \textit{Queue}$

$$\textit{CreateTask_TQ} \hat{=} \textit{ExtendTaskXi} \wedge \textit{CreateTask_T}$$

<i>DeleteTaskN_TQ</i> <i>DeleteTaskN_T</i> $\Delta \textit{TaskQueue}$ $\textit{self} ? : \textit{TASK}$
$\textit{self} ? \notin \text{dom } \textit{release_snd} \cup \text{dom } \textit{release_rcv}$ $\textit{state}(\textit{self}?) = \textit{running}$ $\exists \textit{QueueData}$ $\textit{wait_snd}' = \{ \textit{target} ? \} \triangleleft \textit{wait_snd}$ $\textit{wait_rcv}' = \{ \textit{target} ? \} \triangleleft \textit{wait_rcv}$ $\textit{release_snd}' = \{ \textit{target} ? \} \triangleleft \textit{release_snd}$ $\textit{release_rcv}' = \{ \textit{target} ? \} \triangleleft \textit{release_rcv}$

$$\textit{DeleteTaskS_TQ} \hat{=} \textit{ExtendTaskXi} \wedge \textit{DeleteTaskS_T}$$

$$\textit{DeleteTask_TQ} \hat{=} \textit{DeleteTaskN_TQ} \vee \textit{DeleteTaskS_TQ}$$

$ExecuteRunningTask_TQ \cong ExtendTaskXi \wedge ExecuteRunningTask_T$

$SuspendTaskN_TQ$ $SuspendTaskN_T$ $\Delta TaskQueue$ $self? : TASK$
$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$ $state(self?) = running$ $\exists QueueData$ $wait_snd' = \{target?\} \triangleleft wait_snd$ $wait_rcv' = \{target?\} \triangleleft wait_rcv$ $\exists QReleasingData$

$SuspendTask_TQ \cong SuspendTaskN_TQ$
 $\vee (ExtendTaskXi \wedge (SuspendTaskS_T \vee SuspendTaskO_T))$

$ResumeTask_TQ \cong ExtendTaskXi \wedge ResumeTask_T$

$ChangeTaskPriority_TQ \cong ExtendTaskXi \wedge ChangeTaskPriority_T$

$Migration_TQ \cong ExtendTaskXi \wedge Migration_T$

$CreateQueue_TQ$ $\Delta TaskQueue$ $que? : QUEUE$ $self? : TASK$ $cset? : \mathbb{F} \text{ cores}$ $size? : \mathbb{N}$
$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$ $state(self?) = running$ $que? \notin queue$ $size? > 0$ $cset? \neq \emptyset$ $\exists Task$ $queue' = queue \cup \{que?\}$ $q_max' = q_max \oplus \{(que? \mapsto size?)\}$ $q_size' = q_size \oplus \{(que? \mapsto 0)\}$ $q_ava' = q_ava \oplus \{(que? \mapsto cset?)\}$ $\exists WaitingData$ $\exists QReleasingData$

DeleteQueue_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $self? : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $state(self?) = running$
 $que? \in queue$
 $executable(self?) \in q_ava(que?)$
 $que? \notin \text{ran } wait_snd \cup \text{ran } wait_rcv$
 $que? \notin \text{ran } release_snd \cup \text{ran } release_rcv$
 $\exists Task$
 $queue' = queue \setminus \{que?\}$
 $q_max' = \{que?\} \triangleleft q_max$
 $q_size' = \{que?\} \triangleleft q_size$
 $q_ava' = \{que?\} \triangleleft q_ava$
 $\exists WaitingData$
 $\exists QReleasingData$

QueueSendN_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_rcv$
 $self? \in \text{dom } release_snd \Rightarrow que? = release_snd(self?)$
 $que? \in queue$
 $q_size(que?) < q_max(que?)$
 $state(self?) = running$
 $executable(self?) \in q_ava(que?)$
 $que? \notin \text{ran } wait_rcv$
 $\exists Task$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto q_size(que?) + 1)\}$
 $q_ava' = q_ava$
 $\exists WaitingData$
 $release_snd' = \{self?\} \triangleleft release_snd$
 $release_rcv' = release_rcv$
 $topReady! = self?$

QueueSendF_TQ

$\Delta TaskQueue$
 $que? : QUEUE$

$self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_rcv$
 $self? \in \text{dom } release_snd \Rightarrow que? = release_snd(self?)$
 $que? \in queue$
 $q_size(que?) = q_max(que?)$
 $state(self?) = running$
 $executable(self?) \in q_ava(que?)$
 $self? \notin idles$
 $findTopReady[self?/target?]$
 $\exists st? : STATE \mid st? = blocked$

- $Reschedule[topReady!/target?, tasks/tasks?, executable/executable?, priority/pri?]$

 $\exists QueueData$
 $wait_snd' = wait_snd \oplus \{(self? \mapsto que?)\}$
 $wait_rcv' = wait_rcv$
 $release_snd' = \{self?\} \triangleleft release_snd$
 $release_rcv' = release_rcv$

$QueueSendW_TQ$

$\Delta TaskQueue$
 $que? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_rcv$
 $self? \in \text{dom } release_snd \Rightarrow que? = release_snd(self?)$
 $que? \in queue$
 $state(self?) = running$
 $executable(self?) \in q_ava(que?)$
 $topReady! \in wait_rcv \sim (\{que?\})$
 $\forall wr : wait_rcv \sim (\{que?\}) \mid \bullet priority(topReady!) \geq priority(wr)$
 $priority(topReady!) \leq priority(running_tasks(executable(topReady!)))$
 $\exists TaskData$
 $state' = state \oplus \{(topReady! \mapsto ready)\}$
 $\exists ContextData$
 $\exists PrioData$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto 1)\}$
 $q_ava' = q_ava$
 $wait_snd' = wait_snd$
 $wait_rcv' = \{topReady!\} \triangleleft wait_rcv$
 $release_snd' = \{self?\} \triangleleft release_snd$
 $release_rcv' = release_rcv \oplus \{(topReady! \mapsto que?)\}$

QueueSendWS_TQ

$\Delta TaskQueue$
que? : *QUEUE*
self? : *TASK*
topReady! : *TASK*

self? \notin dom *release_rcv*
self? \in dom *release_snd* $\Rightarrow que? = release_snd(self?)$
que? $\in queue$
state(self?) = running
executable(self?) $\in q_ava(que?)$
self? $\notin idles$
topReady! $\in wait_rcv \sim (\{ que? \} \mid)$
 $\forall wr : wait_rcv \sim (\{ que? \} \mid) \bullet priority(topReady!) \geq priority(wr)$
priority(topReady!) > priority(running_tasks(executable(topReady!)))
 $\exists st? : STATE \mid st? = ready$
 • *Reschedule[topReady!/target?, tasks/tasks?,*
 executable/executable?, priority/pri?]
queue' = queue
q_max' = q_max
q_size' = q_size \oplus \{(que? \mapsto 1)\}
q_ava' = q_ava
wait_snd' = wait_snd
wait_rcv' = \{topReady!\} \triangleleft wait_rcv
release_snd' = \{self?\} \triangleleft release_snd
release_rcv' = release_rcv \oplus \{(topReady! \mapsto que?)\}

QueueSend_TQ $\hat{=} QueueSendN_TQ$
 $\vee QueueSendF_TQ$
 $\vee QueueSendW_TQ$
 $\vee QueueSendWS_TQ$

QueueReceiveN_TQ

$\Delta TaskQueue$
que? : *QUEUE*
self? : *TASK*
topReady! : *TASK*

self? \notin dom *release_snd*
self? \in dom *release_rcv* $\Rightarrow que? = release_rcv(self?)$
que? $\in queue$
q_size(que?) \neq 0
state(self?) = running

$executable(self?) \in q_ava(que?)$
 $que? \notin \text{ran } wait_snd$
 $\exists Task$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(que? \mapsto q_size(que?) - 1)\}$
 $q_ava' = q_ava$
 $\exists WaitingData$
 $release_snd' = release_snd$
 $release_rcv' = \{self?\} \triangleleft release_rcv$
 $topReady! = self?$

QueueReceiveE_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_snd$
 $self? \in \text{dom } release_rcv \Rightarrow que? = release_rcv(self?)$
 $que? \in queue$
 $q_size(que?) = 0$
 $state(self?) = running$
 $executable(self?) \in q_ava(que?)$
 $self? \notin idles$
 $findTopReady[self?/target?]$
 $\exists st? : STATE \mid st? = blocked$

- $Reschedule[topReady!/target?, tasks/tasks?, executable/executable?, priority/pri?]$

 $\exists QueueData$
 $wait_snd' = wait_snd$
 $wait_rcv' = wait_rcv \oplus \{(self? \mapsto que?)\}$
 $release_snd' = release_snd$
 $release_rcv' = \{self?\} \triangleleft release_rcv$

QueueReceiveW_TQ

$\Delta TaskQueue$
 $que? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_snd$
 $self? \in \text{dom } release_rcv \Rightarrow que? = release_rcv(self?)$
 $que? \in queue$
 $state(self?) = running$

$$\begin{aligned}
& executable(self?) \in q_ava(que?) \\
& topReady! \in wait_snd \sim (\{ que? \} \mid) \\
& \forall ws : wait_snd \sim (\{ que? \} \mid) \bullet priority(topReady!) \geq priority(ws) \\
& priority(topReady!) \leq priority(running_tasks(executable(topReady!))) \\
& \exists TaskData \\
& state' = state \oplus \{(topReady! \mapsto ready)\} \\
& \exists ContextData \\
& \exists PrioData \\
& queue' = queue \\
& q_max' = q_max \\
& q_size' = q_size \oplus \{(que? \mapsto q_max(que?) - 1)\} \\
& q_ava' = q_ava \\
& wait_snd' = \{topReady!\} \triangleleft wait_snd \\
& wait_rcv' = wait_rcv \\
& release_snd' = release_snd \oplus \{(self? \mapsto que?)\} \\
& release_rcv' = \{self?\} \triangleleft release_rcv
\end{aligned}$$

QueueReceiveWS_TQ

$$\begin{aligned}
& \Delta TaskQueue \\
& que? : QUEUE \\
& self? : TASK \\
& topReady! : TASK
\end{aligned}$$

$$\begin{aligned}
& self? \notin \text{dom } release_snd \\
& self? \in \text{dom } release_rcv \Rightarrow que? = release_rcv(self?) \\
& que? \in queue \\
& state(self?) = running \\
& executable(self?) \in q_ava(que?) \\
& self? \notin idles \\
& topReady! \in wait_snd \sim (\{ que? \} \mid) \\
& \forall ws : wait_snd \sim (\{ que? \} \mid) \bullet priority(topReady!) \geq priority(ws) \\
& priority(topReady!) > priority(running_tasks(executable(topReady!))) \\
& \exists st? : STATE \mid st? = ready \\
& \quad \bullet Reschedule[topReady!/target?, tasks/tasks?, \\
& \quad \quad \quad executable/executable?, priority/pri?] \\
& queue' = queue \\
& q_max' = q_max \\
& q_size' = q_size \oplus \{(que? \mapsto q_max(que?) - 1)\} \\
& q_ava' = q_ava \\
& wait_snd' = \{topReady!\} \triangleleft wait_snd \\
& wait_rcv' = wait_rcv \\
& release_snd' = release_snd \oplus \{(topReady! \mapsto que?)\} \\
& release_rcv' = \{self?\} \triangleleft release_rcv
\end{aligned}$$

$$\begin{aligned}
\text{QueueReceive_TQ} &\cong \text{QueueReceiveN_TQ} \\
&\vee \text{QueueReceiveE_TQ} \\
&\vee \text{QueueReceiveW_TQ} \\
&\vee \text{QueueReceiveWS_TQ}
\end{aligned}$$

ChangeQueueLevel_TQ

$\Delta \text{TaskQueue}$
 $que? : \text{QUEUE}$
 $self? : \text{TASK}$
 $cset? : \mathbb{F} \text{ cores}$

$self? \notin \text{dom release_snd} \cup \text{dom release_rcv}$
 $que? \in \text{queue}$
 $\text{state}(self?) = \text{running}$
 $\text{executable}(self?) \in q_ava(que?)$
 $cset? \neq q_ava(que?)$
 $cset? \neq \emptyset$
 $\forall t : \text{wait_rcv} \sim (\{ que? \}) \cup \text{wait_snd} \sim (\{ que? \})$

- $\text{executable}(t) \in cset?$

$\exists \text{Task}$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size$
 $q_ava' = q_ava \oplus \{(que? \mapsto cset?)\}$
 $\exists \text{WaitingData}$
 $\exists \text{QReleasingData}$

Appendix J

SPECIFICATION FOR MULTI-CORE TIME

MODEL

Time

$clock : \mathbb{N}$
 $delayed_task : \mathbb{P} \text{ TASK}$
 $time : \text{ TASK} \rightarrow \mathbb{N}$
 $time_slice : \mathbb{N}$

$\forall t : \text{ dom } time \bullet time(t) \geq clock$

Init_Time

Time'

$clock' = 0$
 $delayed_task' = \emptyset$
 $time' = \emptyset$
 $time_slice' = slice_delay$

TaskQueueTime

TaskQueue
Time

$\langle delayed_task, \text{ dom } wait_snd, \text{ dom } wait_rcv \rangle \text{ partition } \text{ dom } time$
 $delayed_task \subseteq state \sim (\{ blocked \})$

$\frac{\text{Init_TaskQueueTime}}{\text{TaskQueueTime}'}$
Init_TaskQueue Init_Time

$$\Delta \text{TaskQueueTime} \hat{=} \text{TaskQueueTime} \wedge \text{TaskQueueTime}' \wedge \Delta \text{Task}$$

$\frac{\text{ExtendTaskQueueXi}}{\Delta \text{TaskQueueTime}}$
$\exists \text{Time}$

$$\text{CreateTask_TQT} \hat{=} \text{ExtendTaskQueueXi} \wedge \text{CreateTask_TQ}$$

$\frac{\text{DeleteTaskN_TQT}}{\text{DeleteTaskN_TQ}} \quad \Delta \text{TaskQueueTime}$
$\text{clock}' = \text{clock}$ $\text{delayed_task}' = \text{delayed_task} \setminus \{\text{target?}\}$ $\text{time}' = \{\text{target?}\} \triangleleft \text{time}$ $\text{time_slice}' = \text{time_slice}$

$$\text{DeleteTaskS_TQT} \hat{=} \text{ExtendTaskQueueXi} \wedge \text{DeleteTaskS_TQ}$$

$$\text{DeleteTask_TQT} \hat{=} \text{DeleteTaskN_TQT} \vee \text{DeleteTaskS_TQT}$$

$$\text{ExecuteRunningTask_TQT} \hat{=} \text{ExtendTaskQueueXi} \wedge \text{ExecuteRunningTask_TQ}$$

$\frac{\text{SuspendTaskN_TQT}}{\text{SuspendTaskN_TQ}} \quad \Delta \text{TaskQueueTime}$
$\text{clock}' = \text{clock}$ $\text{delayed_task}' = \text{delayed_task} \setminus \{\text{target?}\}$ $\text{time}' = \{\text{target?}\} \triangleleft \text{time}$ $\text{time_slice}' = \text{time_slice}$

$$\text{SuspendTask_TQT} \triangleq \text{SuspendTaskN_TQT} \vee \\ (\text{ExtendTaskQueueXi} \wedge (\text{SuspendTaskS_T} \vee \text{SuspendTaskO_T}))$$

$$\text{ResumeTask_TQT} \triangleq \text{ExtendTaskQueueXi} \wedge \text{ResumeTask_TQ}$$

$$\text{ChangeTaskPriority_TQT} \triangleq \text{ExtendTaskQueueXi} \\ \wedge \text{ChangeTaskPriority_TQ}$$

$$\text{Migration_TQT} \triangleq \text{ExtendTaskQueueXi} \wedge \text{Migration_TQ}$$

$$\text{CreateQueue_TQT} \triangleq \text{ExtendTaskQueueXi} \wedge \text{CreateQueue_TQ}$$

$$\text{DeleteQueue_TQT} \triangleq \text{ExtendTaskQueueXi} \wedge \text{DeleteQueue_TQ}$$

$$\text{QueueSendN_TQT} \triangleq \text{ExtendTaskQueueXi} \wedge \text{QueueSendN_TQ}$$

QueueSendF_TQT <hr/> QueueSendF_TQ $\Delta \text{TaskQueueTime}$ $n? : \mathbb{N}$
$n? > \text{clock}$ $\text{clock}' = \text{clock}$ $\text{delayed_task}' = \text{delayed_task}$ $\text{time}' = \text{time} \oplus \{\text{self?} \mapsto n?\}$ $\text{time_slice}' = \text{time_slice}$

QueueSendW_TQT <hr/> QueueSendW_TQ $\Delta \text{TaskQueueTime}$
$\text{clock}' = \text{clock}$ $\text{delayed_task}' = \text{delayed_task}$ $\text{time}' = \{\text{topReady!}\} \triangleleft \text{time}$ $\text{time_slice}' = \text{time_slice}$

QueueSendWS_TQT <hr/> QueueSendWS_TQ $\Delta \text{TaskQueueTime}$
$\text{clock}' = \text{clock}$

$$\begin{array}{l}
\text{delayed_task}' = \text{delayed_task} \\
\text{time}' = \{\text{topReady!}\} \triangleleft \text{time} \\
\text{time_slice}' = \text{time_slice}
\end{array}$$

$$\begin{array}{l}
\text{QueueSend_TQT} \cong \text{QueueSendN_TQT} \\
\quad \vee \text{QueueSendF_TQT} \\
\quad \vee \text{QueueSendW_TQT} \\
\quad \vee \text{QueueSendWS_TQT}
\end{array}$$

$$\text{QueueReceiveN_TQT} \cong \text{ExtendTaskQueueXi} \wedge \text{QueueReceiveN_TQ}$$

$$\text{QueueReceiveE_TQT}$$

$$\begin{array}{l}
\text{QueueReceiveE_TQ} \\
\Delta \text{TaskQueueTime} \\
n? : \mathbb{N}
\end{array}$$

$$\begin{array}{l}
n? > \text{clock} \\
\text{clock}' = \text{clock} \\
\text{delayed_task}' = \text{delayed_task} \\
\text{time}' = \text{time} \oplus \{\text{self?} \mapsto n?\} \\
\text{time_slice}' = \text{time_slice}
\end{array}$$

$$\text{QueueReceiveW_TQT}$$

$$\begin{array}{l}
\text{QueueReceiveW_TQ} \\
\Delta \text{TaskQueueTime}
\end{array}$$

$$\begin{array}{l}
\text{clock}' = \text{clock} \\
\text{delayed_task}' = \text{delayed_task} \\
\text{time}' = \{\text{topReady!}\} \triangleleft \text{time} \\
\text{time_slice}' = \text{time_slice}
\end{array}$$

$$\text{QueueReceiveWS_TQT}$$

$$\begin{array}{l}
\text{QueueReceiveWS_TQ} \\
\Delta \text{TaskQueueTime}
\end{array}$$

$$\begin{array}{l}
\text{clock}' = \text{clock} \\
\text{delayed_task}' = \text{delayed_task} \\
\text{time}' = \{\text{topReady!}\} \triangleleft \text{time} \\
\text{time_slice}' = \text{time_slice}
\end{array}$$

$$\begin{array}{l}
\text{QueueReceive_TQT} \cong \text{QueueReceiveN_TQT} \vee \text{QueueReceiveE_TQT} \\
\quad \vee \text{QueueReceiveW_TQT} \vee \text{QueueReceiveWS_TQT}
\end{array}$$

$ChangeQueueLevel_TQT \hat{=} ExtendTaskQueueXi \wedge ChangeQueueLevel_TQ$

DelayUntil_TQT

$\Delta TaskQueueTime$

$n? : \mathbb{N}$

$self? : TASK$

$topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$state(self?) = running$

$self? \notin idles$

$findTopReady[self?/target?]$

$n? > clock$

$\exists st? : STATE \mid st? = blocked$

- $Reschedule[topReady!/target?, tasks/tasks?, executable/executable?, priority/pri?]$

$\exists Queue$

$clock' = clock$

$delayed_task' = delayed_task \cup \{self?\}$

$time' = time \oplus \{(self? \mapsto n?)\}$

$time_slice' = time_slice$

CheckDelayedTaskN_TQT

$\Delta TaskQueueTime$

$topWaiting! : TASK$

$self? : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$state(self?) = running$

$topWaiting! \in \text{dom } time$

$\forall wt : \text{dom } time \bullet time(topWaiting!) \leq time(wt)$

$\forall wt : \text{dom } time \mid time(wt) = time(topWaiting!)$

- $priority(topWaiting!) \geq priority(wt)$

$priority(topWaiting!)$

$\leq priority(running_tasks(executable(topWaiting!)))$

$\exists TaskData$

$state' = state \oplus \{(topWaiting! \mapsto ready)\}$

$\exists ContextData$

$\exists PrioData$

$\exists QueueData$

$wait_snd' = \{topWaiting!\} \triangleleft wait_snd$

$wait_rcv' = \{topWaiting!\} \triangleleft wait_rcv$

$\exists QReleasingData$

$clock' = time(topWaiting!)$

$delayed_task' = delayed_task \setminus \{topWaiting!\}$

$$time' = \{topWaiting!\} \triangleleft time$$

$$time_slice' = time_slice$$

CheckDelayedTaskS_TQT

$$\Delta TaskQueueTime$$

$$topWaiting! : TASK$$

$$self? : TASK$$

$$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$$

$$state(self?) = running$$

$$topWaiting! \in \text{dom } time$$

$$\forall wt : \text{dom } time \bullet time(topWaiting!) \leq time(wt)$$

$$\forall wt : \text{dom } time \mid time(wt) = time(topWaiting!) \bullet$$

- $priority(topWaiting!) \geq priority(wt)$

$$priority(topWaiting!) > priority(running_tasks(executable(topWaiting!)))$$

$$\exists st? : STATE \mid st? = ready$$

- $Reschedule[topWaiting!/target?, tasks/tasks?, executable/executable?, priority/pri?]$

$$\exists QueueData$$

$$wait_snd' = \{topWaiting!\} \triangleleft wait_snd$$

$$wait_rcv' = \{topWaiting!\} \triangleleft wait_rcv$$

$$\exists QReleasingData$$

$$clock' = time(topWaiting!)$$

$$delayed_task' = delayed_task \setminus \{topWaiting!\}$$

$$time' = \{topWaiting!\} \triangleleft time$$

$$time_slice' = time_slice$$

$$CheckDelayedTask_TQT \cong CheckDelayedTaskN_TQT$$

$$\vee CheckDelayedTaskS_TQT$$

TimeSlicing_TQT

$$\Delta TaskQueueTime$$

$$topReadys! : \mathbb{F} TASK$$

$$\#topReadys! \leq \#cores$$

$$\forall t : topReadys! \bullet state(t) = ready$$

$$\wedge priority(t) = priority(running_tasks(executable(t)))$$

$$\forall t1, t2 : topReadys! \mid executable(t1) = executable(t2) \bullet t1 = t2$$

$$\forall c : cores \mid (\forall t : topReadys! \bullet executable(t) \neq c)$$

- $(\forall t : executable \sim (\{c\}) \mid state(t) = ready$
- $priority(t) < priority(running_tasks(c))$

$$topReadys! \neq \emptyset$$

$$\forall t : \text{dom } time \bullet time_slice \leq time(t)$$

$$\begin{aligned}
& tasks' = tasks \\
& executable(\topReadys! \mid) \triangleleft running_tasks' \\
& \quad = executable(\topReadys! \mid) \triangleleft running_tasks \\
& executable' = executable \\
& (running_tasks(\executable(\topReadys! \mid) \mid) \cup \topReadys!) \triangleleft state' \\
& \quad = (running_tasks(\executable(\topReadys! \mid) \mid) \cup \topReadys!) \triangleleft state \\
& executable(\topReadys! \mid) \triangleleft phys_context' \\
& \quad = executable(\topReadys! \mid) \triangleleft phys_context \\
& running_tasks(\executable(\topReadys! \mid) \mid) \triangleleft log_context' \\
& \quad = running_tasks(\executable(\topReadys! \mid) \mid) \triangleleft log_context \\
& priority' = priority \\
& \forall trt : \topReadys! \\
& \quad \bullet running_tasks'(\executable(trt)) = trt \\
& \quad \quad \wedge state'(trt) = running \\
& \quad \quad \wedge state'(running_tasks(\executable(trt))) = ready \\
& \quad \quad \wedge phys_context'(\executable(trt)) = log_context(trt) \\
& \quad \quad \wedge log_context'(running_tasks(\executable(trt))) \\
& \quad \quad \quad = phys_context(\executable(trt)) \\
& \exists Queue \\
& clock' = clock \\
& delayed_task' = delayed_task \\
& time' = time \\
& time_slice' = time_slice + slice_delay
\end{aligned}$$

NoSlicing_TQT

Δ TaskQueueTime

$$\begin{aligned}
& \forall c : cores \\
& \quad \bullet \forall rt : state \sim (\{ready\} \mid) \mid executable(rt) = c \\
& \quad \quad \bullet priority(rt) < priority(running_tasks(c)) \\
& \forall t : \text{dom } time \bullet time_slice \leq time(t) \\
& \exists TaskQueue \\
& clock' = clock \\
& delayed_task' = delayed_task \\
& time' = time \\
& time_slice' = time_slice + slice_delay
\end{aligned}$$

Appendix K

SPECIFICATION FOR MULTI-CORE

MUTEX MODEL

MutexData

$semaphore : \mathbb{P} \text{ QUEUE}$
 $mutex : \mathbb{P} \text{ QUEUE}$
 $mutex_holder : \text{ QUEUE} \rightarrow \text{ TASK}$
 $mutex_recursive : \text{ QUEUE} \rightarrow \mathbb{N}$

$mutex \cap semaphore = \emptyset$
 $\text{dom } mutex_recursive = mutex$
 $\forall m : mutex \bullet m \notin \text{dom } mutex_holder \Leftrightarrow mutex_recursive(m) = 0$

Init_MutexData

MutexData'

$semaphore' = \emptyset$
 $mutex' = \emptyset$
 $mutex_holder' = \emptyset$
 $mutex_recursive' = \emptyset$

OriginalPrioData

$base_priority : \text{ TASK} \rightarrow \mathbb{N}$

Init_OriginalPrioData

OriginalPrioData'

$base_priority' = \emptyset$

$MReleasingData$

$release_mutex : TASK \rightarrow QUEUE$

$Init_MReleasingData$

$MReleasingData'$

$release_mutex' = \emptyset$

$Mutex$

$MutexData$

$OriginalPrioData$

$MReleasingData$

$dom\ base_priority = ran\ mutex_holder$

$ran\ release_mutex \subseteq mutex$

$Init_Mutex$

$Mutex'$

$Init_MutexData$

$Init_OriginalPrioData$

$Init_MReleasingData$

$TaskQueueTimeMutex$

$TaskQueueTime$

$Mutex$

$semaphore \subseteq queue$

$\forall s : semaphore \bullet q_max(s) = 1$

$mutex \subseteq queue$

$\forall m : mutex \bullet q_max(m) = 1$

$dom\ mutex_holder = \{ m : mutex \mid q_size(m) = 0 \}$

$\forall m : dom\ mutex_holder \bullet executable(mutex_holder(m)) \in q_ava(m)$

$\forall mh : ran\ mutex_holder \bullet priority(mh) \geq base_priority(mh)$

$\forall ms : mutex \cup semaphore \bullet ms \notin ran\ wait_snd \cup ran\ release_snd$

$release_mutex \subseteq release_rcv$

$Init$

$TaskQueueTimeMutex'$

$Init_TaskQueueTime$ $Init_Mutex$
--

$ExtendTQTXi$ $\Delta TaskQueueTimeMutex$
--

$\exists Mutex$

$CreateTask_TQTM \cong ExtendTQTXi \wedge CreateTask_TQT$

$DeleteTask_TQTM$ $DeleteTask_TQT$ $\Delta TaskQueueTimeMutex$
--

$target? \notin \text{ran } mutex_holder$ $\exists MutexData$ $\exists OriginalPrioData$ $release_mutex' = \{target?\} \triangleleft release_mutex$

$ExecuteRunningTask_TQTM \cong ExtendTQTXi \wedge ExecuteRunningTask_TQT$

$SuspendTask_TQTM \cong ExtendTQTXi \wedge SuspendTask_TQT$

$ResumeTask_TQTM \cong ExtendTQTXi \wedge ResumeTask_TQT$

$ChangeTaskPriorityNotHolder_TQTM$ $ChangeTaskPriority_TQT$ $\Delta TaskQueueTimeMutex$

$target? \notin \text{dom } base_priority$ $\exists Mutex$
--

$ChangeTaskPriorityNotInherited_TQTM$ $ChangeTaskPriority_TQT$ $\Delta TaskQueueTimeMutex$
--

$target? \in \text{dom } base_priority$ $base_priority(target?) = priority(target?)$ $\exists MutexData$ $base_priority' = base_priority \oplus \{(target? \mapsto newpri?)\}$

$\exists MReleasingData$

ChangeTaskPriorityInheritedN_TQTM

$\Delta TaskQueueTimeMutex$

$newpri? : \mathbb{N}$

$target? : TASK$

$self? : TASK$

$topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$state(self?) = running$

$target? \in \text{dom } base_priority$

$base_priority(target?) \neq priority(target?)$

$state(target?) \neq nonexistent$

$newpri? \leq priority(target?)$

$target? \in idles \Rightarrow newpri? = 0$

$\exists TaskQueueTime$

$\exists MutexData$

$base_priority' = base_priority \oplus \{(target? \mapsto newpri?)\}$

$\exists MReleasingData$

$topReady! = self?$

ChangeTaskPriorityInheritedU_TQTM

ChangeTaskPriorityN_T

$\Delta TaskQueueTimeMutex$

$target? \in \text{dom } base_priority$

$base_priority(target?) \neq priority(target?)$

$newpri? > priority(target?)$

$\exists Queue$

$\exists Time$

$\exists MutexData$

$base_priority' = base_priority \oplus \{(target? \mapsto newpri?)\}$

$\exists MReleasingData$

ChangeTaskPriorityInheritedS_TQTM

ChangeTaskPriorityS_T

$\Delta TaskQueueTimeMutex$

$target? \in \text{dom } base_priority$

$base_priority(target?) \neq priority(target?)$

$newpri? > priority(target?)$

$\exists Queue$

$\exists Time$

$\exists \text{MutexData}$ $\text{base_priority}' = \text{base_priority} \oplus \{(target? \mapsto \text{newpri?})\}$ $\exists \text{MReleasingData}$

$$\begin{aligned} \text{ChangeTaskPriority_TQTM} &\hat{=} \text{ChangeTaskPriorityNotHolder_TQTM} \\ &\vee \text{ChangeTaskPriorityNotInherited_TQTM} \\ &\vee \text{ChangeTaskPriorityInheritedN_TQTM} \\ &\vee \text{ChangeTaskPriorityInheritedU_TQTM} \\ &\vee \text{ChangeTaskPriorityInheritedS_TQTM} \end{aligned}$$

$$\text{Migration_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{Migration_TQT}$$

$$\text{CreateQueue_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{CreateQueue_TQT}$$

DeleteQueue_TQTM DeleteQueue_TQT $\Delta \text{TaskQueueTimeMutex}$
$que? \notin \text{semaphore} \cup \text{mutex}$ $\exists \text{Mutex}$

QueueSend_TQTM QueueSend_TQT $\Delta \text{TaskQueueTimeMutex}$
$que? \notin \text{mutex} \cup \text{mutex}$ $\exists \text{Mutex}$

QueueReceive_TQTM QueueReceive_TQT $\Delta \text{TaskQueueTimeMutex}$
$\text{self?} \notin \text{dom release_mutex}$ $que? \notin \text{mutex} \cup \text{semaphore}$ $\exists \text{Mutex}$

$$\text{ChangeQueueLevel_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{ChangeQueueLevel_TQT}$$

$$\text{DelayUntil_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{DelayUntil_TQT}$$

$$\text{CheckDelayedTask_TQTM} \hat{=} \text{ExtendTQTXi} \wedge \text{CheckDelayedTask_TQT}$$

$TimeSlicing_TQTM \cong ExtendTQTXi \wedge TimeSlicing_TQT$

$NoSlicing_TQTM \cong ExtendTQTXi \wedge NoSlicing_TQT$

$CreateBinarySemaphore_TQTM$

$\Delta TaskQueueTimeMutex$

$sem? : QUEUE$

$self? : TASK$

$cset? : \mathbb{F} cores$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$state(self?) = running$

$sem? \notin queue$

$cset? \neq \emptyset$

$\exists Task$

$queue' = queue \cup \{sem?\}$

$q_max' = q_max \oplus \{(sem? \mapsto 1)\}$

$q_size' = q_size \oplus \{(sem? \mapsto 1)\}$

$q_ava' = q_ava \oplus \{(sem? \mapsto cset?)\}$

$\exists WaitingData$

$\exists QReleasingData$

$\exists Time$

$semaphore' = semaphore \cup \{sem?\}$

$mutex' = mutex$

$mutex_holder' = mutex_holder$

$mutex_recursive' = mutex_recursive$

$\exists OriginalPrioData$

$\exists MReleasingData$

$DeleteBinarySemaphore_TQTM$

$\Delta TaskQueueTimeMutex$

$sem? : QUEUE$

$self? : TASK$

$sem? \in semaphore$

$DeleteQueue_TQT[sem?/que?]$

$semaphore' = semaphore \setminus \{sem?\}$

$mutex' = mutex$

$mutex_holder' = mutex_holder$

$mutex_recursive' = mutex_recursive$

$\exists OriginalPrioData$

$\exists MReleasingData$

CreateMutex_TQTM

Δ TaskQueueTimeMutex

$mut? : QUEUE$

$self? : TASK$

$cset? : \mathbb{F} \text{ cores}$

$self? \notin \text{dom release_snd} \cup \text{dom release_rcv}$

$state(self?) = \text{running}$

$mut? \notin \text{queue}$

$cset? \neq \emptyset$

$\exists Task$

$queue' = \text{queue} \cup \{mut?\}$

$q_max' = q_max \oplus \{(mut? \mapsto 1)\}$

$q_size' = q_size \oplus \{(mut? \mapsto 1)\}$

$q_ava' = q_ava \oplus \{(mut? \mapsto cset?)\}$

$\exists WaitingData$

$\exists Time$

$semaphore' = semaphore$

$mutex' = mutex \cup \{mut?\}$

$mutex_holder' = mutex_holder$

$mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$

$\exists OriginalPrioData$

$\exists MReleasingData$

DeleteMutex_TQTM

Δ TaskQueueTimeMutex

$mut? : QUEUE$

$self? : TASK$

$mut? \in mutex \setminus \text{dom mutex_holder}$

$DeleteQueue_TQT[mut?/que?]$

$semaphore' = semaphore$

$mutex' = mutex \setminus \{mut?\}$

$mutex_holder' = mutex_holder$

$mutex_recursive' = \{mut?\} \triangleleft mutex_recursive$

$\exists OriginalPrioData$

$\exists MReleasingData$

MutexTakeNnonInh_TQTM

Δ TaskQueueTimeMutex

$mut? : QUEUE$

$self? : TASK$

$topReady! : TASK$

$self? \in \text{dom release_rcv} \Rightarrow self? \in \text{dom release_mutex}$

$mut? \in mutex$
 $self? \notin \text{dom } base_priority$
 $QueueReceiveN_TQT[mut?/que?]$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = mutex_holder \oplus \{(mut? \mapsto self?)\}$
 $mutex_recursive' = mutex_recursive \oplus$
 $\quad \{(mut? \mapsto mutex_recursive(mut?) + 1)\}$
 $base_priority' = base_priority \oplus \{(self? \mapsto priority(self?))\}$
 $release_mutex' = \{self?\} \triangleleft release_mutex$

MutexTakeNInh_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \in \text{dom } release_rcv \Rightarrow self? \in \text{dom } release_mutex$
 $mut? \in mutex$
 $self? \in \text{dom } base_priority$
 $QueueReceiveN_TQT[mut?/que?]$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = mutex_holder \oplus \{(mut? \mapsto self?)\}$
 $mutex_recursive' = mutex_recursive \oplus$
 $\quad \{(mut? \mapsto mutex_recursive(mut?) + 1)\}$
 $\exists OriginalPrioData$
 $release_mutex' = \{self?\} \triangleleft release_mutex$

MutexTakeRecursive_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in mutex$
 $state(self?) = running$
 $executable(self?) \in q_ava(mut?)$
 $self? = mutex_holder(mut?)$
 $\exists TaskQueueTime$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus$

$$\{(mut? \mapsto mutex_recursive(mut?) + 1)\}$$

$\exists OriginalPrioData$
 $\exists MReleasingData$
 $topReady! = self?$

MutexTakeEnonInh_TQTM

$\Delta TaskQueueTimeMutex$
 $mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$
 $n? : \mathbb{N}$

$self? \in \text{dom } release_rcv \Rightarrow self? \in \text{dom } release_mutex$
 $mut? \in \text{dom } mutex_holder$
 $priority(self?) \leq priority(mutex_holder(mut?))$
 $self? \neq mutex_holder(mut?)$
 $QueueReceiveE_TQT[mut?/que?]$
 $\exists MutexData$
 $\exists OriginalPrioData$
 $release_mutex' = \{self?\} \triangleleft release_mutex$

MutexTakeEInheritSameCoreHolder_TQTM

$\Delta TaskQueueTimeMutex$
 $mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$
 $n? : \mathbb{N}$

$self? \notin \text{dom } release_snd$
 $self? \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(self?)$
 $self? \in \text{dom } release_rcv \Rightarrow self? \in \text{dom } release_mutex$
 $mut? \in \text{dom } mutex_holder$
 $state(self?) = running$
 $executable(self?) \in q_ava(mut?)$
 $priority(self?) > priority(mutex_holder(mut?))$
 $n? > clock$
 $executable(mutex_holder(mut?)) = executable(self?)$
 $mutex_holder(mut?) \in state \sim (\{ready\})$
 $topReady! = mutex_holder(mut?)$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $\quad | st? = blocked$
 $\quad \wedge pri? = priority \oplus \{(topReady! \mapsto priority(self?))\}$
 $\quad \bullet Reschedule[topReady!/target?, tasks/tasks?, executable/executable?]$
 $\exists QueueData$
 $wait_snd' = wait_snd$

$wait_rcv' = wait_rcv \oplus \{(self? \mapsto mut?)\}$
 $release_snd' = release_snd$
 $release_rcv' = \{self?\} \triangleleft release_rcv$
 $clock' = clock$
 $delayed_task' = delayed_task$
 $time' = time \oplus \{(self? \mapsto n?)\}$
 $time_slice' = time_slice$
 $\exists MutexData$
 $\exists OriginalPrioData$
 $release_mutex' = \{self?\} \triangleleft release_mutex$

MutexTakeEInheritSameCoreReady_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$
 $n? : \mathbb{N}$

$self? \notin \text{dom } release_snd$
 $self? \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(self?)$
 $self? \in \text{dom } release_rcv \Rightarrow self? \in \text{dom } release_mutex$
 $mut? \in \text{dom } mutex_holder$
 $state(self?) = running$
 $executable(self?) \in q_ava(mut?)$
 $priority(self?) > priority(mutex_holder(mut?))$
 $n? > clock$
 $executable(mutex_holder(mut?)) = executable(self?)$
 $mutex_holder(mut?) \notin state \sim (\{ready\})$
 $findTopReady[self?/target?]$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $\quad | st? = blocked$
 $\quad \quad \wedge pri? = priority \oplus \{(mutex_holder(mut?) \mapsto priority(self?))\}$
 $\quad \quad \bullet Reschedule[topReady!/target?, tasks/tasks?, executable/executable?]$
 $\exists QueueData$
 $wait_snd' = wait_snd$
 $wait_rcv' = wait_rcv \oplus \{(self? \mapsto mut?)\}$
 $release_snd' = release_snd$
 $release_rcv' = \{self?\} \triangleleft release_rcv$
 $clock' = clock$
 $delayed_task' = delayed_task$
 $time' = time \oplus \{(self? \mapsto n?)\}$
 $time_slice' = time_slice$
 $\exists MutexData$
 $\exists OriginalPrioData$
 $release_mutex' = \{self?\} \triangleleft release_mutex$

MutexTakeEInheritDiffCoreN_TQTM

 Δ *TaskQueueTimeMutex* $mut? : QUEUE$ $self? : TASK$ $topReady! : TASK$ $n? : \mathbb{N}$ $self? \notin \text{dom } release_snd$ $self? \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(self?)$ $self? \in \text{dom } release_rcv \Rightarrow self? \in \text{dom } release_mutex$ $mut? \in \text{dom } mutex_holder$ $state(self?) = running$ $executable(self?) \in q_ava(mut?)$ $priority(self?) > priority(mutex_holder(mut?))$ $n? > clock$ $executable(mutex_holder(mut?)) \neq executable(self?)$ $priority(self?) >$ $\quad priority(running_tasks(executable(mutex_holder(mut?))))$ $\quad \Rightarrow mutex_holder(mut?) \notin state \sim (\{ready\})$ $findTopReady[self?/target?]$ $tasks' = tasks$ $running_tasks' = running_tasks \oplus \{(executable(self?) \mapsto topReady!)\}$ $executable' = executable$ $state' = state \oplus \{(self? \mapsto blocked),$ $\quad (topReady! \mapsto running),$ $\quad (mutex_holder(mut?) \mapsto ready)\}$ $phys_context' = phys_context \oplus$ $\quad \{(executable(self?) \mapsto log_context(topReady!))\}$ $log_context' = log_context \oplus \{(self? \mapsto phys_context(executable(self?)))\}$ $priority' = priority \oplus \{(mutex_holder(mut?) \mapsto priority(self?))\}$ $\exists QueueData$ $wait_snd' = wait_snd$ $wait_rcv' = wait_rcv \oplus \{(self? \mapsto mut?)\}$ $release_snd' = release_snd$ $release_rcv' = \{self?\} \triangleleft release_rcv$ $clock' = clock$ $delayed_task' = delayed_task$ $time' = time \oplus \{(self? \mapsto n?)\}$ $time_slice' = time_slice$ $\exists MutexData$ $\exists OriginalPrioData$ $release_mutex' = \{self?\} \triangleleft release_mutex$

MutexTakeEInheritDiffCoreS_TQTM

 Δ *TaskQueueTimeMutex* $mut? : QUEUE$

$self? : TASK$

$topReady! : TASK$

$n? : \mathbb{N}$

$self? \notin \text{dom } release_snd$

$self? \in \text{dom } release_rcv \Rightarrow mut? = release_rcv(self?)$

$self? \in \text{dom } release_rcv \Rightarrow self? \in \text{dom } release_mutex$

$mut? \in \text{dom } mutex_holder$

$state(self?) = running$

$executable(self?) \in q_ava(mut?)$

$priority(self?) > priority(mutex_holder(mut?))$

$n? > clock$

$executable(mutex_holder(mut?)) \neq executable(self?)$

$priority(self?) >$

$\quad priority(running_tasks(executable(mutex_holder(mut?))))$

$mutex_holder(mut?) \in state \sim (\{ready\})$

$findTopReady[self?/target?]$

$tasks' = tasks$

$running_tasks' = running_tasks \oplus \{(executable(self?) \mapsto topReady!),$
 $(executable(mutex_holder(mut?)) \mapsto mutex_holder(mut?))\}$

$executable' = executable$

$state' = state \oplus \{(self? \mapsto blocked),$

$(topReady! \mapsto running),$

$(running_tasks(executable(mutex_holder(mut?))) \mapsto ready),$

$(mutex_holder(mut?) \mapsto running)\}$

$phys_context' = phys_context \oplus$

$\{(executable(self?) \mapsto log_context(topReady!)),$

$(executable(mutex_holder(mut?)) \mapsto$

$log_context(mutex_holder(mut?)))\}$

$log_context' = log_context \oplus \{(self? \mapsto phys_context(executable(self?))),$

$(running_tasks(executable(mutex_holder(mut?)))$

$\mapsto phys_context(executable(mutex_holder(mut?))))\}$

$priority' = priority \oplus \{(mutex_holder(mut?) \mapsto priority(self?))\}$

$\exists QueueData$

$wait_snd' = wait_snd$

$wait_rcv' = wait_rcv \oplus \{(self? \mapsto mut?)\}$

$release_snd' = release_snd$

$release_rcv' = \{self?\} \triangleleft release_rcv$

$clock' = clock$

$delayed_task' = delayed_task$

$time' = time \oplus \{(self? \mapsto n?)\}$

$time_slice' = time_slice$

$\exists MutexData$

$\exists OriginalPrioData$

$release_mutex' = \{self?\} \triangleleft release_mutex$

basePriorityMan

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

self? : *TASK*

$self? \in \text{ran}(\{mut?\} \triangleleft mutex_holder) \Rightarrow \exists OriginalPrioData$
 $self? \notin \text{ran}(\{mut?\} \triangleleft mutex_holder)$
 $\Rightarrow base_priority' = \{self?\} \triangleleft base_priority$

MutexGiveNRecursive_TQTM

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

self? : *TASK*

topReady! : *TASK*

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $self? = mutex_holder(mut?)$
 $state(self?) = running$
 $mutex_recursive(mut?) > 1$
 $\exists TaskQueueTime$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus$
 $\quad \{(mut? \mapsto mutex_recursive(mut?) - 1)\}$
 $\exists OriginalPrioData$
 $\exists MReleasingData$
 $topReady! = self?$

MutexGiveNnonInh_TQTM

Δ *TaskQueueTimeMutex*

mut? : *QUEUE*

self? : *TASK*

topReady! : *TASK*

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $self? = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $base_priority(self?) = priority(self?)$
 $QueueSendN_TQT[mut?/que?]$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$

$mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
basePriorityMan
 $\exists MReleasingData$

MutexGiveNInhN_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$self? : TASK$

$topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $self? = mutex_holder(mut?)$
 $state(self?) = running$
 $mutex_recursive(mut?) = 1$
 $mut? \notin \text{ran } wait_rcv$
 $base_priority(self?) \neq priority(self?)$
 $\forall rt : state \sim (\{ready\} \mid executable(rt) = executable(self?))$

- $base_priority(self?) \geq priority(rt)$

$\exists TaskData$

$\exists StateData$

$\exists ContextData$

$priority' = priority \oplus \{(self? \mapsto base_priority(self?))\}$

$queue' = queue$

$q_max' = q_max$

$q_size' = q_size \oplus \{(mut? \mapsto 1)\}$

$q_ava' = q_ava$

$\exists WaitingData$

$\exists QReleasingData$

$semaphore' = semaphore$

$mutex' = mutex$

$mutex_holder' = \{mut?\} \triangleleft mutex_holder$

$mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$

basePriorityMan

$\exists MReleasingData$

$topReady! = self?$

MutexGiveNInhS_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$self? : TASK$

$topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$mut? \in \text{dom } mutex_holder$

$self? = mutex_holder(mut?)$
 $state(self?) = running$
 $mutex_recursive(mut?) = 1$
 $mut? \notin ran\ wait_rcv$
 $base_priority(self?) \neq priority(self?)$
 $findTopReady[self?/target?]$
 $base_priority(self?) < priority(topReady!)$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $\quad | st? = ready \wedge pri? = priority \oplus \{(self? \mapsto base_priority(self?))\}$
 $\quad \bullet Reschedule[topReady!/target?, tasks/tasks?, executable/executable?]$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(mut? \mapsto 1)\}$
 $q_ava' = q_ava$
 $\exists WaitingData$
 $\exists QReleasingData$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $basePriorityMan$
 $\exists MReleasingData$

MutexGiveWnonInhN_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$self? : TASK$

$topReady! : TASK$

$self? \notin dom\ release_snd \cup dom\ release_rcv$
 $mut? \in dom\ mutex_holder$
 $self? = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $base_priority(self?) = priority(self?)$
 $QueueSendW_TQT[mut?/que?]$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $basePriorityMan$
 $release_mutex' = release_mutex \oplus \{(topReady! \mapsto mut?)\}$

MutexGiveWnonInhS_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $self? = mutex_holder(mut?)$
 $mutex_recursive(mut?) = 1$
 $base_priority(self?) = priority(self?)$
 $QueueSendWS_TQT[mut?/que?]$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $basePriorityMan$
 $release_mutex' = release_mutex \oplus \{(topReady! \mapsto mut?)\}$

MutexGiveWinhN_TQTM

$\Delta TaskQueueTimeMutex$
 $mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $self? = mutex_holder(mut?)$
 $state(self?) = running$
 $mutex_recursive(mut?) = 1$
 $topReady! \in wait_rcv \sim (\{mut?\} \parallel)$
 $\forall wr : wait_rcv \sim (\{mut?\} \parallel) \bullet priority(topReady!) \geq priority(wr)$
 $priority(topReady!) \leq priority(running_tasks(executable(topReady!)))$
 $base_priority(self?) \neq priority(self?)$
 $\forall rt : state \sim (\{ready\} \parallel) \mid executable(self?) = executable(rt)$
 $\bullet base_priority(self?) \geq priority(rt)$
 $executable(self?) = executable(topReady!)$
 $\Rightarrow priority(topReady!) \leq base_priority(self?)$
 $\exists TaskData$
 $state' = state \oplus \{(topReady! \mapsto ready)\}$
 $\exists ContextData$
 $priority' = priority \oplus \{(self? \mapsto base_priority(self?))\}$
 $queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(mut? \mapsto 1)\}$
 $q_ava' = q_ava$
 $wait_snd' = wait_snd$
 $wait_rcv' = \{topReady!\} \triangleleft wait_rcv$

$release_snd' = release_snd$
 $release_rcv' = release_rcv \oplus \{(topReady! \mapsto mut?)\}$
 $clock' = clock$
 $delayed_task' = delayed_task$
 $time' = \{topReady!\} \triangleleft time$
 $time_slice' = time_slice$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $basePriorityMan$
 $release_mutex' = release_mutex \oplus \{(topReady! \mapsto mut?)\}$

MutexGiveWinhSR_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$

$self? : TASK$

$topReady! : TASK$

$topWaiting! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$

$mut? \in \text{dom } mutex_holder$

$self? = mutex_holder(mut?)$

$state(self?) = running$

$mutex_recursive(mut?) = 1$

$topWaiting! \in wait_rcv \sim (\{mut?\})$

$\forall wr : wait_rcv \sim (\{mut?\}) \bullet priority(topWaiting!) \geq priority(wr)$

$findTopReady[self?/target?]$

$base_priority(self?) \neq priority(self?)$

$executable(topWaiting!) \neq executable(self?)$

$\Rightarrow (priority(topWaiting!)$

$\leq priority(running_tasks(executable(topWaiting!)))$

$\wedge base_priority(self?) < priority(topReady!))$

$executable(topWaiting!) = executable(self?)$

$\Rightarrow (priority(topReady!) > base_priority(self?)$

$\wedge priority(topReady!) > priority(topWaiting!))$

$tasks' = tasks$

$running_tasks' = running_tasks \oplus \{(executable(topReady!) \mapsto topReady!)\}$

$executable' = executable$

$state' = state \oplus$

$\{(self? \mapsto ready), (topReady! \mapsto running), (topWaiting! \mapsto ready)\}$

$phys_context' = phys_context \oplus$

$\{(executable(topReady!) \mapsto log_context(topReady!))\}$

$log_context' = log_context \oplus$

$\{(self? \mapsto phys_context(executable(topReady!)))\}$

$priority' = priority \oplus \{(self? \mapsto base_priority(self?))\}$

$queue' = queue$
 $q_max' = q_max$
 $q_size' = q_size \oplus \{(mut? \mapsto 1)\}$
 $q_ava' = q_ava$
 $wait_snd' = wait_snd$
 $wait_rcv' = \{topWaiting!\} \triangleleft wait_rcv$
 $release_snd' = release_snd$
 $release_rcv' = release_rcv \oplus \{(topWaiting! \mapsto mut?)\}$
 $clock' = clock$
 $delayed_task' = delayed_task$
 $time' = \{topWaiting!\} \triangleleft time$
 $time_slice' = time_slice$
 $semaphore' = semaphore$
 $mutex' = mutex$
 $mutex_holder' = \{mut?\} \triangleleft mutex_holder$
 $mutex_recursive' = mutex_recursive \oplus \{(mut? \mapsto 0)\}$
 $basePriorityMan$
 $release_mutex' = release_mutex \oplus \{(topWaiting! \mapsto mut?)\}$

MutexGiveWinhSW_TQTM

$\Delta TaskQueueTimeMutex$

$mut? : QUEUE$
 $self? : TASK$
 $topReady! : TASK$
 $topWaiting! : TASK$

$self? \notin \text{dom } release_snd \cup \text{dom } release_rcv$
 $mut? \in \text{dom } mutex_holder$
 $self? = mutex_holder(mut?)$
 $state(self?) = running$
 $mutex_recursive(mut?) = 1$
 $topWaiting! \in wait_rcv \sim (\{mut?\})$
 $\forall wr : wait_rcv \sim (\{mut?\}) \bullet priority(topWaiting!) \geq priority(wr)$
 $findTopReady[self?/target?]$
 $base_priority(self?) \neq priority(self?)$
 $executable(topWaiting!) \neq executable(self?)$
 $\Rightarrow (priority(topWaiting!) > priority(running_tasks(executable(topWaiting!))) \wedge base_priority(self?) \geq priority(topReady!))$
 $executable(topWaiting!) = executable(self?)$
 $\Rightarrow (priority(topWaiting!) > base_priority(self?) \wedge priority(topWaiting!) \geq priority(topReady!))$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $| st? = ready \wedge pri? = priority \oplus \{(self? \mapsto base_priority(self?))\}$
 $\bullet Reschedule[topWaiting!/target?, tasks/tasks?, executable/executable?]$

$$\begin{aligned}
queue' &= queue \\
q_max' &= q_max \\
q_size' &= q_size \oplus \{(mut? \mapsto 1)\} \\
q_ava' &= q_ava \\
wait_snd' &= wait_snd \\
wait_rcv' &= \{topWaiting!\} \triangleleft wait_rcv \\
release_snd' &= release_snd \\
release_rcv' &= release_rcv \oplus \{(topWaiting! \mapsto mut?)\} \\
clock' &= clock \\
delayed_task' &= delayed_task \\
time' &= \{topWaiting!\} \triangleleft time \\
time_slice' &= time_slice \\
semaphore' &= semaphore \\
mutex' &= mutex \\
mutex_holder' &= \{mut?\} \triangleleft mutex_holder \\
mutex_recursive' &= mutex_recursive \oplus \{(mut? \mapsto 0)\} \\
basePriorityMan & \\
release_mutex' &= release_mutex \oplus \{(topWaiting! \mapsto mut?)\}
\end{aligned}$$

MutexGiveWinhSBoth_TQTM

Δ TaskQueueTimeMutex

mut? : QUEUE

self? : TASK

topReady! : TASK

topWaiting! : TASK

$$\begin{aligned}
self? &\notin \text{dom } release_snd \cup \text{dom } release_rcv \\
mut? &\in \text{dom } mutex_holder \\
self? &= mutex_holder(mut?) \\
state(self?) &= running \\
mutex_recursive(mut?) &= 1 \\
topWaiting! &\in wait_rcv \sim (\{mut?\}) \\
\forall wr : wait_rcv \sim (\{mut?\}) &\bullet priority(topWaiting!) \geq priority(wr) \\
findTopReady[self?/target?] & \\
base_priority(self?) &\neq priority(self?) \\
executable(topWaiting!) &\neq executable(self?) \\
priority(topWaiting!) &> priority(running_tasks(executable(topWaiting!))) \\
priority(topReady!) &> base_priority(self?) \\
tasks' &= tasks \\
running_tasks' &= running_tasks \oplus \\
&\quad \{(executable(topWaiting!) \mapsto topWaiting!), \\
&\quad \quad (executable(topReady!) \mapsto topReady!)\} \\
executable' &= executable \\
state' &= state \oplus \{(self? \mapsto ready), (topReady! \mapsto running), \\
&\quad (topWaiting! \mapsto running), \\
&\quad \quad (running_tasks(executable(topWaiting!)) \mapsto ready)\}
\end{aligned}$$

$$\begin{aligned}
\text{phys_context}' &= \text{phys_context} \oplus \\
&\quad \{(\text{executable}(\text{topWaiting!}) \mapsto \text{log_context}(\text{topWaiting!})), \\
&\quad (\text{executable}(\text{topReady!}) \mapsto \text{log_context}(\text{topReady!}))\} \\
\text{log_context}' &= \text{log_context} \oplus \\
&\quad \{(\text{running_tasks}(\text{executable}(\text{topWaiting!})) \\
&\quad \mapsto \text{phys_context}(\text{executable}(\text{topWaiting!}))), \\
&\quad (\text{self?} \mapsto \text{phys_context}(\text{executable}(\text{self?})))\} \\
\text{priority}' &= \text{priority} \oplus \{(\text{self?} \mapsto \text{base_priority}(\text{self?}))\} \\
\text{queue}' &= \text{queue} \\
\text{q_max}' &= \text{q_max} \\
\text{q_size}' &= \text{q_size} \oplus \{(\text{mut?} \mapsto 1)\} \\
\text{q_ava}' &= \text{q_ava} \\
\text{wait_snd}' &= \text{wait_snd} \\
\text{wait_rcv}' &= \{\text{topWaiting!}\} \triangleleft \text{wait_rcv} \\
\text{release_snd}' &= \text{release_snd} \\
\text{release_rcv}' &= \text{release_rcv} \oplus \{(\text{topWaiting!} \mapsto \text{mut?})\} \\
\text{clock}' &= \text{clock} \\
\text{delayed_task}' &= \text{delayed_task} \\
\text{time}' &= \{\text{topWaiting!}\} \triangleleft \text{time} \\
\text{time_slice}' &= \text{time_slice} \\
\text{semaphore}' &= \text{semaphore} \\
\text{mutex}' &= \text{mutex} \\
\text{mutex_holder}' &= \{\text{mut?}\} \triangleleft \text{mutex_holder} \\
\text{mutex_recursive}' &= \text{mutex_recursive} \oplus \{(\text{mut?} \mapsto 0)\} \\
\text{basePriorityMan} & \\
\text{release_mutex}' &= \text{release_mutex} \oplus \{(\text{topWaiting!} \mapsto \text{mut?})\}
\end{aligned}$$

Appendix L

SPECIFICATION FOR MULTI-CORE TASK MODEL WITH PROMOTION

[*CONTEXT*, *TASK*, *CORE*]

bare_context : *CONTEXT*

idles : \mathbb{F} *TASK*

cores : \mathbb{F} *CORE*

#cores = *#idles*

STATE ::= *nonexistent* | *ready* | *blocked* | *suspended* | *running*

transition == (*blocked* × {*nonexistent*, *ready*, *running*, *suspended*})

∪({*nonexistent*} × {*ready*, *running*})

∪({*ready*} × {*nonexistent*, *running*, *suspended*})

∪({*running*} × {*blocked*, *nonexistent*, *ready*, *suspended*})

∪({*suspended*} × {*nonexistent*, *ready*, *running*})

TaskData

tasks : \mathbb{F} *TASK*

running_task : *TASK*

idle : *TASK*

running_task ∈ *tasks*

idle ∈ *tasks*

idle ∈ *idles*

Init_TaskData

TaskData'

$tasks' = \{idle'\}$
 $running_task' = idle'$
 $idle' \in idles$

StateData

$state : TASK \rightarrow STATE$

ContextData

$phys_context : CONTEXT$
 $log_context : TASK \rightarrow CONTEXT$

Init_ContextData

ContextData'

$phys_context' = bare_context$
 $log_context' = (\lambda x : TASK \bullet bare_context)$

PrioData

$priority : TASK \rightarrow \mathbb{N}$

$\forall i : idles \bullet priority(i) = 0$

Init_PrioData

PrioData'

$priority' = (\lambda x : TASK \bullet 0)$

Task

TaskData

StateData

ContextData

PrioData

$tasks = TASK \setminus (state \sim (\{nonexistent\} \cup))$
 $state \sim (\{running\} \cup) = \{running_task\}$
 $\forall pt : state \sim (\{ready\} \cup) \bullet priority(running_task) \geq priority(pt)$

$\Delta Task$

$Task$
 $Task'$

$\forall st : TASK \mid state'(st) \neq state(st)$
• $state(st) \mapsto state'(st) \in transition$
 $idle' = idle$

$Init_Task$

$Task'$

$Init_TaskData$
 $state' = (\lambda x : TASK \bullet nonexistent) \oplus \{(idle' \mapsto running)\}$
 $Init_ContextData$
 $Init_PrioData$

$Reschedule$

$\Delta Task$

$target? : TASK$
 $tasks? : \mathbb{P} TASK$
 $st? : STATE$
 $pri? : TASK \rightarrow \mathbb{N}$

$tasks' = tasks?$
 $running_task' = target?$
 $state' = state \oplus \{(target? \mapsto running), (running_task \mapsto st?)\}$
 $phys_context' = log_context(target?)$
 $log_context' = log_context \oplus \{(running_task \mapsto phys_context)\}$
 $priority' = pri?$

$CreateTaskN_T$

$\Delta Task$

$target? : TASK$
 $newpri? : \mathbb{N}$

$state(target?) = nonexistent$
 $newpri? \leq priority(running_task)$
 $tasks' = tasks \cup \{target?\}$
 $running_task' = running_task$
 $state' = state \oplus \{(target?, ready)\}$
 $\exists ContextData$
 $priority' = priority \oplus \{(target?, newpri?)\}$

CreateTaskS_T

$\Delta Task$

$target? : TASK$

$newpri? : \mathbb{N}$

$state(target?) = nonexistent$

$newpri? > priority(running_task)$

$\exists st? : STATE; tasks? : \mathbb{P} TASK; pri? : TASK \rightarrow \mathbb{N}$

$| st? = ready \wedge tasks? = tasks \cup \{(target?)\}$

$\wedge pri? = priority \oplus \{(target?, newpri?)\} \bullet Reschedule$

DeleteTaskN_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{ready, blocked, suspended\}$

$tasks' = tasks \setminus \{target?\}$

$running_task' = running_task$

$state' = state \oplus \{(target?, nonexistent)\}$

$phys_context' = phys_context$

$log_context' = log_context \oplus \{(target?, bare_context)\}$

$\Xi PrioData$

$topReady! = running_task$

DeleteTaskS_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{running\}$

$state(topReady!) = ready$

$\forall t : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(t)$

$tasks' = tasks \setminus \{target?\}$

$running_task' = topReady!$

$state' = state \oplus \{(topReady!, running), (target?, nonexistent)\}$

$phys_context' = log_context(topReady!)$

$log_context' = log_context \oplus \{(target?, bare_context)\}$

$\Xi PrioData$

SuspendTaskN_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{ready, blocked\}$

$\exists TaskData$

$state' = state \oplus \{(target?, suspended)\}$

$\exists ContextData$

$\exists PrioData$

$topReady! = running_task$

SuspendTaskS_T

$\Delta Task$

$target? : TASK$

$topReady! : TASK$

$target? \in tasks \setminus \{idle\}$

$state(target?) \in \{running\}$

$state(topReady!) = ready$

$\forall t : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(t)$

$\exists st? : STATE \mid st? = suspended$

$\bullet Reschedule[tasks/tasks?, priority/pri?, topReady!/target?]$

SuspendTaskO_T

$\exists Task$

$target? : TASK$

$topReady! : TASK$

$state(target?) \in \{suspended\}$

$topReady! = running_task$

ResumeTaskN_T

$\Delta Task$

$target? : TASK$

$state(target?) = suspended$

$priority(target?) \leq priority(running_task)$

$\exists TaskData$

$state' = state \oplus \{(target?, ready)\}$

$\exists ContextData$

$\exists PrioData$

ResumeTaskS_T

 $\Delta Task$ $target? : TASK$ $state(target?) = suspended$ $priority(target?) > priority(running_task)$ $\exists st? : STATE \mid st? = ready \bullet Reschedule[tasks/tasks?, priority/pri?]$

ChangeTaskPriorityN_T

 $\Delta Task$ $target? : TASK$ $newpri? : \mathbb{N}$ $topReady! : TASK$ $state(target?) = ready \Rightarrow newpri? \leq priority(running_task)$ $state(target?) = running \Rightarrow (\forall t : state \sim (\{ready\} \mid$

- $\bullet newpri? \geq priority(t)$

 $state(target?) \neq nonexistent$ $target? = idle \Rightarrow newpri? = 0$ $\exists TaskData$ $\exists StateData$ $\exists ContextData$ $priority' = priority \oplus \{(target?, newpri?)\}$ $topReady! = running_task$

ChangeTaskPriorityS_T

 $\Delta Task$ $target? : TASK$ $newpri? : \mathbb{N}$ $topReady! : TASK$ $state(target?) = ready$ $newpri? > priority(running_task)$ $target? = idle \Rightarrow newpri? = 0$ $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$

- $\mid st? = ready$

- $\wedge pri? = priority \oplus \{(target?, newpri?)\}$

- $\bullet Reschedule[tasks/tasks?]$

 $topReady! = target?$

ChangeTaskPriorityD_T

 $\Delta Task$ $target? : TASK$ $newpri? : \mathbb{N}$

$topReady! : TASK$

$state(target?) = running$
 $target? = idle \Rightarrow newpri? = 0$
 $state(topReady!) = ready$
 $\forall t : state \sim (\{ready\}) \bullet priority(topReady!) \geq priority(t)$
 $newpri? < priority(topReady!)$
 $\exists st? : STATE; pri? : TASK \rightarrow \mathbb{N}$
 $| st? = ready$
 $\wedge pri? = priority \oplus \{(target?, newpri?)\}$
 $\bullet Reschedule[tasks/tasks?, topReady!/target?]$

$Multi_Task$

$subTask : cores \rightarrow Task$
 $exeCore : TASK \rightarrow cores$

$\forall c1, c2 : cores \mid c1 \neq c2 \bullet$
 $(subTask\ c1).tasks \cap (subTask\ c2).tasks = \emptyset$
 $dom\ exeCore \in \mathbb{F}\ TASK$
 $dom\ exeCore = \bigcup \{c : cores \bullet (subTask(c)).tasks\}$

$Init$

$Multi_Task'$

$dom\ exeCore' = idles$
 $\forall c : cores \bullet \exists Task'$
 $| Init_Task \bullet subTask'(c) = \theta Task'$
 $\wedge exeCore'((subTask'(c)).idle) = c$

$PromoteC$

$\Delta Multi_Task$
 $\Delta Task$
 $target? : TASK$
 $executeCore : cores$

$target? \notin dom\ exeCore$
 $subTask(executeCore) = \theta Task$
 $subTask' = subTask \oplus \{(executeCore, \theta Task')\}$
 $exeCore' = exeCore \oplus \{(target?, executeCore)\}$

$createTaskN_MT \hat{=} \exists \Delta Task \bullet CreateTaskN_T \wedge PromoteC$

$createTaskS_MT \hat{=} \exists \Delta Task \bullet CreateTaskS_T \wedge PromoteC$

$\underline{\text{findACore_MT}}$

Multi_Task

$\text{target?} : \text{TASK}$

$\text{newpri?} : \mathbb{N}$

$\text{executeCore?} : \text{CORE}$

$\text{executeCore} : \text{CORE}$

$\text{executeCore?} \notin \text{cores}$

$\text{executeCore} \in \text{cores}$

$\exists \text{tcs}, \text{cs} : \mathbb{F} \text{cores}$

| $\text{tcs} = \{ \text{pc} : \text{cores}; \text{subS} : \text{Task} \}$

| $\text{subS} = \text{subTask}(\text{pc})$

$\wedge \text{newpri?} > \text{subS}.\text{priority}(\text{subS}.\text{running_task}) \bullet \text{pc} \}$

• $(\text{tcs} = \emptyset \Rightarrow \text{cs} = \text{cores})$

$\wedge (\text{tcs} \neq \emptyset \Rightarrow \text{cs} = \text{tcs})$

$\wedge (\forall \text{oc} : \text{cs} \bullet \text{executeCore} \in \text{cs})$

$\wedge \#(\text{exeCore} \sim (\{ \text{executeCore} \})) \leq$

$\#(\text{exeCore} \sim (\{ \text{oc} \}))$

$\text{CreateTaskN_MT} \triangleq ([\text{executeCore?}, \text{executeCore} : \text{CORE}$
| $\text{executeCore?} \in \text{cores} \wedge \text{executeCore} = \text{executeCore?}]$
 $\vee \text{findACore_MT}) \wedge \text{createTaskN_MT}$

$\text{CreateTaskS_MT} \triangleq ([\text{executeCore?}, \text{executeCore} : \text{CORE}$
| $\text{executeCore?} \in \text{cores} \wedge \text{executeCore} = \text{executeCore?}]$
 $\vee \text{findACore_MT}) \wedge \text{createTaskS_MT}$

$\underline{\text{PromoteD}}$

$\Delta \text{Multi_Task}$

ΔTask

$\text{target?} : \text{TASK}$

$\text{target?} \in \text{dom } \text{exeCore}$

$\text{subTask}(\text{exeCore}(\text{target?})) = \theta \text{Task}$

$\text{subTask}' = \text{subTask} \oplus \{(\text{exeCore}(\text{target?}), \theta \text{Task}')\}$

$\text{exeCore}' = \{\text{target?}\} \triangleleft \text{exeCore}$

$\text{DeleteTaskN_MT} \triangleq \exists \Delta \text{Task} \bullet \text{DeleteTaskN_T} \wedge \text{PromoteD}$

$\text{DeleteTaskS_MT} \triangleq \exists \Delta \text{Task} \bullet \text{DeleteTaskS_T} \wedge \text{PromoteD}$

$\underline{\text{Promote}}$

$$\begin{array}{l} \Delta Multi_Task \\ \Delta Task \\ target? : TASK \end{array}$$

$$\begin{array}{l} target? \in \text{dom } exeCore \\ subTask(exeCore(target?)) = \theta Task \\ subTask' = subTask \oplus \{(exeCore(target?), \theta Task')\} \\ exeCore' = exeCore \end{array}$$

$$SuspendTaskN_MT \cong \exists \Delta Task \bullet SuspendTaskN_T \wedge Promote$$

$$SuspendTaskS_MT \cong \exists \Delta Task \bullet SuspendTaskS_T \wedge Promote$$

$$SuspendTaskO_MT \cong \exists \Delta Task \bullet SuspendTaskO_T \wedge Promote$$

$$ResumeTaskN_MT \cong \exists \Delta Task \bullet ResumeTaskN_T \wedge Promote$$

$$ResumeTaskS_MT \cong \exists \Delta Task \bullet ResumeTaskS_T \wedge Promote$$

$$\begin{array}{l} ChangeTaskPriorityN_MT \cong \\ \exists \Delta Task \bullet ChangeTaskPriorityN_T \wedge Promote \end{array}$$

$$\begin{array}{l} ChangeTaskPriorityS_MT \cong \\ \exists \Delta Task \bullet ChangeTaskPriorityS_T \wedge Promote \end{array}$$

$$\begin{array}{l} ChangeTaskPriorityD_MT \cong \\ \exists \Delta Task \bullet ChangeTaskPriorityD_T \wedge Promote \end{array}$$

$$MigrationN_MT$$

$$\begin{array}{l} \Delta Multi_Task \\ target? : TASK \\ topReady! : TASK \\ newCore? : cores \\ srcSys, tarSys : Task \end{array}$$

$$\begin{array}{l} target? \in \text{dom } exeCore \\ srcSys = subTask(exeCore(target?)) \\ tarSys = subTask(newCore?) \\ srcSys.state(target?) \in \{ready, blocked, suspended\} \\ srcSys.state(target?) = ready \Rightarrow \\ \quad srcSys.priority(target?) \leq tarSys.priority(tarSys.running_task) \end{array}$$

$$\begin{aligned}
& target? \notin idles \\
& newCore? \neq exeCore(target?) \\
& \{exeCore(target?), newCore?\} \triangleleft subTask' = \\
& \quad \{exeCore(target?), newCore?\} \triangleleft subTask \\
& \exists \Delta Task \bullet \\
& \quad subTask(exeCore(target?)) = \theta Task \\
& \quad \wedge DeleteTaskN_T \\
& \quad \wedge subTask'(exeCore(target?)) = \theta Task' \\
& \exists \Delta Task; tpri : \mathbb{N} | \\
& \quad tpri = (srcSys.priority(target?)) \bullet \\
& \quad \quad subTask(newCore?) = \theta Task \\
& \quad \quad \wedge CreateTaskN_T[tpri/newpri?] \\
& \quad \quad \wedge subTask'(newCore?) = \theta Task' \\
& exeCore' = exeCore \oplus \{(target?, newCore?)\}
\end{aligned}$$

MigrationS_MT

$$\begin{aligned}
& \Delta Multi_Task \\
& target? : TASK \\
& topReady! : TASK \\
& newCore? : cores \\
& srcSys, tarSys : Task
\end{aligned}$$

$$\begin{aligned}
& target? \in \text{dom } exeCore \\
& srcSys = subTask(exeCore(target?)) \\
& tarSys = subTask(newCore?) \\
& srcSys.state(target?) = ready \\
& srcSys.priority(target?) > tarSys.priority(tarSys.running_task) \\
& target? \notin idles \\
& newCore? \neq exeCore(target?) \\
& \{exeCore(target?), newCore?\} \triangleleft subTask' = \\
& \quad \{exeCore(target?), newCore?\} \triangleleft subTask \\
& \exists \Delta Task \bullet \\
& \quad subTask(exeCore(target?)) = \theta Task \\
& \quad \wedge DeleteTaskN_T \\
& \quad \wedge subTask'(exeCore(target?)) = \theta Task' \\
& \exists \Delta Task; tpri : \mathbb{N} | \\
& \quad tpri = (srcSys.priority(target?)) \bullet \\
& \quad \quad subTask(newCore?) = \theta Task \\
& \quad \quad \wedge CreateTaskS_T[tpri/newpri?] \\
& \quad \quad \wedge subTask'(newCore?) = \theta Task' \\
& \quad exeCore' = exeCore \oplus \{(target?, newCore?)\}
\end{aligned}$$

MigrationRuN_MT

$$\Delta Multi_Task$$

$target? : TASK$
 $topReady! : TASK$
 $newCore? : cores$
 $srcSys, tarSys : Task$

$target? \in \text{dom } exeCore$
 $srcSys = subTask(exeCore(target?))$
 $tarSys = subTask(newCore?)$
 $srcSys.state(target?) = running$
 $srcSys.priority(target?) \leq tarSys.priority(tarSys.running_task)$
 $target? \notin idles$
 $newCore? \neq exeCore(target?)$
 $\{exeCore(target?), newCore?\} \triangleleft subTask' =$
 $\{exeCore(target?), newCore?\} \triangleleft subTask$
 $\exists \Delta Task \bullet$
 $subTask(exeCore(target?)) = \theta Task$
 $\wedge DeleteTaskS_T$
 $\wedge subTask'(exeCore(target?)) = \theta Task'$
 $\exists \Delta Task; tpri : \mathbb{N} |$
 $tpri = (srcSys.priority(target?)) \bullet$
 $subTask(newCore?) = \theta Task$
 $\wedge CreateTaskN_T[tpri/newpri?]$
 $\wedge subTask'(newCore?) = \theta Task'$
 $exeCore' = exeCore \oplus \{(target?, newCore?)\}$

MigrationRuS_MT

$\Delta Multi_Task$
 $target? : TASK$
 $topReady! : TASK$
 $newCore? : cores$
 $srcSys, tarSys : Task$

$target? \in \text{dom } exeCore$
 $srcSys = subTask(exeCore(target?))$
 $tarSys = subTask(newCore?)$
 $srcSys.state(target?) = running$
 $srcSys.priority(target?) > tarSys.priority(tarSys.running_task)$
 $target? \notin idles$
 $newCore? \neq exeCore(target?)$
 $\{exeCore(target?), newCore?\} \triangleleft subTask' =$
 $\{exeCore(target?), newCore?\} \triangleleft subTask$
 $\exists \Delta Task \bullet$
 $subTask(exeCore(target?)) = \theta Task$
 $\wedge DeleteTaskS_T$
 $\wedge subTask'(exeCore(target?)) = \theta Task'$
 $\exists \Delta Task; tpri : \mathbb{N} |$

$$\begin{array}{l}
tpri = (srcSys.priority(target?)) \bullet \\
subTask(newCore?) = \theta Task \\
\wedge CreateTaskS_T[tpri/newpri?] \\
\wedge subTask'(newCore?) = \theta Task' \\
exeCore' = exeCore \oplus \{(target?, newCore?)\}
\end{array}$$

$$CN_T \hat{=} CreateTaskN_T$$

$$CS_T \hat{=} CreateTaskS_T$$

$$DN_T \hat{=} DeleteTaskN_T$$

$$DS_T \hat{=} DeleteTaskS_T$$

$$SN_T \hat{=} SuspendTaskN_T$$

$$SS_T \hat{=} SuspendTaskS_T$$

$$SO_T \hat{=} SuspendTaskO_T$$

$$RN_T \hat{=} ResumeTaskN_T$$

$$RS_T \hat{=} ResumeTaskS_T$$

$$ChN_T \hat{=} ChangeTaskPriorityN_T$$

$$ChS_T \hat{=} ChangeTaskPriorityS_T$$

$$ChD_T \hat{=} ChangeTaskPriorityD_T$$

$$\underline{getRunningTask}$$

$$\exists Multi_Task$$

$$core? : cores$$

$$RT! : TASK$$

$$RT! = (subTask(core?)).running_task$$

$$\underline{getPriority}$$

$$\exists Multi_Task$$

$$task? : TASK$$

$$PRIORITY! : \mathbb{N}$$

$$\exists core : cores \mid task? \in (subTask(core)).tasks$$

$$\bullet PRIORITY! = (subTask(core)).priority(task?)$$

Appendix M

VCC ANNOTATED SOURCE CODE

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <string.h>
4
5 #include "FreeRTOS.h"
6 #include "task.h"
7
8 _(dynamic_owns) typedef struct tskTaskControlBlock
9 {
10     xListItem          xGenericListItem;    /*< The list that the state list item
        of a task is reference from denotes the state of that task (Ready, Blocked,
        Suspended ). */
11     xListItem          xEventListItem;     /*< Used to reference a task from an event
        list. */
12     unsigned portBASE_TYPE uxPriority;     /*< The priority of the task. 0 is the
        lowest priority. */
13
14     _(invariant uxPriority < configMAX_PRIORITIES)
15     _(invariant \mine(&xGenericListItem))
16     _(invariant \mine(&xEventListItem))
17 } tskTCB;
18
19 PRIVILEGED_DATA tskTCB * volatile pxCurrentTCB = NULL;
20
21 PRIVILEGED_DATA static xList pxReadyTasksLists[ configMAX_PRIORITIES ]; /*< Prioritised
    ready tasks. */
22 PRIVILEGED_DATA static xList xDelayedTaskList1; /*< Delayed tasks.
    */
23 PRIVILEGED_DATA static xList xDelayedTaskList2; /*< Delayed tasks
    (two lists are used – one for delays that have overflowed the current tick count. */
24 PRIVILEGED_DATA static xList * volatile pxDelayedTaskList ; /*< Points to the
    delayed task list currently being used. */
25 PRIVILEGED_DATA static xList * volatile pxOverflowDelayedTaskList; /*< Points to the
    delayed task list currently being used to hold tasks that have overflowed the current
    tick count. */
26 PRIVILEGED_DATA static xList xPendingReadyList; /*< Tasks that
    have been readied while the scheduler was suspended. They will be moved to the ready
    queue when the scheduler is resumed. */
27
28 PRIVILEGED_DATA static xList xTasksWaitingTermination; /*< Tasks that have
    been deleted – but the their memory not yet freed. */
29
30 PRIVILEGED_DATA static xList xSuspendedTaskList;
31
32 _(def \bool excList(){
33
34     \bool tmp = (
35         \forall int i1, i2; ( 0 <= i1 && i1 < i2 && i2 < configMAX_PRIORITIES )
```

```

36     ==> !((&pxReadyTasksLists[i1]) == (&pxReadyTasksLists[i2]))
37 ) && (
38     \forall int i; ( 0 <= i && i < configMAX_PRIORITIES )
39     ==> (!((&pxReadyTasksLists[i]) == &xDelayedTaskList1)
40         && !((&pxReadyTasksLists[i]) == &xDelayedTaskList2)
41         && !((&pxReadyTasksLists[i]) == &xSuspendedTaskList)
42         && !((&pxReadyTasksLists[i]) == &xTasksWaitingTermination))
43 ) && (
44     &xDelayedTaskList1 != &xDelayedTaskList2
45 ) && (
46     &xDelayedTaskList1 != &xSuspendedTaskList
47 ) && (
48     &xDelayedTaskList1 != &xTasksWaitingTermination
49 ) && (
50     &xDelayedTaskList2 != &xSuspendedTaskList
51 ) && (
52     &xDelayedTaskList2 != &xTasksWaitingTermination
53 ) && (
54     &xSuspendedTaskList != &xTasksWaitingTermination
55 );
56
57 return tmp;
58 })
59
60 PRIVILEGED_DATA static xTaskHandle xIdleTaskHandle = NULL;          /*< Holds the handle
    of the idle task. The idle task is created automatically when the scheduler is
    started. */
61
62 PRIVILEGED_DATA static volatile signed portBASE_TYPE xSchedulerRunning = pdFALSE;
63
64 #define prvGetTCBFromHandle( pxHandle ) ( ( ( pxHandle ) == NULL ) ? ( tskTCB * )
    pxCurrentTCB : ( tskTCB * ) ( pxHandle ) )
65
66 _(ghost typedef tskTCB * TASK;);
67
68 _(ghost _(dynamic_owns) typedef struct{
69     \bool tasks[TASK];
70     STATE state[TASK];
71     //CONTEXT phys_context;
72     //CONTEXT log_context[TASK];
73
74     \natural priority[TASK];
75
76     //READY
77     _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE) ==>
78         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &pxReadyTasksLists[(((
79             tskTCB *) t)->uxPriority] &&
80             ((tskTCB *) t) != pxCurrentTCB) <==> state[t] == ready))
81     //BLOCKED
82     _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE) ==>
83         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &xDelayedTaskList1 ||
84             ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &xDelayedTaskList2)
85             <==> state[t] == blocked))
86     //SUSPENDED
87     _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE) ==>
88         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &xSuspendedTaskList
89             <==> state[t] == suspended))
90     //RUNNING
91     _(invariant \forall TASK t; (tasks[t] && xSchedulerRunning != pdFALSE) ==>
92         (t == (TASK) pxCurrentTCB <==> state[t] == running))
93     _(invariant ((xList *) pxCurrentTCB->xGenericListItem.pvContainer) == &
94         pxReadyTasksLists[pxCurrentTCB->uxPriority])
95     //NONEXISTENT
96     _(invariant \forall TASK t; (t->\closed && xSchedulerRunning != pdFALSE) ==>
97         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
98             xTasksWaitingTermination ||
99             t == NULL) <==> state[t] == nonexistent))
100     _(invariant \forall TASK t; xSchedulerRunning != pdFALSE ==> (state[t] != nonexistent
101         <==> tasks[t]))
102     _(invariant \forall TASK t; tasks[t] ==> \mine(t))
103     _(invariant \forall TASK t; tasks[t] ==> t->\closed)
104     _(invariant \forall TASK t; tasks[t] ==> priority[t] == t->uxPriority)
105     _(invariant \forall TASK t; tasks[t] ==> state[t] <= 4)
106     _(invariant \mine(\embedding(& xIdleTaskHandle)))

```



```

102     _(invariant \mine(\embedding(& pxCurrentTCB)))
103     _(invariant xSchedulerRunning != pdFALSE ==> (tasks[xIdleTaskHandle] && tasks[
        pxCurrentTCB]))
104     _(invariant xSchedulerRunning != pdFALSE && xIdleTaskHandle != NULL ==> priority[
        xIdleTaskHandle] == 0)
105     _(invariant \forall TASK t; xSchedulerRunning != pdFALSE && state[t] == ready ==>
        priority[pxCurrentTCB] >= priority[t])
106 } * FRTOS;)
107
108 //signed portBASE_TYPE xTaskGenericCreate(unsigned portBASE_TYPE uxPriority, xTaskHandle *
        pxCreatedTask _(ghost FRTOS FreeRTOS) _(ghost TASK *newTask))
109 signed portBASE_TYPE xTaskGenericCreate( pdTASK_CODE pxTaskCode, const signed char * const
        pcName, unsigned short usStackDepth, void *pvParameters, unsigned portBASE_TYPE
        uxPriority, xTaskHandle *pxCreatedTask, portSTACK_TYPE *puxStackBuffer, const
        xMemoryRegion * const xRegions _(ghost FRTOS FreeRTOS) _(ghost TASK *newTask) )
110     _(updates FreeRTOS)
111     _(requires \mutable(&xSchedulerRunning))
112     _(requires xSchedulerRunning == pdTRUE)
113     _(requires excList())
114
115     //_(writes newTask, pxCreatedTask,
116     //    &pxReadyTasksLists[uxPriority >= configMAX_PRIORITIES ? configMAX_PRIORITIES - (
        unsigned portBASE_TYPE ) 1U : uxPriority],
117     //    \embedding(&pxReadyTasksLists[uxPriority >= configMAX_PRIORITIES ?
        configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U : uxPriority]))
118     _(writes newTask, pxCreatedTask)
119
120     _(ensures \result == pdPASS ==> uxPriority > \old(pxCurrentTCB)->uxPriority ==> (
        FreeRTOS->state[(TASK)\old(pxCurrentTCB)] == ready && (FreeRTOS->state[(TASK)
        pxCurrentTCB]) == running)
121     _(ensures \result == pdPASS ==> uxPriority <= \old(pxCurrentTCB)->uxPriority ==> (
        FreeRTOS->state[*newTask]) == ready)
122     _(ensures \result == pdPASS ==> FreeRTOS->priority[*newTask] ==
123     (\natural)(uxPriority < configMAX_PRIORITIES ? uxPriority : configMAX_PRIORITIES -
        ( unsigned portBASE_TYPE ) 1U))
124     _(ensures \result == pdPASS ==> pxCurrentTCB->\closed)
125     _(ensures \result == pdPASS ==> \fresh(*newTask))
126     _(ensures \forall TASK t; (FreeRTOS->tasks[t] && \old(FreeRTOS->state[t]) != FreeRTOS
        ->state[t]) ==> transition[\old(FreeRTOS->state[t])][FreeRTOS->state[t]])
127 {
128     _(assert \wrapped(FreeRTOS))
129     _(assert \inv(FreeRTOS))
130     _(assert xSchedulerRunning == pdTRUE)
131     _(assert FreeRTOS->tasks[pxCurrentTCB])
132     _(assert pxCurrentTCB \in FreeRTOS->\owns)
133
134     //_(assert \forall TASK t; FreeRTOS->tasks[t] ==> (t \in FreeRTOS->\owns && t->\closed
        ))
135     //_(assert \forall TASK t; FreeRTOS->tasks[t] ==> t->\closed)
136     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> \inv(t))
137     //_(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->priority[t] == t->
        uxPriority)
138     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->priority[t] <
        configMAX_PRIORITIES)
139
140     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->state[t] <= 4)
141     _(assert pxCurrentTCB->\closed)
142
143     signed portBASE_TYPE xReturn;
144     tskTCB * pxNewTCB;
145
146     pxNewTCB = (tskTCB *) malloc(sizeof(tskTCB));
147
148     _(assert pxNewTCB != pxCurrentTCB)
149     //_(assert \wrapped(pxCurrentTCB))
150
151     if ( pxNewTCB != NULL )
152     {
153         _(assert pxNewTCB)
154         _(assert \fresh(pxNewTCB))
155
156         /* Setup the newly allocated TCB with the initial state of the task. */
157         //prvInitialiseTCBVariables( pxNewTCB, pcName, uxPriority, xRegions,
        usStackDepth );
158

```

```

159     pxNewTCB->uxPriority = (uxPriority < configMAX_PRIORITIES ? uxPriority :
160     configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U);
161     pxNewTCB->xGenericListItem.pvContainer = &pxReadyTasksLists[pxNewTCB->
162     uxPriority];
163     _(assert \writable(&(pxNewTCB->xGenericListItem)))
164     _(assert \writable(&(pxNewTCB->xEventListItem)))
165     _(wrap &(pxNewTCB->xGenericListItem))
166     _(wrap &(pxNewTCB->xEventListItem))
167     _(ghost pxNewTCB->\owns = (\objset) {&(pxNewTCB->xGenericListItem)})
168     _(ghost pxNewTCB->\owns += &(pxNewTCB->xEventListItem))
169
170     _(assert !(FreeRTOS \in pxNewTCB->\owns))
171     _(wrap pxNewTCB)
172
173     if ( ( void * ) pxCreatedTask != NULL )
174     {
175         *pxCreatedTask = ( xTaskHandle ) pxNewTCB;
176     }
177
178     //_(assert \false)
179     taskENTER_CRITICAL();
180     {
181         // ...
182         xReturn = pdPASS;
183     }
184     taskEXIT_CRITICAL();
185 }
186 else
187 {
188     _(assert \inv(FreeRTOS))
189     xReturn = errCOULD_NOT_ALLOCATE_REQUIRED_MEMORY;
190 }
191
192 if( xReturn == pdPASS )
193 {
194     _(unwrapping FreeRTOS){
195         _(assert \inv(FreeRTOS))
196
197         _(ghost {
198             FreeRTOS->priority[(TASK) pxNewTCB] = pxNewTCB->uxPriority;
199             FreeRTOS->state[(TASK) pxNewTCB] = ready;
200             FreeRTOS->tasks[(TASK) pxNewTCB] = \true;
201             FreeRTOS->\owns += pxNewTCB;
202             *newTask = pxNewTCB;
203         })
204
205         _(assert FreeRTOS->state[(TASK) pxNewTCB] == ready)
206         _(assert FreeRTOS->state[(TASK) pxCurrentTCB] == running)
207         _(assert FreeRTOS->priority[(TASK) pxNewTCB] == pxNewTCB->uxPriority)
208         _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxNewTCB)
209             && xSchedulerRunning != pdFALSE) ==>
210             (
211                 (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
212                 pxReadyTasksLists[(((tskTCB *) t)->uxPriority] &&
213                 ((tskTCB *) t) != pxCurrentTCB) <==> FreeRTOS->state[t] == ready
214             )
215
216         _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxNewTCB)
217             && xSchedulerRunning != pdFALSE) ==>
218             (
219                 (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
220                 xDelayedTaskList1 ||
221                 ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
222                 xDelayedTaskList2)
223                 <==> FreeRTOS->state[t] == blocked
224             )
225
226         _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxNewTCB)
227             && xSchedulerRunning != pdFALSE) ==>
228             (

```

```

225             ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
226                 xSuspendedTaskList
227                 <==> FreeRTOS->state[t] == suspended
228         )
229     )
230     _(assert \forall TASK t; (t->\closed && ((tskTCB *)t != pxNewTCB) &&
231         xSchedulerRunning != pdFALSE) ==>
232     (
233         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
234             xTasksWaitingTermination ||
235             t == NULL)
236         <==> FreeRTOS->state[t] == nonexistent
237     )
238
239     if( _(atomic_read \embedding(&pxCurrentTCB))
240         pxCurrentTCB->uxPriority < uxPriority )
241     {
242         _(ghost {
243             FreeRTOS->state[(TASK) pxCurrentTCB] = ready;
244             FreeRTOS->state[(TASK) pxNewTCB] = running;
245         })
246
247         _(assert FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready)
248
249         //portYIELD_WITHIN_API();
250         _(atomic \embedding(&pxCurrentTCB)){
251             pxCurrentTCB = pxNewTCB;
252             _(bump_volatile_version \embedding(&pxCurrentTCB))
253         }
254     }
255
256     _(assert \old(pxCurrentTCB)->uxPriority >= uxPriority ==> (FreeRTOS->state
257         [(TASK) pxNewTCB] == ready && FreeRTOS->state[(TASK) pxCurrentTCB] ==
258         running))
259     _(assert \old(pxCurrentTCB)->uxPriority < uxPriority ==>
260         (FreeRTOS->state[(TASK) pxNewTCB] == running &&
261             pxNewTCB == pxCurrentTCB &&
262             FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready
263         )
264     )
265     _(assert FreeRTOS->priority[(TASK) pxNewTCB] == pxNewTCB->uxPriority)
266     _(assert FreeRTOS->priority[(TASK) \old(pxCurrentTCB)] == \old(
267         pxCurrentTCB)->uxPriority)
268     _(assert FreeRTOS->tasks[(TASK) \old(pxCurrentTCB)])
269
270     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxNewTCB)
271         && xSchedulerRunning != pdFALSE) ==>
272     (
273         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
274             pxReadyTasksLists[[(tskTCB *) t)->uxPriority] &&
275             ((tskTCB *) t) != pxCurrentTCB)
276         <==> FreeRTOS->state[t] == ready
277     )
278
279     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxNewTCB)
280         && xSchedulerRunning != pdFALSE) ==>
281     (
282         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
283             xDelayedTaskList1 ||
284             ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
285             xDelayedTaskList2)
286         <==> FreeRTOS->state[t] == blocked
287     )
288
289     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxNewTCB)
290         && xSchedulerRunning != pdFALSE) ==>
291     (
292         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
293             xSuspendedTaskList
294         <==> FreeRTOS->state[t] == suspended

```

```

287     )
288 )
289
290     _(assume \forall TASK t; (t->\closed && ((tskTCB *)t != pxNewTCB) &&
291         xSchedulerRunning != pdFALSE) ==>
292     (
293         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
294             xTasksWaitingTermination ||
295             t == NULL)
296         <==> FreeRTOS->state[t] == nonexistent
297     )
298     )
299     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> t->\closed)
300
301     //_(assume \inv(FreeRTOS))
302     } //wrapping FreeRTOS
303
304 }
305
306     _(assume Trans())
307     return xReturn;
308 }
309
310 void vTaskDelete( xTaskHandle pxTaskToDelete _(ghost FRTOS FreeRTOS) _(ghost TASK topReady
311 ))
312     _(updates FreeRTOS)
313
314     _(requires FreeRTOS->tasks[topReady])
315     _(requires FreeRTOS->state[topReady] == ready)
316     _(requires \forall TASK rts;
317         (FreeRTOS->tasks[rts] && FreeRTOS->state[rts] == ready)
318         ==> FreeRTOS->priority[topReady] >= FreeRTOS->priority[rts])
319     _(requires topReady != (TASK) pxTaskToDelete)
320     _(requires topReady != (TASK) pxCurrentTCB)
321
322     _(requires FreeRTOS->tasks[(TASK) pxTaskToDelete])
323     _(requires (tskTCB *)pxTaskToDelete != (tskTCB *)xIdleTaskHandle)
324     //_(requires (TASK) pxTaskToDelete->\closed)
325
326     _(requires \mutable(&xSchedulerRunning))
327     _(requires xSchedulerRunning == pdTRUE)
328     _(requires excList())
329
330     _(ensures ((TASK) pxTaskToDelete)->\closed)
331     _(ensures pxTaskToDelete != NULL ==> ! FreeRTOS->tasks[(TASK) pxTaskToDelete])
332     _(ensures pxTaskToDelete == NULL ==> ! FreeRTOS->tasks[(TASK) \old(pxCurrentTCB)])
333     _(ensures pxTaskToDelete == NULL ==> (TASK) pxCurrentTCB == topReady)
334     _(ensures \forall TASK t; (FreeRTOS->tasks[t] && \old(FreeRTOS->state[t]) != FreeRTOS
335         ->state[t]) ==> transition[\old(FreeRTOS->state[t])][FreeRTOS->state[t]])
336
337 {
338     _(assert \wrapped(FreeRTOS))
339     _(assert \inv(FreeRTOS))
340     _(assert xSchedulerRunning == pdTRUE)
341     //_(assert FreeRTOS->tasks[(TASK) pxTaskToDelete])
342     //_(assert pxCurrentTCB \in FreeRTOS->lowms)
343     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> \inv(t))
344     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->priority[t] <
345         configMAX_PRIORITIES)
346     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->state[t] <= 4)
347     _(assert pxCurrentTCB->\closed)
348     _(assert ((TASK)pxTaskToDelete)->\closed)
349
350     tskTCB *pxTCB;
351
352     _(unwrapping FreeRTOS){
353         taskENTER_CRITICAL();
354         {
355             _(assert \inv(FreeRTOS))
356             /**/ Ensure a yield is performed if the current task is being
357             //deleted. */
358             //
359
360             _(atomic \embedding(&pxCurrentTCB)){
361                 if ( pxTaskToDelete == pxCurrentTCB )

```

```

357     {
358         pxTaskToDelete = NULL;
359     }
360 }
361 //
362 /** If null is passed in here then we are deleting ourselves. */
363 /**pxTCB = prvGetTCBFromHandle( pxTaskToDelete );
364 _(atomic \embedding(&pxCurrentTCB)){
365     pxTCB = pxTaskToDelete != NULL ? (tskTCB *) pxTaskToDelete
366         : pxCurrentTCB;
367 }
368 _(assert (TASK) pxTaskToDelete == (TASK) pxCurrentTCB ==>
369 pxTCB == pxCurrentTCB)
370 _(assert \wrapped(pxTCB))
371
372 _(unwrapping pxTCB){
373     /* Remove task from the ready list and place in the termination list.
374     This will stop the task from be scheduled. The idle task will check
375     the termination list and free up any memory allocated by the
376     scheduler for the TCB and stack. */
377
378     _(assert \wrapped(&(pxTCB->xGenericListItem)))
379     if( uxListRemove( ( xListItem * ) &( pxTCB->xGenericListItem ) ) == 0 )
380     {
381         taskRESET_READY_PRIORITY( pxTCB->uxPriority );
382     }
383
384     _(assert \wrapped(&(pxTCB->xGenericListItem)))
385     _(unwrapping &(pxTCB->xGenericListItem)) {
386         pxTCB->xGenericListItem.pvContainer = NULL;
387     }
388
389     _(assert pxTCB->xGenericListItem.pvContainer == NULL)
390
391     /* Is the task waiting on an event also? */
392     if( pxTCB->xEventListItem.pvContainer != NULL )
393     {
394         uxListRemove( &( pxTCB->xEventListItem ) );
395     }
396
397     _(assert pxTCB->xEventListItem.pvContainer == NULL)
398
399     vListInsertEnd( ( xList * ) &xTasksWaitingTermination, &( pxTCB->
400     xGenericListItem ) );
401     _(assert (xList *)(&pxTCB->xGenericListItem.pvContainer) == (xList *) &
402     xTasksWaitingTermination)
403 }
404
405 _(ghost {
406     FreeRTOS->state[(TASK) pxTCB] = nonexistent;
407     FreeRTOS->tasks[(TASK) pxTCB] = \false;
408 })
409 }
410 taskEXIT_CRITICAL();
411
412 _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
413 xSchedulerRunning != pdFALSE) ==>
414     (
415         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
416         pxReadyTasksLists[((tskTCB *) t)->uxPriority] &&
417         ((tskTCB *) t) != pxCurrentTCB)
418         <==> FreeRTOS->state[t] == ready
419     )
420 )
421
422 _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
423 xSchedulerRunning != pdFALSE) ==>
424     (
425         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
426         xDelayedTaskList1 ||
427         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
428         xDelayedTaskList2)
429         <==> FreeRTOS->state[t] == blocked
430     )

```

```

425     )
426
427     _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
428         xSchedulerRunning != pdFALSE) ==>
429         (
430             ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
431             xSuspendedTaskList
432             <==> FreeRTOS->state[t] == suspended
433         )
434     )
435
436     //_(assert \forall TASK t; (t->\closed && ((tskTCB *)t != pxTCB) &&
437         xSchedulerRunning != pdFALSE) ==>
438     // (
439     //     (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
440     //     xTasksWaitingTermination ||
441     //     t == NULL)
442     //     <==> FreeRTOS->state[t] == nonexistent
443     // )
444
445     /* Force a reschedule if we have just deleted the current task. */
446
447     if( xSchedulerRunning != pdFALSE )
448     {
449         if( ( void * ) pxTaskToDelete == NULL )
450         {
451             //portYIELD_WITHIN_API();
452             _(atomic \embedding(&pxCurrentTCB)){
453                 _(ghost pxCurrentTCB = topReady)
454                 _(bump_volatile_version \embedding(&pxCurrentTCB))
455             }
456             _(ghost FreeRTOS->state[topReady] = running)
457
458             _(assert FreeRTOS->state[(TASK) pxTCB] == nonexistent)
459             _(assert (xList *) (pxTCB->xGenericListItem.pvContainer) == &
460                 xTasksWaitingTermination)
461             _(assert !FreeRTOS->tasks[(TASK) pxTCB])
462
463             _(assert FreeRTOS->state[(TASK) pxCurrentTCB] == running)
464             _(assert \old(pxCurrentTCB) == pxTCB)
465             _(assert FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == nonexistent)
466         }
467     }
468
469     _(assert pxTaskToDelete != NULL ==> FreeRTOS->state[topReady] == ready)
470
471     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
472         xSchedulerRunning != pdFALSE) ==>
473     (
474         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
475         pxReadyTasksLists[(((tskTCB *) t)->uxPriority] &&
476         ((tskTCB *) t) != pxCurrentTCB)
477         <==> FreeRTOS->state[t] == ready
478     )
479     )
480
481     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
482         xSchedulerRunning != pdFALSE) ==>
483     (
484         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
485         xDelayedTaskList1 ||
486         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
487         xDelayedTaskList2)
488         <==> FreeRTOS->state[t] == blocked
489     )
490     )
491
492     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
493         xSchedulerRunning != pdFALSE) ==>
494     (
495         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
496         xSuspendedTaskList
497         <==> FreeRTOS->state[t] == suspended
498     )

```

```

488     )
489
490     _(assume \forall TASK t; (t->\closed && ((tskTCB *)t != pxTCB) &&
        xSchedulerRunning != pdFALSE) ==>
491     (
492         (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
            xTasksWaitingTermination ||
493             t == NULL)
494             <==> FreeRTOS->state[t] == nonexistent
495     )
496 )
497
498     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> t->\closed)
499
500 }
501
502 _(assume Trans())
503 //_(assert \false)
504 }
505
506 #if ( INCLUDE_uxTaskPriorityGet == 1 )
507
508     unsigned portBASE_TYPE uxTaskPriorityGet( xTaskHandle pxTask _(ghost FRTOS FreeRTOS))
509     _(maintains \wrapped(FreeRTOS))
510     _(requires FreeRTOS->tasks[(TASK) pxTask])
511     _(requires xSchedulerRunning == pdTRUE)
512     _(requires excList())
513
514     _(writes FreeRTOS)
515
516     _(ensures \result == FreeRTOS->priority[(TASK) pxTask])
517 {
518     tskTCB *pxTCB;
519     unsigned portBASE_TYPE uxReturn;
520
521     taskENTER_CRITICAL();
522     {
523         /* If null is passed in here then we are changing the
524            priority of the calling function. */
525         //pxTCB = prvGetTCBFromHandle( pxTask );
526
527         _(unwrapping FreeRTOS){
528             _(assert \inv(FreeRTOS))
529             pxTCB = pxTask != NULL ? (tskTCB *) pxTask
530                                     : pxCurrentTCB;
531
532             _(assert \wrapped(pxTCB))
533             uxReturn = pxTCB->uxPriority;
534         }
535     }
536     taskEXIT_CRITICAL();
537
538     return uxReturn;
539 }
540
541 #endif
542
543 #if ( INCLUDE_vTaskPrioritySet == 1 )
544
545     void vTaskPrioritySet( xTaskHandle pxTask, unsigned portBASE_TYPE uxNewPriority _(
546         ghost FRTOS FreeRTOS) _(ghost TASK topReady) )
547     _(updates FreeRTOS)
548     _(requires \mutable(&xSchedulerRunning))
549     _(requires xSchedulerRunning == pdTRUE)
550     _(requires excList())
551
552     _(requires FreeRTOS->tasks[topReady])
553     //_(requires FreeRTOS->state[topReady] == ready)
554     //_(requires \forall TASK rts;
555         //    (FreeRTOS->tasks[rts] && FreeRTOS->state[rts] == ready)
556         //    ==> FreeRTOS->priority[topReady] >= FreeRTOS->priority[rts])
557     //_(requires topReady != (TASK) pxTask)
558     //_(requires topReady != (TASK) pxCurrentTCB)
559     _(requires FreeRTOS->tasks[(TASK) pxTask])

```

```

560     _(requires (tskTCB *) pxTask == (tskTCB *) xIdleTaskHandle ==> uxNewPriority == 0)
561
562     _(ensures ((TASK) pxTask)->\closed)
563     _(ensures FreeRTOS->tasks[(TASK) pxTask])
564     _(ensures FreeRTOS->priority[(TASK) pxTask] ==
565       (\natural)(uxNewPriority < configMAX_PRIORITIES ? uxNewPriority :
566         configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U))
567     _(ensures (FreeRTOS->state[(TASK) pxTask] == ready && FreeRTOS->priority[(TASK)
568       pxTask] > \old(pxCurrentTCB)->uxPriority) ==>
569       (FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready && FreeRTOS->state[(TASK)
570         pxTask] == running)
571
572     _(ensures ((pxTask == NULL || (tskTCB *) pxTask == \old(pxCurrentTCB)) && !(
573       \forall TASK t; FreeRTOS->state[t] == ready ==> FreeRTOS->priority[(TASK)
574         pxTask] >= FreeRTOS->priority[t]
575       )) ==> (FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == ready &&
576         FreeRTOS->state[(TASK) pxCurrentTCB] == running &&
577         \old(FreeRTOS->state[(TASK) pxCurrentTCB]) == ready)
578     )
579     _(ensures \forall TASK t; (FreeRTOS->tasks[t] && \old(FreeRTOS->state[t]) !=
580       FreeRTOS->state[t]) ==> transition[\old(FreeRTOS->state[t])][FreeRTOS->state[t]
581       ]))
582
583 {
584     _(assert \inv(FreeRTOS))
585     _(assert xSchedulerRunning == pdTRUE)
586     _(assert FreeRTOS->tasks[pxCurrentTCB])
587     _(assert pxTask != NULL ==> FreeRTOS->tasks[(TASK) pxTask])
588     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> \inv(t))
589     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->priority[t] <
590       configMAX_PRIORITIES)
591     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->state[t] <= 4)
592     _(assert pxCurrentTCB->\closed)
593     _(assert pxCurrentTCB \in FreeRTOS->\owns)
594     _(assert pxTask != NULL ==> ((TASK) pxTask)->\closed)
595     _(assert pxTask != NULL ==> ((TASK) pxTask) \in FreeRTOS->\owns)
596
597     tskTCB *pxTCB;
598     unsigned portBASE_TYPE uxCurrentPriority, uxPriorityUsedOnEntry;
599     portBASE_TYPE xYieldRequired = pdFALSE;
600
601     /* Ensure the new priority is valid. */
602     if ( uxNewPriority >= configMAX_PRIORITIES )
603     {
604         uxNewPriority = configMAX_PRIORITIES - ( unsigned portBASE_TYPE ) 1U;
605     }
606     _(assert uxNewPriority < configMAX_PRIORITIES)
607
608     taskENTER_CRITICAL();
609     {
610         _(unwrapping FreeRTOS){
611             _(assert \inv(FreeRTOS))
612             _(assert \forall TASK t; (t->\closed && xSchedulerRunning != pdFALSE) ==>
613               (((xList*)((tskTCB *) t)->xGenericListItem.pvContainer) == &
614                 xTasksWaitingTermination ||
615                 t == NULL) <==> FreeRTOS->state[t] == nonexistent))
616             _(atomic \embedding(&pxCurrentTCB)){
617                 if ( pxTask == pxCurrentTCB )
618                 {
619                     pxTask = NULL;
620                 }
621             }
622
623             _(atomic \embedding(&pxCurrentTCB)){
624                 pxTCB = pxTask != NULL ? (tskTCB *) pxTask
625                   : pxCurrentTCB;
626             }
627             _(assume \forall TASK t; (t->\closed && xSchedulerRunning != pdFALSE) ==>
628               (((xList*)((tskTCB *) t)->xGenericListItem.pvContainer) == &
629                 xTasksWaitingTermination ||
630                 t == NULL) <==> FreeRTOS->state[t] == nonexistent))
631
632             _(assert (TASK) pxTask == (TASK) pxCurrentTCB ==>
633               pxTCB == pxCurrentTCB)
634         }
635     }

```



```

625     _ (assert \inv(FreeRTOS))
626     _ (assert \wrapped(pxCurrentTCB))
627     _ (assert pxTask != NULL ==> FreeRTOS->tasks[(TASK) pxTask])
628     _ (assert pxTask != NULL ==> \wrapped((TASK) pxTask))
629     _ (assert \wrapped(pxTCB))
630
631     ///if ( configUSE_MUTEXES == 1 )
632     ///{
633         /// // uxCurrentPriority = pxTCB->uxBasePriority;
634         /// }
635     /// #else
636     /// {
637         uxCurrentPriority = pxTCB->uxPriority;
638     /// }
639     /// #endif
640
641     if ( uxCurrentPriority != uxNewPriority )
642     {
643         /// /* The priority change may have readied a task of higher
644         /// priority than the calling task. */
645
646         if ( uxNewPriority > uxCurrentPriority )
647         {
648             if ( pxTask != NULL )
649             {
650                 /// /* The priority of another task is being raised. If we
651                 /// were raising the priority of the currently running task
652                 /// there would be no need to switch as it must have already
653                 /// been the highest priority task. */
654                 xYieldRequired = pdTRUE;
655             }
656         }
657         else if ( pxTask == NULL )
658         {
659             /// /* Setting our own priority down means there may now be another
660             /// task of higher priority that is ready to execute. */
661             xYieldRequired = pdTRUE;
662         }
663     }
664
665     uxPriorityUsedOnEntry = pxTCB->uxPriority;
666
667     /// #if ( configUSE_MUTEXES == 1 )
668     /// /// {
669
670     /// // /// /* Only change the priority being used if the task is not
671     /// // /// currently using an inherited priority. */
672     /// // /// if( pxTCB->uxBasePriority == pxTCB->uxPriority )
673     /// // /// {
674     /// // ///     pxTCB->uxPriority = uxNewPriority;
675     /// // /// }
676
677     /// // /// /* The base priority gets set whatever. */
678     /// // /// pxTCB->uxBasePriority = uxNewPriority;
679     /// /// }
680     /// #else
681     /// /// {
682     /// _ (assert \inv(FreeRTOS))
683     /// _ (unwrapping pxTCB){
684
685         pxTCB->uxPriority = uxNewPriority;
686
687         _ (ghost FreeRTOS->priority [(TASK) pxTCB] = uxNewPriority)
688         _ (assert FreeRTOS->priority [(TASK) pxTCB] == pxTCB->uxPriority)
689     /// }
690     /// #endif
691     /// /// /* If the task is in the blocked or suspended list we need do
692     /// /// nothing more than change it's priority variable. However, if
693     /// /// the task is in a ready list it needs to be removed and placed
694     /// /// in the queue appropriate to its new priority. */
695     /// // /// if( listIS_CONTAINED_WITHIN( &(amp; pxReadyTasksLists[ uxCurrentPriority
696     /// // /// ], &( pxTCB->xGenericListItem ) ) )
697     /// if ( (xList *) pxTCB->xGenericListItem.pvContainer == &(

```

```

698     /*_ (assert (xList *) pxTCB->xGenericListItem.pvContainer == &
699     pxReadyTasksLists[ uxCurrentPriority ] ) )
700     /*_ (assert (xList *) pxTCB->xGenericListItem.pvContainer == &
701     pxReadyTasksLists[ uxCurrentPriority ] ) <==>
702     /*_ (FreeRTOS->state [(TASK) pxTCB] == running || FreeRTOS->state
703     [(TASK) pxTCB] == ready)
704     /*_ (assert FreeRTOS->state [(TASK) pxTCB] == running || FreeRTOS->
705     state [(TASK) pxTCB] == ready)
706     /* The task is currently in its ready list - remove before adding
707     it to it's new ready list. As we are in a critical section we
708     can do this even if the scheduler is suspended. */
709     /*_ (assert \wrapped(&(pxTCB->xGenericListItem)))
710     if ( uxListRemove( ( xListItem * ) &( pxTCB->xGenericListItem ) )
711     == 0 )
712     {
713         /*taskRESET_READY_PRIORITY( uxPriorityUsedOnEntry );
714     }
715     /*prvAddTaskToReadyQueue( pxTCB );
716     /*_ (assert (&(pxTCB->xGenericListItem))->\closed)
717     /*_ (unwrapping &(pxTCB->xGenericListItem)){
718     pxTCB->xGenericListItem.pvContainer = (xList *) &(
719     pxReadyTasksLists[pxTCB->uxPriority] );
720     }
721     /*_ (assert FreeRTOS->tasks [(TASK) pxTCB])
722     /*_ (assert (xList *) pxTCB->xGenericListItem.pvContainer == &(
723     pxReadyTasksLists[ pxTCB->uxPriority ] ) )
724     /*_ (assert FreeRTOS->state [(TASK) pxTCB] == running || FreeRTOS->
725     state [(TASK) pxTCB] == ready)
726     }
727     /*_ (assert (FreeRTOS->state [(TASK) pxTCB] == running || FreeRTOS->state
728     [(TASK) pxTCB] == ready) ==>
729     (xList *) pxTCB->xGenericListItem.pvContainer == &(
730     pxReadyTasksLists[ pxTCB->uxPriority ] ) )
731     if( xYieldRequired == pdTRUE )
732     {
733         /*_ (assert FreeRTOS->tasks[topReady])
734         /*_ (assume FreeRTOS->state[topReady] == ready ||
735         topReady == (TASK) pxCurrentTCB)
736         /*_ (assume \forall TASK t; (FreeRTOS->state[t] == ready || FreeRTOS
737         ->state[t] == running) ==>
738         FreeRTOS->priority[topReady] >= FreeRTOS->priority[t])
739         /*_ (assume FreeRTOS->priority[topReady] == FreeRTOS->priority [(TASK)
740         pxCurrentTCB]
741         ==> topReady == (TASK) pxCurrentTCB)
742         /*_ (ghost {
743         FreeRTOS->state [(TASK) pxCurrentTCB] = ready;
744         FreeRTOS->state[topReady] = running;
745     })
746         /*portYIELD_WITHIN_API();
747         /*_ (atomic \embedding(&pxCurrentTCB)){
748         /*_ (ghost pxCurrentTCB = topReady)
749         /*_ (bump_volatile_version \embedding(&pxCurrentTCB))
750     }
751     /*_ (assert topReady != \old((TASK) pxCurrentTCB) ==> FreeRTOS->state
752     [(TASK) \old(pxCurrentTCB)] == ready)
753     }
754     }
755     /*wrapping pxTCB
756     /*_ (assert \false)
757     /*_ (assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
758     xSchedulerRunning != pdFALSE) ==>
759     (
760     (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
761     pxReadyTasksLists[(((tskTCB *) t)->uxPriority] &&
762     ((tskTCB *) t) != pxCurrentTCB)
763     <==> FreeRTOS->state[t] == ready

```

```

758     )
759   )
760
761   _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
762     xSchedulerRunning != pdFALSE) ==>
763     (
764       (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
765         xDelayedTaskList1 ||
766         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
767         xDelayedTaskList2)
768       <==> FreeRTOS->state[t] == blocked
769     )
770
771   _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
772     xSchedulerRunning != pdFALSE) ==>
773     (
774       ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
775       xSuspendedTaskList
776       <==> FreeRTOS->state[t] == suspended
777     )
778
779   _(assume \forall TASK t; (t->\closed && ((tskTCB *)t != pxTCB) &&
780     xSchedulerRunning != pdFALSE) ==>
781     (
782       (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
783         xTasksWaitingTermination ||
784         t == NULL)
785       <==> FreeRTOS->state[t] == nonexistent
786     )
787
788   } // wrapping FreeRTOS
789
790   _(assert \inv(FreeRTOS))
791 }
792 taskEXIT_CRITICAL();
793 #endif
794
795 #if ( INCLUDE_vTaskSuspend == 1 )
796
797 void vTaskSuspend( xTaskHandle pxTaskToSuspend _(ghost FRTOS FreeRTOS) _(ghost TASK
798   topReady))
799   _(updates FreeRTOS)
800   _(requires \mutable(&xSchedulerRunning))
801   _(requires xSchedulerRunning == pdTRUE)
802   _(requires excList())
803   _(requires FreeRTOS->tasks[(TASK) pxTaskToSuspend])
804   _(requires (tskTCB *)pxTaskToSuspend != (tskTCB *)xIdleTaskHandle)
805
806   _(requires FreeRTOS->tasks[topReady])
807   _(requires FreeRTOS->state[topReady] == ready)
808   _(requires \forall TASK rts;
809     (FreeRTOS->tasks[rts] && FreeRTOS->state[rts] == ready)
810     ==> FreeRTOS->priority[topReady] >= FreeRTOS->priority[rts])
811   _(requires topReady != (TASK) pxTaskToSuspend)
812   _(requires topReady != (TASK) pxCurrentTCB)
813
814   _(ensures FreeRTOS->state[(TASK) pxTaskToSuspend] == suspended)
815   _(ensures pxTaskToSuspend == NULL ==> (TASK) pxTaskToSuspend == \old(pxCurrentTCB))
816   _(ensures pxTaskToSuspend == NULL ==> (TASK) pxCurrentTCB == topReady)
817   _(ensures \forall TASK t; (FreeRTOS->tasks[t] && \old(FreeRTOS->state[t]) != FreeRTOS
818     ->state[t]) ==> transition[\old(FreeRTOS->state[t])][FreeRTOS->state[t]])
819   {
820     _(assert \wrapped(FreeRTOS))
821     _(assert \inv(FreeRTOS))
822     _(assert xSchedulerRunning == pdTRUE)
823     //_(assert FreeRTOS->tasks[(TASK) pxTaskToDelete])
824     //_(assert pxCurrentTCB \in FreeRTOS->owns)

```

```

824  _ (assert \forall TASK t; FreeRTOS->tasks[t] ==> \inv(t))
825  _ (assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->priority[t] <
      configMAX_PRIORITIES)
826  _ (assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->state[t] <= 4)
827  _ (assert pxCurrentTCB->\closed)
828  _ (assert ((TASK)pxTaskToSuspend)->\closed)
829
830  tskTCB *pxTCB;
831
832  _ (unwrapping FreeRTOS){
833      taskENTER_CRITICAL();
834      {
835          _ (assert \inv(FreeRTOS))
836          /* Ensure a yield is performed if the current task is being
837             suspended. */
838          if ( pxTaskToSuspend == _ (atomic_read \embedding(&pxCurrentTCB))
839              pxCurrentTCB )
840          {
841              pxTaskToSuspend = NULL;
842          }
843
844          /* If null is passed in here then we are suspending ourselves. */
845          pxTCB = _ (atomic_read \embedding(&pxCurrentTCB))
846                  prvGetTCBFromHandle( pxTaskToSuspend );
847          _ (assert (TASK) pxTaskToSuspend == (TASK) pxCurrentTCB ==>
848              pxTCB == pxCurrentTCB)
849          _ (assert \wrapped(pxTCB))
850
851          _ (unwrapping pxTCB){
852              /* Remove task from the ready/delayed list and place in the suspended list
853                 . */
854              _ (assert \wrapped(&(pxTCB->xGenericListItem)))
855              if ( uxListRemove( ( xListItem * ) &( pxTCB->xGenericListItem ) ) == 0 )
856              {
857                  //taskRESET_READY_PRIORITY( pxTCB->uxPriority );
858              }
859
860              _ (assert pxTCB->xGenericListItem.pvContainer == NULL)
861
862              /* Is the task waiting on an event also? */
863              if ( pxTCB->xEventListItem.pvContainer != NULL )
864              {
865                  uxListRemove( &( pxTCB->xEventListItem ) );
866              }
867
868              _ (assert pxTCB->xEventListItem.pvContainer == NULL)
869
870              vListInsertEnd( ( xList * ) &xSuspendedTaskList, &( pxTCB->
871                  xGenericListItem ) );
872              _ (assert (xList *) (pxTCB->xGenericListItem.pvContainer) == (xList *) &
873                  xSuspendedTaskList)
874          }
875          }
876          _ (ghost {
877              FreeRTOS->state[(TASK) pxTCB] = suspended;
878          })
879      }
880      taskEXIT_CRITICAL();
881
882      _ (assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
883          xSchedulerRunning != pdFALSE) ==>
884      (
885          (((xList *) ((tskTCB *) t)->xGenericListItem.pvContainer) == &
886              pxReadyTasksLists[(((tskTCB *) t)->uxPriority] &&
887              ((tskTCB *) t) != pxCurrentTCB)
888              <==> FreeRTOS->state[t] == ready
889          )
890      )
891
892      _ (assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
893          xSchedulerRunning != pdFALSE) ==>
894      (
895          (((xList *) ((tskTCB *) t)->xGenericListItem.pvContainer) == &
896              xDelayedTaskList1 ||

```

```

891             ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
892             xDelayedTaskList2)
893         )
894     )
895
896     _(assert \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
897       xSchedulerRunning != pdFALSE) ==>
898     (
899         ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
900         xSuspendedTaskList
901         <==> FreeRTOS->state[t] == suspended
902     )
903 )
904
905 if ( ( void * ) pxTaskToSuspend == NULL )
906 {
907     if ( xSchedulerRunning != pdFALSE )
908     {
909         /* We have just suspended the current task. */
910         //portYIELD_WITHIN_API();
911         _(atomic \embedding(&pxCurrentTCB)){
912             _(ghost pxCurrentTCB = topReady)
913             _(bump_volatile_version \embedding(&pxCurrentTCB))
914         }
915         _(ghost FreeRTOS->state[topReady] = running)
916
917         _(assert FreeRTOS->state[(TASK) pxTCB] ==suspended)
918         _(assert (xList *) (pxTCB->xGenericListItem.pvContainer) == &
919           xSuspendedTaskList)
920
921         _(assert FreeRTOS->state[(TASK) pxCurrentTCB] == running)
922         _(assert \old(pxCurrentTCB) == pxTCB)
923         _(assert FreeRTOS->state[(TASK) \old(pxCurrentTCB)] == suspended)
924     }
925     //else
926     // {
927     //     /* The scheduler is not running, but the task that was pointed
928     //     // to by pxCurrentTCB has just been suspended and pxCurrentTCB
929     //     // must be adjusted to point to a different task. */
930     //     if ( listCURRENT_LIST_LENGTH( &xSuspendedTaskList ) ==
931     //         uxCurrentNumberOfTasks )
932     //     {
933     //         /* No other tasks are ready, so set pxCurrentTCB back to
934     //         // NULL so when the next task is created pxCurrentTCB will
935     //         // be set to point to it no matter what its relative priority
936     //         // is. */
937     //         pxCurrentTCB = NULL;
938     //     }
939     //     else
940     //     {
941     //         vTaskSwitchContext();
942     //     }
943     // }
944 }
945
946 _(assert pxTaskToSuspend != NULL ==> FreeRTOS->state[topReady] == ready)
947
948 _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
949   xSchedulerRunning != pdFALSE) ==>
950 (
951     (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
952     pxReadyTasksLists[(((tskTCB *) t)->uxPriority] &&
953     ((tskTCB *) t) != pxCurrentTCB)
954     <==> FreeRTOS->state[t] == ready
955 )
956 )
957
958 _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
959   xSchedulerRunning != pdFALSE) ==>
960 (
961     (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
962     xDelayedTaskList1 ||
963     ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
964     xDelayedTaskList2)

```

```

956         <==> FreeRTOS->state[t] == blocked
957     )
958 )
959
960     _(assume \forall TASK t; (FreeRTOS->tasks[t] && ((tskTCB *)t != pxTCB) &&
961         xSchedulerRunning != pdFALSE) ==>
962         (
963             ((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
964             xSuspendedTaskList
965             <==> FreeRTOS->state[t] == suspended
966         )
967     )
968     _(assume \forall TASK t; (t->\closed && ((tskTCB *)t != pxTCB) &&
969         xSchedulerRunning != pdFALSE) ==>
970         (
971             (((xList *)((tskTCB *) t)->xGenericListItem.pvContainer) == &
972             xTasksWaitingTermination ||
973             t == NULL)
974             <==> FreeRTOS->state[t] == nonexistent
975         )
976     )
977     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> t->\closed)
978 }
979     _(assume Trans())
980 }
981 #endif
982
983 #if ( INCLUDE_vTaskSuspend == 1 )
984
985 void vTaskResume( xTaskHandle pxTaskToResume _(ghost FRTOS FreeRTOS) )
986     _(updates FreeRTOS)
987     _(requires xSchedulerRunning == pdTRUE)
988     _(requires excList())
989
990     _(requires FreeRTOS->tasks[(TASK) pxTaskToResume])
991     _(requires FreeRTOS->state[(TASK) pxTaskToResume] == suspended)
992
993     _(ensures ((TASK) pxTaskToResume)->\closed)
994     _(ensures pxCurrentTCB->\closed)
995     _(ensures FreeRTOS->tasks[(TASK) pxTaskToResume])
996     _(ensures FreeRTOS->tasks[(TASK) pxCurrentTCB])
997     _(ensures ((tskTCB *)pxTaskToResume)->uxPriority > \old(pxCurrentTCB)->uxPriority
998         ==> (FreeRTOS->state[(TASK)\old(pxCurrentTCB)]) == ready && (TASK)
999         pxCurrentTCB == (TASK) pxTaskToResume && (FreeRTOS->state[(TASK)pxCurrentTCB]
1000         == running)
1001     )
1002     _(ensures ((tskTCB *)pxTaskToResume)->uxPriority <= \old(pxCurrentTCB)->uxPriority
1003         ==> (FreeRTOS->state[(TASK) pxTaskToResume] == ready && FreeRTOS->state[(
1004         TASK)\old(pxCurrentTCB)] == running && \old(pxCurrentTCB) == pxCurrentTCB )
1005     )
1006     _(ensures \forall TASK t; (FreeRTOS->tasks[t] && \old(FreeRTOS->state[t]) !=
1007         FreeRTOS->state[t]) ==> transition [\old(FreeRTOS->state[t])][FreeRTOS->state[t]
1008         ]])
1009 {
1010     _(assert \wrapped(FreeRTOS))
1011     _(assert \inv(FreeRTOS))
1012     _(assert xSchedulerRunning == pdTRUE)
1013     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> FreeRTOS->state[t] <= 4)
1014     _(assert ((TASK) pxTaskToResume)->\closed)
1015     _(assert ((TASK) pxTaskToResume) \in FreeRTOS->\owns)
1016     _(assert ((TASK) pxCurrentTCB)->\closed)
1017     _(assert ((TASK) pxCurrentTCB) \in FreeRTOS->\owns)
1018     _(assert FreeRTOS->tasks[(TASK) pxTaskToResume])
1019     _(assert FreeRTOS->tasks[(TASK) pxCurrentTCB])
1020     _(assert (TASK) pxTaskToResume != (TASK) pxCurrentTCB)
1021
1022     _(assert \forall TASK t; FreeRTOS->tasks[t] ==> \inv(t))
1023
1024     tskTCB *pxTCB;
1025
1026     /* It does not make sense to resume the calling task. */
1027     //configASSERT( pxTaskToResume );
1028 }
1029

```

```

1020     /* Remove the task from whichever list it is currently in, and place
1021     it in the ready list. */
1022     pxTCB = ( tskTCB * ) pxTaskToResume;
1023
1024     _(assert pxTCB != pxCurrentTCB)
1025     _(assert pxTCB == (tskTCB *) pxTaskToResume)
1026
1027     /* The parameter cannot be NULL as it is impossible to resume the
1028     currently executing task. */
1029     _(assert pxCurrentTCB->\closed)
1030
1031
1032     _(unwrapping FreeRTOS){
1033         _(assert \wrapped(pxCurrentTCB))
1034         _(assert \inv(FreeRTOS))
1035
1036         if( ( pxTCB != NULL ) && ( pxTCB != _(atomic_read \embedding(&pxCurrentTCB))
1037             pxCurrentTCB ) )
1038         {
1039             _(assume \inv(FreeRTOS))
1040
1041             taskENTER_CRITICAL();
1042             {
1043                 _(assert \inv(FreeRTOS))
1044                 _(assert FreeRTOS->tasks[(TASK) pxTCB])
1045                 _(assert FreeRTOS->tasks[pxCurrentTCB])
1046                 _(assert \wrapped((TASK) pxTCB))
1047                 _(assert \wrapped((TASK) pxCurrentTCB))
1048                 _(unwrapping pxTCB){
1049                     //if( xTasksTaskSuspended( pxTCB ) == pdTRUE )
1050                     if ((xList *) (pxTCB->xGenericListItem.pvContainer) == &
1051                         xSuspendedTaskList)
1052                     {
1053                         _(assert FreeRTOS->state[(TASK) pxTCB] == suspended)
1054                         //traceTASK_RESUME( pxTCB );
1055                         /* As we are in a critical section we can access the ready
1056                         lists even if the scheduler is suspended. */
1057                         _(assert \wrapped(&(pxTCB->xGenericListItem)))
1058                         uxListRemove( &( pxTCB->xGenericListItem ) );
1059                         _(ghost FreeRTOS->state[(TASK) pxTCB] = ready)
1060                         //prvAddTaskToReadyQueue( pxTCB );
1061
1062                         _(assert \wrapped(pxCurrentTCB))
1063                         /* We may have just resumed a higher priority task. */
1064                         if( pxTCB->uxPriority >= _(atomic_read \embedding(&
1065                             pxCurrentTCB)) pxCurrentTCB->uxPriority )
1066                         {
1067                             _(assert FreeRTOS->priority[(TASK) pxTCB] == pxTCB->
1068                                 uxPriority)
1069                             /* This yield may not cause the task just resumed to
1070                             run, but
1071                             will leave the lists in the correct state for the next
1072                             yield. */
1073                             //portYIELD_WITHIN_API();
1074                             _(atomic \embedding(&pxCurrentTCB)){
1075                                 _(ghost pxCurrentTCB = pxTCB)
1076                                 _(bump_volatile_version \embedding(&pxCurrentTCB))
1077                             }
1078                         }
1079                     }
1080                 }
1081             }
1082             taskEXIT_CRITICAL();
1083             _(assume Trans())
1084         }
1085     }
1086 #endif

```


Bibliography

- [1] E. M. Clarke and J. M. Wing, "Formal methods: state of the art and future directions," *ACM Comput. Surv.*, vol. 28, pp. 626–643, December 1996.
- [2] J.-M. Jazequel and B. Meyer, "Design by contract: the lessons of Ariane," *Computer*, vol. 30, no. 1, pp. 129–130, Jan 1997.
- [3] C. A. R. Hoare and J. Misra, "Verified software: Theories, tools, experiments vision of a grand challenge project," in *Verified Software: Theories, Tools, Experiments*, ser. Lecture Notes in Computer Science, B. Meyer and J. Woodcock, Eds. Springer Berlin / Heidelberg, 2008, vol. 4171, pp. 1–18.
- [4] J. Woodcock, "First steps in the Verified Software Grand Challenge," *IEEE Computer*, vol. 39, no. 10, pp. 57–64, 2006.
- [5] C. Jones, P. O'Hearn, and J. Woodcock, "Verified software: a grand challenge," *IEEE Computer*, vol. 39, no. 4, pp. 93–95, 2006.
- [6] A. Hall, "Seven myths of formal methods," *Software, IEEE*, vol. 7, no. 5, pp. 11–19, Sep. 1990.
- [7] S. King *et al.*, "Is proof more cost-effective than testing?" *Software Engineering, IEEE Transactions on*, vol. 26, no. 8, pp. 675–686, August 2000.
- [8] V. Casey, *Software Testing and Global Industry: Future Paradigms*. United Kingdom: Cambridge Scholars Publishing, 2009.
- [9] E. W. Dijkstra, "Notes on structured programming," in *Structured programming*, O. J. Dahl, E. W. Dijkstra, and C. A. R. Hoare, Eds. London, UK, UK: Academic Press Ltd., 1972, pp. 3–6.
- [10] J. Bowen, *Formal Specification and Documentation Using Z: A Case Study Approach*. London, UK: International Thomson Computer Press, 2003.
- [11] J. Woodcock *et al.*, "Formal methods: Practice and experience," *ACM Comput. Surv.*, vol. 41, pp. 19:1–19:36, October 2009.
- [12] G. Berry, "Synchronous design and verification of critical embedded systems

using SCADE and Esterel,” in *Formal Methods for Industrial Critical Systems*, ser. Lecture Notes in Computer Science, S. Leue and P. Merino, Eds. Springer Berlin / Heidelberg, 2008, vol. 4916, pp. 2–2.

- [13] C. A. R. Hoare, “The verifying compiler: a grand challenge for computing research,” *Journal of the ACM*, vol. 50, no. 1, pp. 63–69, 2003.
- [14] J. Bicarregui, C. Hoare, and J. Woodcock, “The verified software repository: a step towards the verifying compiler,” *Formal Aspects of Computing*, vol. 18, pp. 143–151, 2006.
- [15] C. A. R. Hoare *et al.*, “The Verified Software Initiative: a manifesto,” *ACM Comput. Surv.*, vol. 41, no. 4, 2009.
- [16] J. Woodcock *et al.*, “The certification of the Mondex electronic purse to ITSEC Level E6,” *Formal Aspects of Computing*, vol. 20, no. 1, pp. 5–19, 2008.
- [17] L. Freitas, Z. Fu, and J. Woodcock, “POSIX file store in Z/Eves: an experiment in the verified software repository,” in *Engineering Complex Computer Systems, 2007. 12th IEEE International Conference on*, 2007, pp. 3–14.
- [18] R. Barry, “Using the FreeRTOS Real Time Kernel—A Practical Guide,” PDF book available from www.freertos.org, 2012.
- [19] R. Barry, “FreeRTOS Reference Manual—API functions and configuration options,” PDF book available from shop.freertos.org, 2012.
- [20] E. Cohen *et al.*, “VCC: A Practical System for Verifying Concurrent C,” in *Theorem Proving in Higher Order Logics*, ser. Lecture Notes in Computer Science, S. Berghofer *et al.*, Eds. Springer Berlin Heidelberg, 2009, vol. 5674, pp. 23–42. [Online]. Available: DOI:10.1007/978-3-642-03359-9_2
- [21] M. Moskal *et al.*, *Verifying C Programs: A VCC Tutorial*, MSR Redmond, EMIC, 2012.
- [22] M. Barnett *et al.*, “Boogie: A modular reusable verifier for object-oriented programs,” in *Formal Methods for Components and Objects*, ser. Lecture Notes in Computer Science, F. Boer *et al.*, Eds. Springer Berlin, Heidelberg, 2006, vol. 4111, pp. 364–387. [Online]. Available: DOI:10.1007/11804192_17
- [23] L. Moura and N. Bjørner, “Z3: An efficient SMT solver,” in *Tools and Algorithms for the Construction and Analysis of Systems*, ser. Lecture Notes in Computer Science, C. Ramakrishnan and J. Rehof, Eds. Springer Berlin Heidelberg, 2008, vol. 4963, pp. 337–340. [Online]. Available: DOI:10.1007/978-3-540-78800-3_24
- [24] Intel Corporation. (undated) Intel Museum – The Intel 4004. [Online]. Available: www.intel.com/about/companyinfo/museum/exhibits/4004/facts.htm

- [25] G. Moore, "Cramming more components onto integrated circuits," *Proceedings of the IEEE*, vol. 86, no. 1, pp. 82–85, Jan. 1998.
- [26] Intel Corporation. (undated) Intel® Core™ i7-680UM Processor (4m cache, 1.46 ghz) with spec code(s)slbst. [Online]. Available: ark.intel.com/Product.aspx?id=49664&processor=i7-680UM&spec-codes=SLBST
- [27] P. Bose, D. Albonesi, and D. Marculescu, "Guest editors' introduction: Power and complexity aware design," *Micro, IEEE*, vol. 23, no. 5, pp. 8–11, Sept.–Oct. 2003.
- [28] M. B. Taylor *et al.*, "Tiled multicore processors," in *Multicore Processors and Systems*, S. W. Keckler, K. Olukotun, and H. P. Hofstee, Eds. Springer US, 2009, pp. 1–33.
- [29] R. Camposano and J. Wilberg, "Embedded system design," *Design Automation for Embedded Systems*, vol. 1, pp. 5–50, 1996.
- [30] Intel Research. (undated) Teraflops research chip. [Online]. Available: techresearch.intel.com/ProjectDetails.aspx?id=151
- [31] Intel Corporation. (undated) Intel® Pentium® 4 Processor 531 supporting HT Technology (1M Cache, 3.00 GHz, 800 MHz FSB). [Online]. Available: ark.intel.com/Product.aspx?id=27465
- [32] I. Meisels and M. Saaltink, *Z/Eves 1.5 Reference Manual*, ORA Canada, September 1997, tR-97-5493-03d.
- [33] ProB Documentation. (2010) The ProB Animator and Model Checker. [Online]. Available: www.stups.uni-duesseldorf.de/ProB/index.php5/Main_Page
- [34] Real Time Engineers Ltd. (undated) The FreeRTOS project. [Online]. Available: www.freertos.org
- [35] L. G. Valiant, "A bridging model for multi-core computing," in *Proceedings of the 16th Annual European Symposium on Algorithms*, ser. ESA '08, 2008, pp. 13–28.
- [36] Real Time Engineers Ltd. (undated) Coding Standard and Style Guide. [Online]. Available: <http://www.freertos.org/FreeRTOS-Coding-Standard-and-Style-Guide.html>
- [37] S. Cheng, J. Woodcock, and D. D'Souza, "Using formal reasoning on a model of tasks for FreeRTOS," *Formal Aspects of Computing*, vol. 27, no. 1, pp. 167–192, 2015. [Online]. Available: DOI:10.1007/s00165-014-0308-9
- [38] I. D. Craig, *Formal Models of Operating System Kernels*. Springer, 2006.

- [39] I. D. Craig, *Formal Refinement for Operating System Kernels*. Springer, 2007.
- [40] J. J. Labrosse, *MicroC OS II: The Real Time Kernel*. Newnes, 2002.
- [41] L. Freitas and J. Woodcock, “A chain datatype in Z,” *Int. J. Software and Informatics*, vol. 3, no. 2–3, pp. 357–374, 2009.
- [42] E. Börger and I. Craig, “Modeling an operating system kernel,” in *Informatik als Dialog zwischen Theorie und Anwendung*, V. Diekert, K. Weicker, and N. Weicker, Eds. Vieweg+Teubner, 2009, pp. 199–216. [Online]. Available: DOI:10.1007/978-3-8348-9982-8_17
- [43] G. Klein, “Operating system verification—an overview,” *Sādhanā*, vol. 34, no. 1, pp. 27–69, Feb. 2009.
- [44] G. Klein, “The L4.verified project—next steps,” in *Verified Software: Theories, Tools, Experiments, Third International Conference, VSTTE 2010, Edinburgh, UK, August 16-19, 2010. Proceedings*, ser. Lecture Notes in Computer Science, G. T. Leavens, P. W. O’Hearn, and S. K. Rajamani, Eds., vol. 6217. Springer, 2010, pp. 86–96.
- [45] G. Klein, “A formally verified OS kernel. Now what?” in *Interactive Theorem Proving, First International Conference, ITP 2010, Edinburgh, UK, July 11–14, 2010. Proceedings*, ser. Lecture Notes in Computer Science, M. Kaufmann and L. C. Paulson, Eds., vol. 6172. Springer, 2010, pp. 1–7.
- [46] G. Klein, “From a verified kernel towards verified systems,” in *Programming Languages and Systems—8th Asian Symposium, APLAS 2010, Shanghai, China, November 28–December 1, 2010. Proceedings*, ser. Lecture Notes in Computer Science, K. Ueda, Ed., vol. 6461. Springer, 2010, pp. 21–33.
- [47] T. Sewell *et al.*, “seL4 enforces integrity,” in *Interactive Theorem Proving—Second International Conference, ITP 2011, Berg en Dal, The Netherlands, August 22–25, 2011. Proceedings*, ser. Lecture Notes in Computer Science, M. C. J. D. van Eekelen *et al.*, Eds., vol. 6898. Springer, 2011, pp. 325–340. [Online]. Available: DOI:10.1007/978-3-642-22863-6_24
- [48] R. Buerki and A.-K. Rueegsegger, “Muen—an x86/64 separation kernel for high assurance,” University of Applied Sciences Rapperswil (HSR), Tech. Rep., 2013.
- [49] D. Déharbe, S. Galvão, and A. M. Moreira, “Formalizing FreeRTOS: first steps,” in *Formal Methods: Foundations and Applications, 12th Brazilian Symposium on Formal Methods, SBMF 2009, Gramado, Brazil, August 19–21, 2009, Revised Selected Papers*, ser. Lecture Notes in Computer Science, M. V. M. Oliveira and J. Woodcock, Eds., vol. 5902. Springer, 2009, pp. 101–117. [Online]. Available: DOI:10.1007/978-3-642-10452-7_8

- [50] C. Pronk, “Verifying FreeRTOS: a feasibility study,” Delft University of Technology, Software Engineering Research Group, Tech. Rep. TUD-SERG-2010-042, 2010.
- [51] Y. Lin, “Formal Analysis of FreeRTOS,” Master’s thesis, University of York, 2010.
- [52] R. Arthan, “ProofPower,” Web page www.lemma-one.com/ProofPower/, 2012.
- [53] L. Paulson, “Isabelle,” Web page www.cl.cam.ac.uk/research/hvg/isabelle/, 2012.
- [54] S. Owre, N. Shankar, and J. Rushby, “PVS,” Web page: pvs.csl.sri.com, 2012.
- [55] J. T. Muehlberg and L. Freitas, “Verifying FreeRTOS: from requirements to binary code,” in *Proceedings of the 11th International Workshop on Automated Verification of Critical Systems (AVoCS 2011)*, ser. Electronic Communications of the EASST, J. Bendisposto *et al.*, Eds., vol. X, 2011, pp. 1–2.
- [56] J. T. Mühlberg and G. Lüttgen, “Symbolic object code analysis,” in *Model Checking Software—17th International SPIN Workshop, Enschede, The Netherlands, September 27–29, 2010. Proceedings*, ser. Lecture Notes in Computer Science, J. van de Pol and M. Weber, Eds., vol. 6349. Springer, 2010, pp. 4–21.
- [57] B. Jacobs, J. Smans, and F. Piessens, “A quick tour of the VeriFast program verifier,” in *Proceedings of the 8th Asian Conference on Programming Languages and Systems*, ser. APLAS’10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 304–311. [Online]. Available: dl.acm.org/citation.cfm?id=1947873.1947902
- [58] J. Ferreira, G. He, and S. Qin, “Automated verification of the FreeRTOS scheduler in HIP/SLEEK,” in *6th International Symposium on Theoretical Aspects of Software Engineering (TASE’12)*, 4–6 July 2012.
- [59] W. Su and J.-R. Abrial, May 2011, private communication.
- [60] J. Mistry, “FreeRTOS and multicore,” September 2011, MSc Dissertation, Department of Computer Science, University of York.
- [61] J. Mistry, M. Naylor, and J. Woodcock, “Adapting FreeRTOS for multicores: an experience report,” *Software: Practice and Experience*, vol. 44, no. 9, pp. 1129–1154, 2014. [Online]. Available: DOI:10.1002/spe.2188
- [62] S. Divakaran, D. D’Souza, and N. Sridhar, “Efficient refinement checking in VCC,” in *Verified Software: Theories, Tools and Experiments*, ser. Lecture Notes in Computer Science, D. Giannakopoulou and D. Kroening, Eds. Springer International Publishing, 2014, vol. 8471, pp. 21–36. [Online].

Available: http://dx.doi.org/10.1007/978-3-319-12154-3_2

- [63] I. Sørensen, “A specification language,” in *Program Specification*, ser. Lecture Notes in Computer Science, J. Staunstrup, Ed. Springer Berlin / Heidelberg, 1982, vol. 134, pp. 381–401.
- [64] I. Toyn, Ed., *Information Technology – Z Formal Specification Notation – Syntax, Type System and Semantics*, 1st ed., ser. ISO/ICE 13568:2002. ISO, Oct. 2002.
- [65] J. Woodcock and J. Davies, *Using Z: Specification, Refinement, and Proof*, ser. International Series in Computer Science. Prentice-Hall, 1996.
- [66] J. M. Spivey, *The Z Notation: A Reference Manual*, 2nd ed., ser. Series in Computer Science. Prentice Hall International, 1992.
- [67] C. A. R. Hoare, “An axiomatic basis for computer programming,” *Commun. ACM*, vol. 12, pp. 576–580, October 1969.
- [68] J. M. Spivey. (1990) A guide to the zed style option. [Online]. Available: www.ctan.org/tex-archive/macros/latex/contrib/zed-csp/zed2e.pdf
- [69] M. Saaltink, *Z/Eves 2.0 User's Guide*, ORA Canada, 1999, TR-99-5493-06a.
- [70] M. Saaltink, *Z/Eves 2.0 Mathematical Toolkit*, ORA Canada, October 1999, TR-99-5493-05b.
- [71] G. Klein *et al.*, “seL4: formal verification of an OS kernel,” in *SOSP*, 2009, pp. 207–220.
- [72] Community Z Tools Project. (2013, November) Community Z Tools. [Online]. Available: czt.sourceforge.net
- [73] T. P. Baker, “A comparison of global and partitioned EDF schedulability tests for multiprocessors,” in *Proceedings of the International Conference on Real-Time and Network Systems*, 2006, pp. 119–130.
- [74] M. R. Garey, R. L. Graham, and J. D. Ullman, “Worst-case analysis of memory allocation algorithms,” in *Proceedings of the fourth annual ACM symposium on Theory of computing*, ser. STOC '72. New York, NY, USA: ACM, 1972, pp. 143–150. [Online]. Available: DOI:10.1145/800152.804907