



Modeling and Analysis Suite for Real Time Applications (MAST 1.6.0)

Restrictions

By:	José María Drake	drakej@unican.es
	Michael González Harbour	mgh@unican.es
	José Javier Gutiérrez	gutierjj@unican.es
	José Carlos Palencia	palencij@unican.es
	Julio Luis Medina	medinajl@unican.es

Copyright © 2000-2025 Universidad de Cantabria, SPAIN

1. Introduction

The MAST model has some general restrictions, that affect all application models, and other restrictions that affect specific analysis tools. Some of the restrictions are marked as temporary, because they will be eliminated in the future.

In addition to the restrictions, there are some consistency checks that are made to determine if a MAST description is correct. They are based on a set of rules.

2. Consistency Checks

They are used to determine whether a MAST description is correct and if it covers the most basic general restrictions. They are based on the following set of rules, which are checked for each transaction in the system:

1. At least one external event
2. Each external event link directed at one event handler, and with an external event.
3. Each internal event link comes from an event handler
4. Each simple event handler has an input event link and an output event link
5. Each input event_handler has 2 or more input event links and an output event link
6. Each output event handler has 2 or more output event links and an input event link
7. No circular dependencies in the transaction graph
8. No isolated event links
9. No isolated event handlers
10. All activities have an operation
11. All activities have a scheduling server
12. All scheduling servers have a scheduler
13. All scheduling servers have scheduling parameters
14. All rate divisors, offset and delay event handlers are only followed by activities
15. All primary schedulers have a processing resource



16. All secondary schedulers have a scheduling server
17. All schedulers have a scheduling policy
18. All schedulers with a scheduling policy of the type `FP_Packet_Based` are associated with a primary scheduler located on a network
19. All scheduling servers with parameters of the fixed priority family (excluding interrupts) are associated with schedulers having a policy of the fixed priorities family
20. All scheduling servers with parameters of the EDF family are associated with schedulers having a policy of the EDF family
21. All scheduling servers with parameters of the type `Interrupt_FP_Policy` are associated with primary schedulers
22. All message transmission operations are executed by scheduling servers executing on a network
23. Each processing resource has at most one primary scheduler
24. Each scheduling server has at most one secondary scheduler
25. The size of each message sent through a network driver that does not support message partitioning is smaller than the maximum allowable message size
26. Periods are not null. This includes:
 - 26: external event periods,
 - 26a: minimum interarrival times for sporadic tasks,
 - 26b: ticker periods for timers,
 - 26c: polling periods for polling servers, and
 - 26d: replenishment periods for sporadic servers
27. All scheduling policies of a type that is not `FP_Packet_Based` are associated with a primary scheduler located on a processor

3. Restrictions Related to the Use of Shared Resources

The following restrictions ensure a proper usage of shared resources through the transaction graphs. The following rules are checked for each transaction in the system:

1. All locked resources are unlocked
2. No resource is locked if it was already locked
3. No resource is unlocked if not previously locked
4. All locked resources in a segment are unlocked in that segment¹

The following rules have been added to facilitate the calculation of remote blocking terms:

5. All global shared resources (i.e., those that are shared by operations executed by different processors) use the *Immediate_Ceiling_Protocol*, for mutually-exclusive synchronization.

1. A segment is a set of consecutive activities executed by the same scheduling server



6. The priority ceiling of all global shared resources is such that the tasks that are involved in the calculation of the remote blocking (i.e., those tasks that have priorities higher than or equal to the priority ceiling) do not use shared resources.

The latter restriction is imposed because, otherwise, there would be a circular dependency among the remote blockings of the different tasks, which would make calculations very complex. But we allow some degree of remote blocking, for example by interrupt service routines, or by very high priority tasks that do not need shared resources.

When the *Immediate_Ceiling_Protocol* is used, if the tools are invoked with the “Calculate Ceilings” option set, then the optimum priority ceilings are calculated for each shared resource using this protocol. If the option is not set, then a check is made to determine that all priority ceilings are correct:

7. Consistent priority ceilings: no task locks an immediate ceiling resource having a priority lower than the ceiling.

In addition, we have a temporary restriction caused by the fact that no analysis techniques have been developed to calculate remote blocking effects in multiprocessors, when the blocking terms in one processor depend on the results of the blocking calculations in the other processors. We expect to eliminate this restriction in the future:

8. All resources locked in a code segment are unlocked in that segment. A *segment* is defined as an uninterrupted sequence of activities that are all executed by the same scheduling server (and thus, by the same processing resource).

Finally, a temporary restriction needs to be applied to distributed systems scheduled under EDF, which cannot make use of shared resources. The reason is that current analysis techniques don't know how to calculate the preemption levels of shared resources in distributed systems because preemption levels depend on jitter terms, which in turn depend on the response times which depend on the assigned preemption levels. It is necessary to solve this circular dependency issue in real-time scheduling theory.

9. Distributed systems scheduled under EDF cannot use shared resources.

4. Restrictions Related to Network Drivers

The following restrictions ensure consistent structure for network drivers:

1. There is at most one driver per processor-network pair
2. Every driver has a scheduling server
3. Every driver has a send and receive operation

5. System-Kind Restrictions

The following are system restrictions that apply to some of the tools only:

- *Monoprocessor Only*: only one processing resource, which is a *Processor*
- *Fixed Priority Only*: All Scheduling Servers have Fixed Priority Parameters, all overridden parameters in operations are Fixed Priority, and all priorities are within the appropriate ranges for their processing resources. There are no secondary schedulers.




- *EDF_Only*: All Scheduling Servers have EDF or Interrupt Parameters, there are no overridden parameters in operations, and all interrupt priorities are within the appropriate ranges for their processing resources. There are no secondary schedulers.
- *EDF_Within_Priorities_Only*: The primary schedulers have a fixed priority policy. All secondary schedulers have an EDF policy and are scheduled under a scheduling server that is directly attached to a primary scheduler. All operations with overridden priorities are executed by fixed-priority scheduling servers. All priorities are in range.
- *FP_Or_EDF_Only*: Each node is *Fixed_Priorities_Only* or *EDF_Only*. Each node is *Fixed_Priorities_Only* or *EDF_Only*.
- *Flat_FP_Or_EDF_Only*: Each node is *Fixed_Priorities_Only* or *EDF_Only*. Each node is *Fixed_Priorities_Only* or *EDF_Only* and there are no overridden priorities in operations
- *PCP_Or_Priority_Inheritance_Only*: all resources are PCP (i.e., *Immediate_Ceiling_Resource*) or Priority Inheritance resources
- *PCP_SRP_Or_Priority_Inheritance_Only*: all resources are PCP (i.e., *Immediate_Ceiling_Resource*), SRP, or Priority Inheritance resources. In addition, we restrict SRP resources to be used only by EDF tasks, and Priority inheritance resources to be used only by fixed priority tasks.
- *SRP_Only*: all resources are SRP resources
- *SRP_Or_PCP_Only*: All resources are SRP or PCP.
- *Referenced_Events_Are_External_Only*: No internal events are referenced by global timing requirements.
- *Simple_Transactions_Only*: Checks that every transaction has only one segment. A segment is a continuous sequence of activities executed by the same server.
- *Linear_Transactions_Only*: checks that every transaction only has one external event and that its event handlers are all *Activities*.
- *Linear_Plus_Transactions_Only*: Checks that every transaction is regular and is not multipath, i.e., and has no *Concentrators* or *Delivery_Servers*, or *Query_Servers* or *Multicasts*, or *Barriers*
- *Regular_Transactions_Only*: Checks that every transaction is regular.
- *No_Permanent_Overridden_Priorities*: Checks that there are no operations with permanent overridden priorities.
- *No_Permanent_FP_Inside_Composite_Operations*: Checks that there are no operations with overridden priorities inside composite operations. If this would be required, the composite operation would have to be broken up into several activities.
- *No_Intermediate_Timing_Requirements*: No timing requirements are attached to intermediate events, i.e., all timing requirements are attached to output events
- *No_Rate_Divisors_In_Multipath_Transactions*: There are no rate divisors in multipath transactions.
- *No_Shared_Resources*: There are no shared resources in the system.
- *No_Hard_Local_Deadlines*: There are no hard local deadlines




- *Restricted_Multipath_Transactions_Only*: Checks that every transaction has a single input event, has no branch elements (delivery or query servers), and has no rate divisors. It also checks that the transaction follows the set of allowed constructs mentioned below.

Allowed multipath constructs:


A. Fork at the beginning




B. Fork followed by join



C. Fork followed by merge




D. Merge followed by fork






Forbidden multipath constructs:


- E. Join followed by fork: This construct might be convenient as a modelling element, but is currently difficult to implement. A workaround for this construct is to implement a dummy processor (with zero context switch time) that has a dummy task (with zero execution time). By using the dummy task between the join and the fork you ensure that the response time of the output events of the join and the fork have the same value, since both the execution time and the context switch of the processor are zero.



- F. Fork followed by fork: This can be implemented using a single fork element




- G. Join followed by join: This can be implemented using a single join element







H. Merge followed by merge: This can be implemented using a single merge element




I. Merge at the end: Unnecessary as a modelling element



J. Join at the end: Unnecessary as a modelling element




K. Fork at the end: Unnecessary as a modelling element





- L. Join after merge. This case is difficult to model internally in the analysis tools. If necessary, this restriction could be removed in the future



All schedulability analysis tools require:

- *Referenced_Events_Are_External_Only.*
- *No_Hard_Local_Deadlines*

In addition, some tools require other restrictions, as shown in the following table:

TABLE 1. Response Time Analysis Tools

Tool	Restrictions
Classic_RM	Monoprocessor_Only, Simple_Transactions_Only, Fixed_Priorities_Only, No_Permanent_Overridden_Priorities
Varying_Priorities	Monoprocessor_Only, Linear_Transactions_Only, No_Permanent_FP_Inside_Composite_Operations, Fixed_Priorities_Only, No_Intermediate_Timing_Requirements
EDF_Monoprocessor	Monoprocessor_Only, Simple_Transactions_Only, EDF_Only, SRP_Only
EDF_Within_Priorities	Monoprocessor_Only, Simple_Transactions_Only, EDF_Within_Priorities_Only
Holistic	Restricted_Multipath_Transactions_Only, FP_Or_EDF_Only ^{ab} , No_Permanent_Overridden_Priorities, No_Rate_Divisors_In_Multipath_Transactions
Offset_Based_Approx	Restricted_Multipath_Transactions_Only, FP_Or_EDF_Only ^{ab} , No_Permanent_Overridden_Priorities, No_Rate_Divisors_In_Multipath_Transactions
Offset_Based_Approx_W_Pr	Linear_Plus_Transactions_Only ^{ac} , Fixed_Priority_Only, No_Permanent_Overridden_Priorities, No_Rate_Divisors_In_Multipath_Transactions



TABLE 1. Response Time Analysis Tools

Tool	Restrictions
Offset_Based_Slanted	Restricted_Multipath_Transactions_Only, FP_Or_EDF_Only ^{ab} , No_Permanent_Overridden_Priorities, No_Rate_Divisors_In_Multipath_Transactions
Offset_Based_Brute_Force	Restricted_Multipath_Transactions_Only, FP_Or_EDF_Only ^{ab} , No_Permanent_Overridden_Priorities, No_Rate_Divisors_In_Multipath_Transactions

a. The analysis in EDF nodes does not yet support shared resources. Restriction: No_Shared_Resources

b. The analysis in multipath transactions is not available for EDF

c. The analysis with precedence constraints is not available for global EDF, only for local EDF.

The Shared resources analysis tools have the following restrictions:

- *PCP_SRP_Or_Priority_Inheritance_Only*

In addition, each of the following tools has the following additional restrictions:

TABLE 2. Shared Resource Analysis Tools

Tool	Restrictions
Calculate_Ceilings_And_Levels	Linear_Plus_Transactions_Only
Calculate_Blocking_Times	Linear_Plus_Transactions_Only

The scheduling parameters assignment tools have the same restrictions as the response time analysis tools they use. In addition, each tool has the following additional restrictions:

TABLE 3. Scheduling Parameters Assignment Tools

Tool	Restrictions
Monoprocessor_Priority_Assignment	Monoprocessor_Only, Simple_Transactions_Only
HOSPA	Linear_Plus_Transactions_Only, FP_Or_EDF_Only
Proportional distribution	Linear_Plus_Transactions_Only, FP_Or_EDF_Only
Normalized distribution	Linear_Plus_Transactions_Only, FP_Or_EDF_Only
Linear_Simulated_Annealing_Assignment	Linear_Plus_Transactions_Only, Fixed_Priorities_Only
EQS (Equal Slack)	Linear_Plus_Transactions_Only, FP_Or_EDF_Only
EQF (Equal Flexibility):	Linear_Plus_Transactions_Only, FP_Or_EDF_Only



TABLE 3. Scheduling Parameters Assignment Tools

Tool	Restrictions
UD (Ultimate Deadline)	Linear_Plus_Transactions_Only, FP_Or_EDF_Only
ED (End Deadline)	Linear_Plus_Transactions_Only, FP_Or_EDF_Only