

Example Specifications of Non-functional Properties of a Simple Counter Application

Technical Report COMP-006-2008, Computing Department, Lancaster University

Steffen Zschaler
Computing Department
Lancaster University
Lancaster, United Kingdom
szschaler@acm.org

November 3, 2008

1 Introduction

This document lists the complete TLA⁺ specifications for the main example from [3]. It is meant to serve as an external appendix to that paper in order to improve understanding and provide additional detail that could not be included in the journal-paper version for space reasons. We begin by giving some more TLA⁺ background before discussing the individual specification modules one after the other.

2 Background

This section provides additional TLA⁺ background beyond what has been given in [3]. Of course, we cannot go into all detail of TLA⁺; for this, the interested reader is referred to [1], which is an excellent textbook on the language and logics. Here, we focus on the things necessary for understanding our specifications.

TLA⁺ specifications are divided into modules. Each module starts with a line containing the `MODULE` keyword and the name of the module. Modules may contain arbitrarily many horizontal divider lines. These are only used for visual structuring and have no formal semantics. Modules may extend other modules (using the `EXTENDS` keyword on the first line), importing all definitions of all extended modules. Modules may also instantiate other modules, by using the `INSTANCE` keyword and mapping all variables and constants of the instantiated module to variables and constants of the instantiating module. Modules may also contain inner modules. These are only available within their containing module and cannot be instantiated from anywhere else. Inner modules can use all definitions of the outer module directly. The outer module can only make use of the definitions within an inner module by instantiating the inner module. In such an instantiation, only the variables and constants defined in the inner module must be mapped. In the specifications below we will use inner modules to allow us to hide helper variables from users of the outer module. The basic pattern is to define all externally visible variables as variables of the outer module and all helper variables

as variables of the inner module. The inner module is then instantiated in the outer module, using existential quantification to provide values for the helper variables. Understanding a TLA⁺ module is best done beginning from the end. Typically, the last formulas in a TLA⁺ module are the ones that are really of interest. Everything before is often defined to help with the definition of these interesting formulas. Often, a TLA⁺ module defines a state machine. Such a definition looks like this:

$$Spec \triangleq \wedge Init \\ \wedge \square [Next]_{vars}$$

Spec is the name of the state-machine specification defined. *Init* and *Next* refer to a previously defined predicate and a previously defined action. *Init* describes the possible initial states of the state machine and *Next* describes what can happen in a step, using a disjunction of individual actions describing individual step alternatives. *vars* is a collection of all variables relevant for the state machine. Often, this is given directly as a sequence of the relevant variables, written $\langle a, b, c \rangle$.

In defining measurements and other forms of specifications, we will use a form of specification that is very close to aspect-oriented programming. The final formula looks very similar to the definition of a state machine as described above. However, *Next* is a *conjunction* of alternatives and each alternative is defined as an implication $A \Rightarrow B$, where *A* is an action describing an alternative step from some base state-machine specification and *B* is an action that should be executed whenever *A* is executed. As has been discussed in [3] and in more detail in [2], this form of specification effectively adds *B* to the base state machine whenever *A* holds.

3 Specification of Time

The first module defines the notion of *time*. It has been taken and slightly modified from [1]. The main modification is that we have separated safety and liveness parts of the specification so that we can use the safety part of the definition independently. Time is captured by the new variable *now*.

```

1  ┌────────────────────────── MODULE RealTime ───────────────────────────┐
  │ This is based on the original RealTime specification from the TLA toolkit, but we removed the liveness │
  │ part—that is, NZ(v)—from RTnow.                                     │
6  EXTENDS Reals
  │
  │ Variables:
  │ now - the current system time.
13 VARIABLE now
  │
15  ┌────────────────────────── A helper definition ───────────────────────────┐
16  LOCAL NowNext(v) ≜ ∧ now' ∈ {r ∈ Real : r > now}
17  └────────────────────────── ∧ UNCHANGED v ───────────────────────────┘
  │
  │ RTnow(v) asserts two things: a) time never runs backward, b) steps changing now do not change any other
  │ variable in v, and vice versa
  │
  │ RTnow(v) is a safety property, that is, it allows systems in which time stops altogether. This is useful
  │ for certain proofs. If one needs to explicitly exclude this possibility, one conjoins NZ(v), which adds the
  │ required fairness constraints.
29 RTnow(v) ≜ ∧ now ∈ Real

```




Figure 1: State-machine representation of the service-operation context model

30 $\wedge \square [NowNext(v)]_{now}$

The so called *NonZeno* condition, which asserts that time will eventually exceed every bound. This liveness constraint is only required under certain circumstances.

37 $NZ(v) \triangleq \forall r \in Real :$
 38 $\quad WF_{now}(NowNext(v) \wedge (now' > r))$
 39

4 Context Model Definition

We define two context models: *Service* defines a service operation and will be the basis for defining the response-time measurement, *Component* defines a component operation and will be the basis for defining the execution-time measurement. These context models have already been discussed in their state-machine form in the main paper, but here we show the TLA⁺ specifications.

4.1 A Context Model for Service Operations

Figure 1 gives the state-machine view of this context model again. This has already been shown in [3], but we show it here again to simplify understanding of the formal specification. Notice that the TLA⁺ specification differs from the graphical rendering in two respects:

1. State information has been divided into two parts: Variable *inState* captures if the service is currently idle or is handling a request. Additionally, variable *unhandledRequest* captures if a request is currently waiting to be handled.
2. The specification has been split into an environment specification and a service specification. This has been done to simplify proofs of feasibility further down the line. It can be shown, however, that this form of specification can be transformed to a form that uses only one integrated state machine.

1 $\overline{\text{MODULE } Service}$

Service Context Model

Variables:

inState – the current state of the service execution machinery.
unhandledRequest – TRUE indicates a fresh request has been placed in the system.

13 VARIABLES *inState*, *unhandledRequest*

15 vars $\triangleq \langle inState, unhandledRequest \rangle$

17 |

18 | **The environment model**

20 | **Initially there are no requests.**

21 | $InitEnv \triangleq unhandledRequest = FALSE$

The environment sets the *unhandledRequest* flag at some arbitrary moment to indicate a new request.

27 | $RequestArrival \triangleq \wedge unhandledRequest = FALSE$

28 | $\wedge unhandledRequest' = TRUE$

29 | $\wedge UNCHANGED\ inState$

31 | **Somebody, but not the environment, will collect the request**

32 | **Also, *inState* changes independently of the environment**

33 | $ServAgent \triangleq \vee \wedge unhandledRequest = TRUE$

34 | $\wedge unhandledRequest' = FALSE$

35 | $\vee \neg UNCHANGED\ inState$

37 | $EnvSpec \triangleq \wedge InitEnv$

38 | $\wedge \square [RequestArrival \vee ServAgent]_{vars}$

41 | **The actual service.**

43 | **Initially we start out in the *Idle* state**

44 | $InitServ \triangleq inState = \text{"Idle"}$

The transition from idle to handling request is triggered by an incoming request

50 | $StartRequest \triangleq \wedge inState = \text{"Idle"}$

51 | $\wedge unhandledRequest = TRUE$

52 | $\wedge inState' = \text{"HandlingRequest"}$

53 | $\wedge unhandledRequest' = FALSE$

55 | **Request handling can finish any time**

56 | $FinishRequest \triangleq \wedge inState = \text{"HandlingRequest"}$

57 | $\wedge inState' = \text{"Idle"}$

58 | $\wedge UNCHANGED\ unhandledRequest$

60 | $NextServ \triangleq StartRequest \vee FinishRequest$

62 | **The environment occasionally provides new requests**

63 | $EnvAgent \triangleq \wedge unhandledRequest = FALSE$

64 | $\wedge unhandledRequest' = TRUE$

66 | $ServiceSpec \triangleq \wedge InitServ$

67 | $\wedge \square [\vee NextServ$

68 | $\vee EnvAgent]_{vars}$

70 |

72 | $Service \triangleq EnvSpec \multimap ServiceSpec$

74 |




Figure 2: State-machine representation of the component-operation context model

4.2 A Context Model for Component Operations

Figure 2 gives the state-machine view of this context model for comparison with the specification. The same notes as in the previous subsection apply also for this context model specification.

```

1 |----- MODULE Component -----|
  | Context Model of a component implementation. |
6 |-----|
  | Variables: |
  | unhandledRequest – set to TRUE by the environment to indicate that a new request has arrived and should |
  | be handled. |
  | inState – the state in which the component is. |
15 VARIABLE inState
16 VARIABLE unhandledRequest
18 vars  $\triangleq$   $\langle inState, unhandledRequest \rangle$ 
20 |-----|
  | The environment specification. |
  | The environment in particular influences the unhandledRequest variable by entering new requests into the |
  | system. |
29 | Initially there are no requests in the system |
30 InitEnv  $\triangleq$  unhandledRequest = FALSE
  | The environment sets the unhandledRequest flag at some arbitrary moment to indicate a new request. |
36 RequestArrival  $\triangleq$   $\wedge$  unhandledRequest = FALSE
37  $\wedge$  unhandledRequest' = TRUE
38  $\wedge$  UNCHANGED inState
40 | Somebody, but not the environment, will collect the request |
41 | Also, inState changes independently of the environment |
42 CompAgent  $\triangleq$   $\vee$   $\wedge$  unhandledRequest = TRUE
43  $\wedge$  unhandledRequest' = FALSE
44  $\vee$   $\neg$ UNCHANGED inState
46 EnvSpec  $\triangleq$   $\wedge$  InitEnv

```

47 $\wedge \square [RequestArrival \vee CompAgent]_{vars}$

49

The actual component.

It mainly specifies changes to the *inState* variable, however it communicates with the environment via the *unhandledRequest* variable.

58 Initially we start out in the idle state

59 $InitComponent \triangleq inState = \text{"Idle"}$

61 The transition from idle to handling request is triggered by an

62 incoming request

63 $StartRequest \triangleq \wedge inState = \text{"Idle"}$
64 $\wedge unhandledRequest = \text{TRUE}$
65 $\wedge \vee inState' = \text{"HandlingRequest"}$
66 $\vee inState' = \text{"Blocked"}$
67 $\wedge unhandledRequest' = \text{FALSE}$

69 Request handling can finish any time

70 $FinishRequest \triangleq \wedge inState = \text{"HandlingRequest"}$
71 $\wedge inState' = \text{"Idle"}$
72 $\wedge \text{UNCHANGED } unhandledRequest$

74 Also, the runtime environment may at any time take away the

75 CPU from us and assign it to someone else.

76 $SwitchToOther \triangleq \wedge inState = \text{"HandlingRequest"}$
77 $\wedge inState' = \text{"Blocked"}$
78 $\wedge \text{UNCHANGED } unhandledRequest$

80 But, it may also at any time give back the CPU to us

81 $SwitchBack \triangleq \wedge inState = \text{"Blocked"}$
82 $\wedge inState' = \text{"HandlingRequest"}$
83 $\wedge \text{UNCHANGED } unhandledRequest$

85 $NextComponent \triangleq \vee StartRequest \vee FinishRequest$
86 $\vee SwitchToOther \vee SwitchBack$

88 The environment occasionally provides new requests

89 $EnvAgent \triangleq \wedge unhandledRequest = \text{FALSE}$
90 $\wedge unhandledRequest' = \text{TRUE}$

92 $ComponentSpec \triangleq \wedge InitComponent$
93 $\wedge \square [\vee NextComponent$
94 $\vee EnvAgent]_{vars}$

96

The complete specification

102 $Component \triangleq EnvSpec \multimap ComponentSpec$

104

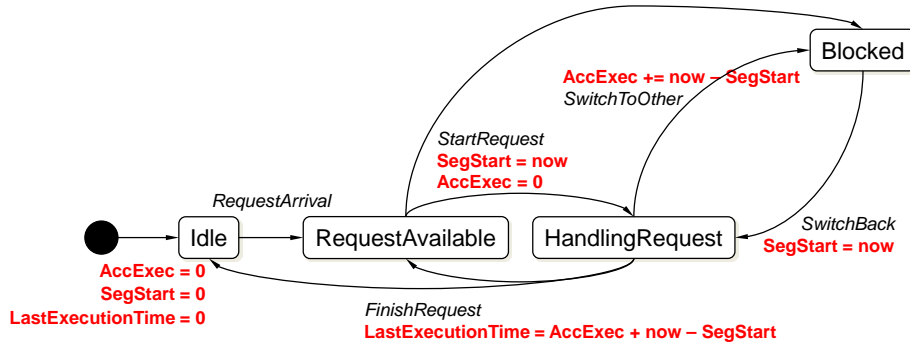


Figure 3: Definition of the execution-time measurement

5 Measurement Definition

Based on the context models defined in the previous section, we can now define measurements. In this section, we will define three different measurements:

1. *Execution time*: This intrinsic measurement is based on the component context model from Sect. 4.2 and represents the execution time of the last invocation of a component operation.
2. *Response time*: This extrinsic measurement is based on the service context model from Sect. 4.1 and represents the response time of the last invocation of a service operation.
3. *Inter-request time*: This extrinsic measurement is based on the service context model from Sect. 4.1 and represents the time between the two last invocations of a service operation.

The following subsections present these measurement specifications in full detail. Additionally, each module will also predefine a parametrised property based on the measurement.

5.1 Execution Time

Based on the `Component` context model, we define the execution-time measurement. This specification uses the `Component` module (Line 47) and attaches actions measuring the time the component spends actually computing (this is specified on Lines 55–78). The execution time of the last completed invocation of the operation is stored in variable `LastExecutionTime`. In addition to defining the measurement, Line 106 adds a constraint on execution time. This is parametrised by an upper-bound value passed in through the `ExecutionTime` parameter defined on Line 14.

Figure 3 shows the corresponding state-machine representation. The new variables and actions introduced by the measurement definition are high-lighted in red.

```

1 |----- MODULE ExecTimeConstrainedComponent -----|
  | Specification of a component which offers one operation the execution time of which can be constrained. |
6 | EXTENDS RealTime
  |
  | Parameters:
  |-----|
  | ExecutionTime – an upper bound for the execution time of the component’s operation.
  |-----|
14 | CONSTANT ExecutionTime
15 | ASSUME
16 | (ExecutionTime ∈ Real) ∧ (ExecutionTime > 0)
  |
  | Variables:
  |-----|
  | inState – the state in which the component currently is.
  | unhandledRequest – TRUE if the environment put another request into the system.
  | LastExecutionTime – the execution time of the last service execution.
  |-----|
26 | VARIABLES inState, unhandledRequest
27 | VARIABLE LastExecutionTime
  |
  |-----|
31 |----- MODULE Inner -----|
  | Internal module containing the actual specification.
  |-----|
  | Variables:
  |-----|
  | AccExec – The accumulated execution time of the current service execution.
  | SegStart – The start time of the current service execution.
  |-----|
43 | VARIABLE AccExec
44 | VARIABLE SegStart
  |
  | Based on the component context model
46 | BasicComponent ≜ INSTANCE Component
  |
  |-----|
49 |-----|
51 | Init ≜ ∧ AccExec = 0
52 |           ∧ SegStart = 0
53 |           ∧ LastExecutionTime = 0
  |
  | StartNext reacts to a StartRequest step
55 | StartNext ≜ BasicComponent!StartRequest ⇒
56 |           ∧ SegStart' = now
57 |           ∧ AccExec' = 0
58 |           ∧ UNCHANGED LastExecutionTime
59 |
  | RespNext reacts to a FinishRequest step
61 | RespNext ≜ BasicComponent!FinishRequest ⇒
62 |           ∧ LastExecutionTime' =
63 |             AccExec + now – SegStart
64 |           ∧ UNCHANGED ⟨SegStart, AccExec⟩
65 |
  | STONext reacts to a SwitchToOther step
67 | STONext ≜ BasicComponent!SwitchToOther ⇒
68 |           ∧ AccExec' =
69 |

```



```

70           $AccExec + now - SegStart$ 
71           $\wedge \text{UNCHANGED } \langle LastExecutionTime,$ 
72              $SegStart \rangle$ 

74  SBNext reacts to a SwitchBack step
75   $SBNext \triangleq BasicComponent!SwitchBack \Rightarrow$ 
76      $\wedge SegStart' = now$ 
77      $\wedge \text{UNCHANGED } \langle LastExecutionTime,$ 
78         $AccExec \rangle$ 

80   $ExcludeOtherChange \triangleq$ 
81      $(\neg \vee BasicComponent!StartRequest$ 
82         $\vee BasicComponent!FinishRequest$ 
83         $\vee BasicComponent!SwitchToOther$ 
84         $\vee BasicComponent!SwitchBack)$ 
85      $\Rightarrow \text{UNCHANGED } \langle AccExec, SegStart, LastExecutionTime \rangle$ 

88   $Next \triangleq \wedge StartNext$ 
89      $\wedge RespNext$ 
90      $\wedge STONext$ 
91      $\wedge SBNext$ 
92      $\wedge ExcludeOtherChange$ 

94   $ctrvars \triangleq \langle inState, unhandledRequest \rangle$ 
95   $vars \triangleq \langle AccExec, SegStart, LastExecutionTime,$ 
96      $inState, unhandledRequest \rangle$ 

98   $Spec \triangleq \wedge Init$ 
99      $\wedge \square [Next \wedge \neg \text{UNCHANGED } ctrvars]_{vars}$ 

101 Compose the various partial specifications
102  $Component \triangleq$ 
103     $\wedge BasicComponent!Component$ 
104     $\wedge RTnow(vars)$ 
105     $\wedge Spec$ 
106     $\wedge \square (LastExecutionTime \leq ExecutionTime)$ 

108 |-----|
110 |-----|

112  $\_Component(AccExec, SegStart) \triangleq \text{INSTANCE } Inner$ 
113  $Component \triangleq$ 
114     $\exists ae, ss : \_Component(ae, ss)!Component$ 

116 |-----|

```

5.2 Response Time

Based on the `Service` context model, we define the response-time measurement, in a similar fashion to execution time. Here, too we already added the definition of a constraint on response time on Line 70.




Figure 4: Definition of the response-time measurement

Figure 4 shows the corresponding state-machine representation. The new variables and actions introduced by the measurement definition are high-lighted in red.

```

1 |----- MODULE ResponseTimeConstrainedService -----|
2 | EXTENDS RealTime
  |
  | Parameter:
  | ResponseTime – Maximum response time a request should exhibit.
  |
9 | CONSTANT ResponseTime
10 | ASSUME (ResponseTime ∈ Real) ∧ (ResponseTime > 0)
  |
  | Variables:
  | inState      – the current state of the service machinery.
  | unhandledRequest – TRUE indicates the arrival of a new request.
  | LastResponseTime – the response time of the last request serviced.
  |
19 | VARIABLES inState, unhandledRequest
20 | VARIABLE LastResponseTime
  |
22 |-----|
24 |----- MODULE Inner -----|
  | The actual specification.
  |
  | Variables:
  | Start – the start of the last request.
  |
34 | VARIABLE Start
  |
36 | Based on the Service context model
37 | Serv ≜ INSTANCE Service
  |
39 |-----|
41 | Init ≜ ∧ Start = 0
42 |      ∧ LastResponseTime = 0
  |
44 | StartNext reacts to a StartRequest step
45 | StartNext ≜ Serv!StartRequest ⇒
46 |           ∧ Start' = now
47 |           ∧ UNCHANGED LastResponseTime
  |
49 | RespNext reacts to a FinishRequest step
50 | RespNext ≜ Serv!FinishRequest ⇒
  
```

```

51           $\wedge \text{LastResponseTime}' = \text{now} - \text{Start}$ 
52           $\wedge \text{UNCHANGED Start}$ 
54   $\text{ExcludeOtherChange} \triangleq$ 
55     $\neg(\text{Serv!StartRequest} \vee \text{Serv!FinishRequest})$ 
56     $\Rightarrow \text{UNCHANGED} \langle \text{Start}, \text{LastResponseTime} \rangle$ 
58   $\text{Next} \triangleq \text{StartNext} \wedge \text{RespNext} \wedge \text{ExcludeOtherChange}$ 
60   $\text{ctxvars} \triangleq \langle \text{inState}, \text{unhandledRequest} \rangle$ 
61   $\text{vars} \triangleq \langle \text{Start}, \text{LastResponseTime}, \text{inState},$ 
62     $\text{unhandledRequest} \rangle$ 
64   $\text{RespSpec} \triangleq \wedge \text{Init}$ 
65     $\wedge \square[\text{Next} \wedge \neg \text{UNCHANGED ctxvars}]_{\text{vars}}$ 
67   $\text{Service} \triangleq \wedge \text{Serv!Service}$ 
68     $\wedge \text{RTnow}(\text{vars})$ 
69     $\wedge \text{RespSpec}$ 
70     $\wedge \square(\text{LastResponseTime} \leq \text{ResponseTime})$ 
72  ───────────────────────────────────────────────────────────────────────────────────┘
74  ───────────────────────────────────────────────────────────────────────────────────┘
76   $\_Service(\text{Start}) \triangleq \text{INSTANCE Inner}$ 
77   $\text{Service} \triangleq \exists s : \_Service(s)!Service$ 
79  ───────────────────────────────────────────────────────────────────────────────────┘

```

5.3 Inter-Request Time

The following module describes an additional measurement, that we will use to describe environment behaviour in later modules. It measures the time between individual requests for a service sent by the environment. This is later used to define a constraint on the frequency with which the environment sends request for an operation to a given service (see Line 65). The specification is parametrised: The desired minimum time between requests should be passed to the constant *RequestPeriod*. The actual specification is encapsulated in module *Inner* on Lines 63–65. It makes use of the specification of a service from above.

Figure 5 shows the corresponding state-machine representation. The new variables and actions introduced by the measurement definition are high-lighted in red. There is a subtle difference between the simplified state-machine diagram and the actual specification: Because we have separated environment specification and service specification in the service context model (see Sect. 4.1), the TLA⁺ specification actually measures all incoming requests including those arriving during request handling.

```

1  ───────────────────────────────────────────────────────────────────────────────────┘
  Specification of a system environment which sends service request with a certain minimum time between
  individual requests.
  Note that this is not a specification of what we expect from an environment but actually a description of a
  behaviour of one specific system environment. It only becomes a specification of an expectation the way it
  is used in the system specification.
  ───────────────────────────────────────────────────────────────────────────────────┘

```




Figure 5: Definition of the inter-request-time measurement

11 EXTENDS *RealTime*

Parameters:

RequestPeriod – the lower limit for the time between individual requests that should be observed by the environment.

19 CONSTANT *RequestPeriod*

20 ASSUME ($RequestPeriod \in Real$) \wedge ($RequestPeriod > 0$)

Variables:

LastDeltaTime – The amount of time between the last two requests.

inState – Current state of the service invoked.

unhandledRequest – TRUE signals that a new request has been put into the system.

30 VARIABLES *LastDeltaTime*

31 VARIABLES *inState*, *unhandledRequest*

33|-----|
 34|-----| MODULE *Inner* -----|
 The actual specification.

Variables:

StartDelta – Start time of the last request.

44 VARIABLE *StartDelta*

46 *TheService* \triangleq INSTANCE *Service*

48|-----|

50 $vars \triangleq \langle inState, unhandledRequest, LastDeltaTime, StartDelta \rangle$

52 $Init \triangleq \wedge LastDeltaTime = RequestPeriod$

53 $\wedge StartDelta = now$

55 $NewRequest \triangleq TheService!RequestArrival$

56 $\Rightarrow \wedge LastDeltaTime'$

57 $= now - StartDelta$

58 $\wedge StartDelta' = now$

60 $ReqPeriod \triangleq \wedge Init$

61 $\wedge \square [NewRequest]_{vars}$

63 $Service \triangleq \wedge TheService!Service$

```

64          $\wedge ReqPeriod$ 
65          $\wedge \square (LastDeltaTime \ge RequestPeriod)$ 
67     |-----|
69 |-----|
70  $\_Environment(StartDelta) \triangleq INSTANCE Inner$ 
71  $Environment \triangleq \exists sd : \_Environment(sd)!Service$ 
73 |-----|

```

6 Resource specification

The following three specifications deal with the resource CPU. Each of the modules specifies one of the layers of a resource specification (see [3, p. 13]):

1. The *resource-service layer* models the service provided by the resource. Here, the corresponding module models the essential service provided by a CPU: to be available to tasks for a certain time and to be assigned from one task to another, eventually serving all tasks.
2. The *resource-measurement layer* provides measurement definitions that allow quantitative statements to be made about a resource.
3. The *resource-property layer* defines constraints over the measurements defined in the resource-measurement layer. Here, we define a RMS-scheduled CPU and its schedulability criterion.

6.1 Resource-Service Layer

The first module defines what a CPU is: It is a resource that is allocated to tasks one at a time in some fashion. Constant *TaskCount* is used to identify the number of tasks to be scheduled, variable *AssignedTo* indicates the task to which the resource has currently been assigned.

```

1 |-----| MODULE CPUScheduler |-----|
  | A CPU Scheduler allocates the resource CPU to various tasks. We model this through a variable
  | AssignedTo holding in each state the number of the task which has currently been allocated the resource.
7 | EXTENDS Naturals
  |
  | Parameters:
  | TaskCount – the number of tasks which need to share the resource.
14 | CONSTANT TaskCount
15 | ASSUME (TaskCount ∈ Nat) ∧ (TaskCount > 0)
  |
  | Variables:
  | AssignedTo – holds the number of the task currently assigned the resource
22 | VARIABLE AssignedTo
24 | AssignedToType  $\triangleq \{1 .. TaskCount\}$ 
26 |-----|

```

```

28 Initially, an arbitrary task has been assigned the CPU.
29 Init  $\triangleq$  AssignedTo  $\in$  AssignedToType

31 The Switch action reassigns the resource from from to to.
32 Switch(from, to)  $\triangleq$   $\wedge$  AssignedTo = from
33  $\wedge$  AssignedTo' = to

35 The CPU can be switched from any task to any other task.
36 Next  $\triangleq$   $\exists i \in$  AssignedToType :
37  $\exists j \in$  AssignedToType :
38 Switch(i, j)

40 CPUScheduler  $\triangleq$   $\wedge$  Init
41  $\wedge \square[Next]_{AssignedTo}$ 
43

```

6.2 Resource-Measurement Layer

The next specification adds some history-determined variables (quite similar to measurements) that allow to determine for what amount of time each task has been allocated the resource. It is based on the previous specificatio, which it imports on Line 38. In addition to the *TaskCount* parameter, it introduces the parameter *Periods* storing the requested period length per task, so that times can be determined per period. Lines 175–177, finally, provide a boolean measurement formalising the situation where all tasks get a sufficiently large share of the resource. To this end, an additional parameter *Wcets* is introduced. This parameter captures the requested amount of time per period for each task.

The newly defined variables are all array variables. We, therefore, need to use TLA⁺'s syntax for array definition and update:

- $[k \in K \mapsto e(k)]$ represents an array that is defined for all $k \in K$. The value associated to a specific k is defined by $e(k)$.
- $A[k]$ represents the value associated with k in array A .
- $[A \text{ EXCEPT } ![k] = e]$ represents an array that is identical to array A except that value k is mapped to the result of expression e . e may use the special identifier $@$, which stands for $A[k]$.

```

1 |----- MODULE TimedCPUScheduler -----|
  | A CPU scheduler for which the time each task is assigned can be measured.
  | The corresponding formulae are derived by conjoining history variables to the CPU scheduler specification.
8 EXTENDS RealTime

  | Parameters:
  | TaskCount – the number of tasks which need to share the resource.
  | Periods – the periods of each task This is an array with one entry per task.
  | Wcets – the worst case execution times of the tasks to be scheduled. This is an array with one entry per
  | task.
19 CONSTANT TaskCount
20 ASSUME (TaskCount  $\in$  Nat)  $\wedge$  (TaskCount > 0)

```

```

22 CONSTANT Periods
23 ASSUME Periods ∈ [{1 .. TaskCount} → Real]

25 CONSTANT Wcets
26 ASSUME Wcets ∈ [{1 .. TaskCount} → Real]

```

Variables:

MinExecTime – records for each task the minimum amount of execution time per period it has been allocated over all periods so far.
AssignedTo – holds the number of the task currently assigned the resource

```

35 VARIABLE MinExecTime
36 VARIABLE AssignedTo

38 CPUSched ≜ INSTANCE CPUScheduler

```

```

41 |----- MODULE Inner -----|

```

Inner module with the actual specification. This is done so that we can hide some of the helper variables.

Variables:

ExecTimeStart – Records for each task the time when it last started executing
LastExecTime – Records the last accumulated execution time for each task.
LastPeriodStart – Records for each task when it last started a period.

```

58 VARIABLES ExecTimeStart, LastExecTime
59 VARIABLE LastPeriodStart

```

```

63 A little helper function
64 Min(a, b) ≜ CASE a ≤ b → a
65                   □ a > b → b

```

```

67 Init ≜ ∧ ExecTimeStart =
68         [k ∈ CPUSched!AssignedToType ↦ 0]
69 ∧ LastExecTime =
70   [k ∈ CPUSched!AssignedToType ↦ 0]
71 ∧ IF (TaskCount > 1) THEN
72   We start out with infinity, so that any real
73   execution time will definitely be smaller
74   MinExecTime =
75     [k ∈ CPUSched!AssignedToType
76       ↦ Infinity]
77 ELSE
78   We need to handle this case specially for
79   technical reasons
80   MinExecTime =
81     [k ∈ CPUSched!AssignedToType
82       ↦ Periods[k]]
83 ∧ LastPeriodStart =
84   [k ∈ CPUSched!AssignedToType ↦ 0]

```

Next we define what happens when a *CPUSched!Switch* occurs

```

89 OnSwitch(from, to)  $\triangleq$ 
90   Cumulate the time the CPU was allocated to Task from
91    $\wedge$  LastExecTime' = [LastExecTime EXCEPT
92     ![from] = @ + now
93     - ExecTimeStart[from]]
94   Remember when Task to received the CPU
95    $\wedge$  ExecTimeStart' = [ExecTimeStart EXCEPT
96     ![to] = now]
97    $\wedge$  UNCHANGED  $\langle$ MinExecTime, LastPeriodStart $\rangle$ 

```

The *OSNext* action binds *OnSwitch* to corresponding *Switch* actions

```

102 OSNext  $\triangleq$   $\forall i \in$  CPUSched!AssignedToType :
103    $\forall j \in$  CPUSched!AssignedToType :
104     CPUSched!Switch(i, j)
105      $\Rightarrow$  OnSwitch(i, j)

```

The *ExecTime* action determines the accumulated execution time for task *i* in the next state, but at most to the end of its current period. A helper action used by action *PeriodEnd* below.

```

112 ExecTime(i)  $\triangleq$  LastExecTime[i] +
113   IF (AssignedTo = i) THEN
114     Min(now',
115       LastPeriodStart[i] +
116       Periods[i] -
117       ExecTimeStart[i]
118     ELSE 0

```

The *PeriodEnd* action reacts to the end of a period for task *i*

```

123 PeriodEnd(i)  $\triangleq$ 
124    $\wedge$  A period is going to end
125   (now' - LastPeriodStart[i])  $\geq$  Periods[i]
126    $\wedge$  The following is the measurement we are really
127   interested in
128   MinExecTime' = [MinExecTime EXCEPT
129     ![i] = Min(@, ExecTime(i))]
130    $\wedge$  But we also need to perform some cleanup to prepare for
131   the next period
132   LastPeriodStart' = [LastPeriodStart EXCEPT
133     ![i] = @ + Periods[i]]
134    $\wedge$  ExecTimeStart' = [ExecTimeStart EXCEPT
135     ![i] = LastPeriodStart'[i]]
136    $\wedge$  LastExecTime' = [LastExecTime EXCEPT ![i] = 0]

```

CheckPeriods catches all period ends of all tasks

```

141 CheckPeriods  $\triangleq$ 
142   IF (TaskCount > 1) THEN
143      $\forall k \in$  CPUSched!AssignedToType : PeriodEnd(k)
144   ELSE
145     If there's only one process it will be allowed to run
146     for the whole period

```



```

147      $MinExecTime'[1] = Periods[1]$ 
149      $Next \triangleq OSNext$ 
151      $vars \triangleq \langle AssignedTo, ExecTimeStart, LastExecTime \rangle$ 
153      $timeVars \triangleq \langle LastPeriodStart, MinExecTime, now \rangle$ 
155      $TimingSpecification \triangleq \wedge RTnow(vars)$ 
156          $\wedge Init$ 
157          $\wedge \square[Next]_{vars}$ 
158          $\wedge \square[CheckPeriods]_{timeVars}$ 
160      $TimedCPUScheduler \triangleq \wedge CPUSched!CPUScheduler$ 
161          $\wedge TimingSpecification$ 
163 |-----|
165 |-----|
167  $\_TimedCPUScheduler(ExecTimeStart, LastExecTime,$ 
168      $LastPeriodStart)$ 
169      $\triangleq INSTANCE Inner$ 
170  $TimedCPUScheduler$ 
171      $\triangleq \exists ets, let, lps :$ 
172      $\_TimedCPUScheduler(ets, let, lps)$ 
173      $!TimedCPUScheduler$ 
175  $ExecutionTimesOk \triangleq$ 
176      $\forall k \in CPUSched!AssignedToType :$ 
177      $(MinExecTime[k] \geq Wcets[k])$ 
179 |-----|

```

6.3 Resource-Property Layer

Finally, *RMSScheduler* below defines an actual CPU which is scheduled using RMS. The main contribution of this specification is the schedulability criterion defined on Lines 59–64. This is the standard RMS schedulability criterion.

```

1 |-----| MODULE RMSScheduler |-----|
  | A CPU Scheduler using RMS. |
5 |-----|
  EXTENDS Reals
  Parameters:
  TaskCount – the number of tasks to be scheduled on the CPU.
  Periods – the periods to be scheduled for these tasks. This is an array with one entry per task.
  Wcets – the worst case execution times of the tasks to be scheduled. This is an array with one entry per
  task.
16 CONSTANT TaskCount
17 ASSUME ( $TaskCount \in Nat$ )  $\wedge$  ( $TaskCount > 0$ )
19 CONSTANT Periods
20 ASSUME  $Periods \in \{1 .. TaskCount\} \rightarrow Real$ 

```

```

22 CONSTANT Wcets
23 ASSUME  $Wcets \in \{1 \dots TaskCount\} \rightarrow Real$ 

Variables:
MinExecTime – records for each task the minimum amount of execution time it has been allocated over
all periods so far.
AssignedTo – holds the number of the task currently assigned the resource
now – the current time.

33 VARIABLE MinExecTime
34 VARIABLE AssignedTo
35 VARIABLE now

37 TimedCPUSched  $\triangleq$  INSTANCE TimedCPUScheduler

39 |-----|
40 A few helpers

42 both root of a
43  $sqrt(b, a) \triangleq a^{(1/b)}$ 

45 Sum of all the elements in an array (function)
46 Copied from Bags.tla
47  $Sum(f) \triangleq$ 
48 LET  $DSum[S \in SUBSET DOMAIN f] \triangleq$ 
49 LET  $elt \triangleq CHOOSE e \in S : TRUE$ 
50 IN IF  $S = \{\}$ 
51 THEN 0
52 ELSE  $f[elt] + DSum[S \setminus \{elt\}]$ 
53 IN  $DSum[DOMAIN f]$ 

55 |-----|
Schedulable is TRUE if the given task load can be scheduled using RMS.

59 Schedulable  $\triangleq$ 
60 LET  $usage \triangleq [k \in \{1 \dots TaskCount\}$ 
61  $\mapsto (Wcets[k]/Periods[k])]$ 
62 IN
63  $Sum(usage) \leq (TaskCount * (sqrt(TaskCount, 2)$ 
64  $- 1))$ 

The actual specification: A TimedCPUScheduler which will meet all deadlines provided the RMS
schedulability is met by the tasks to be scheduled.

70 RMSScheduler  $\triangleq$ 
71  $\wedge TimedCPUSched!TimedCPUScheduler$ 
72  $\wedge \square Schedulable$ 
73  $\stackrel{\pm}{\triangleright} \square TimedCPUSched!ExecutionTimesOk$ 

75 |-----|

```

7 Container Strategy Specification

Resource allocations, intrinsic component properties, and extrinsic service properties must be related by a container strategy. The following module defines such a container

strategy. It is structured into four major parts:

1. Import of measurements and abstract resource specifications required. This is on Lines 49–176.
2. Definition of container expectations. This is on Lines 179–198.
3. Definition of services guaranteed by the container. This is on Lines 201–215.
4. The actual container strategy specification. This is on Lines 217–218.

The container strategy is parametrised by the response time it should provide and the worst-case execution time it can expect. To ensure that only sensible parameter values are provided, a sanity check is performed on Line 180.

We require container strategies to be *functionality preserving*; that is, the functionality offered by the service should be the same as the functionality provided by the underlying component. This is formally expressed on Line 194. Notice, that the predicates used to express component and service behaviour are left open as parameters to the specification by defining them as abstract predicates on Lines 102–110 and 168–176. This way, the specification will be applicable to arbitrary concrete components and services. All that needs to be done is to associate these abstract predicates with concrete predicates when instantiating the container-strategy module.

```

1 |----- MODULE SimpleContainer -----|
  | A container specification for a very simple container. This container manages just one component instance |
  | and tries to achieve a certain response time with it. |
7 | EXTENDS RealTime |
  | Parameters: |
  | ResponseTime – the response time the container should achieve. |
  | ExecutionTime – the execution time of the component available. |
15 | CONSTANT ResponseTime |
16 | ASSUME (ResponseTime ∈ Real) ∧ (ResponseTime > 0) |
18 | CONSTANT ExecutionTime |
19 | ASSUME |
20 | (ExecutionTime ∈ Real) ∧ (ExecutionTime > 0) |
  | Variables: |
  | TaskCount – the number of tasks the container would want to execute on the CPU. |
  | Periods – the periods the container associates with these tasks. |
  | Wcets – the worst case execution times the container associates with these tasks. |
31 | VARIABLES TaskCount, Periods, Wcets |
33 |-----|
  | Specification of required CPU scheduling behaviour. Note that this does not make any statement about the |
  | actual scheduling regime, but only states what tasks need to be scheduled. |
  | Variables: |
  | CPUMinExecTime – records for each task the minimum amount of execution time it has been allocated |
  | over all periods so far. |
  | CPUAssignedTo – holds the number of the task currently assigned the resource. |

```

49 VARIABLES *CPUMinExecTime*, *CPUAssignedTo*

51 *_SomeCPUScheduler*(*TaskCountConstraint*,
52 *PeriodsConstraint*,
53 *WcetsConstraint*)
54 \triangleq INSTANCE *TimedCPUScheduler*
55 WITH *MinExecTime* \leftarrow *CPUMinExecTime*,
56 *AssignedTo* \leftarrow *CPUAssignedTo*,
57 *TaskCount* \leftarrow *TaskCountConstraint*,
58 *Periods* \leftarrow *PeriodsConstraint*,
59 *Wcets* \leftarrow *WcetsConstraint*
60 *CPUCanSchedule*(*TaskCountConstraint*,
61 *PeriodsConstraint*,
62 *WcetsConstraint*)
63 \triangleq \wedge *_SomeCPUScheduler*(*TaskCount*,
64 *Periods*,
65 *Wcets*)
66 *!TimedCPUScheduler*
67 \wedge \square *_SomeCPUScheduler*(*TaskCount*,
68 *Periods*,
69 *Wcets*)
70 *!ExecutionTimesOk*

72

Specification of required component behaviour.

Variables:

CmpInState – the state in which the component currently is.
CmpUnhandledRequest – TRUE if the environment put another request into the system.
CmpLastExecutionTime – the execution time of the last service execution.

85 VARIABLES *CmpInState*, *CmpUnhandledRequest*
86 VARIABLE *CmpLastExecutionTime*

88 *_Component*(*ExecutionTimeConstraint*)
89 \triangleq INSTANCE *ExecTimeConstrainedComponent*
90 WITH
91 *ExecutionTime* \leftarrow *ExecutionTimeConstraint*,
92 *inState* \leftarrow *CmpInState*,
93 *unhandledRequest* \leftarrow *CmpUnhandledRequest*,
94 *LastExecutionTime* \leftarrow *CmpLastExecutionTime*
95 *ComponentMaxExecTime*(*ExecutionTimeConstraint*)
96 \triangleq *_Component*(*ExecutionTimeConstraint*)
97 *!Component*

This predicate represents the functionality of the component.

102 CONSTANT *CompFun*
103 ASSUME *CompFun* \in BOOLEAN

This predicate represents the mapping between functionality and context model of the component.

109 CONSTANT *CompModelMapping*
110 ASSUME *CompModelMapping* \in BOOLEAN

112 | Specification of required request interarrival time.

Variables:

EnvLastDeltaTime – The amount of time between the last two requests.
EnvInState – Current state of the service invoked.
EnvUnhandledRequest – TRUE signals that a new request has been put into the system.

125 VARIABLES *EnvLastDeltaTime*, *EnvInState*
126 VARIABLE *EnvUnhandledRequest*

128 *_MinInterrequestTime*(*RequestPeriodConstraint*)
129 \triangleq INSTANCE *MaxRequPeriodEnv*
130 WITH
131 *RequestPeriod* \leftarrow *RequestPeriodConstraint*,
132 *LastDeltaTime* \leftarrow *EnvLastDeltaTime*,
133 *inState* \leftarrow *EnvInState*,
134 *unhandledRequest* \leftarrow *EnvUnhandledRequest*
135 *MinInterrequestTime*(*RequestPeriodConstraint*)
136 \triangleq *_MinInterrequestTime*(*RequestPeriodConstraint*)
137 *!Environment*

139 | Specification of guaranteed service behaviour.

Variables:

ServLastResponseTime – the response time of the last request serviced.
ServInState – the current state of the service machinery.
ServUnhandledRequest – TRUE indicates the arrival of a new request.

151 VARIABLES *ServLastResponseTime*, *ServInState*
152 VARIABLE *ServUnhandledRequest*

154 *_ServiceResponseTime*(*ResponseTimeConstraint*)
155 \triangleq INSTANCE *ResponseTimeConstrainedService*
156 WITH
157 *ResponseTime* \leftarrow *ResponseTimeConstraint*,
158 *LastResponseTime* \leftarrow *ServLastResponseTime*,
159 *inState* \leftarrow *ServInState*,
160 *unhandledRequest* \leftarrow *ServUnhandledRequest*
161 *ServiceResponseTime*(*ResponseTimeConstraint*)
162 \triangleq *_ServiceResponseTime*(*ResponseTimeConstraint*)
163 *!Service*

This predicate represents the functionality of the service.

168 CONSTANT *ServFun*
169 ASSUME *ServFun* \in BOOLEAN

This predicate represents the mapping between functionality and context model of the service.

175 CONSTANT *ServModelMapping*
176 ASSUME *ServModelMapping* \in BOOLEAN

178 |
179 *ContainerPreCond* \triangleq

```

180  $\wedge ExecutionTime \leq ResponseTime$ 
181  $\wedge$  The CPU must be able to schedule exactly one task with a
182 period equal to the requested response time and a wcet
183 equal to the specified execution time of the available
184 component.
185  $\wedge CPUCanSchedule(1,$ 
186  $\quad [n \in \{1\} \mapsto ResponseTime],$ 
187  $\quad [n \in \{1\} \mapsto ExecutionTime])$ 
188  $\wedge$  A component with the required execution time is available.
189  $\wedge ComponentMaxExecTime(ExecutionTime)$ 
190  $\wedge CompFun$ 
191  $\wedge CompModelMapping$ 
192  $\wedge$  The component functionality implements the service
193 functionality.
194  $CompFun \Rightarrow ServFun$ 
195  $\wedge$  Requests arrive with a constant period, the length of
196 which is somehow related to the period length requested
197 from the CPU.
198  $\wedge MinInterrequestTime(ResponseTime)$ 

201  $ContainerPostCond \triangleq$ 
202  $\wedge$  The promised response time can be guaranteed
203  $\wedge ServiceResponseTime(ResponseTime)$ 
204  $\wedge ServFun$ 
205  $\wedge ServModelMapping$ 
206  $\wedge$  The container will allocate exactly one task for the
207 component.
208  $\square \wedge TaskCount = 1$ 
209  $\quad \wedge Periods = [n \in \{1\} \mapsto ResponseTime]$ 
210  $\quad \wedge Wcets = [n \in \{1\} \mapsto ExecutionTime]$ 
211  $\wedge$  State that the container will hand requests directly
212 to the component, without buffering them in any way. If
213 the container provides buffering, this would need to go
214 here
215  $\square(CmpUnhandledRequest = EnvUnhandledRequest)$ 

217  $Container \triangleq$ 
218  $ContainerPreCond \overset{\pm}{\triangleright} ContainerPostCond$ 
219
```

8 The Counter Application

So far, we have been discussing the non-functional properties in the abstract. In the following specifications, we define a sample Counter application and provide model mappings to apply our measurements to this application.

8.1 Application Model

The next two modules define the Counter application itself. We begin with the definition of its interface. Notice that this is just a helper module that we will later use to hide the actual implementation of the Counter application. The interface module uses abstract actions to define the interactions with the environment that can be observed of a Counter application without defining how they are realised. An abstract action is defined by a *Boolean* constant, possibly with open parameter slots (indicated by $_$). More details on abstract actions can be found in [1].

```

1 |----- MODULE CounterInterface -----|
   |
   | A global representation of the counter's state. We do not say anything about what this state looks like.
   |
7 | VARIABLE counterState
   |
   | A DoInc (counterState, counterState') step represents an incoming request to increment the internal
   | counter of the component
   |
13 | CONSTANT DoInc(., .)
   |
   | A GetData (counterState, counterState') step represents an incoming request for the current value.
   |
19 | CONSTANT GetData(., .)
   |
   | A SendData (value, counterState, counterState') step represents a response to a GetData step.
   |
25 | CONSTANT SendData(., ., .)
   |
27 | CONSTANT InitialCounterStates
   |
29 | ASSUME  $\forall v, csOld, csNew :$ 
30 |      $\wedge DoInc (csOld, csNew) \in \text{BOOLEAN}$ 
31 |      $\wedge GetData (csOld, csNew) \in \text{BOOLEAN}$ 
32 |      $\wedge SendData(v, csOld, csNew) \in \text{BOOLEAN}$ 
   |
34 |-----|

```

The next module defines the actual Counter implementing this interface. This is the application model of our example. It is a normal TLA^+ specification. However, note how it binds the Counter interface from the previous module by referencing the abstract actions on Lines 14, 20, and 33.

```

1 |----- MODULE CounterApp -----|
   |
3 | EXTENDS CounterInterface, Naturals
   |
5 | Internal variables:
6 | VARIABLE internalCounter
7 | VARIABLE doHandle
   |
9 | Init  $\triangleq \wedge internalCounter = 0$ 
10 |      $\wedge doHandle = 0$ 
11 |      $\wedge counterState \in InitialCounterStates$ 
   |
14 | IncrementReq  $\triangleq \wedge DoInc(counterState, counterState')$ 
15 |      $\wedge doHandle = 0$ 
   |

```

```

16           $\wedge$  internalCounter'
17             = internalCounter + 1
18           $\wedge$  UNCHANGED doHandle
20 ReceiveGetData  $\triangleq$   $\wedge$  GetData(counterState,
21                               counterState')
22           $\wedge$  doHandle = 0
23           $\wedge$  doHandle' = 1
24           $\wedge$  UNCHANGED internalCounter
26 HandleGetData  $\triangleq$   $\wedge$  doHandle = 1
27           $\wedge$  doHandle' = 2
28           $\wedge$  UNCHANGED  $\langle$ internalCounter,
29                               counterState $\rangle$ 
31 ReplyStep  $\triangleq$   $\wedge$  doHandle = 2
32           $\wedge$  doHandle' = 0
33           $\wedge$  SendData(internalCounter,
34                               counterState,
35                               counterState')
36           $\wedge$  UNCHANGED internalCounter
38 Next  $\triangleq$   $\vee$  IncrementReq
39           $\vee$  ReceiveGetData  $\vee$  HandleGetData
40           $\vee$  ReplyStep
42 vars  $\triangleq$   $\langle$ counterState, internalCounter, doHandle $\rangle$ 
44 Spec  $\triangleq$   $\wedge$  Init
45           $\wedge$  [Next]vars
47 |

```

8.2 Model Mappings

The following two specifications define the model mappings for execution time of the `Counter` component and for response time of the `Counter` service, resp. Both specifications work in a similar manner: They extend the `CounterApp` specification, so that all specifications and variables from that specification are directly available. Then, they import the measurement specification. Eventually, they define the *Model-Mapping* formula, the actual model-mapping relation ϕ_{App}^{Ctx} by relating states of the `Counter` application to states of the context model. Lines 55–57 in Module `CounterAppExecTime` and Lines 49–51 in Module `CounterAppResponseTime` finally encode the model mapping as given by Equation (2) in the main paper:

$$\Pi_{Ctx}^{App} \triangleq \Pi_{App} \wedge \Pi_{Ctx} \wedge \square (\langle v_{App}, v_{Ctx} \rangle \in \phi_{App}^{Ctx})$$

```

1 |----- MODULE CounterAppExecTime -----|

```

```

  A module defining execution time of the GetData() operation.

```

```

5 EXTENDS CounterApp, Realtime

```


Variables:

inState – the state in which the component currently is.
unhandledRequest – TRUE if the environment put another request into the system.
LastExecutionTime – the execution time of the last service execution.

```
15 VARIABLE inState
17 VARIABLE unhandledRequest
18 VARIABLE ExecutionTime
20 ExecTimeSpec(ExecutionTimeConstr)
21    $\triangleq$  INSTANCE ExecTimeConstrainedComponent
22     WITH LastExecutionTime  $\leftarrow$  ExecutionTime,
23         ExecutionTime  $\leftarrow$  ExecutionTimeConstr
25 CompSpec  $\triangleq$  ExecTimeSpec(20)!Component
27 |-----|
```

Definition of the context-model-application-model mapping

Note how this maps the *GetData/SendData* operation, but not *DoInc*.

```
34 ModelMapping  $\triangleq$ 
35    $\wedge$  doHandle = 0  $\Rightarrow$ 
36      $\wedge$  inState = "Idle"
37      $\wedge$  unhandledRequest = FALSE
38    $\wedge$  doHandle = 1  $\Rightarrow$ 
39      $\wedge$  inState = "Idle"
40      $\wedge$  unhandledRequest = TRUE
41    $\wedge$  doHandle = 2  $\Rightarrow$ 
42      $\wedge$  inState  $\in$  {"HandlingRequest",
43                   "Blocked"}
44      $\wedge$  unhandledRequest = FALSE
45   Dummy mapping for completeness' sake
46    $\wedge$  (doHandle  $\notin$  {0, 1, 2})  $\Rightarrow$ 
47      $\wedge$  inState = "Idle"
48      $\wedge$  unhandledRequest = FALSE
50 |-----|
```

Final model of the counter component.

```
55 CounterComponent  $\triangleq$   $\wedge$  Spec
56      $\wedge$  CompSpec
57      $\wedge$   $\square$  ModelMapping
59 |-----|
```

```

1 |----- MODULE CounterAppResponseTime -----|
  | A module defining response time of the GetData() operation. |
5 | EXTENDS CounterApp, Realtime |
  | Variables: |
  | ResponseTime - the response time of the last request serviced. |
  | inState      - the current state of the service machinery. |
  | unhandledRequest - TRUE indicates the arrival of a new request. |
14 | VARIABLES ResponseTime, inState, unhandledRequest |
16 | ResponseTimeSpec(ResponseTimeConstr) |
17 |   ≜ INSTANCE ResponseTimeConstrainedService |
18 |     WITH LastResponseTime ← ResponseTime, |
19 |        ResponseTime ← ResponseTimeConstr |
21 | ServSpec ≜ ResponseTimeSpec(50)!Service |
23 |-----|
  | Definition of the context-model-application-model mapping |
  | Note how this maps the GetData/SendData operation, but not DoInc. |
30 | ModelMapping ≜ |
31 |   ∧ doHandle = 0 ⇒ |
32 |     ∧ inState = "Idle" |
33 |     ∧ unhandledRequest = FALSE |
34 |   ∧ doHandle = 1 ⇒ |
35 |     ∧ inState = "Idle" |
36 |     ∧ unhandledRequest = TRUE |
37 |   ∧ doHandle = 2 ⇒ |
38 |     ∧ inState = "HandlingRequest" |
39 |     ∧ unhandledRequest = FALSE |
40 |   ∧ (doHandle ∉ {0, 1, 2}) ⇒ |
41 |     ∧ inState = "Idle" |
42 |     ∧ unhandledRequest = FALSE |
44 |-----|
  | Final model of the counter service . |
49 | CounterService ≜ ∧ Spec |
50 |                   ∧ ServSpec |
51 |                   ∧ □ ModelMapping |
53 |-----|

```

9 System Specification

Finally, we are ready to pull everything together. This we do in the system specification. The important bit is on Lines 237–264, where the system specification is composed from the individual elementary specifications. Everything before that is mainly of technical relevance, importing the previous specifications.

The actual connections between the component, the resource, the container, and the service are expressed by means of shared flexible variables. This can be seen in two ways in the specification: 1) on Lines 239–245 we explicitly pass parameters that perform part of the connection between container and resource and between component and resource; 2) on Lines 246–264 we use explicit constraints to relate other variables, relating the rest of the system parts to each other.

The complete system composition is then defined by formula *System* on Lines 237–264. Formula *ExternalService* on Lines 269–270 defines the service we expect the system to provide. Notice that this is conditional based on environment behaviour. Lines 278 and 279, finally, define what it means for the system to be feasible. This is the property we need to prove to show that we have indeed specified a feasible system. As explained in the main paper, we can make use of Abadi/Lamport’s composition theorem for this proof.

```

1 |----- MODULE SystemSpecification -----|
  | A sample system specification.
  | The system contains one counter with an execution time of 20 milliseconds, a RMS
  | scheduled CPU, and a simple container.
8 | EXTENDS Reals, CounterInterface
  |
  | Parameters:
  | RequestPeriod – Part of an environment assertion: The environment promises to
  | send requests with a minimum distance of RequestPeriod milliseconds.
17 | CONSTANT RequestPeriod
18 | ASSUME (RequestPeriod ∈ Real) ∧ (RequestPeriod > 0)
  |
  | Variables:
  | now – the current time.
25 | VARIABLE now
  |
27 |-----|
  | The counter component. The only intrinsic property offered by this component
  | is its execution time, which is always less than 20ms.
  |
  | Variables:
  | MyCompExec – The last execution time of a service request handled by
  | MyComponent.
  | MyCompInState – The current state of component MyComponent
  | MyCompUnhandledRequest – Set to TRUE to send a request to MyComponent.
41 | VARIABLES MyCompExec, MyCompInState
42 | VARIABLE MyCompUnhandledRequest
43 | VARIABLES MyInternalCounter, MyDoHandle
  |
45 | _MyComponent ≜ INSTANCE CounterAppExecTime
46 | WITH
47 |   ExecutionTime ← MyCompExec,
48 |   inState ← MyCompInState,
49 |   unhandledRequest ← MyCompUnhandledRequest,
50 |   internalCounter ← MyInternalCounter,
51 |   doHandle ← MyDoHandle
  |
53 | The actual component specification.
54 | MyComponent ≜ _MyComponent!CounterComponent

```

```

55 CompMap  $\triangleq$   $\square$  _MyComponent!ModelMapping
57 _MyCompFunc  $\triangleq$  INSTANCE CounterApp WITH
58   internalCounter  $\leftarrow$  MyInternalCounter,
59   doHandle  $\leftarrow$  MyDoHandle
60 MyCompFunc  $\triangleq$  _MyCompFunc!Spec

```

62

A *CPU*. The parameters of the specification can be used to indicate the number of tasks to be scheduled, their respective periods as well as their respective worst case execution times.

Variables:

MYCPU_MinExecTime – records for each task the minimum amount of execution time it has been allocated over all periods so far.

MYCPU_AssignedTo – holds the number of the task currently assigned the resource

```

78 VARIABLE MYCPU_MinExecTime
79 VARIABLE MYCPU_AssignedTo

81 _MyCPU(TaskCount, Periods, Wcets)
82    $\triangleq$  INSTANCE RMSScheduler WITH
83     MinExecTime  $\leftarrow$  MYCPU_MinExecTime,
84     AssignedTo  $\leftarrow$  MYCPU_AssignedTo
85 MyCPU(TaskCount, Periods, Wcets)
86    $\triangleq$  _MyCPU(TaskCount, Periods, Wcets)
87   !RMSScheduler

```

89

Environment specification.

Variables:

EnvLastDeltaTime – The amount of time between the last two requests.

EnvInState – Current state of the service invoked.

EnvUnhandledRequest – TRUE signals that a new request has been put into the system.

```

103 VARIABLES EnvLastDeltaTime, EnvInState
104 VARIABLE EnvUnhandledRequest

106 _Environment(RequestPeriodConstraint)
107    $\triangleq$  INSTANCE MaxRequPeriodEnv WITH
108     RequestPeriod  $\leftarrow$  RequestPeriodConstraint,
109     LastDeltaTime  $\leftarrow$  EnvLastDeltaTime,
110     inState  $\leftarrow$  EnvInState,
111     unhandledRequest  $\leftarrow$  EnvUnhandledRequest
112 Environment(RequestPeriodConstraint)
113    $\triangleq$  _Environment(RequestPeriodConstraint)
114   !Environment

```

116

The service the system is to perform.

Variables:

ServResponseTime – the response time of the last request serviced.
ServInState – the current state of the service machinery.
ServUnhandledRequest – TRUE indicates the arrival of a new request.

129 VARIABLES *ServResponseTime*, *ServInState*
130 VARIABLE *ServUnhandledRequest*
131 VARIABLES *ServInternalCounter*, *ServDoHandle*

133 *_Service*

134 \triangleq INSTANCE *CounterAppResponseTime* WITH
135 *ResponseTime* \leftarrow *ServResponseTime*,
136 *inState* \leftarrow *ServInState*,
137 *unhandledRequest* \leftarrow *ServUnhandledRequest*,
138 *internalCounter* \leftarrow *ServInternalCounter*,
139 *doHandle* \leftarrow *ServDoHandle*
140 *Service* \triangleq *_Service!CounterService*
141 *ServMap* \triangleq \square *_Service!ModelMapping*

143 *_MyServFunc*

144 \triangleq INSTANCE *CounterApp* WITH
145 *internalCounter* \leftarrow *ServInternalCounter*,
146 *doHandle* \leftarrow *ServDoHandle*
147 *MyServFunc* \triangleq *_MyServFunc!Spec*

149

Container specification.

Variables:

SCCPUMinExecTime – records for each task the minimum amount of execution time it has been allocated over all periods so far.

SCCPUAssignedTo – holds the number of the task currently assigned the resource.

SCCmpInState – the state in which the component currently is.

SCCmpUnhandledRequest – TRUE if the environment put another request into the system.

SCCmpLastExecutionTime – the execution time of the last service execution.

SCEnvLastDeltaTime – The amount of time between the last two requests.

SCEnvInState – Current state of the service invoked.

SCEnvUnhandledRequest – TRUE signals that a new request has been put into the system.

SCServLastResponseTime – the response time of the last request serviced.

SCServInState – the current state of the service machinery.

SCServUnhandledRequest – TRUE indicates the arrival of a new request.

177 VARIABLES *SCCPUMinExecTime*, *SCCPUAssignedTo*
178 VARIABLES *SCCmpInState*, *SCCmpUnhandledRequest*
179 VARIABLES *SCCmpLastExecutionTime*, *SCEnvInState*,
180 VARIABLE *SCEnvLastDeltaTime*
181 VARIABLE *SCEnvUnhandledRequest*
182 VARIABLES *SCServLastResponseTime*, *SCServInState*,
183 VARIABLE *SCServUnhandledRequest*

185 *_MyContainer*(*ExecutionTimeConstr*,
186 *ResponseTimeConstr*,
187 *TaskCount*, *Periods*,

```

188            $Wcets) \triangleq$ 
189   INSTANCE SimpleContainer
190     WITH
191     ExecutionTime  $\leftarrow$  ExecutionTimeConstr,
192     ResponseTime  $\leftarrow$  ResponseTimeConstr,
193     CPUMinExecTime  $\leftarrow$  SCCPUMinExecTime,
194     CPUAssignedTo  $\leftarrow$  SCCPUAssignedTo,
195     CmpInState  $\leftarrow$  SCCmpInState,
196     CmpUnhandledRequest  $\leftarrow$  SCCmpUnhandledRequest,
197     CmpLastExecutionTime  $\leftarrow$  SCCmpLastExecutionTime,
198     EnvLastDeltaTime  $\leftarrow$  SCEnvLastDeltaTime,
199     EnvInState  $\leftarrow$  SCEnvInState,
200     EnvUnhandledRequest  $\leftarrow$  SCEnvUnhandledRequest,
201     ServLastResponseTime  $\leftarrow$  SCServLastResponseTime,
202     ServInState  $\leftarrow$  SCServInState,
203     ServUnhandledRequest  $\leftarrow$  SCServUnhandledRequest,
204     CompFun  $\leftarrow$  MyCompFunc,
205     CompModelMapping  $\leftarrow$  CompMap,
206     ServFun  $\leftarrow$  MyServFunc,
207     ServModelMapping  $\leftarrow$  ServMap
209   MyContainer(ExecutionTimeConstr,
210               ResponseTimeConstr,
211               TaskCount, Periods,
212               Wcets)
213    $\triangleq$  _MyContainer(ExecutionTimeConstr,
214                     ResponseTimeConstr,
215                     TaskCount, Periods,
216                     Wcets)! Container

```

218 |-----|

The complete system.

Variables:

TaskCount – the number of tasks to be scheduled on the *CPU* as determined by the container.

Periods – the periods to be scheduled for those tasks as determined by container.

Wcets – the worst case execution times to be considered when scheduling. As determined by the container.

```

234 VARIABLES CPUTaskCount, CPUPeriods, CPUWcets
235 VARIABLES SCTaskCount, SCPeriods, SCWcets
237 System  $\triangleq$ 
238    $\wedge$  MyComponent
239    $\wedge$  MyCPU(CPUTaskCount,
240             CPUPeriods,
241             CPUWcets)
242    $\wedge$  MyContainer(20, 50,
243                 SCTaskCount,
244                 SCPeriods,
245                 SCWcets)

```

```

246   $\wedge \square \wedge \text{ServResponseTime} =$ 
247       $\text{SCServLastResponseTime}$ 
248       $\wedge \text{ServInState} = \text{SCServInState}$ 
249       $\wedge \text{ServUnhandledRequest} =$ 
250           $\text{SCServUnhandledRequest}$ 
251   $\wedge \square \wedge \text{MYCPU\_MinExecTime} =$ 
252       $\text{SCCPUMinExecTime}$ 
253       $\wedge \text{MYCPU\_AssignedTo} = \text{SCCPUAssignedTo}$ 
254       $\wedge \text{CPUTaskCount} = \text{SCTaskCount}$ 
255       $\wedge \text{CPUPeriods} = \text{SCPeriods}$ 
256       $\wedge \text{CPUWcets} = \text{SCWcets}$ 
257   $\wedge \square \wedge \text{SCCmpLastExecutionTime} = \text{MyCompExec}$ 
258       $\wedge \text{SCCmpInState} = \text{MyCompInState}$ 
259       $\wedge \text{SCCmpUnhandledRequest} =$ 
260           $\text{MyCompUnhandledRequest}$ 
261   $\wedge \square \wedge \text{SCEnvLastDeltaTime} = \text{EnvLastDeltaTime}$ 
262       $\wedge \text{SCEnvInState} = \text{EnvInState}$ 
263       $\wedge \text{SCEnvUnhandledRequest} =$ 
264           $\text{EnvUnhandledRequest}$ 

```

The external behaviour we require of the system.

```

269  ExternalService
270       $\triangleq \text{Environment}(\text{RequestPeriod}) \dashv\triangleright \text{Service}$ 

```

272

This is the property we need to prove to ensure that we have a feasible system.

```

278  IsFeasible
279       $\triangleq \text{System} \Rightarrow \text{ExternalService}$ 

```

281

References

- [1] Leslie Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002.
- [2] Steffen Zschaler. *A Semantic Framework for Non-functional Specifications of Component-Based Systems*. PhD thesis, Technische Universität Dresden, Germany, April 2007.
- [3] Steffen Zschaler. Formal specification of non-functional properties of component-based software systems: A semantic framework and some applications thereof. *Software and Systems Modelling (SoSyM)*, 2008. To appear.