



CISTER

Research Center in
Real-Time & Embedded
Computing Systems

Technical Report

Logic-based Schedulability Analysis for Compositional Hard Real-Time Embedded Systems

André Pedro

David Pereira

Luis Miguel Pinho

Jorge Sousa Pinto

CISTER-TR-131201

Version:

Date: 12/6/2013

Logic-based Schedulability Analysis for Compositional Hard Real-Time Embedded Systems

André Pedro, David Pereira, Luis Miguel Pinho, Jorge Sousa Pinto

CISTER Research Unit

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8340509

E-mail: anmap@isep.ipp.pt, dmrpe@isep.ipp.pt, Imp@isep.ipp.pt,

<http://www.cister.isep.ipp.pt>

Abstract

Over the past decades several approaches for schedulability analysis have been proposed for both uni-processor and multi-processor real-time systems. Although different techniques are employed, very little has been put forward in using formal specifications, with the consequent possibility for mis-interpretations or ambiguities in the problem statement. Using a logic based approach to schedulability analysis in the design of hard real-time systems eases the synthesis of correct-by-construction procedures for both static and dynamic verification processes. In this paper we propose a novel approach to schedulability analysis based on a timed temporal logic with time durations. Our approach subsumes classical methods for uni-processor scheduling analysis over compositional resource models by providing the developer with counter-examples, and by ruling out schedules that cause unsafe violations on the system. We also provide an example showing the effectiveness of our proposal.

Logic-based Schedulability Analysis for Compositional Hard Real-Time Embedded Systems

André de Matos Pedro, David Pereira, Luís Miguel Pinho, and Jorge Sousa Pinto*
CISTER/INESC TEC, ISEP, Polytechnic Institute of Porto, Portugal

{anmap, dmrpe, lmp}@isep.ipp.pt

* HASLab/INESC TEC & Universidade do Minho, Portugal
jsp@di.uminho.pt

Abstract—Over the past decades several approaches for schedulability analysis have been proposed for both uni-processor and multi-processor real-time systems. Although different techniques are employed, very little has been put forward in using formal specifications, with the consequent possibility for misinterpretations or ambiguities in the problem statement. Using a logic based approach to schedulability analysis in the design of hard real-time systems eases the synthesis of correct-by-construction procedures for both static and dynamic verification processes. In this paper we propose a novel approach to schedulability analysis based on a timed temporal logic with time durations. Our approach subsumes classical methods for uni-processor scheduling analysis over compositional resource models by providing the developer with counter-examples, and by ruling out schedules that cause unsafe violations on the system. We also provide an example showing the effectiveness of our proposal.

I. INTRODUCTION

Schedulability analysis is a very important part of the research that is carried out in real-time systems. Due to the complex nature of the scenarios that real-time systems face, functional properties must be coupled with a predictable response time, so that the operations are performed safely and within the expected constraints. Relaxing any of these two conditions, in the case of hard real-time systems, might lead to catastrophic events, including the loss of human lives.

Along almost forty years, a bewildering diversity of schedulability tests for hard real-time systems has been proposed to address the constraints imposed by the required predictability. These tests vary considerably in their complexity, expressivity, and target scheduling policies (e.g., fixed task or job priority, preemptive or non-preemptive). The literature [1], [12] reveals that generally schedulability test works by assuming a worst-case scenario and checking that each of the involved task gets a sufficient allocation of shared resources or jobs always complete before their deadlines. Naturally, cases that are not "the worst" will also succeed.

In this paper we consider periodic resource models [23], [24] for the composability of components each one with its own set of real-time tasks, providing a rigorous definition of their timing properties, intending to be able to formally verify their composition. These definitions are established by the language and semantics of timed temporal logic, an approach that is not new in the context of real-time systems verification [5].

In this paper we also consider a variant of MTL- f [16] that is well-suited to analyze sequences and durations of timed executions. This type of analysis is sufficient to solve the schedulability decision problem of periodic resource models, and compositional periodic resource models. The reasons for adopting a logic-based paradigm towards schedulability analysis are: it becomes more comprehensive and expressive; rules out potential specification incoherences typical from informal specifications; and it has some benefits relatively to the available analysis, not in terms of efficiency but in terms of easy extension for monitoring approaches such as the acquisition of the maximum detection delay of a task as in [26]. As further context to the work, we note that:

- 1) the outcome of a classical schedulability analysis is typically a verdict for a certain set of tasks, but no counter-examples are shown if the set of tasks is not schedulable;
- 2) the behavior of the scheduler is assumed rather than being explicitly included in the schedulability test;
- 3) the timing description of the tasks is the unique data provided by classical analysis methods (i.e., offsets, jitters, periods, deadlines);
- 4) standard approaches are not possible to extend with other useful properties such as monitoring and enforcement of real-time properties [22], [21], due to the restricted definition of their sets of tasks (e.g., defining a bound for two consecutive instructions, the inter-arrival time of an event);
- 5) some real-time literature [26], [27] commonly considers the estimation of an arrival rate, which implies minimization and produces significant issues (e.g., under and over estimations, local minimums and maximums, etc.).

Our work intends to integrate the description of the scheduling behavior with the schedulability analysis, which enables to draw counter-examples when the system is not schedulable. These counter-examples are fundamental for the system designer to understand and adapt the design accordingly.

Another key point of our approach is that the rigorous of our definitions enable the integration of formal-verification techniques such as runtime verification or model checking, by subsuming the corresponding computational artifacts in a correct-by-construction way. It opens the possibility of adopting mature and experimental formal verification tools [2], [6], [19], [4] that are already available for the scenarios we intend to certify in future development of our work.

Although this paper’s focus is solely on the schedulability analysis of compositional periodic resource models under the *rate monotonic* (RM) policy, we introduce this work as a foundational approach for schedulability analysis of compositional resource models, on which we intend to use more advanced schedulability policies and principles in the future. Moreover, this research work is part of a long term project whose aim is the development of novel approach for the unified specification of hard real-time systems (functional and non-functional requirements), supported by the combination of off-the-shelf static verification and runtime verification methods.

We provide a fragment of the *metric temporal logic with durations* (MTL- \int) which we named *restricted metric temporal logic with durations* (RMTL- \int), a schedulability analysis for periodic resource models and coupled periodic resource models, as well as the encoding of both models in RMTL- \int . Our encoding allows us to isolate by construction cases where the *worst-case execution time* (WCET) violations are unsafe to the schedulability of the system, and to analyze each component knowing only the high-level specifications (not the internals) of the other components, which excludes the dynamic increment of components. A synthetic workload is also described and its schedulability test exemplified using our analysis. For the best of our knowledge this is the first approach that combines MTL- \int with schedulability analysis.

The paper is organized as follows. Section II introduces research work that relates to the one presented in this paper. Section III introduces the preliminary concepts that are necessary for our schedulability analysis as a background. Section IV describes the syntax and semantics of the MTL- \int logic, including a set of necessary axioms. Section V introduces the new concept of schedulability analysis using MTL- \int and timed execution traces. Section VI exemplifies how to use runtime monitors through a practical application of the method of schedulability analysis that we propose. Finally, Section VII draws some conclusions and points to further work directions.

II. RELATED WORK

So far, not many alternative approaches for schedulability analysis of real-time systems have been proposed nor specific formalisms for tests have emerged. We will describe an alternative schedulability analysis based on *timed automata* [9], [14] and *process algebraic* [20].

A. Automata-based

One interesting research effort that discards the classic schedulability analysis is the one proposed by Fersman et al. [9], [10]. They use *timed automata extended with real-time tasks* to specify the system behavior plus the scheduler behavior. Basically, the authors rewrite any hard real-time system into these type of automaton which is further coupled with another automaton that behaves according to a particular scheduling policy (e.g., RM, or EDF). In this sense, the schedulability test remains a reachability analysis problem which is normally solved by model checker tools such as UPPAAL [2] or NuSMV [6]. In order to check the model, this type of automaton shall be translated into one similar timed automaton.

B. Task Automata

The schedulability of a task automaton is undecidable [11], [15]. Recently, some progress has been made to show that only a small class of task automata is undecidable. Yi et al. [9] proved that task automata for a single-processor system that have preemption, variable task execution time (i.e., the best case execution time is different from the worst-case execution time) and feedback (i.e., the precise finishing time of a task may influence the new task releases) are undecidable. Then, it is shown that (even with only one processor) schedulability becomes undecidable if a task automaton holds these properties. In turn, it was also shown that if there is no preemption, the problem becomes decidable. The same holds if there is no variable execution time. The open question remained, whether schedulability is decidable holding a non-feedback property.

For multi-processor systems [14] the last presented variant has also been proven decidable for certain types of schedulers, but there is no result for the general case and therefore for multi-processor systems. It is shown, in [14], that the schedulability for multi-processor system is possible for non-preemptive schedulers and preemptive scheduler with constant execution time.

C. Other Approaches

Another interesting research effort was the process algebraic approach for schedulability analysis as initially proposed by Ben-Abdallah et al [3]. Philippou et al [20] formalizes the problem of compositional hierarchical scheduling by introducing a process algebraic framework for modeling resource demand and supply. This was inspired in the *timed process algebra*.

Timed petri nets (TPNs) are presented in [25] as a model for schedulability analysis in real-time systems. The authors prove that schedulability analysis is reduced to a state reachability problem. The authors describe that timing and behavioral properties should be formalized in different levels. Zonghua and Shin [13] describes a translation from TPN to timed automata.

III. PRELIMINARIES

In this section we introduce the main concepts that support our proposal.

A. Basic Notions

In the rest of the paper we will assume *tasks sets* $\Gamma = \{\tau_1, \tau_2, \dots, \tau_n\}$, such that $n \in \mathbb{N}^+$ is the number of tasks $\tau_i = (p_i, e_i)$ where p_i and e_i are, respectively, the period and the worst-case execution time of τ_i . Each task $\tau_i \in \tau$ is periodic. A *periodic resource model* $\omega = (\tau, \pi, \theta, rm)$, where $\tau \subseteq \Gamma$, π is the *replenishment period*, θ is the *server budget*, and rm is the RM scheduling algorithm. The set of periodic resource models is denoted by Ω .

The outputs of a resource model ω are *sequences of events*. Considering a par (ω, τ_i) with $\omega \in \Omega$ and $\tau_i \in \tau$, each event can be of one of the following types: a *release-event* $\text{erelease}(\omega, \tau_i)$; a *start-event* $\text{estart}(\omega, \tau_i)$; a *sleep-event*

esleep(ω, τ_i); a *resume-event* eresume(ω, τ_i); or a *stop-event* estop(ω, τ_i). In addition, we assume a parameterized event $\varepsilon(\omega_j, \tau_i, id)$ that denotes the critical events of a task, where id is the event identifier, and $erenewal(\omega)$ denotes the budget release of a resource model. We denote sets of events by \mathcal{E} .

A sequence of events, also known as *execution trace*, is an infinite sequence

$$\rho = (e_1, t_1)(e_2, t_2) \cdots$$

of time-stamped events (e_i, t_i) with $e_i \in \mathcal{E}$ and $t_i \in \mathbb{R}^+$. The sequence satisfies monotonicity and progresses, *i.e.*, $t_i \leq t_{i+1}$ for all $i \in \mathbb{N}^+$, and for all $t \in \mathbb{R}^+$ there is some $i > 0$ such that $t_i > t$, respectively.

B. Schedulability Analysis of Periodic Resource Models

The schedulability analysis for periodic resource models is provided by Shin and Lee [23], [24]. The authors formulate an analysis based on *resource model supply*. The *supply bound function* $sbf_\omega(t)$ is defined to calculate the minimum resource supply for every interval of length t as follows:

$$sbf_\omega(t) = \begin{cases} t - (k+1)(p-e) & \text{if } t \in \mathcal{I}, \\ (k-1)e & \text{otherwise,} \end{cases}$$

where $\mathcal{I} = [(k+1)p - 2e, (k+1)p - e]$. The value k is given by

$$k = \begin{cases} x & \text{if } x > 1 \\ 1 & \text{otherwise} \end{cases},$$

where $x = \left\lceil \frac{t-(p-e)}{p} \right\rceil$.

Lehoczy et al. [17] proposed a demand-bound function $dbf_{RM}(\tau, t, i)$ for RM that computes the worst-case cumulative response demand of a task $\tau_i \in \tau$ for any interval of length t . It is defined by

$$dbf_{RM}(\tau, t, i) = \sum_{\tau_k \in \gamma_\tau(i)} \left\lceil \frac{t}{pk} \right\rceil \cdot e_k,$$

where $\gamma_\tau(i) = \{\tau_1, \dots, \tau_i\}$ is a function that returns a set of tasks with higher-priority (and including) than task τ_i . The *demand-bound function* for resource models is the same since the set of tasks is schedulable using the RM policy. This means that the supply of a resource model shall be greater than the demand of the set of tasks that a resource model contains.

The tasks set τ of a resource model is said schedulable according to a RM policy if, and only if,

$$\forall \tau_i \in \tau, \exists t_i \in [0, p_i] \text{ s.t. } dbf_{RM}(\tau, t_i, i) \leq sbf_\omega(t_i).$$

This approach is the state of the art on schedulability analysis for periodic resource models. We will subsume this approach with one based on timed temporal logics. Our approach allows to ensure response time guarantees about the composition with runtime monitors without employing great efforts to find more adequate optimization techniques to find the schedulability answer.

Operator	Abbreviation	Equivalent Formula
Eventually	$\diamond_{\sim\alpha}\phi$	$true U_{\sim\alpha} \phi$
Always	$\square_{\sim\alpha}\phi$	$\neg(\diamond_{\sim\alpha} \neg\phi)$
Next	$\bigcirc_{\phi_1}\phi_2$	$\phi_1 U_{\sim\infty} \phi_2$
Implies Next	$\phi_1 \implies \phi_2$	$\neg\phi_1 \vee \bigcirc_{\phi_1}\phi_2$

Table III: Syntactic abbreviations for RMTL- f

IV. RESTRICTED METRIC TEMPORAL LOGIC WITH DURATIONS

In this section we introduce the RMTL- f , a fragment of MTL- f [16] where the evaluation is carried out with respect to sequences of events produced by resource models.

The main motivation for proposing RMTL- f comes from the fact that restricting some terms and relations over terms, the logic is suitable for generation of runtime monitors as well as to thinking statically over real-time constraints. The main difference between RMTL- f and MTL- f is that the former uses only the relation \leq , $<$, and $=$ over terms, and excludes functions from the language of terms. This restriction allows us to turn our logic terms always Riemann integrable.

Definition 1 (RMTL- f): Let \mathcal{P} be a set of propositions and \mathcal{V} a set of logical variables (interpreted over \mathbb{R}). The syntax of RMTL- f is inductively defined according to the rules depicted in Table I, where δ are terms, $\int^\delta \varphi$ is the duration of the formula φ in the interval $[0, \delta]$, $x \in \mathcal{V}$, $p \in \mathcal{P}$, $\rho \in \mathbb{R}_{\geq 0}$, $\sim \in \{<, \leq, =\}$, and $\alpha \in \mathbb{R}$.

We are now able to define the semantic of the MTL- f . The semantic of MTL- f is separated in two parts: *terms* and *formulas*. The semantic of terms is defined using the notation $\mathcal{T}[\tau](\sigma, \vartheta)t$ in Table II. All terms represent numerical values in \mathbb{R}_0^+ . The term $\int^\delta \varphi$ is the integral over the Boolean function $B_{\phi(\sigma, \vartheta)}(t)$ (whose return value is 1 if $(\sigma, \vartheta, t) \models \phi$, and 0 otherwise). Since $B_{\phi(\sigma, \vartheta)}(t)$ behaves as a step function, it is always Riemann integrable. The same is not true in the MTL- f logic. The semantic of the MTL- f formula is defined inductively in Table II, where the satisfiability of a formula ϕ in a model (σ, ϑ) at time t is defined by $(\sigma, \vartheta, t) \models \phi$.

Along the remaining of the paper we will frequently refer to the abbreviations presented in Table III in order to ease the presentation of specific schedulability related specification. For illustrative purposes, we now introduce a practical example of the expressive power of RMTL- f 's language.

Example 1: To ensure that a task responds in a bounded response time, the formula $\psi_1 \implies \diamond_{\leq\alpha} \psi_2$ is sufficient. The proposition ψ_1 describes a set of events that may violate the system, the proposition ψ_2 describes the task invocation, and α is the maximum expected response time bound. Informally, the formula means that if a fault event occurs, then the task executes within α time units.

V. SCHEDULABILITY ANALYSIS USING MTL- f

Our schedulability analysis consists in the evaluation of a logic formula over a trace (or a set of traces) produced by a periodic resource model. Regarding these our approach remains

Language of RTML- \int terms	
δ	$::= \alpha \mid x \mid \int^\delta \varphi$
Language of RTML- \int formulae	
φ	$::= p \mid \delta_1 \sim \delta_2 \mid \varphi_1 \vee \varphi_2 \mid \neg \varphi \mid \varphi_1 U_{\sim \rho} \varphi_2 \mid \varphi_1 S_{\sim \rho} \varphi_2 \mid \exists x \varphi$

Table I: Syntax of RMTL- \int

Evaluation of RMTL- \int terms	
$\mathcal{T}[\alpha](\sigma, \vartheta)t$	$= \alpha$
$\mathcal{T}[x](\sigma, \vartheta)t$	$= \vartheta(x)$
$\mathcal{T}[\int^\delta \phi](\sigma, \vartheta)t$	$= \begin{cases} \int_t^{t+\mathcal{T}[\delta](\sigma, \vartheta)t} B_{\phi(\sigma, \vartheta)}(t_*) dt_* & \text{if } \mathcal{T}[\delta](\sigma, \vartheta)t \geq 0 \\ 0 & \text{otherwise} \end{cases}$
Evaluation of RMTL- \int formulae	
$(\sigma, \vartheta, t) \models p$	iff $\sigma(p)(t) = true$ and $t < \sigma $
$(\sigma, \vartheta, t) \models \delta_1 \sim \delta_2$	iff $\mathcal{T}[\delta_1](\sigma, \vartheta)t \sim \mathcal{T}[\delta_2](\sigma, \vartheta)t$
$(\sigma, \vartheta, t) \models \phi_1 \vee \phi_2$	iff $(\sigma, \vartheta, t) \models \phi_1$ or $(\sigma, \vartheta, t) \models \phi_2$
$(\sigma, \vartheta, t) \models \neg \phi$	iff $(\sigma, \vartheta, t) \not\models \phi$
$(\sigma, \vartheta, t) \models \phi_1 U_{\sim \rho} \phi_2$	iff $\exists t' \in \mathbb{R}_{\geq 0}$ st. $t \leq t' \sim t + \rho \wedge (\sigma, \vartheta, t') \models \phi_2$, and st. $\forall t'' \in \mathbb{R}_{\geq 0}, t \leq t'' < t', (\sigma, \vartheta, t'') \models \phi_1$
$(\sigma, \vartheta, t) \models \phi_1 S_{\sim \rho} \phi_2$	iff $\exists t' \in \mathbb{R}_{\geq 0}$ st. $t - \rho \sim t' \leq t \wedge (\sigma, \vartheta, t') \models \phi_2$, and $\forall t'' \in \mathbb{R}_{\geq 0}, t' < t'' \leq t, (\sigma, \vartheta, t'') \models \phi_1$
$(\sigma, \vartheta, t) \models \exists x \varphi$	iff $\exists v \in \mathbb{R}$ st. $(\sigma, \vartheta_x^v, t) \models \varphi$

Table II: Semantic of RMTL- \int

a model-checking problem [7], where the model checks a set of logic properties, and otherwise generates counter-examples.

In order to decrease the state space search we can assume for uni-processor systems the critical instant theorem [18]. This assumption reduces our problem to just one trace acceptance for a set of logic properties. This assumptions allows us to identify the relevant traces and combine our approach with the foundational real-time systems theory.

We will describe the encoding of the schedulability test for periodic resource models [23], [24] as well as their composition using our MTL- \int fragment.

A. Encoding Notations

To ease the encoding of schedulability analysis properties we first introduce some syntactical notations and formulae abbreviations. Let ω be a resource model in Ω and let τ_i be a task in τ . The set of tasks with higher-priority (and including) than τ_i for ω is denoted by $\gamma_{\omega}^{\tau_i}$. The set of resource models with higher-priority (and excluding) than ω is denoted by γ_{Ω}^{ω} . For events, we have the following notations: $\varepsilon(\omega, \cdot)$ denotes the set of events that can be generated by the resource model ω ; $evs^+(\omega_j, \tau_i)$ specifies all events that a task τ_i in the resource model ω_j can trigger. It is defined by

$$evs(\omega_j, \tau_i) \vee estop(\omega_j, \tau_i) \vee \varepsilon(\omega_j, \tau_i, \cdot) \vee erelease(\omega_j, \tau_i),$$

with $evs(\omega_j, \tau_i) \stackrel{def}{=} estart(\omega_j, \tau_i) \vee eresume(\omega_j, \tau_i) \vee erenewal(\omega)$; $evs^-(\omega_j, \tau_i)$ denotes the formula resulting from the removal of the $erelease(\omega_j, \tau_i)$ and $estop(\omega_j, \tau_i)$ events from $evs^+(\omega_j, \tau_i)$; and $evs^*(\omega_j, \tau_i)$ denotes the formula resulting from the removal of the $estart(\omega_j, \tau_i)$ and $estop(\omega_j, \tau_i)$ events from $evs^+(\omega_j, \tau_i)$.

For event occurrences we establish a MTL- \int formula that specifies the exact number of times that a proposition p holds through the following recurrent function:

$$occur(p, n) \stackrel{def}{=} \begin{cases} \Box \neg e & \text{if } n = 0 \\ \neg e U (e U occur(p, (n - 1))) & \text{otherwise,} \end{cases}$$

where $p \in \mathcal{P}$ and $n \in \mathbb{N}$ is the number of occurrences to check. Furthermore, we also introduce the definition of a function that restricts occur in the sense that it captures the period for the event under consideration. Such function is the following:

$$\mu(e, t_e, p_e, 0) \stackrel{def}{=} \Box \neg e, \\ \mu(e, t_e, p_e, (n + 1)) \stackrel{def}{=} \neg e U_{=(p_e - t_e)} (e U_{\leq t_e} o(e, t_e, p_e, n)),$$

where t_e is the time that event e consumes, and p_e the period of the event e . This definition allows us to restrict the occurrence of e by a the period p_e . In the following we introduce an example to give, using MTL- \int , the occurrences of a certain event in a trace.

Example 2: Suppose that we require to minimize the parameter z of the following formula

$$\diamond_{\leq \alpha} \text{occur}(\varepsilon(\omega, \tau_i, \cdot), z).$$

Informally, the formula indicates that there exists at most z occurrences of $\varepsilon(\omega, \tau_i, \cdot)$ until α time units. The maximal value to which z can be assigned is the positive infinity (∞). We select this maximum for z as the initial point, and decrease successively the variable z until a the formula holds or zero is achieved. This allows to find a value for z in the interval $[0, \infty)$ and to obtain the exact number of events that a trace contains. Note that this is not trivially solved, and some assumptions about the trace must be made, such as the number of events that are required to minimize z in practice (e.g., ∞ is replaced by $|\rho|$, the length of trace ρ).

B. Encoding Periodic Resource Models

Schedulability analysis over the language of RMTL- f is divided in two parts: the encoding of the scheduler's behavior – including their scheduling policy and workload parameters – and the consequent schedulability test. With both parts holding, we are able to evaluate if a given set of workload parameters is enough to be schedulable over a certain scheduler policy. We begin by detailing the encoding phase and the schedulability test.

The behavior of the scheduler is specified by several formulas within capture the budget supply, the schedulability policy, the task durations, and some intrinsic settings of the scheduler. Assuming a correct release of events, the budget supply is specified by the formula

$$\phi(\omega) \equiv \square_{\leq \infty} (\text{erenewal}(\omega) \implies rp(\omega)),$$

where

$$rp(\omega) \equiv (\diamond_{=\pi} \text{erenewal}(\omega)) \wedge \int^{\pi} \bigvee_{\tau_i \in \tau} \text{evs}^+(\omega, \tau_i) \leq \theta,$$

ω is one resource model, π and θ their renewal period and budget, and $\text{erenewal}(\omega)$ is the budget renewal event. This formula states that for each occurrence of the event $\text{erenewal}(\omega)$ in the resource model ω , the duration of the other events until π time units does not overpasses the budget θ per period π .

For the partial order of the task releases we introduce the MTL- f formula

$$\eta(\omega) \equiv \square_{\leq \infty} \bigwedge_{\tau_i \in \tau} \left(\text{erelease}(\omega, \tau_i) \implies sq(\omega, \tau_i) \right),$$

where

$$sq(\omega, \tau_i) \equiv \text{ev}(\omega, \tau_i) U_{\leq p_i} \text{estop}(\omega, \tau_i),$$

$$\text{ev}(\omega, \tau_i) \equiv \left(\bigvee_{\tau_k \in \gamma_{\omega}^{(\tau_i-1)}} \text{evs}^+(\omega, \tau_k) \right) \vee \text{evs}^-(\omega, \tau_i)$$

and $\gamma_{\omega_j}^{(\tau_i-1)}$ denotes the set of higher-priority tasks, excluding events triggered by the task τ_i . This formula means that for every event $\text{erelease}(\omega, \tau_i)$ there is always an event $\text{estop}(\omega, \tau_i)$,

and that the events occurring before $\text{estop}(\omega, \tau_i)$ should be any event from τ_i 's higher-priority tasks.

The duration of tasks allocated to one resource model is specified by the formula

$$\psi^{\leq}(\omega) \equiv \square_{\leq \infty} \bigwedge_{\tau_i \in \tau} \left(\text{erelease}(\omega, \tau_i) \implies \bigcirc du^{\leq}(\omega, \tau_i) \right),$$

where

$$du^{\leq}(\omega, \tau_i) \equiv \int^{p_i} \bigvee_{\tau_k \in \gamma_{\omega}^{(\tau_i)}} \text{evs}^+(\omega, \tau_k) \leq e_i.$$

Note that the \leq operator could be changed to \geq in order to specify the absolute WCET of the task set. We denote the duration of a task by the \geq operator as $\psi^{\geq}(\omega)$.

In order for our formalization to work, we still specify some other features such as the precedence of the event $\text{estop}(\omega, \tau_i)$ (i.e., each event $\text{estart}(\omega, \tau_i)$ is always followed by an event $\text{estop}(\omega, \tau_i)$, and vice-versa), the number of release events, and the time period at which the release of events its triggered. The precedence of the event $\text{estop}(\omega, \tau_i)$ is specified by the formula

$$\xi^1(\omega) \equiv \bigwedge_{\tau_i \in \tau} \left(\text{estop}(\omega, \tau_i) \implies \bigcirc pr(\omega, \tau_i) \right),$$

where

$$pr(\omega, \tau_i) \equiv \text{es}(\omega, \tau_i) S_{\leq p_i} \text{estart}(\omega, \tau_i),$$

and

$$\text{es}(\omega, \tau_i) \equiv \left(\bigvee_{\tau_k \in \gamma_{\omega}^{(\tau_i-1)}} \text{evs}^+(\omega, \tau_k) \right) \vee \text{evs}^*(\omega, \tau_i).$$

The release of events is captured by the recursive function $\mu(e, t_e, p_e, n)$. To ensure the periodicity of all events (erelease and r^*) for certain t time units, we introduce the formula

$$\xi^2(\omega, t) \equiv \text{oc}(\omega, t) \wedge \mu \left(\text{erenewal}(\omega), \pi, 0, \left\lfloor \frac{t}{\pi} \right\rfloor \right),$$

where

$$\text{oc}(\omega, t) \equiv \bigwedge_{\tau_i \in \tau} \mu \left(\text{erelease}(\omega, \tau_i), p_i, 0, \left\lfloor \frac{t}{p_i} \right\rfloor \right),$$

$\left\lfloor \frac{t}{\pi} \right\rfloor$ is the number of occurrences of $\text{erenewal}(\omega)$ in t , $\left\lfloor \frac{t}{p_i} \right\rfloor$ is the number of occurrences of $\text{erelease}(\omega, \tau_i)$ in t , π is the period of the budget renewal of the resource model ω , and p_i is the period of the task τ_i . Note that this formula is able to specify the number of events that can be released in t units of time for a task or a periodic resource model.

The encoding of the periodic resource model is given by

$$\text{PRM}(\omega, t) \equiv \phi(\omega) \wedge \eta(\omega) \wedge \psi^{\geq}(\omega) \wedge \xi^1(\omega) \wedge \xi^2(\omega, t),$$

where ω is defined according to certain parameters and a workload, which allows us to unroll the sub-formulas. This concludes the formalization of the periodic resource model's behavior in RMTL- f .

C. Encoding Coupled Periodic Resource Models

Here we propose an encoding of coupled periodic resource models and an analysis that ensures non-interference, and avoids priority inversion between resource models due to WCET violations.

The budgets that each resource model is allowed to use is specified by the formula

$$\phi_{\Omega}(t) \equiv \bigwedge_{\omega \in \Omega} \left(\phi(\omega) \wedge \mu \left(\text{renewal}(\omega), \pi, 0, \left\lfloor \frac{t}{\pi} \right\rfloor \right) \right).$$

With this formula, each periodic resource model meets the settings assigned to it for a given time t .

Other two definitions need to be formulated. One describes the fixed priority behavior of the periodic resource models, and the other describes the execution time allowed to a given set of resource models and their respective task sets.

The partial order of the task events for a set of resource models Ω is specified by the formula

$$\eta(\Omega) \equiv \square_{\leq \infty} \bigwedge_{\tau_i \in \tau} \bigwedge_{\omega \in \Omega} \left(\text{erelease}(\omega, \tau_i) \stackrel{\circ}{\implies} \text{or}(\Omega, \omega, \tau_i) \right),$$

where

$$\text{or}(\Omega, \omega, \tau_i) \equiv (\text{re}(\Omega, \omega, \tau_i) \vee \text{evs}^-(\omega, \tau_i)) \ U_{\leq p_i} \ \text{estop}(\omega, \tau_i),$$

and

$$\text{re}(\Omega, \omega, \tau_i) \equiv \bigvee_{\omega \in \gamma_{\Omega}^{\omega}} \bigvee_{\tau_j \in \tau} \text{evs}^+(\omega, \tau_j) \vee \bigvee_{\tau_k \in \gamma_{\omega}^{(\tau_i-1)}} \text{evs}^+(\omega, \tau_k).$$

The formula $\text{re}(\Omega, \omega, \tau_i)$ describes the resource events that can occur before an event $\text{estop}(\omega, \tau_i)$ event.

The execution time allowed for a set of resource models Ω is defined by

$$\psi_{\Omega}^{\leq}(t) = \bigwedge_{\omega \in \Omega} (\psi^{\leq}(\omega) \wedge \text{oc}(\omega, t)).$$

To specify the worst-case we can use $\psi^{\geq}(\omega)$ instead of $\psi^{\leq}(\omega)$. Substituting the above formula we have $\psi_{\Omega}^{\geq}(t)$.

The composition of periodic resource models is encoded by the RMTL- f formula

$$CPRM(\Omega, t) \equiv \phi_{\Omega}(t) \wedge \eta(\Omega) \wedge \psi_{\Omega}^{\leq}(t) \wedge \left(\bigwedge_{\omega \in \Omega} \xi^1(\omega) \right).$$

D. Schedulability Test

To provide schedulability tests for our encodings we need to find a model that satisfies the PRM formula for resource models, and the $CPRM$ formula for a composition of resource models. By the semantic definition of RMTL- f we need to find an observation, a logical environment, and a duration in accordance with the specified properties. This behavior is formulated by the following definitions.

Definition 2: Let ω be a resource model in Ω . The resource model ω is schedulable if and only if, there exists a trace ρ of

duration t such that $PRM(\omega, t)$ is satisfied, and the duration of ρ is greater or equal to the maximum value of p_i in τ .

Definition 3: Let Ω be a set of resource models, and ω a resource model in Ω . The composition of resource models Ω is schedulable if and only if, there exists a trace ρ with duration t such that $CPRM(\Omega, t)$ holds and t of trace ρ is greater or equal than the maximum $p_i \in \tau$, of all resource models in ω .

Summarizing our definitions states that the our schedulability decision problem is a satisfiability problem of a trace regarding a RMTL- f formula.

E. Feasible Tests

Encoding the schedulability test is not enough to ensure that we always obtain a positive or negative answer. In order to cope with this problem, we make the necessary assumption on the structure of traces so that their evaluation indeed produces some verdict, *i.e.*, if the system under consideration is schedulable or not.

To find the worst execution trace we begin by the introduction of the following definition.

Definition 4: The worst execution of a resource model is a trace that complains the budgets supply, the schedulability policy, the WCET of a task set, and a restricted set of intrinsic formulas.

To generate the worst execution trace we can adopt two distinct strategies. On one hand, we can assume some theorem that gives freely and by construction the worst trace that a system can generate (*e.g.*, for uni-processor systems we could adopt the critical instant theorem [18]). On another hand, we can rewrite our schedulability decision test into a Boolean satisfiability problem. In this paper we will focus only the first one, and address the second one to further work. However, we believe that the second one is able to extend schedulability analysis for multi-processor systems.

Given a worst execution trace we are able to evaluate such formula using our semantics which decides if the trace is valid or invalid among the logic formulas that describe the scheduler behavior. The process remains the check of the logical formula PRM or $CPRM$ for a given trace.

VI. EXEMPLIFICATION OF OUR SCHEDULABILITY ANALYSIS

Our schedulability analysis for several period resource models relax the truth notion of the WCET. This means that the WCET of a task (or several tasks) can be erroneously estimated, and ensures that the remain resource models are schedulable. In the following, we present an example of how to use the Definition 2 and Definition 3 for periodic resource models, and a composition of resource models, respectively.

To demonstrate in practice the schedulability analysis using our logic fragment, a synthetic workload will be described. Suppose, for example, a workload composed by three components, four tasks, and two monitors as depicted in Table IV. By Definition 3 we may conclude that the workload is schedulable if there exists a trace that complains our logic restrictions.

Before introducing the example we need to make two notes. Our schedulability analysis does not strongly assume the

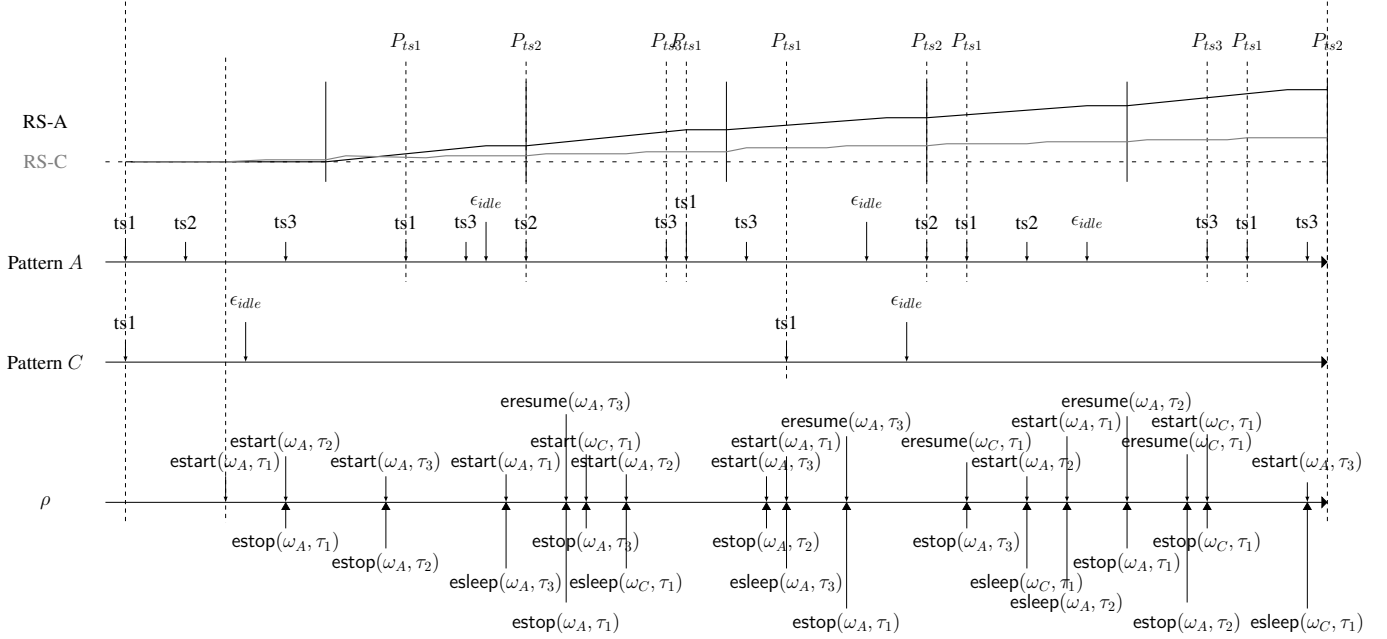


Figure 1: Example of patterns and the global trace generated by the composition of resource models defined in the Table IV

	Resource Model		Task				Priority
	π	θ	id	period	priority	WCET	
RS-A	10	8	ts1	14	1	3	2
			ts2	20	2	5	
			ts3	27	3	7	
RS-C	5	1	ts1	33	1	6	1

Table IV: A synthetic workload scheme

behavior of the scheduler (e.g., the periodic resource model), and if the decision is negatively affirmed, a counter-example is returned (i.e., a trace). We can also state that for every trace generated by a scheduler if the behavior does not correspond to the specified one then the scheduler is not a periodic resource model.

We will introduce our example assuming the critical instant theorem. Assuming this theorem we can find the *worst execution trace* for a certain workload settings. This problem is converted to an acceptance problem. We only need to apply our evaluation of $\text{RMTL-}\int$ formulas to draw a verdict about the schedulability.

A. Unfold the $\text{RMTL-}\int$ formulas

Our schedulability analysis provides two definitions for schedulability testing. According to Figure 1, we will explain step-by-step how the evaluation is performed. Beginning by unfolding the $\phi(\omega)$ of the PRM formula for the resource model RS-C, we have the formula on Table V. The example is for a trace with duration 4 but the truth value is the same for the Figure 1. Note that this formula needs to be fed with traces

whose durations are multiples of 5. Otherwise the meaning of the formula is false due to the eventually operator. The $\psi^{\geq}(\omega)$ formula ensures that the task can be executed by their WCET. Since the resource model RS-C only contains one task, we need to ensure the worst duration for this task as specified in Table V. We evaluate the formula with a trace of duration 10. The remaining formulas $\eta(\omega)$, $\xi^1(\omega)$, $\xi^2(\omega, t)$ are trivially satisfied since we have only one task in our resource model RS-C. Note that the release events of the periodic resource models and the tasks are not considered in trace ρ to decrease the trace complexity, but they are considered when the formulas are evaluated.

Considering the resource model RS-A, which has a more elaborated formula after its unfolding, we are able to exemplify why the composition of two schedulable resource models are not schedulable when coupled. In Figure 1 we have generated a counter-example, namely, the trace ρ . Applying the formula CPRM for the composition, the formula $\eta(\Omega)$ does not hold since it is impossible to force the consumption of the WCET of a task until its next period (only five units can be assigned until the period of the next execution). If we change the period of the task τ_1 in RS-C to value 50 instead of 33 our composition of resource models RS-A and RS-C is schedulable. We can check this informally by looking at Figure 1, unfold our formula CPRM and draw a verdict for each formula of the resulting conjunction.

VII. CONCLUSION AND FURTHER WORK

In this paper we have introduced a novel approach to schedulability analysis based on timed temporal logics. Compared with classical methods, our approach has a built-in

$\phi(\text{RS-C})$	True	$\Box_{\leq 4} (\text{erenewal}(\omega) \implies (\Diamond_{=5} \text{erenewal}(\omega)) \wedge \int^5 \text{evs}^+(\text{RS-C}, \tau_1) \leq 1)$
$\psi^{\geq}(\text{RS-C})$	True	$\Box_{\leq 10} \neg (\text{erelease}(\text{RS-C}, \tau_1)) \vee (\neg (\text{erelease}(\text{RS-C}, \tau_1)) U_{\leq 1} \int^{P1} \text{evs}^+(\text{RS-C}, \tau_1) \geq e_1)$
Trace	erenewal(RS-A), erenewal(RS-C), estart(RS-A, τ_1), estop(RS-A, τ_1), estart(RS-A, τ_2), erenewal(RS-C), estop(RS-A, τ_2), estart(RS-A, τ_3), erenewal(RS-A), erenewal(RS-C), esleep(RS-A, τ_3), \dots	

Table V: Unfold of the formulas $\phi(\text{RS-C})$ and $\psi^{\geq}(\text{RS-C})$ and their truth value in accordance with trace ρ .

scheduler behavior; avoids the rate approximations of events as experienced in [26], allows us to extend this analysis for runtime monitoring architectures by ensuring the maximum detection delay of the monitors with a simple response time RMTL- \int formula; and supplies a predictable trace set of traces that can be analyzed prior to the execution and provide counterexamples.

In terms of future work, our aim is to implement a verification platform that incorporates the ideas presented in this paper with primary focus on the automatic synthesis of RMTL- \int specifications into runtime monitors and corresponding program instrumentation. Alternatively, we are also interested in encoding our schedulability test into reachability analysis and use model checking tools such as e.g. the NUSMV model checker tool [6] to check them, or by using statistical methods to solve such issue, which is commonly natural in cases where the previous methods do not have enough resources to do the job. Yet another alternative is to encode our language in some formal verification framework, such as the Why3 or BoogiePL [8] intermediate verification languages, and rely on their backend provers to improve the chances of automatically proving the properties.

VIII. ACKNOWLEDGMENTS

This work was partially supported by National Funds through FCT (Portuguese Foundation for Science and Technology) and by ERDF (European Regional Development Fund) through COMPETE (Operational Programme 'Thematic Factors of Competitiveness'), within projects Ref. FCOMP-01-0124-FEDER-022701 (CISTER), FCOMP-01-0124-FEDER-015006 (VIPCORE) and FCOMP-01-0124-FEDER-020486 (AVIACC).

REFERENCES

- [1] N. C. Audsley, A. Burns, R. I. Davis, K. W. Tindell, and A. J. Wellings. Fixed priority pre-emptive scheduling: an historical perspective. *Real-Time Syst.*, 8(2-3):173–198, March 1995.
- [2] G. Behrmann, A. David, K. G. Larsen, J. Hakansson, P. Petterson, W. Yi, and M. Hendriks. UPPAAL 4.0. QEST '06, pages 125–126, 2006.
- [3] H. Ben-Abdallah, J. Choi, D. Clarke, Y. S. Kim, I. Lee, and H. Xie. A process algebraic approach to the schedulability analysis of real-time systems. *Real-Time Syst.*, 15(3):189–219, 1998.
- [4] F. Bobot, J. C. Filliâtre, C. Marché, and A. Paskevich. Why3: Shepherd Your Herd of Provers. In *Boogie 2011: First International Workshop on Intermediate Verification Languages*, pages 53–64, 2011.
- [5] A. Burns and T. M. Lin. An engineering process for the verification of real-time systems. *Form. Asp. Comput.*, 19(1):111–136, 2007.
- [6] A. Cimatti, E. M. Clarke, F. Giunchiglia, and M. Roveri. Nusmv: a new symbolic model checker. *International Journal on Software Tools for Technology Transfer*, 2(4):410–425, 2000.
- [7] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model checking*. MIT Press, Cambridge, MA, USA, 1999.
- [8] Robert Deline, K. Rustan, and M. Leino. Boogiepl: A typed procedural language for checking object-oriented programs. Technical report, 2005.
- [9] E. Fersman, P. Krčál, P. Pettersson, and W. Yi. Task automata: Schedulability, decidability and undecidability. *Inf. Comput.*, 205(8):1149–1172, 2007.
- [10] E. Fersman, L. Mokrushin, P. Pettersson, and W. Yi. Schedulability analysis of fixed-priority systems using timed automata. *Theor. Comput. Sci.*, 354(2):301–317, 2006.
- [11] E. Fersman, P. Pettersson, and W. Yi. Timed Automata with Asynchronous Processes: Schedulability and Decidability. TACAS '02, pages 67–82, 2002.
- [12] C. J. Fidge. Real-time schedulability tests for preemptive multitasking. *Real-Time Syst.*, 14(1):61–93, January 1998.
- [13] Z. Gu and K. G. Shin. Analysis of event-driven real-time systems with time petri nets: A translation-based approach. DIPES '02, pages 31–40, 2002.
- [14] P. Krčál, M. Stigge, and W. Yi. Multi-processor schedulability analysis of preemptive real-time tasks with variable execution times. FORMATS'07, pages 274–289, 2007.
- [15] P. Krčál and W. Yi. Decidable and undecidable problems in schedulability analysis using timed automata. volume 2988 of *TACAS'04*, pages 236–250. Springer-Verlag, 2004.
- [16] Y. Lakhneche and J. Hooman. Metric temporal logic with durations. *Theor. Comput. Sci.*, 138(1):169–199, 1995.
- [17] J. Lehoczky, Lui Sha, and Y. Ding. The rate monotonic scheduling algorithm: exact characterization and average case behavior. In *Real Time Systems Symposium*, pages 166–171, 1989.
- [18] C. L. Liu and James W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM*, 20(1):46–61, January 1973.
- [19] Dillon Pariente and Emmanuel Ledinot. Formal Verification of Industrial C Code using Frama-C: a Case Study. FoVeOOS'10, 2010.
- [20] Anna Philippou, Insup Lee, Oleg Sokolsky, and Jin-Young Choi. A process algebraic framework for modeling resource demand and supply. FORMATS'10, pages 183–197, 2010.
- [21] L. Pike, A. Goodloe, R. Morisset, and S. Niller. Copilot: a hard real-time runtime monitor. RV'10, pages 345–359, 2010.
- [22] S. Pinisetty, Y. Falcone, T. Jérón, H. Marchand, A. Rollet, and O. Nguena Timo. Runtime enforcement of timed properties. RV'13, pages 229–244. 2013.
- [23] I. Shin and I. Lee. Periodic Resource Model for Compositional Real-Time Guarantees. RTSS '03, pages 2–13, 2003.
- [24] I. Shin and I. Lee. Compositional real-time scheduling framework with periodic model. *ACM Trans. Embed. Comput. Syst.*, 7(3):30:1–30:39, 2008.
- [25] J. P. Tsai, S. J. Yang, and Y. Chang. Timing constraint petri nets and their application to schedulability analysis of real-time system specifications. *IEEE Trans. Softw. Eng.*, 21(1):32–49, 1995.
- [26] H. Zhu, M. B. Dwyer, and S. Goddard. Predictable Runtime Monitoring. ECRTS '09, pages 173–183, 2009.
- [27] H. Zhu, S. Goddard, and M.B. Dwyer. Selecting server parameters for predictable runtime monitoring. RTAS'10, pages 227–236, 2010.