



Proceedings of the  
14th International Workshop on  
Automated Verification of Critical Systems (AVoCS 2014)

Adaptive Task Automata with Earliest-Deadline-First Scheduling

Leo Hatvani, Alexandre David, Cristina Seceleanu and Paul Pettersson

15 pages

## Adaptive Task Automata with Earliest-Deadline-First Scheduling

Leo Hatvani<sup>1</sup>, Alexandre David<sup>2</sup>, Cristina Seceleanu<sup>3</sup> and Paul Pettersson<sup>4</sup>

<sup>1</sup>[leo.hatvani@mdh.se](mailto:leo.hatvani@mdh.se)

<sup>3</sup>[cristina.seceleanu@mdh.se](mailto:cristina.seceleanu@mdh.se)

<sup>4</sup>[paul.pettersson@mdh.se](mailto:paul.pettersson@mdh.se)

Mälardalen University  
Västerås, Sweden

<sup>2</sup>[adavid@cs.aau.dk](mailto:adavid@cs.aau.dk)

Aalborg University  
Aalborg, Denmark

**Abstract:** Adjusting to resource changes, dynamic environmental conditions, or new usage modes are some of the reasons why real-time embedded systems need to be adaptive. This requires a rigorous framework for designing such systems, to ensure that the adaptivity does not result in invalidating the system's real-time constraints.

To address this need, we have recently introduced adaptive task automata, a framework for modeling, verification, and schedulability analysis in adaptive, hard real-time embedded systems, assuming a fixed-priority scheduler.

In this work, we extend the adaptive task automata framework to incorporate the earliest-deadline-first scheduling policy, as well as enable implementation of any other dynamic scheduling policy. To prove the decidability of our model, and at the same time maintain a manageable degree of conciseness, we show an encoding of our model as a network of timed automata with clock updates. To support this, we also show that reachability in our class of timed automata with updates is decidable. Our contribution helps to streamline the process of designing safety critical adaptive embedded systems.

**Keywords:** model-checking, task automata, earliest-deadline-first scheduling

## 1 Introduction

One way to enable real-time embedded systems to cope with environment, application, or platform changes is to introduce adaptivity at the design phase of system development. Adaptivity lets the system adjust to a new situation, but at the same time may introduce new errors such as breached timing constraints or other extra-functional requirements. Our goal is to propose a way to streamline modeling and verification of adaptive embedded systems (AES) in order to minimize the introduction of such errors at the design stage.

In the framework of *adaptive task automata* (ATA) that we have recently proposed [HPS12], we have started to address this need by providing formal support for modeling the AES behavior, simulation of the system execution, and verification of the schedulability. By formally verifying the

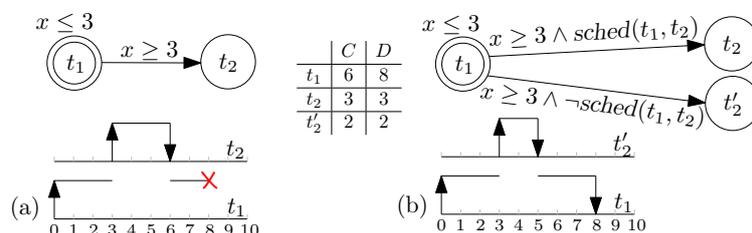


Figure 1: An adaptation example: (a) task automaton model, and (b) ATA model.

system's schedulability, we ensure that the system is going to meet its hard real-time specifications as well as satisfy any other extra-functional properties.

In our previous work on adaptive task automata, we have assumed fixed priority scheduling (FPS) policy. In this work we are extending the framework to support dynamic scheduling policies by incorporating the earliest-deadline-first (EDF) scheduling policy into the framework. Hereinafter we will refer to the specific variant of ATA with the EDF scheduling policy as  $ATA_{EDF}$ .

The main contribution of this work is to find solutions to the challenges of verifying the EDF schedulability of hard real-time tasks, in ATA. To tackle this, we show that verification of schedulability in  $ATA_{EDF}$ , described in Section 2, is decidable, by proposing an encoding of the framework as a network of timed automata with (clock) updates (Section 3). We present a summary of the proof of bisimilarity between the model and its encoding as well as decidability of reachability for our class of timed automata with updates (Section 4).

## 2 Adaptive Task Automata

The adaptive task automata framework builds on top of task automata [FKPY07] by providing predicates that influence task release patterns based on the content of the ready queue. The task automata framework, in turn, is based on timed automata [Alu99] extended with: tasks that can be released upon entering locations, a queue, and a scheduler to handle the released tasks and simulate their execution. Since the current work elaborates on ATA extensively, we refer the reader to the cited literature for more in-depth information.

In our model, we assume a uniprocessor system with independent, non-suspending tasks. For each task, computation time and relative deadline are known and are specified as natural numbers. At any point in time, there can be at most one task instance (job) per task in the queue and will be also referred to as task.

### 2.1 Introductory Example

As a simple example, consider the set of tasks in Figure 1. Each task is characterized by its execution time  $C$  and a relative deadline  $D$ . Figure 1(a) models the release of the task  $t_1$  at time 0 by annotating the initial location (double concentric circle) with the task. Task  $t_2$  is released in the second location after 3 time units. The delay is modeled by adding a zero-initialized clock ( $x$ ) to the system, annotating the initial location with the invariant  $x \leq 3$  that models that the location will be exited after at most 3 time units, and adding a guard  $x \geq 3$  on the edge, denoting that the

edge will not be taken until at least 3 time units have passed.

If we schedule the model in Figure 1(a) using EDF, the deadline of the task  $t_1$  will be reached before the task has a chance to complete. Assuming that we have  $t'_2$ , a lower quality alternative to task  $t_2$ , having a lower computation time, we could release  $t'_2$  instead. To be able to choose the variant of the task to be released, we have introduced the following predicates in our previous work [HPS12]:

- $inqueue(t_i)$  which is true iff the task  $t_i$  is waiting in the ready queue or currently executing.
- $sched(t_i)$  evaluates whether the task  $t_i$  is going to complete its execution by the deadline.
- $sched(t_i, t_j)$ , assuming that the task  $t_i$  is already in the queue, evaluates whether it will complete in time if the task  $t_j$  is released into the queue.

By incorporating the predicate  $sched(t_i, t_j)$  into the model of Figure 1(a), we get the model presented in Figure 1(b). Here, task  $t_2$  is released only if it will not disrupt task  $t_1$ , otherwise, task  $t'_2$  is released. With this modification, which can be seen as adaptive behavior, both tasks can successfully complete.

## 2.2 Overview of the Existing Framework

In ATA, the ready queue is a sequence of tasks ordered by the scheduling policy. Each task  $t_i$  in the ready queue is defined by two real values  $c_i$  and  $d_i$ . They represent the remaining execution time until completion ( $c_i$ ) and the time until the task reaches its deadline ( $d_i$ ).

Let us denote by  $T$  the set of tasks, and by  $P(T)$ , ranged over by  $p$ , the set of various Boolean combinations of the above predicates over the set of tasks. Utilizing this notation, an adaptive task automaton can be defined as follows.

**Definition 1** [HPS12] An adaptive task automaton over actions  $Act$ , clocks  $X$ , invariants  $\Phi(X)$ , guard constraints  $B(X)$ , tasks  $T$ , and predicates over tasks  $P(T)$  (Definition 3) is a tuple  $\langle Act, X, L, l_0, E, I, M \rangle$  where  $L$  is a finite set of locations,  $l_0 \in L$  is the initial location,  $E \subseteq L \times B(X) \times P(T) \times Act \times 2^X \times L$  is the set of edges,  $I : L \mapsto \Phi(X)$  is a function assigning each location an invariant, and  $M : L \mapsto T$  is a function annotating locations with tasks.

Guard constraints  $B(X)$  are a set of conjunctions of atomic constraints of the type  $x \sim C$  or  $x - y \sim C$  where  $x, y \in X$  are clocks,  $C$  is a natural number, and  $\sim \in \{<, \leq, =, \geq, >\}$ . Invariants  $\Phi(X)$  are a set of conjunctions of atomic constraints of the type  $x \sim C$  where  $x \in X$  is a clock,  $C$  is a natural number, and  $\sim \in \{<, \leq\}$ .

In the case of  $(l, g, p, a, r, l') \in E$ , we write  $l \xrightarrow{g, p, a, r} l'$ , where  $g \in B(X)$  is a guard constraint,  $a \in Act$  is an action, and  $r$  is the subset of clocks that will be reset on taking the edge.  $\square$

We can represent the state of an adaptive task automaton as a triple  $\langle l, u, q \rangle$ , where  $l \in L$  is the current location,  $u \mapsto \mathbb{R}_{\geq 0}$  is a function mapping clocks to non-negative real values, and  $q = [t_0(c_0, d_0), \dots, t_n(c_n, d_n)]$  is the current ready task queue.  $Sch(q)$  is a function that returns the ready queue sorted according to the scheduling policy, and  $Run_{Sch}(q, \delta)$  is a function that returns the ready queue after it was executed for  $\delta$  time units.

**Definition 2** [HPS12] Given an adaptive task automaton  $\langle Act, X, L, l_0, E, I, M \rangle$  with an initial state  $\langle l_0, u_0, q_0 \rangle$ , and a scheduling strategy Sch, its semantics is a transition system defined as:

$$\begin{aligned} \langle l, u, q \rangle &\xrightarrow{a}_{\text{Sch}} \langle l', r(u), \text{Sch}(M(l') :: q) \rangle \text{ if } l \xrightarrow{g,p,a,r} l' \in E, q \models p, \text{ and } u \models g \\ \langle l, u, q \rangle &\xrightarrow{\delta}_{\text{Sch}} \langle l, u \oplus \delta, \text{Run}_{\text{Sch}}(q, \delta) \rangle \text{ if } (u \oplus \delta) \models I(l) \end{aligned}$$

where  $r(u)$  is 0 for all  $x_i \in r$  and  $u(x_i)$  otherwise,  $t :: q$  is the result of releasing  $t$  into the queue  $q$ , and  $u \oplus \delta$  is the result of adding  $\delta \in \mathbb{R}_{\geq 0}$  to all clock values in  $u$ . If both transitions are enabled, the choice is non-deterministic.  $\square$

Intuitively, in the context of tasks, transitions are possibilities to release new tasks, while delays in locations correspond to the execution of tasks.

**Definition 3** [HPS12] Given a task automaton state  $\langle l, u, q \rangle$ , with  $q = [t_0(c_0, d_0), \dots, t_n(c_n, d_n)]$ , a scheduling policy Sch, and two distinct tasks,  $t_i$  and  $t_j$ , let  $P$  be the set of predicates  $\{\text{inqueue}(t_i), \text{sched}(t_i), \text{sched}(t_i, t_j)\}$  satisfied as follows:

$$\begin{aligned} \langle l, u, q \rangle &\models \text{inqueue}(t_i) \text{ if } t_i \in q \\ \langle l, u, q \rangle &\models \text{sched}(t_i) \text{ if } \text{inqueue}(t_i) \wedge (c_i + \sum_{j \in \text{HP}(t_i)} c_j) \leq d_i \vee \\ &\quad \neg \text{inqueue}(t_i) \wedge \langle l, u, \text{Sch}(t_i :: q) \rangle \models \text{sched}(t_i) \\ \langle l, u, q \rangle &\models \text{sched}(t_i, t_j) \text{ if } \text{inqueue}(t_i) \wedge \langle l, u, \text{Sch}(t_j :: q) \rangle \models \text{sched}(t_i) \end{aligned}$$

where  $\text{HP}(t_i)$  is the set of all tasks that have higher priority than  $t_i$ , and  $\text{Sch}(t_j :: q)$  is the queue ordered by the scheduling policy Sch after the release of the task  $t_j$ .

Boolean combinations of the above predicates over a set of tasks  $T$  give us the set of all possible combinations of predicates denoted by  $P(T)$ .  $\square$

### 3 Encoding of $\text{ATA}_{\text{EDF}}$

In order to show the decidability of the  $\text{ATA}_{\text{EDF}}$  framework, we have encoded the universal  $\text{ATA}_{\text{EDF}}$  model as a network of timed automata with (clock) updates (TAU). First we present the framework of timed automata with updates. The framework was introduced previously by Bouyer et al. [BDFP04], yet we use a variant whose decidability has to be proven for our result to hold. Then the encoding itself is laid out in three steps. The first step shows a way to encode task releases, the second provides the intuition behind the encoding of the predicates used for adaptivity, and the third introduces the encoding of the scheduler. After we have encoded the system as timed automata with updates, we provide a proof that the reachability problem for our class of timed automata with updates is decidable and that the encoding is bisimilar to the original model. The  $\text{ATA}_{\text{EDF}}$  is more challenging than ATA as the task priorities are decided online.

#### 3.1 Timed Automata with Updates

The timed automata framework, as defined by Alur and Dill [AD94], has served as the basis for several modeling variations proposed in order to fit specific design purposes [FKPY07, BDFP04, LBB<sup>+</sup>01]. Along the same line, our approach also relies on a variant of timed automata.

To concisely encode the scheduler model as timed automata, we need to allow for “clock to clock” assignments. Although such clock assignments are already present in the updatable timed automata framework [BDFP04], they are defined on models without invariants on locations. Since our work depends on location invariants, let us define the extension of timed automata that supports clock to clock assignments as well as location invariants.

**Definition 4** A *timed automaton with updates* (TAU) over clocks  $X$  and actions  $Act$  is a tuple  $\langle Act, X, L, l_0, E, I \rangle$ , where  $L$  is a finite set of locations,  $l_0$  is the initial location,  $E \subseteq L \times B(X) \times Act \times 2^X \times 2^{X^2} \times L$  is the set of edges, and  $I : L \rightarrow \Phi(X)$  assigns invariants to locations. In the set of edges  $E$ ,  $B(X)$  is the set of guard constraints,  $2^X$  represents the set of clock resets, and  $2^{X^2}$  represents the set of clock assignments of the form  $x := y$ , where  $x, y \in X$ .

The set of invariants  $\Phi(X)$  is a set of conjunctions of atomic expressions of the type  $x \sim C$  where  $x \in X$  is a clock,  $C$  is a natural number, and  $\sim \in \{<, \leq\}$ . The set of guard constraints  $B(X)$  can be defined as a set of Boolean combinations of atomic expressions of the type  $x \sim C$  or  $x - y \sim C$  where  $x, y \in X$  are clocks, and  $\sim \in \{<, \leq, =, \geq, >\}$ .

In the case of  $(l, g, a, r, s, l') \in E$ , we write  $l \xrightarrow{g, a, r, s} l'$ , where  $r$  is the subset of clocks that will be reset on taking the edge, and  $s$  the set of clock assignments.  $\square$

The semantics of TAU is defined in terms of a timed transition system over states of the form  $(l, u)$ , where  $l$  is a location,  $u \mapsto \mathbb{R}_{\geq 0}$  is an assignment of clocks to non-negative real values, and the initial state is  $(l_0, u_0)$ , where  $u_0$  assigns all clocks in  $X$  to 0.

**Definition 5** Given a timed automaton with updates  $\langle Act, X, L, l_0, E, I \rangle$  with an initial state  $(l_0, u_0)$ , its semantics is a transition system defined as:

- $\langle l, u \rangle \xrightarrow{a} \langle l', r(s(u)) \rangle$  if  $l \xrightarrow{g, a, r, s} l' \in E$  and  $u \models g$
- $\langle l, u \rangle \xrightarrow{\delta} \langle l, u \oplus \delta \rangle$  if  $(u \oplus \delta) \models I(l)$

where  $s(u)$  performs the assignments  $x_i := x_j$  for every  $(x_i, x_j) \in s$ ,  $r(u)$  is 0 for all  $x_i \in r$  and  $u(x_i)$  otherwise, and  $u \oplus \delta$  is the result of adding  $\delta \in \mathbb{R}_{\geq 0}$  to all clock values in  $u$ . If both transitions are enabled, the choice is non-deterministic.  $\square$

A timed trace  $\sigma$  of a TAU, as is also the case with timed automata [AD94], is a sequence of delay and action transitions  $\sigma = (l_0, u_0) \xrightarrow{a_1} (l_1, u_1) \xrightarrow{a_2} \dots \xrightarrow{a_n} (l_n, u_n)$  where  $a_i$  can be either action ( $\xrightarrow{a}$ ) or delay ( $\xrightarrow{\delta}$ ) transition, and a location  $l$  is said to be *reachable* if there exists a timed trace ending in the state  $(l, u)$ .

A *network* of TAU,  $A_1 || \dots || A_n$  over  $X$  and  $Act$  is defined as the parallel composition of  $n$  TAU over  $X$  and  $Act$ . Semantically, a network of TAU again describes a timed transition system obtained from those components, by requiring action transitions to synchronize on complementary actions (i.e.,  $a?$  is complementary to  $a!$ ) [BY04].

### 3.2 Earliest-Deadline-First Scheduling Policy

To encode the scheduler, we need to clearly define the EDF policy in the context of this paper. Since the strategy for choosing the next task between two or more tasks with equal deadlines does

not impact the optimality of the EDF algorithm [GD99], we can give the following definition of EDF with deterministic tie resolution.

**Definition 6** According to the EDF scheduling policy with deterministic tie resolution, the priority  $P_i$  of task  $t_i$  is greater than the priority  $P_j$  of task  $t_j$  if the time left until the absolute deadline  $d_i$  of task  $t_i$  is smaller than the time left until the absolute deadline  $d_j$  of task  $t_j$ , or their absolute deadlines are equal and  $i > j$  holds. This can be expressed as

$$P_i > P_j \iff d_i < d_j \vee (d_i = d_j \wedge i > j)$$

where  $i$  and  $j$  represent strictly ordered task indices. □

### 3.3 Task Releases

In ATA, tasks are released on changing to locations that are annotated with sets of tasks. A straightforward method to realize instant task triggering upon entering a location is to use synchronization channels on the edges of the corresponding TAU representation. This is demonstrated in Figure 2.

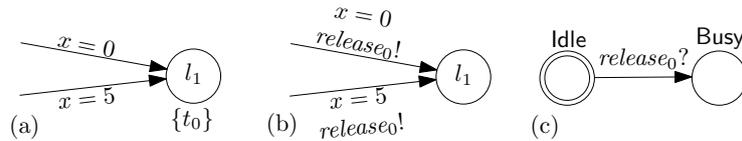


Figure 2: (a) task automaton, (b) (a)'s encoding, (c) part of (b)'s scheduler

In Figure 2(a), we have a basic task automaton location with two disjunctive edges leading to it. Location  $l_1$  is annotated with the task set  $\{t_0\}$ . By entering the location via any of the edges, the task  $t_0$  should be released and handled by the scheduler.

Modeling this behavior in TAU requires annotating every edge entering the location  $l_1$  with a synchronization channel that creates a network of timed automata between the observed automaton presented in Figure 2(b) and the corresponding edges in the scheduler automaton as seen in Figure 2(c). In some cases, additional committed locations [BGK<sup>+</sup>02] might be needed to accomplish this.

### 3.4 Schedulability Predicates

The ATA model implements adaptivity via a set of scheduling predicates that may restrict edge guards:  $sched(t_i)$ ,  $sched(t_i, t_j)$ , and  $inqueue(t_i)$ . All predicates are evaluated within the context of the current ready queue.

To express the predicates in timed automata with updates, we need to define an adequate encoding of the relevant variables that describe tasks in ATA models. The task automata model and consequently the adaptive task automata model define the task  $t_i$  in terms of remaining computation time  $c_i$  and time left until the deadline  $d_i$ . We encode the remaining computation time as the difference between the response time  $r_i$  and the computation time  $c_i$ :  $c_i = r_i - c_i$ .

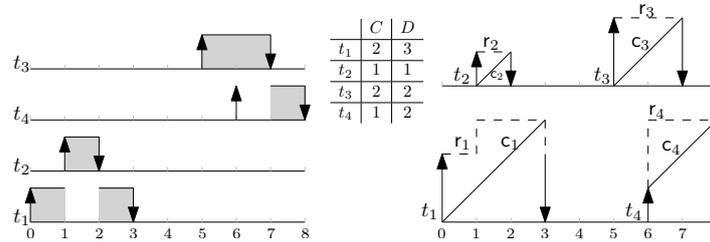


Figure 3: Gantt chart and the encoding specific representation of tasks

To illustrate this encoding, let us observe Figure 3. The left side of the figure presents a Gantt chart of task releases, while the right side presents a graph of the values of the variables  $c$  and  $r$  for the same set of tasks. Note that, in the graph, the tasks  $t_2$  and  $t_3$ , as well as  $t_1$  and  $t_4$  are presented on the same level to conserve vertical space.

At time 0, task  $t_1$  is released. A higher priority task  $t_2$  preempts it at time 1. At the moment of preemption, the response time  $r_1$  is increased by  $C_2$ , the computation time of  $t_2$ , while the response time of task  $t_2$  is equal to its computation time. Both tasks complete when their computation time becomes equal to their response time, respectively.

Two time units after task  $t_1$  completes, task  $t_3$  is released. It is already executing when task  $t_4$  is released. Although task  $t_4$  has computation time of only 1 time unit, its response time already accounts for  $t_3$ . Due to the continuous nature of timed automata clocks, we cannot extract information on how much of the computation time of task  $t_3$  has been already used, so we have to use the full response time of task  $t_3$  increased by the response time of task  $t_4$ . In order for this response time to be in context, we also need to copy the clock value of  $c_3$  to  $c_4$ , hence the clock  $c_4$  starts from 1.

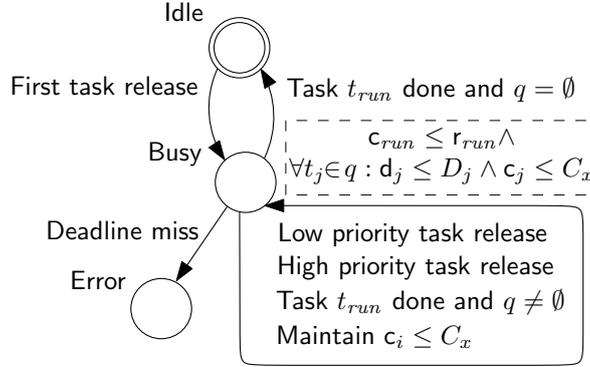
The time until deadline is encoded by simply comparing an increasing clock to the relative deadline, but it is not shown here.

### 3.5 Scheduler and Queue

Next, we encode the EDF scheduling policy together with the queue as a single automaton, which we will hereafter refer to as the scheduler automaton.

Our scheduler is created assuming the encoding of predicates outlined in the previous section and the EDF policy presented earlier. These two constraints, addressed at the same time, have significantly increased the complexity of the encoding. In Figure 4, we show the entire scheduler model encoded as a timed automaton with updates, using synchronization channels to release tasks.

To reduce the presentation complexity of the encoding and make it more accessible to human readers, we have used a number of shorthands. For example, the queue is encoded as the set  $q$ . Since this set is referenced in every location, we need to replicate each location for every possible value of  $q$ . Since the number of tasks in the system ( $N$ ) is finite and known in advance, this means that there will be  $2^N$  replications of every location to reflect the set  $q$ . Next, only those locations that imply values that satisfy the incoming guards are connected by the edges to the originating location. The same approach can be applied to all other integer variables and translate



<b>First task release</b>	<b>Deadline miss</b>
Sync $release_i?$	Guard $t_i \in q \wedge c_i < r_i \wedge d_i \geq D_i$
Update $q := q \cup \{t_i\}; t_{run} := t_i; r_i := C_i;$ $c_i := 0; d_i := 0; P_i := N$	<b>Low priority task release</b>
<b>Task <math>t_{run}</math> done and <math>q = \emptyset</math></b>	Guard $t_{next}^{EDF}(t_i) = t_j \wedge c_{run} < r_{run}$
Guard $c_{run} = r_{run} \wedge d_{run} \leq D_{run} \wedge q = \{t_{run}\}$	Sync $release_i?$
Update $q := q \setminus \{t_{run}\}$	Update $q := q \cup \{t_i\};$ $\forall k \in q   P_k < P_j : P_k := P_k - 1;$ $P_i := P_j - 1; r_i := r_j; c_i := c_j; d_i := 0;$ $\forall t_k \in q   P_k < P_j : r_k := r_k + C_i$
<b>Task <math>t_{run}</math> done and <math>q \neq \emptyset</math></b>	<b>High priority task release</b>
Guard $c_{run} = r_{run} \wedge d_{run} \leq D_{run} \wedge t_i \in q \wedge$ $t_i \neq t_{run} \wedge P_i = P_{run} - 1$	Guard $t_{next}^{EDF}(t_i) = \emptyset \wedge c_{run} < r_{run}$
Update $q := q \setminus \{t_{run}\}; t_{run} := t_i;$ $\forall t_j \in q : P_j := P_j + 1$	Sync $release_i?$
<b>Maintain <math>c_i \leq C_x</math></b>	Update $q := q \cup \{t_i\}; c_i := 0; d_i := 0; r_i := 0;$ $\forall t_j \in q \setminus \{t_i\} : P_j := P_j - 1; P_i := N;$ $\forall t_j \in q : r_j := r_j + C_i; t_{run} := t_i$
Guard $c_i = C_x \wedge t_i \in q$	
Update $c_i := 0; r_i := r_i - C_x$	

Figure 4: Overview of the encoding  $E(\text{Sch})$ .

this representation into a pure TAU. The exception to this approach is the function  $t_{next}^{EDF}()$  that will be addressed later.

The scheduler consists of three locations: Idle, Busy, and Error. The edges are classes of edges that are instantiated by iterating the variable  $t_i$  over the set of tasks. Task identifiers such as  $t_i$  and  $i$  are used interchangeably to reduce the maximum subscript level.

Since a task can be in the queue or not, the queue is encoded as a set  $q$ . Tasks themselves are represented via a number of variables:  $t_i$  represents the  $i$ -th task,  $t_{run}$  keeps track of the currently running task,  $c_i$  represents task computation clock explained in Subsection 3.4,  $r_i$  contains the current response time of the task, and  $c_i$  is compared to  $r_i$  to evaluate if the task has completed its execution;  $d_i$  is a clock that is reset when a task is released, and is compared to the natural  $D_i$  to check if the task's deadline has passed,  $P_i$  is the current priority of the task. The priority  $N$ , equal to the number of tasks in the system, is the highest priority and it corresponds to the currently

executing task.

The scheduler starts in the location *Idle*. This location corresponds to an empty task queue and it will be reentered on any occasion when there are no tasks left in the queue.

The edge going out of the location *Idle* is *First task release*. This edge is taken whenever the encoding of the adaptive task automaton synchronizes on *release<sub>i</sub>* channel without any additional constraints. Consequently, the task  $t_i$  is added to the queue, the currently running task is set to  $t_i$ , the response time is set to the computation time, the deadline clock is reset, and the task is assigned the highest priority.

In *Busy* location, there are four edges looping in the state, one returning to *Idle* and one leading to *Error* location. The invariant on *Busy* location, shown in dashed rectangle in [Figure 4](#), ensures that, in the *Busy* location, the currently running task will not execute longer than its computation time, and that all of the tasks in the system have not missed their deadlines.

In case that a deadline is missed, the edge *Deadline miss* is taken. The deadline is considered missed when the task is in the queue, still has some execution left, and has reached or exceeded its deadline. In such case, the system enters the *Error* location and deadlocks.

To explain the looping edges on the *Busy* location, let us first define the selector  $t_{next}^{EDF}()$ .

**Definition 7** The selector  $t_{next}^{EDF}(t_i) = t_j$  selects the task  $t_j$  that has the next higher priority in the queue relative to the task  $t_i$ , regardless of whether the task  $t_i$  is in the queue or not according to the deterministic EDF policy ([Definition 6](#)).

The selector returns the empty set if it is invoked for the highest priority task in the queue or any not-yet-released task that would become the highest priority task if it were released.  $\square$

Due to the nature of the EDF algorithm, a pure TAU implementation of this selector requires replication of any edge annotated with this selector into several edges. For the current permutation of tasks in the queue (implied by the current pure TAU location, and expressed via  $P_i$  and  $q$  variables in the representation), edges are created to test whether the new task will fit into any of the given possible positions in the queue. During verification, due to the determinism outlined in [Definition 7](#), only one of those edges will be enabled at any time.

The edge *High priority task release* employs this selector to check if the newly released task has higher priority than any of the tasks in the queue. The edge guard also checks whether the currently running task is still running. This check ensures that whenever a task completes it is removed from the queue before any further actions are taken. Since the newly released task has a higher priority than any other task in the queue, its response time is equal to its computation time. All of the other tasks' response times need to be increased by the computation time of the newly released task. Priorities of other tasks are reduced and the newly released task acquires the highest priority.

On the other hand, when the newly released task has lower priority than the currently running task, it needs to be placed at the correct place in the queue via *Low priority task release* edge. This is where the determinism of our EDF implementation via  $t_{next}^{EDF}$  selector comes into play. We need to ensure that the tasks added to the queue via this edge will be executed in the same sequence as they are added to the queue. Otherwise, the computed response times would be invalidated. As with the previous edge, we add the task to the queue, but this time we need to copy the response time and computation clock from the higher priority task. Then, we increase the response time of

the new task, as well as any of lower priority, with the computation time of the released task.

As the time passes in Busy state, tasks are executing and will be removed from the queue when they complete, by one of the Task  $t_{run}$  done edges. The edge Task  $t_{run}$  done and  $q \neq \emptyset$  is taken if the task has completed its execution before the deadline and if the next task is present in the queue. To switch the currently running task, the latter is taken out of the queue, a new task is set to currently running task and all of the active tasks' priorities are increased by one to keep priority values bound between 1 and  $N$ .

If the task is the last task in the queue, the edge Task  $t_{run}$  done and  $q = \emptyset$  is enabled and removes the task from the queue while moving the automaton into the Idle location.

To keep all clocks and response times bound the edge Maintain  $c_i \leq C_x$ , resets the clock  $c_i$  to 0 every time an active clock reaches the maximum clock value  $C_x$  and the corresponding response time is decreased by  $C_x$ . While this edge alters the value of clocks, it does not influence the relevant difference  $r_i - c_i$ . This mechanism resolves the potential unboundedness of the system caused by the inheritance of  $c_i$  and  $r_i$  values in Low priority task release. Without it, any system that repeatedly releases tasks of lower priority than the currently running task can become unbounded.

## 4 Decidability

The decidability of schedulability verification for our model depends on two things: decidability of reachability for our variant of timed automata with updates (Subsection 4.1) and that the encoding of  $ATA_{EDF}$  model into timed automata with updates represents the original model correctly (Subsection 4.2).

### 4.1 Decidability of Timed Automata with Updates

Alur and Dill [AD94] observe that we can partition the state space of a timed automaton into a finite number of discrete regions that can be exhaustively explored in a finite amount of time. Hence, the location reachability problem is decidable.

Our refined region equivalence relation is based on the relation given in [AD94] and extended by the region equivalence relation for timed automata with diagonal constraints presented by Bengtsson and Yi [BY04], and Fersman et al. [FKPY07].

**Definition 8** (Refined region equivalence  $\approx$  [FKPY07, AD94, BY04]) For a clock  $x \in X$ , let  $C_x$  be a natural number. For a positive real number  $t$ , let  $\{t\}$  denote the fractional part of  $t$ , and  $\lfloor t \rfloor$  its integer part. Let  $u, v \in \mathcal{V}$  be two regions,  $\mathcal{G}$  a finite set of diagonal constraints in the form  $x - y \bowtie \mathbb{Z}_{\geq 0}$  where  $\mathbb{Z}_{\geq 0}$  is the set of non-negative integers, and  $\bowtie \in \{<, \leq, =, \geq, >\}$ .

We define  $u \approx v$ , i.e.  $u$  and  $v$  are refined-region-equivalent iff

1. for each clock  $x$ , either  $\lfloor u(x) \rfloor = \lfloor v(x) \rfloor$  or  $u(x) > C_x$  and  $v(x) > C_x$ ,
2. for each clock  $x$ , if  $u(x) \leq C_x$ , then  $\{u(x)\} = 0$  iff  $\{v(x)\} = 0$ ,
3. for all clocks  $x, y$ , if  $u(x) \leq C_x$  and  $u(y) \leq C_x$  then  $\{u(x)\} \leq \{u(y)\}$  iff  $\{v(x)\} \leq \{v(y)\}$ , and
4.  $u \models g$  iff  $v \models g$  for all  $g \in \mathcal{G}$ . □

Given [Definition 8](#) of refined region equivalence, we can postulate that operations over regions will not disrupt the refined region equivalence relationship on TAU.

**Lemma 1** *Given a timed automaton with updates, let  $\mathcal{G}$  denote the set of diagonal constraints in the automaton and  $C_x$  be the maximum of  $M_x$  (the ceiling of  $x$ ) and all constants appearing in the guards and invariants of the automaton involving clock  $x$ . Let  $u, v \in \mathcal{V}$  and  $t, t' \in \mathbf{R}_{\geq 0}$ . Then  $u \approx v$  implies*

1.  $u + t \approx v + t'$  for some real number  $t'$  such that  $\lfloor t \rfloor = \lfloor t' \rfloor$ ,
2.  $u[x \mapsto 0] \approx v[x \mapsto 0]$  for a clock  $x$ , and
3.  $u[x \mapsto y] \approx v[x \mapsto y]$  for all pairs of clocks  $x$  and  $y$ .

*Proof Outline.* [Lemma 1](#) can be trivially proven for the case when only one clock is assigned a new value. The case with multiple clocks being assigned new values can be proven by observing that we can reduce the problem to relative ordering of fractional parts of clocks which are consistent for all clocks between regions based on the third criterion of [Definition 8](#). The full proof is given in [[HDSP14](#)].  $\square$

**Lemma 2** (*Bisimulation of TAU*) *Let us assume a timed automaton with updates, a location  $l$  and clock assignments  $u$  and  $v$ . Then  $u \approx v$  implies that:*

1. when  $(l, u) \rightarrow (l', u')$  then  $(l, v) \rightarrow (l', v')$  for some  $v'$  such that  $u' \approx v'$ , and
2. when  $(l, v) \rightarrow (l', v')$  then  $(l, u) \rightarrow (l', u')$  for some  $u'$  such that  $u' \approx v'$ .

*Proof Outline.* The proof follows from [Lemma 1](#). Assume a location  $l$  and clock assignments  $u$ , and  $v$ , such that  $u \approx v$ . The refined region equivalence relation  $\approx$  defines that the guards will evaluate in both  $u$  and  $v$  to the same truth values. Therefore, the set of enabled transitions is equal in both valuations.  $\square$

**Lemma 3** (*Location Reachability*) *The location reachability problem for timed automata with updates and invariants is decidable if the bound  $M_x$  for each clock  $x$  is known.*

*Proof.* [Lemma 1](#) shows that for each location  $l$  of the automaton, there is a finite number of equivalence classes derived from the bisimulation relation  $\approx$ . Since the number of locations is finite, the entire state space of an automaton can be partitioned into a finite number of equivalence classes and these equivalence classes can be effectively generated and searched.  $\square$

## 4.2 Model Bisimulation

Once we have encoded the entire  $\text{ATA}_{\text{EDF}}$  system as a network of TAU, we need to show that there exists a bisimulation between the original model and the encoding.

Our main result is described by [Lemma 4](#) below, for which we outline the proof. In [Definition 9](#), we first introduce the concept of schedulability as reachability.

**Definition 9** (Schedulability) *The adaptive task automaton  $A$  with initial state  $(l_0, u_0, q_0)$  and scheduling strategy  $\text{Sch}$  is not schedulable iff there exists a trace  $(l_0, u_0, q_0)(\xrightarrow{\text{Sch}})^*(l', u', q')$*

such that in the state  $(l', u', q')$  there is a task  $t_i$  with more than zero computation time left,  $c_i > 0$ , and no more time to execute, that is  $d_i \leq 0$ . The state  $(l', u', q')$  is marked as  $(l', u', \text{Error})$ .  $\square$

**Lemma 4** *Let  $A$  be an adaptive task automaton and  $Sch$  the EDF scheduling strategy presented in Definition 6. Assume that  $(l_0, u_0, q_0)$  and  $(\langle l_0, \text{Idle} \rangle, u_0 \cup v_0)$  are the initial states of  $A$ , and the product automaton  $E(A) \parallel E(Sch)$ , respectively, where  $l_0$  is the initial location of  $A$ ,  $u_0$  and  $v_0$  are clock assignments assigning all clocks with 0, and  $q_0$  is the empty task queue. Then:*

*For all  $l$  and  $u$ :  $(l_0, u_0, q_0) \rightarrow^* (l, u, \text{Error})$  implies  
 $(\langle l_0, \text{Idle} \rangle, u_0 \cup v_0) \rightarrow^* (\langle l, \text{Error} \rangle, u \cup v)$  for some  $v$ .*

*For all  $l$ ,  $u$ , and  $v$ :  $(\langle l_0, \text{Idle} \rangle, u_0 \cup v_0) \rightarrow^* (\langle l, \text{Error} \rangle, u \cup v)$  implies  
 $(l_0, u_0, q_0) \rightarrow^* (l, u, \text{Error})$ .*

*Proof Outline.* The encoding of the  $\text{ATA}_{\text{EDF}}$  automaton differs from the automaton itself in two key aspects: invocations of tasks and adaptivity predicates.

Since the task releases in  $\text{ATA}_{\text{EDF}}$  can be said to be of a non-blocking nature, we essentially verify that the scheduler will be non-blocking as well. Indeed, the only situation when there are no enabled edges annotated with a synchronization channel is when the scheduler automaton enters the Error state.

Our encoding exposes the values required for the evaluation of the encoded schedulability predicates directly. In order to check that the  $\text{ATA}_{\text{EDF}}$  adaptivity predicates are going to evaluate to the same results, we establish a correlation between tasks' parameters in  $\text{ATA}_{\text{EDF}}$  and the encoding. Once these mappings have been properly established, we check edge-by-edge that they are maintained. The full proof is given in [HDSP14].  $\square$

Since we have proven that the reachability problem is decidable for TAU, stated by Lemma 3, also that every  $\text{ATA}_{\text{EDF}}$  can be translated into a bisimilar TAU, we can conclude that the problem of checking schedulability of  $\text{ATA}_{\text{EDF}}$  is decidable as well.

## 5 Related Work

Our work tries to unify schedulability analysis with modeling and analysis of adaptive embedded systems. At the same time, a number of works address problems in those two separate fields, as well as non-modeling methods for analysis of schedulability in adaptive contexts. While this is by no means an exhaustive list of the works in these areas, we will try to list those that are closest to ours.

In the following works, verification of adaptive embedded systems is done on a more coarse scale than in our approach. Most of these approaches could be used in synergy with ours to provide system level verification, while ours provides task level granularity. Adler et al. [ASSV07] use Kripke structures as the underlying presentation of the system and specify the system's properties using LTL. Schneider et al. [SST06] have proposed a method to describe and analyze adaptation behavior in embedded systems in which the data flow is augmented with quality descriptions used by configuration rules to determine potential adaptations. Goldsby et al. [GCZ08] provide the AMOEBA-RT model focused on run-time verification and monitoring.

In the area of adaptive scheduling, most work [JPG04, LRK03] was done to achieve a lower

energy consumption by exploiting dynamic voltage scaling features of modern CPUs. While such approaches can be used to analyze schedulability in some adaptive contexts, our approach makes it possible to model and analyze more precisely task release patterns of non-periodic tasks.

Finally, other works have approached verification of schedulability by means of timed automata for uniprocessors [DILS09, MLR<sup>+</sup>10], and multiprocessors [YLX10] without explicit inclusion of adaptive functionality.

## 6 Conclusion

In this work, we have shown that the verification of adaptive task automata with earliest-deadline-first scheduling policy is decidable. To support our claim, we have encoded our adaptive task automata model as timed automata with updates and presented that the model and its encoding are bisimilar, as well as given a proof that reachability in our variant of timed automata with updates is decidable.

Our main result is the proof of decidability of our ATA extensions. Using ATA, it is possible to model the environment of an embedded system as well as behavior of functional and extra-functional properties in response to internal or environmental changes. Thus we verify the behavior of specified properties throughout the execution of the system.

In this work, we have implemented the EDF scheduling policy. However, by replacing the selector  $t_{next}^{EDF}()$ , we can implement any other policy that is deterministic and does not change relative task priorities after their release into the queue. A non-deterministic selector would invalidate the schedulability testing predicates ( $sched()$ ) since the response times predicted when testing a task would not necessarily correspond to the actual response times after the task is released.

During the encoding, we have faced a number of challenges. To support dynamic scheduling policies and schedulability predicates, we have required dynamic construction of task response times, which, in turn, have required a clock copying mechanism that had to be added as an extension of timed automata.

As future work, we plan to further explore removal of the assumptions, specifically extend the framework to support modeling of multi-core systems, smart handling of tasks with variable execution time, shared resources, as well as create a set of templates that correctly model the most commonly utilized task release patterns.

**Acknowledgements:** This research has been supported by the Swedish Research Council, which is gratefully acknowledged.

## Bibliography

- [AD94] R. Alur, D. L. Dill. A theory of timed automata. *Theoretical Computer Science* 126:183–235, April 1994.  
[doi:10.1016/0304-3975\(94\)90010-8](https://doi.org/10.1016/0304-3975(94)90010-8)

- [Alu99] R. Alur. Timed Automata. In Halbwachs and Peled (eds.), *Computer Aided Verification*. Lecture Notes in Computer Science 1633, pp. 8–22. Springer Berlin Heidelberg, 1999.  
[doi:10.1007/3-540-48683-6\\_3](https://doi.org/10.1007/3-540-48683-6_3)
- [ASSV07] R. Adler, I. Schaefer, T. Schuele, E. Vecchié. From Model-Based Design to Formal Verification of Adaptive Embedded Systems. In Butler et al. (eds.), *Formal Methods and Software Engineering*. Lecture Notes in Computer Science 4789, pp. 76–95. Springer Berlin Heidelberg, 2007.  
[doi:10.1007/978-3-540-76650-6\\_6](https://doi.org/10.1007/978-3-540-76650-6_6)
- [BDFP04] P. Bouyer, C. Dufourd, E. Fleury, A. Petit. Updatable timed automata. *Theoretical Computer Science* 321(23):291 – 345, 2004.  
[doi:10.1016/j.tcs.2004.04.003](https://doi.org/10.1016/j.tcs.2004.04.003)
- [BGK<sup>+</sup>02] J. Bengtsson, W. D. Griffioen, K. J. Kristoffersen, K. G. Larsen, F. Larsson, P. Pettersson, W. Yi. Automated verification of an audio-control protocol using Uppaal. *The Journal of Logic and Algebraic Programming* 52 – 53(0):163 – 181, 2002.  
[doi:10.1016/S1567-8326\(02\)00036-X](https://doi.org/10.1016/S1567-8326(02)00036-X)
- [BY04] J. Bengtsson, W. Yi. Timed Automata: Semantics, Algorithms and Tools. In Desel et al. (eds.), *Lectures on Concurrency and Petri Nets*. Lecture Notes in Computer Science 3098, pp. 87–124. Springer Berlin Heidelberg, 2004.  
[doi:10.1007/978-3-540-27755-2\\_3](https://doi.org/10.1007/978-3-540-27755-2_3)
- [DILS09] A. David, J. Illum, K. Larsen, A. Skou. *Model-Based Framework for Schedulability Analysis Using UPPAAL 4.1*. Pp. 93–119. CRC Press, 2011/12/27 2009.  
[doi:10.1201/9781420067859-c4](https://doi.org/10.1201/9781420067859-c4)
- [FKPY07] E. Fersman, P. Krcal, P. Pettersson, W. Yi. Task automata: Schedulability, decidability and undecidability. *Information and Computation* 205(8):1149 – 1172, 2007.  
[doi:10.1016/j.ic.2007.01.009](https://doi.org/10.1016/j.ic.2007.01.009)
- [GCZ08] H. J. Goldsby, B. H. Cheng, J. Zhang. AMOEBA-RT: Run-Time Verification of Adaptive Software. In Giese (ed.), *Models in Software Engineering*. Lecture Notes in Computer Science 5002, pp. 212–224. Springer Berlin Heidelberg, 2008.  
[doi:10.1007/978-3-540-69073-3\\_23](https://doi.org/10.1007/978-3-540-69073-3_23)
- [GD99] J. Goossens, R. Devillers. Feasibility intervals for the deadline driven scheduler with arbitrary deadlines. In *Real-Time Computing Systems and Applications, 1999. RTCSA '99. Sixth International Conference on*. Pp. 54 –61. 1999.  
[doi:10.1109/RTCSA.1999.811193](https://doi.org/10.1109/RTCSA.1999.811193)
- [HDSP14] L. Hatvani, A. David, C. Seceleanu, P. Pettersson. Adaptive Task Automata with Earliest-Deadline-First Scheduling. Technical report ISSN 1404-3041 ISRN MDH-MRTC-287/2014-1-SE, Mälardalen Real-Time Research Centre, Mälardalen University, August 2014.  
<http://www.es.mdh.se/publications/3661->

- [HPS12] L. Hatvani, P. Pettersson, C. Seceleanu. Adaptive Task Automata: A Framework for Verifying Adaptive Embedded Systems. In Lara and Zisman (eds.), *Fundamental Approaches to Software Engineering*. Lecture Notes in Computer Science 7212, pp. 115–129. Springer Berlin Heidelberg, 2012.  
[doi:10.1007/978-3-642-28872-2\\_9](https://doi.org/10.1007/978-3-642-28872-2_9)
- [JPG04] R. Jejurikar, C. Pereira, R. Gupta. Leakage Aware Dynamic Voltage Scaling for Real-time Embedded Systems. In *Proceedings of the 41st Annual Design Automation Conference*. DAC '04, pp. 275–280. ACM, New York, NY, USA, 2004.  
[doi:10.1145/996566.996650](https://doi.org/10.1145/996566.996650)
- [LBB<sup>+</sup>01] K. Larsen, G. Behrmann, E. Brinksma, A. Fehnker, T. Hune, P. Pettersson, J. Romijn. As Cheap as Possible: Efficient Cost-Optimal Reachability for Priced Timed Automata. In Berry et al. (eds.), *Computer Aided Verification*. Lecture Notes in Computer Science 2102, pp. 493–505. Springer Berlin Heidelberg, 2001.  
[doi:10.1007/3-540-44585-4\\_47](https://doi.org/10.1007/3-540-44585-4_47)
- [LRK03] Y.-H. Lee, K. Reddy, C. Krishna. Scheduling techniques for reducing leakage power in hard real-time systems. In *Real-Time Systems, 2003. Proceedings. 15th Euromicro Conference on*. Pp. 105–112. July 2003.  
[doi:10.1109/EMRTS.2003.1212733](https://doi.org/10.1109/EMRTS.2003.1212733)
- [MLR<sup>+</sup>10] M. Mikučionis, K. Larsen, J. Rasmussen, B. Nielsen, A. Skou, S. Palm, J. Pedersen, P. Hougaard. Schedulability Analysis Using Uppaal: Herschel-Planck Case Study. In Margaria and Steffen (eds.), *Leveraging Applications of Formal Methods, Verification, and Validation*. Lecture Notes in Computer Science 6416, pp. 175–190. Springer Berlin / Heidelberg, 2010.  
[doi:10.1007/978-3-642-16561-0\\_21](https://doi.org/10.1007/978-3-642-16561-0_21)
- [SST06] K. Schneider, T. Schuele, M. Trapp. Verifying the adaptation behavior of embedded systems. In *Proceedings of the 2006 international workshop on Self-adaptation and self-managing systems*. SEAMS '06, pp. 16–22. ACM, New York, NY, USA, 2006.  
[doi:10.1145/1137677.1137681](https://doi.org/10.1145/1137677.1137681)
- [Y LX10] F. Yu, G. Li, N. Xiong. Schedulability analysis of multi-processor real-time systems using Uppaal. In *Information Science and Engineering (ICISE), 2010 2nd International Conference on*. Pp. 1–6. dec. 2010.  
[doi:10.1109/ICISE.2010.5689944](https://doi.org/10.1109/ICISE.2010.5689944)