

Multi-Processor Schedulability Analysis of Preemptive Real-Time Tasks with Variable Execution Times

Pavel Krcal¹, Martin Stigge², and Wang Yi¹

¹ Uppsala University, Sweden

Email: {pavelk,yi}@it.uu.se

² Humboldt University Berlin, Germany

Email: mstigge@informatik.hu-berlin.de

Abstract In this paper, we study schedulability analysis problems for multi-processor real-time systems. Assume a set of real-time tasks whose execution times and deadlines are known. We use timed automata to describe the non-deterministic arrival times of tasks. The schedulability problem is to check whether the released task instances can be executed within their given deadlines on a multi-processor platform where each processor has a task queue to buffer task instances scheduled to run on the processor. On the positive side, we show that the problem is decidable for systems with non-preemptive schedulers or tasks with fixed execution times. A surprising negative result is that for multi-processor systems with variable task execution times and a non-preemptive scheduler, the schedulability analysis problem is undecidable, which is still an open problem in the single-processor setting.

1 Introduction

Real-time systems are often designed as a collection of real-time tasks and a scheduling strategy implemented as a scheduler. Each of the tasks may have a set of task parameters such as release rate (or release pattern), best and worst-case execution times on the target hardware, priority, and deadline etc. In the operation of a system, the tasks will be released according to the pre-designed release rates, and the released task instances are scheduled to execute on a processor by the scheduler. Here the scheduler, namely the scheduling strategy, is the critical component for the correct functioning of a system. It makes the decision about in which order the released task instances should be executed, based on the task parameters and the current system state.

An important design problem is to analyze whether all the task instances can be executed within the given deadlines, which is essentially to estimate the worst-case response times of the tasks. This is the so-called schedulability analysis based on (1) the task release patterns, (2) the task parameters and (3) the scheduling policy, all of which are from the system design phase. The source of complexity in solving the analysis problem is in dealing with task releases

and preemptions. A newly released task instance may preempt a running task, which influences the response time for the preempted task. Classic solutions, e.g., Rate-Monotonic Analysis often assume deterministic task release patterns such as periodic tasks [LL73] or deterministic patterns with fixed type of non-determinism such as jitters [But97] and offsets [RC01]. A challenge is to solve the schedulability analysis problem for systems with dynamic and non-deterministic task releases and preemptions. In recent years, in a series of work, we have used timed automata [AD94] to model task release patterns and solved the problem for a large class of single-processor systems [FKPY07,FMPY06].

In [FKPY07], timed automata are extended with asynchronous tasks. Each location of a timed automaton may be associated with a task. As soon as the automaton visits a location, an instance of the associated task is released and put (scheduled) into a task queue for execution. Compared with classic task models studied in the literature on scheduling, extended timed automata provide a much more expressive model which, in fact, inherits the full expressive power of timed automata for modeling of dynamic and non-deterministic task releases and preemptions. In our previous work, the classic notion of schedulability analysis has been extended to the automata model, and it is shown that for a large class of systems, the problem can be solved automatically using algorithmic methods. However, the study has been restricted to the single-processor setting.

In this paper, we study the schedulability analysis problem of the extended automata model in the multi-processor setting. Basic scheduling algorithms for multi-processor systems can be found in, e.g., [Liu00]. For recent work on schedulability analysis for multi-processor systems and further references, see [ABJ01,BCL05]. We assume that a system has a fixed number of processors available. Each processor is associated a task queue, where the released tasks wait to be processed. A scheduling policy (modeled as a function) decides for a newly released task instance into which queue and at which position in the queue it will be inserted. However, task migration is not allowed, that is, once a task instance is assigned to a processor, it will remain in the associated queue until it finishes.

On the positive side, we show that the problem is decidable as for the single-processor case if

1. the scheduler runs a non-preemptive scheduling strategy, or
2. the tasks have fixed execution times, that is, the best and worst-case execution times coincide.

It is an open problem, whether for the class of systems with variable task execution times and a preemptive scheduler, the schedulability problem is decidable in the single-processor setting. This problem becomes undecidable when the scheduling policy has at least two processors with their task queues available. More precisely as a main technical result of this paper, we show that the schedulability problem will be undecidable for multi-processor systems if

1. the scheduler runs a preemptive scheduling strategy, and
2. the tasks have variable execution times ranging over an interval between the best and worst-case execution times.

2 Preliminaries

In this section, we introduce the concept of *task automata* (timed automata extended with tasks) developed in [FKPY07], and the multi-processor schedulability problem.

2.1 Task Automata

Let \mathcal{C} be a finite set of clocks. A function $\nu : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ is called a *clock valuation* while the set of all clock valuations over the clocks \mathcal{C} is denoted by $\mathcal{V}(\mathcal{C})$. With $\nu[r]$ we denote the clock valuation which is equal to ν , except that all clocks in $r \subseteq \mathcal{C}$ are reset to zero. $\mathcal{B}(\mathcal{C})$ is the set of *clock guards* g , which are conjunctions of expressions of the form $x_1 \bowtie N$ and $x_1 - x_2 \bowtie N$ with $x_1, x_2 \in \mathcal{C}$, $N \in \mathbb{N}_{\geq 0}$ and $\bowtie \in \{<, =, >\}$. If g contains only $x_1 \bowtie N$ expressions, we say it is *diagonal-free*. A valuation ν satisfies a guard g (written $\nu \models g$) if for all expressions $x_1 \bowtie N$ and $x_1 - x_2 \bowtie N$ in g it holds that $\nu(x_1) \bowtie N$ and $\nu(x_1) - \nu(x_2) \bowtie N$, respectively.

Tasks and task queues. We define a task type as a tuple (P, B, W, D) written $P(B, W, D)$ where P is the task name (unique for each task type), $B, W \in \mathbb{N}_{\geq 0}$ the best and worst case calculation times (with $B \leq W$ and $W \geq 1$) and $D \in \mathbb{N}_{\geq 1}$ the relative deadline. A task instance $P_i(b_i, w_i, d_i)$ of type $P_i(B_i, W_i, D_i)$ is a released copy of this task type with $b_i, w_i, d_i \in \mathbb{R}$ being the remaining values. A task queue q is a list $[P_1(b_1, w_1, d_1), \dots, P_n(b_n, w_n, d_n)]$ of task instances (of possibly the same type). By a discrete part of a queue $[P_1(b_1, w_1, d_1), \dots, P_n(b_n, w_n, d_n)]$ we mean the list of the corresponding task names $[P_1, \dots, P_n]$ (the information about the remaining computation times and deadline is projected out). Let \mathcal{P} be the finite set of task types and $Q_{\mathcal{P}}$ be the set of queues over this task type set. We use a function $\text{Run} : Q_{\mathcal{P}} \times \mathbb{R}_{\geq 0} \rightarrow Q_{\mathcal{P}}$ which given a non-negative real number t and a task queue q returns the task queue after t time units of execution on a processor. The result of $\text{Run}(q, t)$ for $q = [P(b_1, w_1, d_1), Q(b_2, w_2, d_2), \dots, R(b_n, w_n, d_n)]$ and $t \leq w_1$ is defined as $q' = [P(b_1 - t, w_1 - t, d_1 - t), Q(b_2, w_2, d_2 - t), \dots, R(b_n, w_n, d_n - t)]$. For an empty queue, denoted by $[], \text{Run}([], t) = []$.

Definition 1. A task automaton over actions \mathcal{Act} , clocks \mathcal{C} , and task types \mathcal{P} is a tuple $\langle N, l_0, E, I, M, x_{done} \rangle$ where

- N is a finite set of locations,
- $l_0 \in N$ is the initial location,
- $E \subseteq N \times \mathcal{B}(\mathcal{C}) \times \mathcal{Act} \times 2^{\mathcal{C}} \times N$ is the set of edges,
- $I : N \rightarrow \mathcal{B}(\mathcal{C})$ is a function assigning a clock constraint to each location which is called location invariant,
- $M : N \hookrightarrow \mathcal{P}$ is a partial function assigning locations with task types,¹ and
- $x_{done} \in \mathcal{C}$ is the clock which is reset whenever a task finishes.

We write $l \xrightarrow{g a r} l'$ for $\langle l, g, a, r, l' \rangle \in E$.

¹ Note that M is a partial function meaning that some of the locations may have no tasks.

Semantics. An important part of the operational semantics is the scheduling function, which decides, into which queue and at which position a newly released task should be inserted. For the sake of presentation, we first introduce the scheduling function $\text{Sch} : \mathcal{P} \times Q_{\mathcal{P}} \rightarrow Q_{\mathcal{P}}$ for the single-processor case. Given a task instance and a task queue, it returns the task queue with the task instance inserted and the order of the other task instances preserved. Depending on whether the scheduler is preemptive or non-preemptive, the function may insert new tasks as the first element or not. Since we want this function to be encodable in timed automata, the definition has the following important condition.

Definition 2. $\text{Sch} : \mathcal{P} \times Q_{\mathcal{P}} \rightarrow Q_{\mathcal{P}}$ is a scheduling function, if for each task type $P(B, W, D)$ and discrete part $[P_1, \dots, P_n]$ of a queue there can be effectively constructed a diagonal-free timed automaton with

- Clocks $y_1^b, y_1^w, y_1^d, \dots, y_n^b, y_n^w, y_n^d$,
- $n + 2$ locations l_0, l_1, \dots, l_{n+1} , and
- $n + 1$ edges from l_0 to l_i for $1 \leq i \leq n + 1$,

such that $\text{Sch}(P(B, W, D), [P_1(b_1, w_1, d_1), \dots, P_n(b_n, w_n, d_n)])$ inserts $P(B, W, D)$ into the queue at the k -th position if and only if l_k is the only location reachable from (l_0, u) where $u(y_i^b) = b_i, u(y_i^w) = w_i, u(y_i^d) = d_i$ for all $1 \leq i \leq n$.

This definition generalizes the notion of a scheduling policy. Most of the standard scheduling policies such as EDF (Earliest Deadline First), FPS (Fixed Priority Scheduling), FIFO, etc. satisfy this condition, but also ad hoc policies such as combinations of the standard ones are also included. Therefore, the results hold for a very general class of scheduling policies.

In the multi-processor case, the scheduling function takes k queues and a task type as an input and returns the queues with the new task instance inserted at some position in one of the queues. It can use the information from all the queues for its decision. Each timed automaton corresponding to a scheduling policy then contains clocks for all the instances in all the queues. To simplify the notation, we assume that for k queues q_1, \dots, q_k with discrete parts $\tilde{q}_1, \dots, \tilde{q}_k$, all the task instances are indexed by one index i ranging from 1 to $\sum_{1 \leq j \leq k} |\tilde{q}_j|$, where $|\tilde{q}_j|$ denotes the number of the task instances in the queue q_j . E.g., the index of the first task instance in q_2 is $|\tilde{q}_1| + 1$. Also, all possible positions where a new task can be inserted are indexed by one index ranging from 1 to $\sum_{1 \leq j \leq k} |\tilde{q}_j| + k$. E.g., the $|\tilde{q}_1| + 2$ -th position denotes the first (head) position of the second queue (the first task instance in q_2 after insertion). By b_i, w_i and d_i we denote as before the continuous queue information for the task instance p_i .

Definition 3. (*Multi-processor Scheduler*) Let $k \in \mathbb{N}$ be the number of processors. Then $\text{Sch}_k : \mathcal{P} \times (Q_{\mathcal{P}})^k \rightarrow (Q_{\mathcal{P}})^k$ is a multi-processor scheduling function, if for each task type $P(B, W, D)$ and discrete parts of queues $\tilde{q}_1, \dots, \tilde{q}_k$, there can be effectively constructed a diagonal-free timed automaton with

- Clocks $y_1^b, y_1^w, y_1^d, \dots, y_K^b, y_K^w, y_K^d$ where $K = \sum_{1 \leq j \leq k} |\tilde{q}_j|$,

- $K + k + 1$ locations l_0, l_1, \dots, l_{K+k} and
- $K + k$ edges from l_0 to l_j for $1 \leq j \leq K + k$,

such that the function $\text{Sch}(P(B, W, D), q_1, \dots, q_k)$ inserts $P(B, W, D)$ at the m -th position if and only if l_m is the only location reachable from (l_0, u) where $u(y_i^b) = b_i, u(y_i^w) = w_i, u(y_i^d) = d_i$ for all $1 \leq i \leq K + k$.

Note that this definition does not allow for moving tasks between processors after they got assigned to one processor at the moment when they are released, that is, we do not allow task migration.

A task automaton may – just as a timed automaton – perform event and delay transitions, and additionally *task finishing* transitions. An event transition corresponds to the arrival of a new task, and a delay transition corresponds to active tasks being executed while the others are waiting, or just processor idling in case some queue is empty. Transitions of the third type remove finished tasks from the queues and reset the clock x_{done} , giving thus a feedback to the automaton (x_{done} can be checked in the guards). We give now the formal definition of the operational semantics for a system with k processors as a labeled transition system (LTS), where $S := N \times \mathcal{V}(\mathcal{C}) \times (Q_{\mathcal{P}})^k$ is the state space, thus incorporating the queues into the state information. Let ν_0 be a clock valuation assigning all clocks the value 0, and $\Sigma := \text{Act} \cup \mathbb{R}_{\geq 0} \cup \{\text{fin}\}$ a set of labels (for events, time pass values and task finishing).

Definition 4. *Given a scheduling strategy Sch_k over k processors, the semantics of a task automaton $A = \langle N, l_0, E, I, M, x_{done} \rangle$ is a labeled transition system $\llbracket A_{\text{Sch}_k} \rrbracket = \langle S, s_0, \Sigma, T \rangle$ with $s_0 = (l_0, \nu_0, \square)$ and T the set of transitions defined by the following rules:*

- $(l, \nu, q_1, \dots, q_k) \xrightarrow{a}_{\text{Sch}_k} (l', \nu[r], \text{Sch}_k(M(l'), q_1, \dots, q_k))$ if $l \xrightarrow{g \text{ aT}} l', \nu \models g$, and $\nu[r] \models I(l')$,
- $(l, \nu, q_1, \dots, q_k) \xrightarrow{t}_{\text{Sch}_k} (l, \nu+t, \text{Run}(q_1, t), \dots, \text{Run}(q_k, t))$ if $t \in \mathbb{R}_{\geq 0}, (\nu+t) \models I(l)$, and for all i with $q_i = P(b, w, d) :: q'_i$ it holds that $t \leq w$, and
- $(l, \nu, q_1, \dots, P(b, w, d) :: q_i, \dots, q_k) \xrightarrow{\text{fin}}_{\text{Sch}_k} (l, \nu[x_{done}], q_1, \dots, q_i, \dots, q_k)$ if $b \leq 0 \leq w$ and $\nu[x_{done}] \models I(l)$,

where $P(b, w, d) :: q$ denotes the queue with the task instance $P(b, w, d)$ on the first position and q being the (possibly empty) tail.

Finally, we can also define the main question this work deals with, namely whether a task automaton models schedulable sequences of task releases. As all deadlines in our model are hard, we say that a task automaton is schedulable for a given scheduling strategy if no matter how does the (non-deterministic) computation evolve, all deadlines are met. We use q_{err} to denote queues containing a task instance $P(b, w, d)$ with $d < 0$.

Definition 5. (Schedulability) *A task automaton A with initial state $(l_0, u_0, \square, \dots, \square)$ is unschedulable with Sch_k if $(l_0, u_0, \square, \dots, \square) \xrightarrow{*}_{\text{Sch}_k} (l, u, q_1, \dots, q_{err}, \dots, q_k)$ for some l and u . Otherwise, we say that A is schedulable with Sch_k .*

We call a queue *unschedulable* if it will inevitably lead to a deadline miss provided that all tasks take their worst case computation times. Otherwise, a queue is said to be *schedulable*. The important observation for schedulable queues is that their length is bounded ([FKPY07]):

Lemma 1. *Given a finite task type set \mathcal{P} , one can effectively construct a natural number $B_{\mathcal{P}}$ such that $|q| \leq B_{\mathcal{P}}$ for all schedulable queues q .*

2.2 Decidability for Task Automata

A system, modeled by a task automaton and a scheduling function, can have the following three properties:

Preemption: The scheduler may insert a newly released task to the head of a non-empty queue, thus preempting all tasks already in the queue (a non-preemptive scheduler may insert tasks only at other positions).

Variable task execution times: There can be task types $P_i(B_i, W_i, D_i)$ such that $B_i < W_i$, meaning that the task may non-deterministically finish execution at a time within the interval $[B_i, W_i]$.

Feedback: The precise finishing time of a task may influence the new task releases (by means of using x_{done} in a guard).

Note that the more of these properties are dropped for a task automaton, the easier the schedulability analysis problem becomes.

In [FKPY07], the schedulability problem is studied for single-processor systems. It is shown that (even with only one processor) schedulability becomes undecidable if a task automaton has all three 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 remains, whether schedulability is decidable if x_{done} is not used in the guards for creating feedback. This last variant has also been proven decidable for certain types of schedulers in [FKPY07], but there is no result for the general case.

We study these questions for multi-processor scheduling. Since single-processor systems are just a special case of multi-processor systems, the negative result, namely that the schedulability becomes undecidable if a task automaton has all three properties, transfers to our setting. We show that the problem remains decidable for the following classes:

1. A non-preemptive scheduler (but possibly variable execution time and feedback) or
2. Fixed execution time tasks ($B = W$ for all task types, but possibly a preemptive scheduler and feedback).

On the other hand, we show that the schedulability analysis problem becomes undecidable for the third case:

3. No feedback (but possibly a preemptive scheduler and variable execution time).

3 Decidable Cases of Multi-Processor Scheduling

In this section, we show that the multi-processor schedulability problem is decidable for the first two cases. The proof is done by reduction of this problem to the reachability problem for timed automata. For a given number of processors, a task automaton, and a multi-processor scheduling policy, we construct a timed automaton with an error location such that the system is unschedulable if and only if the error location is reachable. Our construction is based on the construction for the single-processor case from [FKPY07] and [EWY99].

3.1 Fixed computation time

First, we handle the case where all tasks have fixed computation time (but the scheduler can be preemptive and there can be feedback from the scheduler back to the automaton).

Theorem 1. *The problem of checking whether a task automaton A with fixed computation times of tasks together with a multi-processor scheduler Sch_k is schedulable, is decidable.*

Proof (Sketch). The given task automaton A is transformed into a standard timed automaton $E(A)$ by taking the underlying timed automaton of A , removing the labels and adding new labels release_P on edges where A would release an instance of task type P . Because of Definition 2 and Lemma 1, the whole scheduling strategy for schedulable queues can be encoded as a timed automaton $E(\text{Sch}_k)$ (the *scheduler automaton*) as follows. The continuous part of the queues, namely the best/worst case computation times and remaining relative deadlines, are encoded in clocks. There are two clocks x_i^c and x_i^d for each task instance p_i in the queues, which measure how long this task instance has been computed and how long it has been released. The discrete parts of the queues (the order of the task instances) are encoded in locations of $E(\text{Sch}_k)$. Since the queue length of schedulable queues is bounded (Lemma 1) and the number of the queues is fixed, there are only finitely many locations needed. Once a queue becomes unschedulable, $E(\text{Sch}_k)$ will enter a dedicated error location. The edges and their guards correspond to the comparisons which the scheduling function uses for its decisions.

The values of b_i , w_i and d_i for each task instance p_i in the queue can be expressed using the clocks x_i^c and x_i^d , which is ensured by the construction in the following way. Whenever a new instance p_i of type P_j is released (event release_{P_j}), the clock x_i^d is reset. Before a task instance p_i is put to the head of its queue (which models the beginning of the task execution), the values b_i , w_i and d_i from the queue can be expressed as B_j , W_j and $D_j - x_i^d$, respectively. The clock x_i^c is not used at all.

As soon as the task instance p_i is scheduled for execution for the first time, the clock x_i^c is reset, keeping track of its computation progress. For a running task instance p_i , the values b_i , w_i and d_i from its queue can be expressed as $B_j - x_i^c$,

$W_j - x_i^c$ and $D_j - x_i^d$, respectively. Therefore, this task instance may finish whenever " $B_j \leq x_i^c \leq W_j$ " holds. To model task finishing, $E(\text{Sch}_k)$ removes p_i from its queue and resets x_{done} . Whenever a constraint " $x_i^d > D_j$ " is met for a released task instance p_i , an error location "unschedulable" is entered.

Because of the preemption, tasks waiting in the queues may have already been executed for some time, but they are stopped now. Since it is not possible to stop their x_i^c clocks, the time for which a preempted task p_i was already computed is represented not just by its x_i^c clock, but by a difference $x_i^c - x_j^c$. Here, p_j is the task which directly preempted p_i .

When a running task p_m finishes, all x_i^c of the preempted tasks p_i are updated by subtracting x_m^c , i.e., the computation time which was needed by p_m . In general, the reachability for timed automata with such updates is undecidable, but because of the fixed computation time property, $x_m^c = B_j = W_j$ is a constant natural number (assume that the type of p_m is P_j). Therefore, clocks in the resulting timed automaton $E(\text{Sch}_k)$ (and in the product automaton $E(\text{Sch}_k) \parallel E(A)$) are updated only by subtractions of integers. There is also a bound on the values of the clocks which are subtracted (the deadline), which makes the reachability problem for this type of automata (*Timed Automaton with Bounded Subtraction*) decidable, as proven in [FKPY07].

The product automaton $E(\text{Sch}_k) \parallel E(A)$ described above is a timed automaton with bounded subtraction for which it holds that the "unschedulable" error state is reachable if and only if A is unschedulable with scheduling strategy Sch_k .

The proof of this fact for the multi-processor case is the same as for the single-processor case in [FKPY07], with the difference that there can be several tasks running at the same timepoint. Therefore, it is necessary to check that Condition C_3 from the proof in [FKPY07], i.e., correctness of the invariant $b_i = B_m - (x_i^c - x_j^c)$, also holds here.

But from the fact that the scheduling function cannot move tasks from one queue to another one it follows that whenever a task instance p_i of type P_m is preempted by another instance p_j of type P_n , it will stay preempted (i.e., not being computed) until p_j finishes. Moreover, for p_j it holds that $b_j = B_n, w_j = W_n$ at the timepoint when it preempts p_i . This together implies the correctness of C_3 and thus also of the updates at the end of preemption.

□

Note that here, the diagonal-freeness in guards from the scheduler (see Definition 2) is important, because it is necessary to use clock differences to express the computation time of tasks. (A comparison of a clock difference to a constant is already a diagonal constraint which cannot be further extended by another clock.)

3.2 Non-preemptive scheduler

In this case, tasks are not allowed to preempt other already running tasks, which makes handling the computation time of all tasks in the queues easier. Therefore, even variable computation time of tasks can be allowed.

Theorem 2. *The problem of checking whether a task automaton together with a non-preemptive multi-processor scheduler is schedulable, is decidable.*

Proof (Sketch). Here, the same construction as above is used, i.e., two automata $E(A)$ and $E(\text{Sch}_k)$ are created. Again, $E(\text{Sch}_k)$ keeps track of computation progress and time pass since the release in clocks x_i^c and x_i^d for each released task p_i .

As in the previous case, whenever a new instance p_i of type P_j is released (event release_{P_j}), the clock x_i^d is reset. Before a task instance p_i is put to the head of a queue, the values b_i , w_i and d_i from the queue can be expressed as B_j , W_j and $D_j - x_i^d$, respectively. The clock x_i^c is not used at all.

As soon as the task instance p_i is scheduled for execution, the clock x_i^c is reset, keeping track of the computation progress of the task. For a running task instance p_i , the values b_i , w_i and d_i from its queue can be expressed as $B_i - x_i^c$, $W_i - x_i^c$ and $D_i - x_i^d$, respectively.

Because the scheduler is non-preemptive and the scheduling function cannot move tasks from one queue to another one, only the running task in each queue has already started its computation. The tasks which are not at the head of some queue did not start their computations yet. For this reason, we do not need to use x_i^c of the waiting tasks to compute b_i and w_i .

□

Note that for this result, the definition of the scheduling function (Definition 2) does not have to be that restrictive in the sense that also diagonal constraints can be allowed for the decisions in the scheduling function. The reason is that the computation time and deadline values in the queues are encoded into (at most) one clock each. The possibility of using diagonal constraints in the scheduler's decisions makes (for the non-preemptive case) encoding even of the Least Slack First [But97] scheduling policy possible.

4 Undecidability

In this section we show, that the schedulability problem is undecidable for multi-processor systems with at least two scheduling queues. This holds even if the precise task finishing times cannot influence releases of new tasks.

Theorem 3. *The problem of checking whether a task automaton without feedback using a k -multi-processor scheduler is schedulable, is undecidable for $k \geq 2$.*

To develop the proof, we first sketch a construction used in [FKPY07] for the single-processor case, where the automaton was allowed to use task feedback. We then extend the result to the multi-processor case, even if no task feedback is involved.

4.1 Undecidability on single-processor using feedback

The undecidability proof for the general single-processor schedulability problem is done by a reduction from the halting problem for two-counter machines. A *two-*

counter machine consists of a finite state control unit and two unbounded non-negative integer counters. The three possible instructions are counter-increment, counter-decrement and conditional branching (checking whether a counter value is zero). After each step, the machine state is changed deterministically. One of the states is a dedicated *halt state*. It is known, that the problem whether this halt state is reachable (the *halting problem for two-counter machines*) is undecidable.

The idea is, given a two-counter machine M , to construct a task automaton A_M and a scheduling policy such that a dedicated halt location in A_M is reachable if and only if the halt state of M is reachable. It will be ensured that no task can miss its deadline as long as the halt location is not visited. In turn, the queue can unboundedly grow in the halt location, making tasks miss their deadlines. The result of these two properties is that A_M will be unschedulable if and only if M can reach its halt state.

For each state of M 's control unit there is one corresponding location l_i in A_M . These locations are connected depending on the operation which M would execute at the corresponding state (increment, decrement, branching), through auxiliary locations "executing" this instruction.

To encode the counters into clock values, an N -*wrapping* construction from [HKPV98] is used. All clocks x stay within the interval $[0, N]$ for a constant N by resetting each clock x as soon as $x = N$ (*wrapping reset*). For a dedicated system clock x_{sys} these are the only resets (which makes x_{sys} periodic). In this way, *wrapping values* for all other clocks can be defined as their values at the (periodic) times where $x_{sys} = 0$. Thus, between any two consecutive *non-wrapping resets* of the clock x (i.e., resets when $x < N$ holds), its wrapping value does not change.

Using this construction, the value v of a counter C can be kept as the wrapping-value 2^{1-v} of a clock x_C . Therefore, this wrapping-value of x_C is always smaller than or equal to 2. The conditional branching is done by directly comparing the value of x_C to 2, and the increment (decrement) operation is implemented as dividing (multiplying) the wrapping value of x_C by 2. For the implementation of the decrement, a task Q with fixed computation time is released at a non-deterministically chosen time and a clock $x_{C_{new}}$ is reset at the same time. This task is then preempted by a task P of higher priority with variable execution time. P is released when x_C is reset to zero by a wrapping edge. A guard with $x_{done} = 0$ is used to check if its execution time is equal to the wrapping value of x_C , i.e., if it finishes when x_{sys} is reset. This preemption is repeated, and by using the $x_{done} = 0$ guard again afterwards to check that Q finishes when x_{sys} is reset, the wrapping-value of $x_{C_{new}}$ is forced to be the double of the wrapping value of x_C . The response time of Q is a constant time (its computation time) plus two times the wrapping value of x_C – because of the two preemptions from P . If any step in the construction fails (some non-deterministic choice was wrong), the automaton enters a sink location where no task is released. Figure 1 illustrates a successful run of the whole decrement procedure.




Figure 1. Time chart of the doubling procedure using the N -wrapping construction.

An increment for C needs to halve x_C . To achieve this, the wanted new value is simply guessed in a clock $x_{C_{new}}$ and then checked using the above decrement procedure. Only if the double of $x_{C_{new}}$ is x_C , the computation will continue.

During all three operations, all used tasks meet their deadline (using a fixed priority scheduler) and therefore the automaton is schedulable as long as it executes these instructions. A special location l_{halt} is used to represent the halt state of M . This location releases the task Q and it has an unguarded selfloop. Therefore, if A_M reaches l_{halt} , it will be able to release an unbounded amount of task instances at the same timepoint, making the system unschedulable.

Lemma 2 ([FKPY07]). *For a given two-counter machine M , the constructed task automaton A_M is unschedulable with the fixed priority scheduler if and only if M halts.*

Note that the construction uses all three properties (preemption, variable execution time, feedback) given in Section 2.

4.2 Undecidability on multi-processor without feedback

We extend this result to the multi-processor case for the systems with preemptive schedulers and variable computation time of tasks, but the finishing time of a task is never tested in the guards of a task automaton, that is, task feedback is not allowed (x_{done} is never used in guards). To achieve the same effect as with the task feedback, we develop two mechanisms based on the status of task queues and the multi-processor scheduling policy. First, we describe how the guards $x_{done} = 0$ can be replaced by certain task releases so that the system is still able to detect if the task in question finished at the right time. Secondly, a construction will be given to store this information in a task queue.

Both mechanisms will ensure that the demand for task feedback in the construction is removed. There is no branching upon the fact whether $x_{done} = 0$, the automaton continues its computation even if some non-deterministic choice of the system is wrong. However, to store the information about the fact that