

Formal verification of multitasking applications based on timed automata model

Libor Waszniowski · Zdeněk Hanzálek

Published online: 28 September 2007
© Springer Science+Business Media, LLC 2007

Abstract The aim of this paper is to show, how a multitasking application running under a real-time operating system compliant with an OSEK/VDX standard can be modeled by timed automata. The application under consideration consists of several non-preemptive tasks and interrupt service routines that can be synchronized by events. A model checking tool is used to verify time and logical properties of the proposed model. Use of this methodology is demonstrated on an automated gearbox case study and the result of the worst-case response time verification is compared with the classical method based on the time-demand analysis. It is shown that the model-checking approach provides less pessimistic results due to a more detailed model and exhaustive state-space exploration.

Keywords Formal methods · Verification · Model-checking · Timed automata · OSEK/VDX · Multitasking

1 Introduction

This paper deals with formal modeling of applications running under real-time operating system (OS). The typical application under assumption, shown as a case study in Sect. 7, is a complex controller consisting of periodic and aperiodic tasks constrained by deadlines and synchronized via inter-task communication primitives. The objective is to use model-checking approach (Larsen et al. 1995; Berard et al. 2001) for automatic verification of the model described in this paper.

L. Waszniowski (✉) · Z. Hanzálek
Centre for Applied Cybernetics, Department of Control Engineering, Faculty of Electrical Engineering, Czech Technical University, Karlovo nám 13, 12135 Prague 2, Czech Republic
e-mail: xwasznio@fel.cvut.cz

Z. Hanzálek
e-mail: hanzalek@fel.cvut.cz

The model based on timed automata (Alur and Dill 1994) considers an operating system, application tasks and a controlled environment behavior. It assumes a *fine grain model* of the task internal structure consisting of computations, OS calls, selected variables, code branching and loops. Therefore the model combines both, logic and timing characteristics of the discrete event system enabling one to check rather complex properties (safety and bounded liveness properties, state reachability or schedulability) by model checking tools (e.g. UPPAAL; Larsen et al. 1997 and Kronos; Daws et al. 1996) in finite time. Deadlock freeness of the application, occurrence of the race condition during access to shared data structures, a concrete value of some essential variable under certain conditions, end to end response time of an arbitrary event, proper ordering and timing of events in the control application or the controlled environment can be verified, for example.

Due to the composability of timed automata, models produced by different authors can be directly combined together. For example, a single processor system model can be simply expanded to a distributed system model by adding a communication layer model (Krákora et al. 2004).

Even though timed automata and model-checking (analogous to other formal methods) allow one to model and verify almost everything, it is generally known, that they are susceptible to the state space explosion. This fact restricts the size of verified application to a small size that seems to be unusable in practice (compared with matured response time analysis methods (Klein et al. 1993) or the offset-based analysis proposed by Palencia and Harbour (1998)). Therefore we try to show in this paper, how to build a compromised model of a reasonable size on one side and reasonable granularity on the other side, allowing for a detailed formal analysis of real-time properties that can not be made by response time analysis.

Methods for response time analysis based on time-demand analysis (Buttazzo 1997; Klein et al. 1993; Liu 2000) are well known and used in practice. These methods, e.g. rate monotonic analysis (RMA) (Sha et al. 1991), are straightforward for systems with independent periodic tasks but incorporation of non-periodic tasks and inter-task communication primitives can lead to pessimistic results (Bailey et al. 1995). This is caused by limited information inherent in the simple model of a task consisting of the worst-case execution time (WCET), the worst-case inter-arrival period and the worst-case blocking by lower-priority tasks. Analysis of end-to-end response time of transactions in a distributed system has been successfully solved by Tindell and Clark (1994) and later extended by Palencia and Harbour (1998). Also, these approaches however do not consider the detailed model of the controlled environment and the tasks internal structure as our approach does.

The response time analysis based on an exhaustive analysis of the fine grain model provides more precise (less pessimistic) results in some cases as is shown in Sect. 8. The price paid for this is higher memory requirements and time complexity of the model-checking method. Therefore the model-checking-based response time analysis cannot be seen as a universal method but as a less pessimistic and more demanding alternative to classical scheduling theory-based methods.

This paper focuses on a non-preemptive scheduling since tasks consisting of non-preemptive blocks of code can be modeled by timed automata, for which effective verification algorithms based on symbolic and reduction methods (see e.g. Larsen et al. 2003) exist.

Modeling of preemptive tasks has been studied by Corbett in (1996). This work provides a method for constructing models of real-time Ada tasking programs based on constant slope linear hybrid automata. Even though the author reports that the analyzing algorithm does usually terminate in practice, the reachability problem for hybrid automata is undecidable and therefore the analyzing algorithm termination is not guaranteed in general. The termination of the timed automata model verification is guaranteed, which is the advantage of our approach.

When modeling preemption in a multitasking application, it is necessary to stop a clock variable measuring the execution time of a preempted task and remember its value until the task is scheduled again. This can be done in hybrid automata, but not in timed automata. On the other hand the reachability problem is decidable for timed automata. This is a motivation of work (Waszniowski and Hanzálek 2005) providing a timed automata based over-approximate model of preemptive tasks. The over-approximation of the model means that besides the real behavior of the system, also some additional behavior is modeled. Therefore only properties preserved by this approximation (e.g. safety and bounded liveness properties) can be verified by a model-checking tool. Similarly, the model presented in this paper is over-approximate (in some cases, see Sect. 6) due to interrupts.

There are also extensions of Time Petri Nets allowing one to model systems with preemption; Preemptive Time Petri Nets (pTPN) (Bucci et al. 2004) and Scheduling Extended Time Petri Nets (SETPN) (Lime and Roux 2004). However, states of these formalisms are represented by a general convex polyhedra and the problem of state reachability is undecidable. It has been shown in Henzinger et al. (1998) that the problem of state reachability is undecidable for any formalisms that is expressive enough for modelling preemption. Therefore, decidable (finite state space) over-approximations that preserve safety and bounded liveness properties are usually used for verification of preemptive systems.

Timed automata are used to model primitives of Ravenscar run-time kernel for Ada in Lundqvist and Asplund (2003). However, the variable used to measure the execution time of tasks (modeling the system clock) is an integer, periodically incremented by a timed automaton after each “tick”. Therefore the notion of time in the application is discrete opposite to our approach where time is dense.

Discrete time for modeling a real-time application is also used in Campos and Clarke (1999) presenting a modeling language and a symbolic algorithm for quantitative analysis (providing minimum and maximum time between events) of synchronous real-time systems. Discrete time is also used in Fredette and Cleaveland (1993) where a generalized approach to schedulability analysis based on process algebra is proposed. Even though these approaches consider the task internal structure, the controlled environment affecting release times of tasks is not modeled. Our approach considers the controlled environment model.

Another interesting approach to schedulability analysis is based on timed automata extended by asynchronous tasks (i.e. tasks triggered by events) that provide a model for event-driven systems (Fersman et al. 2002, 2003). Each task specified by its execution time is associated to one timed automaton location. A transition leading to the location denotes an event releasing the task. Released tasks are stored in a queue and they are assumed to be executed according to a given scheduling strategy. The

problem of the system schedulability is transformed to the reachability problem in a timed automaton. This approach provides good results for aperiodic tasks (due to the detailed model of the environment releasing the tasks) but it does not consider the task internal structure. It would be possible to model the task internal structure by splitting the task to blocks of code and assigning them to locations of extended timed automaton representing control structure of the original task. Shared variables can be used to synchronize the end of one block of code execution, with the transition of the control structure timed automaton that starts the next block of code. However, the reachability problem of such a model is decidable only for non-preemptive scheduling or when all tasks have constant execution times (Krčál and Yi 2004).

In recent years several approaches integrating the schedulability analysis to some formal description methods have been published. Alvarez et al. in (2003) developed a method for computation of response time of tasks integrated to specification and description language (SDL). Similarly, Wang and Tsai (2004) present an approach to extend message sequence chart (MSC) by tasks parameters, and by response time analysis. Both these methods are an application of standard response time analysis without considering the internal structure of tasks and controlled environment.

This paper is organized as follows: Sect. 2 gives an overview on the basic features of OSEK compliant operating system. Readers familiar with OSEK specification do not need to read this section. Section 3 describes the *fine grain model* used in this paper. Sections 4, 5, and 6 presents the main results of this paper—timed automata models of tasks, OSEK compliant OS kernel and interrupt service routine (ISR). Section 7 demonstrates the proposed approach on an automated gearbox case study and Sect. 8 compares the task's response time analysis made by the model-checking approach and by the classical time-demand analysis. The paper is concluded with Sect. 9.

2 OSEK/VDX overview

This section surveys the basic features of an operating system compliant with OSEK/VDX Operating System specification, version 2.2.3 (OSEK 2005) (further called OSEK). OSEK is a simple static multitasking singleprocessor executive for electronic control units (ECU) used in automotive applications. Small memory demand requires simple services, which can be modeled by timed automata of reasonable size. All objects of the system are created in compilation time. Therefore they can be modeled by timed automata and static data structures.

2.1 Task management

OSEK provides static priority based, preemptive and non-preemptive scheduling (OSEK 2005), but we consider only non-preemptive scheduling in this paper. Even though OSEK distinguishes basic and extended tasks, we consider only extended ones, since basic tasks are only a subset variant of extended ones and both are modeled in the same way.

Tasks, created as *suspended* at the system generation time, become *ready* after activation by the OS service *ActivateTask* called from ISR or another task. The highest

priority ready task starts *running*. The running task may terminate its execution by calling the service *TerminateTask* and become suspended or it may voluntarily relinquish the processor by calling the service *Schedule* and become ready. If there is no higher-priority ready task, calling of the service *Schedule* does not affect the task execution. Extended tasks are, moreover, allowed to use the system call *WaitEvent*, which may result in a *waiting* state. To become ready, the waiting task requires the event (which it is waiting for) to be set.

2.2 Event management

OSEK provides an event management for task synchronization. The event is represented by one bit in a byte assigned to an extended task—the event’s owner. The event is therefore identified by its owner and its name (or mask specifying more than one event). The event owner may wait for the event and clear the event (services *WaitEvent* and *ClearEvent*). All tasks may set or get the binary value of a non-suspended task event (services *SetEvent* and *GetEvent*).

2.3 Resource management

Resource management is used to coordinate the access of several tasks (and interrupt service routines) to the critical section. The resource access protocol is used to provide mutually exclusive access, to prevent priority inversion and deadlock. According to this protocol, the priority ceiling is statically assigned to the resource at the system generation time. Its value is equal to the highest priority of all tasks (or ISR) accessing the resource. At run time, the priority of the task occupying the resource is increased to the resource priority ceiling. Task priority is reset to the previous value after releasing the resource. Consequently, no task (or ISR) ever tries to access the occupied resource and therefore no task can be blocked on the resource (notice that the OSEK specification does not allow any blocking OS services inside the critical section).

In non-preemptive scheduling, the mutually exclusive access of several tasks to the critical section is provided just by its non-preemptability and by the restriction of calling the OS services *Schedule*, *WaitEvent* and *TerminateTask* from the critical section. Simultaneous access of a task and an ISR to the critical section can be prevented by disabling interrupts within the critical section or by using resources.

2.4 Interrupt management

OSEK distinguishes interrupt service routines (ISRs) of category 1 that do not use any OS services (no influence on the task management) and ISR of category 2 allowing all OS services except some services dedicated entirely to tasks (*WaitEvent*, *TerminateTask*). There is no difference between both categories from the modeling point of view. When a task with priority higher than the interrupted one is activated by an OS service called from the ISR, the interrupted task is not preempted due to non-preemptive scheduling. Even when the processor is idling, when an interrupt occurs, no rescheduling takes place at the OS service called from the ISR, but the OS

service only changes states of tasks, and the rescheduling takes place at the end of the ISR. Therefore, when several tasks are activated in the ISR, it does not depend on their activation order, but the highest priority one is scheduled at the end of the ISR (when no task is currently running).

3 Fine grain model of multitasking application

The *fine grain model* treats the internal structure of the tasks and the interrupt service routines (ISR), the OS functionality and the controlled environment behavior. All components are modeled by timed automata synchronized via channels and by shared variables. The task model consists of several blocks of code called *computations*, calls of OS services, selected variables, code branching and loops (affected by values of selected variables). *Computations* are defined by the BCET (the best-case execution time) and the WCET (the worst-case execution time). Considering the execution time as an interval $\langle \text{BCET}, \text{WCET} \rangle$, it allows one to incorporate the uncertainty of the execution time due to non-modeled code branching inside the *computations*, cycle stealing by a DMA device, etc. When a general property of the model is analyzed by an exhaustive state space search (made by a model checking tool), the execution time of a task must be specified by an interval covering all possible cases, i.e. $\langle \text{BCET}, \text{WCET} \rangle$. Due to the possibility of a scheduling anomaly, the WCET of *computations* does not necessarily lead to the worst-case response time of the whole task.

The structure of the entire model is shown in Fig. 1. Rectangular blocks represent particular timed automata (e.g. task automaton in Fig. 5b) or OS service automaton in Fig. 8). Synchronization of timed automata is expressed by arcs labeled by the name of the synchronization channel (*ActivateTask*, *EndSysCall*, etc.). The most important data structures (e.g. *Q*, *P*, *State*) are shown on the right side of the figure. The essential components are explained in Sects. 4, 5 and 6.

Since the schedulability is one of the most often analyzed properties, it is attractive to compare our *fine grain model* based on timed automata with the classical scheduling theory task model based on the WCET of the whole task, its minimum inter-arrival time and its blocking time related to resources (Buttazzo 1997; Sha et al. 1991; Liu 2000). Classical response time analysis based on such model computes the worst case finishing time of the task by adding together its worst case execution time, duration of preemption by higher-priority tasks in the worst case inter-arrival times and phasing, and the worst case blocking by lower-priority tasks on shared resources. Such worst-case finishing time is a conservative abstraction of all possible finishing times but it could be a too pessimistic abstraction in many applications, since all the mentioned worst cases do not occur at the same time (Bailey et al. 1995).

Exhaustive analysis of *fine grain model* behavior (automatically completed by model checking tool UPPAAL) considers the task finishing times corresponding to the realistic phasing, the realistic blocking and the realistic execution time in relation to the modeled code branching. Therefore, the result is as precise as the model. The price paid for the exhaustive analysis is higher complexity.

We will demonstrate the advantage of the model containing the tasks internal structure in a simple example. Let us consider two tasks, higher-priority *Task₁* and

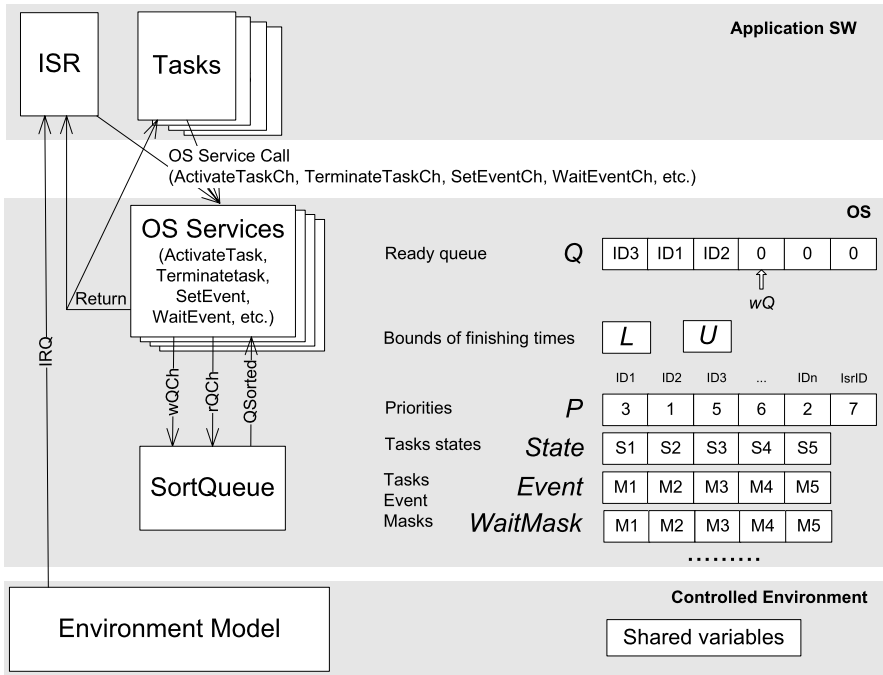


Fig. 1 Overview of entire timed automata model

Fig. 2 Tasks pseudocode

```

Task1
{
  if (Data==OK)
  {
    LongComputation;
  }
  else
  {
    ShortComputation;
    ActivateTask (Task2);
  }
}

Task2
{
  Computation;
}
    
```

lower-priority *Task₂* listed in Fig. 2. *Task₁* is activated with a period 12. Depending on the variable *data*, it performs either *LongComputation* taking 8 time units or *ShortComputation* taking only 2 time units. *Task₂* is activated if and only if *ShortComputation* is executed. *Task₂* execution takes 10 time units.

Figure 3 shows that both tasks are finished prior to their next activation in both cases, *Data == OK* and *Data! = OK*.

Let us analyze the application by demand analysis (Liu 2000) based on a simple model considering only the WCET and the period of tasks. Without knowledge of *Task₁*'s internal structure, we must consider its WCET₁ to be equal to 8 and activation of *Task₂* must be considered at each period of *Task₁*.

Fig. 3 Schedule considering tasks internal structure

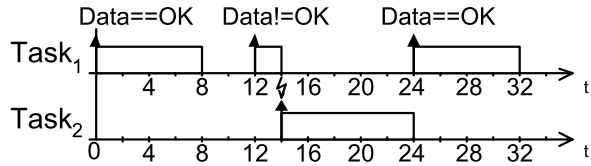
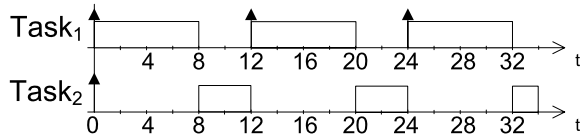


Fig. 4 Schedule ignoring tasks internal structure



In the case of preemptive scheduling, *Task*₁ is always finished 8 time units after the beginning of the period, but the worst-case response time of *Task*₂ is $10 + 3 \times 8 = 34$ (see Fig. 4), which exceeds its period.

In the case of non-preemptive scheduling, the worst-case response time of both tasks is $10 + 8 = 18$, which is longer than the activation period.

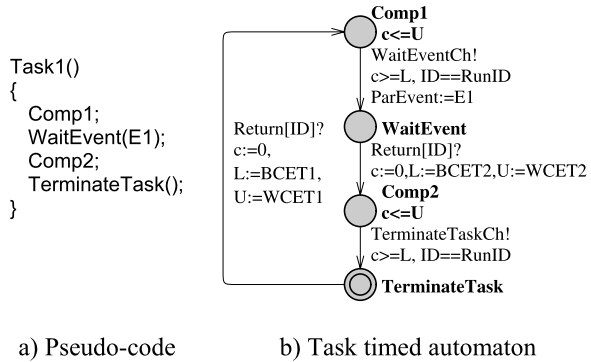
4 Task model

Each task instance is modeled by one timed automaton that is synchronized with the OS model via channels depicted as arrows in Fig. 1. Figure 5 demonstrates the modeling methodology in the example of a simple task executing *computations Comp1* and *Comp2* and calling OS services *WaitEvent(task, event)* and *TerminateTask*.

Timed automata are depicted in the UPPAAL notation (Larsen et al. 1997), where the location with double circles represents the initial location. Each location can be labeled by its name and time invariant (both in bold font). Invariant in the form “ $c \leq U$ ”, allows it to stay in the location only when the valuation of the clock variable c is smaller than or equal to integer U . Each transition can be labeled by synchronization (channel name with ‘?’ or ‘!’), guard (logical terms separated by a comma, e.g., $c \geq L[1], State[1] == RUNNING$) and assignment (assignments using the sign ‘:=’ separated by a comma).

Each *computation* is represented by one location of the same name (e.g. *Comp1*). Time spent in this location (measured by clock c) represents the *computation’s* finishing time (i.e. time necessary for its execution including interrupts) and is bounded by values stored as integer variables L and U (provided by time invariant $c \leq U$ and guard $c \geq L$). Clock c is reset and variables L and U are initialized by the BCET and the WCET of the *computation* on the transition incoming at the corresponding location (e.g. $c := 0, L := BCET1, U := WCET1$ on the transition leading to *Comp1*). Variables L and U are increased when the task is interrupted (provided by timed automaton modeling ISR; see Sect. 6). Notice that due to non-preemptive scheduling, only one clock c can be shared by all task timed automata, which considerably reduces the size of the state space. The guard $ID == RunID$ prevents the task timed

Fig. 5 Simple task example



automaton from progressing when it is not scheduled (i.e. an ISR is executed). Constant *ID* is a unique identifier (0, 1, 2, ...) of tasks and ISR.

OS service call is modeled by the transition synchronized via the channel of the corresponding name (e.g. *WaitEventCh!*) with the automaton modeling the OS service functionality, and by the location of the corresponding name (e.g. *WaitEvent*) in which the task is waiting a return from the service (channel *Return[ID]?*). OS service parameters, if they are required, are delivered through shared variables *ParTask* and *ParEvent*. Notice that some OS services (e.g. *Schedule* or *WaitEvent*) can cause rescheduling. In this case, the return from the called service occurs after finishing all higher-priority tasks.

Realize that the task can be running or interrupted when its model is in a location corresponding to a computation. When the task model is in location corresponding to an OS service, OS executes the service and then the task can continue its execution (e.g. service *SetEvent*), or the task is blocked while waiting for an external event (service *WaitEvent*), or the task is ready to execute but a higher priority task is executed (service *Schedule*) or the task is suspended (service *TerminateTask*). Even though the task code in Fig. 5a is linear (it does not contain any loops), the timed automaton in Fig. 5b is cyclic since the suspended task (task timed automaton is in the location *TerminateTask*) can be activated and start its execution from the beginning. This is also the reason why location *TerminateTask* is the initial one.

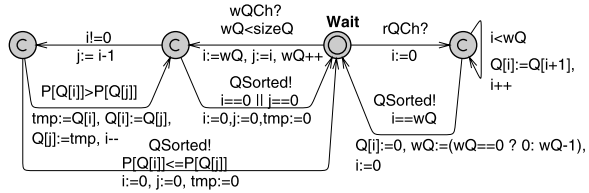
5 OS kernel model

The OS kernel model consists of integer variables representing the OS objects (e.g. ready queue), the timed automata representing OS services functionality, and the timed automaton sorting the ready queue according to priorities (*SortQueue*). See the model overview in Fig. 1.

5.1 Kernel variables

The task priority is stored in a global array *P*, indexed by *ID*. Higher number represents higher priority. *P* can be a constant when the resource management is not modeled.

Fig. 6 *SortQ* automaton



The task state is stored in the array *State* at the index corresponding to task *ID*. The task state is either *SUSPENDED*, *WAITING*, *READY* or *RUNNING*. However it is necessary to distinguish only the *SUSPENDED* from all others in the proposed model. Therefore the symbols *WAITING*, *READY* or *RUNNING* have the same value in the model.

The variable *RunID* stores *ID* of the currently running task or interrupt service routine.

IDs of all tasks, which are ready for execution, are stored in the ready queue modeled as a global array *Q* (see Fig. 1). *IDs* of tasks are stored at the lowest positions in the array. The variable *wQ* contains the first empty position in *Q*. Tasks are ordered in descending order according to their priorities in *Q*. The ready task with the highest priority is always at the position zero; the ready task with the lowest priority is always at the position *wQ*-1.

The queue must be reordered according to tasks priorities after writing a new task and all elements of the queue must be shifted to the left after reading the highest priority ready task from the zero position. Both these mechanisms are provided by the automaton *SortQueue* depicted in Fig. 6.

The reordering mechanism is started by the synchronization channel *wQCh* after writing a new *ID* to *Q[wQ]*. The pointer *wQ* is then increased and priorities of tasks in neighboring position in *Q* are compared (started from *wQ*) and if there is a higher priority task in a higher position, *IDs* are swapped. The termination of this mechanism is announced by the channel *QSorted*. Shifting of *Q* after reading *Q[0]* is started by the synchronization channel *rQCh* and its finishing is announced by the channel *QSorted*.

Notice that it would be possible to implement the ready queue as a circular buffer. The top of the queue would not always be at position zero, but it would be pointed by the pointer (lets call it *rQ*) that is increased after reading the highest priority task. It is not necessary to shift elements of *Q* in this case. A circular buffer would, therefore, be a more elegant approach from the programming point of view, but it is not appropriate for verification purposes, since such a model generates a bigger state space. Realize that two different configurations of a circular buffer containing the same tasks but stored in different positions (different *rQ* and *wQ*) are represented by two different states in the state space, but they represent the same situation from the application point of view. Contrary to that, all situations when *Q* contains the same tasks are represented by only one state in our approach, since the same set of tasks is always stored in the same position in *Q* (from zero to *wQ*-1).

Events are represented by the integer array *Event* associating one byte *Event[ID]* to each task. Each bit in *Event[ID]* represents one event that can be set or cleared.

Fig. 7 *WaitEvent* OS service pseudo-code

```

WaitEvent (Mask)
{
  if ((Event[RunID] & Mask) == 0)
  {
    State[RunID] := WAITING;
    WaitMask[RunID] := Mask;
    RunID := Extract Top of ReadyQ;
    ContextSwitch;
    State[RunID] := RUNNING;
  }
  return E_OK;
};

```

Moreover, the integer array *WaitMask* represents events, which the corresponding task is waiting for.

Variables *L* and *U* are necessary for model execution of the task code. As it has been already stated in Sect. 4, the variables are used to store lower and upper margins of the finishing time of the current started *computation*. These integers are initialized by the BCET and WCET respectively, at the start of the *computation* and they are increased by an ISR BCET and WCET respectively when an interrupt occurs (provided by a timed automaton modeling the ISR).

5.2 OS services

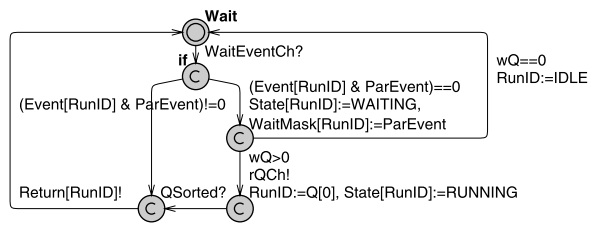
Each OS service is modeled by a timed automaton representing its functionality defined by OSEK specification (OSEK 2005). The automaton is waiting in its initial location until its function is called from the task model (by synchronization via the corresponding channel e.g. *WaitEventCh*). Then it manipulates the tasks' states, ready queue and other operating system objects (e.g. events). OS services *Schedule*, *Wait-Event* and *TerminateTask* can, moreover, reschedule the current tasks. This is done by choosing the highest priority ready task and storing its *ID* in the variable *RunID*. The next *computation* of the *RunID* task is then started by taking a transition synchronized by channel *Return[RunID]*. Models of all OSEK services can be found on: <http://dce.felk.cvut.cz/waszniewski/RTVerif/RTVerif.htm>

As an example of a service model, we introduce *WaitEvent(Mask)* service that causes the task to wait for events specified by *Mask*. Figure 7 shows *WaitEvent* OS service functionality in pseudo-code. The corresponding automaton is in Fig. 8. It is supposed that interrupts are disabled within the whole service. Locations marked by “C” are so called committed locations. The committed location is left immediately without any interference from another automaton that is not in committed location. Since all locations in the OS services automata, except the initial one, are committed locations, the whole service is atomic from the point of view of the tasks and the controlled environment models. The execution time of the OS services is involved in the execution times of *computations* calling them.

6 Interrupt service routine model

In this section we present a timed automata model of an ISR. The ISR can interrupt the execution of a task. It is modeled by increasing the bounds of the interrupted task

Fig. 8 WaitEvent OS service automaton



execution time (stored in variables L, U) by the bounds of the ISR execution time ($BCET_ISR, WCET_ISR$). We show that this approach brings an over-approximation to the model.

For reasons of simplicity, only one ISR is assumed to exist in this article, but it can be generalized when all hardware details are considered. The ISR is modeled by a timed automaton modeling application dependent code in the same way as the task code. Moreover, there is an initialization part preventing a task scheduling inside the ISR. $RunID$ is stored in the local variable *InterruptedID* and the ID of the ISR ($IsrID$) is written to the variable $RunID$. Therefore the OS services called from the *ISR* do not schedule any task (because $RunID! = IDLE$). Values of L and U are increased by the $BCET_ISR$ and $WCET_ISR$ respectively. Further, there is a finalization part providing task scheduling at the end of the ISR (as it is required by OSEK specification (OSEK 2005)). Either *InterruptedID* or, if it is equal to $IDLE$, the ID from the top of the ready queue is written to the variable $RunID$. An example of ISR pseudocode is in Fig. 13 and the corresponding timed automaton is in Fig. 19.

Lets us explore the approach for taking the ISR execution time into account in the interrupted task execution time. When the interrupt occurs the execution time bounds of the interrupted *computation* (stored in variables L and U) should be prolonged by the duration of the ISR execution. Since the right duration of the interruption cannot be measured in timed automata (a clock variable cannot be stopped or stored), the bounds L and U are increased by bounds of the possible ISR execution time $BCET_ISR$ and $WCET_ISR$. This introduces an additional non-determinism to the model since the modeled duration of the interrupted task interruption is not necessary equal to the ISR execution (what holds in the real system). Therefore the set of real system behaviors is a subset of the modeled behaviors, i.e. the model is an over-approximation of the real system.

To illustrate the over-approximation let us consider for example a *computation* of task T with the execution time $C_T \in [1, 4]$ interrupted by an ISR with the execution time $C_{ISR} \in [2, 4]$. All possible relative finishing times of the interrupted *computation* versus the ISR execution time in the real system and in the proposed model are depicted in Fig. 9. Finishing time of the interrupted *computation* is equal to its execution time C_T plus interruption duration. Interruption duration is bounded by bounds of C_{ISR} in the model but it is equal to the actual execution time of *ISR* in the real system.

Figure 9 shows that not all modeled behaviors can occur in the real system. It is very important to keep this fact in mind during the verification process, since the over-approximation does not preserve a general property. It means that it cannot be automatically concluded that a general property satisfied by the model is also satisfied

Fig. 9 Possible values of relative finishing time F_T of the interrupted computation of the task T versus the ISR execution time C_{ISR}

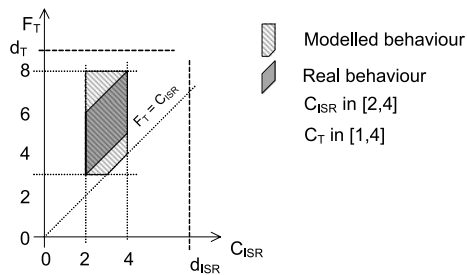
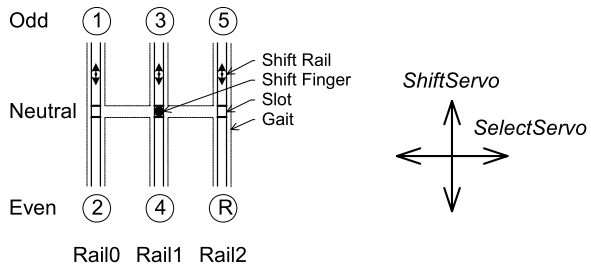


Fig. 10 Gear box mechanism



by the real system. On the other hand, it is important from the practical point of view, that over-approximation preserves safety and bounded liveness properties (Berard et al. 2001). A safety property states that, under certain conditions, an undesirable event never occurs. A bounded liveness property states that, under certain condition, some desirable event will occur within some deadline. See examples in Sect. 7.

Please realize that the model is over-approximate only in the case that the $WCET_ISR$ differs from the $BCET_ISR$.

Schedulability is an often verified property, exploring whether *computations* are finished prior to their deadlines (d_{ISR} and d_T in Fig. 9) in all situations. Figure 9 shows that the worst case finishing time of the task or ISR is the same in the model and in the real system. A result of the schedulability analysis based on this model is therefore correct and corresponds to reality (it is not pessimistic).

7 Gear box case study

7.1 System description

The proposed modeling methodology is demonstrated on an automated gearbox control system. The controlled system consists of a five-speed gearbox and a dry clutch. The gearbox mechanics are depicted in Fig. 10. They consist of three shift rails and a shift finger actuated by *SelectServo* and *ShiftServo*. *SelectServo* can move the shift finger from a slot of one rail to another one. *ShiftServo* engages one of two gears (odd or even) or neutral by moving the selected rail by the shift finger. The direction of the shift finger movement is limited by a gait.

Fig. 11 Slip control task pseudocode

```

SlipCtrlTask() // Activated periodically
{
    // Slip control, torque tracking, clutch protection
    if (ClutchState == CLOSED)
        CompSlipCtrl;
    TerminateTask();
};

```

Fig. 12 Select Gear task pseudocode

```

SelectGearTask() // Activated periodically
{
    if (GBReady)
        // Select gear according to current car conditions
        CompDesiredGear;
        if (DesiredGear != CurrentGear)
            ActivateTask (GearBoxCtrlTask);
    }
    TerminateTask();
};

```

The gearbox is controlled by a single processor control unit running an OSEK compliant OS. The application software consists of three tasks (*SlipCtrlTask*, *SelectGearTask*, *GearBoxCtrlTask*) and one ISR.

The ISR (see pseudocode in Fig. 13) is periodically invoked by a timer (with the period 10) and by the clutch, *ShiftServo* or *SelectServo* when their position changes. The source of the interrupt is specified by bits *bTimerInt*, *bClutchInt*, *bShiftServoInt* and *bSelectServoInt*. According to the source of the interrupt, tasks are activated (*bTimerInt*) or an event is set (*bClutchInt*, *bShiftServoInt* and *bSelectServoInt*).

Task *SlipCtrlTask* (see pseudocode in Fig. 11) is periodically activated by ISR. Its priority is 2 and its period is 10. It provides slip control and torque tracking but its detailed functionality is not relevant to verification, therefore, it is not considered here. Only its computation time is modeled.

Task *SelectGearTask* (see pseudocode in Fig. 12) is periodically activated by ISR. Its priority is 0 and its period is 500. It selects the appropriate transmission rate, writes it to the variable *DesiredGear*, and if the desired gear differs from the current one, it activates task *GearBoxCtrlTask* that controls changing of the gear. Also the model of this task is very rough.

Task *GearBoxCtrlTask* (see pseudocode in Fig. 14) has priority 1. It sends a command to open the clutch first, then it waits for the event *ClutchEvent* signaling that the clutch is open. If NEUTRAL is not currently engaged, it disengages the current gear by sending the command to the *ShiftServo* to move the shift finger to the neutral position and waits for the event *ShiftServoEvent*. Then the new gear, stored in variable *DesiredGear*, can be engaged. First, the rail and shift direction corresponding to the *DesiredGear* are computed and stored in the variable *DesiredRail* (Rail 0, 1, or 2) and *DesiredShift* (ODD or EVEN). Then if the *DesiredRail* is not currently selected, the command to the *SelectServo* is sent to move the shift finger to the position of the *DesiredRail*. When the *DesiredRail* is selected (signalled by the event *SelectServoEvent*) the *DesiredGear* is engaged by sending the command to the *ShiftServo* to move the selected rail to *DesiredShift* position. After finishing the *ShiftServo* movement (signalled by the event *ShiftServoEvent*), the command to close the clutch is

```

ISR() // IRQ sources - TIMER, Clutch, ShiftServo, SelectServo
{
  // Initialization part – for modeling purposes
  InterruptedID := RunID;
  RunID := IsrID;
  U += U_IsrID; // Prolong Interrupted task by constants U_IsrID and L_IsrID
  L += L_IsrID;

  // User defined code
  while (bTimerInt || bClutchInt || bShiftServoInt || bSelectServoInt)
  {
    Comp;
    if (bTimerInt)
    {
      bTimerInt:=0;
      clk:=(clk>MAX_CLK ? 1 : clk+1);
      if (clk% SlipCtrlTaskPeriod == 0)
        ActivateTask(SlipCtrlTask);
      if (clk% SelectGearTaskPeriod == 0)
        ActivateTask(SelectGearTask);
    }
    else if (bClutchInt)
    {
      bClutchInt:=0;
      SetEvent (GearBoxCtrlTask, ClutchEvent);
    }
    else if (bShiftServoInt)
    {
      bShiftServoInt:=0;
      SetEvent (GearBoxCtrlTask, ShiftServoEvent);
    }
    else if (bSelectServoInt)
    {
      bSelectServoInt:=0;
      SetEvent (GearBoxCtrlTask, SelectServoEvent);
    }
  }

  // Finalization part (Scheduling point) – for modeling purposes
  if (InterruptedID != IDLE)
  {
    RunID := InterruptedID;
    Return to Interrupted Task;
  }
  else if (ReadyQueue is Empty)
  {
    RunID := IDLE;
    Return to IDLE task;
  }
  else
  {
    RunID := Extract Top of ReadyQ;
    Return to RunID task;
  }
};

```

Fig. 13 Interrupt service routine pseudocode

```

GearBoxCtrlTask() // Activated by SelectGearTask
{
  GBReady := 0;
  ClearEvent (ClutchEvent);
  OpenClutch; // Send command to clutch
  WaitEvent (ClutchEvent);
  if (CurrentShift != NEUTRAL)
  {
    // Disengage
    ClearEvent (ShiftServoEvent);
    ShiftServo_Goto (NEUTRAL); // Send command to ShiftServo
    WaitEvent (ShiftServoEvent);
  }
  if (DesiredGear != NEUTRAL)
  {
    // Select shifting rail
    DesiredRail := (DesiredGear-1)/2; // integer division
    DesiredShift := (DesiredGear-1)%2+1; // modulo operation
    if (DesiredRail != CurrentRail)
    {
      // Select
      ClearEvent (SelectServoEvent);
      SelectServo_Goto (DesiredRail); // Send command to SelectServo
      WaitEvent (SelectServoEvent);
    }
    // Shift
    ClearEvent (ShiftServoEvent);
    ShiftServo_Goto (DesiredShift); // Send command to ShiftServo
    WaitEvent (ShiftServoEvent);
  }
  ClearEvent (ClutchEvent);
  CloseClutch; // Send command to clutch
  WaitEvent (ClutchEvent);
  GBReady := 1;
  CurrentGear:=DesiredGear;
  TerminateTask();
};

```

Fig. 14 Gear Box Control task pseudocode

sent and when the clutch is closed (signaled by the event *ClutchEvent*), the variable *CurrentGear* is updated and the task is terminated.

7.2 Model

A model of the whole system consists of timed automata representing the controlled system (*Clutch*, *SelectServo* and *ShiftServo*), a hardware of a control unit (periodic timer generating interrupts), an OS (services *ActivateTask*, *TerminateTask*, *SetEvent*, *WaitEvent* and automaton *SortQueue*), three application tasks (*SlipCtrlTask*, *SelectGearTask*, *GearBoxCtrlTask*) and one ISR. An overview of the whole model is depicted in Fig. 15 and Fig. 16. Figure 15 shows the timed automata synchronization via the channels and Fig. 16 shows events and variables shared by the timed automata in the model. The variables and timed automata modeling OS have been explained in the previous section, therefore they are omitted in both figures.

Fig. 15 Model overview—timed automata (rectangles) and synchronizations via channels (arrows)

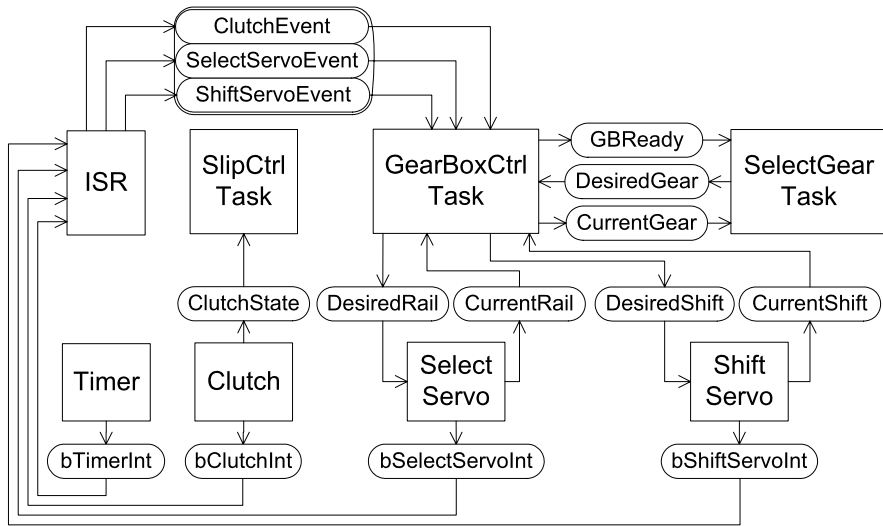
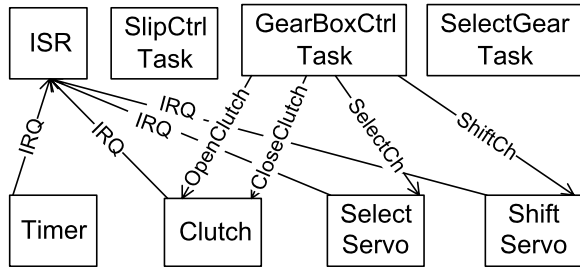


Fig. 16 Model overview—timed automata (rectangles) and shared variables (ovals)

Since the *Timer* timed automaton is very simple, it is not depicted here. It only waits in its initial location and periodically generates interrupt via the channel *IRQ* and the variable *bTimerInt*.

The *Clutch* timed automaton is depicted in Fig. 17. It is in the location *Closed* or *Opened* in a steady state. When the *Clutch* receives the command to open or close (via channel *OpenClutch* or *CloseClutch* respectively), it moves to *Opening* or *Closing* respectively. After the time bounded by *ShiftTime* and *ShiftTimeU*, the *Clutch* reaches a new steady state and generates an interrupt request (*IRQ*) via the channel *IRQ*.

The *SelectServo* timed automaton is depicted in Fig. 18. The automaton is in one location corresponding to *Rail0*, *Rail1* or *Rail2* in a steady state. When the command to select a new rail is received via the channel *SelectCh*, the *SelectServo* automaton moves to the *DesiredRail* (*BetweenRail0andRail1* and *BetweenRail1andRail2*). When the *DesiredRail* is reached, an *IRQ* is generated via the channel *IRQ*.

The *ShiftServo* timed automaton (not depicted here) differs from the *SelectServo* timed automaton only in a few details. Locations *Rail0*, *Rail1*, *Rail2*, *BetweenRail0andRail1* and *BetweenRail1andRail2* are replaced by *OddPos*, *NeutralPos*,

Fig. 17 Clutch timed automaton

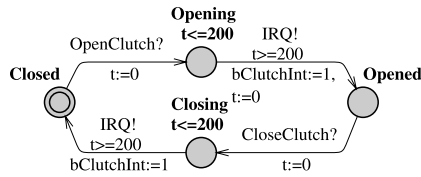


Fig. 18 SelectServo timed automaton

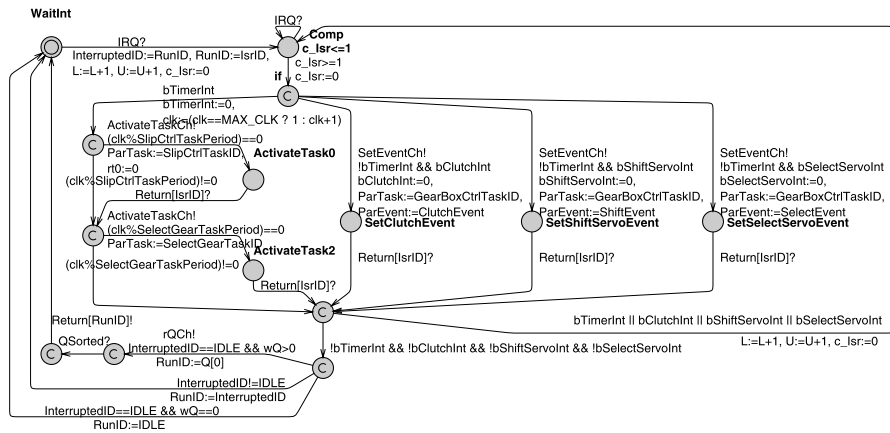
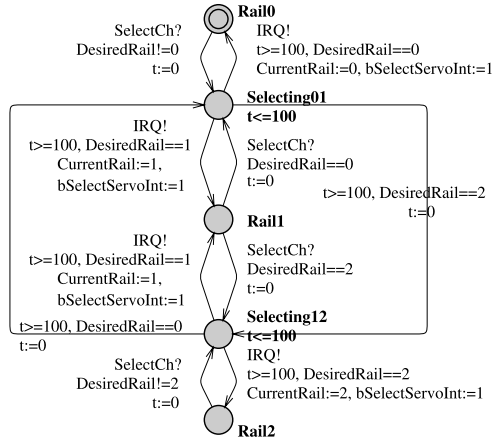


Fig. 19 ISR timed automaton

EvenPos, *BetweenNeutralAndOdd* and *BetweenNeutralAndEven*. Variables and channels related to selecting (*DesiredRail*, *CurrentRail*, *SelectCh*, *bSelectServoInt*) are replaced by variables and channels related to shifting (*DesiredShift*, *CurrentShift*, *ShiftCh*, *bShiftServoInt*). Moreover, the time spent in locations *BetweenNeutralAndOdd* and *BetweenNeutralAndEven* is not 100 time units but 200 time units.

Timed automata modeling tasks *SlipCtrlTask*, *SelectGearTask* and *GearBoxCtrlTask* and *ISR* function are depicted in Fig. 20, Fig. 21, Fig. 22 and Fig. 19, re-

Fig. 20 SlipCtrlTask timed automaton

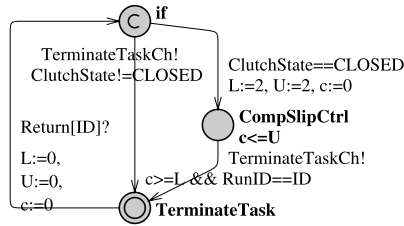
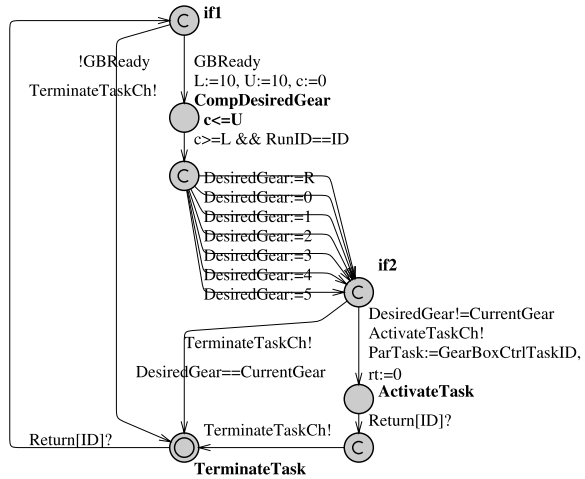


Fig. 21 SelectGearTask timed automaton



spectively. They have been obtained by translating tasks pseudocodes from Fig. 11, Fig. 12, Fig. 14 as well as Fig. 13 to timed automata according to the methodology described in Sects. 4 and 6, respectively.

Notice that the computation of the *DesiredGear* in the *SelectGearTask* is modeled by non-deterministic choice in the *SelectGearTask* timed automaton. Therefore, all possibilities are explored by the model-checking tool. Notice also that the OS service *ClearEvent* is very simple ($Event[ID] := Event[ID] \&! ClutchEvent$); it is not modeled by a special automaton but it is modeled directly in the *GearBoxCtrlTask* timed automaton (see Fig. 22).

7.3 Formal verification

The following properties are required for proper function of the system:

Safety properties:

- P1. Shifting is allowed only when the clutch is open
- P2. Selecting is allowed only when the shift servo is in neutral
- P3. Shifting is allowed only when a rail is selected
- P4. Clutch cannot be open longer than 650 time units

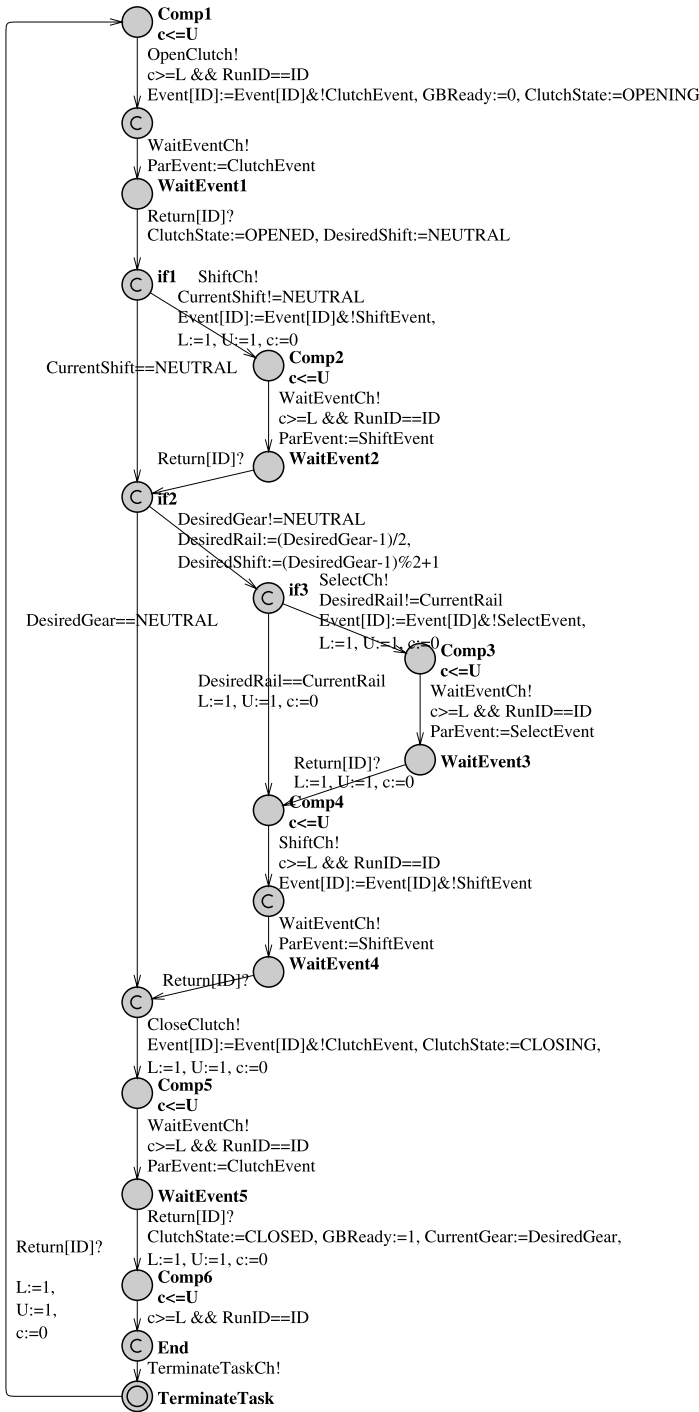


Fig. 22 GearBoxCtrlTask timed automaton

Bounded liveness:

P5–P11. When new desired gear (NEUTRAL, 1...5, R) is selected, it is engaged in 1020 time units

Deadlock-freeness:

P12. The system is deadlock free

Notice that since multiple tasks activation is disabled in the model, deadlock-freeness (P12) guarantees that all tasks are finished prior to their new activation in the next period. Therefore the property P12 can be interpreted as schedulability (deadlines at the end of periods).

Please realize that the deadlock-freeness is not proved to be preserved by the over-approximation described in Sect. 6. It can be therefore verified only in the case of the WCET_ISR equal to the BCET_ISR that holds in this model.

The above listed properties have been formalized in UPPAAL requirement specification language as follows:

P1. $A[] \text{Clutch.Closed} \text{ imply } (\text{ShiftServo.OddPos} \text{ or } \text{ShiftServo.NeutralPos} \text{ or } \text{ShiftServo.EvenPos})$

P2. $A[] \text{ not } (\text{SelectServo.Rail0} \text{ or } \text{SelectServo.Rail1} \text{ or } \text{SelectServo.Rail2}) \text{ imply } \text{ShiftServo.NeutralPos}$

P3. $A[] \text{ not } \text{ShiftServo.NeutralPos} \text{ imply } (\text{SelectServo.Rail0} \text{ or } \text{SelectServo.Rail1} \text{ or } \text{SelectServo.Rail2})$

P4. $A[] \text{Clutch.Opened} \text{ imply } \text{Clutch.t} \leq 650$

P5. $(\text{DesiredGear} == 0 \text{ and } \text{SelectGearTask.ActivateTask}) \rightarrow (\text{ShiftServo.NeutralPos} \text{ and } \text{rt1} \leq 1020)$

P6. $(\text{DesiredGear} == 1 \text{ and } \text{SelectGearTask.ActivateTask}) \rightarrow (\text{ShiftServo.OddPos} \text{ and } \text{SelectServo.Rail0} \text{ and } \text{rt1} \leq 1020)$

P7. $(\text{DesiredGear} == 2 \text{ and } \text{SelectGearTask.ActivateTask}) \rightarrow (\text{ShiftServo.EvenPos} \text{ and } \text{SelectServo.Rail0} \text{ and } \text{rt1} \leq 1020)$

P8–P11. Similar to P6 and P7

P12. $A[] \text{ not deadlock}$

In UPPAAL requirement specification language the syntax $A[] f$ represents the computation tree logic (CTL) formula $\forall \square f$ (i.e. “invariantly holds f ”), and the syntax $p \rightarrow q$ denotes a CTL property $\forall \square (p \Rightarrow \forall \diamond q)$ (i.e. “whenever p holds, eventually q will hold as well”). Notice that the clock $rt1$ measuring the response time in all bounded liveness properties P5–P11 is reset when the new *DesiredGear* is selected in the *SelectGearTask* timed automaton depicted in Fig. 21.

All the above mentioned properties of the system have been successfully verified by model-checker UPPAAL 3.4.7 running on Windows 2000 on PC AMD Athlon 1 GHz, with 1.3 GB RAM. The time required for verification of all of these twelve properties is 58 seconds. The required memory is 78 MB. The most demanding properties are the bounded liveness properties. Realize that time and memory requirements drastically grow with the complexity of the model. Even though the memory requirement of 78 MB is acceptable, an augmentation of the model by other tasks can easily make the verification impossible.

8 Response time analysis

The previous section presents the verification possibilities of the model-checking methods. One of the many properties that can be verified by model-checking is, whether a task response time satisfies its deadline. This section compares a task response time analysis based on the classical scheduling theory (Klein et al. 1993) and (Palencia and Harbour 1998) to the one based on the model-checking approach proposed in this paper. Both these approaches are applied on the *SlipCtrlTask* and the *GearBoxCtrlTask* and the worst-case response times (WCRT) obtained by both approaches are compared. The results show, that the worst-case response times obtained by the model-checking approach are, due to a more detailed model and exhaustive state space analysis, less pessimistic. In the case of the *SlipCtrlTask*, the advantage of the model-checking approach is based on the information contained in the controlled environment model and the model of the task internal structure. In the case of the *GearBoxCtrlTask* containing self-suspension, the advantage of the model-checking approach is based on exhaustive analysis of the model state space.

On the other hand, exhaustive analysis of the state-space (very quickly growing with the number of tasks, computations and variables) limits the size of applications for which the model-checking method can be used. Therefore, we do not want to present the model-checking approach as a universal method for response time analysis but as an alternative approach providing some reasonable advantages in some cases.

8.1 WCRT of the *SlipCtrlTask* by scheduling theory based approach

Let us explore the response time of the *SlipctrlTask* from its activation within the ISR (see pseudo code in Fig. 13) till its termination. The scheduling theory based response time analysis deals with the following information.

The WCET of the *SlipCtrlTask* is 2 time units. Even though the *SlipCtrlTask* is executed at the highest priority, its response time can be prolonged by the longest non-preemptable part of any lower-priority task, and by an ISR. The longest non-preemptable block in the system is the task *SelectGearTask* whose WCET is 10. The ISR is invoked by the timer (with period 10), the clutch (with minimal inter-arrival time 200), the *SelectServo* (with minimal inter-arrival time 100), and the *ShiftServo* (with minimal inter-arrival time 200). Servicing each of the mentioned requests takes one time unit.

The WCRT of SlipCtrlTask in the worst-case phasing (all interrupts occurred and the longest non-preemptable block just began) is $17 (2 + 10 + 2 \times 1 + 1 + 1 + 1)$.

Notice that the timer ISR is considered twice, since it is invoked twice prior the *SlipCtrlTask* is finished (after 17 time units).

8.2 WCRT of the *SlipCtrlTask* by model-checking approach

Contrary to the scheduling theory based approach presented in the previous section, the model-checking approach considers only the possible phasing determined by the

detailed model of the controlled environment (the gearbox mechanism: *Clutch*, *SelectServo* and *ShiftServo*) and the control algorithm (*GearBoxCtrlTask*). It is impossible in this phasing that all interrupts are invoked by the *Clutch*, *SelectServo* and *ShiftServo* simultaneously. The model-checking tool, moreover, explores the right blocking time of the *SlipCtrlTask* by lower priority tasks.

We explore the WCRT of the *SlipCtrlTask* in the following way: The response time of the task is measured by the clock variable $rt0$, which is reset when the *SlipCtrlTask* is activated in *ISR*, at the transition leading to the location *ActivateTask0*. Then the following property is verified: “Always, when the end of the *SlipCtrlTask* (location *End*) is reached, the inequality $rt0 \leq \text{WCRT}$ holds”. This property is formalized in the UPPAAL requirement specification language as follows:

$A[] \text{SlipCtrlTask.End imply } rt0 \leq \text{WCRT}$.

Then the verification is made by UPPAAL for the particular value of the WCRT. First, the value of the WCRT must be estimated by the designer and then, if the formula is satisfied, its value can be decreased. The smallest value of the WCRT can be found in several iterations. Notice that algorithms also exist for parametric model-checking verifying whether a state is reachable in a model with an uncertain parameter (the WCRT in the observer automaton). However, this problem is undecidable in general (Alur et al. 1993). The smallest value of the WCRT is therefore found by the interval bisection. *The WCRT of SlipCtrlTask obtained by the model-checking approach is 5*. This result is a significantly smaller value than in the case of the scheduling theory approach (see Sect. 8.1). This result is valid only in the case when the *SlipCtrlTask* is never activated more than once before it is finished. This requirement is expressed by the property $A[] \text{nActivated[SlipCtrlTaskID]} < 1$ and it has been successfully verified in the proposed model.

Both of these properties have been successfully verified by the model-checker UPPAAL 3.4.9 running on Windows 2000 on PC AMD Athlon 1 GHz, with 1.3 GB RAM. The time required for verification of both of these properties is 7 seconds. The required memory is 57 MB.

8.3 WCRT of GearBoxCtrlTask by scheduling theory based approach

It is clear from the *GearBoxCtrlTask* pseudocode listed in Fig. 14 that the task suspends itself several times while waiting for external events. This fact must be considered in the WCRT analysis since the suspended task had to compete for the processor again after the end of the suspension. The worst-case execution path of *GearBoxCtrlTask* is in Fig. 23. It consists of six *computations* (*Comp1* to *Comp6*) separated by five self-suspensions (*WaitEvent(...)*). The WCET of all *computations* are 1 and the worst-case self-suspension times (WCSST) are 200.

Palencia and Harbour (1998) proposed an offset-based approach for response-time analysis of tasks with static and dynamic offsets that can be successfully applied to tasks with self-suspension. The response-time analysis based on this approach has been implemented in the tool MAST (González Harbour et al. 2001) that can be downloaded at <http://mast.unican.es/mast.html>. We have used this tool to compute the WCRT of *GearBoxCtrlTask*. *The WCRT of GearBoxCtrlTask obtained by offset-based approach is 1021*.

Fig. 23 *GearBoxCtrlTask*
worst-case execution path

```

GearBoxCtrlTask – worst-case execution path
{
  Comp1;                               // WCET1 = 1
  WaitEvent (ClutchEvent);             // WCSST1 = 200
  Comp2;                               // WCET2 = 1
  WaitEvent (ShiftServoEvent);         // WCSST2 = 200
  Comp3;                               // WCET3 = 1
  WaitEvent (SelectServoEvent);        // WCSST3 = 2*100
  Comp4;                               // WCET4 = 1
  WaitEvent (ShiftServoEvent);         // WCSST4 = 200
  Comp5;                               // WCET5 = 1
  WaitEvent (ClutchEvent);             // WCSST5 = 200
  Comp6;                               // WCET6 = 1
};

```

Realize however that the offset-based approach does not consider branching affected by values of variables. Since the contribution of the higher priority task *SlipCtrlTask* to the WCRT of the *GearBoxCtrlTask* strongly depends on the value of variable *ClutchState* (see Fig. 11), and the blocking by the lower priority task *SelectGearTask* strongly depends on value of variable *GBReady* (see Fig. 12), the model used for analysis by the tool MAST has been manually modified in the following way: The task *SelectGearTask* is not considered in the model, since it never blocks any computation of *GearBoxCtrlTask* (*GearBoxCtrlTask* is started at the end of *SelectGearTask* and the execution of *SelectGearTask* is blocked by the value of the variable *GBReady* that is zero during the execution of the whole *GearBoxCtrlTask*). The task *SlipCtrlTask* cannot affect computations *Comp2* to *Comp5* of the task *GearBoxCtrlTask* due to value of the variable *ClutchState* (see Fig. 14). The contribution of the task *SlipCtrlTask* is therefore involved in the WCET of computations *Comp1* and *Comp6* and the task *SlipCtrlTask* is not considered in the model.

This modification of the model prevents the pessimisms of the analysis caused by not considering the branching affected by values of variables. Realize however that such modification would be hard or even impossible for a more complex system structure and cannot be therefore understood as a systematic approach to analysis. We have done it only for a fair comparison with the model-checking approach. The WCRT of *GearBoxCtrlTask* obtained by offset-based approach without described modification is 1034.

8.4 WCRT of the *GearBoxCtrlTask* by model-checking approach

We explore the WCRT of the *GearBoxCtrlTask* in a similar way as in the case of *SlipCtrlTask*. The response time of *GearBoxCtrlTask* is measured by the clock variable *rtI*, which is reset when the *GearBoxCtrlTask* is activated in *SelectGearTask*, at the transition from location *if2* to location *ActivateTask*. Then the following property is verified: “Always, when the end of *GearBoxCtrlTask* (location *End*) is reached, the inequality $rt \leq \text{WCRT}$ holds”. This property is formalized in the UPPAAL requirement specification language as follows:

A[] *GearBoxCtrlTask.End* imply $rt \leq \text{WCRT}$.

Then the verification is made by UPPAAL for the particular value of the WCRT. The smallest value of the WCRT of *GearBoxCtrlTask* is 1011. This result is valid only in the case when *GearBoxCtrlTask* is never activated more than once before it is finished. This requirement is expressed by the property $A[] nActivated[GearBoxCtrlTaskID] < 1$ and it has been successfully verified in the proposed model.

Both of these properties have been successfully verified by model-checker UPPAAL 3.4.9 running on Windows 2000 on PC AMD Athlon 1 GHz, with 1.3 GB RAM. The time required for verification of both of these properties is 8 seconds. The required memory is 61 MB.

The difference between the WCRT found by model-checking (1011) and scheduling theory (1021) seems to be minor. Realize however that the biggest part of the WCRT corresponds to the duration of the self-suspension (5×200). The pessimism of the scheduling theory approach (1021–1011) is therefore comparable with the exact time when the task is ready or executed ($1011 - 5 \times 200$).

9 Conclusions

In this paper, we have demonstrated, how timed automata can be used for modeling of multitasking, non-preemptive applications. The complex time and logical properties of the proposed model, considering the values of variables, behavior of the controlled environment and an internal structure of the control system tasks (e.g. “Shifting is allowed only when the clutch is opened”, “Desired gear is engaged in 1020 time units”, “System is deadlock free”, “A task is finished within X time units”, etc.), can be automatically verified by a model-checking tool.

Even when comparing to the task response time analysis, where classical scheduling theory can be applied, an advantage of the model-checking approach based on a fine grain timed automata model is that it considers the task internal structure and the controlled environment and it exhaustively analyzes its state space by a symbolic model-checking algorithm. Consequently, a more precise (less pessimistic) analysis is provided by the model-checking approach in the cases, when the analyzed application contains features that make the response time analysis pessimistic (e.g. branching in the tasks code, tasks self-suspension), or when the worst-case behavior, considered by the classical scheduling theory, can never occur in the controlled environment. It is clear however that the high memory requirements of the model-checking are preventing this method from becoming a universal response time analysis method.

An exhaustive analysis of the detailed timed automata model subjects to state space explosion (which is a general property of most formal methods (Corbett 1996)). Therefore, the proposed model is abstract as much as possible and it contains only information necessary for correct verification of the system specification. The operating system model uses only modest data structures, it does not use any clock variables, it does not allow any non-determinism and all locations are committed, which prevents paths interleaving and therefore restricts the explored state space. Moreover the OS model is scalable. Therefore only features used in the modeled application are used in the OS model. Notice also that OSEK is one of the most appropriate operating

systems to be modeled by timed automata since it is static (all objects are created at compilation time) and it is designed for modest runtime environment of embedded devices. For example events and resources do not require any queue of waiting tasks. The model of an application tasks must be designed as a compromise between the model precision and its state space size. It is necessary to limit the size of modeled data, non-determinism and number of *computations* in order to obtain the model of reasonable size. In spite of these restrictions, the model-checking approach is applicable for formal verification of realistic applications whose verification made manually by humans would be hard and error prone.

Acknowledgements We would like to thank Gerd Behrmann from the Department of Computer Science at Aalborg University, Denmark for his advice related to reducing the complexity of the ready queue model.

This work was supported by the Ministry of Education of the Czech Republic under Project 1M0567 and by Academy of Science of the Czech Republic under Project IET400750406.

References

- Alur R, Dill DL (1994) A theory of timed automata. *Theor Comput Sci* 126:183–235
- Alur R, Henzinger TA, Vardi MY (1993) Parametric real-time reasoning. In: *Proceedings of the 25th ACM symposium on theory of computing*, pp 592–601
- Alvarez JM, Diaz M, Lopis L, Pimentel E, Troya JM (2003) Integrating schedulability analysis and design techniques in SDL. *Real-Time Syst* 24:267–302
- Bailey CM, Burns A, Wellings AJ, Forsyth CH (1995) A performance analysis of a hard real-time system. *Control Eng Pract* 3(4):447–464
- Berard B, Bidoit M, Finkel A, Laroussinie F, Petit A, Petrucci L, Schnoebelen P, McKenzie P (2001) *Systems and software verification: model-checking techniques and tools*. Springer, New York
- Bucci G, Fedeli A, Sassoli L, Vicario E (2004) Timed state space analysis of real-time preemptive systems. *IEEE Trans Softw Eng* 30(2):97–111
- Buttazzo G (1997) *Hard real-time computing systems: predictable scheduling algorithms and applications*. Kluwer Academic, Boston
- Campos S, Clarke E (1999) Analysis and verification of real-time systems using quantitative symbolic algorithms. *J Softw Tools Technol Transf* 2(3):260–269
- Corbett JC (1996) Timing analysis of Ada tasking programs. *IEEE Trans Softw Eng* 22(7):461–483
- Daws C, Olivero A, Tripakis S, Yovine S (1996) The tool kronos. In: *Proceedings of hybrid systems III, verification and control*. Lecture notes in computer science, vol 1066. Springer, New York, pp 208–219
- Fersman E, Pettersson P, Yi W (2002) Timed automata with asynchronous processes: schedulability and decidability. In: *Proceedings of 8th international conference on tools and algorithms for the construction and analysis of systems, TACAS 2002*. Lecture notes in computer science, vol 2280. Springer, New York, pp 67–82
- Fersman E, Pettersson P, Yi W (2003) Schedulability analysis using two clocks. In: *Proceedings of TACAS'03*. Lecture notes in computer science, vol 2619. Springer, New York, pp 224–239
- Fredette AN, Cleaveland R (1993) RTSL: a language for real-time schedulability analysis. In: *Proceedings of the real-time systems symposium*. IEEE Computer Society Press, pp 274–283
- González Harbour M, Gutiérrez García JJ, Palencia Gutiérrez JC, Drake Moyano JM (2001) MAST: modeling and analysis suite for real time applications. In: *Proceedings of 13th euromicro conference on real-time systems*, Delft, The Netherlands. IEEE Computer Society Press, pp 125–134, June
- Henzinger T, Kopke P, Puri A, Varaiya P (1998) What's decidable about hybrid automata? *J Comput Syst Sci* 57:94–124
- Klein M, Ralya T, Pollak B, Obenza R, Harbour MG (1993) *A practitioner's handbook for real-time systems analysis*. Kluwer Academic, Dordrecht
- Králkora J, Waszniowski L, Píša P, Hanzálek Z (2004) Timed automata approach to real time distributed system verification. In: *Proceedings of 5th IEEE international workshop on factory communication systems, WFCs*, Vienna, September 22–24, 2004, pp 407–410

- Krčál P, Yi W (2004) Decidable and undecidable problems in schedulability analysis using timed automata. In: Proceedings of TACAS'04. Lecture notes in computer science, vol 2988. Springer, New York, pp 236–250
- Larsen KG, Pettersson P, Yi W (1995) Model-checking for real-time systems. In: Proceedings of the 10th international conference on fundamentals of computation theory. Lecture notes in computer science, vol 965. Springer, New York, pp 62–88
- Larsen KG, Pettersson P, Yi W (1997) UPPAAL in a nutshell. *Int J Softw Tools Technol Transf* 1(1/2):134–152
- Larsen KG, Larson F, Pettersson P, Yi W (2003) Compact data structures and state-space reduction for model-checking real-time systems. *Real-Time Syst* 25:255–275
- Lime D, Roux OH (2004) A translation based method for the timed analysis of scheduling extended time Petri nets. In: Proceedings of the 25th IEEE international real-time systems symposium, December 2004, Lisbon, Portugal, pp 187–196
- Liu JWS (2000) Real-time systems. Prentice-Hall, Upper Saddle River
- Lundqvist K, Asplund L (2003) A ravenscar-compliant run-time kernel for safety-critical systems. *Real-Time Syst J* 24(1):29–54
- OSEK (2005) OSEK/VDX operating system specification 2.2.3, <http://www.osek-vdx.org/>
- Palencia JC, Harbour G (1998) Schedulability analysis for tasks with static and dynamic offsets. In: Proceedings of the 19th IEEE real-time systems symposium. IEEE Computer Society Press, p 26
- Sha L, Klein M, Goodenough J (1991) Rate monotonic analysis for real-time systems. In: Foundations of real-time computing: scheduling and resource management. Kluwer Academic, Boston, pp 129–155
- Tindell K, Clark J (1994) Holistic schedulability analysis for distributed hard real-time systems. *Microprocess Microprogram* 50(2–3):117–134
- Wang S, Tsai G (2004) Specification and timing analysis of real-time systems. *Real-Time Syst* 28:69–90
- Waszniowski L, Hanzálek Z (2005) Over-approximate model of multitasking application based on timed automata using only one clock. In: Proceedings 19th IEEE international parallel and distributed processing symposium IPDPS 2005 (Workshop 2: parallel and distributed real-time systems). IEEE Computer Society, p 128



Libor Waszniowski graduated in technical cybernetics from the Czech Technical University in Prague (CTU) in 2000 and received the PhD degree in electrical engineering and informatics from CTU in 2006. Currently, he is a researcher with the Center for Applied Cybernetics at CTU. His main research activities are in formal methods for real-time systems design and analysis and rapid application development.



Zdeněk Hanzálek is a member of the IEEE and the IEEE Computer Society. He obtained the Diploma in Electrical Engineering from the Czech Technical University (CTU) in Prague in 1990. He obtained his PhD degree in Control Engineering from the CTU in Prague and PhD degree in Industrial Informatics from the Université Paul Sabatier Toulouse. He was with LAAS—Laboratoire d'Analyse et d'Architecture des Systèmes in Toulouse (1992 to 1997) and with LAG INPG—Institut National Polytechnique de Grenoble (1998 to 2000). In 2005, he obtained Doc. degree at the Czech Technical University in Prague. His research interests include real-time systems and scheduling.