

Property-Preservation Synthesis for Unified Control- and Data-Oriented Models*

Oana Florescu, Jeroen Voeten, Henk Corporaal
Information and Communication Systems Group, Faculty of Electrical Engineering
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands

Abstract

In the Software/Hardware Engineering model-driven design methodology, the preservation of real-time system properties can be guaranteed up to a small time-deviation in the model synthesis. Therefore, this methodology is well suited for the design of control-systems in which execution times of actions are small; thus the time-deviations obtained are small. However, in systems containing time-intensive computations, the time-deviations become large and, consequently, the real-time properties are much weakened. This paper gives an initial idea for obtaining stronger property-preservation by abstracting from the internal actions of a system and counting only the observable actions for the time-deviation. In this way, a unified way of analysis and synthesis of a larger area of real-time applications can be obtained, which would allow designers to reason about different properties of systems.

1 Introduction

The main purpose of modelling is to help engineers understand the relevant aspects of a system, while avoiding the expense and trouble of actually building it. Whereas traditional forms of engineering have a well-established modelling methodology, software engineering, and particularly real-time embedded software, is still an emerging discipline. Although it is applied to increasingly complex systems, its modelling techniques are neither mature nor reliable yet [SM03]. Nevertheless, software models have a unique and remarkable advantage over most other engineering models: they can be used to automatically generate the realisation of the system modelled, which is an executable program for a particular platform. Starting with a simplified and highly abstract model, refinements can be carried out until a complete specification is obtained, including all the details necessary in the final product. From such a detailed specification adequate computer tools can generate an implementation.

The Model-Driven Architecture (MDA) initiative of the Object Management Group (OMG) [OMG01] shows that the interest in technologies for supporting model-driven development has increased. In the development trajectory proposed in MDA, system models are made in very early stages, as shown in Figure 1, to help designers in reasoning about different trade-offs. By making design decisions and adding the corresponding details to the model, the design space is narrowed. The software models are kept independent from the platform on which their implementation would run, as long as possible in this design trajectory. The platform-independence provides the flexibility of reusing the design model and/or of targeting it to a different platform. Moreover, it may allow the prediction from the model of a suitable platform.

*This work is being carried out as part of the Boderc project under the responsibility of the Embedded Systems Institute. This project is partially supported by the Netherlands Ministry of Economic Affairs under the Senter TS program.

Going lower in the design pyramid by increasing the number of details, a complete specification can be obtained, from which the software implementation can be automatically generated.

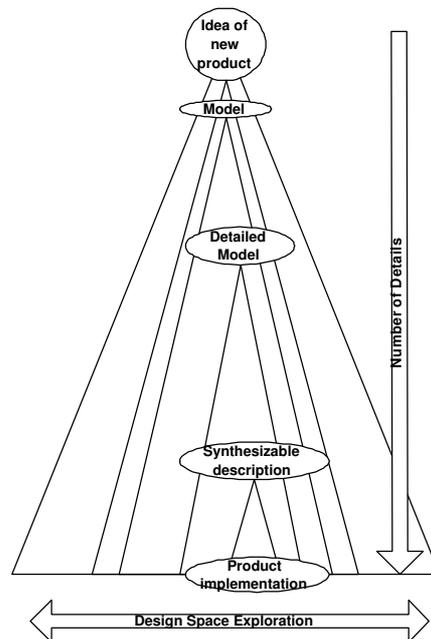


Figure 1: Model-Driven Approach

The software components employed in the embedded systems, like the ones in cars, airplanes, printer/copier machines, or medical devices, are supposed to synchronise and coordinate different processes and activities. Therefore, their behaviour must meet hard timing constraints, either for people's safety or simply to ensure things work correctly. Usually, a hard real-time software component must work together with other software and hardware components to obtain the specified behaviour. Its correctness depends on both the logical result obtained and the moment in time when the result was ready. Experience showed that existing model-driven development approaches for software systems are not suited to cope with real-time system design. Traditional design approaches proved themselves unable to capture adequately both functional and non-functional (timing) characteristics of a system, while abstracting from low-level details. For predictably designing such systems, an appropriate methodology needs to provide [HVVvB03]: (i) a suitable modelling technique that can appropriately capture functional and timing properties in models in order to formally analyse them, and (ii) a mechanism to generate the implementation from the model while preserving the properties analysed, phase also known as model synthesis.

The Software/Hardware Engineering [vdPV97] is a model-driven methodology suitable for analysis and synthesis of real-time systems in which actions need small execution time. In this paper, we give an initial idea for the synthesis, using the same methodology, of system models containing both actions and time-intensive computations while still preserving the real-time properties analysed. We make observations regarding the possibility of code generation from models which are equivalent from the perspective of an external user. By applying this idea, we can have a predictable and unified trajectory from a model towards a property-preserving system realisation for a large area of real-time applications (both control-oriented and data-oriented).

The paper is organised as follows. Section 2 discusses related research. Section 3 presents the technique used for formal modelling of systems. Section 4 shows how the properties of control-oriented systems models are preserved in their implementations. Section 5 discusses a way to synthesise models of applications that contain time-intensive computations. Conclusions are drawn in Section 6.

2 Related Research

In the context of model-driven approaches for software development, the Unified Modelling Language (UML) [OMG03] has been adopted as a standard facility for constructing models of object-oriented software. UML proved to be suitable for modelling the functional aspects of a system, which can also be correctly synthesised. However, extensions were defined to it to provide a standardised way of denoting non-functional (timing) aspects for real-time systems as well [OMG05]. Although UML is largely used both in industry and academia, it has shown to be inappropriate to model and synthesise hard real-time systems as it lacks the formal semantics, restraining it from a proper analysis of a system behaviour, and the properties-preservation code generation mechanism.

For modelling purposes, a number of techniques and theories were proposed, targeting a certain view over a system, e.g. correctness analysis, scheduling analysis. We give here one example in order to emphasise the benefits of these techniques, but also the fact they are not suited for synthesis.

Classic scheduling theory [But97] provides techniques for the analysis of timing behaviour of a system and for the scheduling of its tasks onto the target platform, such that the timing constraints are satisfied. Real-time tasks are usually assumed to be periodic, i.e. tasks arrive (and will be computed) with fixed periodicity; therefore, well-studied methods, like rate monotonic scheduling, can be applied. Nevertheless, analysis of such models often yields pessimistic results and it is not able to handle non-periodic tasks with non-deterministic behaviours. Moreover, the models analysed by classical scheduling analysis do not incorporate information about the functionality of tasks which makes them unsuitable for the model synthesis.

A way to relax the stringent constraints on task arrival times is by using automata with timing constraints to model task arrival patterns. The model can describe concurrency and synchronisation of periodic, sporadic, preemptive or non-preemptive real-time tasks with or without precedence constraints. An automaton is schedulable if there exists a scheduling strategy such that all possible sequences of events accepted by the automaton are schedulable (all associated tasks can be computed within deadlines). Based on the results obtained for schedulability analysis on timed automata, the TIMES tool [AFM⁺03] has been designed for schedulability analysis and synthesis of real-time systems. A model consists of (i) a set of application tasks whose executions may be required to meet different timing, precedence and resource constraints, (ii) a network of timed automata describing the task arrival patterns and (iii) a preemptive or non-preemptive scheduling policy. From such a model, the TIMES tool can generate a scheduler and compute the worst-case response time for all tasks. Nevertheless, TIMES tool does not have enough expressive power to describe the data computations involved in a system. Because it relies on exhaustive analysis, all the things that might lead to state space explosion problems need to be left out. Therefore, TIMES analysis and synthesis do not scale up and cannot be applied to any kind of system development.

3 Real-Time Systems Models

The Software/Hardware Engineering (SHE) [vdPV97] is a system-level design methodology that uses a UML profile to formulate the concepts needed for the realisation of the requested functionality of a system. The UML profile smoothes the application of the Parallel Object-Oriented Specification Language (POOSL) [vdPV97] to develop an executable model, as shown in Figure 2. POOSL formalises the behaviour specified in informal UML diagrams, establishing a formal executable model. The realisation of the system can be generated from this model using the Rotalumis [vB02] tool.

POOSL is equipped with mathematical semantics that can formally describe concurrency,

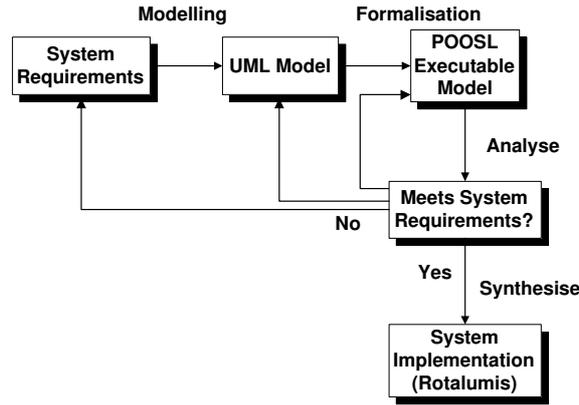


Figure 2: SHE method for real-time systems design

distribution, communication, timing and functional features of a system in a single executable model, using a small set of very powerful primitives. Primitives can be combined in an unrestricted fashion and any combination has a precisely defined meaning. The formal semantics guarantees a unique and unambiguous interpretation of a POOSL model, guided by semantical axioms and rules, independent of the underlying execution platform. The importance of the formal semantics of a modelling language in supporting the predictability of the system design process is investigated in [HVF⁺].

POOSL consists of a process part and a data part. The process part (processes and clusters), based on a real-time extension of the process algebra CCS [Mil89], is used to specify the real-time behaviour of active components. The data part, based upon traditional concepts of sequential object-oriented programming, is used to specify the information that is generated, exchanged, interpreted or modified by the active components.

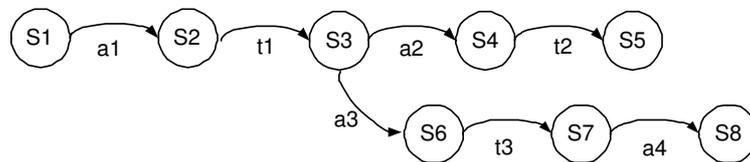


Figure 3: A timed labelled transition system

The semantics of POOSL is defined as a timed labelled transition system, as the example in Figure 3 shows, where $S1 - S8$ represent states of the system, $a1 - a4$ action transitions and $t1 - t3$ time transitions. A timed labelled transition system represents an abstract view over the system, considering it as an entity having some internal state and, depending on that state, it can engage in transitions leading to other states. Such a transition might be autonomous or stimulated by the environment. When action transitions take place, the state of the system changes by changing its content (for example, when an event happens, certain parameters of the system get a different value). In case of time transitions, the system actually resides in the same state, without changing its content, for the time interval specified.

In a model based on a timed labelled transition system, the execution has two phases, as shown in Figure 4: the state of a system changes either by asynchronously executing atomic actions, such as communication or data computation, without passage of time (phase 1), or by letting time pass synchronously without any action being performed (phase 2).

A run over a transition system represents a timed trace, as the one in Figure 5, where each action is executed at a particular time. As there are many possible runs due to the parallelism and non-deterministic choices that can be expressed, a POOSL model represents, in fact, a set of

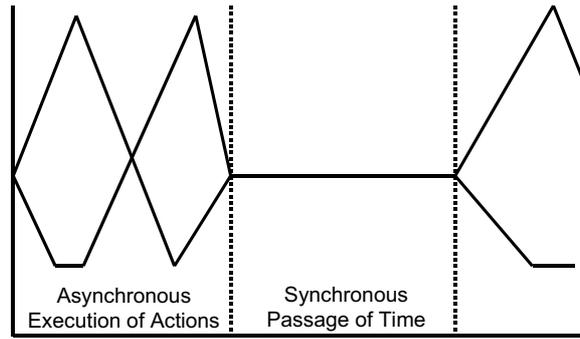


Figure 4: Two phases of the execution of a POOSL model

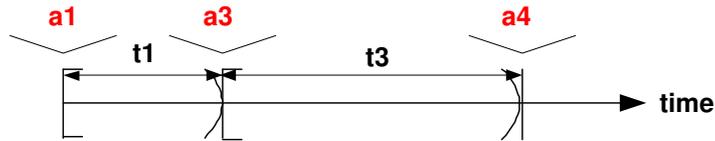


Figure 5: A timed trace of the system

timed traces. If all the traces of the model satisfy a real-time property (for example, a particular event happens at a certain moment), then the model of a system has that particular real-time property.

An example of a POOSL specification¹ for the control of a simple system is given below. The controller reads some data x from the environment, performs computations with it and delivers the result y back to the environment at a certain time.

```

in?input(x);      /* x is received as a message */
computation(x)(y); /* x is the input & y is the output of computation */
delay deadline;   /* wait for an amount of time "deadline" */
out!output(y);    /* y is sent as a message */

```

The model can be graphically represented using the UML stereotype `<<capsule>>`, as depicted in Figure 6, where the small black squares represent output ports, and the white ones input ports.

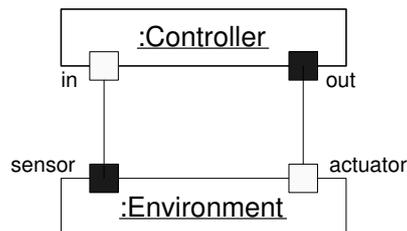


Figure 6: The graphical representation of a POOSL model

For the POOSL specification given as example, the timed labelled transition system looks like in Figure 7. According to the semantics of the language, a timed trace of the model is the one shown in Figure 8. `in?input(x)` and `computation(x)(y)` are executed in this exact ordering, without consuming any time and at the same instant $t1$. Then, time passes for `deadline` units ($t2 = t1 + \text{deadline}$) and, finally, `out!output(y)` is instantly performed at $t2$.

¹Note that the notations in a POOSL specification are CCS alike.

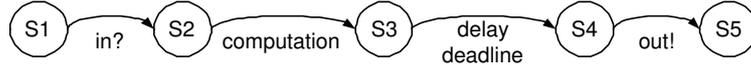


Figure 7: The timed labelled transition system of the model

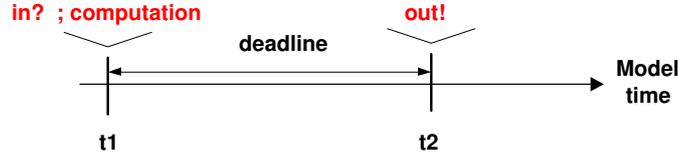


Figure 8: The meaning of the model

4 From a Model to Its Realisation

As mentioned in the previous section, a real-time system can be formalised as a set of timed traces. If two timed traces have the same sequence of actions, a notion of distance between them is defined. The distance represents the largest deviation between the ending points of corresponding time intervals, as shown in Figure 9. Two timed traces whose distance between them is equal to ϵ are called ϵ -close. If two execution traces are ϵ -close and one of the traces satisfies a real-time property², then this property, weakened up to ϵ^3 , is satisfied in the second trace as well. This result was mathematically proved in [HVG03].

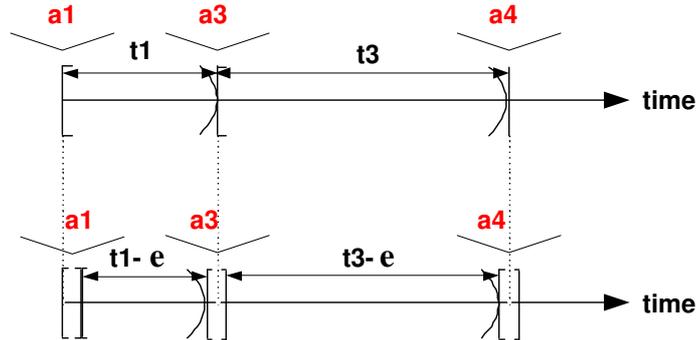


Figure 9: Timed traces ϵ -close

Both the model and the realisation of a system can be viewed as sets of timed traces. To obtain an implementation of a system which preserves the properties analysed in its model, thus an implementation consistent with the model, two things must be achieved: (i) to generate a trace in the implementation from the set of execution traces of the model; (ii) to make the corresponding traces in the model and in the implementation to be ϵ -close.

A mechanism of generating a trace from a POOSL model was proposed and proved correct in [Gei02]. The data part of a POOSL model is directly translated into corresponding C++ expressions. Each process in the model is represented by a C++ data structure named process execution tree (PET) whose nodes represent statements in the specification of behaviour. During the evolution of the system, a PET scheduler makes choices for granting actions or time transitions, while each PET adjusts its internal state according to the choice of the PET scheduler. This mechanism guarantees that the realisation of the model generated by the code generation tool is a trace from the model.

²An example of a real-time property is that a certain action happens at a particular moment in time.

³If a property P is true in the first trace in the interval $[t1, t2]$, the other trace satisfies P in the interval $[t1 - \epsilon/2, t2 + \epsilon/2]$.

However, as actions in a model are timeless, whereas, in reality, it will always take a certain amount of time to execute them, between the corresponding traces there appears a *time-deviation*. If the distance between these two traces is ϵ (ϵ -hypothesis), then *all* the properties of the model are preserved up to ϵ in the implementation.

To generate the implementation ϵ -close to the POOSL model, the code generation tool, Rotalumis [vB02], synchronises the *model time* with the *physical time*. As shown in the UML sequence diagram from Figure 10, all the actions that happen instantly in the model at a certain time t (in Figure 8, `in?` and `computation` happen at model time t_1 , `out!` at t_2), are executed within a small ϵ amount of time around the corresponding moment in physical time (`in?` and `computation` are executed in ϵ_1 around physical time t_1 , `out!` in ϵ_2 around t_2). To maintain the synchronisation between model time and physical time, `delays` cannot be executed in the implementation exactly as specified in the model (*deadline* units of time), but physical time passes until the next corresponding moment in the model time is reached (the delay *deadline* = $t_2 - t_1$ is shortened to $t_2 - t_1 - \epsilon_1$).

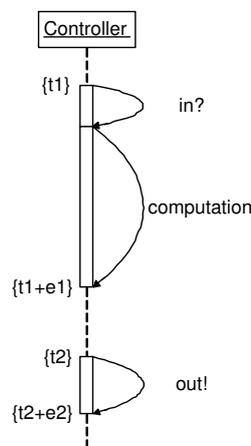


Figure 10: Implementation in physical time

The size of the maximum time-deviation between a model and its implementation can be obtained at the time of generation and execution of code by using measurements. Another approach is to estimate it from the model, based on the Y-chart scheme depicted in Figure 11, that contains the models of both the real-time application and the target platform [FVHC04]. The time-deviation depends on how many actions need to be executed at the same time in the model, as well as on their execution times. If the value obtained for ϵ is considered too large, either the implementation is generated for a higher performance platform, on which the execution of all the actions takes less time, the mapping is changed or the model is re-designed.

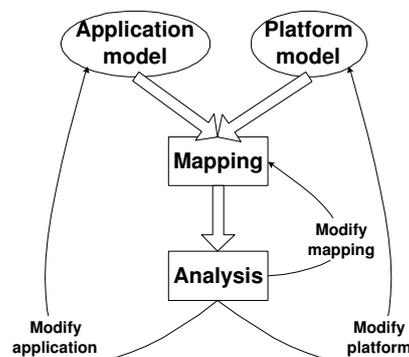


Figure 11: Y-chart scheme for hard real-time systems

5 Realisation of Systems with Time-Intensive Computations

In a model of a real-time system, usually, a distinction can be made between *actions* and *time-intensive computations*. *Action* is the name given to an "activity" specified in the model that needs small execution time on the target platform (e.g. a control action). On the other hand, *time-intensive computations* are the "activities" specified in a model that usually need a considerable amount of time for execution, as it is the case of the `computation` in Figure 10. In case of real-time systems containing such computations, the time-deviation between the model and the implementation is usually large. Therefore, with the current generation of code, the properties analysed in the model will be much weakened in the implementation.

Nevertheless, in data-oriented real-time applications, many time-intensive computations are needed to be modelled (for example, different multimedia algorithms must be applied on a stream of data). For this kind of system, it is not intended for the computations to be instantaneous, but to be finished before a *deadline*, when the results must be given to the environment (like in the example given in Section 3).

Two systems are called observational equivalent if they cannot be distinguished between them through the interaction of a user with each of them. They have the same observable properties⁴ and the same set of timed traces with respect to these properties. Therefore, an implementation preserving the observable properties of a model preserves the observable properties of the observational equivalent one.

Based on this insight, in the case of systems with time-intensive computations, instead of generating an implementation for the original model, we could generate the implementation for an observational equivalent one. Below, we give a specification which is observational equivalent with the example given in Section 3:

```
in?input(x);
computation1(x)(y1);
delay deadline1;
computation2(y1)(y2);
delay deadline2;
computation3(y2)(y);
delay deadline3;
out!output(y);
```

The `computation` is split, for example, into three smaller parts (`computation1`, `computation2` and `computation3`) that, put in sequence, form the original computation specified in the model. After each small computation, a certain amount of time delay follows (`deadline1`, `deadline2` and `deadline3`) and the sum of all delays makes the original delay amount (`deadline = deadline 1 + deadline2 + deadline3`). A timed trace of this system is given in Figure 12.

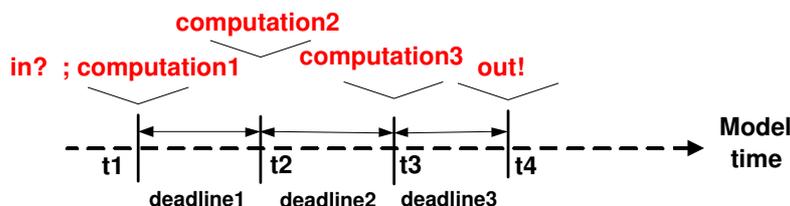


Figure 12: Model timed trace

The two systems modelled are, obviously, observational equivalent for a user for whom it is important when the input data `x` is read from the environment, what is the flow of computations

⁴A user can see the same properties by interacting with the systems.

performed on x , and when the final result y is available. For this system, the existing synthesis mechanism for POOSL models, which relies on the ϵ -closeness between traces for the properties-preservation, can be applied. To obtain an implementation trace ϵ -close to its corresponding trace in the model, as shown in the previous section, a synchronisation of each moment in the model time when an action happens with the corresponding physical time, up to ϵ , is realised, as shown in Figure 13. For the implementation of the original model, there are only two synchronisation points, t_1 and t_2 from Figure 10, and the time-deviation is large. For the observational equivalent model, in Figure 13, there are four synchronisation points, t_1 , t_2 , t_3 and t_4 , and the time-deviation for each of them is smaller. Therefore, over the whole system, the properties are stronger preserved in the realisation of the second model.

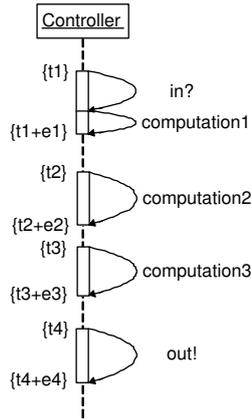


Figure 13: Implementation in physical time

From the perspective of the code generation tool, looking at Figure 13, what it actually has to do is to generate a trace in which the execution of the `computation` (made of `computation1`, `computation2` and `computation3`) starts immediately after reading x from the environment, continues more or less without stopping, and finishes before the moment the result y must be given back to the environment. In other words, `deadline` represents the *deadline* of the computation, and only the observable actions of the system are synchronised in the physical time, as depicted in Figure 14. The time needed for the execution of computation does not have to count against the size of the time-deviation between model and implementation because, for the observational equivalent model in Figure 13, the value of ϵ is small and it is determined by the execution time of `out!output(y)`.

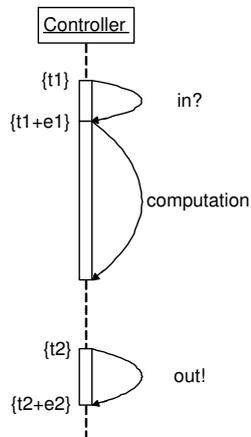


Figure 14: A possible execution that still preserves the properties

As shown in this simple example, the implementation of a model containing time-intensive

computations can be generated from an equivalent model that has the same observational behaviour and the same properties as the original one. Under these circumstances, we can define *actions* and *computations* slightly different than at the beginning of this section. We name *actions* those activities that can be observed by a user interacting with the system and, therefore, their moments of execution in the model time must be synchronised with the physical time. On the other hand, *computations* are the internal activities of a system that a user cannot observe and who need to be scheduled for execution such that they can meet their deadlines. Moreover, if they are still running, they can be preempted by an action whose model time must be synchronised with the physical time. By abstracting from the internal actions of the model and synchronising the model time with the physical time only for the moments when observable actions happen, the observable properties of a model can be preserved; thus the code generation tool can handle the model synthesis of data-oriented applications as well.

However, it is not always possible to execute a computation until a *deadline*. For a specification like

```

abort
  (computation(x)(y); delay deadline)
with p?urgentMessage;

```

according to the formal semantics of the language, the urgent message can arrive either before the execution of `computation(x)(y)` or during the `delay`, which means after the execution of `computation` has finished. If the `computation` has a `deadline` in the implementation, then, at the time the `urgentMessage` appears, the execution of `computation` must be preempted. The `computation` will not be allowed to continue; thus the state in which the system realisation will be in that moment will not be a state present in the model. In this case, the relaxation of the timing constraints is not possible because there is no equivalence relation between this model and another one that has the computation split into smaller parts. Therefore, the execution time of the computation contributes to the total time-deviation between model and implementation of this system.

Nevertheless, the mechanism that we propose in this paper for the synthesis of real-time systems with time-intensive computations has the benefit of using an existing methodology, without changing the syntax, the semantics of the modelling language or anything else, just by relaxing the constraints on the properties to be preserved. However, work needs to be done in order to formalise these ideas and to mathematically prove they are true. Moreover, a mechanism of identification of the observational equivalent system whose implementation is the same with the one of the given model is required.

To analyse a model with different kinds of activities (taking longer or shorter execution time), the Y-chart scheme can be used again. Such a unified model helps designers in reasoning about aspects like what is the largest time-deviation (ϵ) that the system can allow, or what is an appropriate scheduling of the time-intensive computations, as shown in [FVC04].

6 Conclusions and Future Work

To achieve a predictable design of real-time embedded systems, a unified executable model, capturing both functional and timing aspects of a system, is suitable to allow engineers to reason about different properties in a unified manner. Moreover, such a model must be easily refinable towards a complete system specification, from which the implementation can be automatically obtained.

In this paper, we have presented how the Software/Hardware Engineering methodology can be used for the modelling, analysis and synthesis of a large area of real-time systems (control-oriented, data-oriented applications). The POOSL modelling language allows specification of

both timing and functional aspects of systems, while the ϵ -hypothesis guarantees the preservation of properties between two timed systems with a small time-deviation. By satisfying the ϵ -hypothesis, the code generation tool, Rotalumis, succeeds in synthesising an implementation of a model preserving *all* the properties, in case the actions specified are not time-consuming. For the data-oriented applications, we give an initial idea on generating the realisation from a model which is observational equivalent with the original one, but who has the advantage that the time-deviation obtained for it is smaller. In fact, we suggest that it is possible to make an abstraction from the internal actions of the system and synchronise the physical time with the model time only for the observable actions. Moreover, this realisation would preserve the observable properties of the original real-time system.

As future research, we aim at formalising and proving this mechanism and, moreover, at giving a mathematical definition for the circumstances when computations can be safely pre-empted by actions. Furthermore, we want to prove that such an implementation still preserves the properties of the model, and to adapt the code generation tool to work according to the proposed mechanism.

References

- [AFM⁺03] Tobias Amnell, Elena Fersman, Leonid Mokrushin, Paul Pettersson, and Wang Yi. TIMES: a tool for schedulability analysis and code generation of real-time systems. In *1st International Workshop on Formal Modeling and Analysis of Timed Systems, FORMATS 2003*, September 2003.
- [But97] Giorgio C. Buttazzo. *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Kluwer Academic Publishers, Boston MA, 1997.
- [FVC04] Oana Florescu, Jeroen P.M. Voeten, and Henk Corporaal. A unified model for analysis of real-time properties. In *1st International Symposium on Leveraging Applications of Formal Methods (ISoLa 2004)*, pages 220–227, November 2004.
- [FVHC04] Oana Florescu, Jeroen P.M. Voeten, Jinfeng Huang, and Henk Corporaal. Error estimation in model-driven development for real-time software. In *Forum on Specification & Design Languages 2004 (FDL'04)*, pages 228–239, September 2004.
- [Gei02] Marc G.W. Geilen. *Formal Techniques for Verification of Complex Real-Time Systems*. PhD thesis, Eindhoven University of Technology, Eindhoven NL, 2002.
- [HVF⁺] Jinfeng Huang, Jeroen P.M. Voeten, Oana Florescu, Piet van der Putten, and Henk Corporaal. *Advances in Design and Specification Languages for SoCs (Best of FDL'04)*, chapter Predictability in real-time system development. Kluwer Academic Publishers. to be published in 2005.
- [HVG03] Jinfeng Huang, Jeroen P.M. Voeten, and Marc C.W. Geilen. Real-time property preservation in approximations of timed systems. In *1st ACM & IEEE International Conference on Formal Methods and Models for Codesign (MEMOCODE'2003)*, June 2003.
- [HVVvB03] Jinfeng Huang, Jeroen P.M. Voeten, Andre Ventevogel, and Leo J. van Bokhoven. Platform-independent design for embedded real-time systems. In *Forum on Specification & Design Languages 2003 (FDL'03)*, September 2003.
- [Mil89] Robin Milner. *Communication and Concurrency*. Prentice Hall International Series in Computer Science. Prentice Hall, 1989.

- [OMG01] OMG. *Model Driven Architecture (MDA)*. OMG document ormsc/2001-07-01, Needham MA, 2001.
- [OMG03] OMG. *Unified Modeling Language (UML) - Version 1.5*. OMG document formal/2003-03-01, Needham MA, 2003.
- [OMG05] OMG. *UML Profile for Schedulability, Performance, and Time Specification - Version 1.1*. OMG document formal/2005-01-02, 2005.
- [SM03] Bran Selic and Leo Motus. Using models in real-time software design. *IEEE Control Systems Magazine*, 23(3):31–42, June 2003.
- [vB02] Leo J. van Bokhoven. *Constructive Tool Design for Formal Languages: From Semantics to Executing Models*. PhD thesis, Eindhoven University of Technology, Eindhoven NL, 2002.
- [vdPV97] Piet H.A. van der Putten and Jeroen P.M. Voeten. *Specification of Reactive Hardware/Software Systems*. PhD thesis, Eindhoven University of Technology, Eindhoven NL, 1997.