Aalborg Universitet

AALBORG UNIVERSITY

DENMARK

Co-Simulation of Hybrid Systems with SpaceEx and Uppaal

Bogomolov, Sergiy; Greitschus, Marius; Jensen, Peter Gjøl; Larsen, Kim Guldstrand; Mikuionis, Marius; Strump, Thomas; Tripakis, Stavros

Co-Simulation of Hybrid Systems with SpaceEx and Uppaal

Sergiy Bogomolov[1]   Marius Greitschus[2]   Peter G. Jensen[3]   Kim G. Larsen[3]
Marius Mikučionis[3]   Thomas Strump[2]   Stavros Tripakis[4]

[1]IST Austria, Austria
[2]University of Freiburg, Germany
[3]Aalborg University, Denmark
[4]Aalto University, Finland, and University of California, Berkeley, USA

Abstract

The Functional Mock-up Interface (FMI) is an industry standard which enables co-simulation of complex heterogeneous systems using multiple simulation engines. In this paper, we show how to use FMI in order to co-simulate hybrid systems modeled in the model checkers SPACEEX and UPPAAL. We show how FMI components can be automatically generated from SPACEEX and UPPAAL models. We also validate the co-simulation approach by comparing the simulations of a room heating benchmark in two cases: first, when a single model is simulated in SPACEEX; and second, when the model is split in two submodels, and co-simulated using SPACEEX and UPPAAL. Finally, we perform a measurement experiment on a composite model to show a potential for statistical model checking using stochastic co-simulations.

Keywords: FMI, hybrid system, timed automaton

1   Introduction

Despite advances in model checking and other formal verification techniques, simulation remains the workhorse for system analysis. A plethora of simulation tools are available today, from academia as well as from industry. These tools support a large variety of modeling languages, targeted at different types of systems from various disciplines (e.g., mechanical, electrical, digital, continuous or discrete, or mixes thereof). Unfortunately, these tools can rarely interoperate. This is a problem because modern cyber-physical systems are highly complex and multidisciplinary, requiring specialized modeling languages and tools from several domains.

The *Functional Mock-up Interface*[1] (FMI) is a standard developed to address this problem. FMI defines an XML schema for describing simulation components and a C API that these components must implement. The components are called *functional mock-up units*,

or FMUs. An FMU is typically generated automatically *(exported)* from some simulation tool, and corresponds to a (sub-)model designed in that tool. The sub-models/FMUs are then *imported* into a *host* simulator. The host commands the simulation by calling the API methods of the FMUs, thus effectively achieving integration of the original simulation environments. FMI supports two integration modes: (a) *model exchange*, where the host simulator is handles the numerical integration; and (b) *co-simulation*, where each FMU implements its own numerical integration mechanism (or any other internal mechanism to advance its state in time). Because each mode imposes its own requirements on FMUs (for instance, in model exchange, the FMUs must provide the host with information such as state derivatives, which are not necessary for co-simulation) the FMI APIs for the two modes are different.

In this paper, we use FMI in order to connect two state-of-the-art modeling and verification tools for cyber-physical systems: SPACEEX (Frehse et al., 2011) and UPPAAL (Larsen et al., 1997). SPACEEX is a tool for modeling and verifying *hybrid systems* (Alur et al., 1995). UPPAAL is primarily a model-checker for *timed automata* (Alur and Dill, 1994), however, it also supports statistical model-checking of hybrid systems (David et al., 2011).

Our goal is to integrate these two tools for co-simulation. That is, we want to be able to: (a) build a sub-model of the system (e.g., the model of the *plant* under control) in SPACEEX; (b) build another sub-model (e.g., the *controller*) in UPPAAL; (c) automatically generate an FMU for each sub-model; (d) import the FMUs, connect and co-simulate them in a host environment.

The motivations for connecting SPACEEX and UPPAAL in this manner are numerous. First, although both SPACEEX and UPPAAL support simulation of hybrid systems, each tool offers its own modeling language, which is not compatible with that of the other tool. Translating from one language to the other is limited to common features supported by the tools. For example, even though the frameworks CIF (Agut et al.,

---

[1]See https://www.fmi-standard.org/ for more details.

2013; Beohar et al., 2010) and HSIF (Pinto et al., 2006) solve the complexity problem of one format translation to another by performing at most two translations, the approach still suffers from the fact that UPPAAL features like committed locations and C-like function code are not supported in SPACEEX and UPPAAL has limited support for ODEs. Moreover, by using co-simulation, we are able to take advantage not just of the specific strengths of the language of each tool, but also of their native simulation engines, since each FMU is internally running essentially a "copy" of the simulation algorithm of the original tool.

As host environment we use the tool Ptolemy[2]. Ptolemy is a modeling and simulation environment for heterogenous systems (Eker et al., 2003). Recently, support has been implemented in Ptolemy for using it as a host environment for co-simulation based on FMI. FMUs (developed by other tools) can be imported into Ptolemy, connected using Ptolemy's graphical user interface, and co-simulated using an implementation of the co-simulation algorithm described by Broman et al. (2013). This algorithm has desirable properties, such as *determinacy*, namely, the fact that the results of the simulation are independent of arbitrary factors such as names of the FMUs, order of creation, or order of evaluation in the diagram.

The contributions of this paper are the following:

1. We show how FMUs can be generated automatically from models of hybrid and timed automata built in SPACEEX and UPPAAL. There are several subtleties involved in this, as hybrid and timed automata are models designed primarily with verification in mind, whereas FMI is designed for simulation and therefore imposes certain properties on FMUs, such as determinism.

2. We report on the implementation and case studies. In particular, we apply our co-simulation framework to a room heating benchmark (Fehnker and Ivancic, 2004).

3. We validate the co-simulation algorithm proposed by Broman et al. (2013) by comparing the results of the case study in two settings: (a) when the case study is modeled and simulated in a single tool, and (b) when the various components of the case study are modeled in two tools and co-simulated using our framework. We show that our co-simulation framework computes the same simulation trajectories as the setting (b) provided that the maximum simulation step size of co-simulation is sufficiently small.

4. We demonstrate how stochastic simulations can be included into the composite model with hybrid systems and applied a simple statistical measurement

---

to show the potential for statistical model checking using FMI co-simulations.

The rest of the paper is organized as follows. In Sec. 2, we introduce the necessary background on FMI for this work. Afterwards, we present our translation of SPACEEX and UPPAAL models into FMUs in Sec. 3. This is followed by the case study in Sec. 4. We discuss related work in Sec. 5. Finally, we conclude the paper in Sec. 6.

2 Background on FMI

Conceptually an FMU can be seen as a (timed) state machine. This machine has a set of input variables (or *ports*), a set of output variables, and a set of internal states. The machine interacts with its environment only by means of a clearly defined set of *interface methods*. These methods are specified in the FMI standard. For the purposes of this paper, and following the formalization presented by Broman et al. (2013), the key interface methods of FMI (for co-simulation) are:

- A method to initialize the state of the FMU. If S is the set of states of the FMU, then $\texttt{init} \in S$.

- A method $\texttt{set}$ to set a given input variable to a certain value. The signature of $\texttt{set}$ is $\texttt{set} : S \times U \times \mathbb{V} \to S$, where U is the set of input variables of the FMU, and $\mathbb{V}$ is the set of all possible values (for simplicity we ignore typing and use a single universe $\mathbb{V}$ of values for all variables). Given state s, input variable $u \in U$, and value $v \in \mathbb{V}$, $\texttt{set}(s,u,v)$ returns the new state obtained after setting u to v.

- A method $\texttt{get}$ which returns the value of a given output variable. Its signature is $\texttt{get} : S \times Y \to \mathbb{V}$, where Y is the set of output variables of the FMU. Given state s and output variable $y \in Y$, $\texttt{get}(s,y)$ returns the value of y in s.

- A method $\texttt{doStep}$ which advances the state of the machine in time. Its signature is $\texttt{doStep} : S \times \mathbb{R}_{\geq 0} \to S \times \mathbb{R}_{\geq 0}$, where $\mathbb{R}_{\geq 0}$ is the set of non-negative real numbers. The behavior of $\texttt{doStep}$ is explained below.

As said above, an FMU is essentially a state machine: the $\texttt{get}$ method corresponds to the output function of the machine, while the $\texttt{doStep}$ method corresponds to the transition function. The difference is that $\texttt{doStep}$ takes as input a *time step* $h \in \mathbb{R}_{\geq 0}$: in that sense, an FMU is a timed state machine.

The behavior of $\texttt{doStep}$ is as follows. Given state $s \in S$, and time step $h \in \mathbb{R}_{\geq 0}$, a call to $\texttt{doStep}(s,h)$ is interpreted as the co-simulation algorithm "asking" the FMU to perform a simulation step of length h. For a

---

[2]See http://ptolemy.eecs.berkeley.edu/.

number of reasons, including numerical integration issues, the FMU may "accept" or "reject" this request. If it rejects, it means that it was not able to advance time by h (but may have been able to advance time by a smaller delay $h' < h$). Formally, $\texttt{doStep}(s,h)$ returns a pair (s',h') where $s' \in S$ is a state and $h' \in \mathbb{R}_{\geq 0}$ is a time step, such that:

- either $h' = h$, which is interpreted as F having *accepted* h, and having moved to a new state s';

- or $0 \leq h' < h$, which is interpreted as F having *rejected* h, but having made partial progress up to h', and having reached a new state s'.

It is worth noting that FMUs are *deterministic* machines, in the sense that for a given sequence of inputs (i.e., a sequence of input values and time steps), the sequence of states and outputs that the machine produces is unique. This is because there is a unique initial state $\texttt{init} \in S$, and $\texttt{set}, \texttt{get}, \texttt{doStep}$ are all *total functions*. Moreover, the fact that these functions are total implies that the machine is able to accept any input at any time, therefore, it is implicitly *input-enabled*.

We also rely on zero-time steps in a sense of allowing $\texttt{doStep}(s,h)$ calls with $h = 0$ (despite that version 2.0 of the FMI standard forbids this), because they are essential for modeling discrete transitions like instantaneous mode switches in hybrid automata models.

In addition to the above, each FMU comes with a set of *input-output dependencies*, $D \subseteq U \times Y$. D specifies for each output variable which input variables it depends upon (if any): $(u,y) \in D$ means that output variable y depends on input variable u. This information is used to ensure that a network of FMUs has no cyclic dependencies, and also to determine the order in which all network values are computed during a simulation step (Broman et al., 2013).

FMI specifies the methods that every FMU must implement, but it does *not* specify the co-simulation algorithm (also called a *master algorithm*). In fact, devising such an algorithm with good properties is not a trivial problem, and has been the topic of previous work (Broman et al., 2013). In that work, two co-simulation algorithms were proposed and proved to have desirable properties, such as termination of a simulation step, and *determinacy*. The determinacy property says that the results of a simulation do not depend on the order in which the algorithm chooses to call $\texttt{doStep}$ over a set of FMUs. This ensures that the simulation results are well-defined and are not influenced by arbitrary factors such as FMU names, order of creation, geometrical position in the diagram of a graphical model, etc., as is often the case with simulation tools.

In a nutshell, the co-simulation method proposed by Broman et al. (2013) relies on the following principle. First, the co-simulation algorithm chooses a default time step, $h_{\max}$, called the *maximum step size*. Second, the algorithm saves the state of each FMU in the model (FMI specifies methods for an FMU to export and import its state, although these are optional). Assuming there are n FMUs, $F_1, ..., F_n$, the algorithm maintains n states, $s_1, ..., s_n$. Third, the algorithm calls $F_i.\texttt{doStep}(s_i, h_{\max})$ on each FMU F_i, and collects the returned time steps $h'_1, ..., h'_n$. There are two cases: either all FMUs accepted the proposed time step, i.e., $h'_1 = h'_2 = \cdots = h'_n = h_{\max}$, in which case this simulation step is over, and the algorithm proceeds to the next one; or at least one FMU F_i rejected $h_{\max}$, i.e., $h'_i < h_{\max}$ for some i. In the latter case, the algorithm computes the minimum of $h'_1, ..., h'_n$, $h_{\min} = \min\{h'_1, ..., h'_n\}$, restores the saved state of each FMU, and tries again with new step size $h_{\min}$.

Assuming that the FMUs satisfy the reasonable "monotonicity" property that if they were able to advance time by h'_i then they are also able to advance time by any smaller step, and by the fact that $h_{\min}$ is smaller than all h'_i, the second attempt is guaranteed to succeed. That is, $h_{\min}$ will be accepted by all FMUs. As a result, at most after two attempts, a co-simulation step is successful, and the algorithm proceeds with the next step, repeating the same procedure as above.

The FMI standard sets out a framework where FMUs share the notion of time and exchange variable values via input-output ports: outputs from one FMU are mapped as inputs to other FMU(s) and so on. The output port values are said to be owned and controlled by the emitting FMU, whereas the inputs are computed and provided by another (outputting) FMU. The framework foresees that before producing an output an FMU may first need some input values and thus input-output dependency information is introduced. Overall the I/O port connectivity graph derived from the model of interconnected FMUs, together with the local I/O dependencies of each individual FMU, result in a global I/O dependency graph for the entire model (Broman et al., 2013).

Time and I/O values are synchronized by the co-simulation algorithm: the time is agreed by repeatedly consulting each FMU and the I/O values are propagated according to dependencies. The co-simulation algorithm assumes that each FMU provides a static dependency list of its ports before simulation starts, and that the resulting global I/O dependency graph is acyclic, and therefore there exists a schedule for computing the value of every input port before the value of a dependent output port is requested (Broman et al., 2013).

3 Translating Models into FMUs

The behavior of individual FMUs is provided by the model-checker's simulation engines based on the guidelines described by Tripakis (2015). In particular, the report distinguishes continuous and discrete dynamics. The continuous behavior is modeled by differential equations over continuous variables whose values can be

shared among FMUs by the means of port connections. The output ports of an FMU are mapped to the owned/controlled variables which are read and written to, whereas input ports map to read-only variables within the FMU.

The discrete behavior is modeled by discrete transitions in the timed/hybrid automata control flow structure. The discrete transitions are designed to be executed with micro-steps of zero delay. Transitions can also be decorated with event labels and each tool supports its own kind(s) of synchronizing compositions internally and therefore the discrete transition synchronization is also handled individually within the tools. Tripakis (2015) provides the means of discrete transition synchronization by allocating two special port variables: one for incoming (input) synchronization and one for outgoing (output) synchronization. The domain of discrete input (output) ports coincides with the set of input (output) labels plus a special value *absent* which denotes no synchronization or an internal discrete transition.

3.1 Uppaal

UPPAAL uses timed automata models (Alur and Dill, 1994), extended with discrete variables over structured types to describe behaviors of a timed system. In timed automata, the continuous dynamics is controlled by real-valued clock variables (with derivatives set to one) and discrete states complemented with integer variables – both of which are candidates for exchange via FMU input-output ports. Statistical model checking (SMC) extensions (David et al., 2011, 2015) allow a finer control of the clock derivatives by means of ordinary differential equations, moreover the discrete transitions are stochastic where the execution is determinized by probability distributions over time and over branching edges. The stochastic semantics of a parallel composition is similar to the FMI co-simulation algorithm (Broman et al., 2013): the way the minimum delay is negotiated and thus the timed composition within the FMI framework is straightforward, and task is to find a systematic way of handling discrete synchronizations. UPPAAL also supports the maximal progress or ASAP semantics on edges labeled with urgent channels.

UPPAAL supports the notion of discrete I/O synchronization natively by means of input and output channel labels. Thus, its discrete input and output transitions can be mapped directly to the input/output port variables of an FMU that is dedicated to transfer the synchronization label name. Nonetheless, we distinguish the following kinds of transitions: internal (transitions without I/O channel synchronization or internally synchronized transitions for which channels are not marked as input or output), input transitions (labeled by an input synchronization where the channel name is marked as an FMU input), and output transitions (labeled by an output synchronization where channel is marked as an FMU output). The marked outputs are controlled by the UPPAAL simulation and are executed asynchronously irrespective of whether the receiving FMU is ready to synchronize. Meanwhile, the input transitions are executed only when there is a corresponding input label set on a discrete input port. At most one (internal, input or output) transition is allowed at a time, hence fine-grained simulation control can be achieved by the co-simulation algorithm.

UPPAAL FMUs do not introduce I/O dependencies between continuous variables because the models do not use algebraic expressions to compute variable values. Instead of algebraic expressions the automata use discrete transitions to update the variable values. However, only one discrete transition is allowed at a time, therefore all discrete outputs have dependencies on the inputs dedicated to synchronization labels which restrict the selection of a particular discrete transition and hence specific variable update.

3.2 SpaceEx

SPACEEX (Frehse et al., 2011) uses hybrid automata to describe system behavior where the continuous variable derivatives are constrained by differential equations. The continuous variables are candidates for input and output exchange via FMU ports. The discrete transitions of hybrid automata can be decorated with labels. Synchronization may involve multiple participating processes, but there is no notion of input and output – all processes are equal contributors, therefore the simulator needs to implement the input/output semantics required by FMI. We use a special label naming notation to mark input and output labels (see Fig. 6). The transitions with input labels are only executed when the discrete input variable of FMU is set to the corresponding label name. Meanwhile, the transitions with an output label are controlled by SPACEEX' simulation, and are executed asynchronously by setting the discrete output variable with the label name irrespectively of whether the receiving FMU can synchronize with it. We ensure the SPACEEX FMU determinism by enforcing the must-semantics of discrete transitions in a hybrid automaton. In other words, a discrete transition is taken as soon as its guard is enabled. Finally, we resolve the non-determinism between input, output, and internal transitions in the following way: input transitions have priority over output transitions and output transitions are preferred over the internal ones.

Both UPPAAL and SPACEEX translations simulate the source models as they are without intermediate transformations, except of the following additions: 1) input enabledness is ensured by broadcast channels in UPPAAL modeling and asynchronous I/O is implemented for SPACEEX synchronization labels, 2) for determinization SPACEEX uses maximal progress whereas UPPAAL uses stochastic semantics with a possibility of urgent channels for maximal progress.

Figure 1. An example of four timed automata chain.

3.3 Discussion on Co-Simulation Semantics

In this section, we discuss the co-simulation semantics and contrast it to those typically used by a model-checking tool. In particular, we demonstrate by example how the FMI co-simulation algorithm resolves input/output dependencies and contrast it with execution analysed in a model checker. Our goal is to offer insights in the differences of the two semantics.

Consider a system model shown in Fig. 1 which consists of four timed automata composed in parallel. Labels of the form $a!$ denote sending output a, whereas $a?$ denotes receiving an input a. The variable x is a *clock* measuring time starting from zero. The constraint $x = 1$ is a *guard* which allows the corresponding transition of the automaton to occur only if the guard is satisfied, i.e., in this case only when x equals 1. The automata synchronize in a chain: the first can output a to the second one, the second one can output b to the third one and so on.

In principle, the system can be loaded into an FMI model in any combination: individually (one automaton per FMU) or collectively (multiple automata per FMU), but before an FMU can be loaded into an FMI model, it must declare its input/output dependencies. According to Broman et al. (2013) each automaton should expose an input/output variable which will contain the synchronization label value. Automaton A_1 in the example above will have only an output variable, which may have values $\{a, absent\}$. Automaton A_2 will have an input variable ranging over $\{a, absent\}$ and an output variable ranging over $\{b, absent\}$, and so on. The special value *absent* denotes that currently there is no synchronization. Timed automata must declare a dependency between its input and output label variable in order to avoid simultaneous input and output synchronizations.

In addition, it is assumed that each FMU is *input-enabled*, meaning that it can handle (i.e., it is able to receive) any declared input at any time. If a component is not input-enabled and an input synchronization is triggered then simulation is aborted, to avoid such situation we allow only broadcast channels, which do not block the sender process and receiver may simply ignore the synchronization if has no receiving edge.

Suppose the automata from Fig. 1 are loaded within separate FMUs and connected according to synchronization labels. That is, the output of $FMU(A_1)$ is connected to the input of $FMU(A_2)$, the output of $FMU(A_2)$ is connected to the input of $FMU(A_3)$, and so on. The co-simulation algorithm would detect that it has to fulfill inputs values for the $FMU(A_4)$, $FMU(A_3)$, and $FMU(A_2)$

in order to proceed, therefore the input/output value propagation will have to start with $FMU(A_1)$ and then proceed to the $FMU(A_2)$ etc.. Once the values of all input and output variables are propagated, the algorithm proceeds with advancing each FMU in time by calling $doStep()$. It is this dynamic behavior in time which interests us in this example.

In particular, observe that $A_{2,3,4}$ automata are non-deterministic in the sense that, according to UPPAAL semantics, at time $x = 1$ an automaton can either delay, or take an outputting transition, or synchronize on inputs. For instance, at time $x = 1$, A_2 can either emit b, or receive a (which will be available in this case, because it is sent by A_1 at exactly that time), or let the time pass. In timed automata semantics, all these options are possible at the individual component level. Moreover, not only individual components can be non-deterministic, but their composition is non-deterministic as well, based on so-called *interleaving semantics*. This means that when multiple automata are enabled at a given time, the choice of which one to execute is arbitrary. Non-determinism is a useful abstraction and thus model reduction technique in verification and model-checking. The same is true when these tools are used for simulation, i.e. different simulations in UPPAAL may yield different results.

In FMI, the situation is very different, as all FMUs are treated as deterministic components, and their composition, ensured by the co-simulation algorithm, is guaranteed to yield deterministic results as well. Interestingly, in this example, if all automata decide to output at time $x = 1$, some of them will succeed outputting in parallel, while others will be preempted by incoming inputs. In particular, the master algorithm will request $FMU(A_1)$ to produce its output, and thus $FMU(A_2)$ will be busy handling an input and will not be producing output at that time. Since $FMU(A_2)$ is not sending anything, then $FMU(A_3)$ will be free to produce an output and hence preempt $FMU(A_4)$.

As witnessed from above, such FMI system selects a particular sequence of steps (which is expected) but is not able to simulate all possible execution orders as in original semantics even if we allow FMUs to determinize their actions by themselves, which means that FMI simulations are selecting a particular subset of all possible behaviors and some behaviors may not be reproducible in FMI. Also FMI simulations may contain parallel synchronizations (e.g. actions $A_1 \overset{a}{\rightarrowtail} A_2$ and $A_3 \overset{c}{\rightarrowtail} A_4$ at the same computation step) which are possible only in several steps in timed automata semantics (action a and only then action c within zero-time), hence the intermediate state between a and c actions might not be accessible in FMI without very fine grained control over individual $doStep()$ calls in one zero-time computation step. However, the successor state of such parallel executions can be matched with a state after multiple transitions in the given automata semantics, hence the FMI simulation

states in between system computation steps are included in the original semantics, albeit definite proof requires more formal insight to examine all scenarios.

4 Case Study

We have implemented the FMI standard in the UPPAAL (Larsen et al., 1997) and SPACEEX (Frehse et al., 2011) model checkers by providing model export to FMU[3]. In this section, we present and evaluate the performance of the resulting FMI framework on a case study inspired by the well-known room heating benchmark originally proposed by Fehnker and Ivancic (2004). Our model consists of a room with a heater (Fig. 2) and a controller (Fig. 3) which regulates the heater behavior. We model the room and the controller as a SPACEEX and UPPAAL FMU, respectively (see Fig. 4). Our bang-bang controller turns the heater on and off as soon as some temperature thresholds T_{low} and T_{high} have been reached. The as-soon-as-possible behavior is enforced by using urgent channels which effectively make the controller deterministic. The room temperature T evolves according to the following differential equation:

$$\dot{T} = k \cdot (T_{env} - t) + h_{power}$$
$$\dot{T}_{env} = 0$$
$$\dot{h}_{power} = 0$$

In other words, the room temperature depends linearly on the difference between the current room temperature T and outside temperature T_{env}. We assume the outside temperature T_{env} and heater power h_{power} to be constant. The constant k defines the heat exchange rate between the room and outside environment. If the heater is off, the heater power is set to zero.

4.1 Evaluation

We evaluate our FMU framework by comparing simulation trajectories of the FMUs with the ones produced by a SPACEEX model consisting of both the controller and room components. We consider three different simulation step values: 1 (see Fig. 5a), 0.1 (see Fig. 5b) and 0.01 (see Fig. 5c). Considering the simulations, we observe that the FMU trajectories *overshoot* the controller constraints in the sense that the controller exhibits a delayed reaction when the room temperature crosses the temperature thresholds. The behavior is justified by the fact that the method call `doStep` for every FMU relies only on the *local* information about the state evolution when making decisions, e.g., the controller FMU does not have any information about the room temperature evolution beyond the value which can be provided when

---

[3]A package containing the benchmarks is available for download at `http://swt.informatik.uni-freiburg.de/tool/spaceex/co-simulation`.

Figure 2. Room component modelled in SPACEEX. The component switches between "on" and "off" modes. The temperature variable T is exported as output and synchronizations labels h_{on} and h_{off} as inputs.


Figure 3. Controller in UPPAAL uses urgent channels to ensure as-soon-as-possible transition trigger. Temperature T is an input and labels h_{on} and h_{off} are outputs.

the method `doStep` is called. Therefore, the controller FMU detects that the guard is enabled only a *simulation iteration later* after this event has already happened. We observe that the impact of the overshooting can be made arbitrary small by choosing a small enough simulation step (see Fig. 5c vs. Fig. 5a and Fig. 5b).

We note that the overshooting problem is *inherent* to the considered master algorithm and can be circumvented by incorporating additional cross-component knowledge into the master algorithm. Overall, our experiments validate that on this case study our co-simulation framework based on SPACEEX and UPPAAL provides equivalent simulation results compared to the setting where all components are modelled in one tool.

4.2 Supervisory Control Example

In this section, we show how supervisory control systems similar to the benchmarks presented by Fehnker and Ivancic (2004) can be modeled using the FMI paradigm. Compared to Section 4.1, we consider a model of the building with two rooms sharing a common wall and a heater. In this setting, the room temperature is influ-

**Figure 4.** SPACEEX and UPPAAL FMUs connected using the room temperature T and heater mode h_{mode}.

(a) Maximum step size 1.

(b) Maximum step size 0.1.

(c) Maximum step size 0.01.
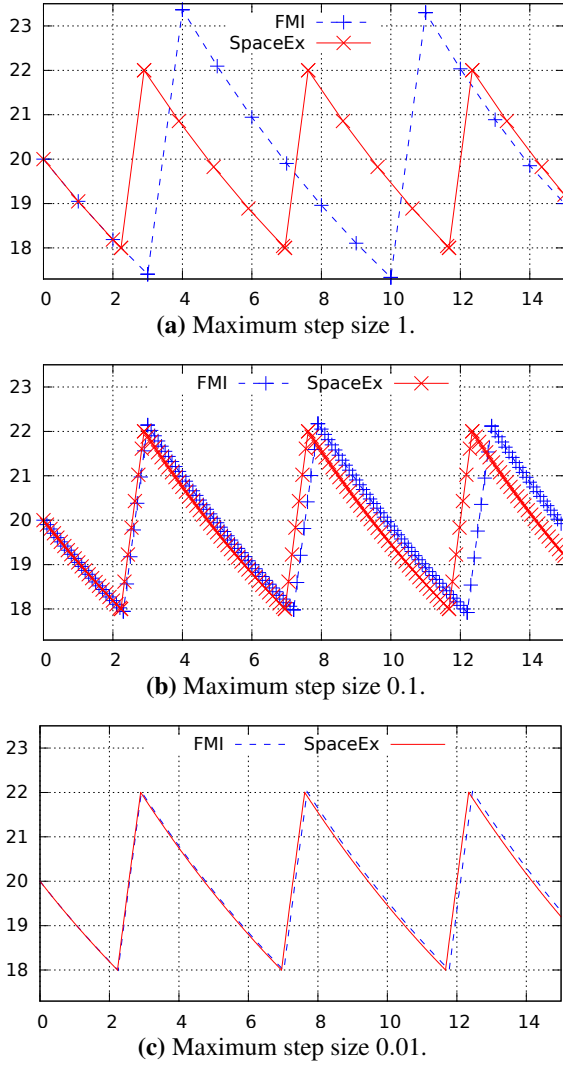
Figure 5. Simulation trajectories: each red × is a data point reported by SPACEEX, and blue + reported by the co-simulation.

enced by both the outside temperature and heat transfer between the rooms. Figure 6 shows a hybrid automaton from SPACEEX modeling the room temperature dynamics. The difference from the previous example here is an extra term $(Tother - t) * 0.2$ denoting a contribution from another room. Another room is modeled analogously except that it responds to *heater2_on* and *heater2_off* signals instead of *heater1_on* and *heater1_off*.

Our controller consists of two parts: local bang-bang controller and a supervisor shown in Fig. 7. In order to model the transitions of the heaters between the rooms, we assume that the controllers can be turned on/off by the supervising controller. Therefore, the local controller has an extra mode besides *On* and *Off* which stands for the controller being currently deactivated. The supervising controller has two kinds of stochastic behavior: it can pick any pair of rooms (one recipient and another donor) to transfer the heater, and it can choose the timing of transfer. When a pair of rooms is selected (by choosing concrete room identifiers for *rec* and *donor* variables) the donor is disabled by moving from location *decide* to lo-

cation *move* and the recipient is enabled by going from *move* to *idle*. The supervisor may stay in location *idle* arbitrary long, but the exact duration is decided by an exponential probability distribution of rate 1 which means the duration of 1/1 time units on average. Similarly the supervisor may stay in *decide* and *move* but the duration will be 1/10000 on average, i.e. denoting that the heater is moved rather quickly.

Figure 8 shows the overall component connectivity diagram where the supervisor is reading temperatures from each room and controls the local movable heater controllers. The movable heaters then may either turn on the heat in their room or let them cool off giving the heat to
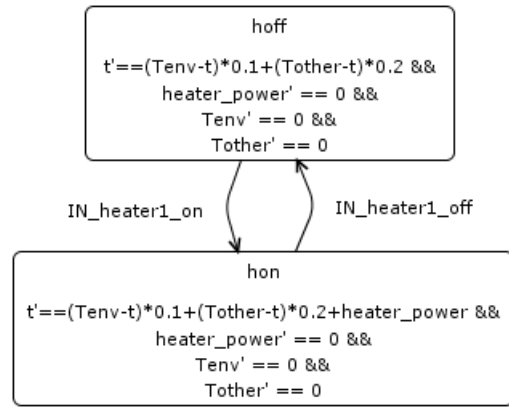
Figure 6. Hybrid automaton for a heated room connected to another room. Inputs are temperatures *Tenv*, *Tother* and labels *IN_heater1_on* and *IN_heater1_off*, while output is temperature *t*. We use the prefix *IN* to mark input labels.

(a) Local bang-bang controller which can be moved (disabled). The inscribed U means urgent location where time delay is not allowed. The inputs are temperature variable *T[id]* and labels *enable[id]* and *disable[id]*, while outputs are labels *heater1_on* and *heater1_off*.
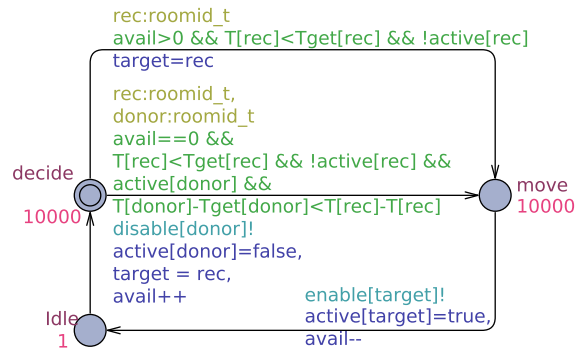
(b) Supervising controller moves the heaters between rooms by reading inputs on *T[i]* and sending outputs on labels *enable[i]* and *disable[i]* where *i* is the room index.

Figure 7. Two layers of UPPAAL controllers.

outside. The individual heated rooms are then connected to the outside temperature and to each other denoting the heat exchange. The splitter FMUs are repeaters needed to connect multiple components to the same signal.

In the following, we discuss the behavior of the resulting composed model. Figure 9 shows the temperature dynamics in each room. In particular, the plot shows that in the beginning the temperature drops until the supervisor detects a room temperate below $Tget = 17°$, then around 6 time units a heater raises the temperature in room 1. The local controller keeps rising the temperature until it goes over $22°$ bound at around 7.5 time units. Notice that the temperature in room 2 also rises due to heat exchange between the rooms. Around 10 time units the supervisor decides to hand over the heater to room 2. At 14 time units the heater is switched back to room 1 and so on. We can conclude that even though the temperature drops well below $18°$ overall it seems that the controllers manage to sustain the temperature at the similar level without loosing control (without dropping to outside temperature level).

4.3 Stochastic Simulations and SMC

The following is a demonstration of statistical model checking (SMC) using the FMI framework. We show how the performance of two stochastic controllers simulated by UPPAAL can be compared using SMC approach together with the heated room simulation provided by SPACEEX. Figure 10 shows two controllers: (a) reacting within 1 time unit to $18.0°$ and $22.0°$ temperature bounds and (b) reacting within 2 time units to $19.0°$ and $21.0°$ temperature bounds. The channels used in these controllers are not urgent and therefore the delay between temperature detection and heater activation is decided stochastically based on uniform distribution over the allowed delay by invariants, i.e. the concrete delay will be chosen from $[0, 1]$ for the first controller and from $[0, 2]$ for the second one. The *On* and *Off* locations do not have any invariant and therefore in principle the process may stay there forever. In such cases UPPAAL uses an exponential (Poisson) probability distribution to decide a particular time delay and hence asks to provide a rate of the exponential. The higher the exponential rate, the shorter the delays, hence we can provide a high rate to ensure that the detecting transition is fired arbitrary quickly.

In our setup, we would like to know which controller is better at keeping the room temperature within $18.0°$ and $22.0°$ bounds. In order to answer this question we setup two FMI models for each controller with an equal room, run 100 simulations with 100 time units in length and 0.05 granularity, compute the amount of time spent outside the temperature range for each simulation and then compute the confidence intervals for both models. Table 1 shows a summary of amounts of time during which the temperature was either below or above the range. The estimated time duration use confidence interval (CI) notation which means that if we repeat the measurement experiment then the real mean (which is unknown) will fall into the interval with a probability of 95%. The results show that the second controller was more successful at maintaining the lower bound of the temperature, but was more overshooting beyond the upper bound. In total, the first controller kept the temperature in good range longer by 8.57 time units on average, which is much larger than confidence interval, hence the first controller is better.

Table 1. Time with temperature outside the range (95% CI).

Controller	Time below	Time above	Total
Wide and fast	7.56 ± 0.20	32.69 ± 3.36	40.26 ± 0.59
Narrow and slow	2.40 ± 0.19	46.43 ± 0.82	48.83 ± 0.79

5 Related Work

The FMI standard and corresponding documentation are constantly evolving, as new versions of the standard are developed. The web site[4] also contains a list of tools supporting FMI. Descriptions of FMI can also be found in the academic literature (Blochwitz et al., 2011).

Discussions about the limitations of FMI can be found in the works by Broman et al. (2013, 2015). Broman et al. (2013) also formalize the main methods of FMI (`get`, `set`, `doStep`) by establishing a *contract* (pre-/post-conditions) for each method and propose a *master algorithm* (i.e., a co-simulation algorithm). Furthermore, the authors proves its termination, determinacy, and other properties. However, the paper does not discuss how FMUs can be created. A different, master-slave based, co-simulation approach is proposed by Bastian et al. (2011), but formal properties such as determinacy are not discussed in this work.

Broman et al. (2015) defines a suite of test models that should be supported by a hybrid co-simulation environment, giving a mathematical model of an ideal behavior, plus a discussion of practical implementation considerations. Furthermore, the paper describes a set of basic modeling components in the spirit of Ptolemy actors (constant, gain, adder, integrator, etc.). Finally, the authors provide a kind of denotational description for each component (input and output signals), but no encoding into FMUs is discussed.

The FMU generation problem for various formalisms is discussed by Tripakis (2015). This work only refers to a generic model of timed machines which does not include the particularities of UPPAAL's timed automata. In addition, hybrid automata are not considered in this work.

Recently, the co-simulation algorithm presented by Broman et al. (2013) has been implemented in the

---

[4] `https://www.fmi-standard.org/`
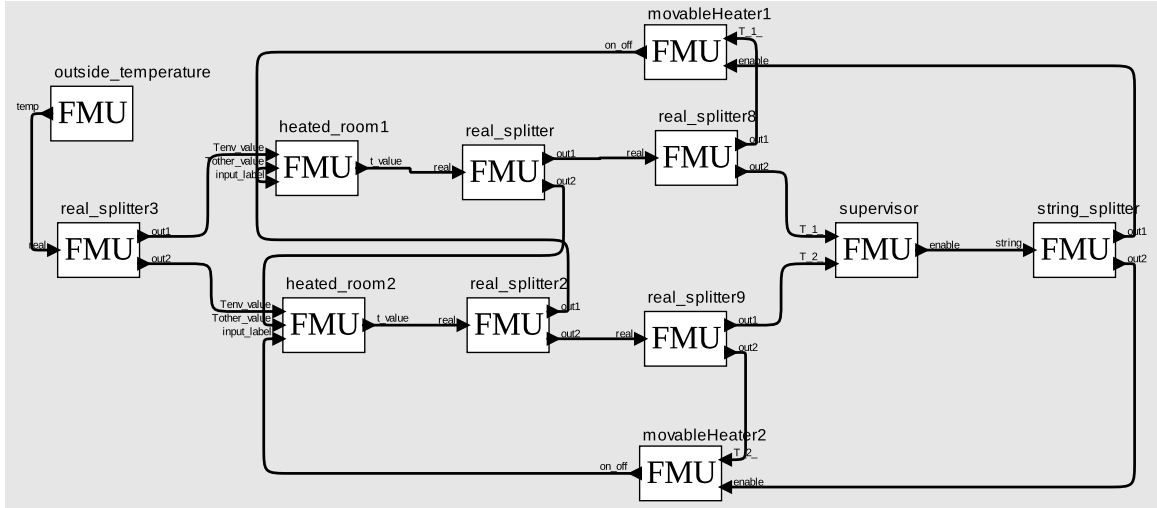
Figure 8. Ptolemy diagram for supervisory control of two heated rooms.

open-source Ptolemy tool. As mentioned above, we use this framework in order to import FMUs into Ptolemy and co-simulate them. However, this framework does not address the FMU generation problem.

Feldman et al. (2014) present a plugin for Rhapsody for generating FMUs from Statechart SysML blocks. They provide high level guidelines for how to generate Statechart FMUs, but do not provide a formalization. Pohlmann et al. (2012) also discuss how to encode statechart models, described in MechatronicUML.

Co-simulation is one, but not the only approach to solve the tool interoperability problem. A further attempt to solve this problem is the *Hybrid Systems Interchange Format* (HSIF), designed with the goal of being "a sort of 'maximum common denominator' among all hybrid system modeling environments" (Pinto et al., 2005). HSIF aims therefore at defining a "maximal syntax" where all the syntax of different languages could be translated into. It could be seen as a type of XML schema for hybrid systems. HSIF is primarily aimed at enabling model *translation* between different hybrid system tools. Bak et al. (2015) present the tool Hyst which provides an automatic source-to-source model translation between a number of up-to-date hybrid model checkers. In their approach, Bak et al. (2015) do not use any intermediate format like HSIF. Both model translation-based approaches

outlined above provide support only of the *common subset* of the tool features. Our co-simulation framework does not limit the model designer to use the "maximal syntax" among all the tools because every FMU takes care of its features independently.

6 Conclusions

We have shown how two state-of-the-art modeling and verification tools for hybrid systems, SPACEEX and UPPAAL, can be integrated using the FMI co-simulation standard. The result is a powerful framework which allows users to build submodels separately in each of the two tools (as well as in other tools potentially), then generate individual FMUs for each submodel, and then combine all the FMUs into a single model, which can be co-simulated within the Ptolemy FMI implementation. We demonstrated the feasibility of our framework by comparing the co-simulation results of a simple two-FMU model to the simulation results that are obtained when we model and simulate the entire system in a single tool. By empirical evaluation on case studies we found that, provided time steps are small enough, individual components can ensure timely reactions to continuous signals, and the simulations can be made arbitrary close to self-contained model simulation. In addition to individual tool export to FMUs, we showed

Figure 9. Temperature trajectories for each of the rooms composed with stochastic supervising controller.

(a) Wide range and fast.   (b) Narrow range and slow.

Figure 10. Stochastic controllers use regular (non-urgent) channels, therefore timings are stochastic: delays are distributed uniformly (when clock invariant is used) and exponentially (in locations *On* and *Off*).

how the non-deterministic models can be determinized using stochastic semantics and included into FMI co-simulation. We also provided an example how statistical model checking can be performed using numerous FMI simulations which is an essential feature evaluating stochastic behavior. The integration of model-checkers into co-simulation frameworks provides further possibilities of analyzing early design models like conformance monitoring by checking that a simulation trace of a refined (e.g. hybrid) model is included in a more a abstract (e.g. timed automata) specification. We envision our work being a further step towards integrating tools developed in the formal methods community into the industrial system design and modeling workflow of cyber-physical systems.

7 Acknowledgments

References

D.E. Nadales Agut, Dirk A. van Beek, and J.E. Rooda. Syntax and semantics of the compositional interchange format for hybrid systems. *The Journal of Logic and Algebraic Programming*, 82(1):1 – 52, 2013. ISSN 1567-8326. doi:10.1016/j.jlap.2012.07.001.

R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138(1):3–34, 1995.

Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.

Stanley Bak, Sergiy Bogomolov, and Taylor T. Johnson. HYST: a source transformation and translation tool for hybrid automaton models. In *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, HSCC, Seattle, WA, USA, April 14-16, 2015*, pages 128–133. ACM, 2015.

Jens Bastian, Christoph Clauß, Susann Wolf, and Peter Schneider. Master for Co-Simulation Using FMI. In *8th International Modelica Conference*, 2011.

Harsh Beohar, D. E. Nadales Agut, Dirk A. van Beek, and Pieter J. L. Cuijpers. Hierarchical states in the compositional interchange format. In Luca Aceto and Pawel Sobocinski, editors, *Proceedings Seventh Workshop on Structural Operational Semantics, SOS 2010, Paris, France, 30 August 2010.*, volume 32 of *EPTCS*, pages 42–56, 2010. doi:10.4204/EPTCS.32.4.

T. Blochwitz, M. Otter, M. Arnold, C. Bausch, C. Clauß, H. Elmqvist, A. Junghanns, J. Mauss, M. Monteiro, T. Neidhold, D. Neumerkel, H. Olsson, J.-V. Peetz, and S. Wolf. The Functional Mockup Interface for Tool independent Exchange of Simulation Models. In *8th International Modelica Conference*, Dresden, Germany, March 2011. Modelica Association.

D. Broman, C. Brooks, L. Greenberg, E. A. Lee, S. Tripakis, M. Wetter, and M. Masin. Determinate Composition of FMUs for Co-Simulation. In *13th ACM & IEEE International Conference on Embedded Software (EMSOFT'13)*, 2013.

D. Broman, L. Greenberg, E. A. Lee, M. Masin, S. Tripakis, and M. Wetter. Requirements for Hybrid Cosimulation Standards. In *Hybrid Systems: Computation and Control (HSCC)*, 2015.

Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikučionis, Danny Bøgsted Poulsen, Jonas van Vliet, and Zheng Wang. Statistical model checking for networks of priced timed automata. In Uli Fahrenberg and Stavros Tripakis, editors, *Formal Modeling and Analysis of Timed Systems*, volume 6919 of *Lecture Notes in Computer Science*, pages 80–96. Springer Berlin Heidelberg, 2011.

Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikučionis, and Danny Bøgsted Poulsen. Uppaal SMC tutorial. *International Journal on Software Tools for Technology Transfer*, 2015.

J. Eker, J. Janneck, E. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuendorffer, S. Sachs, and Y. Xiong. Taming heterogeneity – the Ptolemy approach. *Proceedings of the IEEE*, 91(1):127–144, January 2003.

Ansgar Fehnker and Franjo Ivancic. Benchmarks for hybrid systems verification. In *In Hybrid Systems: Computation and Control (HSCC)*, pages 326–341. Springer, 2004.

Y. A. Feldman, L. Greenberg, and E. Palachi. Simulating Rhapsody SysML Blocks in Hybrid Models with FMI. In *10th Modelica Conference*, pages 43–52, 2014.

Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. SpaceEx: Scalable Verification of Hybrid Systems. In Shaz Qadeer Ganesh Gopalakrishnan, editor, *23rd International Conference on Computer Aided Verification (CAV)*, LNCS. Springer, 2011.

Kim G. Larsen, Paul Pettersson, and Wang Yi. Uppaal in a nutshell. *International Journal on Software Tools for Technology Transfer*, 1(1-2):134–152, 1997.

Alessandro Pinto, Alberto L. Sangiovanni-Vincentelli, Luca P. Carloni, and Roberto Passerone. Interchange formats for hybrid systems: review and proposal. In *Hybrid Systems: Computation and Control*, HSCC. Springer, 2005.

Alessandro Pinto, Luca P. Carloni, Roberto Passerone, and Alberto Sangiovanni-Vincentelli. Interchange format for hybrid systems: Abstract semantics. In Joao P. Hespanha and Ashish Tiwari, editors, *Hybrid Systems: Computation and Control*, volume 3927 of *LNCS*, pages 491–506. Springer Berlin Heidelberg, 2006.

Uwe Pohlmann, Wilhelm Schäfer, Hendrik Reddehase, Jens Röckemann, and Robert Wagner. Generating Functional Mockup Units from Software Specifications. In *9th Modelica Conference*, pages 765–774, 2012.

Stavros Tripakis. Bridging the Semantic Gap Between Heterogeneous Modeling Formalisms and FMI. In *International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation – SAMOS XV*, 2015.