

Representation of Synchronous, Asynchronous, and Polychronous Components by Clocked Guarded Actions

Submitted to Special Issue on
Languages, Models and Model Based Design for Embedded Systems

Jens Brandt · Mike Gemünde · Klaus Schneider ·
Sandeep K. Shukla · Jean-Pierre Talpin

Received: date / Accepted: date

Abstract For the design of embedded systems, many languages are in use, which are based on different models of computation such as event-, data-, and clock-driven paradigms as well as paradigms without a clear notion of time. Systems composed of such heterogeneous components are hard to analyze so that mainly co-simulation by coupling different simulators has been considered so-far. In this article, we propose clocked guarded actions as a unique intermediate representation that can be used as a common basis for simulation, analysis, and synthesis. We show how synchronous, (untimed) asynchronous, and polychronous languages can be translated to clocked guarded actions to demonstrate that our intermediate representation is powerful enough to capture rather different models of computation. Having a unique and composable intermediate representation of these components at hand allows one a simple composition of these components. Moreover, we show how clocked guarded actions can be used for verification by symbolic model checking and simulation by SystemC.

Keywords Models of Computation · Co-simulation · Synchronous vs. Asynchronous Models · Guarded Command Language

1 Introduction

For the design of embedded systems, a plethora of languages based on different models of computation (MoC) [3, 29, 46, 53, 54] have been proposed over the years. For example, languages like Verilog [42], VHDL [44], and SystemC [43] are based on an event-driven paradigm [18], synchronous languages [4, 32] such as Esterel [7, 9, 11], Lustre [17, 34], Quartz [64], and some statechart variants are based on clock-driven paradigms, polychronous languages like Signal [2, 28, 49] are based on a declarative and non-deterministic paradigm

Jens Brandt, Mike Gemünde, Klaus Schneider
Department of Computer Science, University of Kaiserslautern, <http://es.cs.uni-kl.de>

Sandeep K. Shukla
Electrical and Computer Engineering, Virginia Tech, <http://www.fermat.ece.vt.edu>

Jean-Pierre Talpin
INRIA, Unité de Recherche Rennes-Bretagne-Atlantique, <http://www.irisa.fr/espresso>

using several clocks, data flow languages like CAL [25] are based on data-driven paradigms [50–52], and others like SHIM [24] or most multi-threaded languages are based on asynchronous threads with a rendezvous-style communication [38, 39].

Depending on a particular application domain such as digital signal processing or reactive controllers, depending on the design task such as modeling, simulation, analysis or synthesis, and depending on the synthesis target such as digital hardware circuits, multi-threaded software or software for heterogeneous MPSoCs, the one or the other language might be preferable. For this reason, many existing components are given in different languages using different MoCs. The co-simulation of such heterogeneous systems has been widely considered [10, 55, 59, 67, 69] and covers also languages with different MoCs [20]. Co-simulation is also used to create virtual prototypes that are required to achieve hard time-to-market constraints.

However, co-simulation alone is not sufficient for a seamless design flow. Formal verification and a common synthesis of the different components require an *integration* that has gained a lot of interest in recent years [26, 31, 36, 37, 48, 56, 61, 62, 68]. Moreover, having a common description for the different components at hand, one can easily combine the components, e.g. one can create new components by using existing ones in a hierarchical way as known in block/schematic-oriented languages such as Simulink or those used by many tools for digital hardware circuit design. This way, one is no longer restricted in a parallel composition of the heterogeneous components by using appropriate wrappers. Instead, one can create a hierarchy of modules that combine modules based on different MoCs. This allows one to establish a design flow using several steps of refinement where asynchronous descriptions can become synchronous by adding a schedule to clocks, and synchronous ones may become asynchronous again when one considers the actions scheduled to one clock tick. Finally, one can re-use existing backend tools for synthesis and verification and of course, simulation would be greatly simplified: Instead of coupling different simulators, a single one could consider the entire system in a way it will later on be verified and synthesized.

A classical solution used e.g. in most compilers is to use a common intermediate representation, which bridges the gap between powerful programming languages with complex semantics and the low-level description of the target code. This intermediate representation must be based on a common model of computation which covers all MoCs that should be integrated without complex translations. Such an intermediate representation has many advantages. It not only achieves the above mentioned integration, but it also allows designers to share the tool infrastructure: new input languages can be added by simply implementing a corresponding front-end while existing back-end tools can be re-used.

In this article, we propose *clocked guarded actions* as an intermediate representation to cover various MoCs used for the design of embedded systems. This representation is in the spirit of guarded commands, a well-established concept for the description of concurrent systems. With a theoretical background in conditional term rewriting systems [21, 30, 47], guarded actions have been not only used in many specification and verification formalisms (e. g. Dijkstra’s guarded commands [22], Unity [19], Murphi [23]) but they have also shown their power in hardware and software synthesis (e. g. Bluespec [1, 40] and Concurrent Action-Oriented Specifications [41]).

To demonstrate the power of our approach, we sketch a design flow based on our intermediate representation (see Figure 1): In particular, we show how different languages based on different MoCs can be translated to clocked guarded actions, and systems given as a set of clocked guarded actions can be used for formal verification by symbolic model checking using SMV and simulation using SystemC simulators.

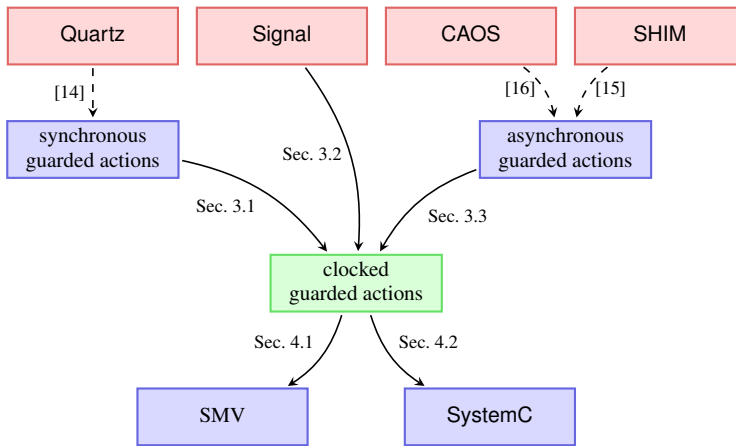


Fig. 1 Design Flow for Clocked Guarded Actions

This article is an extension of a previous paper [13], which introduced clocked guarded actions as a common intermediate representation. In this article, we extend our previous approach by introducing control-flow contexts, which allows us to nest components in an arbitrary behavioral hierarchy, i. e. we can *call* components from other components. Obviously, this has not only consequences for the intermediate representation itself but also related tasks: the translation to the intermediate format must be revised, and the procedure to link several components must be generalized.

The rest of this article is organized as follows: Section 2 defines the core of our approach, i. e., clocked guarded actions, components represented by a set of clocked guarded actions and an interface, and their meaning in terms of a denotational semantics. Section 3 shows how synchronous languages like Quartz [64], polychronous languages like Signal [2, 28, 49], and asynchronous languages like CAOS [66] can be translated to clocked guarded actions. Section 4 shows how clocked guarded actions can be used for verification and simulation. Section 5 illustrates our intermediate representation with the help of a case study. Finally, Section 6 considers some related work on the integration of components with different MoCs before we conclude the article with a short summary in Section 7.

2 Intermediate Representation

This section presents our intermediate representation of components by clocked guarded actions. To this end, we define an extension of our Averest Intermediate Format AIF that is based on guarded actions without clocks and is used in our Averest framework¹ as intermediate representation for simulation, verification, analysis, transformations, and synthesis. Clocked guarded actions lead to AIF⁺ which generalizes and extends the single-clocked synchronous intermediate format AIF [14] by clocks to cover polychronous and asynchronous systems.

In the following, we only show fundamental design decisions of AIF⁺. We start with some fundamental definitions in Section 2.1 before Section 2.2 describes the basic artifacts in our intermediate representation, namely components and their interfaces. Clocked

¹ <http://www.averest.org>

guarded actions, which are the basis for the description of the behavior, are introduced in Section 2.3.

2.1 Foundations

Guarded commands or guarded actions are a well-established concept for the description of concurrent systems. As already stated in the introduction, they have already been used for many purposes. Usually, guarded actions are seen as an asynchronous model as follows: In the current state, the guards of all actions are evaluated and then an arbitrary subset of the enabled actions is selected for execution. The action to be executed may consist of several statements which are then executed in parallel. This kind of semantics defines an asynchronous system since the enabled actions may be executed, but have no need to be executed. As a result the execution is by definition in general non-deterministic.

Guarded actions have also been successfully used for synchronous languages [14, 33, 63]. In contrast to the traditional guarded actions as described above, synchronous guarded actions required that *all enabled actions are executed*. Thus, they define a deterministic model of computation.

Causality problems may occur in synchronous and asynchronous guarded actions if the actions are allowed to take immediate effects, i.e. if the executed action modify variables that are used in their guards. Causality analysis checks whether such cyclic dependencies can be constructively resolved which means in practice that the execution follows the data dependencies between the actions.

Clocked guarded actions (CGA), which we propose in this article, provide a basis to integrate both variants. As the name suggests, they are defined over a set of explicitly declared *clocks* \mathcal{C} , which define logical timescales the system uses to synchronize its computations so that asynchrony and synchrony in the system can be precisely described. The basis of the whole temporal model are so-called instants, i. e. the points of time where some event in the system occurs.

Definition 1 (Instants) Let \mathcal{I} be the set of *instants* and let $\preceq \subseteq \mathcal{I} \times \mathcal{I}$ be a preorder on \mathcal{I} such that for any two instants $I_1, I_2 \in \mathcal{I}$, we say $I_1 \preceq I_2$ iff I_1 occurs before I_2 , or if both of them occur together. Let \approx be the equivalence relation induced by \preceq : thus $I_1 \approx I_2$ iff $I_1 \preceq I_2 \wedge I_2 \preceq I_1$. We also define a precedence relation $\prec \subseteq \mathcal{I} \times \mathcal{I}$ on events such that $I_1 \prec I_2$ iff $I_1 \preceq I_2 \wedge \neg(I_1 \approx I_2)$.

The actions of programs are scheduled to a set of such instants (also called reactions or macro steps [35]). The actions that take place within an instant (sometimes called micro steps) are not explicitly ordered. Instead, micro steps are assumed to happen simultaneously, i. e. in the same variable environment. Hence, variables seem to be constant during the execution of the micro steps and only change synchronously at macro steps. From the semantical point of view, which postulates that a reaction is atomic, neither communication nor computation take time in this sense. In reality, all actions within an instant are executed according to their data dependencies to establish the illusion of zero-time computations.

With the help of this temporal framework, we define the behavior of the system. It is given by a set of signals \mathcal{V} , which reflect the state changes in the course of the execution.

Definition 2 (Data Values and Signals) Let \mathcal{D} be the set of data values, i. e. all the values that program expressions may be evaluated to. A *signal* $x \in \mathcal{V}$ is a function that maps a totally ordered sequence of instants $C = \langle I_0, I_1, I_2, \dots \rangle \subseteq \mathcal{I}$ with $I_i \prec I_{i+1}$ to the data values \mathcal{D} .

Each variable $x \in \mathcal{V}$ is related to a clock $c \in \mathcal{C}$, which will be referred to as its clock \hat{x} in the following.

Definition 3 (Clocks) For each signal x , let $\text{Instants}(x) \subseteq \mathcal{I}$ be the domain of a signal x . This set gives rise to the *clock* \hat{x} of a signal x , which is the characteristic function of this set, i. e. it holds in instant $I \in \mathcal{I}$ iff $I \in \text{Instants}(x)$. The signal x is said to be *present* in an instant iff \hat{x} holds, otherwise the signal x is *absent*. Furthermore, two signals x_1 and x_2 are *synchronous to each other* iff they have the same domain, i. e. $\text{Instants}(x_1) = \text{Instants}(x_2)$. The union (intersection) of two clocks C_1 and C_2 is the union (intersection) of the set of their instances.

Clocks are fundamental elements of all synchronous languages, and they have a similar role in our approach. They define the possibly infinite set of instants at which the signal communicates a data value. As we have the underlying notion of perfectly synchronous instants [4], it is always possible to detect the absence of a signal (i. e. to be sure that an event for this signal does not arrive in the current instant) and based on this knowledge, initiate an action (the so-called reaction to absence).

2.2 Interfaces

The purpose of the proposed intermediate representation is to describe heterogeneous systems, i. e. systems that consist of components based on different MoCs. Thus, the basic structural unit of AIF⁺ is a *component*. Components are hierarchically organized in a tree structure, i. e. a component can aggregate a set of child components. Thus, in a system a component is either the top component (communicating with the environment) or it has a unique parent component.

In order to cooperate, components need to exchange information. To this end, each component has an *interface*, which defines how it interacts with its parent module. In AIF⁺, the interface consists of a list of signals $X = \langle x_1, \dots, x_n \rangle$ exposed to the parent module (or the environment in the case of the top module). As described in Section 2.1, a signal does not only consist of a sequence of values but its domain also fixes the instants in which the data is transmitted. Thus, the components at either end of the interface know when to exchange information. In order to define the direction of the communication, the *information flow* of each signal in the interface is explicitly declared. *Inputs* can be only read by a component, *outputs* only written, and *inouts* are shared between components.

For common data-flow models, this type of interface would be sufficient. The modules are usually hierarchically organized, but this is only used for structuring the system. Finally, all modules functionally run in parallel from the start of the system forevermore, and there are no means to start or stop a module during the execution. Figure 2 (a) illustrates this. When the whole system (module M_0 in the figure) is started, all its submodules are started, and they run in parallel. In this case, the *start* functionality is only given implicitly and not addressed in the system model.

However, since we want to integrate control-flow oriented models (such as hierarchical statecharts or imperative synchronous languages) as first-class citizens in our intermediate representation, we have to address starting, suspending and stopping a module. This is done by an additional control-flow context in the interface of an AIF⁺ module, e. g. activation or preemption conditions given by the surrounding component are passed to the child component and status information (e. g. about its activation or termination) are passed to the calling

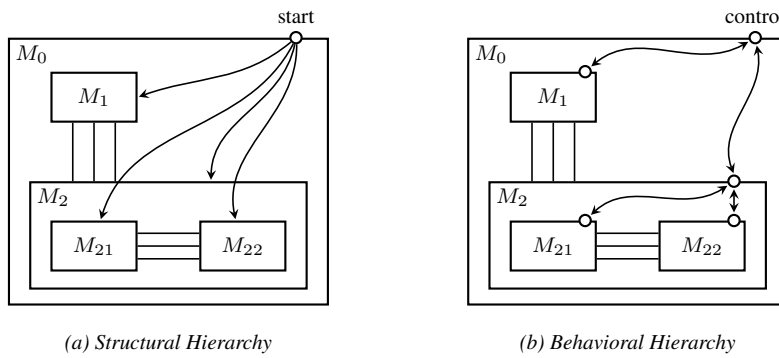


Fig. 2 Structural vs. Behavioral Hierarchy

module in return. This situation is illustrated in Figure 2 (b). In contrast to the structural approach, each module has a control interface, which does not only allow to start it but also to control its execution. In turn, the module itself provides some status information to the parent module e. g. it notifies the outer module whether it is still running or whether it terminates. The path of the control signals follows the hierarchy of the modules, and a module controls its submodules. Thus, in the example, M_2 can start and abort its submodule M_{21} . In this way, a module is e. g. able to execute its submodules in a sequence which is pretty common in control-flow oriented languages. The first submodule is started and not until it terminates the second one is started (and so on).

To this end, we implicitly extend the interface by eight additional signals of Boolean type, which indicate the control-flow context of a component. This is in the spirit of the compilation of imperative synchronous languages, which use similar parameters for of their compilation [8, 14, 64]. Each component is given the following five inputs.

- $\text{strt}(M)$ is the start or reset signal, which holds iff M should be (re)started in the given instant. With the help of this signal, the internal state of the component is initialized and its behavior is (re)started. Please not that the data variables appearing in the interface are not reset since they are shared with the surrounding component. Not also that a module may only be started if it is currently inactive or it currently terminates.
- $\text{prmt}(M)$ is a preemption signal for the initial instant. This signal can be used together with $\text{strt}(M)$ to immediately abort the behavior of a component so that only a the initialization is executed in a weakly preempted start.
- $\text{abrt}(M)$ is a signal that holds if an already running component M should be aborted in the given instant. As the compilation of imperative synchronous programs shows, we need to distinguish $\text{prmt}(M)$ and $\text{abrt}(M)$: the first one aborts a starting component, while the latter one aborts an already active one.
- $\text{susp}(M)$ similarly describes the suspension context: if this signal holds in an instant, M will be suspended. Suspension means that the current state is sustained as long as this control signal is set and the execution resumes from that state as soon as the suspension is over.
- $\text{strg}(M)$ is an enable condition. The signal is used to prevent that the component sets its outputs. With its help several mutually exclusive components may drive the same signals depending on control-flow context.

$\text{strt}(M)$ and $\text{prmt}(M)$ control the initial step of a component, while the other inputs are only read by an already running component. All the control inputs are mostly orthogonal, i. e. they can be set arbitrarily. The only potential conflict is a simultaneous set of $\text{abrt}(M)$ and $\text{susp}(M)$. If both are set for a component M , $\text{susp}(M)$ has a higher priority, i. e. then the suspension takes place. Combinations of the other signals can be used to control a specific control-flow behavior: e. g. $\text{abrt}(M) \wedge \text{strg}(M)$ means that the module is strongly aborted (the macro-step is abort before executing its actions), while $\text{abrt}(M) \wedge \neg\text{strg}(M)$ aborts it after the execution of its actions. To illustrate the interaction between control inputs for the initial step and the rest, consider another example: $\text{strt}(M) \wedge \text{abrt}(M)$ means that a currently running module should be aborted and immediately restarted in the current step.

While the above list shows the additional input signals, each component also provides additional output signals for its control flow:

- $\text{inst}(M)$ is a signal that holds if the execution of M will immediately terminate if it would now be started (M is often said to be instantaneous).
- $\text{insd}(M)$ holds at some point of time if the control flow is currently at some location inside M , i. e., if M is active.
- $\text{term}(M)$ is signal that holds if the control flow is currently somewhere inside M and wants to leave M voluntarily.

Similar to the data signals in the interface, we need to determine a clock for the additional control-flow signals so that this information can be aligned to the data streams. AIF⁺ implicitly assumes that all control-flow signals share the same clock and that this clock is the union of the clocks of the data signals of the interface. Therefore, the intermediate format does not store this information in its data structure.

This choice for the definition of the clock of the control-flow signals is motivated as follows: First, all control signals should have the same clock as they all refer to the same control-flow. Obviously, this clock should consist of the instants when control moves through the program. Second, as component-based design always aims at proving a well-defined functionality while abstracting from internal design decisions, the control-flow signals should be at least on the level of the data signals. Otherwise, internal intermediate computations may become visible to the environment, which violates the common principle of information hiding. We also get the advantage that components can be only preempted at well-defined instants (and not during internal computations), which avoids inconsistent system configurations. Finally, the control-flow should not be unnecessarily coarse-grained. These three considerations lead to our choice that the control-flow signals have the same clock as the union of the data-flow signals of the interface.

2.3 Behavior

As the name suggests, guarded actions basically consist of a guard γ and an action A , i. e. they have the form $\langle \gamma \Rightarrow A \rangle$. The intuition behind a guarded action is that the body A is executed if its guard evaluates to true in the current instant.

The guard γ is a Boolean expression over the program variables and their clocks. There are no syntactical restrictions for the guards. As they must be evaluated in all instants, our semantic model ensures that all variables can be evaluated in all instants of a program execution (see below).

In our intermediate representation, an action A can be either an assignment or a constraint: assignments are an operational description, which completely fix the value of the

written signal, while constraints may leave several possibilities: every behavior that complies to them is considered to be a valid one.

In our intermediate representation clocked guarded actions have one of the following forms:

- (a) $\gamma \Rightarrow x = \tau$
- (b) $\gamma \Rightarrow \text{next}(x) = \tau$
- (c) $\gamma \Rightarrow \text{assume}(\sigma)$
- (d) $\gamma \Rightarrow \text{assert}(\sigma)$

The first action (a) is an *immediate assignment*, which set the signal x at the given instant to the value of the expression τ . It implicitly imposes the constraint $\gamma \rightarrow \hat{x}$: the clock of x must hold whenever x is assigned. The delayed assignment (b) evaluates the expression τ in the given instant but changes the value of the variable x the next time clock \hat{x} ticks. Thus, no additional constraint is imposed by a delayed assignment since the instant where the variable is updated is defined by the next tick of its clock. Assumptions (c) and assertions (d) define constraints. They determine a Boolean condition which has to hold whenever the guard γ is true. The difference between them is that assumptions are guaranteed by the programmer, whereas assertions represent verification tasks, which have to be addressed in the following design flow.

In designs, the first instant usually differs from all other ones since additional behavior for initialization must be done. In the following, we use the expression $\text{init}(C)$ for any clock C , which exactly holds the first time C ticks.

Furthermore, there are generally several guarded actions that write a variable x . For each variable, the behavior part of our intermediate representation also defines a *default assignment* if no action determines its value in the current step. For a variable x , this is the case iff the guards of all immediate assignments to x are false in the current step and the guards of all delayed assignments to x have also been false in the preceding step. In this case, it takes the default value according to its intended storage mode: *event* variables $\text{eventVars} \subseteq \mathcal{V}$ are reset (like wires in hardware circuits), while *memorized* variables $\text{memVars} \subseteq \mathcal{V}$ store their previous values (like registers in hardware circuits). In general, this default assignment can be given by

$$(e) \text{ default}(x) = \tau$$

The intended meaning is that x will be given the value x in the *next* step if there is no delayed action in the current step and if there is no immediate action in the next step that write to x . Thus, the expression τ is 0 (or false) for event variables, and it is x for memorized variables.

Before we give a formal denotational semantics of the guarded actions, we first explain the evaluation of an expression τ at instant I , which will be denoted by $\llbracket \tau \rrbracket_I$ in the following. In contrast to some synchronous languages [4, 32], we can evaluate a variable even if its clock does not hold. Thus, the clock of a signal does identify when there is a value and when not. In AIF⁺, every variable can be read in every instant and the variable will have the value that was assigned to it when its clock has held the last time.² Thus, in AIF⁺ a clock ticks corresponds to a potential value change. Based on this evaluation of expressions we can define a formal denotational semantics.

Definition 4 (Consistency of Guarded Actions) A guarded action A is defined to be *consistent* (written $\text{Consistent}(A)$) w. r. t. a set of instants \mathcal{I} of a program execution as follows:

² This is somehow similar as reading the value $?x$ of Signal x in Esterel[7], or implicitly using the `cell` operator in Signal [27].

- $\text{Consistent}(\gamma \Rightarrow x = \tau) := \forall I \in \mathcal{I}. \llbracket \gamma \rrbracket_I \rightarrow (\llbracket \hat{x} \rrbracket_I \wedge \llbracket x \rrbracket_I = \llbracket \tau \rrbracket_I)$
- $\text{Consistent}(\gamma \Rightarrow \text{next}(x) = \tau) := \forall I_1 \prec I_2 \in \mathcal{I}. \llbracket \gamma \rrbracket_{I_1} \rightarrow (\llbracket \hat{x} \rrbracket_{I_2} \wedge \llbracket x \rrbracket_{I_2} = \llbracket \tau \rrbracket_{I_1})$
 where I_2 such that $\nexists I' \in \mathcal{I}. (I_1 \prec I' \prec I_2) \wedge \llbracket \hat{x} \rrbracket_{I'}$
- $\text{Consistent}(\gamma \Rightarrow \text{assume}(\sigma)) := \forall I \in \mathcal{I}. \llbracket \gamma \rrbracket_I \rightarrow \llbracket \sigma \rrbracket_I$
- $\text{Consistent}(\gamma \Rightarrow \text{assert}(\sigma)) := \forall I \in \mathcal{I}. \llbracket \gamma \rrbracket_I \rightarrow \llbracket \sigma \rrbracket_I$
- $\text{Consistent}(\text{default}(x) = \tau) := \forall I_1 \prec I_2 \in \mathcal{I}. \xi \rightarrow (\llbracket \hat{x} \rrbracket_{I_2} \wedge \llbracket x \rrbracket_{I_2} = \llbracket \tau \rrbracket_{I_1})$
 where I_2 such that $\nexists I' \in \mathcal{I}. (I_1 \prec I' \prec I_2) \wedge \llbracket \hat{x} \rrbracket_{I'}$
 and $\xi = \left(\llbracket \bigwedge_{(\gamma \Rightarrow x = \tau)} \neg \gamma \rrbracket_{I_2} \wedge \llbracket \bigwedge_{(\gamma \Rightarrow \text{next}(x) = \tau)} \neg \gamma \rrbracket_{I_1} \right)$

A set of guarded actions \mathcal{A} is consistent if all its elements are consistent:

- $\text{Consistent}(\mathcal{A}) := \forall A \in \mathcal{A}. \text{Consistent}(A)$

We are convinced that the representation of the behavior by clocked guarded actions is exactly at the *right level of abstraction* for an intermediate code format, since guarded actions provide a good balance between (1) removal of complexity from the source code level and (2) the independence of a specific synthesis target. (This will be illustrated by Section 4, which gives straightforward translations targeting symbolic model checking and simulation). On the one hand, problems such as schizophrenia [8, 57, 65] and the semantics of complex control flow statements like preemption statements can be completely solved during the translation to guarded actions, so that subsequent analysis, optimization and synthesis become much simpler. On the other hand, despite their very simple structure, efficient translation to both software and hardware is efficiently possible from guarded actions.

Guarded actions allow *many analyses and optimizations*. In particular, causality analysis can be effectively performed on guarded actions. If the causality analysis determined that a set of guarded actions does always have a dynamic schedule to compute the variables without first reading them in each macro step, then even an acyclic set of guarded actions can be determined. Other transformations on guarded actions are the grouping of guarded actions with regard to the variable they modify, which corresponds to the generation of static single-assignment form in the compilation of sequential languages. For synchronous languages this is often called an equational code generation, since for every variable, a single equation is generated.

2.4 Structure

Having explained the parts contained in the AIF⁺ intermediate representation, we can now put all parts together to describe the whole structure. In the representation, we distinguish between *modules* and *systems*. Thereby, an AIF⁺ module (keyword **module**) contains the control interface, which allows the composition with other modules. In contrast, an AIF⁺ system (keyword **system**) is considered as the result of the composition of modules. The control interface is bound, and all module calls resolved, i. e. it is fully linked. The only control signal which is still used in a system is `strt (M)` to start the system.

The structure of the AIF⁺ intermediate format is presented in Figure 3. The data interface contains the declarations of inputs, outputs and local variables. The behavior is given by clocked guarded actions. The control interface is described in the last part of the structure. These are the conditions which are defined by the module. The control signals which come from the outside, are simply used as expressions in the behavior, i. e. the clocked guarded actions. For a system, the control part is simply omitted.

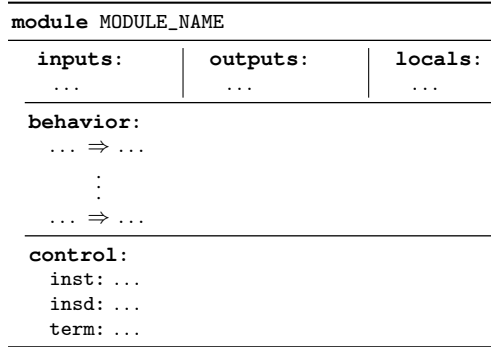


Fig. 3 AIF⁺ Module Structure

3 Translating to Clocked Guarded Actions

In this section, we show how different models can be compiled to clocked guarded actions. For some of them, we make use of existing translations (as cited in the subsections) which transform systems to some kind of guarded actions. Hence, the starting point are models which have a similar syntax (see Figure 1). Hence, we only need a few adaptations to achieve the intended integration, which lets us focus on the core: the mapping of clocks and synchronizations.

The rest of this section describes the foundations of different classes of systems which we will consider in this article, namely single-clocked synchronous programs in Section 3.1, polychronous specifications in Section 3.2 and finally concurrent action-oriented specifications in Section 3.3. For each of them, we briefly describe their semantics before we show how they can be represented by clocked guarded actions as introduced in the previous section.

3.1 Synchronous Programs

The synchronous *model of computation* [4, 32] assumes that the execution of programs consists of a totally ordered sequence of instants. In each of these instants, the system reads its inputs and computes and writes its outputs. In the single-clocked case, which we will consider in the following, all signals have a value in every instant. The introduction of this logical time scale is not only the key for a straightforward translation of synchronous programs to hardware circuits [5, 58, 64]; it also provides a very convenient programming model, which allows compilers to generate *deterministic* single-threaded code from multi-threaded synchronous programs [6].

In general, synchronous programs such as Esterel, Lustre or Quartz can be translated to synchronous guarded actions [14]. This translation extracts all actions (assignments, assumptions and assertions) of the program and computes for each of them a trigger condition from its context in the program. As already stated in the previous section, the *semantics of synchronous guarded actions* implements the synchronous model of computation and fires all activated actions simultaneously in each macro step. Synchronous guarded actions without causality problems are always deterministic since there is no choice due to the firing of all actions.

$$\left[\begin{array}{c} \gamma_1 \Rightarrow A_1 \\ \vdots \\ \gamma_n \Rightarrow A_n \end{array} \right] \Rightarrow \left[\begin{array}{c} C \wedge \gamma_1 \Rightarrow A_1 \\ \vdots \\ C \wedge \gamma_n \Rightarrow A_n \\ (\text{strt}(\mathbb{M}) \vee \text{insd}(\mathbb{M}) \wedge \neg \text{strg}(\mathbb{M})) \Rightarrow \text{assume}(\widehat{x}_1 = C) \\ \vdots \\ (\text{strt}(\mathbb{M}) \vee \text{insd}(\mathbb{M}) \wedge \neg \text{strg}(\mathbb{M})) \Rightarrow \text{assume}(\widehat{x}_m = C) \end{array} \right]$$

Fig. 4 Translating Synchronous Guarded Actions of Module \mathbb{M} to Clocked Guarded Actions

```

module Sequence
  (bool ?switch, int ?i,!o,!mode)
{
  abort {
    mode = 1;
    M1(i, o); // call module M1
  } when (switch);
  ℓ: pause;
  abort {
    mode = 2;
    M2(i, o); // call module M2
  } when (switch);
}

```

Fig. 5 Quartz Example: Sequence

As expected, translating synchronous guarded actions to clocked guarded actions is straightforward. We only need to introduce a single clock C , which serves as the clock for all variables of the system. This clock C is then added as an additional clause to the guard of all actions. Additional clock constraints ensure that this clock is the clock of all variables: Whenever C holds, the original system performs a computation step. The general principle of the translation is shown in Figure 4 (a). A system in the AIF-format, which is based on synchronous guarded actions, contains a set of guarded actions which are executed in any instant. The guards of the guarded actions are strengthened by the clock C . Thus, they are now executed at any tick of C . The clock constraints ensure that all variables have the same clock. The clock constraints are just taken into account when the module is running which is the case when the module is started ($\text{strt}(\mathbb{M})$), or when the control flow is inside the module ($\text{insd}(\mathbb{M})$). In the second case, it also has to be ensured that the module is not aborted or suspended ($\neg \text{strg}(\mathbb{M})$).

An AIF component already stores a control-flow context, which consists of the signals described in Section 2.2. Since all the variables of the interface have the same clock, its translation to AIF⁺ is trivial: the conditions are stored without any modification in AIF⁺.

We will now give an example AIF⁺ representation of an synchronous module. Since the representation in guarded actions is pretty simple to the representation with clocked guarded actions, we will omit this version. However, we will give the synchronous example in Quartz code which provides a more readable representation and the translation to guarded actions is given in [14] and also mentioned in Figure 1. The behavior of the example Quartz module Sequence which is given in Figure 5 is the following. It gets the input i and produces the output o . However, the real computation of o is done either by the submodule M1 or M2. The module Sequence does *just* control which of these both submodules drive the output. First, the output is computed by the submodule M1 until it is finished or its execution is aborted

module Sequence		
inputs: bool switch int i	outputs: int o, mode	locals: clock C label ℓ
behavior:		
$C \wedge \text{str}(\text{Sequence}) \Rightarrow \text{mode} = 1$		
$C \wedge \text{str}(\text{Sequence}) \Rightarrow \left\{ \begin{array}{l} \text{M1}(i, o) \{ \\ \text{prmt}: \text{prmt}(\text{Sequence}) \\ \text{abrt}: \text{abrt}(\text{Sequence}) \vee \text{switch} \\ \text{susp}: \text{susp}(\text{Sequence}) \\ \text{strg}: \text{strg}(\text{Sequence}) \vee \text{switch} \\ \} \end{array} \right.$		
$C \wedge \text{str}(\text{Sequence}) \wedge \text{inst}(\text{M1}) \Rightarrow \text{next}(\ell) = \text{true}$		
$C \wedge \ell \wedge \text{susp}(\text{Sequence}) \Rightarrow \text{next}(\ell) = \text{true}$		
$C \wedge \left(\begin{array}{l} \text{term}(\text{M1}) \wedge \\ \neg \text{strg}(\text{Sequence}) \wedge \\ \neg (\text{susp}(\text{Sequence}) \vee \\ \text{abrt}(\text{Sequence})) \end{array} \right) \Rightarrow \text{next}(\ell) = \text{true}$		
$C \wedge \ell \wedge \neg \text{strg}(\text{Sequence}) \Rightarrow \text{mode} = 2$		
$C \wedge \ell \wedge \neg \text{strg}(\text{Sequence}) \Rightarrow \left\{ \begin{array}{l} \text{M2}(i, o) \{ \\ \text{prmt}: \text{abrt}(\text{Sequence}) \vee \text{susp}(\text{Sequence}) \\ \text{abrt}: \text{abrt}(\text{Sequence}) \vee \text{switch} \\ \text{susp}: \text{susp}(\text{Sequence}) \\ \text{strg}: \text{strg}(\text{Sequence}) \vee \text{switch} \\ \} \end{array} \right.$		
$\left(\begin{array}{l} \text{str}(\text{Sequence}) \vee \\ \text{insd}(\text{Sequence}) \wedge \neg \text{strg}(\text{Sequence}) \end{array} \right) \Rightarrow \text{assume}(\widehat{\text{switch}} = C)$		
$\left(\begin{array}{l} \text{str}(\text{Sequence}) \vee \\ \text{insd}(\text{Sequence}) \wedge \neg \text{strg}(\text{Sequence}) \end{array} \right) \Rightarrow \text{assume}(\widehat{i} = C)$		
$\left(\begin{array}{l} \text{str}(\text{Sequence}) \vee \\ \text{insd}(\text{Sequence}) \wedge \neg \text{strg}(\text{Sequence}) \end{array} \right) \Rightarrow \text{assume}(\widehat{o} = C)$		
$\left(\begin{array}{l} \text{str}(\text{Sequence}) \vee \\ \text{insd}(\text{Sequence}) \wedge \neg \text{strg}(\text{Sequence}) \end{array} \right) \Rightarrow \text{assume}(\widehat{\text{mode}} = C)$		
$\text{true} \Rightarrow \text{assume}(\widehat{\ell} = C)$		
control:		
inst: false		
insd: ℓ ∨ insd(M1) ∨ insd(M2)		
term: ℓ ∧ inst(M2) ∨ term(M2) ∨ switch ∧ insd(M2)		

Fig. 6 AIF⁺ of Example: Sequence

by the signal `switch`. Then, M2 computes the output until it is finished or its execution is aborted by the signal `switch`. The resulting AIF⁺ representation is given in Figure 6. According to the translation described above, except for the clocks it is the same as the AIF representation. If the module is started, i. e. `str(Sequence)`, the assignment `mode=1` is executed and also the module M1 is called. For calling the module, also the control interface has to be defined. In this case it is combined from the control interface of `Sequence` and the **abort** statement. Thus, e. g. the module M1 is aborted, when the internal abort condition (`switch`) holds or an abort condition is given from the outside (`abrt(Sequence)`). The module M2 is started, when the module M1 terminates or when it is aborted with the input `switch`.

We now introduce the first part of a *running example*, which we will use in the remainder of this paper. It shows how modules and processes written in different languages and based on different models can be combined on the basis of AIF⁺. Thereby, it will show how a

```

module OuterQuartz
  (bool ?susp, int ?i, !o)
{
  suspend {
    InnerSignal(i, o);
  } when (susp);
}

module InnerQuartz
  (int ?y, !x)
{
  loop {
    x = y;
    pause;
    x = 2 * y;
    pause;
  }
}

```

Fig. 7 Running Example: Quartz Modules OuterQuartz and InnerQuartz

module OuterQuartz		
inputs: bool susp int i	outputs: int o	locals: clock C1
behavior:		
$ \begin{aligned} & \text{true} \Rightarrow \text{assume}(\widehat{i} = C1) \\ & \text{true} \Rightarrow \text{assume}(\widehat{\text{susp}} = C1) \\ & \text{true} \Rightarrow \text{assume}(\widehat{o} = C1) \\ & C \wedge \text{strt}(\text{OuterQuartz}) \Rightarrow \left\{ \begin{array}{l} \text{InnerSignal}(i, o) \{ \\ \quad \text{prmt: prmt}(\text{OuterQuartz}) \\ \quad \text{abrt: abrt}(\text{OuterQuartz}) \\ \quad \text{susp: susp}(\text{OuterQuartz}) \vee \text{susp} \\ \quad \text{strg: strg}(\text{OuterQuartz}) \vee \text{susp} \\ \} \end{array} \right. \end{aligned} $		
control:		
inst: inst(InnerSignal) insd: insd(InnerSignal) term: term(InnerSignal)		

Fig. 8 AIF⁺ of Example: OuterQuartz

hierarchical structure as it is described in Section 2.2 can be established. In total, the running example consists of two Quartz modules, which are described here, and another module, which will be explained later.

The Quartz modules OuterQuartz and InnerQuartz are shown in Figure 7. OuterQuartz is the outermost module of the hierarchy and forms the interface to the environment. It takes two inputs `susp` and `i` and produces an output `o`. Inside, it calls a module named InnerSignal, which will be described later. OuterQuartz just calls the sub-module and provides a suspension context based on the input `susp`. InnerSignal by itself will make use of the second Quartz module InnerQuartz. It takes one input `y` and produces one output `x`, which is set either to $x = y$; or $x = 2 * y$;

The AIF⁺ descriptions of both modules are shown in Figure 8 and Figure 9. Thereby, OuterQuartz synchronizes the clocks of the inputs and outputs and gives the interface signals to InnerSignal. The suspension context is strengthened by the input `susp`. InnerQuartz does not call a submodule, but it shows the usage of the interface signals for its own behavior.

module InnerQuartz		
inputs: int y	outputs: int x	locals: clock $C2$ label ℓ_1, ℓ_2
behavior:		
		$\text{true} \Rightarrow \text{assume}(\widehat{y} = C2)$
		$\text{true} \Rightarrow \text{assume}(\widehat{x} = C2)$
		$\text{true} \Rightarrow \text{assume}(\widehat{\ell}_1 = C2)$
		$\text{true} \Rightarrow \text{assume}(\widehat{\ell}_2 = C2)$
	$C2 \wedge \left(\text{str}(\text{InnerQuartz}) \vee \ell_2 \wedge \neg \text{str}(\text{InnerQuartz}) \right)$	$\Rightarrow x = y$
	$C2 \wedge (\ell_1 \wedge \neg \text{str}(\text{InnerQuartz}))$	$\Rightarrow x = 2 * y$
$C2 \wedge \left(\begin{array}{l} \text{str}(\text{InnerQuartz}) \wedge \neg \text{pr}(\text{InnerQuartz}) \vee \\ \ell_1 \wedge \text{susp}(\text{InnerQuartz}) \vee \\ \ell_2 \wedge \neg \text{str}(\text{InnerQuartz}) \wedge \\ \neg(\text{susp}(\text{InnerQuartz}) \vee \text{abrt}(\text{InnerQuartz})) \end{array} \right)$		$\Rightarrow \text{next}(\ell_1) = \text{true}$
$C2 \wedge \left(\begin{array}{l} \ell_2 \wedge \text{susp}(\text{InnerQuartz}) \vee \\ \ell_1 \wedge \neg \text{str}(\text{InnerQuartz}) \wedge \\ \neg(\text{susp}(\text{InnerQuartz}) \vee \text{abrt}(\text{InnerQuartz})) \end{array} \right)$		$\Rightarrow \text{next}(\ell_2) = \text{true}$
control:		
inst: false		
insd: $\ell_0 \vee \ell_1$		
term: false		

Fig. 9 AIF⁺ of Example: InnerQuartz

3.2 Polychronous Programs

Polychronous specifications [27, 28, 49] as implemented by Signal use several clocks, which means that signals do not need to be present in all instants. Furthermore, in contrast to synchronous systems, polychronous models are not based on a linear model of time, so that the reactions of a polychronous system are partially ordered. Two instants are only compared on the time scale if both contain events on a shared signal x .

Polychronous specifications are usually considered to be *relational* and not *functional*: even in the presence of the same input values, various temporal alignments, which comply to the clock constraints, may lead to different output values. Hence, polychronous models are generally nondeterministic and should be seen as specifications, which describe a set of acceptable implementations. State-of-the-art tools check for determinism before synthesis [27, 28].

Signal programs consist of several processes, where each process is either given by a set of equations or a composition of other processes. Each process has an input interface consisting of input signals, an output interface consisting of output signals and several possible internal signals. The equations can be built from one of the following primitive operators:

- *Function*: A function $y := f(x_1, \dots, x_n)$ determines the output y by applying the given function f to the input values. Additionally, this process requires that all inputs and the output have the same clock.
- *Delay*: The delay operator $y := x \ \$ \ \text{init } d$ has exactly one input x and one output y . Each time a new incoming value arrives, it outputs the previously stored value and stores the new one. Initially, the buffer simply returns the given value d . By definition, the input and the output have the same clock, i. e. $\widehat{x} = \widehat{y}$.

$$\begin{array}{l}
y := f(x_1, \dots, x_n) \Rightarrow \left[\begin{array}{l} \widehat{y} \wedge \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow y = f(x_1, \dots, x_n) \\ \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = \widehat{x}_1) \\ \dots \\ \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = \widehat{x}_n) \end{array} \right] \\
y := x \$ \text{init } d \Rightarrow \left[\begin{array}{l} \widehat{y} \wedge \text{str}(\mathbf{M}) \Rightarrow y = d \\ \widehat{y} \wedge \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \Rightarrow y = y' \\ \widehat{y} \Rightarrow \text{next}(y') = x \\ \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = \widehat{x}) \\ \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = \widehat{y}') \end{array} \right] \\
y := x \text{ when } z \Rightarrow \left[\begin{array}{l} \widehat{y} \wedge \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow y = x \\ \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = (z \wedge \widehat{z} \wedge \widehat{x})) \end{array} \right] \\
y := x \text{ default } z \Rightarrow \left[\begin{array}{l} \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow y = x \\ \widehat{z} \wedge \neg \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow y = z \\ \left(\begin{array}{l} \text{str}(\mathbf{M}) \vee \\ \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = (\widehat{x} \vee \widehat{z})) \end{array} \right]
\end{array}$$

Fig. 10 Translation from Signal to Clocked Guarded Actions

- *When*: The downsample operator $y := x$ **when** z has two inputs, x of arbitrary type and z of Boolean type, and one output port y . Each time a new x arrives, it checks whether there is an input at z . If there is one and if it is true, a new output event with the value of x is emitted for y . In all other cases, no event will be produced.
- *Default*: The merge operator $y := x$ **default** z has two input ports x and z and a single output port y . Each time inputs arrive at x and z , they will be forwarded to y . If there are events present at both ports in a particular instant, x will be forwarded, and z will be discarded.

Programs may contain more clock constraints to restrict the behaviors of clocks. They are very general: for example, a clock can be declared to be a subclock of another one $\widehat{x} \rightarrow \widehat{y}$.

Polychronous Signal programs can be structurally translated to clocked guarded actions by translating each operator separately. Figure 10 shows the translation where the condition:

$$\text{str}(\mathbf{M}) \vee \text{insd}(\mathbf{M}) \wedge \neg \text{strg}(\mathbf{M})$$

indicates that the module is currently executed and is added to the guards.

- *Function*: The function operator is applied to the inputs and produces the output value. All variables are forced to have the same clock for a function application.
- *Delay*: The delay operator violates the rule to add the above condition to all guards, because its behavior is split into two cases. (1) The first value that is produced by this operator when the process is started is the value that is given by the constant d . (2) In all other steps the value of x of the last tick is used. Therefore, we model this behavior by

transferring the value of x to the following step in every tick of the additional signal y' . The constraint ensures that all variables have the same clock.

- *When*: The sample operator **when** transfers the value of x to y whenever it is needed. The clock constraint ensures that \hat{y} only holds when both inputs are present and z holds.
- *Default*: The **default** operator merges two signals with priority for the first one. Therefore, if the first input is present, it is passed to x . If it is not present, but the second one is, the second one is passed through. The clock constraint ensures that \hat{y} only holds when at least one of the inputs is present.

Additional clock constraints ϕ should hold in every instant when the process is running. Note that there is an elementary difference between AIF^+ and Signal (and thereby also DC^+). In Signal , it is not possible to read a variable when its clock does not hold. Thus, in its context the clock of a signal identifies when there is a value and when not. As explained in Section 2.3, in AIF^+ every variable can be read in every instant and the variable will have the value that was assigned to it when its clock has held the last time. Thus, in AIF^+ the clock means a potential value change. Nevertheless, the translation of Signal to AIF^+ works because it ensures that every variable is read or set if and only if its clock holds. This can be easily checked in Figure 10 (b). However, interaction with other computational models is done in the model of AIF^+ by allowing to read the variable in every instant.

The translation above uses the control signal $\text{insd}(M)$ which is usually defined by the process itself but not yet explained for Signal processes. In order to fully translate Signal programs to AIF^+ , we have to define a rudimentary control-flow. We introduce a local Boolean event variable ℓ , which simply models the activation of the component: if ℓ is set, the control-flow is currently inside the component and all the equations of the data-flow description are activated; if ℓ is not set, the component is inactive. Its value is set by the following guarded actions:

- $\hat{\ell} \wedge \text{strt}(M) \wedge \neg \text{prmt}(M) \Rightarrow \text{next}(\ell) = \text{true}$
- $\hat{\ell} \wedge \ell \wedge \neg \text{abrt}(M) \Rightarrow \text{next}(\ell) = \text{true}$

The first action considers the activation of the component. If it is started and not immediately preempted, the component becomes active and therefore ℓ true. The second action models a running component: it will remain active unless an abortion takes place (suspension does not have an effect on ℓ). We use this flag to provide the necessary control-flow information to its parent component:

- $\text{inst}(M) = \text{false}$
- $\text{insd}(M) = \ell$
- $\text{term}(M) = \text{false}$

As the component is considered to run the given SINGAL equations, it is never instantaneous. We modeled its activation by ℓ so that we can use it for the $\text{insd}(M)$ signal. Finally, it can only terminate if an explicit abortion is triggered from the output.

Based on the control-flow context the data-flow of the component can then be guarded: each guarded action $\gamma \Rightarrow A$ is only executed if the module is started or if it is running and the data-flow enabled, i. e. its guard is strengthened to $\gamma \wedge (\text{strt}(M) \vee \ell \wedge \text{strg}(M))$.

A first example illustrates the translation of a Signal specification to an AIF^+ system. The Signal process `Counter` is shown in Figure 11. The intention of the process is that for each input value n , the output values $n, n-1, \dots, 0$ are produced. To this end, the local signal c stores the last value of the produced output, whereas o is produced by subtraction of 1 from c . However, when a new value for the input n arrives, the output is updated by this value.


```

process Counter =
  (? integer n;
   ! integer o;)
  (| c := o $ init 0
   | o := n default (c-1)
   | n ^= (when (c=0))
   |)
  where
    integer c;
  end;

```

$n = [2, \square, \square, 1, \square]$

$c = [0, 2, 1, 0, 1]$

$o = [2, 1, 0, 1, 0]$

Trace 1

$n = [2, \square, 1, \square]$

$c = [0, 2, 1, 1]$

$o = [2, 1, 1, 0]$

Trace 2

$n = [2, \square, \square, 1, \square]$

$c = [0, 2, 1, 0, 1]$

$o = [2, 1, 0, 1, 0]$

Trace 1

$n = [2, \square, 1, \square]$

$c = [0, 2, 1, 1]$

$o = [2, 1, 1, 0]$

Trace 2

Fig. 11 Signal Example: Counter

system Counter		
inputs: int n	outputs: int o	locals: int c, c' label ℓ
behavior:		
$\widehat{c} \wedge \text{str}(\text{Counter}) \Rightarrow c = 0$		
$\widehat{c} \wedge \ell \Rightarrow c = c'$		
$\widehat{c}' \wedge (\text{str}(\text{Counter}) \vee \ell) \Rightarrow \text{next}(c') = o$		
$(\text{str}(\text{Counter}) \vee \ell) \Rightarrow \text{assume}(\widehat{c} = \widehat{o})$		
$(\text{str}(\text{Counter}) \vee \ell) \Rightarrow \text{assume}(\widehat{c} = \widehat{c}')$		
$\widehat{n} \wedge (\text{str}(\text{Counter}) \vee \ell) \Rightarrow o = n$		
$\widehat{c} \wedge \neg \widehat{n} \wedge (\text{str}(\text{Counter}) \vee \ell) \Rightarrow o = c$		
$(\text{str}(\text{Counter}) \vee \ell) \Rightarrow \text{assume}(\widehat{o} = \widehat{x} \vee \widehat{z})$		
$(\text{str}(\text{Counter}) \vee \ell) \Rightarrow \text{assume}(\widehat{n} = (\widehat{c} \wedge c == 0))$		
$\widehat{\ell} \wedge \text{str}(\text{Counter}) \Rightarrow \text{next}(\ell) = \text{true}$		
$\widehat{\ell} \wedge \ell \Rightarrow \text{next}(\ell) = \text{true}$		
$\text{true} \Rightarrow \text{assume}(\widehat{\ell} = \text{str}(\widehat{\text{Counter}}))$		

Fig. 12 AIF⁺ of Example: Counter

The clock constraint $n \wedge = (\text{when } (c = 0))$ ensures that a new input is only allowed to arrive when the local signal c reaches 0. On the right side of the figure, two sample traces for this example are shown. Thereby, \square indicates the absence of a signal, i. e. it is not present in the instant. The first trace is a valid one and shows the desired behavior. First, 2 arrives as input and the output produces sequentially the values 2, 1, 0. After that, the local signal c is 0 and a new input is allowed to arrive. The second trace is an invalid one, because the second input value of n arrives too early and thus, the clock constraint is not fulfilled by this execution. Note, that without the given clock constraint both traces are valid but the constraint selects just the first one to be valid. The translation to an AIF⁺ system of this example is shown in Figure 12. Each given Signal equation is translated. The additional label ℓ stores the activation state of the process.

The second Signal example `InnerSignal`, which is shown in Figure 13 is part of the running example, and it illustrates the translation to an AIF⁺ module. It takes the input i and produces the output o . The local signal s alternates the values `sigtrue` and `false` each time a new input i is given (its clock is set to be the same as the clock of i). The signal y holds the value of the input i whenever s is `true`, thus, for every second input. The Quartz module `InnerQuartz` is used to compute the value of x for every value of y . Finally, the output o is set to x if it is present, otherwise the given input i is passed through. In summary, every second input value of i is passed to `InnerQuartz` to compute the output o , all other inputs

```

process InnerSignal =
  (? integer i;
   ! integer o;)
  (| s := (s $ init true) xor true
   | y := i when s
   | o := x default i
   | i ^= s
   | InnerQuartz(y, x)
  )
where
  integer x, y; boolean s;
end;

```

Fig. 13 Running Example: Signal Process InnerSignal

are directly passed to o . The resulting AIF⁺ module of the process InnerSignal is shown in Figure 14.

3.3 Asynchronous (Untimed) Programs

The third class of programs we consider in this article are the asynchronous (untimed) programs, which do not have an underlying notion of synchrony. They focus on describing the causalities of actions, i. e. which event happens after an other one, thereby, deferring the difficult task of global scheduling and coordination to compilers or runtime environments.

In particular, this class includes languages such as Concurrent Action-Oriented Specifications [1, 40, 41] (CAOS), a language intended to describe the data-flow of hardware circuits, or SHIM [24], a multi-threaded language based on asynchronous threads with a rendezvous-style communication [38, 39].

Similar to the synchronous and polychronous programs before, we do not consider these languages as the starting point of our translation to AIF. Instead, we start with simple asynchronous guarded actions, which can be used as an intermediate representation of the compiler. The step from full CAOS to this intermediate representation is straightforward: it only consists of a simple dismantle step [16]. SHIM can be also efficiently translated to asynchronous guarded actions as described in [15].

In our approach, we take asynchronous guarded actions in the form of rules and methods [16]. Thereby, the behavior is described by a set of *rules*, which are guarded atomic actions of the form:

$$\mathbf{rule} \ r_i \ \mathbf{when}(\gamma_{r_i}) \ B_i$$

Thereby, γ_{r_i} is the guard and B_i the body of rule r_i . CAOS provides two kinds of assignments: while wire assignments are immediately visible, register assignments are committed with the current state update. Hence, the body of a rule B_i is a set of synchronous guarded actions of the form $\langle \gamma \Rightarrow x = \tau \rangle$ (for an immediate assignment) or $\langle \gamma \Rightarrow \text{next}(x) = \tau \rangle$ (for a delayed assignment) as known from synchronous programs (see Section 3.1).

For the interaction with the environment, the target model makes use of so-called methods, which are parameterized rules. In addition to the local variables, the actions of a method have access to the variables specified in its parameter list, which may contain inputs and outputs.

$$\mathbf{method} \ m_i(p_{i1}, p_{i2}, \dots) \ \mathbf{when}(\gamma_{m_i}) \ B_1$$

module InnerSignal		
inputs: int i	outputs: int o	locals: int x, y bool s, s' label ℓ
behavior:		
$\widehat{s} \wedge \text{str}(\text{InnerSignal}) \Rightarrow s = \text{true}$ $\widehat{s} \wedge \ell \wedge \neg \text{susp} \Rightarrow s = s'$		
$\widehat{s}' \wedge \left(\begin{array}{l} \text{str}(\text{InnerSignal}) \vee \\ \ell \wedge \neg \text{strg}(\text{InnerSignal}) \end{array} \right) \Rightarrow \text{next}(s') = s \odot \text{true}$		
$\widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{InnerSignal}) \vee \\ \ell \wedge \neg \text{strg}(\text{InnerSignal}) \end{array} \right) \Rightarrow \text{assume}(\widehat{s} = \widehat{s}')$		
$\widehat{i} \wedge \neg \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{InnerSignal}) \vee \\ \ell \wedge \neg \text{strg}(\text{InnerSignal}) \end{array} \right) \Rightarrow o = x$		
$\widehat{i} \wedge \neg \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{InnerSignal}) \vee \\ \ell \wedge \neg \text{strg}(\text{InnerSignal}) \end{array} \right) \Rightarrow y = i$		
$\widehat{i} \wedge \neg \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{InnerSignal}) \vee \\ \ell \wedge \neg \text{strg}(\text{InnerSignal}) \end{array} \right) \Rightarrow \text{assume}(\widehat{i} = \widehat{s})$		
$\widehat{i} \wedge \neg \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{InnerSignal}) \vee \\ \ell \wedge \neg \text{strg}(\text{InnerSignal}) \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = s \wedge \widehat{s} \wedge \widehat{i})$		
$\widehat{\ell} \wedge \text{str}(\text{InnerSignal}) \wedge \neg \text{prmt}(\text{InnerSignal}) \Rightarrow \text{next}(\ell) = \text{true}$		
$\widehat{\ell} \wedge \ell \Rightarrow \text{next}(\ell) = \text{true}$		
$\ell \Rightarrow \text{assume}(\widehat{\ell} = \text{str}(\widehat{\text{Counter}}))$		
$C \wedge \text{str}(\text{OuterQuartz}) \Rightarrow \left\{ \begin{array}{l} \text{InnerSignal}(i, o) \{ \\ \text{prmt}: \text{prmt}(\text{InnerSignal}) \\ \text{abrt}: \text{abrt}(\text{InnerSignal}) \\ \text{susp}: \text{susp}(\text{InnerSignal}) \\ \text{strg}: \text{strg}(\text{InnerSignal}) \\ \} \end{array} \right.$		
control:		
inst: false		
insd: ℓ		
term: false		

Fig. 14 AIF⁺ of Example: InnerSignal

As already described in Section 2, the semantics of the asynchronous guarded actions is as follows: first the guards of all actions are evaluated with respect to the current state, then an arbitrary activated one is chosen and its body is executed. Inside the body, there are multiple synchronous actions, which are considered to execute in parallel. Hence, let q_0 be the initial state of the system, and $q \xrightarrow{S} q'$ indicate that action S transforms the system in state q to state q' . Then, a run of a model is a sequence of system states $\langle q_0, q_1 \dots \rangle$ where $q_i \xrightarrow{S_x} q_{i+1}$ and $\text{when}(\gamma_x) C_x$ is an arbitrary action which is activated in state q_i , i. e. $q_i(\gamma_x) = \text{true}$. Obviously, the system description is nondeterministic: even in the presence of the same inputs, which lead to the same activation of guards, the system can produce different outputs by choosing different activated actions. Models consisting of asynchronous guarded actions are generally intended to be specifications, which describe a set of acceptable implementations.

The translation of CAOS to AIF⁺ is illustrated in Figure 15. In order to model the non-determinism, a clock C_r for each rule r and a clock C_m for each method m is introduced. A tick of such a clock models an execution of the rule or method. First, the rules and the methods are translated on their own as shown in the figure. The guard of each action of a

$$\begin{array}{l}
\mathbf{rule} \ r_i \ \mathbf{when}(\gamma_{r_i}) \\
\alpha_{r_i1} \Rightarrow A_{r_i1} \\
\alpha_{r_i2} \Rightarrow A_{r_i2} \\
\dots
\end{array}
\Rightarrow
\begin{cases}
C_{r_i} \wedge \alpha_{r_i1} \Rightarrow A_{r_i1} \\
C_{r_i} \wedge \alpha_{r_i2} \Rightarrow A_{r_i2} \\
\dots \\
\mathbf{true} \Rightarrow \mathbf{assume}(C_{r_i} \rightarrow \gamma_{r_i}) \\
\mathbf{true} \Rightarrow \mathbf{assume}(C_{r_i} \rightarrow \neg \bigvee_{C \in \{C_{m_j}, C_{r_j}\} \setminus C_{r_i}} C)
\end{cases}$$

$$\begin{array}{l}
\mathbf{method} \ m_i(p_{i1}, \dots) \\
\mathbf{when}(\gamma_{m_i}) \\
\alpha_{m_i1} \Rightarrow A_{m_i1} \\
\alpha_{m_i2} \Rightarrow A_{m_i2} \\
\dots
\end{array}
\Rightarrow
\begin{cases}
C_{m_i} \wedge \alpha_{m_i1} \Rightarrow A_{m_i1} \\
C_{m_i} \wedge \alpha_{m_i2} \Rightarrow A_{m_i2} \\
\dots \\
\mathbf{true} \Rightarrow \mathbf{assume}(C_{m_i} \rightarrow \gamma_{m_i}) \\
\mathbf{true} \Rightarrow \mathbf{assume}(C_{m_i} \rightarrow \neg \bigvee_{C \in \{C_{m_j}, C_{r_j}\} \setminus C_{m_i}} C) \\
\mathbf{true} \Rightarrow \mathbf{assume}(C_{m_i} \leftrightarrow \widehat{p_{i1}}) \\
\dots \\
\mathbf{true} \Rightarrow \mathbf{assume}(\widehat{x_1} \wedge \widehat{x_2} \wedge \dots)
\end{cases}$$

Fig. 15 Translation from CAOS to AIF⁺

rule r is strengthened by the clock C_r that is associated with the rule. Thus, all actions of the rule are just executed when the clock ticks and the clock constraint $C_r \rightarrow \gamma_r$ ensures that the clock can only tick when the rule is enabled, i. e. γ_r holds. The reference semantics requires that at most one rule is executed at once. The second clock constraint for a rule ensures that its clock can only tick when no clock of another rule or method does. Methods are translated accordingly, but the input and output variables of a method have a different clock than all internal variables: they only change their value when the method is executed. This restriction is added by clock constraints for the clocks of the variables. In this way new parameters can only be given if and only if the method is executed. After translating the rules and methods, the clocks for the local variables need to be fixed. The clock constraint ensures that the clock of all local variables (identified with x_1, x_2, \dots) ticks at each instant. This is because the semantics of register assignments in CAOS require that the changes are visible right after the execution of the rule or method, thus for the next execution instant.

For the control-flow context we follow the translation of Signal. As we have an underlying data-flow model again, we add the same rudimentary as described in Section 3.2 and guard all clocked guarded actions similarly. This concludes the translation of CAOS to AIF⁺.

Now, we illustrate our approach by a CAOS example. The CAOS model is given in Figure 16, and the description derived from our translation is given in Figure 17. The system describes a token ring, where messages can be only exchanged between neighbors by a common single-place buffer. Communication is directed and its direction is static, i. e. each buffer is always the input of the following node and the output of the preceding node. In this example, we have three nodes connected to the ring, which all have the same behavior: if the output buffer is empty and the message in the input buffer is not for the node itself, the packet is forwarded. This part of the behavior is described by the rules `node1`, `node2`, and `node3`. Packets are inserted and removed from the ring by `send` and `receive` methods. For the sake of simplicity, this example contains these methods only for the first node. By firing a `send1`, a new packet is inserted into the ring, which can be only done if the current output buffer of the first node is empty. Packets can be received by a call to `receive1`, which requires that there is a packet waiting in the input buffer of the node. Obviously, all rules write to different buffers. The only resource conflict is between the forwarding of the first node and the introduction of a new packet, which both write to `31`. The translation is done according to the rules described above. The last constraint is interesting, which is due

```

module TokenRing {
  int buf_addr_12 = 0, buf_addr_23 = 0, buf_addr_31 = 0;
  int buf_data_12 = 0, buf_data_23 = 0, buf_data_31 = 0;

  rule node1
    when((buf_addr_31 != 1) & (buf_data_12 == 0)) {
      next(buf_addr_12) = buf_addr_31;
      next(buf_data_12) = buf_data_31;
      next(buf_addr_31) = 0;
      next(buf_data_31) = 0;
    }
  rule node2
    when((buf_addr_12 != 2) & (buf_data_23 == 0)) {
      next(buf_addr_23) = buf_addr_12;
      next(buf_data_23) = buf_data_12;
      next(buf_addr_12) = 0;
      next(buf_data_12) = 0;
    }
  rule node3
    when((buf_addr_23 != 3) & (buf_data_31 == 0)) {
      next(buf_addr_31) = buf_addr_23;
      next(buf_data_31) = buf_data_23;
      next(buf_addr_23) = 0;
      next(buf_data_23) = 0;
    }
  method send1(int ?a , int ?d)
    when(buf_data_31 == 0) {
      next(buf_addr_31) = a;
      next(buf_data_31) = d;
    }
  method receive3(int !d)
    when ((buf_data_23 != 0) & (buf_addr_23 == 3)) {
      next(buf_data_23) = 0;
      d = buf_data_23;
    }
}

```

Fig. 16 CAOS Example: TokenRing

to the CAOS semantics. It forbids that all four nodes fire in parallel, since this cannot be represented by any sequential firing.

3.4 Composition

As already highlighted in Section 2.2, AIF⁺ is not only an intermediate format for different languages but it also aims at composing modules obtained from different languages. To this end, there is an AIF⁺ linker, which can substitute module calls by the appropriate instance and the connects the signals of the control-flow and the data-flow interfaces.

We illustrate the linking with the help of our running example. While the previous sections showed the individual modules `OuterQuartz`, `InnerSignal` and `InnerQuartz`, this section describes their composition. To avoid confusion, the variable names in the running example have been chosen that the usual renaming step is not necessary.

The resulting AIF⁺ system after linking the modules is shown in Figure 18. The only control signal that remained in the behavior is `strt` (`OuterQuartz`), which is still needed

```

module TokenRing
  inputs:
    int n
  outputs:
    int n
  locals:
    int buf_addr_12, buf_addr_23, buf_addr_31
    int buf_data_12, buf_data_23, buf_data_31
    clock Cnode1, Cnode2, Cnode3
    clock Csend1, Creceive3
    bool ℓ

```

```

behavior:
  /* initialization */
    init(buf_addr_12) ⇒ buf_addr_12 = 0
    ⋮
    init(buf_data_31) ⇒ buf_data_31 = 0

  /* rule: node1 */
  (
    Cnode1 ∧
    (buf_addr_31 ≠ 1) ∧
    (buf_data_12.data == 0)
  ) ⇒ next(buf_addr_12) = buf_addr_31
  (
    Cnode1 ∧
    (buf_addr_31 ≠ 1) ∧
    (buf_data_12.data == 0)
  ) ⇒ next(buf_data_12) = buf_data_31
  (
    Cnode1 ∧
    (buf_addr_31 ≠ 1) ∧
    (buf_data_12.data == 0)
  ) ⇒ next(buf_addr_31) = 0
  (
    Cnode1 ∧
    (buf_addr_31 ≠ 1) ∧
    (buf_data_12.data == 0)
  ) ⇒ next(buf_data_31) = 0

  true ⇒ assume(Cnode1 → (
    (buf_addr_31 ≠ 1) ∧
    (buf_data_12.data == 0)
  ))
  true ⇒ assume(Cnode1 → ¬ (
    Cnode2 ∧
    Cnode3 ∧
    Csend1 ∧
    Creceive3
  ))

  /* rule: node2 */
  ...

  /* method: send1 */
  Csend1 ∧ (buf_data_31 == 0) ⇒ next(buf_addr_31) = a
  Csend1 ∧ (buf_data_31 == 0) ⇒ next(buf_data_31) = d
  true ⇒ assume(Csend1 → (buf_data_31 == 0))

  true ⇒ assume(Csend1 → ¬ (
    Cnode1 ∧
    Cnode2 ∧
    Cnode3 ∧
    Creceive3
  ))

  true ⇒ assume(Csend1 ↔  $\hat{a}$ )
  true ⇒ assume(Csend1 ↔  $\hat{d}$ )

  /* method: receive3 */
  ...

  /* overall constraint */
  true ⇒ assume(
    (
      buf_addr_12 ∧ buf_data_12 ∧
      buf_addr_23 ∧ buf_data_23 ∧
      buf_addr_31 ∧ buf_data_31
    )
  )

  strt(TokenRing) ∧
  ¬prmt(TokenRing) ⇒ next(ℓ) = true
  ℓ ∧
  ¬abrt(TokenRing) ⇒ next(ℓ) = true

```

```

control:
  inst: false
  insd: ℓ
  term: false

```

Fig. 17 AIF⁺ of Example: TokenRing

to start the module. All other control signals have been bound by the linker. The example also illustrates that many guards of the intermediate representation, which look complicated due to the control-flow interface signals, have become very simple after inserting the actual control-flow context in the course of linking.

The clock of the outermost module `OuterQuartz` is twice as often present as the clock of inner module `InnerQuartz` due to the sampling within the `Signal` part. Note that also the other direction would be possible if oversampling is used in the `Signal` process (like it is shown in the example in Figure 12).

4 Translating From Clocked Guarded Actions

From our intermediate representation of guarded actions, many synthesis targets can be thought of. In the following, we sketch the translation to two exemplary targets, a symbolic transition system, which is suitable for formal verification of program properties by symbolic model checking, and the translation to SystemC code, which can be used for an integrated simulation of the system. Similar to the previous section, we adapt previous work [12, 64] for synchronous languages and extend it by multiple clocks.

This section should serve two purposes: first, it illustrates the usage of clocked guarded actions in design flows and shows how modeled systems can be translated to formats processed by existing tools. Second, as the presented translations are very efficient, it also supports our argument that the representation of the behavior by clocked guarded actions is at an appropriate level of abstraction, providing a good balance between (1) removal of complexity from the source code level and (2) the independence of a specific synthesis target.

4.1 Symbolic Model Checking

For symbolic model checking, the system generally needs to be represented by a transition system. This basically consists of a triple $(\mathcal{S}, \mathcal{I}, \mathcal{T})$ with set of states \mathcal{S} , initial states $\mathcal{I} \subseteq \mathcal{S}$ and a transition relation $\mathcal{T} \subseteq \mathcal{S} \times \mathcal{S}$. Each state s is a mapping from variables to values, i. e. s assigns to each variable a value of its domain. As we aim for a symbolic description, we describe the initial states and the transition relation by propositional formulas $\Phi_{\mathcal{I}}$ and $\Phi_{\mathcal{T}}$, which are their characteristic functions.

For the presentation of the translation, assume that our intermediate representation contains immediate and delayed actions for each variable x of the following form

$$\begin{aligned} &(\gamma_1, x = \tau_1), \quad \dots, (\gamma_p, x = \tau_p) \\ &(\chi_{1, \text{next}}(x) = \pi_1), \dots, (\chi_{q, \text{next}}(x) = \pi_q) \end{aligned}$$

Figure 19 sketches the translation of the immediate and delayed actions writing variable x to clauses used for the description of a symbolic transition system.

As one might expect first, the construction of a transition system is not straightforward. Since delayed actions generally predetermine a new value for the next point of time \hat{x} is present while other actions still read its current value. To circumvent this problem, we introduce an auxiliary variable x' called the *the carrier of x* to capture delayed assignments at the previous point of time [64].

Before considering the constraints for x' , let us consider the invariant for x (Invar_x): clearly, we have to demand that x equals to τ_i whenever the guard γ_i of an immediate assignment $x = \tau_i$ holds. In case no guard of an immediate assignment to x holds, we have

system OuterQuartz		
inputs: bool susp int i	outputs: int o	locals: int x, y bool s, s' clock C1, C2 label ℓ, ℓ_1, ℓ_2
behavior:		
/* OuterQuartz */ true \Rightarrow assume($\widehat{i} = C1$) true \Rightarrow assume($\widehat{susp} = C1$) true \Rightarrow assume($\widehat{o} = C1$) /* Signal */ $\widehat{s} \wedge \text{str}(\text{OuterQuartz}) \Rightarrow s = \text{true}$ $\widehat{s} \wedge \ell \wedge \neg \text{susp} \Rightarrow s = s'$ $\widehat{s'} \wedge \left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell \wedge \neg \text{susp} \end{array} \right) \Rightarrow \text{next}(s') = s \odot \text{true}$ $\widehat{s} \wedge \left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell \wedge \neg \text{susp} \end{array} \right) \Rightarrow \text{assume}(\widehat{s} = \widehat{s'})$ $\widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell \wedge \neg \text{susp} \end{array} \right) \Rightarrow o = x$ $\widehat{i} \wedge \neg \widehat{x} \wedge \left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell \wedge \neg \text{susp} \end{array} \right) \Rightarrow y = i$ $\left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell \wedge \neg \text{susp} \end{array} \right) \Rightarrow \text{assume}(\widehat{i} = \widehat{s})$ $\left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell \wedge \neg \text{susp} \end{array} \right) \Rightarrow \text{assume}(\widehat{y} = s \wedge \widehat{s} \wedge \widehat{i})$ $\widehat{\ell} \wedge \text{str}(\text{OuterQuartz}) \Rightarrow \text{next}(\ell) = \text{true}$ $\widehat{\ell} \wedge \ell \Rightarrow \text{next}(\ell) = \text{true}$ $\ell \Rightarrow \text{assume}(\widehat{\ell} = \text{str}(\text{OuterQuartz}))$ /* InnerQuartz */ $C2 \wedge \left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell_2 \wedge \neg \text{susp} \end{array} \right) \Rightarrow x = y$ $C2 \wedge (\ell_1 \wedge \neg \text{susp}) \Rightarrow x = 2 * y$ $C2 \wedge \left(\begin{array}{l} \text{str}(\text{OuterQuartz}) \vee \\ \ell_2 \wedge \neg \text{susp} \end{array} \right) \Rightarrow \text{next}(\ell_1) = \text{true}$ $C2 \wedge (\ell_1 \wedge \neg \text{susp}) \Rightarrow \text{next}(\ell_2) = \text{true}$ true \Rightarrow assume($\widehat{x} = C2$) true \Rightarrow assume($\widehat{x} = C2$) true \Rightarrow assume($\widehat{\ell_1} = C2$) true \Rightarrow assume($\widehat{\ell_2} = C2$)		

Fig. 18 AIF⁺ of Example: OuterQuartz

to distinguish whether x is expected to tick or not: if this is the case, its default reaction determines the value, which is covered by the equations for the carrier variable x' . If x is not set by any action, it just keeps its old value so that other actions can still read it - which is covered by the clause Trans_x .

The meaning of x' is as follows: x' captures all of the delayed assignments $\text{next}(x) = \pi_j$ to x , that is whenever $\text{next}(x) = \pi_j$ is executed, we evaluate the right hand side π_j at the current point of time and assign this value to x' (not yet to x) at the next point of time.

$$\begin{aligned}
\text{Invar}_x &::= \left(\bigwedge_{j=1}^p (\gamma_j \rightarrow x = \tau_j \wedge \widehat{x}) \wedge \right. \\
&\quad \left. \left(\bigwedge_{j=1}^p \neg \gamma_j \right) \wedge \widehat{x} \rightarrow x = x' \right) \\
\text{Init}_{x'} &::= \text{Default}(x) \\
\text{Trans}_x &::= \neg \text{next}(\widehat{x}) \rightarrow \text{next}(x) = x \\
\text{Trans}_{x'} &::= \left(\bigwedge_{j=1}^q (\chi_j \rightarrow \text{next}(x') = \pi_j) \wedge \right. \\
&\quad \left(\bigwedge_{j=1}^q \neg \chi_j \right) \wedge \text{next}(\widehat{x}) \rightarrow \text{next}(x') = x \\
&\quad \left(\bigwedge_{j=1}^q \neg \chi_j \right) \wedge \neg \text{next}(\widehat{x}) \rightarrow \text{next}(x') = x' \right)
\end{aligned}$$

Fig. 19 Transition Relation for x

Hence, x' is determined by the delayed assignments to x . This leaves open what the initial value of x' should be, so we additionally define the initial value of x' as the default value of x .

By this definition of the initial value of x' , the initial value of x is correct. In later macro steps, if one of the immediate assignments to x is enabled, then this assignment determines the value of x at this point of time as given by the invariant equation for x . Otherwise, a delayed assignment $\text{next}(x) = \pi_j$ may have been executed at some previous point of time. If so, then x' has now the value that has been obtained by evaluating π_j at the previous point of time, and the invariant equation takes this value via x' .

For the additional assumptions, the translation is straightforward. Assume that the intermediate representation contains the following set of additional assumptions

$$(\delta_1, \text{assume}(\sigma_1)), \dots, (\delta_r, \text{assume}(\sigma_r))$$

They are translated to the following clause:

$$\text{Assume} ::= \bigwedge_{i=1, \dots, r} \delta_i \rightarrow \sigma$$

The final result is then the conjunction of the clauses of all writable variables \mathcal{V}_W together with the additional assumptions, i. e.

$$\begin{aligned}
\Phi_{\mathcal{I}} &= \text{Assume} \wedge \bigwedge_{\mathcal{V}_W} (\text{Init}_{x'} \wedge \text{Invar}_x) \\
\Phi_{\mathcal{T}} &= \text{Assume} \wedge \bigwedge_{\mathcal{V}_W} (\text{Trans}_x \wedge \text{Trans}_{x'} \wedge \text{Invar}_x)
\end{aligned}$$

This general format can be transformed by a straightforward syntactic translation to input which can be used by a model checker such as SMV. Figure 20 shows a part of the code

```

...
DEFINE _grd18 := C2 & (_strt_OuterQuartz | ell2 & !susp)
DEFINE _grd19 := C2 & (ell1 & !susp)
...
INVAR _grd18 -> (x = y) & _clk_x
INVAR _grd19 -> (x = 2*y) & _clk_x
INVAR !_grd18 & !_grd19 & _clk_x -> x = _carrier_x
TRANS next(_clk_x) -> next(_carrier_x) = x
TRANS !next(_clk_x) -> next(_carrier_x) = _carrier_x

INVAR _clk_ell1 -> ell1 = _carrier_ell1 & _clk_ell1
TRANS _grd18 -> next(_carrier_ell1) = TRUE
TRANS !_grd18 & next(_clk_x) -> next(_carrier_ell1) = FALSE
TRANS !_grd18 & !next(_clk_x) -> next(_carrier_ell1) = _carrier_ell1
...

SPEC AG AF (_clk_x)
SPEC AG ((x=y) | (x=2*y))
SPEC AG (_clk_susp & susp) -> !_clk_x
...

```

Fig. 20 Model Checking AIF⁺ with SMV

for our running example `OuterQuartz` (Figure 8), which describes the transitions for the variables x and ℓ_1 . In the SMV file, we first define all the guards in order to share them among many clauses in the rest of the description. The next part is a simple mapping of the clauses as described in Figure 19. Thereby, `_clk_x` represents \hat{x} and `_carrier_x` the variable x' . Finally, some specifications can be given, which are verified by SMV. In our example given in Figure 20, we check three properties: whether the clock of variable x is always live, whether x is always equal or the double of y , and whether a suspension really deactivates the emission of x .

From the theoretical point of view, the translation to transition systems does not have any limitations. All features of the AIF⁺ system are translated to a symbolic transition system (including assignments, assumptions and assertions). Practically, state-space explosion is always an issue, which might make model checking unusable for large systems.

4.2 SystemC Simulation

The simulation semantics of SystemC is based on the discrete-event model of computation [18], where reactions of the system are triggered by events. All threads that are sensitive to a specified set of events are activated and produce new events during their execution. Updates of variables are not immediately visible, but become visible in the next delta cycle.

We start the translation by the definition of a global clock that ticks in each instant and drives all the computation. Thus, we require that the processed model is *endochronous* [27, 28], i. e. there is a signal which is present in all instants of the behavior and from which all other signals can be determined. In SystemC, this clock is implemented by a single `sc_clock` at the uppermost level, and all other components are connected to this clock. Hence, the translations of the macro steps of the synchronous program in SystemC are triggered by this clock, while the micro steps are triggered by signal changes in the delta cycles. For this reason, input and output variables of the synchronous program are mapped to input signals (`sc_in`) and output signals (`sc_out`) of SystemC of the corresponding type.

```

void Module
  :: compute_x ()
{
  while(true) {
    if(_clk_x.read() &&  $\gamma_1$ )
      x.write( $\tau_1$ );
    ...
    else if(_clk_x.read() &&  $\gamma_n$ )
      x.write( $\tau_n$ );
    else if(_clk_x.read())
      x.write(_carrier_x.read());
    wait();
  }
}

void Module
  :: compute_delayed_x ()
{
  while(true) {
    if( $\xi_1$ )
      _carrier_x.write( $\pi_1$ );
    ...
    else if( $\xi_n$ )
      _carrier_x.write( $\pi_n$ );
    else
      _carrier_x.write(x.read());
    wait();
  }
}

```

Fig. 21 SystemC: Translation of Immediate and Delayed Actions

Additionally, we declare signals for all other clocks of the system. They are inputs since the clock constraints (as given by **assume**) do not give an operational description of the clocks, but can be only checked in the system. The clock calculus for Signal [2, 27, 28] or scheduler creation for CAOS [16] aim at creating exactly these schedulers which give an operational description of the clocks. Although not covered in the following, their result can be linked to the system description so that clocks are driven by the system itself.

The translation of the synchronous guarded actions to SystemC processes is however not as simple as one might expect. The basic idea is to map guarded actions to methods which are sensitive to the read variables so that the guarded action is re-evaluated each time one of the variables it depends on changes. For a *constructive* model it is guaranteed that the simulation does not hang up in delta-cycles.

The translation to SystemC must tackle the following two problems: (1) As SystemC does not allow a signal to have multiple drivers, all immediate and delayed actions must be grouped by their target variables. (2) The SystemC simulation semantics can lead to spurious updates of variables (in the AIF⁺ context), since threads are always triggered if some variables in the sensitivity list have been updated - even if they are changed once more in later delta cycles. As actions might be spuriously activated, it must be ensured that at least one action is activated in each instant, which sets the final value. Both problems are handled in a similar way as the translation to the transition system presented in the previous section: we create an additional variable `_carrier_x` for each variable x to record values from their delayed assignments, and group all actions in the same way as for the transition system.

With these considerations, the translation of the immediate guarded actions $\langle \gamma \Rightarrow x = \tau_i \rangle$ is straightforward: We translate each group of actions into an asynchronous thread in SystemC, which is sensitive to all signals read by these actions (variables appearing in the guards γ_i or in the right-hand sides τ_i). Thereby, all actions are implemented by an **if**-block except for the last one, which handles the case that no action fires. Since the immediate actions should become immediately visible, the new value can be immediately written to the variable with the help of a call to `x.write(...)`. Analogously, the evaluations of the guard γ_i and the right-hand side of the assignment τ_i make use of the `read` methods of the other signals. The left-hand side of Figure 21 shows the general structure of such a thread.

Delayed actions $\langle \gamma \Rightarrow \text{next}(x) = \pi_j \rangle$ are handled differently: While the right-hand side is immediately evaluated, the assignment should only be visible in the following macro

```

void OuterQuartz
::compute_x()
{
  while(true) {
    if(_grd18)
      x.write(y.read());
    else if(_grd19)
      x.write(2*y.read());
    else if(_clk_x.read())
      x.write(_carrier_x.read());
    wait();
  }
}

void OuterQuartz
::compute_delayed_x()
{
  while(true) {
    _carrier_x.write(x.read());
    wait();
  }
}

void OuterQuartz
::compute_delayed_e111()
{
  while(true) {
    if(_grd18)
      e111.write(true);
    else
      e111.write(false);
    wait();
  }
}

```

Fig. 22 SystemC Code for OuterQuartz

step and not yet in the current one. Hence, they do not take part in the fixpoint iteration. Therefore, we write their result to `_carrier_x` in a *clocked thread*, which is triggered by the master trigger. Thereby, signals changed by the delayed actions do not affect the current fixpoint iteration and vice versa.

Consider once again our running example given in Figure 8. The procedure describe above can be used to generate cycle-accurate SystemC code, which can be used for a detailed simulation of the given system. Each simulation step of the SystemC kernel corresponds to an instant of our system.

Figure 22 gives the SystemC code for the part that simulates the variables `x`. Similar to the SMV translation, we abbreviate guards for reuse in different SystemC processes. Then, the translation of the immediate actions writing `x` is straightforward; they correspond to the first two cases in method `OuterQuartz::compute_x()`. The last case is responsible for setting `x` to its previous value if neither of the two immediate actions fires. As already stated above, we need to do this explicitly. To this end, the previous value of variable `x` is always stored in `_carrier_x`. For variables, which are only set by delayed actions, we can simplify the general scheme of Figure 21. In this case, we can combine the two threads, as the clocked thread `compute_delayed_e111` in Figure 22 illustrates: we only need a single variable, which is set by this thread.

5 Case Study

In order to evaluate the modeling capabilities of our intermediate representation we used it for a case study taken from the automotive domain. The considered system describes a vehicle stability system, which controls the steering, traction and yaw rate of a car. Structurally, it consists of three components (see Figure 23): the first one (LogicalSensors) encapsulates the sensors and sensor fusion, the second one (Controllers) contains the actual control software, and the last one (LogicalActuators) is responsible for the actuators.

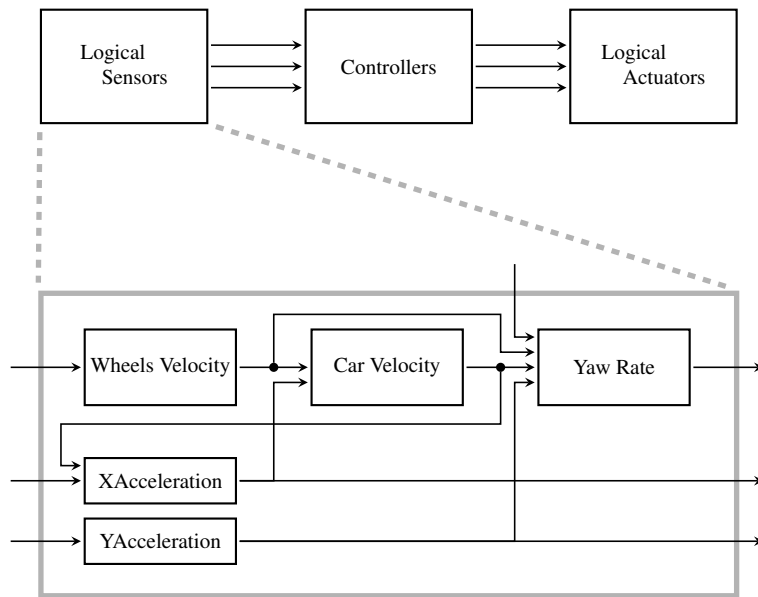


Fig. 23 Structure of VehicleStabilitySystem and LogicalSensors

In order to demonstrate the usability of our intermediate representation, we modeled the components of the vehicle stability system in different languages. The apparent data-flow layer is described by a Signal specification, which is a good choice: equations are good to describe pure data-flow, and the polychronous style does not constrain the rates of independent sensors, which allows us to instantiate concrete sensors at the end.

In contrast, the control-flow parts are very hard to encode in a data-flow language such as Signal. Therefore, we wrote Quartz programs for them. The rich set of precise control-flow statements makes it possible to write concise and readable models. In our example, the control-flow parts are on two different layers: below and on top of the data flow. Thus, we have a *hierarchy of different descriptions*: Quartz programs used by Signal data-flow specifications, which again are managed by mode automata described in Quartz.

To see the control-flow layer below of the data flow, consider the steering control. It is responsible to limit the steering movements at high speeds so that instable situations of the vehicle are avoided. It has three states: after the initialization, which sets up the parameters and checks all required parts, it can be activated and deactivated – depending on the speed of the vehicle. This part of the behavior is perfectly modeled by the following Quartz skeleton:

```

/* initialize device */
...
suspend {
  loop {
    /* control steering */
    ...
  }
} when(speed < SPEED_THRESHOLD);

```

The module has the input `steerIn` and the output `steerOut`, which represent the uncontrolled and controlled values of the steering angle. As they are only emitted in certain instants (which defines their clocks), the internal steps of the steering control are completely invisible to the outside. For example, during the initialization phase, no `steerOut` is computed, which hides this phase. To the rest of the system, the module is just a component that reads a steering input and writes a steering output each step.

To see the other control-flow layer on top of the data flow, consider the first component, which is responsible for sensor fusion. It must be very adaptive due to many reasons: e. g. a changing environment requires different strategies or transient hardware failures make sensor data unreliable. Therefore, on top of the data-flow, there is a control-flow part, which is responsible for switching between various modes of the system. For example, the component for the yaw rate determination is able to return data in many different ways: if the dedicated sensor fails, the lateral acceleration or the steering angle and the velocity of the rear wheels can be used to estimate the data. Similarly, the velocity can be determined from the acceleration and vice versa. As all possibilities cannot be combined arbitrarily (assume that speed and acceleration cannot be measured, then they cannot be determined from each other), the control is *on top of* the individual components.

In the actuator part of the system, we control several dual servo motors (our vehicle has an electric drive), which are given the number of revolutions per minute. If the control fails for some reason, we have to halt them so that the car stops safely. However, transient failures are quite common, we tolerate a short time (in which the old values are sustained) before we disrupt the drive. This adaptive behavior is modeled by the following Quartz fragment:

```
await(!error);
loop {
  abort {
    loop {
      /* normal mode */
      ...
    }
  } when(error);
  abort {
    /* sustain mode */
    loop {
      /* off mode */
      ...
    }
  } when(!error);
}
```

In this context, the control-flow interface (see end of Section 3.2) plays an important role. It is responsible that only the actions of the current mode are activated. All others are not active since the control-flow label representing the mode are inactive or a preemption have been invoked (indicated by `prmt(M)`).

Recall that preemption interrupts the execution of modules only at well-defined points. If a particular mode is aborted, it is made synchronous to the signals visible at the interface (and not between internal steps). In the vehicle stability system, this has the consequence that a mode change does not take place in the middle of a computation. For example, the steering controller will not be interrupted during initialization since these internal steps are invisible to the rest of the system.

This case study supports our claim that the proposed model is suitable to represent heterogeneous systems hierarchically composed of different descriptions. After modeling all parts and compiling them to our intermediate representation, we get the desired integrated model, which can be subsequently used in the design flow.

6 Related Work

Our integration significantly differs from previous approaches: Whereas the tagged signal model [3, 29, 53, 54] only defined a general framework for comparing different MoCs, frameworks such as Ptolemy [26], ForSyDe [45], HetSC [37], SystemC-H [56] or System-MoC [36] embed various different MoCs in a host language like Java, Haskell or SystemC. Most of the approaches support a composition of components based on different MoCs but typically only offer simulation of the combined systems.

For example, *Ptolemy* endows components which follow a particular MoC with a so-called *domain director*, which schedules the actors in the model and their actions according to the general rules of the MoC. When modules of different MoCs are composed, each one runs a director of its own, and the scheduling of actions by the actors of the submodules is orchestrated by another director which coordinates between the directors at the individual MoC level. The whole system model is therefore endowed with a hierarchy of directors, which has a certain overhead but allows an orthogonal use of MoCs. Determinism is often a burden of the writer of the model, as the director only schedules actors and their certain pre-designated functions (such as *pre-fire*, *fire*, *post-fire* etc.).

For an integrated simulation, all the approaches mentioned above rely on the operational semantics of their host languages. System models are executable by construction, which is also their main purpose. However, formal verification and hardware-software (co)synthesis may become harder since the internal representation primarily aims at an efficient or/and flexible simulation. For example, synthesis from a Ptolemy model must either synthesize the directors together with the models, or has to provide a run-time system for the synthesized code.

In addition to frameworks to integrate components with different MoCs, there are also related approaches for intermediate representations. In particular, the multi-clock declarative code DC+ [60] adds clock hierarchies to traditional DC [33], the privileged interchange format of synchronous languages for hardware circuit synthesis, symbolic verification and optimization. The format reflects declarative or data flow synchronous programs like Lustre, as well as the equational representations of imperative synchronous programs. The underlying idea is the definition of flows by equations governed by clock hierarchies. Clocked guarded actions as used in this article share the same general idea, but have a more general scope: we are not only interested in integrating synchronous models but also asynchronous ones, such as CAOS [16] or SHIM [24]. To this end, we have to generalize DC+ with respect to the definitions of flows and clocks (covered by the clocked guarded actions) and with respect to the semantic properties they need to fulfill.

Furthermore, AIF⁺ was designed to support a better compositionality. As already shown in Section 2.2, the interfaces of AIF⁺ make it possible to have a complete behavioral hierarchy (Figure 2 (b)), i. e. components can be called at arbitrary (control-flow) contexts in other components. In contrast, DC+ only supports the following two variants of composition.

- Basically, systems are hierarchically organized as nodes [60], which are components with a pure data interface. They can be arbitrarily nested in DC+, which corresponds to our structural hierarchy (Figure 2 (a)).

- Another way to integrate behavior from other components are functions and procedures (basically functions with several results). Functions can be used in expressions (since they have a single return value), and procedures as actions. Hence, both of them may be generally influenced by a surrounding control-flow context (e. g. action in the scope of an abort or suspend statement). However, DC+ requires that the interfaces of functions and procedures are monochronous [60], i. e. all the input and output parameters must have the same clock.

Thus, it is impossible to implement the proposed behavioral hierarchy (Figure 2 (b)) in DC+. For imperative programs, essential information such as the control-flow context, local variable context [14] is missing. Only with these features, it is possible to write a system, which consists of a Quartz components that calls a Signal component, which itself instantiates a Quartz program. For such a system, we expect that an abortion/suspension of the outer Esterel component is forwarded through the Signal specification, which itself also become inactive until the outer Esterel component resumes the execution again.

7 Summary

In this article, we proposed clocked guarded actions as a common intermediate representation of components that were written in languages based on different models of computation. Using this common intermediate representation we are able to create new components that consist of other components so that new ways to combine heterogeneous components are achieved. Moreover, the common intermediate representation is the basis for a seamless design flow that covers simulation, verification and other kinds of analyses, as well as synthesis to different targets. In addition to the definition of a precise intermediate language based on clocked guarded actions with its formal semantics, we also described the translation of components that were written in synchronous, polychronous or asynchronous languages to clocked guarded actions. Moreover, we showed how formal verification by means of symbolic model checking and a combined simulation based on a translation to SystemC can be easily achieved with our intermediate representation.

References

1. Arvind. Bluespec: A language for hardware design, simulation, synthesis and verification invited talk. In *Formal Methods and Models for Codesign (MEMOCODE)*, pages 249–254, Mont Saint-Michel, France, 2003. IEEE Computer Society.
2. A. Benveniste, P. Bournai, T. Gautier, and P. Le Guernic. SIGNAL: A data flow oriented language for signal processing. Research Report 378, INRIA, Rennes, France, March 1985.
3. A. Benveniste, B. Caillaud, L.P. Carloni, and A.L. Sangiovanni-Vincentelli. Tag machines. In W. Wolf, editor, *Embedded Software (EMSOFT)*, pages 255–263, Jersey City, New Jersey, USA, 2005. ACM.
4. A. Benveniste, P. Caspi, S. Edwards, N. Halbwegs, P. Le Guernic, and R. de Simone. The synchronous languages twelve years later. *Proceedings of the IEEE*, 91(1):64–83, 2003.
5. G. Berry. A hardware implementation of pure Esterel. In *Formal Methods in VLSI Design*, Miami, Florida, USA, 1991.

6. G. Berry. Synchronous languages for hardware and software reactive systems. In C. Delgado Kloos and E. Cerny, editors, *International Conference on Hardware Description Languages and their Applications (CHDL)*, page 3, Toledo, Spaine, 1997. Chapman and Hall.
7. G. Berry. The foundations of Esterel. In G. Plotkin, C. Stirling, and M. Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press, 1998.
8. G. Berry. The constructive semantics of pure Esterel, July 1999.
9. G. Berry and L. Cosserat. The Esterel synchronous programming language and its mathematical semantics. In S.D. Brookes, A.W. Roscoe, and G. Winskel, editors, *Seminar on Concurrency (CONCUR)*, volume 197 of *LNCS*, pages 389–448, Pittsburgh, Pennsylvania, USA, 1985. Springer.
10. M. Bombana and F. Bruschi. SystemC-VHDL co-simulation and synthesis in the HW domain. In *Design, Automation and Test in Europe (DATE)*, pages 20101–20105, Munich, Germany, 2003. IEEE Computer Society.
11. F. Boussinot and R. de Simone. The Esterel language. *Proceedings of the IEEE*, 79(9):1293–1304, 1991.
12. J. Brandt, M. Gemünde, and K. Schneider. From synchronous guarded actions to SystemC. In M. Dietrich, editor, *Methoden und Beschreibungssprachen zur Modellierung und Verifikation von Schaltungen und Systemen (MBMV)*, pages 187–196, Dresden, Germany, 2010. Fraunhofer Verlag.
13. J. Brandt, M. Gemünde, K. Schneider, S. Shukla, and J.-P. Talpin. Integrating system descriptions by clocked guarded actions. In A. Morawiec, J. Hinderscheit, and O. Ghenassia, editors, *Forum on Specification and Design Languages (FDL)*, pages 1–8, Oldenburg, Germany, 2011. IEEE Computer Society.
14. J. Brandt and K. Schneider. Separate translation of synchronous programs to guarded actions. Internal Report 382/11, Department of Computer Science, University of Kaiserslautern, Kaiserslautern, Germany, March 2011.
15. J. Brandt, K. Schneider, and S.A. Edwards. Translating SHIM to guarded actions. Technical Report 387/12, University of Kaiserslautern, Kaiserslautern, Germany, February 2012.
16. J. Brandt, K. Schneider, and S.K. Shukla. Translating concurrent action oriented specifications to synchronous guarded actions. In J. Lee and B.R. Childers, editors, *Languages, Compilers, and Tools for Embedded Systems (LCTES)*, pages 47–56, Stockholm, Sweden, 2010. ACM.
17. P. Caspi, N. Halbwachs, D. Pilaud, and J.A. Plaice. LUSTRE: A declarative language for programming synchronous systems. In *Principles of Programming Languages (POPL)*, pages 178–188, Munich, Germany, 1987. ACM.
18. C.G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Springer, 2 edition, 2008.
19. K.M. Chandy and J. Misra. *Parallel Program Design*. Addison-Wesley, Austin, Texas, USA, May 1989.
20. A. Chatelain, Y. Mathys, G. Placido, A. La Rosa, and L. Lavagno. High-level architectural co-simulation using Esterel and C. In *Hardware-Software Codesign (CODES)*, Copenhagen, Denmark, 2001.
21. N. Dershowitz, M. Okada, and G. Sivakumar. Canonical conditional rewrite systems. In E.L. Lusk and R.A. Overbeek, editors, *Conference on Automated Deduction (CADE)*, volume 310 of *LNCS*, pages 538–549, Argonne, Illinois, USA, 1988. Springer.
22. E.W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM (CACM)*, 18(8):453–457, 1975.

23. D.L. Dill. The Murphi verification system. In R. Alur and T.A. Henzinger, editors, *Computer Aided Verification (CAV)*, volume 1102 of *LNCS*, pages 390–393, New Brunswick, New Jersey, USA, 1996. Springer.
24. S.A. Edwards and O. Tardieu. SHIM: a deterministic model for heterogeneous embedded systems. In W. Wolf, editor, *Embedded Software (EMSOFT)*, pages 264–272, Jersey City, New Jersey, USA, 2005. ACM.
25. J. Eker and J.W. Janneck. CAL actor language – language report. EECS Department, University of California at Berkeley, May 2002.
26. J. Eker, J.W. Janneck, E.A. 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.
27. A. Gamatié. *Designing Embedded Systems with the SIGNAL Programming Language*. Springer, 2010.
28. A. Gamatié, T. Gautier, P. Le Guernic, and J.P. Talpin. Polychronous design of embedded real-time applications. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 16(2), April 2007.
29. A. Girault, B. Lee, and E.A. Lee. Hierarchical finite state machines with multiple concurrency models. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T-CAD)*, 18(6):742–760, June 1999.
30. B. Gramlich and C. Wirth. Confluence of terminating conditional rewrite systems revisited. In H. Ganzinger, editor, *Rewriting Techniques and Applications (RTA)*, volume 1103 of *LNCS*, pages 245–259, New Brunswick, New Jersey, USA, 1996. Springer.
31. G. Gössler and A. Sangiovanni-Vincentelli. Compositional modeling in Metropolis. In A.L. Sangiovanni-Vincentelli and J. Sifakis, editors, *Embedded Software (EMSOFT)*, volume 2491 of *LNCS*, pages 93–107, Grenoble, France, 2002. Springer.
32. N. Halbwachs. *Synchronous programming of reactive systems*. Kluwer, 1993.
33. N. Halbwachs. *The Declarative Code DC, version 1.2a*, April 1998.
34. N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud. The synchronous dataflow programming language LUSTRE. *Proceedings of the IEEE*, 79(9):1305–1320, September 1991.
35. D. Harel and A. Naamad. The STATEMATE semantics of Statecharts. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 5(4):293–333, 1996.
36. C. Haubelt, J. Falk, J. Keinert, T. Schlichter, M. Streubühr, A. Deyhle, A. Hadert, and J. Teich. A SystemC-based design methodology for digital signal processing systems. *EURASIP Journal on Embedded Systems*, 2007(1):15–15, January 2007.
37. F. Herrera and E. Villar. A framework for heterogeneous specification and design of electronic embedded systems in SystemC. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 12(3):1–31, August 2007.
38. C.A.R. Hoare. Communicating sequential processes. *Communications of the ACM (CACM)*, 21(8):666–677, 1978.
39. C.A.R. Hoare. Communicating sequential processes. *Communications of the ACM (CACM)*, 26(1):100–106, January 1983.
40. J.C. Hoe and Arvind. Hardware synthesis from term rewriting systems. Technical Report CSG-MEMO 421-a, Computer Science and Artificial Intelligence Laboratory, Cambridge, Massachusetts, USA, 1999.
41. J.C. Hoe and Arvind. Operation-centric hardware description and synthesis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T-CAD)*, 23(9):1277–1288, September 2004.

42. IEEE. *IEEE Standard Hardware Description Language Based on the Verilog Hardware Description Language*, 2005. IEEE Std. 1394-2005.
43. IEEE. *IEEE Standard SystemC Language Reference Manual*. New York, New York, USA, December 2005. IEEE Std. 1666-2005.
44. IEEE. *IEEE Standard VHDL Language Reference Manual*, 2008. IEEE Std. 1076-2008.
45. A. Jantsch. *Modeling Embedded Systems and SoCs*. Morgan Kaufmann, 2004.
46. A. Jantsch. Models of computation for networks on chip. In *Application of Concurrency to System Design (ACSD)*, pages 165–178, Turku, Finland, 2006. IEEE Computer Society.
47. S. Kaplan. Conditional rewrite rules. *Theoretical Computer Science (TCS)*, 33:175–193, 1984.
48. T.J. Koo, J. Liebman, C. Ma, B. Horowitz, A. Sangiovanni-Vincentelli, and S. Sastry. Platform-based embedded software design for multi-vehicle multi-modal systems. In A.L. Sangiovanni-Vincentelli and J. Sifakis, editors, *Embedded Software (EMSOFT)*, volume 2491 of *LNCS*, pages 32–45, Grenoble, France, 2002. Springer.
49. P. Le Guernic, T. Gauthier, M. Le Borgne, and C. Le Maire. Programming real-time applications with SIGNAL. *Proceedings of the IEEE*, 79(9):1321–1336, 1991.
50. E.A. Lee and D.G. Messerschmitt. Static scheduling of synchronous data flow programs for digital signal processing. *IEEE Transactions on Computers (T-C)*, 36(1):24–35, January 1987.
51. E.A. Lee and D.G. Messerschmitt. Synchronous data flow. *Proceedings of the IEEE*, 75(9):1235–1245, September 1987.
52. E.A. Lee and T. Parks. Dataflow process networks. *Proceedings of the IEEE*, 83(5):773–801, May 1995.
53. E.A. Lee and A. Sangiovanni-Vincentelli. Comparing models of computation. In *International Conference on Computer-Aided Design (ICCAD)*, pages 234–241, San Jose, California, USA, 1996. ACM/IEEE Computer Society.
54. E.A. Lee and A. Sangiovanni-Vincentelli. A framework for comparing models of computation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T-CAD)*, 17(12):1217–1229, December 1998.
55. C. Passerone, L. Lavagno, M. Chiodo, and A.L. Sangiovanni-Vincentelli. Fast Hardware/Software co-simulation for virtual prototyping and trade-off analysis. In *Design Automation Conference (DAC)*, pages 389–394, Anaheim, California, USA, 1997. ACM.
56. H.D. Patel, S.K. Shukla, and R.A. Bergamaschi. Heterogeneous behavioral hierarchy extensions for SystemC. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T-CAD)*, 26(4):765–780, April 2007.
57. A. Poigné and L. Holenderski. Boolean automata for implementing pure Esterel. Arbeitspapiere 964, GMD, Sankt Augustin, Germany, 1995.
58. F. Rocheteau and N. Halbwachs. Implementing reactive programs on circuits: A hardware implementation of LUSTRE. In J.W. de Bakker, C. Huizing, W.P. de Roever, and G. Rozenberg, editors, *Real-Time: Theory in Practice*, volume 600 of *LNCS*, pages 195–208, Mook, The Netherlands, 1992. Springer.
59. J.A. Rowson. Hardware/Software co-simulation. In *Design Automation Conference (DAC)*, pages 439–440, San Diego, California, USA, 1994. ACM.
60. ESPRIT Project: Safety Critical Embedded Systems (SACRES). *The Declarative Code DC+, version 1.4*, November 1997.
61. I. Sander, A. Jantsch, and Z. Lu. Development and application of design transformations in ForSyDe. In *Design, Automation and Test in Europe (DATE)*, pages 10364–10369,

- Munich, Germany, 2003. IEEE Computer Society.
62. A.L. Sangiovanni-Vincentelli, L.P. Carloni, F. De Bernardinis, and M. Sgroi. Benefits and challenges for platform-based design. In S. Malik, L. Fix, and A.B. Kahng, editors, *Design Automation Conference (DAC)*, pages 409–414, San Diego, California, USA, 2004. ACM.
 63. K. Schneider. Embedding imperative synchronous languages in interactive theorem provers. In *Application of Concurrency to System Design (ACSD)*, pages 143–154, Newcastle Upon Tyne, England, UK, 2001. IEEE Computer Society.
 64. K. Schneider. The synchronous programming language Quartz. Internal Report 375, Department of Computer Science, University of Kaiserslautern, Kaiserslautern, Germany, December 2009.
 65. K. Schneider, J. Brandt, and T. Schuele. A verified compiler for synchronous programs with local declarations. *Electronic Notes in Theoretical Computer Science (ENTCS)*, 153(4):71–97, 2006.
 66. G. Singh and S.K. Shukla. Algorithms for low power hardware synthesis from concurrent action oriented specifications CAOS. *International Journal of Embedded Systems (IJES)*, 3(1/2):83–92, 2007.
 67. A.R.W. Todesco and T.H.Y. Meng. Symphony: A simulation backplane for parallel mixed-mode co-simulation of VLSI systems. In *Design Automation Conference (DAC)*, pages 149–154, Las Vegas, Nevada, USA, 1996. ACM.
 68. C. Zebelein, J. Falk, C. Haubelt, and J. Teich. Classification of general data flow actors into known models of computation. In *Formal Methods and Models for Codesign (MEMOCODE)*, pages 119–128, Anaheim, California, USA, 2008. IEEE Computer Society.
 69. V. Zivojnovic and H. Meyr. Compiled HW/SW co-simulation. In *Design Automation Conference (DAC)*, pages 690–695, Las Vegas, Nevada, USA, 1996. ACM.