A compositional behavioral modeling framework for embedded system design and conformance checking

Jean-Pierre Talpin, Paul Le Guernic, Sandeep Kumar Shukla, Rajesh K. Gupta

► To cite this version:

Jean-Pierre Talpin, Paul Le Guernic, Sandeep Kumar Shukla, Rajesh K. Gupta. A compositional behavioral modeling framework for embedded system design and conformance checking. International Journal of Parallel Programming, 2005, 33 (6), pp.613-643. 10.1007/s10766-005-8907-y . hal-00541986

HAL Id: hal-00541986

<https://hal.science/hal-00541986>

Submitted on 1 Dec 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A compositional behavioral modeling framework for embedded system design and conformance checking

Jean-Pierre Talpin¹, Paul Le Guernic¹, Sandeep Kumar Shukla² and Rajesh Gupta³

¹ INRIA-IRISA, Rennes, France. E-mail: talpin@irisa.fr

² Virginia Tech, Blacksburg, USA. E-mail: shukla@vt.edu

³ University of California San Diego, La Jolla, USA. E-mail: rgupta@ucsd.edu

We propose a framework based on a synchronous multi-clocked model of computation to support the inductive and compositional construction of scalable behavioral models of embedded systems engineered with *de facto* standard design and programming languages. Behavioral modeling is seen under the paradigm of type inference. The aim of the proposed type system is to capture the behavior of a system under design and to re-factor it by performing global optimizing and architecture-sensitive transformations on it. It allows to modularly express a wide spectrum of static and dynamic behavioral properties and automatically or manually scale the desired degree of abstraction of these properties for efficient verification. The type system is presented using a generic and language-independent static single assignment intermediate representation.

KEY WORDS: Embedded system design, formal methods, models of computation, program transformation, verification.

1. INTRODUCTION

The popular slogan “*write once, run anywhere*” effectively renders the expressive capabilities of general purpose programming languages for developing, deploying, and reusing target-independent applications. Generality and simplicity has driven most attention of the compiler technology community to developing local and compositional compiler optimization techniques. When it comes to the implementation of embedded software, this

approach is however far from satisfactory, especially in hard real-time system design (e.g. airborne systems, digital circuits) where conformance to real-time specifications is critical.

Domain-specific models and languages, such as these proposed under the synchronous programming paradigm, provides the necessary formal engineering models and design methodologies to allow for a program written once to be mapped on any distributed execution architecture by using global transformation and optimization techniques. Our aim is to relate this domain-specific model to embedded software development using general-purpose environments. To this end, we set the methodological framework of our synchronous model of computation within the general and reusable concept of a type system targeting the generic programming language setting of GCC's intermediate representations (three-address code and static single assignment). We give formal semantics to both our type system and the functional subset of SSA under consideration, define a type inference system and prove its correctness, before to depict the applications of our technique as developed in our project and presented in previous works.

A functional application domain. We consider embedded software implemented by resource-constrained⁴ multi-threaded programs on a specific runtime sub-system (e.g. , the real-time JVM, an RTOS, or simply hardware) which we call its execution architecture. Our technique consists of a type inference system that relates threads (imperative programs in intermediate form) to propositions expressed by synchronous transition systems that describe their behaviour.

Example. On the right, we outline the extent of our technique by depicting a test-case studied in ⁽¹⁴⁾. We consider modeling a real-time Java program consisting of three threads (right), a scheduler (top-left) and shared resources control (bottom-left).

This decomposition is obtained by partitioning the executable program and its environment into:

- *the execution architecture*: a hardware platform, a middle-ware library, a real-time operating system, a virtual machine (e.g. in Java), a simula-

⁴ It is common sense to restrict ourselves to programs where all objects are first created and initialized to elaborate the application architecture. Then, threads implement reactions to inputs in the nominal phase of execution and do not allocate any new object (to comply with certification requirement in software design or simply with common sense in SoC design).

tion kernel (e.g. in SystemC). The execution architecture describes an API of generic process and communication management services.

- *the application architecture*: a program, starting from the `main()` procedure, which initializes and links objects to form a hierarchical structure of shared data and communicating threads. The application mapping constructively describes the architecture of the system.
- *the application functionalities*: a set of program threads which periodically or sporadically react to inputs from the environment by interacting with each other for the access to shared data.

Our methodology consists of considering the three elements of an embedded system (its execution and application architectures, its application functionalities) in specific ways.

- *modeling*: the execution architecture, viewed through an application programming interface (API) of generic services, is modeled by template propositions. For instance, the procedure for thread creation in an RTOS API corresponds to a template proposition in the RTOS model whose parameters are the number of threads supported by the application scheduler, the period and deadline of the thread (for a real-time thread), etc.
- *analysis*: the application architecture, viewed as a hierarchical structure, is interpreted to elaborate a model by the instantiation of generic API services to the parameters and initial values provided in the program (e.g. thread parameters).
- *translation*: each thread consists of a sequential program that describes a functionality to be periodically or sporadically executed by the scheduler and corresponds to a particular model.

This allows for a complete separation of the virtual (threading or functional) architecture of an application from its actual, real-time and resource-constrained implementation: it provides an implementation of the “*write once run anywhere*” slogan in embedded system design.

Context. Our methodology arises from previous work on real-time operating systems modeling, embedded systems modeling and verification in the Polychrony workbench⁵, a tool-set for embedded system design based on a multi-clocked synchronous model of computation and implemented by the

⁵ URL: <http://www.irisa.fr/espresso/Polychrony>

data-flow notation Signal ⁽³⁾. In ⁽⁷⁾, the authors describe the implementation of a real-time operating system standard for avionics application: ARINC ⁽⁷⁾. The commercial implementation of this library, RT-Builder from TNI-Valiosys, is used for industrial-scale embedded software engineering project in avionics.

In ⁽¹⁴⁾, this model is used to describe key services of the real-time Java virtual machine. It is applied to rethreading multi-threaded real-time Java programs by global optimization. In ⁽¹⁵⁾, the application of our methodology to system-level design is further developed by studying its application to checking behavioral conformance between embedded systems described in SpecC and at heterogeneous levels of abstraction. In ⁽¹⁶⁾, a generic translation scheme of SystemC programs to the Polychrony workbench is described by considering a static single assignement intermediate representation due to the GCC project ⁽¹¹⁾. It is applied to design checking (e.g. race and lock detection). In ⁽⁵⁾, it is applied to modular verification by model checking and component-wise model abstraction.

We set our methodological framework within the general paradigm of a behavioral type system that associates meaning to software functionalities. The type system is cast in the generic programming language-oriented context of the three-address code (TAC) and static single assignment (SSA) intermediate representations (IR) of GCC.

2. RATIONALE

To allow for an easy grasp on the type system proposed for modeling behaviors, we outline the analysis of an imperative program, Figure 1, and depict the construction of its type, Figure 2. Figure 1 depicts a simple C code fragment consisting of an iterative program that counts the number of bits set to one in the variable `idata`. While `idata` is not equal to zero, it adds its right-most bit to an output count variable `ocount` and shifts it right in order to process the next bit. In the intermediate representation (IR) of the program (Figure 1, second column), all variables (`idata` and `ocount`) are read and written once per cycle.

<pre>while (idata != 0) { ocount = ocount + (idata & 1); idata = idata >> 1; }</pre>	<pre>L2:T1 = idata; T0 = T1 == 0; if T0 then goto L3; T2 = ocount;</pre>	<pre>T3 = T1 & 1; ocount = T2 + T3; idata = T1 >> 1; goto L2;</pre>
--	--	---

Fig. 1. From a C-like program to its intermediate representation.

This IR can equally be one of the TAC and SSA formats of GCC. Label L2 is the entry point of the block associated with the while loop. The first

instruction loads the input variable `idata` into the register `T1`. The second instruction stores the result of its comparison with 0 in the register `T0`. If `T0` is false, control is passed to block `L3`. Otherwise, the next instruction is executed: the variable `ocount` is loaded into `T2`, the last bit of `T1` is loaded into `T3`, the sum of `T2` and `T3` assigned to `ocount` and the right-shift of `T1` assigned to `idata`. The block terminates with an unconditional branch back to label `L2`.

A behavioral type system. The meaning of this C program fragment is given in a minimalist formalism akin to Pnueli's synchronous transition systems⁽¹²⁾. It not only describes a behavior of the program suitable for its formal verification but also allows for global model transformations to be performed on it. Let us zoom on the block `L2` in the example of Figure 2. The behavioral type of the block `L2`, middle, consists of the simultaneous composition of logical propositions that form a synchronous transition system. Each proposition is associated with one instruction: it specifies its *invariants*: it tells when the instruction is executed, what it computes, when it passes control to the next statement, when it branches to another block.

<code>L2:T1 = idata;</code>	<code>L2⇒T1 :=idata</code>
<code>T0 = T1 == 0;</code>	<code>T0 := (T1 = 0)</code>
<code>if T0 then goto L3;</code>	<code>T0 ⇒L3'</code>
<code>T2 = ocount;</code>	<code>¬T0⇒T2 := ocount</code>
<code>T3 = T1 & 1;</code>	<code>T3 := T1&1</code>
<code>ocount = T2 + T3;</code>	<code>ocount' := T2 + T3</code>
<code>idata = T1 >> 1;</code>	<code>idata' := T1 >> 1</code>
<code>goto L2;</code>	<code>L2'</code>

Fig. 2. From a generic intermediate representation to propositions.

On line 1 for instance, we associate the instruction `T1 := idata` to the proposition `L2 ⇒ T1 := idata`. The variable `L2` is a boolean that is true iff the block of label `L2` is being executed. Hence, the proposition says that, if the label `L2` is being executed, then `T1` is equal to `idata`. All propositions are conditioned by `L2` to mean that they hold when block `L2` is executed. The extent of a proposition is the duration of a reaction.

A reaction can be an arbitrarily long yet finite period of time provided that every variable or register changes its value at most once during that period. For instance, consider the instruction `if T0 then L3`. It is likely that label `L3` will, just as `L2`, perform some operation on the input `idata`. Therefore, its execution is delayed until after the current reaction. We refer to `L3'` as the next value of the state variable `L3`, to indicate that it will be active

during the next reaction. Hence, the proposition $L2 \Rightarrow T0 \Rightarrow L3'$ says that control will be passed to L3 at the next reaction when control is presently at L2 and when T0 is true. The instructions that follow this test are conditioned by the negative $\neg T0$, this means: "in the block L2 and not in its branch to L3".

3. A BEHAVIORAL TYPE SYSTEM

The central element of the type system is a process. It consists of simultaneous propositions that manipulate signals. A signal is an infinite flow of values that is sampled by a discrete series of reactions. This series is called a clock. An event corresponds to the value carried by a signal during a reaction. The formal syntax of propositions in the behavioral type system is defined by the inductive grammar P . A proposition or process P manipulates boolean values noted $v \in \{\text{false}, \text{true}\}$ and signals noted x, y, z . A location l refers to the initial value x^0 , the present value x and the next value x' of a signal. A reference r is either a value v or a signal x .

$$\text{(reference)} \quad r ::= x | v \quad \text{(location)} \quad l ::= x^0 | x | x'$$

A clock expression e is a proposition on boolean values that, when true, defines a particular period in time. The clocks 0 and 1 denote events that never/always happen. The clock $x = r$ denotes the proposition: "x is present and holds the value r". Particular instances are: the clock $\hat{x} \stackrel{\text{def}}{=} (x = x)$, which stands for "x is present"; the clock $x \stackrel{\text{def}}{=} (x = \text{true})$ for "x is true", and the clock $\neg x \stackrel{\text{def}}{=} (x = \text{false})$ for "x is false". Clocks are propositions combined using the logical combinators of conjunction $e \wedge f$, to mean that both e and f hold, disjunction $e \vee f$, to mean that either e or f holds, and symmetric difference $e \setminus f$, to mean that e holds and not f .

$$\text{(clock)} \quad e, f ::= 0 | x = r | e \wedge f | e \vee f | e \setminus f | 1$$

A process P consists of the simultaneous composition of elementary propositions. 1 is the process that does nothing. The proposition $l = r$ means that "l holds the value r". The process $e \Rightarrow P$ is a guarded command. It means: "if e is present then P holds". Processes are combined using synchronous composition $P | Q$ to denote the simultaneity of the propositions P and Q . Restricting a signal name x to the lexical scope of a process P is written P/x .

$$\text{(process)} \quad P, Q ::= 1 | l = r | x \rightarrow l | e \Rightarrow P | (P | Q) | P/x$$

An order of execution is imposed to a proposition by a scheduling constraint, noted $x \rightarrow l$, to mean that "l cannot happen before x". Consequently, a proposition, e.g. $x = y$, is seen as the abstraction of an assignment, written $x := y$, defined by $x = y | y \rightarrow x$.

3.1. A synchronous model of computation

The meaning of our notation is given in the synchronous model of computation of ⁽⁸⁾. We consider a partially-ordered set $(\mathcal{T}, \leq, 0)$ of tags. A tag $t \in \mathcal{T}$ denotes a symbolic period in time. The relation $\leq$ denotes a partial order and its minimum is noted 0. We note $C \in \mathcal{C}$ a *chain* of tags (a totally ordered subset of $\mathcal{T}$). We define an *event* $e \in \mathcal{T} \times \mathcal{V}$ by the pair of a value and a tag, a *signal* $s \in \mathcal{S} = \{C \rightarrow \mathcal{V} \mid C \in \mathcal{C}\}$ by a function from a *chain* of tags C to values, a *behavior* $b \in \mathcal{B} = \mathcal{X} \rightarrow \mathcal{S}$ by a finite map from signal names $\mathcal{X}$ to signals $\mathcal{S}$, a *process* $p \in \mathcal{P}$ by a set of behaviors of same domain. We write $\text{tags}(s)$ for the tags of a signal s , $b|_X$ for the projection of a behavior b on $X \subset \mathcal{X}$ and $b/X = b|_{\text{vars}(b) \setminus X}$ for its complementary, $\text{vars}(b)$ and $\text{vars}(p)$ for the domains of b and p .

Example 1. Figure 3 depicts a behavior b over three signals named x , y and z . Two frames depict timing domains formalized by chains of tags. Signal x and y belong to the same timing domain: x is a down-sampling of y . Its events are synchronous to odd occurrences of events along y and share the same tags, e.g. t_1 . Even tags of y , e.g. t_2 , are ordered along its chain, e.g. $t_1 < t_2$, but absent from x (we write $t < t'$ if $t \leq t'$ and $t' \not\leq t$). Signal z belongs to a different timing domain. Its tags, e.g. t_3 are not ordered with respect to the chain of y , e.g. $t_1 \not\leq t_3$ and $t_3 \not\leq t_1$.

Fig. 3. behavior b over three signals x , y and z in two clock domains.

Scheduling structure. To schedule the occurrence of events during a period or an instant t , we consider the fact that the pair x_t of a time tag t and of a signal name x renders its very date d . The tag t represents the period during which the event takes place and the signal x its location. This consideration defines scheduling $\rightarrow$ by a pre-order relation between dates. Figure 4 depicts such a relation superimposed to the signals x and y of Figure 3. The relation $y_{t_1} \rightarrow x_{t_1}$, for instance, requires y to be calculated before x at the period t_1 . A scheduling relation naturally satisfies containment with respect to the timing partial order $\leq$ of every signal x in a behavior b , in

that for all $t, t' \in \text{tags}(b(x))$, $t < t'$ naturally implies $x_t \rightarrow^b x_{t'}$ and, conversely $x_t \rightarrow^b x_{t'}$ implies $t' \not< t$. A scheduling relation is implicitly transitive ($x_t \rightarrow^b y_{t'} \rightarrow^b z_{t''}$ implies $x_t \rightarrow^b z_{t''}$) and its closure for restriction b/X is defined by $x_t \rightarrow^{b/X} y_{t'}$ iff $x_t \rightarrow^b y_{t'}$ and $x, y \notin X$.

Fig. 4. Scheduling relations between simultaneous events.

Synchronous composition is noted $p \parallel q$ and defined by the union of all behaviors b (from p) and c (from q) which are synchronous. All signals x shared by b and c belong to $I = \text{vars}(p) \cap \text{vars}(q)$ and are equal i.e. $b|_I = c|_I$: $p \parallel q = \{b \cup c \mid (b, c) \in p \times q, I = \text{vars}(p) \cap \text{vars}(q), b|_I = c|_I\}$.

Fig. 5. Synchronous composition of $b \in p$ and $c \in q$.

3.2. Meaning of clocks.

The denotation $\llbracket e \rrbracket_b$ of a clock expression e (table 6) is defined relatively to a given behavior b and consists of the set of tags satisfied by the proposition e in the behavior b .

$$\begin{aligned}
 \llbracket 0 \rrbracket_b &= \emptyset & \llbracket 1 \rrbracket_b &= \text{tags}(b) \\
 \llbracket e \wedge f \rrbracket_b &= \llbracket e \rrbracket_b \cap \llbracket f \rrbracket_b \\
 \llbracket e \vee f \rrbracket_b &= \llbracket e \rrbracket_b \cup \llbracket f \rrbracket_b \\
 \llbracket e \setminus f \rrbracket_b &= \llbracket e \rrbracket_b \setminus \llbracket f \rrbracket_b \\
 \llbracket x = v \rrbracket_b &= \{t \in \text{tags}(b(x)) \mid b(x)(t) = v\} \\
 \llbracket x = y \rrbracket_b &= \{t \in \text{tags}(b(x)) \cap \text{tags}(b(y)) \mid b(x)(t) = b(y)(t)\}
 \end{aligned}$$

Fig. 6. Denotational semantics of clocks.

In Figure 6, the meaning of the clock $x = v$ (resp. $x = y$) in b is the set of tags $t \in \text{tags}(b(x))$ (resp. $t \in \text{tags}(b(x)) \cap \text{tags}(b(y))$) such that $b(x)(t) = v$ (resp. $b(x)(t) = b(y)(t)$). In particular, $\llbracket \hat{x} \rrbracket_b = \text{tags}(b(x))$. The meaning of a conjunction $e \wedge f$ (resp. disjunction $e \vee f$ and difference $e \setminus f$) is the intersection (resp. union and difference) of the meaning of e and f . Clock 0 has no tags.

3.3. Meaning of propositions.

The meaning $\llbracket P \rrbracket^e$ of a proposition P is defined with respect to a clock expression e . Where this information is absent, we assume $\llbracket P \rrbracket = \llbracket P \rrbracket^1$ to mean that P is an invariant (and is hence independent of a particular clock). The meaning of an initialization $\llbracket x^0 = v \rrbracket^e$ consists of all behaviors defined on x , written $b \in \mathcal{B}|_x$ such that the initial value of the signal $b(x)$ equals v . Notice that it is independent from the clock expression e provided by the context. We write $\mathcal{B}|_X$ for the set of all behaviors of domain X , $\min(C)$ for the minimum of the chain of tags C , $\text{succ}_t(C)$ for the immediate successor of t in the chain C , $\text{vars}(P)$ and $\text{vars}(e)$ for the sets of signal names of P and e .

$$\begin{aligned}
\llbracket x = y \rrbracket^e &= \{b \in \mathcal{B}|_{\text{vars}(e) \cup \{x, y\}} \mid \forall t \in \llbracket e \rrbracket_b, \\
&\quad t \in \text{tags}(b(x)) \wedge t \in \text{tags}(b(y)) \wedge b(x)(t) = b(y)(t)\} \\
\llbracket y \rightarrow x \rrbracket^e &= \{b \in \mathcal{B}|_{\text{vars}(e) \cup \{x, y\}} \mid \forall t \in \llbracket e \rrbracket_b, \\
&\quad t \in \text{tags}(b(x)) \Rightarrow t \in \text{tags}(b(y)) \wedge y_t \rightarrow^b x_t\} \\
\llbracket x' = y \rrbracket^e &= \{b \in \mathcal{B}|_{\text{vars}(e) \cup \{x, y\}} \mid \forall t \in \llbracket e \rrbracket_b, \\
&\quad t \in C = \text{tags}(b(x)) \wedge t \in \text{tags}(b(y)) \wedge b(x)(\text{succ}_t(C)) = b(y)(t)\} \\
\llbracket y \rightarrow x' \rrbracket^e &= \{b \in \mathcal{B}|_{\text{vars}(e) \cup \{x, y\}} \mid \forall t \in \llbracket e \rrbracket_b, \\
&\quad t \in C = \text{tags}(b(x)) \Rightarrow t \in \text{tags}(b(y)) \wedge y_t \rightarrow^b x_{\text{succ}_t(C)}\} \\
\llbracket x^0 = v \rrbracket^e &= \{b \in \mathcal{B}|_x \mid b(x)(\min(\text{tags}(b(x)))) = v\} \\
\llbracket f \Rightarrow P \rrbracket^e &= \llbracket P \rrbracket^{e \wedge f} \quad \llbracket P \mid Q \rrbracket^e = \llbracket P \rrbracket^e \mid \llbracket Q \rrbracket^e \quad \llbracket P/x \rrbracket^e = \llbracket P \rrbracket^e / x
\end{aligned}$$

Fig. 7. Denotational semantics of propositions.

The meaning of a proposition $x = y$ at the clock e consists of all behaviors b defined on $\text{vars}(e) \cup \{x, y\}$ such that all tags $t \in \llbracket e \rrbracket_b$ at the clock e belong to $b(x)$ and $b(y)$ and are associated with the same value. A scheduling specification $y \rightarrow x$ at the clock e denotes the set of behaviors b defined on $\text{vars}(e) \cup \{x, y\}$ which, for all tags $t \in \llbracket e \rrbracket_b$, requires x to precede y : if t is in $b(x)$ then it is necessarily in $b(y)$ and satisfies $y_t \rightarrow^b x_t$. The propositions

$x' = y$ and $y \rightarrow x'$ is interpreted similarly by considering the tag t' that is the successor of t in the chain C of x . The behavior of a guarded command $f \Rightarrow P$ at the clock e is equal to the behavior of P at the clock $e \wedge f$. The behavior of $P|Q$ consists the synchronous composition of the behaviors of P and Q .

4. AN INTERMEDIATE REPRESENTATION

We are now equipped with the required mathematical framework to address the modeling of embedded systems described by communicating program threads. This model is described in terms of a type inference system and extended to the structuring elements of a generic module system. This framework allows to give a behavioral signature of the component of the system, compositionally check the correct composition of such components to form architecture, to optimize the described software elements from the imposed hardware elements by, first, detaching the formal model from the functional architecture description and, second, using the model to regenerate an optimized software matching the requirements of the execution architecture.

Formal syntax. Imperative programs are represented in an intermediate form that is common to the TAC and SSA IRs of GCC which provides language-independence and local optimization. A program pgm consists of a sequence of labeled blocks $L:blk$. Each block consists of a label L and of a sequence of statements stm terminated by a return statement rtn .

```

(program)   $pgm ::= L:blk | pgm; pgm$ 
(block)     $blk ::= stm; blk | rtn$ 
(instruction)  $stm ::= x = f(y_{1..n})$ 
           |  $\text{if } x \text{ then } L$ 
(return)    $rtn ::= \text{goto } L$ 
           |  $\text{return } x$ 
           |  $\text{throw } x$ 
           |  $\text{catch } x \text{ from } L \text{ to } L \text{ using } L$ 

```

Fig. 8. Syntax for an intermediate representation of imperative programs.

Block instructions consist of native method invocations $x = f(y_{1..n})$, lock monitoring and branches $\text{if } x \text{ then } L$. Blocks are returned from by either a $\text{goto } L$, a return or an exception $\text{throw } x$. The declaration $\text{catch } x$

from L_1 to L_2 using L_3 that matches an exception x raised at block L_1 activates the exception handler L_3 and continues at block L_2 .

In the remainder, we only assume that a block always starts with a label and finishes with a return statement: $stm_1; L:stm_2$ is rewritten as $stm_1; goto L; L:stm_2$. A call $x = f(y)$ to a possibly blocking external method f , such as `waitx` in SystemC or Java, is always placed at the beginning of a block L . For instance, $stm_1; waitx; stm_2$ is rewritten as $stm_1; goto L; L:waitv; stm_2$. By contrast, primitive operations $x = f(y, z)$ are assumed to take an insignificant amount of time and are executed with the normal control-flow of the block.

Example 2. To outline the construction of the intermediate representation of a program, let us reconsider the example of Section 2 and detail the function that counts the number of bits set to 1 in a bit-array data (Figure 9). It consists of three blocks. The block labeled L1 waits for the signal lock before initializing the local state variable `idata` to the value of the input signal data and `ocount` to 0. Label L2 corresponds to a loop that shifts `idata` right to add its right-most bit to `ocount` until termination (condition T0). In the block L3, `ocount` is sent to the signal `count` and lock is unlocked before going back to L1.

L1:wait (lock); idata=data; ocount=0; goto L2; L2:T1 = idata;	T0 = T1 == 0; if T0 then goto L3; T2 = ocount; T3 = T1 & 1; ocount = T2 + T3;	idata = T1 >> 1; goto L2; L3:notify (lock); count = ocount; goto L1;
---	---	--

Fig. 9. From three address code ...

The SSA form of the program differs in the function-wise guarantee that all variable be assigned once during an execution cycle. It consists of performing assignments to `idata` and `ocount` in blocks L1 and L2 to temporary variables and branch to a merge block L4 where the appropriate copy is assigned to the variable upon the value of a boolean condition ϕ (to mean from L1 or not).

Meaning of instructions. The denotation of instructions for programs which strictly adhere either of the TAC or SSA requirements (i.e. all variables are written at most once per block) is given figure 11. To lighten notations, we write $C = \text{chain}_b(X)$ iff for all $x \in X$, $C = \text{tags}(b(x))$ and write $b(x)(t)$ for $b(x)(t) = \text{true}$ and $\neg b(x)(t)$ for $b(x)(t) = \text{false}$. The denotation of

L1: ... idata1=data; ocount1=0; goto L4;	L2: ... ocount2 = T2 + T3; idata2 = T1 >> 1; goto L4;	L4: idata=φ?idata1,idata2; ocount=φ?ocount1,ocount2; goto L2;
--	---	---

Fig. 10. ... to static single assignment.

a program $\langle\langle pgm \rangle\rangle^E$ takes an environment giving the meaning of external functions f using call-by-name λ -expressions and returns the set of behaviors b corresponding to the execution of pgm .

For an instruction stm , the function $\langle\langle stm \rangle\rangle_{L_1 L_2}^E$ takes two labels which represent the entry label L_1 of the statement and its continuation by the pseudo-label L_2 . The denotation of a function call $x = f(x_{1..k})$ is that given by E for the variable names $x_{1..k}x$ and the entry and exit labels L_1 and L_2 .

The meaning of an `if  $x$  then  $L_1$` instruction consists of all behaviors b defined on x , L_1 , L_2 and L_3 which share the same chain of tags C and such that, if $b(L_1)(t)$ is true, then the continuation label L_3 is active iff x is false, i.e. $b(L_3)(t) = \neg b(x)(t)$; and if x is true then L_2 is active next, i.e. $b(x)(t)$ true implies $b(L_2)(\text{succ}_t(C))$ true. For a return instruction rtn , the denotation function $\langle\langle rtn \rangle\rangle_L^E$ only takes one (entry) label L . The meaning of `return  $x$` , `goto  $L$` and `throw  $x$` instructions are given using the same principle as for the `if  $x$  then  $L$` .

$$\begin{aligned}
& \langle\langle x = f(x_{1..k}) \rangle\rangle_{L_1 L_2}^E = E(f)(x_{1..k}xL_1 L_2) \\
& \langle\langle \text{if } x \text{ then } L_1 \rangle\rangle_{L_2 L_3}^E = \{b \in \mathcal{B}|_{xL_{123}} \mid \forall t \in C = \text{chain}_b(xL_{123}), \\
& \quad b(L_1)(t) \Rightarrow (b(L_3)(t) = \neg b(x)(t)) \\
& \quad b(x)(t) \Rightarrow b(L_2)(\text{succ}_t(C))\} \\
& \langle\langle \text{return } x \rangle\rangle_L^E = \{b \in \mathcal{B}|_{Lx} \mid E(\text{return}) = y, \forall t \in C = \text{chain}_b(Lx), \\
& \quad b(L)(t) \Rightarrow b(y)(t) = b(x)(t)\} \\
& \langle\langle \text{goto } L_1 \rangle\rangle_{L_2}^E = \{b \in \mathcal{B}|_{L_1 L_2} \mid \forall t \in C = \text{chain}_b(L_1 L_2), \\
& \quad b(L_2)(t) \Rightarrow b(L_1)(\text{succ}_t(C))\} \\
& \langle\langle \text{throw } x \rangle\rangle_L^E = \{b \in \mathcal{B}|_{Lx} \mid \forall t \in C = \text{chain}_b(Lx), \\
& \quad b(L)(t) \Rightarrow b(x)(t)\} \\
& \langle\langle stm; blk \rangle\rangle_{L_1}^E = (\langle\langle stm \rangle\rangle_{L_1 L_2}^E \mid \langle\langle blk \rangle\rangle_{L_2}^E) / L_2 \\
& \langle\langle L : blk; pgm \rangle\rangle^E = \langle\langle blk \rangle\rangle_L^E \mid \langle\langle pgm \rangle\rangle^E \\
& \langle\langle m f(x_{1..k}) \{pgm\} \rangle\rangle^E = E[f : \lambda x_{1..k}xyy^{\text{exit}}.(p/L_{1..j})] \mid \\
& \quad p = \langle\langle pgm \rangle\rangle^{E[\text{return}:y]} \wedge \text{labs}(pgm) = L_{1..j}
\end{aligned}$$

Fig. 11. Denotational semantics of instructions.

Notice the introduction of a pseudo-label to handle a sequence of instructions. The meaning of a sequence $stm; blk$ starting at block L_1 is defined by using a local pseudo-label L_2 to denote the continuation of stm by $\langle\langle stm \rangle\rangle_{L_1 L_2}^E$ and hence the entry point of blk by $\langle\langle blk \rangle\rangle_{L_2}^E$. The meaning of the sequence is finalized by synchronous composition and the scope of L_2 restricted to it. The meaning of a program $L : blk; pgm$ is similar yet simpler as there is no continuation between blocks. The meaning of a function declaration $m f(x_{1..k}) \{pgm\}$ is listed just to show the order in which the argument, result, entry and exit label names are used to parameterize the meaning of the function body.

5. BEHAVIORAL TYPE INFERENCE

The behavioral type inference system is defined by induction on the formal syntax of programs pgm . To define it, we assume that the finite set $\mathcal{L}$ of program labels L . To each block of label L , the inference system associates a boolean proposition L of the same name, called the *input clock*, and a boolean proposition L^{exit} , called its *output clock*. The proposition L is true iff the block L is active during a given transition. The proposition L^{exit} is true iff the execution of block L terminates during a given transition. The relation defined by the behavioral type system has the form:

$$e_0, \mathcal{E} \vdash L : blk : \langle P, e_1 \rangle$$

where e_0 denotes the input clock of the block of instructions blk , L is its label, P the proposition to denote its behavior, and e_1 its output or continuation clock. The type environment $\mathcal{E}$ gives the behavior of methods and functions defined in the context of the program. It associates a variable x to a type m (a class name), a class name m to a class type $\mathcal{T}$ (described in the next section) and a method f to a proposition P and an output clock e parameterized by the sequence $x_{1..n}$ formed of its input and output variables and input clock name (see rule (8) below).

$$\mathcal{E} ::= [] \mid \mathcal{E}[x : m] \mid \mathcal{E}[m : \mathcal{T}] \mid \mathcal{E}[f : \lambda_{(1..n)}. \langle P, e \rangle]$$

Rules (1–8) define the behavioral type inference system. Rules (1–2) are concerned with the iterative decomposition of a program pgm into blocks blk and with the decomposition of a block into statements stm and return instruction rtn .

$$(1) \frac{L, \mathcal{E} \vdash L : blk : P \quad \mathcal{E} \vdash pgm : Q}{\mathcal{E} \vdash L : blk; pgm : P \mid Q}$$

Notice that, in rule (2), the input clock e of the block $stm; blk$ is passed to stm . The output clock e_1 of stm becomes the input clock of blk . The input and output clocks of an instruction may differ.

$$(2) \frac{e_1, \mathcal{E} \vdash L : stm : P, e_2 \quad e_2, \mathcal{E} \vdash L : blk : Q}{e_1, \mathcal{E} \vdash L : stm; blk : P \mid Q}$$

This is the case, rule (3), for instruction `if  $x$  then  $L_1$` . Let e be the input clock of the instruction. If x is false then control is passed to the continuation of this instruction in the block, at the output clock $e \wedge \neg x$. Otherwise, control is passed to block L_1 , at the clock $e \wedge x$. Hence the type $(e \wedge x) \Rightarrow L'_2$ to mean that the next value of L_2 is true when e is active and when x is true.

$$(3) e, \mathcal{E} \vdash L : \text{if } x \text{ then } L_1 : \langle (e \wedge x) \Rightarrow (L^{exit} \mid L'_1), e \wedge \neg x \rangle$$

All return instructions, rules (4–7), define the output clock L^{exit} of the current block L by the input clock e . This is the right place to do that: e defines the condition upon which the block actually reaches its return statement. A `goto  $L_1$` instruction, rule (4), passes control to block L_1 unconditionally at the input clock e .

$$(4) e, \mathcal{E} \vdash L : \text{goto } L_1 : e \Rightarrow (L^{exit} \mid L'_1)$$

A return instruction, rule (5), fetches the variable y used as return variable for the current method or function and sets y^{exit} to true at clock e in order to notify the caller that the method terminates execution.

$$(5) \frac{\mathcal{E}(\text{return}) = y}{e, \mathcal{E} \vdash L : \text{return } x : e \Rightarrow (L^{exit} \mid y^{exit} \mid y := x)}$$

A `throw  $x$` instruction, rule (6), produces an event along the signal x at the input clock e by $e \Rightarrow \hat{x}$.

$$(6) e, \mathcal{E} \vdash L : \text{throw } x : e \Rightarrow (L^{exit} \mid \hat{x})$$

Example 3. Let us zoom on the block L2 of Figure 2. On the first line, for instance, we associate the instruction `T1 = idata` of block label L2 to the proposition $L2 \Rightarrow T1 = \text{idata}$. In this proposition, the variable L2 is a boolean that is true iff the block L2 is being executed. So, the proposition says that, if L2 is being executed, then T1 is always equal to idata. If it not, another proposition may hold. All subsequent propositions are conditioned by L2 to mean that they hold when L2 is executed. Next, consider

the instruction `if T0 then L3`. Its invariant $L2 \Rightarrow T0 \Rightarrow L3'$ says that control passes to L3 when control is presently at L2 and when T0 is true. The instructions that follow this test are conditioned by the negative $\neg T0$, this means: "in the block L2 and not in its branch to L3".

$$\begin{array}{c|c} L2: \dots \text{if } T0 \text{ then goto } L3; & L2 \Rightarrow \dots T0 \Rightarrow L3' \\ \vdots & \neg T0 \Rightarrow \dots L2' \\ \text{goto } L2; & \end{array}$$

Fig. 12. Modeling control flow in an imperative program.

The catch statement `catch  $x$  from  $L$  to  $L_1$  using  $L_2$` matching rule (6), passes control in rule (7) to the exception handler L_2 and then to the block L_1 upon termination of L_2 notified by L_2^{exit} . This requires, first, to activate L_2 from L when x is present and then to pass the control to L_1 upon termination of the handler.

$$(7) \ e, \mathcal{E} \vdash L : \text{catch } x \text{ to } L_1 \text{ using } L_2 : (\hat{x} \wedge L^{exit} \Rightarrow L_2' \mid L_2^{exit} \Rightarrow L_1')$$

Rule (8) is concerned with type assignment for native and external method invocations $x = f(x_{1..k})$. The generic type of f is taken from an environment $\mathcal{E}(f)$. It is given the name of the actual parameters $x_{1..k}$, of the result x and of the input clock e . $\mathcal{E}(f)(x_{1..k}xe)$ yields the corresponding behavioral type $\langle P, e_1 \rangle$.

$$(8) \ e, \mathcal{E} \vdash L : x = f(x_{1..k}) : \mathcal{E}(f)(x_{1..k}, x, e)$$

Example 4. As an example, the wait-notify protocol used in SystemC of Java to arbiter access to shared data is modeled using a boolean flip-flop variable x . The `notify` method defines the next value of the lock x by the negation of its current value at the input clock e . The `wait` method continues activates iff the value of the lock x has changed at the input clock L : $L \wedge (x \neq x')$. Otherwise, at the clock $L \wedge (x = x')$, the control is passed to L by a delayed transition $e \setminus \hat{y} \Rightarrow L'$.

$$\begin{aligned} \mathcal{E}(\text{notify}) &= \lambda xe. \langle e \Rightarrow (x' = \neg x), e \rangle \\ \mathcal{E}(\text{wait}) &= \lambda xL. \langle L \wedge (x = x') \Rightarrow L', L \wedge (x \neq x') \rangle \end{aligned}$$

Consider the wait-notify protocol at blocks L1 and L3, Figure 13. The wait instruction continues if L1 receives control and if the lock is toggled (proposition $\text{lock} \neq \text{lock}'$). If so, the block is executed and control passes to the block L2 and, if not, to the block L1.

Fig. 13. Modeling the access to locks.

Completion. By definition, a proposition L holds the value `true` iff the block L is active during execution. Otherwise, L should be `false`. This default value requires a completion of the next-state logic for the type P of a given program pgm . We write $\bar{P}$ this completion. It is simply defined by considering the proposition $e_L \Rightarrow L'$ implied by the type P for all labels L of a given program pgm . The clock e_L is defined by the union (disjunction) of all clocks $e \Rightarrow L'$ present in P . The default rule is defined by $\bar{L} \setminus e_L \Rightarrow \neg L'$. The same holds for output clocks L^{exit} .

Correspondence. The correspondance between instructions and propositions defined through our type system $\mathcal{E} \vdash pgm : P$ can now be formally established by stating Property 1. We write $\llbracket \mathcal{E} \rrbracket$ for the interpretation of the environment $\mathcal{E}$ defined by

$$\llbracket \mathcal{E} \rrbracket [f : \lambda(x_1..kxL). \langle P, x^{exit} \rangle] = \llbracket \mathcal{E} \rrbracket [f : \lambda(x_1..kxLx^{exit}). \llbracket P \rrbracket]$$

Property 1 established a classical soundness property by stating that whenever pgm has type P and the typing environment $\mathcal{E}$ has meaning E then b is a behavior of P (guarded by 1 to mean always) if and only if it is a behavior of pgm with the environment E . Notice that the top-level environment $\mathcal{E}$ defines the model of the runtime communication and processes management API for the application program pgm . The proof of property 1 consists of showing that both implications $\llbracket P \rrbracket \subseteq \langle\langle pgm \rangle\rangle$ and $\llbracket P \rrbracket \supseteq \langle\langle pgm \rangle\rangle$ hold by induction on the structure of pgm ending up in a case analysis on the correspondence between each instruction.

Property 1.

If $\mathcal{E} \vdash pgm : P$ and $E = \llbracket \mathcal{E} \rrbracket$ then $b \in \llbracket P \rrbracket_1$ iff $b \in \langle\langle pgm \rangle\rangle^E$

From TAC to SSA. The type system and its semantics rely on the property of the TAC IR that every variable is defined at once within a block (this hypothesis is sound for a program in SSA form as well). As a consequence, each block delimits an atomic reaction in the type system and, therefore, transition from a block to another cannot be immediate (by saying L for "label L is active") but delayed (by saying L' for "label L will be active next time"). In SSA, this guarantee is provided for the whole "text" of the

function. In particular, for a goto from a block L_1 to a block L_2 textually after L_1 (written $L_1 < L_2$), SSA guarantees that all variables defined in L_1 are different from those in L_2 . This is of course not the case for a loop, in which case we have $L_1 \geq L_2$. To take advantage of this additional guarantee, our type inference system can be refined by considering the following rule to handle *gotos* (and similarly, *if-thens* and *throw-catches*). It consists of activating the target block L_2 immediately.

$$(4b) \frac{L_1 < L_2}{e, \mathcal{E} \vdash L_1 : \text{goto } L_2 : e \Rightarrow (L_1^{exit} \mid L_2)}$$

The translation of the EPC in SSA form using rule (3b) outlines the benefits of this optimization. The resulting type has strictly fewer delayed transitions: one to L2 in L3 and another to L1 in L4. All other transitions are immediate and considered within the same reaction.

L1:wait (lock);	L1 $\Rightarrow$ lock=lock' $\Rightarrow$ L1'
idata1=data;	lock $\neq$ lock' $\Rightarrow$ idata1:=data
ocount1=0;	ocount1:=0
goto L3;	L3
L2:T1 = idata;	L2 $\Rightarrow$ T1 := idata
T0 = T1 == 0;	T0 := T1 == 0
if T0 then goto L3;	T0 $\Rightarrow$ L4
T2 = ocount;	$\neg$ T0 $\Rightarrow$ T2 := ocount
T3 = T1 & 1;	T3 := T1 & 1
ocount2 = T2 + T3;	ocount2 := T2 + T3
idata2 = T1 >> 1;	idata2 := T1 >> 1
goto L3;	L3
L3:idata= ϕ ?idata1,idata2;	L3 $\Rightarrow$ L1 $\Rightarrow$ idata:=idata1
ocount= ϕ ?ocount1,ocount2;	ocount:=ocount1
goto L2;	L2 $\Rightarrow$ idata:=idata2
L4:notify (lock);	ocount:=ocount2
count = ocount;	L2'
goto L1;	L4 $\Rightarrow$ lock':= $\neg$ lock
	count := ocount
	L1'

Fig. 14. Model of the even-parity checker in SSA form.

6. CONFORMANCE CHECKING

Just as the multi-clocked synchronous formalism Signal it is based upon, our type system allows for the refinement-based design methodologies

considered in ⁽¹⁵⁾ to be easily implemented. Checking the correct refinement of an initial module, of type P , by its upgrade, of type Q , amounts to checking that the final guarantee Q satisfies the initial assumptions P . In ⁽¹⁵⁾, this is implemented by compositionally model checking that Q is finitely flow-equivalent to P .

Fig. 15. Conformance-checking the refinement of an even-parity checker.

Figure 15 describes a typical case study of conformance checking. We consider the refinement of the C model of an even parity checker (EPC) from a high-level design abstraction, left, where communication is abstracted by shared variables and a lock, to an architecture-level design abstraction, right, where the communication medium is refined by the insertion of a channel implementing a double handshake protocol, Figure 16.

Fig. 16. Refinement of locks with a double handshake protocol.

Checking conformance of the architecture-level design with respect to its system-level abstraction amounts to checking that both designs are flow equivalent. The very notion of flow equivalence under consideration con-

sists is defined in the asynchronous structure of our model of computation that is presented next.

6.1. Asynchronous structure

The asynchronous structure of polychrony is modeled by weakening the clock-equivalence relation to allow for comparing behaviors whose successive values match regardless of time: two behaviors are flow-equivalent iff their signals hold the same values in the same order. The *relaxation* relation allows to individually stretch the signals of a behavior in a way preserving scheduling constraints. A behavior c is a *relaxation* of b , written $b \sqsubseteq c$, iff $\text{vars}(b) = \text{vars}(c)$ and, for all $x \in \text{vars}(b)$, $b|_{\{x\}} \leq c|_{\{x\}}$.

Fig. 17. Relating asynchronous behaviors by relaxation.

Relaxation is a partial-order relation which defines flow-equivalence: b and c are *flow-equivalent*, written $b \approx c$, iff there exists a behavior d s.t. $d \sqsubseteq b$ and $d \sqsubseteq c$. Figure 17 illustrates two asynchronously equivalent behaviors related by relaxation. The first event along x has been shifted (and its scheduling constraint with an initially synchronous event along y lost) as the effect of finitely delaying its transmission. Asynchronous composition is noted $p \parallel q$ and defined using the partial-order structure induced by the relaxation relation. The composition of p and q consists of behaviors d that are relaxations of behaviors b and c from p and q along shared signals $I = \text{vars}(p) \cap \text{vars}(q)$, i.e. $b|_I \sqsubseteq d|_I \sqsubseteq c|_I$, and that are stretching of b and c along the independent signals of p and q , i.e. $b/I \leq d/I \leq c/I$.

$$\left(\begin{array}{c} x : \bullet^{t_1} \quad \bullet^{t_3} \\ y : \bullet^{t_2} \quad \bullet^{t_4} \end{array} \right) \parallel \left(\begin{array}{c} x : \bullet^{t_1} \quad \bullet^{t_2} \\ y : \bullet^{t_1} \quad \bullet^{t_2} \\ z : \bullet \quad \bullet \quad \bullet \end{array} \right) \ni \left(\begin{array}{c} x : \bullet \quad \bullet \quad \bullet \\ y : \bullet \quad \bullet \quad \bullet \\ z : \bullet \quad \bullet \quad \bullet \end{array} \right)$$

Fig. 18. Asynchronous composition.

Figure 18 illustrates the asynchronous composition of a behavior b and of a behavior c . Signals x and y are alternated in b , left, and synchronous in c , middle. Asynchronous composition allows x and y to be independently

stretched in b and c in order to find a common flow in the asynchronous composition, right.

6.2. Flow preservation

To check the existence of a flow-preserving timing relation between the two systems outlined in the previous section, the refinement-based methodology similar of ⁽¹⁵⁾ shows that the types P and Q of Figure 15 are finitely flow-equivalent. To this end, we formulate the timing deformation allowed by finite buffering protocols starting from the model of a one-place FIFO buffer which we will use to draw the spectrum of possible timing relations under consideration. Figure 19 depicts the timing deformation allowed along a signal x by a one place buffer.

Fig. 19. Relation between events through a one place buffer along x .

Finite relaxation. Definition 1 formalizes this relation by considering the timing deformation between an initial behavior b and a final behavior c performed by a one-place FIFO buffer of internal signal m and behavior d . The behavior d is defined by stretching $b \leq d/m$ and c/x by d/mx . Let us write $\text{pred}_C(t)$ (resp. $\text{succ}_C(t)$) for the immediate predecessor (resp. successor) of the tag t in the chain C .

Definition 1 (finite relaxation). *The behavior c is a 1-relaxation of x in b , written $b \sqsubseteq_1^x c$ iff $\text{vars}(b) = \text{vars}(c)$ and there exists a signal m , a behavior d and a chain $C = \text{tags}(d(m)) = \text{tags}(d(x)) \cup \text{tags}(c(x))$ such that $d/m \geq b$, $d/mx = c/x$ and, for all $t \in C$,*

- (1) $t \in \text{tags}(d(x)) \Rightarrow d(x)(t) = d(m)(t) \wedge x_t \rightarrow_d m_t$
- (2) $t \notin \text{tags}(d(x)) \Rightarrow d(m)(t) = d(m)(\text{pred}_C(t))$
- (3) $t \in \text{tags}(c(x)) \Rightarrow c(x)(t) = d(m)(t) \wedge \forall y \in \text{vars}(d) \setminus m, y_t \rightarrow_d x_t$
- (4) $t \in \text{tags}(c(x)) \Rightarrow c(x)(t) = d(x)(t) \vee c(x)(\text{succ}_C(t)) = d(x)(t)$

For all $t \in C$, rule (1) says that, when an input $d(x)$ is present at some time t , then $d(m)$ takes its value. If no input is present along x at t , rule (2),

then $d(m)$ takes its previous value. Rule (3) says that, if the output $c(x)$ is present at t , then it is defined by $d(m)(t)$. Finally, rule (4) requires this value to either be the present or previous value of the input signal $d(x)$, binding the size of the buffer to one place.

Fig. 20. Timing and scheduling relations through finite relaxation.

Definition 1 accounts for the behavior of bounded FIFOs in a way that preserves scheduling relations. It implies a series of (reflexive-anti-symmetric) relations $\sqsubseteq_n$ (for $n > 0$) which yields the (series of) reflexive-symmetric flow relations $\approx_n$ to identify processes of same flows up to a flow-preserving first-in-first-out buffer of size n . We write $b \sqsubseteq_1 c$ iff $b \sqsubseteq_1^x c$ for all $x \in \text{vars}(b)$, and, for all $n > 0$, $b \sqsubseteq_{n+1} c$ iff there exists d such that $b \sqsubseteq_1 d \sqsubseteq_n c$. The largest equivalence relation modeled in the polychronous model of computation consists of behaviors equal up to a timing deformation performed by a finite FIFO protocol: b and c are *finitely flow-equivalent*, written $b \approx^* c$, iff there exists $n > 0$ and d s.t. $d \sqsubseteq_n b$ and $d \sqsubseteq_n c$.

6.3. A compositional methodology

We say that a process P is finitely flow-preserving iff given finitely flow-equivalent inputs, it can only produce behaviors that are finitely flow equivalent.

Definition 2 (finite flow-preservation).

P is finitely flow-preserving with $I \subset \text{in}(P)$ iff for all behaviors b, c of $\llbracket P \rrbracket$, if $(b|_I) \approx (c|_I)$ then $b/I \approx^* c/I$.

Example of finitely flow-preserving processes are endochronous processes⁽⁸⁾. An endochronous process which receives flow equivalent inputs produces clock-equivalent outputs. It hence forms a restricted subclass of finitely-flow preserving processes. Furthermore, notice that flow-preservation is stable to the introduction of a wrapper of P consisting of a finite FIFO buffering protocol. A refinement-based design methodology based on the property of finite flow-preservation consists of characterizing sufficient invariants for a given model transformation to preserve flows.

Definition 3 (finite flow-invariance).

The transformation of P into Q such that $I \subset \text{in}(P) = \text{in}(Q)$ is finitely flow-invariant iff $\forall b \in \llbracket P \rrbracket, \forall c \in \llbracket Q \rrbracket, (b|_I) \approx^*(c|_I) \Rightarrow b \approx^* c$

The property of finite flow-invariance is a very general methodological criterion. For instance, it can be applied to the characterization of correctness criteria for model transformations such as protocol insertion or desynchronization. Let P and Q be two finitely flow-preserving processes and R a protocol to link P and Q , such as a finite FIFO buffer, or a double hand-shake protocol, or a relay station ⁽⁶⁾, or a loosely time-triggered architecture ⁽⁴⁾. In definition 4, we write $b[x/y]$ for the behavior resulting of substitution of the signal name y by the signal name x in the domain of the behavior b and $[x_i/y_i]_{0 \leq i \leq n}$ for the composition of n substitutions.

Definition 4 (flow-preserving protocol).

The process R is a flow-preserving protocol iff there exists $n > 0$ such that inputs $\text{in}(R) = \{x_{1..n}\}$ are finitely flow-equivalent to outputs $\text{out}(R) = \{y_{1..n}\}$, i.e., $\forall b \in \llbracket R \rrbracket, b|_{x_{1..n}} \approx^* (b|_{y_{1..n}}[x_i/y_i]_{0 \leq i \leq n})$

The wrapper $R\langle P \rangle$ of a process P with a protocol R is defined by redirecting the signals of P to R . In definition 5, this redirection is modeled by substituting signal names: we write $P[x/y]$ for the process resulting of substituting y by x in P .

Definition 5 (wrapper).

Let P be a process such that $\text{in}(P) = \{x_{1..n}\}$ and $\text{out}(P) = \{x_{m+1..n}\}$. Let R be a flow-preserving protocol such that $\text{in}(R) = \{y_{1..n}\}$ and $\text{out}(R) = \{z_{1..n}\}$. The wrapper of P with R is the template process noted $R\langle P \rangle$ and defined by:

$$R\langle P \rangle \stackrel{\text{def}}{=} \left(\begin{array}{l} ((R[x_i/z_i]_{m < i \leq n})[x_i/y_i]_{0 < i \leq m}) \\ | (P[y_i/x_i]_{m < i \leq n})[z_i/x_i]_{0 < i \leq m} \end{array} \right) / y_{1..n} z_{1..n}$$

A sufficient condition for the insertion of a protocol between two synchronous processes P and Q to finitely preserve flow is to guaranty that $P|_I | Q|_I$ is finitely flow preserving for $I = \text{vars}(P) \cap \text{vars}(Q)$, meaning that all communications between P and Q via a shared signal $x \in I$ should be flow preserving and that P and Q may otherwise evolve independently.

Property 2 (protocol insertion).

If R is a flow-preserving protocol and P is finitely flow-preserving then $R\langle P \rangle$ is finitely flow-preserving. If R is a flow-preserving protocol and $P, Q, P|_I | Q|_I$ are finitely flow-preserving then $R\langle P \rangle | R\langle Q \rangle$ is finitely preserving ($I = \text{vars}(P) \cap \text{vars}(Q)$).

7. FURTHER APPLICATIONS

We have introduced a type system allowing to model the control and data flow graphs of a given imperative program in intermediate form. Applications of the proposed type system encompass optimization and verification issues encountered in system design.

7.1. Rethreading

Because our type system entirely model the control and data-flow of application components and architecture functionalities, one can operate global optimization on the whole model of the application. Signal, in particular, implements the notation of our type system using data-flow equations and allows for the generation of sequential code by employing a global control-flow graph transformation called hierarchization⁽²⁾. Hierarchization consists of hooking elementary control flow graphs (in the form of if-then-else structures). For instance, let $h3$ be a clock computed using $h1$ and $h2$ and h be the head of a tree in which $h1$ and $h2$ are computed. Then $h3$ can be computed after $h1$ and $h2$ and placed under h .

Example 5. The implications of hierarchization for code generation can be outlined by considering the specification of one-place buffer. The process buffer has input x , output y and implements two functionalities.

$$\text{buffer} \langle x, y \rangle \stackrel{\text{def}}{=} \text{alternate} \langle x, y \rangle \mid \text{current} \langle x, y \rangle$$

One is the process **alternate** which desynchronizes the signals x and y by synchronizing them to the true and false values of an alternating boolean signal s .

$$\text{alternate} \langle x, y \rangle \stackrel{\text{def}}{=} (s^0 = \text{true} \mid \hat{x} = s \mid \hat{y} = \neg s \mid s' := \neg s)$$

The other functionality is the process **current**. It defines a cell in which values are stored at the input clock $\hat{x}$ and loaded at the output clock $\hat{y}$.

$$\text{current} \langle x, y, b \rangle \stackrel{\text{def}}{=} (r^0 = b \mid r' := x \mid \hat{x} \Rightarrow y := x \mid \hat{y} \setminus \hat{x} \Rightarrow y := r)$$

We observe that s defines the master clock of **buffer**. There are two other synchronization classes, x and y , that corresponds to the true and false values of the boolean flip-flop variable s , respectively. This defines three nodes in the control-flow graph of the generated code (Figure 21). At the master clock $\hat{s}$, the value of s is calculated from zs , its previous value. At the sub-clock $s = \hat{x}$, the input signal x is read. At the sub-clock $\neg s = \hat{y}$ the output signal y is written. Finally, the new value of zs is determined.

```
buffer_iterate () {
  s = !zs;
  cy = !s;
  if (s) { if (!r_buffer_i(&x)) return FALSE; }
  if (cy) { y = x; w_buffer_o(y); }
  zs = s;
  return TRUE; }
```

Fig. 21. C code generated for the one-place buffer specification.

Operating this transformation on the model of a multi-threaded application results in merging all threads into a single control-flow graph whose scheduler foot-prints sequentially processes each elementary execution block upon a particular condition. In ⁽¹⁴⁾, we report a 300% average speedup resulting of applying this optimization to real-time Java programs compared to their execution using a commercial compiler.

7.2. Module checking

In ⁽¹⁶⁾, we define a behavioral module checking algorithm based on similar principles as those exposed in the previous section. This system allows to give guarantees As an example, consider a SystemC class m_0 whose virtual fields are the clocks x, y and a procedure f . Assume an explicit behavioral type declaration $\#TYPE(f, Q)$ which associates f with a description of its behavior: the proposition Q denotes its expected functionality. Let us associate the interface m_0 with the class parameter m_1 of a template class m_2 . The interface m_0 now gives a behavioral type to the method f in the class parameter m_1 expected by the module m_2 . The assumption Q on the

behavior of $m_1.f$ is required to provide a guarantee on the behavior of the module m_2 produced by the template class. Module m_3 is a candidate parameter for m_2 . It structurally implements the interface m_0 and is annotated with the guarantee $\#TYPE(f, P)$, where P is the type of pgm . Now, let m_4 be the class defined by the instantiation of the template m_2 and the parameter m_3 . To check the compatibility of the actual parameter m_3 with the formal parameter m_0 , we check the containment of the behaviors denoted by the proposition P (the type of the actual parameter) in the proposition Q (the type of the formal parameter). This amounts to check that P implies Q , either by model checking (if Q contains state transitions) or by static checking (if Q is a "stateless" property).

```
class  $m_0$  { virtual sc_clock  $x, y$ ; virtual void  $f()$  {} #TYPE( $f, Q$ ) };
template <class  $m_1$ > #TYPE( $m_1, m_0$ )
  SC_MODULE( $m_2$ ) { SC_CTOR( $m_2$ ) { SC_THREAD( $m_1.f$ ) sensitive <<  $x$  } };
class  $m_3$  { sc_clock  $x, y$ ; void  $f()$  {  $pgm$  } #TYPE( $f, P$ ) };
 $m_2$  < $m_3$ >  $m_4$ ;
```

Fig. 22. Type assumptions and guarantees in the SystemC module system.

We consider a simple and minimalistic module system model for the purpose of exemplifying the scalability of our technique to structuring elements of general-purpose languages such as Java, C++ or SystemC. A component mod in an architecture is a class definition $class\ m\{dec\}$, a template declaration $template\ <class\ x : m\>\ mod$ or a sequence of modules $mod; mod$. A class consists of a sequence of declarations. The keyword $use\ m$ allows to use the members of class m within the current module (hence name elaboration is assumed to be explicit for simplification purposes). Declarations dec associate locations x with native classes m or template class instances $m\langle m_{1..k} \rangle$ and methods with a name f and a definition pgm . For instance, `integer  $x$` defines an integer variable x (in Java or C) while `sc_signal<boolean>  $x$` defines a boolean signal x in SystemC. As we focus on typing program module behaviors we assume no sub-typing relation between data-types.

$$mod ::= class\ m\{dec\} \mid template\ <class\ x : m\>\ mod \mid mod; mod$$

$$dec ::= m\langle n_{1..k} \rangle\ x \mid use\ m \mid m\ f(x_{1..k})\ \{pgm\} \mid dec; dec$$

Fig. 23. Abstract syntax for declarations and modules.

We define our module system starting from the behavioral type system of Section 5. The type $\mathcal{T}$ of a module m consists of an environment $\mathcal{E}$ (that associates functions f with behaviors and variables x with data-types) and of a proof obligation $\mathcal{C}$. The type $\mathcal{T}_1 \rightarrow \mathcal{T}_2$ denotes a template class that produces a module of type $\mathcal{T}_2$ given a parameter of type $\mathcal{T}_1$.

$$(\text{type}) \mathcal{T} ::= \mathcal{E} / \mathcal{C} \mid \Lambda x : \mathcal{T}_1. \mathcal{T}_2$$

A proof obligation is a conjunction of propositions of the form $P \Rightarrow Q$. A proof obligation $P \Rightarrow Q$ is incurred by the instantiation of a template class, whose formal parameter has type P and by an actual class parameter, of type Q .

$$(\text{obligation}) \mathcal{C} ::= \text{true} \mid P \Rightarrow Q \mid \mathcal{C} \wedge \mathcal{C}$$

The synthesis of proof obligations pertaining on the correctness of module composition is defined by the relation $\mathcal{E} \vdash \text{mod} : \mathcal{E} / \mathcal{C}$ and by induction on the syntax of modules and declarations. Rule (a) associates the location x with the type name m in the class-field type $[x : m]$. Rule (b) allows to use or open a module m .

$$(a) \mathcal{E} \vdash mx : [x : m] \quad (b) \mathcal{E}[m : \mathcal{T}] \vdash \text{use } m : \mathcal{T}$$

Rule (c) associates a method definition f with the class-field type $[f : \lambda x_{1..k}. xL. \langle P, x^{exit} \rangle]$. Its side-condition $(*)$ is that $\mathcal{L} = \text{labs}(pgm)$ is the set of labels defined in pgm and that $L = \text{start}(pgm)$ is the entry point of pgm . It defines the proposition P and the continuation or output clock x^{exit} of the method f parameterized by its sequence $x_{1..k}$ of input variables, its result variable x , and the label L that defines its input clock. To process the function, we associate its return value, denoted by `return` to a signal x used to carry its value.

$$(c) \frac{L, \mathcal{E}[\text{return} : x] \vdash L : \text{blk}; pgm : P^{(*)}}{\mathcal{E} \vdash m f(x_{1..k}) \{pgm\} : [f : \lambda x_{1..k}. xL. \langle P / \mathcal{L}, x^{exit} \rangle]}$$

Rule (d) sequentially processes the declarations dec in a module. The constraint `true` is omitted in rules (a) and (c).

$$(d) \frac{\mathcal{E} \vdash dec_1 : \mathcal{E}_1 / \mathcal{C}_1 \quad \mathcal{E} \uplus \mathcal{E}_1 \vdash dec_2 : \mathcal{E}_2 / \mathcal{C}_2}{\mathcal{E} \vdash dec_1; dec_2 : \mathcal{E}_1 \uplus \mathcal{E}_2 / \mathcal{C}_1 \wedge \mathcal{C}_2}$$

Class-field declarations contribute to building the type $\mathcal{T}$ of a module. We write $\mathcal{E} \vdash m : \mathcal{T}$ iff $\mathcal{E}$ contains $[m : \mathcal{T}]$. An extension noted $\mathcal{E}_1 \uplus \mathcal{E}_2$ is defined by $\mathcal{E}_2$ and all class names and class-field names of $\mathcal{E}_1$ not overridden

by a declaration in $\mathcal{E}_2$. Rule (α) defines the type $\mathcal{T}$ of a class by that of its field declarations.

$$(\alpha) \frac{\mathcal{E} \vdash dec : \mathcal{T}}{\mathcal{E} \vdash \text{class } m \{dec\} : [m : \mathcal{T}]}$$

Rule (β) defines the type of a template instance $m_1 \langle m_2 \rangle m$. Let $\Lambda(x : \mathcal{T}_1). \mathcal{T}$ be the type of the functor m_1 . Let $\mathcal{T}_2$ be the type of the parameter m_2 . If the subtyping relation $\mathcal{T}_1 \leq \mathcal{T}_2$ implies the proof obligation $\mathcal{C}$ then the type of m is $\mathcal{T}[m_2/x]$ (x is substituted by m_2).

$$(\beta) \frac{\mathcal{E} \vdash m_1 : \Lambda(x : \mathcal{T}_1). \mathcal{T} \quad \mathcal{E} \vdash m_2 : \mathcal{T}_2 \quad \mathcal{T}_1 \leq \mathcal{T}_2 \Rightarrow \mathcal{C}}{\mathcal{E} \vdash m_1 \langle m_2 \rangle m : ([m : \mathcal{T}] / \mathcal{C})[m_2/x]}$$

Rule (γ) defines the type of a template declaration $\langle \text{class } m_1 : n_1 \rangle \text{mod}$. Provided the *assumption* that the formal parameter m_1 of the template has the type $\mathcal{T}_1$ (that of the virtual class name n_1) the template *guarantees* that the module m_2 it defines has type $\mathcal{T}_2$. Hence the type $\Lambda(m_1 : \mathcal{T}_1). \mathcal{T}_2$ for module m_2 .

$$(\gamma) \frac{\mathcal{E} \vdash n_1 : \mathcal{T}_1 \quad \mathcal{E}[m_1 : \mathcal{T}_1] \vdash \text{mod} : [m_2 : \mathcal{T}_2]}{\mathcal{E} \vdash \text{template } \langle \text{class } m_1 : n_1 \rangle \text{mod} : [m_2 : \Lambda(m_1 : \mathcal{T}_1). \mathcal{T}_2]}$$

Rule (δ) processes module declarations in sequence.

$$(\delta) \frac{\mathcal{E} \vdash \text{mod}_1 : \mathcal{E}_1 / \mathcal{C}_1 \quad \mathcal{E} \uplus \mathcal{E}_1 \vdash \text{mod}_2 : \mathcal{E}_2 / \mathcal{C}_2}{\mathcal{E} \vdash \text{mod}_1; \text{mod}_2 : \mathcal{E}_1 \uplus \mathcal{E}_2 / \mathcal{C}_1 \wedge \mathcal{C}_2}$$

Finally, the resolution of the behavioral sub-typing relation $\mathcal{T}_1 \leq \mathcal{T}_2$ is defined by structural induction. It reduces to the proof of a conjunction of propositions $P_1 \Rightarrow P_2$.

7.3. Design checking

Properties pertaining on common design errors can easily be expressed and checked using our type system. Whereas related approaches consist of proposing an ad-hoc type system for analyzing a specific pattern of design errors: race conditions, deadlocks, threads termination; and in a given programming language: Java, C, SystemC, our type system provides a generic framework to perform verification via model checking of behavioral properties of embedded systems described using imperative programming languages.

Termination. A common design error found in embedded system design is the unexpected termination of a thread due to, e.g., an uncaught exception. Here, the termination of a thread f can simply be expressed by the accessibility of the property $f^{exit} = 1$. Unexpected termination can hence be avoided by checking that f satisfies $f^{exit} = 0$.

Deadlocks. Another common design error is a wait statement that does not match a notification and yields the thread to block. Let $L_{1..n}$ be the clocks of the blocks $L_{1..n}$ in which a lock x is notified. Waiting for x at a given label L eventually terminates if P satisfies $L \wedge \neg(\bigwedge_{i=1}^n L_i) = 0$.

Races. Similarly, concurrent write accesses to a variable x shared by parallel threads can be checked exclusive by considering the input clocks $e_{1..n}$ of all write statements $x = f(y, z)$ by verifying that P satisfies $(e_i \wedge (\bigvee_{j \neq i} e_j)) = 0$ for all $0 < i \leq n$.

Larger case-studies reporting applications of our technique in system design and verification are the complete model of a finite input response (FIR) filter starting from the SystemC 2.0.1 distribution ⁽⁵⁾. In this case study, we demonstrate the benefits of modularly associating each System module to a behavioral type interface to perform optimizations and verifications which are modular and yet sensitive to the architecture in which modules or components are placed as reflected by the architecture's behavioral type and by application of an assumption-guarantee reasoning principle. A more recent and larger experiment applies the principles presented in this article to co-modeling by considering predefined SystemC components and connecting them around a bus architecture by giving a synchronous data-flow model of the interconnection wrappers.

8. RELATED WORK

By contrast to traditional type systems, which focus on rendering data-structure abstractions, behavioral type systems ^(10,13) are concerned with the abstraction of control structure in concurrent programs.

A related direction of research is software model checking using popular tools like Bandera ⁽¹⁷⁾, Mops ⁽¹⁸⁾, Verisoft ⁽¹⁹⁾, Modex ^(9,20), Slam ⁽²¹⁾, CBMC ⁽²²⁾, Magic ⁽²³⁾, Blast ⁽²⁴⁾, Pathfinder ⁽²⁵⁾. Most software model checking tools proceed by extracting temporal logic models of source programs (either Java or C but rarely both) and perform sophisticated and efficient abstractions to drastically accelerate property verification.

Our approach contrasts with the software model checking trend in that it is primarily aimed at *modeling* software and then perform either of global model transformations (desynchronization, rethreading, etc) and code generation ⁽¹⁴⁾, conformance checking by finite-flow equivalence using model checking techniques ⁽¹⁵⁾ or modular state-less abstraction for

efficient property verification ⁽⁵⁾. As such, our approach most closely relates to that of Modex ⁽²⁰⁾ in which temporal property models are extracted for later verification with Spin. We experienced that representing such models using executable specifications expressed in a multi-clocked synchronous model offers the additional benefit of operating correctness-preserving model transformations such as protocol synthesis (desynchronization ⁽¹⁵⁾) or static scheduling (rethreading ⁽¹⁴⁾). Finally, and unlike most related approaches in SMC, which are geared towards a particular programming language, we focus on a language-independent intermediate representation of Gnu's GCC.

We share the aim of a scalable and correct-by-construction exploration of abstraction-refinement of system behaviors with the work of Henzinger et al. on interface automata ⁽¹⁾. Our approach primarily differs from interface automata in the data-structure used in the Polychrony workbench: clock equations, boolean propositions and state variable transitions express the multi-clocked synchronous behavior of a system. Compared to an automata-based approach, our declarative approach allows to hierarchically explore abstraction capabilities and to cover design exploration with the methodological notion of refinement along the whole design cycle of the system, ranging from the early requirements specification to the latest sequential and distributed code-generation ⁽⁸⁾.

9. CONCLUSIONS

Our contribution contrasts from related studies by the capability to capture a complete behavioral model of the type-checked system as well as model abstractions expressed at a scalable degree of precision. In our type system, scalability ranges from the capability to express the exact meaning of the program, in order to make structural transformations and optimizations on it (just as in a traditional type system), down to properties expressed by boolean equations between clocks, allowing for a rapid static-checking of design correctness properties. Our system allows for a wide spectrum of design abstraction and refinement patterns to be applied on a model, e.g. abstraction of states by clocks, abstraction of existentially quantified clocks, hierarchic abstraction, in the aim of choosing a better degree of abstraction for faster verification.

The main novelty in our approach is the use of a multi-clocked synchronous formalism to support the construction of a scalable behavioral type inference system for *de facto* standard design and programming languages, and the materialization of a companion refinement-based design methodology imposed through the strong typing policy of a module system, that reduces compositional design correctness verification to the vali-

dation of synthesized proof obligations. The proposed type system allows to capture the behavior of an entire system-level design and to re-factor it, allowing to modularly express a wide spectrum of static and dynamic behavioral properties, and to automatically or manually scale the desired degree of abstraction of these properties for efficient verification. The type system is presented using a generic and language-independent intermediate representation. It operates transformations implemented in the platform Polychrony, to perform refinement-based design exploration. It yields to SAT and model checking verification tools for an efficient verification of expected design properties and an early discovery of design errors.

REFERENCES

1. DE ALFARO, L., HENZINGER, T. A. "Interface theories for component-based design". *International Workshop on Embedded Software*. Lecture Notes in Computer Science v. 2211. Springer-Verlag, 2001.
2. AMAGBEGNON, T. P., BESNARD, L., LE GUERNIC, P. "Implementation of the data-flow synchronous language SIGNAL". In *Conference on Programming Language Design and Implementation*. ACM Press, 1995.
3. BENVENISTE, A., LE GUERNIC, P., JACQUEMOT, C. "Synchronous programming with events and relations: the SIGNAL language and its semantics". In *Science of Computer Programming*, v. 16. Elsevier, 1991.
4. BENVENISTE, A., CASPI, P., LE GUERNIC, P., MARCHAND, H., TALPIN, J.-P., TRIPAKIS, S. "A protocol for loosely time-triggered architectures". In *Embedded Software Conference*. Lecture Notes in Computer Science, Springer Verlag, October 2002.
5. BERNER, D., TALPIN, J.-P., LE GUERNIC, P., SHUKLA, S. K., "Modular design through component abstraction". In *International conference on compilers, architectures and synthesis for embedded systems*. ACM Press, September 2004.
6. CARLONI, L. P., McMILLAN, K. L., SANGIOVANNI-VINCENTELLI, A. L. "Latency-Insensitive Protocols". In *Proceedings of the 11th. International Conference on Computer-Aided Verification*. Lecture notes in computer science v. 1633. Springer Verlag, July 1999.
7. GAMATIÉ, A., GAUTIER, T. "Synchronous modeling of avionics applications using the SIGNAL language". In *Real-time embedded technology and applications symposium*. IEEE Press, 2002.
8. LE GUERNIC, P., TALPIN, J.-P., LE LANN, J.-C. Polychrony for system design. In *Journal of Circuits, Systems and Computers. Special Issue on Application-Specific Hardware Design*. World Scientific, 2002.
9. HOLZMANN G.J. Software Model Checking. In *Journal of Circuits, Systems and Computers. Special Issue on Application-Specific Hardware Design*. World Scientific, 2002.
10. NIELSON, F., NIELSON, H. "Type and Effect Systems: Behaviours for Concurrency". IC Press, 1999.

11. NOVILLO, D. "Tree SSA, a new optimization infrastructure for GCC". GCC developers summit, 2003.
12. PNUELI, A., SHANKAR, N., SINGERMAN, E. Fair synchronous transition systems and their liveness proofs. In *Symposium on Formal Techniques in Real-time and Fault-tolerant Systems*. Lecture Notes in Computer Science v. 1468. Springer, 1998.
13. S. K. RAJAMANI AND J. REHOF. "A behavioral module system for the π -calculus". *Static Analysis Symposium*. Lecture Notes in Computer Science. Springer, July 2001.
14. TALPIN, J.-P., GAMATIÉ, A., LE DEZ, B., BERNER, D., LE GUERNIC, P. Hard real-time implementation of embedded software in JAVA. In *FIDJI'2003*. Lecture Notes in Computer Science, Springer, November 2003.
15. TALPIN, J.-P., LE GUERNIC, P., SHUKLA, S., GUPTA, R., DOUCET, F. "Formal refinement checking in a system-level design methodology". *Fundamenta Informaticae*. IOS Press, August 2004
16. TALPIN, J.-P., BERNER, D., SHUKLA, S. K., LE GUERNIC, P., GUPTA, R. "A behavioral type inference system for compositional system design". *Application of Concurrency to System Design*. IEEE Press, 2004.
17. HATCLIFF, J., DWYER, M. "Using the Bandera Tool Set to Model-check Properties of Concurrent Java Software". *Invited tutorial, conference on concurrency theory*. Lecture Notes in Computer Science V. 2154. Springer, 2001.
18. CHEN, H., DEAN, D., WAGNER, D. "Model Checking One Million Lines of C Code". *Network and Distributed System Security*. ISOC, February 2004.
19. GODEFROID, P. "Software Model Checking: The VeriSoft Approach". *Technical Memorandum ITD-03-44189G*. Bell Labs, March 2003.
20. HOLZMANN, G., SMITH, M. "An automated verification method for distributed systems software based on model extraction". *IEEE Transactions on Software Engineering*, v. 28. IEEE Press, April 2002.
21. BALL, T., COOK, B., DAS, S., RAJAMANI, S. "Refining Approximations in Software Predicate Abstraction". In *Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, v. 2988. Springer, 2004.
22. KROENING, D., CLARKE, E., LERDA, F. "A Tool for Checking ANSI-C Programs". In *Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science, v. 2988. Springer, 2004.
23. CHAKI, S., CLARKE, E., GROCE, A., JHA, S., VEITH, H. "Modular Verification of Software Components in C". In *Transactions on Software Engineering*, v. 30. IEEE Press, JUNE 2004.
24. BEYER, D., CHLIPALA, A., HENZINGER, T. "The Blast query language for software verification". *International Static Analysis Symposium*. Lecture Notes in Computer Science, v. 3148. Springer, 2004.
25. VISSER, W., HAVELUND, K., BRAT, G., PARK, S., LERDA, F. "Model Checking Programs". *Automated Software Engineering Journal*, v. 10, April 2003.