

Design and Verification of SystemC Transaction-Level Models

Ali Habibi, *Member, IEEE*, and Sofiène Tahar, *Member, IEEE*

Abstract—Transaction-level modeling allows exploring several SoC design architectures, leading to better performance and easier verification of the final product. In this paper, we present an approach to design and verify SystemC models at the transaction level. We integrate the verification as part of the design flow where we first model both the design and the properties (written in Property Specification language) in Unified Modeling Language (UML); then, we translate them into an intermediate format modeled with AsmL [language based on Abstract State Machines (ASM)]. The AsmL model is used to generate a finite state machine of the design, including the properties. Checking the correctness of the properties is performed on the fly while generating the state machine. Finally, we translate the verified design to SystemC and map the properties to a set of assertions (as monitors in C#) that can be reused to validate the design at lower levels by simulation. For existing SystemC designs, we propose to translate the code back to AsmL in order to apply the same verification approach. At the SystemC level, we also present a genetic algorithm to enhance the assertions coverage. We will ensure the soundness of our approach by proving the correctness of the SystemC-to-AsmL and AsmL-to-SystemC transformations. We illustrate our approach on two case studies including the PCI bus standard and a master/slave generic architecture from the SystemC library.

Index Terms—SystemC, system-level design, system-level verification, transaction-level modeling.

I. INTRODUCTION

MODELING hardware architectures usually requires pin-level descriptions, typically at the register transfer level (RTL). The design and verification of models at this level requires great effort, and the simulation is tediously slow. Therefore, modeling at a higher abstraction level is unavoidable, considering the growth in complexity and size of the systems. SystemC [27] is a relatively new system-level language which has been proposed to support modeling at transaction level.

From the verification point of view, one needs expressive languages to specify assertions and properties of systems exhibiting complex behavior. The Property Specification language (PSL) [1] from Accellera is meant to address this issue. However, the language by itself is not not enough to improve the design and verification flows.

In this paper, we first provide a methodology to design SystemC transactional models starting from the Unified Modeling Language (UML) level where a more precise specification of

the system and its properties are defined. We used a modified sequence diagram representation to capture more precise properties description. The UML model is then mapped to an Abstract State Machine (ASM), which is a formal specification method for software and hardware systems that has become successful for specifying and verifying complex systems. ASMs provide features to capture the behavioral semantics of programming and modeling languages where large systems are modeled at a high level of abstraction allowing easier validation and verification operations. We wrote the ASM models in the AsmL language [23], which is one of the very latest languages developed for ASM [10] supporting object-oriented modeling at higher level of abstraction in comparison to C++ or Java. Besides, using the AsmL Tester, from an AsmL specification, it is possible to generate finite state machines (FSMs).

To enable the integration of both the model and the properties at the ASM level, we modeled the PSL semantics in AsmL. At this level, it is possible to verify these properties using model checking. For instance, we encode the property's evaluation in every state which enables evaluating its correctness on the fly while executing the FSM generation algorithm (which is part of the AsmL tool). An incorrect property detection stops the reachability algorithms and outputs a subportion from the complete FSM which represents a complete scenario for a counter-example. Eventually, not all of the properties can be verified due to the state explosion problem. For this reason, we complement our verification methodology by integrating the properties as assertion monitors in the final SystemC design. We compile the PSL property (using the AsmL compiler [23]) to C# while we translate the design from AsmL to SystemC. Both codes, SystemC and C#, are then combined to form a single model, thus enabling the verification of the assertions by simulation.

The objective of the verification process is not only to write assertions but to verify them. This latter task is usually performed using test vector generation tools mostly based on random processes. This kind of blind simulation does not guarantee that the assertion will be covered during the test execution. Therefore, it is very important to consider a smarter and more efficient test vector generation approach. To do so, we propose first to use static code analysis to extract a dependency relation between the design inputs and the assertion's variables. This analysis will also define for every input the range of possible values that may trigger the assertion which provides some very useful information to improve the assertion's coverage.

In order to enhance the coverage even more, we also propose to use a genetic algorithm based on a community of random generators having a variety of DNA information [6]. This latter

Manuscript received January 8, 2005. This work was supported by the IEEE. The authors are with the Department of Electrical and Computer Engineering, Concordia University, Montréal, QC H3K 2H7, Canada (e-mail: habibi@ece.concordia.ca).

Digital Object Identifier 10.1109/TVLSI.2005.863187

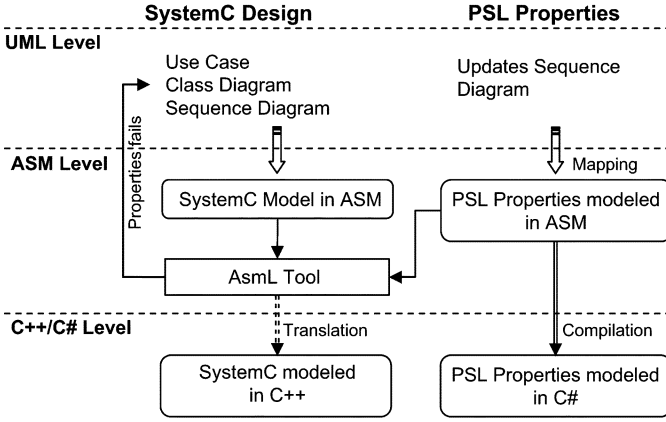


Fig. 1. Design and verification methodology.

will help define the list of variables considered in the test generation, their possible values, and a weighted probability over the previous range [14]. The DNA update/mutation rules will be defined according to the coverage each generator offers. At the end of the genetic procedure, we expect the final DNA to provide an identification of a generator that offers a better coverage than a random one.

The soundness of our approach relies on proving the correctness of the SystemC to AsmL and vice versa transformations. The basic concept of this proof of soundness is based on the systematic design of program transformation frameworks defined in [4]. For instance, we provide a formalization of the SystemC and AsmL semantics in fixpoint based on the OO general case given in [22]. Then, we prove that, for every SystemC (respectively, AsmL) program, there exists an AsmL (respectively, SystemC) program preserving the same properties, with respect to an observation function α_o .

The remainder of this paper is organized as follows. Section II describes the proposed design methodology. Section III discusses the proposed verification approach. Section IV presents the proofs of correctness of our approach. Section V illustrates our approach on two case studies. Section VI discusses the related work. Finally, Section VII concludes the paper.

II. DESIGN METHODOLOGY

Our design methodology, as displayed in Fig. 1, includes two parallel paths concerning the design and its properties. We model the design in the classical way that a C++ design is modeled using UML (i.e., using use cases and class diagrams). Then, we translate the UML model to an ASM model in AsmL in order to perform model checking of certain properties. These latter are obtained from the UML sequence diagram and encoded in the PSL syntax. The verification process ends to: 1) a completion either with a success or failure of the property or 2) a state explosion. Both tasks (UML update and UML-to-AsmL translation) are repeated until all the properties pass (either proved to be correct or do not complete). Then, we compile the PSL properties into a set of C# classes, using the AsmL tool to be used as assertion monitors. The design in AsmL is, from the other side, translated to SystemC and co-integrated with the assertions for verification by simulation.

A. Modeling PSL Properties

PSL is an implementation independent language to define properties. PSL is a hierarchical language, where every layer is built on top of the layer below. This approach allows the expressing of complex properties from simple primitives.

1) *UML Model*: Using UML as a high level of abstraction for design showed a lot of success when applied to Software. Main proposals consider either using UML as new system-level design [5] or as a top layer in combination with existent languages (such as SystemC) [31]. Nevertheless, the proposals completely neglected the properties of the system (PSL-like properties in particular) while sequence diagrams, for example, include very useful information to set transaction properties for TLM in particular [13].

Unfortunately, sequence diagrams do not allow a direct mapping to PSL due to two reasons: 1) the complexity of the PSL property which may include temporal operators and 2) the need for instantiation in PSL. In fact, PSL was defined for real instances formed from objects, from the design, while the sequence diagram considers only classes. For these facts, UML will not present completely and precisely all PSL property. However, it can be used to provide a general skeleton of the property that could be refined and instantiated at the ASM level.

In order to make the UML sequence diagram more adequate for PSL representation, we introduced the following operators.

Clocks: We use the operator to specify the clock that activates the current action.

Number of Cycles: Every action can include the information about after how many cycles the action will start executing (for, e.g., $Mtd[5]()$ says that the action Mtd is executed for exactly five consecutive cycles).

Temporal Operators: These includes operators specifying if, for example, the action will be always executed (A), eventually executed (E), or executed Until a condition is fulfilled (U), etc. These, in fact, represent a mapping to the PSL temporal operators (second layer of PSL).

Sequence Operations: These includes information about the order of executing certain sequences (e.g., *next* and *prev*)

Text Output: This refers to a message that is displayed in case the action fails.

Action Duration: Certain actions are supposed to execute for a certain number of cycles (e.g., reading cycle takes four cycles). We added an operator to specify this information.

"()" Operator: This specifies the set of arguments for a specific action.

Fig. 2 gives an example of a sequence diagram describing a PSL property saying that, if a nonblocking master sends a new request, then, in the next cycle, the arbiter will be notified. After an additional cycle, the arbitration will take place and the master starts sending. The bus is released in the fourth cycle, and a notification will be eventually sent by the slave to the bus who will forward it in the next cycle to the Master.

2) *ASML Model*: There are two ways to embed PSL properties into the design, either as part of the design code itself or by adding them as external monitors. We adopted the second approach, where all of the parameters of PSL properties are defined as objects. The objective of the embedding is to reuse PSL

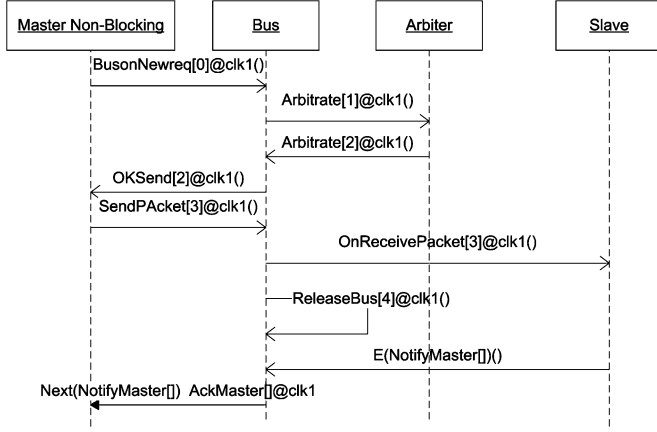


Fig. 2. Example of a modified UML sequence diagram.

properties, as modeled in AsmL, at lower design levels since the AsmL tool can automatically compile them into a C# or .NET code, which can be executed with the concrete SystemC level or as a stand-alone module.

PSL properties are defined in a hierarchical way inspired from the hardware design modular concept. For this reason we defined the embedding in a similar structure, where all the components are defined as objects and every PSL layer *extends* its lower layer using the inheritance feature of AsmL [13].

Boolean Layer: This layer is the basic layer of PSL. Even though it is called the *Boolean layer*, it includes types other than Boolean such as integers and bit vectors. We modeled this layer in AsmL by defining classes for all types and expressions, including their actions. Our embedding is based on the semi-formal semantics presented in the reference manual [1] and the formal semantics definition in HOL [7]. Our embedding of the PSL Boolean layer mainly includes the following.

- 1) *Expression type class* includes the basic five types: *Boolean*, *PSLBit*, *PSLBitVector*, *Numeric*, and *String*. Both *Boolean* and *String* types are directly inherited from the AsmL's *AsmL.Boolean* and *AsmL.String*, respectively.
- 2) *PSL Expressions* includes constructing properties using the implication and equivalence operators.
- 3) *PSL Built Functions* include all the functions defined by PSL to operate at the Boolean layer. We distinguish two actions: a action providing the previous values of a variable (e.g., *prev()*) and an action providing the future values of a variable (e.g., *next()*).

Temporal Layer: PSL supports sequential extended regular expressions (SEREs). The syntax is derived from standard UNIX regular expressions and, hence, the name SERE. The curly braces around the sequence mark the beginning and ending of a SERE. In real life, the delay between two such expressions can be: 1) more than one; 2) a range; and 3) not necessarily occurring in contiguous clock cycles. PSL supports all of these requirements via its repetition operators.

Modeling the SERE feature in AsmL refers to modeling the following.

- 1) *Sequential Expressions*, where a SERE is defined as an AsmL sequence of Boolean. It offers several operations to construct, manipulate, and evaluate the SERE expression.

```

class PSL_SERE
// Memebers
var m_size as Integer = 0 var m_seq as Seq of Boolean
var m_cycle as Seq of Integer var m_actualState as Integer = 0
var m_evaluation as SERE_Evaluation = NOT_STARTED
// Methods
public Evaluate() as SERE_Evaluation
require m_evaluationState = INIT
if(m_seq(m_actualState) = false)
m_evaluation := FAILED return FAILED
else
if m_actualState = m_size
m_actualState := m_actualState + 1 return IN_PROGRESS
else
m_actualState := 0 return SUCCEEDED

```

Fig. 3. Embedding PSL SERE in AsmL.

- 2) *Properties* including the operations necessary to create properties from sequential expressions. It also controls when and how the sequence is to be verified (i.e., the property “verify the sequence is true after n states” is defined as *PSL_Property.EvaluateNext(n)*).

Fig. 3 shows the example of the *PSL_SERE.Evaluate()*, which checks if a sequence is true in a certain path. This action is activated according to an *INIT* signal that must be set by the property.

Verification Layer: This layer is intended to tell the verification tool how to perform the verification process. It allows to construct assertions from properties and to specify relations between them. The embedding mainly includes the following.

- 1) *Verification Directives* to specify how the property will be interpreted (e.g., assertion or requirement). This class extends the *PSL_Property* class from the temporal layer.
- 2) *Verification Unit* is a compact way to include several properties together. The modeled class includes a set operations to add, remove, and update the unit's list of properties.

Modeling Layer: This layer is not used in our verification approach since it is intended for VHDL and Verilog flavors of PSL. Thus, we did not consider it in our embedding.

B. Modeling SystemC

SystemC is built on standard C++. The core language consists of an event-driven simulator as the base. It works with events and processes. The other core language elements consist of modules and ports for representing structures. Interfaces and channels are used to describe communications. SystemC provides data types for hardware modeling and certain types of software programming as well.

1) *AsmL Model:* Our design model at the ASM level is purely OO where every class includes a set of parameters and methods. The particularity of this model resides in the fact that it will be used to generate an FSM using the reachability algorithm part of AsmL tool. Thus, a specific style of programming is required in addition to a precise configuration of the algorithm. This latter generates the FSM by executing the model program in a special execution environment, keeping track of the actions it performs and recording the states it visits. This process is called *exploration*.

```

class PCI_Arbiter
private var m_ActiveMaster as Integer = -1
private var m_req as Boolean = false
private var m_gnt as Boolean = false
public PCI_Arbiter()
public PCI_ArbiterUpdate_m_req()
  require (SystemInit = true) and m_gnt = false and m_req = false
    m_ActiveMaster := min id | id in Masters_Range where
      (MASTERS(id).m_req = true)

  m_req := true

```

Fig. 4. An example of an AsmL model.

The FSM generation algorithm requires as input: domains, methods, actions, and variables (optional inputs are filters, action groups, and properties). The transitions in the FSM are the method calls (including argument values) in the test sequences. The methods in the model program that appear in the transitions are called actions. The states in the FSM are determined by the values of selected variables in the model program. In order for the *exploration* to succeed to generate the FSM, the algorithm requires: 1) initializing all of the model's objects; 2) defining a set of preconditions for every action considered in the exploration process; and 3) providing for every state variable an exploration domain.

For a model M including a set of classes $C = \{c_1, \dots, c_n\}$, where n is the total number of classes in M . For every class c_i in C , we denote its set of methods by c_i^{meth} and the set of members by c_i^{mem} . We defined a set of rules (called R_{FSM}) to guarantee the generation of an FSM representing a portion of the complete system's FSM; these include the following.

Rule R_{FSM}^1 : For every class c_i in C , we have to define a list of instantiations of the class. This ensures that the algorithm will not throw an exception.

Rule R_{FSM}^2 : The first executed method in the design must verify that all of the objects from the class domains were correctly instantiated. This ensures that the algorithm will not misbehave.

Rule R_{FSM}^3 : For every class c_i in C , every method in c_i^{meth} must include a list of *preconditions* to specify when the algorithm considers this method in the exploration process. This ensures that, in every state, we only explore the involved methods.

Rule R_{FSM}^4 : For every class c_i in C , domains for all members in c_i^{mem} must be inherited from AsmL types and restricted to the possible values the system can accept (in particular for inputs). This will allow exploring known types and limits the risks of state explosion.

The optimal scenario is to explore all of the methods and domains in the model; nevertheless, this is not possible all of the time due to the state space explosion. For this reason, working carefully, the domains and the set of actions is the very critical path in the FSM generation process. For illustration purposes, Fig. 4 shows a generic AsmL model with a method including a precondition (denoted by the *require* keyword) setting that the method needs the system to be initialized (*SystemInit = true*) and that it has both variables *m_gnt* and *m_req* set to *false* before it can be executed. Such conditions define strictly at which state the system can execute a particular set of actions.

2) *Translation to SystemC*: Once the AsmL model is verified using the properties describing its behavior, we translate it

to SystemC according to a set of rules to ensure that the final SystemC model preserves the original AsmL code properties. The transformation is purely syntactical, it is performed to certain rules (that we call R_{C++}) that could be summarized in the following.

Rule R_{C++}^1 : “Basic types”: AsmL basic types are all mapped to their equivalent SystemC types (e.g., *Integer* to *int*, *Byte* to *unsigned char*). AsmL includes the same types as C++ which are used for SystemC also.

Rule R_{C++}^2 : “Class Translation”: this includes two separate rules for variables and methods.

Rule $R_{C++}^{2.1}$: “Class Members”: are translated into SystemC signals having the same basic type, e.g., *var m_val as Integer* is translated to *sc_signal<int>m_val*.

Rule $R_{C++}^{2.2}$: “Class Methods”: in AsmL, these contain two parts, with the first one defining the postconditions/preconditions for its execution and the method itself. The first part is integrated in the SystemC module's *constructor*. For instance, a method *Send* defined in AsmL with the following precondition *require clk = true* is inserted into the SystemC module constructor area as “*SC_THREAD(Send); sensitive << clk*”. The method itself is integrated as it is in the SystemC module (we just modify the basic types according to Rule 1).

Rule R_{C++}^3 : “Global Modules” are integrated in the SystemC's main procedure *sc_main*. The naming mapping is used to link different modules together.

III. VERIFICATION METHODOLOGY

The verification process is decomposed into two parts: 1) by model checking at the ASM level and 2) by assertion-based verification at the SystemC (C++)/C# level.

A. Model Checking

PSL properties are modeled in AsmL as specific objects providing a unique view of the property in every system's state (we will refer to these particular object as *A-Property*: AsmL Property). It also simulates the design with the property as a monitor. We build the object starting from basic Boolean components, sequences, and then verification units. We encapsulate sequences in the verification unit as an A-Property which is embedded in the design. Given a set of Boolean items x_1, x_2, \dots, x_n , and y_1, y_2, \dots, y_m belonging to the Boolean layer, and the sequences S_1 and S_2 belonging to the temporal layer, we can define: $S_1 = \{x_1, x_2, \dots, x_n\}$, and $S_2 = \{y_1, y_2, \dots, y_m\}$ and then use a A-Property to check any PSL operation between S_1 and S_2 such as $S_1 \text{ OP } S_2$, where OP is a PSL operator (e.g., implication (\Rightarrow) or equivalence (\Leftrightarrow)). The A-Property is built as follows.

Step 1) Add all of the Boolean items to the sequences: $\forall i$ in 1 to n : $S_1.\text{AddElement}(x_i)$
 $\forall j$ in 1 to m : $S_2.\text{AddElement}(y_j)$

Step 2) Create the property: $P := S_1 \text{ OP } S_2$

Step 3) Define the *verification unit* as an A-Property, A , that includes the property P : $A.\text{Add}(P)$

This property is monitored in every state in the FSM generated by the AsmL tool and is represented by two Boolean state variables P_eval and P_value (saying, respectively, if the property

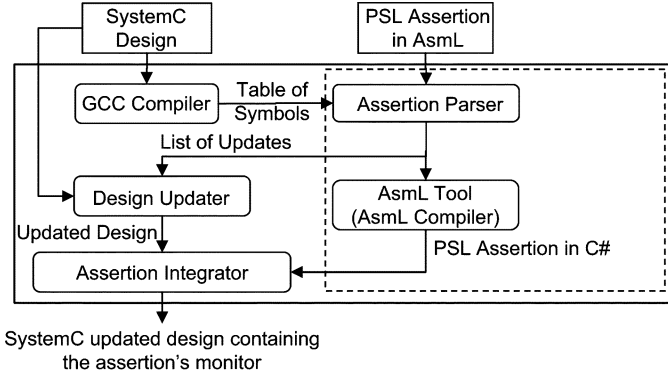


Fig. 5. Methodology to verify PSL assertions for SystemC designs.

can be evaluated and what the value of the property in the current state is). A violated property is detected once $P_{eval} = true$ and $P_{value} = false$. We set the previous condition as filter for the FSM generation algorithm. This way, the generation stops when an error is detected. The generated portion of the state machine, at this point, can be used to identify the problem through a scenario of a counter-example. For multiple properties, the filter is set as a conjunction of all of the conditions for the separate properties. This technique minimizes radically the number of state variables (the FSM size and its generation time). A correct verification process results on the generation of the system's FSM (according the configuration file constraints).

B. Assertion-Based Verification

Fig. 5 describes our methodology to integrate and verify PSL assertions for SystemC designs, which consists of the following three main steps: 1) updating the SystemC design to interface to the assertion monitor; 2) generating the assertion as a C# code from its AsmL description; and 3) integrating the assertion into the design.

Generating the table of symbols from the SystemC design is important in order to validate the variables (names and types) that are used in the assertion. While compiling the assertion, we are concerned with, first, its syntactical correctness, and second, its semantical validity where we check the type and the naming of the assertion variables.

Once the assertion's structure is verified, we translate it to its equivalent AsmL code. In our embedding of the assertion in AsmL, we defined a one-to-one mapping between the PSL assertion and their AsmL embedding. Hence, the transformation is purely syntactical, which guarantees the correctness of the modeled assertion.

In the validation phase of the assertion structure, we also generate a list of updates required to prepare the design to integrate the assertion. For instance, the signals (variables) that are used in the assertion have to be input to the assertion monitor. For this reason, we provide the Design Updater with a list of variables as defined by their unique identifier in the table of symbols. Then, the Design Updater modifies the SystemC design to make the needed variables visible to the monitor. This transformation does not affect the behavior of the code because these variables will be accessed in a read-only mode.

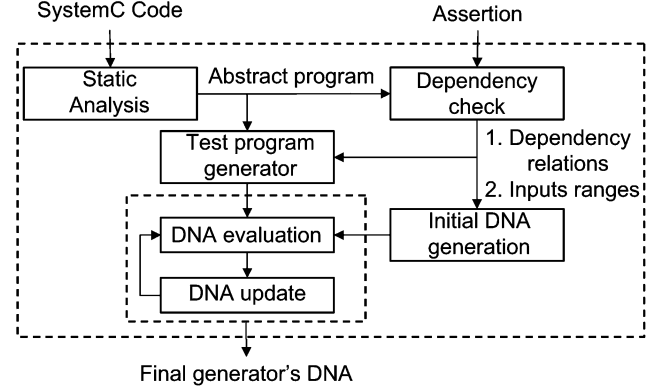


Fig. 6. Enhancing the assertion's coverage.

Once the code is updated and the assertion is generated, the Design Integrator will add the required instantiation of the assertion to bind it to the existing SystemC design modules. The assertion monitor, acting as part of the design, can do the following: 1) stop the simulation when the assertion is fired; 2) write a report about the assertion status and all its variables; and 3) send a warning signal to other modules (if required). We note that the internal code of the assertion is C# so the designer can update it or do any other functionalities that can be coded in C#.

C. Assertions' Coverage Enhancement

Once the assertion integrated in with the design, our next goal is to define a test generation approach that offers better coverage of the assertions. To do so, we first start by statically analyzing the design in order to define a dependency relation between the system inputs and the assertions variables. Such a relation is very useful to omit the inputs that are not affecting the assertion. It serves also to identify the required inputs and the range of their possible values that may affect the assertion. We also identify which processes need to be activated in order to get the assertion fired. Fig. 6 gives an overview of our methodology, including the following steps.

- Step 1) *Static analysis*: We apply a static analysis technique to generate an abstract representation of the design modeled as graph, called *hypergraph* [32], that will include a representation of both the program's environment and the process's environment.
- Step 2) *Dependency check*: From the hypergraph representation, we extract the dependency graph and the range of inputs that may affect the assertion.
- Step 3) *Test program generator*: Using the abstract program (modeled as a hypergraph structure) and the dependency graph, we generate a reduced model containing only the units involved in the assertion.
- Step 4) *Initial DNA generation*: Considering the list of input variables of interest for the assertion and their ranges, we create a DNA structure that will serve as starting point for the genetic algorithm.
- Step 5) *DNA evaluation/update*: Using the initial DNA, the algorithm will update the generators' community starting from the initial DNA to obtain an optimal DNA using the assertion coverage as selection criteria.

1) *Static Code Analysis*: In order to analyze SystemC designs statically and extract the required information to generate the “inputs/assertions variables” dependency relation, we considered an approach based on abstract interpretation [3], which is a formal technique that has proven to be efficient with object-oriented languages and large programs.

At the end of the analysis, the program is represented as a hypergraph [32], which can be interpreted as a general automata connecting its states by branches (also called hyper-branches). These branches can be seen as an extension to binary decision diagrams (BDDs), but more adapted to programs representation. We augmented this work to support the SystemC library and simulator in the form of specific classes to extract information related to SystemC processes and events from the design [17].

2) *Genetic Algorithm*: Genetic algorithms belong to a family of computational models inspired by evolution [20]. They encode a potential solution to a specific problem on a simple chromosome-like data structure and apply recombination operators to these structures to preserve critical information. Since their introduction by Holland [20], genetic algorithms have been applied to a broad range of learning and optimization problems [29]. Typically, a genetic algorithm starts with a random population of encoded candidate solutions (test generators for our case), called chromosomes. The objective is to maximize the likelihood of generating an optimal solution. This can be guaranteed by: 1) evaluating the *fitness* of each candidate solution in the current population; 2) selecting the fittest candidate solutions to act as parents of the next generation of candidate solutions; and 3) recombining and mutating selected parents to generate offsprings.

In our context, the search space to be explored is the state space of the system that may trigger the assertion(s) under verification. Candidate solutions are finite sequences of input ranges and probability weights. Each candidate solution is encoded by a chromosome (a finite string of bits). The information encoded in the DNA includes: 1) the list of input variables; 2) their ranges (possible values); and 3) a weighted probability to their random generation. The algorithm evaluates the fitness of the candidate by executing a test generation based on the information embedded in the corresponding chromosome. A coverage report is then generated to serve in the fitness evaluation phase.

The chromosome encoding is the most important aspect of our algorithm. During the static analysis phase, we obtain the list of variables of the program and their types. Each variable is given a unique identifier. Each type is also given a space of possible values (for example, for the type *char*, the range is [0..255]). The chromosome encodes the list of variables, their types, and a weight relation over the range of possible values. This latter varies according to the type and its interpretation. For every basic type, we defined a list of possible weight relations, e.g., for *Integer*, we use the following window relation: $I < -50$ or $I > 50$ for $w = 0.2$ and $-50 \leq I \leq 50$ for $w = 0.8$.

This relation states that the integer variable I is generated randomly in the interval $[-50, 50]$ with a probability of 80% and 20% inside and outside the interval, respectively.

The proposed fitness function serves to guide the genetic search toward firing the assertion’s variables. Its intuitive idea is to reduce the range of possible values of the input variables

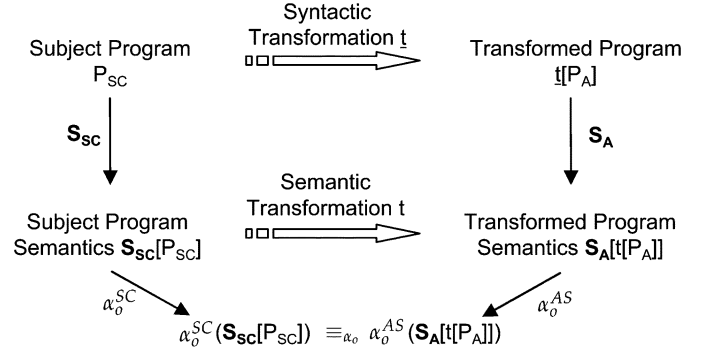


Fig. 7. Online program transformation.

and to find the best probability distribution of the random test generation that will modify the assertion’s variables. This way, we maximize the assertion evaluations, since the evaluation of the chromosomes is defined as an award bonus that is proportional to the number of assertion evaluations. In order to improve the efficiency of the algorithm, we keep track of the best and worst chromosome fitness in each generation; if both fitness values become equal, we increase the mutation rate, in order to help the genetic evolution get out of local maxima. Once there is an improvement in the overall fitness, we restore the original mutation rate to continue the evolution normally.

IV. CORRECTNESS OF THE SYSTEMC/ASML AND ASML/SYSTEMC TRANSFORMATIONS

The work of Patrick and Radhia Cousot in [4] is the essence for any program transformation using abstract interpretation. The tactical choice of using semantics to link the subject program to the transformed program is very smart in the sense that it enables proving the soundness proof of the transformation, related to an observational semantics. The transformation from SystemC to AsmL, and vice versa, represents an online program transformation which corresponds to the approach described in [4, Sec. 3.9]. Fig. 7 displays a projection of that generic methodology on a SystemC subject program and an AsmL-transformed program. The same figure can be used to perform the soundness of a transformation and to construct it. In both cases, we need to define the syntax, semantics, and observation functions for both AsmL and SystemC.

A. Systemc Fixpoint Semantics

1) *Syntactical Domains*: SystemC has a large number of syntactical domains. However, they are all based on the single `SC_Module` domain. Hence, the minimum representation for a general SystemC program is as a set of modules.

Definition 4.1: (SystemC Module: SC_Module): A SystemC Module is a set $\langle DMem, Ports, Chan, Mth, SC_Ctr \rangle$, where $DMem$ is a set of the module data members, $Ports$ is a set of ports, $Chan$ a set of SystemC Chan, Mth is a set of methods (functions) definition, and SC_Ctr the module constructor.

Definition 4.2: (SystemC Port: SC_Port): A SystemC Port is a set $\langle IF, N, SC_In, SC_Out, SC_InOut \rangle$, where IF is a set of the virtual methods declarations, N is the number of interfaces

that may be connected to the port, `SC_In` is an input port (provides only a `Read` method), `SC_Out` is an output port (provides only a `Write` method), and `SC_InOut` is an input/output port (provides `Read` and `Write` methods).

In contrast to default-class constructors for OO languages, the SystemC module constructor `SC_Ctr` contains the information about the processes and threads that will be executed during simulation.

Definition 4.3: (SystemC Constructor: `SC_Ctr`): A SystemC Constructor is a set $\langle \text{Name}, \text{Init}, \text{SC_Pr}, \text{SC_SSt} \rangle$, where `Name` is a string specifying the module name, `Init` is a default class constructor, `SC_Pr` a set of processes, and `SC_SSt` is a set of sensitivity statements (to set the process sensitivity list `SC_SL`).

Definition 4.4: (SystemC Process: `SC_Pr`): A SystemC process is a set $\langle \text{PMth}, \text{PTh}, \text{PCTh} \rangle$, where `PMth` is a method process (defined as a set $\langle \text{Mth}, \text{SC_SL} \rangle$ including the method and its sensitivity list), `PTh` is a thread process (accepts a wait statement in comparison to the method process), and `PCTh` is a clocked thread process (sensitive to the clock event).

Definition 4.5: (SystemC Program: `SC_Pg`): A SystemC program is a set $\langle \text{L_SC_Mod}, \text{SC_main} \rangle$, where `L_SC_Mod` is a set of SystemC modules and `SC_main` is the main function in the program that performs the simulator initialization and contains the modules declarations.

2) *Fixpoint Semantics:* In this section, we define the semantics of the whole SystemC program, $\mathbb{W}[[\text{SC_Pg}]]$, and the SystemC module, $\mathbb{M}_{\text{SC}}[[\text{m_sc}]]$. Then, we present the proofs (or proof sketches) of the soundness and completeness of $\mathbb{M}_{\text{SC}}[[\text{m_sc}]]$.

Definition 4.6: (Delta Delay: δ_d): The SystemC simulator considers two phases *evaluate* and *update*. The separation between these two phases is called *delta delay*.

Definition 4.7: (SystemC Environment: `SC_Env`): The SystemC environment is the summation of the default C++ environment (`Env`) as defined in [22] and the signal environment (`Sig_Store`) specific to SystemC: $\text{SC_Env} = \text{Env} + \text{Sig_Env} = [\text{Var} \rightarrow \text{Addr}] + [\text{SC_Sig} \rightarrow \text{Addr}, \text{Addr}]$, where `Var` is a set of variables, `SC_Sig` is a set of SystemC signals, and $\text{Addr} \subseteq \mathbb{N}$ is a set of addresses.

Definition 4.8: (SystemC Store: `SC_Store`): The SystemC store is the summation of the default C++ store (`Store`) as defined in [22] and the signal store (`Sig_Store`): $\text{SC_Store} = \text{Store} + \text{Sig_Store} = [\text{Addr} \rightarrow \text{Val}] + [(\text{Addr}, \text{Addr}) \rightarrow (\text{Val}, \text{Val})]$, where `Val` is a set of values such that $\text{SC_Env} \subseteq \text{Val}$.

Let $R_0 \in \mathcal{P}(\text{SC_Env} \times \text{SC_Store})$ be a set of initial states, `pc_in` be the entry point of the main function, and $\rightarrow_{\subseteq} : (\text{SC_Env} \times \text{SC_Store}) \times (\text{SC_Env} \times \text{SC_Store})$ be a transition relation.

Definition 4.9: (Whole SystemC Program Semantics: $\mathbb{W}[[\text{SC_Pg}]]$): Let $\text{SC_Pg} = \langle \text{L_SC_Mod}, \text{SC_main} \rangle$ be a SystemC program. Then, the semantics of `SC_Pg`, $\mathbb{W}[[\text{SC_Pg}]] \in \mathcal{P}(\text{SC_Env} \times \text{SC_Store}) \rightarrow \mathcal{P}(\mathcal{T}(\text{SC_Env} \times \text{SC_Store}))$ is

$$\begin{aligned} \mathbb{W}[[\text{SC_Pg}]](R_0) = & \\ \text{lf}_{\text{P}\emptyset} \lambda X. & (R_0) \cup \{ \rho_0 \rightarrow \dots \rho_n \rightarrow \rho_{n+1} \\ & | \rho_{n+1} \in (\text{SC_Env} \times \text{SC_Store}) \wedge \\ & \{ \rho_0 \rightarrow \dots \rho_n \} \in X \wedge \rho_n \rightarrow \rho_{n+1} \}. \end{aligned}$$

Both definitions of the semantics of process declaration ($\mathbb{P}_R[[\text{SC_Pr}]]$) and SystemC module constructor ($\mathbb{P}_{\text{Ctr}}[[\text{SC_Ctr}]]$) are given in [15]. In contrast to the semantics definition of an OO object in [22], a SystemC method can be activated either by the default context or by the SystemC simulator through the sensitivity list of the process. A complete definition of the semantics of a SystemC module object ($\mathbb{O}_{\text{SC}}[[\text{o_sc}]]$) through the definition of a transition function $\text{next}_{\text{SC}}(\sigma) = \text{next}(\sigma) \cup \text{next}_{\text{sig}}(\sigma)$, including both parts C++ related and SystemC specific functions, can be found in [15].

Definition 4.10 (SystemC Module Semantics: $\mathbb{M}_{\text{SC}}[[\text{m_sc}]]$): Let $\text{m_sc} = \langle \text{DMem}, \text{Ports}, \text{Chan}, \text{Mth}, \text{SC_Ctr} \rangle$ be a SystemC module, and then its semantics $\mathbb{M}_{\text{SC}}[[\text{m_sc}]] \in \mathcal{P}(\mathcal{T}(\Sigma))$ can be given as

$$\begin{aligned} \mathbb{M}_{\text{SC}}[[\text{m_sc}]] = & \{ \mathbb{O}_{\text{SC}}[[\text{o_sc}]](v_{\text{sc}}, s_{\text{sc}}) | \text{o_sc} \text{ is an} \\ & \text{instance of m_sc, } v_{\text{sc}} \in \text{D_in}, \\ & s_{\text{sc}} \in \text{SC_Store} \}. \end{aligned}$$

Theorem 4.1: (SystemC Module Semantics in Fixpoint) Let

$$\begin{aligned} G_{\text{sc}}(S) = & \lambda T. \{ \mathbb{S}_0 \langle v, s \rangle | \langle v, s \rangle \in S \} \cup \\ & \{ \sigma_0 \xrightarrow{l_0} \dots \xrightarrow{l_{n-1}} \sigma_n \xrightarrow{l'} \sigma' | \sigma_0 \xrightarrow{l_0} \dots \\ & \xrightarrow{l_{n-1}} \sigma_n \in T, \text{next}_{\text{SC}}(\sigma_n) \ni \langle \sigma', l' \rangle \}. \end{aligned}$$

Then, $\mathbb{M}_{\text{SC}}[[\text{m_sc}]](v_{\text{sc}}, s_{\text{sc}}) = \text{lf}_{\text{P}\emptyset} G_{\text{sc}}(\text{D_in} \times \text{Store})$.

Proof: Although the SystemC model presents some additional functionalities on top of C++, the proof of this theorem is similar to [22, Proof of Theorem 3.2]. For instance, considering the definition of \mathbb{M}_{SC} and applying in order Definition of a SystemC module object in [15], [18, Theorem 2.10], and the fixpoint theorem in [2], the proof is straightforward. ■

The last step in the SystemC fixpoint semantics is to relate the module semantics to the whole SystemC program semantics. Hence, we consider an updated version of the function *abstract* (α°) as defined in [22]. The new function is upgraded to support the SystemC simulation semantics, environment, and store. The complete definitions of α_{SC}° can be found in [15].

Theorem 4.2: (Soundness of $\mathbb{M}_{\text{SC}}[[\text{m_sc}]]$) Let M_{SC} be a whole SystemC program and let $\text{m}_{\text{SC}} \in \mathbb{M}_{\text{SC}}$. Then

$$\begin{aligned} \forall R_0 \in & \text{SC_Env} \times \text{SC_Store}. \\ \forall \tau \in & \mathcal{T}(\text{SC_Env} \times \text{SC_Store}). \\ \tau \in & \mathbb{W}[[\text{SC_Pg}]](R_0) \Rightarrow \\ \exists \tau' \in & \mathbb{M}_{\text{SC}}[[\text{m}_{\text{SC}}]]. \alpha_{\text{SC}}^\circ(\{\tau\}) = \{\tau'\}. \end{aligned}$$

Proof: (Sketch) We have to consider both cases when τ contains an object `oSC`, instantiation of `mSC`, and when it does not include any `oSC`. For the second situation, the proof of the theorem is trivial considering that τ will be an empty trace. In the first case, the trace is not empty (let it be τ''). Since SystemC modules are initialized in the main program `sc_main` before the simulation starts, there exist an initial environment, store and set of variables that define the initial trace $\sigma_0 \in \tau''$. The rest of the traces in τ'' are interaction states of `oSC` because they are obtained by applying α_{SC}° on τ . Therefore, $\tau'' \in \mathbb{M}_{\text{SC}}[[\text{m}_{\text{SC}}]]$. ■

Theorem 4.3: (Completeness of $\mathbb{M}_{\text{SC}}[[\cdot]]$) Let m_{SC} be a SystemC module. Then

$$\begin{aligned} \forall \tau \in \mathcal{T}(\Sigma). \quad & \tau \in \mathbb{M}_{\text{SC}}[[\text{m}_{\text{SC}}]] \Rightarrow \exists \text{SC_P} \in \langle \text{L}_{\text{SC_Pg}} \rangle. \\ & \exists \rho_0 \in \text{SC_Env} \times \text{SC_Store}. \\ & \exists \text{o}_{\text{SC}} \text{ instance of } \text{m}_{\text{SC}}. \\ & \exists \tau' \in \mathcal{T}(\text{SC_Env} \times \text{SC_Store}). \\ & \tau' \in \mathbb{W}[[\rho_0]] \wedge \alpha_{\text{SC}}^{\circ}(\{\tau'\}) = \{\tau\}. \end{aligned}$$

Proof: (Sketch) A SystemC program satisfying the previous theorem can be constructed by creating an instance of m_{SC} in the sc_main function, where the initial state corresponds to the state when the module's constructor SC_Ctr was executed. An execution of a method of m_{SC} corresponds to executing a method thread (setting of the events in its sensitivity list to *Active*), and a change of a port corresponds to updating its internal signal by the new values. Hence, it is always possible to construct both SC_P and ρ_0 . For instance, there exist many other possible constructions involving SystemC threads or clocked threads. ■

B. AsmL Fixpoint Semantics

1) Syntactical Domains:

Definition 4.11 (AsmL Class: AS_C): An AsmL class is a set $\langle \text{AS_DMem}, \text{AS_Mth}, \text{AS_Ctr} \rangle$, where AS_DMem is a set of the module data members, AS_Mth a set of methods (functions) definition and AS_Ctr is the module constructor.

One of the important features that we are going to use in AsmL corresponds to the methods preconditions (the Boolean proposition was verified before the execution of the method).

Definition 4.12 (AsmL Method: AS_Mth): An AsmL method is a set $\langle \text{AS_M}, \text{AS_Pre}, \text{AS_Pos}, \text{AS_Cst} \rangle$, where AS_M is the method's core, AS_Pre is a set of preconditions, AS_Pos is a set of postconditions and AS_Cst is a set of constraints.

Note that AS_Pre , AS_Pos and AS_Cst share the same structure. They are differentiated in the methods by using a specific keyword for each of them (e.g., *require* for preconditions).

Definition 4.13: (AsmL Program: AS_Pg): An AsmL Program is a set $\langle \text{L}_{\text{AS_C}}, \text{INIT} \rangle$, where $\text{m}_{\text{AS_C}}$ is a set of AsmL classes and INIT is the main function in the program.

2) *Fixpoint Semantics:* Similar to the notion of delta delay (δ_d) of SystemC, AsmL considers two phases: *evaluate* and *update*. The program will be always running in the *evaluate* mode except if an update is requested. There are two types of updates, total and partial (usually performed using the Step instruction).

Definition 4.14: (AsmL Environment: AS_Env): The AsmL Environment is a modified OO environment $\text{AS_Env} = [\text{Var} \rightarrow \text{Addr}, \text{Addr}]$, where Var is a set of variables and $\text{Addr} \subseteq \mathbb{N}$ is as set of addresses (two addresses store the current and new values of $v \in \text{Var}$).

Definition 4.15: (AsmL Store: AS_Store): The AsmL store is $\text{AS_Store} = [(\text{Addr}, \text{Addr}) \rightarrow (\text{Val}, \text{Val})]$, where Val is a set of values such that $\text{AS_Env} \subseteq \text{Val}$.

The whole AsmL program semantics ($\mathbb{W}_{\text{AS}}[[\text{AS_Pg}]]$), method semantics ($\mathbb{M}_{\text{AS}}[[\cdot]]$) and object semantics ($\mathbb{O}_{\text{AS}}[[\text{o_AS}]]$) through the definition of a transition function $\text{next}_{\text{AS}}(\sigma)$ can be found in [16]. The AsmL class constructor can be defined according to [22, Definition 3.8].

Definition 4.16: (AsmL Class Semantics: $\mathbb{C}_{\text{AS}}[[\text{c_as}]]$): Let $\text{c_as} = \langle \text{as_dmem}, \text{as_mth}, \text{as_ctr} \rangle$ be an AsmL class, then its semantics $\mathbb{C}_{\text{AS}}[[\text{c_as}]] \in \mathcal{P}(\mathcal{T}(\Sigma))$ is: $\mathbb{C}_{\text{as}}[[\text{c_as}]] = \{\mathbb{O}_{\text{AS}}[[\text{o_as}]](v_{\text{as}}, s_{\text{as}}) | \text{o_as} \text{ is an instance of } \text{c_as}, v_{\text{as}} \in \text{D_in}, s_{\text{as}} \in \text{SC_Store}\}$.

Theorem 4.4: (AsmL Class Semantics in Fixpoint) Let

$$\begin{aligned} \text{H}_{\text{as}}\langle S \rangle &= \lambda T. \quad \{\text{S}_0\langle v, s \rangle | \langle v, s \rangle \in S\} \\ &\cup \{\sigma_0 \xrightarrow{l_0} \dots \xrightarrow{l_{n-1}} \sigma_n \xrightarrow{l'} \sigma' | \sigma_0 \xrightarrow{l_0} \dots \xrightarrow{l_{n-1}} \\ &\sigma_n \in T, \text{next}_{\text{as}}(\sigma_n) \ni \langle \sigma', l' \rangle\}. \end{aligned}$$

Then, $\mathbb{C}_{\text{AS}}[[\text{c_as}]](v_{\text{as}}, s_{\text{as}}) = \text{lfp}_{\emptyset}^{\subseteq} \text{H}_{\text{as}}\langle \text{D}_{\text{in}} \times \text{Store} \rangle$.

Proof: (See [18]) ■

The function $\alpha_{\text{AS}}^{\circ}$ is an updated version of the function *abstract* (α°) defined in [22]. The complete definition of $\alpha_{\text{AS}}^{\circ}$ is given in [16].

Theorem 4.5: (Soundness of $\mathbb{C}_{\text{AS}}[[\text{c_as}]]$) Let P_{AS} be a whole AsmL program and let $\text{c}_{\text{AS}} \in \mathbb{C}_{\text{AS}}$. Then $\forall \text{R}_0 \in \text{AS_Env} \times \text{AS_Store}, \forall \tau \in \mathcal{T}(\text{AS_Env} \times \text{AS_Store})$

$$\begin{aligned} \tau \in \mathbb{W}[[\text{AS_Pg}]](\text{R}_0) : \exists \tau' \in \mathbb{C}_{\text{AS}}[[\text{c}_{\text{AS}}]]. \\ \alpha_{\text{AS}}^{\circ}(\{\tau\}) = \{\tau'\}. \end{aligned}$$

Proof: (See [18]). ■

Theorem 4.6: (Completeness of $\mathbb{C}_{\text{AS}}[[\cdot]]$) Let c_{AS} be a AsmL class. Then

$$\begin{aligned} \forall \tau \in \mathcal{T}(\Sigma). \quad & \tau \in \mathbb{C}_{\text{SC}}[[\text{c}_{\text{SC}}]] : \exists \text{AS_P} \in \langle \text{L}_{\text{AS_Pg}} \rangle. \\ & \exists \rho_0 \in \text{AS_Env} \times \text{AS_Store}. \exists \text{o}_{\text{AS}} \text{ instance} \\ & \text{of } \text{c}_{\text{AS}}. \exists \tau' \in \mathcal{T}(\text{AS_Env} \times \text{AS_Store}). \\ & \tau' \in \mathbb{W}[[\rho_0]] \wedge \alpha_{\text{AS}}^{\circ}(\{\tau'\}) = \{\tau\}. \end{aligned}$$

Proof: (See [18]). ■

C. Program Transformation

The equivalence in behavior, with respect to an observation α_o , between the source SystemC program and the target AsmL program is required to ensure that the soundness of any verification results at the AsmL level. Our objective is to define a relation between the SystemC processes active for certain delta cycles and the set of methods allowed to be executed in the AsmL model. Hence, we will map every thread (method, sensitivity list) in the SystemC design to a method (e.g., method core or preconditions) in the AsmL model.

The SystemC observation function needs to see all the active processes at the beginning of a delta-cycle by checking for the end of the update phase.

Definition 4.17: (SystemC Observation Function: α_o^{SC}): Let $\text{SC_Pg} = \langle \text{L}_{\text{SC_Mod}}, \text{SC_main} \rangle$ be a SystemC program, the observation function $\alpha_o^{\text{SC}} \in \mathcal{P}(\text{SC_Env} \times \text{SC_Store}) \rightarrow \mathcal{P}(\mathcal{T}(\text{SC_Env} \times \text{SC_Store}))$ is

$$\begin{aligned} \alpha_o^{\text{SC}}[[\text{SC_Pg}]](\text{R}_0) &= \\ \text{lfp}_{\emptyset}^{\subseteq} \lambda X. \quad & \text{R}_0 \cup \{\tilde{\rho}_0 \rightarrow \dots \tilde{\rho}_n | \tilde{\rho}_i \in (\text{SC_Env} \times \\ & \text{SC_Store}) \exists \{\rho_0^i \rightarrow \dots \rho_m^i\} \in X \wedge \\ & \rho_m^i \rightarrow \tilde{\rho}_i \wedge \{\text{m_scinM}_{\text{SC}}\} \\ & \exists \text{o_sc} \in \mathbb{M}_{\text{SC.o_sc}}(\rho_m^i()) \neq \{\epsilon\} = \emptyset\}. \end{aligned}$$

In the previous definition, α_o^{SC} is only tracing the initial states of a simulation cycle. For instance, the third condition ensures

that the list of process ready to run is empty. Similarly, we define an observation function α_o^{AS} for an AsmL program.

Definition 4.18: (AsmL Observation Function: α_o^{AS}): Let $\text{AS_Pg} = \langle L_{\text{AS_C}}, \text{INIT} \rangle$ be an AsmL program, the observation function $\alpha_o^{\text{AS}} \in \mathcal{P}(\text{AS_Env} \times \text{AS_Store}) \rightarrow \mathcal{P}(T(\text{AS_Env} \times \text{AS_Store}))$ is

$$\begin{aligned} \alpha_o^{\text{AS}}[[\text{AS_Pg}]](R_0) = \\ \text{If}_{\tilde{\rho}} \lambda X. (R_0) \cup \{ \tilde{\rho}_0 \rightarrow \dots \tilde{\rho}_n \mid \forall \tilde{\rho}_i \in (\text{SC_Env} \times \\ \text{AS_Store}) \exists \{ \rho_0^i \rightarrow \dots \rho_m^i \} \in X \wedge \\ \rho_m^i \rightarrow \tilde{\rho}_i \wedge \{ m_{\text{as}} \text{ in } \mathbb{C}_{\text{AS}} \mid \exists o_{\text{as}} \in \mathbb{C}_{\text{AS}}. \\ o_{\text{as}}(\rho_m^i()) \neq \{\epsilon\} \} = \emptyset \}. \end{aligned}$$

Next, we define the notion of equivalence between the two observations. Although SystemC and AsmL have different environment and store structures, it is possible to ensure that they contain the same information.

Definition 4.19 (Equivalence w.r.t. α_o : \equiv_{α_o}): Let SC_Pg be a SystemC program, V_{sc} a set of its variables, AS_Pg be an AsmL program and Dout_{as} a set of its output variables

$$\begin{aligned} \text{prog}_{\text{sc}} \equiv_{\alpha_o} \text{prog}_{\text{as}} \text{ if} \\ \forall R_0^{\text{SC}} \text{ set of initial states of SC_Pg.} \\ \forall R_0^{\text{AS}} \text{ set of initial states of AS_Pg.} \\ \forall \tilde{\rho} \in \{ \tilde{\rho}_0 \rightarrow \dots \rightarrow \tilde{\rho}_n \} \in \alpha_o^{\text{SC}}[[\text{SC_Pg}]](R_0^{\text{SC}}). \\ \exists \hat{\rho} \in \{ \hat{\rho}_0 \rightarrow \dots \rightarrow \hat{\rho}_n \} \in \alpha_o^{\text{AS}}[[\text{AS_Pg}]](R_0^{\text{AS}}) \\ \mid \forall v_{\text{sc}} \in V_{\text{sc}}. \exists v_{\text{as}} \in V_{\text{as}} \text{ such that} \\ \text{if } v_{\text{sc}} \in \text{SC_Sig} \text{ then } (\tilde{\rho}(v_{\text{sc}}) = (v11, v12)) \wedge \\ (\hat{\rho}(v_{\text{as}}) = (v11, v12)) \\ \text{if } v_{\text{sc}} \in \text{AS_DMem} \text{ then } (\tilde{\rho}(v_{\text{sc}}) = v11) \wedge \\ (\hat{\rho}(v_{\text{as}}) = (v11, v11)). \end{aligned}$$

The observation function ensures that the AsmL program is mimicking the *evaluate* and *update* phases (same length n of the ρ sets). The first if condition takes care of the SystemC signals while the second one concerns basic C++ variables.

Theorem 4.7: (Existence of transformed AsmL program w.r.t. α_o^{SC}) Let SC_Pg be a whole SystemC program, let SC_Din be a set of inputs, and let SC_Dout be a set of outputs. Then, $\exists \text{AS_Pg}$, an AsmL program, such that $\text{SC_Pg} \equiv_{\alpha_o} \text{AS_Pg}$.

Proof: (Sketch) The proof is done by constructing the AsmL program. For instance, for every SystemC module, we affect an AsmL class having the same data members and methods. We set the preconditions AS_Ctr for the AsmL methods as a conjunction of the state of the events present in the sensitivity list, SC_SL , of the SystemC program processes. The tricky point in the construction is when to make the updates in the AsmL program. We have two possibilities: 1) C++ variables update: whenever a C++ variable is involved in an instruction, a partial update can be applied using the notion of *binders* in AsmL and 2) SystemC signals: all signals are updated when all methods preconditions are false. Once the set of AsmL classes are defined, Theorem 4.6 ensures the existence of the AsmL program. ■

Theorem 4.8: (Existence of transformed SystemC program w.r.t. α_o^{A}) Let AS_Pg be a whole AsmL program, let AS_Din be a set of inputs, and let AS_Dout be a set of outputs. Then, $\exists \text{SC_Pg}$, a SystemC program, such that $\text{AS_Pg} \equiv_{\alpha_o} \text{SC_Pg}$.

Proof: (Sketch) Similar to Theorem 4.7, the proof is done by constructing the transformed program (in SystemC for this case). For instance, for every AsmL class, we affect a SystemC module having the same data members and methods. In the SystemC module, for every method, we affect a process method having as sensitivity list SC_SL the list of preconditions AS_Ctr of the corresponding AsmL method. Updates are set whenever a step is found in the AsmL program. Once the set of SystemC classes defined, Theorem 4.3 ensures the existence of the transformed SystemC program. ■

Theorem 4.9: (Soundness of the transformations) Let SC_Pg be a whole SystemC program and let AS_Pg be a whole AsmL program. Then

$$\begin{aligned} \text{SC_Pg} \equiv_{\alpha_o} \text{AS_Pg} : \\ \forall \text{Prop}(V_{\text{sc}}, \tilde{\rho}) \mid \tilde{\rho} \in \alpha_o^{\text{SC}}[[\text{SC_Pg}]] : \\ \text{SC_Pg} \vdash \text{Prop}(V_{\text{sc}}, \tilde{\rho}) : \\ \text{AS_Pg} \vdash \text{Prop}(V_{\text{as}}, \hat{\rho}) \mid \hat{\rho} \in \alpha_o^{\text{AS}}[[\text{AS_Pg}]] \end{aligned}$$

where Prop is a program's property, V_{sc} is a set of variables of the SystemC program, and V_{as} are their corresponding variables in the AsmL program.

Proof: The proof is straightforward from the construction of equivalence relation \equiv_{α_o} in Definition 4.19. ■

V. EXPERIMENTAL RESULTS

In order to illustrate the proposed design and verification methodology, we considered two models: 1) peripheral component interconnect (PCI) [9] local bus standard and (2) an extension of the master/slave bus structure provided by the SystemC distribution [27]. Both models include certain properties, such as liveness, that cannot be verified using simulation, which requires using formal verification techniques such as model checking. Moreover, we aim to evaluate the performance of the overall approach according to the system's size, which can be performed by varying the number of masters and slaves. The experiments were conducted on a Pentium IV processor (2.4 GHz) with 512 MB.

Additional case studies can be found in [12], [28], [19], and [11]. In particular, in [11], we compared the performances of our approach to the RuleBase [21] model checker on a look-aside interface [26]. Experiments showed the performance of our approach to handle cases that classical HDL model checkers cannot support (due to state explosion problem).

A. Models Description

PCI boasts a 32-b-data-path 33-MHz clock speed and a maximum data transfer rate of 132 MB/s. Each PCI master has a pair of arbitration lines that connect it directly to the PCI bus arbiter. In the PCI environment, bus arbitration can take place while another master is still in control of the bus. Data are transferred between an initiator which is the bus master, and a target, which is the bus slave. PCI supports several masters and slaves and allows stopping transactions.

The SystemC master/slave bus represents a more generic bus structure including a set of masters, a set of slaves, an arbiter, and a shared bus. The arbiter is responsible for choosing the appropriate master (if more than one are connected to the bus).

TABLE I
PCI BUS: MODEL CHECKING AND SIMULATION RESULTS

Number of		Model Checking			Simul.
Mast.	Slav.	CPU	Num. of FSM		δ (10^{-9} s)
		Time (s)	Nod.	Tran.	
1	1	2.31	20	25	24.31
1	2	2.93	39	53	29.32
3	1	26.01	236	341	29.76
2	2	26.84	293	449	30.89
2	3	101.37	658	1117	32.74
3	2	574.18	1881	3153	34.03
3	3	6836.01	6346	12097	36.82

Two modes are supported by the bus: 1) *blocking mode*: data are moved through the bus in a burst-mode and 2) *non-blocking mode*: the master reads or writes a single data word.

B. Model Checking

For model checking, we consider, for both models, a set of properties describing all of the possible scenarios of transactions over the bus (e.g., reading, writing, and arbitration). The machine time (user time) needed for verifying the properties depends on the complexity of the original model as well as the property parameters. The CPU time required for the model checking of the properties of the PCI bus for different numbers of masters and slaves is given in Table I. We note that the numbers of states and transitions increase exponentially as a function of the number of masters and slaves connected to the bus, which explains the need for a sharp definition of the exploration domains and active actions. Similar results for the generic master/slave case study are given in Table II, where, for masters, “B” refers to blocking and “NB” to nonblocking.

C. Assertion-Based Verification

The last column in Table I shows a simulation evaluation of the PCI bus when implemented in SystemC including the assertions monitors. We display the average execution time per clock cycle (δ given in nanoseconds) as a function of the number of masters and slaves connected to the bus. This shows a very short time (few seconds) to simulate million of cycles, which offers good coverage for the assertions. In this case also, we obtained similar results for the generic master/slave model, as shown in the last column of Table II.

D. Assertions Coverage Enhancement

In order to evaluate the proposed genetic algorithm, we considered the master/slave bus structure model and a set of ten assertions,¹ Table III compares the assertion coverage results obtained: 1) with a blind random generation; 2) in the initialization phase of the genetic algorithm (GA), i.e., just after the first DNA was generated from the static analysis phase; and 3) after 35 generations of the GA. We used 10^9 simulation cycles for every generation. The coverage is measuring the percentage of test vectors that updated at least one of the variables of the as-

¹For more details, we refer the reader to <http://hvg.ece.concordia.ca/Research/SoC/GeneticAlgo/>.

TABLE II
MASTER/SLAVE BUS: MODEL CHECKING AND SIMULATION RESULTS

Num. of			Model Checking			Simul.
Slv.	Mast.		CPU	Num. of FSM		δ (10^{-9} s)
	B	NB	Time(s)	Nod.	Tran.	
2	1	1	3.54	14	22	27.04
2	3	3	142.32	146	531	31.44
2	3	4	402.32	276	1174	33.02
2	4	4	1192.57	530	2584	35.41
3	1	1	4.32	15	27	28.01
3	3	3	186.64	147	723	36.85
3	3	4	518.73	278	1622	38.82
3	4	4	1541.32	535	3606	40.08
4	1	1	5.21	17	31	29.92
4	3	3	214.46	148	915	39.41
4	3	4	630.48	280	2070	41.11
4	4	4	2002.54	538	4630	43.25

TABLE III
ASSERTIONS’ COVERAGE ANALYSIS

Assert.	Rand. Simul. (%)	Genetic Algorithm		
		GA (%)	after 35 iter.	Avg. Exec per iter. (s)
A1	10	34	92	46.60
A2	8	42	93	55.18
A3	12	32	85	37.19
A4	11	37	89	52.08
A5	14	41	87	31.07
A6	16	46	91	49.22
A7	10	41	94	43.91
A8	17	33	83	39.88
A9	16	31	82	25.91
A10	14	45	97	35.00

sertion. We clearly notice that the static analysis phase already offers a better initial state than starting with totally random generation. The last column in Table III illustrates the average execution time per iteration of the genetic algorithm. For assertion A6, for example, the execution time required to raise the coverage from 16% for the random generation to 91% is 1722.7s.

Fig. 8 gives more details about the evolution of the algorithm for the three assertions (A1, A2, and A3). Typically, a genetic algorithm makes relatively quick progress in the beginning stages of evolution. We noted that there exist some phases, where the algorithm hits local maxima before mutating further, which improves its performance. We even noticed that the coverage sometimes decreases slowly from generation to generation (e.g., generation 20 for A3). This is due to the fact that the evaluation of the assertion is based on weighted random generation. In other terms, since the number of tests is finite, a generator may have two different coverage results for two different test trials.

VI. RELATED WORK

Related work to ours concerns mainly: 1) defining system-level design methodologies using SystemC; 2) writing




Fig. 8. Assertion coverage evolution as function of the population generation.

the formal semantics of PSL, SystemC and AsmL; 3) using genetic algorithms to enhance assertions coverage; and 4) performing program transformation from SystemC to AsmL and vice versa.

Several proposals for system-level design, in particular, [31], used a combination of UML and SystemC for SoC design in general (TLM in particular). We are not aware of any other work that considered ASM as an intermediate layer between UML and SystemC to enable model checking. Besides, we focused on extracting and defining the system properties at the early design stages (from the UML sequence diagram), which makes our study complementary to existing approaches by offering an in-design verification solution.

Genetic algorithms have already been used for a broad range of applications. The most related work to ours is that of Godefroid *et al.* [6], which, in contrast to other approaches, addressed in particular the exploration of large state spaces of concurrent reactive systems as defined for model checking. Nevertheless, this work was restricted to simple Boolean assertions and was based on BDDs, which is not suitable for high-level languages like SystemC. We added to [6] a static analysis phase before applying the genetic algorithm. We also considered a chromosome-encoding based on weighted probability over the space of the possible values of the program variables. We are not aware of any other work where genetic algorithms have been combined with static code analysis to optimize test vector generators in order to improve assertions coverage for SystemC.

In [7], Gordon used the semiformal semantics in the PSL/Sugar documentation to create a deep embedding of the whole language in the HOL theorem prover [8]. The author described how to “execute” the formal semantics of PSL using HOL to see if it is feasible to implement useful tools that work directly from the formal semantics by a mechanized proof. However, he did not provide any framework for the verification of PSL for any implementation language. Besides, he does not offer any approach to reuse the PSL properties as assertion (which is a very important feature in PSL).

Several approaches have been used to write the SystemC semantics (e.g., using ASM in [24]). Denotational semantics [25] are found to be the most effective, since objects can be expressed as fixpoints on suitable domains. Salem in [30] proposed a denotational semantics for SystemC, but his proposal was very shallow, neglecting to relate the semantics of the whole SystemC program to the semantics of its classes. Therefore, in order to

construct a transformation relation between SystemC and AsmL and to prove its soundness, we defined, in this paper, our own SystemC denotational semantics.

Regarding the program transformation, the work of Cousot and Cousot in [4] is the essence for any program transformation using abstract interpretation. We used a projection of that generic approach, described in [4, Sec. 3.9] on a SystemC subject program and an AsmL transformed program can be used to perform the soundness of a transformation and to construct it. In both cases, we defined the syntax, semantics, and observation functions for both AsmL and SystemC.

VII. CONCLUSION

In this paper, we presented a methodology to design and verify SystemC transactional models starting from a UML system specification and integrating an intermediate layer using the AsmL language. We proposed to upgrade the UML sequence diagram in order to capture transaction-related system properties. Then, both the design and its properties are modeled in AsmL to enable performing model checking. On the other hand, to cover for the state explosion problem that may result due to the system’s complexity, we completed our approach by offering a methodology to apply assertion-based verification reusing the already defined PSL properties. To do so, we defined a set of translation rules to transform the design’s model in AsmL to its implementation in SystemC.

In order to efficiently verify assertions in SystemC, we further apply a static code analysis technique based on abstract interpretation. This phase generates an abstracted version of the initial design modeled as a hypergraph that helps defining the dependency between the system inputs and the assertion’s variables, as well as restricting the possible values of the inputs to certain ranges that may update the assertion. Although experiments showed that this approach improves the assertion’s coverage, we proposed to use a genetic algorithm that optimizes the probability distribution of the inputs over the space of their possible values.

We also presented the fixpoint semantics of the SystemC library including, in particular, the semantics of a SystemC Module that we proved to be sound and complete w.r.t. a trace semantics of a SystemC program. We provided also the semantics of a subset of AsmL, and we proved the soundness and completeness of an AsmL class w.r.t. to a trace semantics of the AsmL program. Then, we proved the existence, for every SystemC program, of an AsmL program having similar behavior w.r.t. an observation function that we set to consider the traces of the system just after the update phase of the SystemC simulator. We have used this SystemC to AsmL transformation to reduce the complexity of SystemC models and enabled their formal verification using model checking and theorem-proving approaches used with AsmL and ASM languages in general.

Experimental results showed good model checking results even for complex systems such as the PCI bus standard. The final SystemC models also were running a quite fast simulation, thus enabling to offer better coverage for the whole system state space. Our genetic algorithm showed an improvement of the assertions coverage by a factor of eight in comparison to the

random case. As future work in this direction, we target to optimize the genetic algorithm to improve various coverage metrics.

REFERENCES

- [1] Accellera Property Specification Language Reference Manual (2004). [Online]. Available: <http://www.accellera.org>
- [2] P. Cousot and R. Cousot, "Systematic design of program analysis frameworks," in *Proc. Symp. Principles Programming Languages*, San Antonio, TX, 1979, pp. 269–282.
- [3] —, "Abstract interpretation frameworks," *J. Logic Computat.*, vol. 2, no. 4, pp. 511–547, Aug. 1992.
- [4] —, "Systematic design of program transformation frameworks by abstract interpretation," in *Proc. Symp. Principles Programming Languages*, Portland, OR, Jan. 2002, pp. 178–190.
- [5] R. Damasevicius and V. Stuijks, "Application of UML for hardware design based on design process model," in *Proc. Asia South Pacific Design Automat. Conf.*, 2004, pp. 244–249.
- [6] P. Godefroid and S. Khurshid, "Exploring very large state spaces using genetic algorithms," in *Proc. Tools Algorithms Construction Anal. Syst.*, 2002.
- [7] M. Gordon, "Validating the PSL/Sugar semantics using automated reasoning," in *Proc. Formal Aspects Comput.*, 2003, pp. 306–421.
- [8] M. Gordon and T. Melham, *Introduction to HOL: A Theorem Proving Environment for Higher-Order Logic*. Cambridge, U.K.: Cambridge Univ. Press, 1993.
- [9] PCI Special Interest Group (2004). [Online]. Available: <http://www.pcisig.com/>
- [10] Y. Gurevich, "Evolving algebras 1993: Lipari guide," in *Specification and Validation Methods*. Oxford, U.K.: Oxford Univ. Press, 1995, pp. 9–36.
- [11] A. Habibi, A. I. Ahmed, O. A. Mohamed, and S. Tahar, "On the design and verification of the look-aside interface," in *Proc. Design Automat. Test Europe*, vol. 3, Munich, Germany, 2005, pp. 290–295.
- [12] A. Habibi, A. Gawanmeh, and S. Tahar, "Assertion based verification of PSL for SystemC designs," in *Proc. IEEE Int. Symp. System-on-Chip*, Tampere, Finland, Nov. 2004, pp. 177–180.
- [13] A. Habibi and S. Tahar, "Design for verification of SystemC transaction level models," in *Proc. Design Automat. Test Europe*, vol. 1, Munich, Germany, 2005, pp. 560–565.
- [14] —, "Toward an efficient assertion based verification of SystemC designs," in *Proc. 9th Int. High Level Design Validation Test Workshop*, Sonoma Valley, CA, Nov. 2004, pp. 19–22.
- [15] —, (2004, Dec.) SystemC Fixpoint Semantics. Dept. Elect. Comput. Eng., Concordia Univ.. [Online]. Available: www.ece.concordia.ca/~habibi/techrp/TR0401/
- [16] —, (2004, Dec.) AsmL Fixpoint Semantics. Dept. Elect. Comput. Eng., Concordia Univ.. [Online]. Available: www.ece.concordia.ca/~habibi/techrp/TR0402/
- [17] —, "Abstract Interpretation of SystemC Designs," Dept. Elect. Comput. Eng., Concordia Univ., www.ece.concordia.ca/~habibi/techrp/TR0404/, Tech. Rep., Jun. 2004.
- [18] —, "On the transformation of SystemC to AsmL using abstract interpretation," *Electronic Notes Theoret. Computer Sci.*, vol. 131, pp. 39–49, May 2005.
- [19] A. Habibi, S. Tahar, and L. Halleb, "Formal verification of a bus structure modeled in SystemC," in *Proc. 2nd Annu. IEEE Northeast Workshop Circuits Syst.*, Montreal, QC, Canada, Jun. 2004, pp. 61–64.
- [20] J. Holland, *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: Univ. Michigan Press, 1975.
- [21] *RuleBase Formal Verification Tool (Version 1.5). User's Guide*, May 2003. IBM Haifa Res. Labs..
- [22] F. Logozzo, "Analyse statique modulaire de langages a objets," Ph.D. dissertation, Math. Comput. Sci. Dept., Ecole Polytechnique, Paris, France, June 2004.
- [23] AsmL for Microsoft .NET Framework [Online]. Available: research.microsoft.com
- [24] W. Müller, J. Ruf, and W. Rosenstiel, *SystemC Methodologies and Applications*. Norwell, MA: Kluwer, 2003.
- [25] *Denotational Semantics, Volume B of Handbook of Theoretical Computer Science*, Elsevier Science B.V., 1990, pp. 575–631. P. D. Mosses.
- [26] *Look-Aside (LA-1) Interface, Implementation Agreement, Revision 1.1*, Apr. 2004. Network Processing Forum/Kluwer.
- [27] Open SystemC Initiative (2004). [Online]. Available: <http://www.systemc.org>
- [28] K. Oumalou, A. Habibi, and S. Tahar, "Design for verification of a PCI bus in SystemC," in *Proc. Symp. System-on-Chip*, Tampere, Finland, Nov. 2004, pp. 201–204.
- [29] H. Rudin, "Protocol development success stories: Part I," in *Proc. 12th Int. Symp. Protocol Spec., Testing Verificat.*, vol. C-8, Jun. 1992, pp. 149–160.
- [30] A. Salem, "Formal semantics of synchronous SystemC," in *Proc. Design, Automat. Test Europe Conf.*, Munich, Germany, Mar. 2003, pp. 376–381.
- [31] C. Schulz-Key, M. Winterholer, T. Schweizer, T. Kuhn, and W. Rosenstiel, "Object-oriented modeling and synthesis of SystemC specifications," in *Proc. Conf. Asia South Pacific Design Automation*, 2004, pp. 238–243.
- [32] F. Vederine, "Analyzes totales de programmes par interpretation abstraite," Ph.D. dissertation, Math. Comput. Sci. Dept., Ecole Polytechnique, Paris, France, 2000.



Ali Habibi (S'04–M'06) received the Diploma degree in computer engineering from the École Supérieure des Communication (Sup'Com), Tunis, Tunisia, in 1999, the Diplôme des Études Approfondies (DEA) in communications from École Des Ingénieurs de Tunis (ENIT), in 2001, and the Ph.D. degree in electrical engineering from Concordia University, Montréal, QC, Canada, in 2005.

He is currently a Postdoctoral Fellow with Concordia University. His research interests include system-level verification, formal verification, and microprocessor verification.



Sofiene Tahar (M'96) received the Diploma degree in computer engineering from the University of Darmstadt, Darmstadt, Germany, in 1990 and the Ph.D. degree in computer science from the University of Karlsruhe, Karlsruhe, Germany, in 1994.

From 1995 to 1996, he was a Postdoctoral Fellow with the Université de Montréal, Montréal, QC, Canada. Currently, he is an Associate Professor with the Department of Electrical and Computer Engineering, Concordia University, Montréal, QC, Canada. He has made contributions and published papers in the areas of formal hardware verification, microprocessor verification, modeling and verification of communications architectures and protocols, and mobile CDMA communications. In 2001, he was appointed Concordia Research Chair in Formal Verification of Microelectronics Systems.

Dr. Tahar has been involved in various international conference program committees as well as national research grant selection committees. In 1998, he was the recipient of a Canada Foundation for Innovation (CFI) Researcher Award. He is a Professional Engineer in the Province of Quebec, Canada.