Generation of embedded Hardware/Software from SystemC

Dominique Houzet, Salim Ouadjaout

► To cite this version:

Dominique Houzet, Salim Ouadjaout. Generation of embedded Hardware/Software from SystemC. EURASIP Journal on Embedded Systems, 2006, 2006, pp.ID18526. 10.1155/ES/2006/18526 . hal-00127973

HAL Id: hal-00127973

<https://hal.science/hal-00127973>

Submitted on 5 Oct 2008

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.

Generation of Embedded Hardware/Software From SystemC

Salim Ouadjaout, Dominique Houzet

*Institut of Electronics and Telecommunications of Rennes (IETR) (UMR CNRS 6164), INSA, 20 avenue des Buttes de Coësmes, 35043 Rennes Cedex, France
Email : salim.ouadjaout@insa-rennes.fr, houzet@insa-rennes.fr*

Designers increasingly rely on reusing of Intellectual Property (IP) and on raising the level of abstraction to respect System-on-Chip (SoC) market characteristics. However, most hardware and embedded software codes are recoded manually from system level. This recoding step often results in new coding errors that must be identified and debugged. Thus, shorter time to market requires automation of the system synthesis from high level specifications. In this paper, we propose a design flow intended to reduce the SoC design cost. This design flow unifies hardware and software using a single high level language. It integrates hardware/software (HW/SW) generation tools and an automatic interface synthesis through a custom library of adapters. We have validated our interface synthesis approach on a hardware producer/consumer case study and on the design of a given software radiocommunication application.

Keywords and phrases: System-level synthesis, Hardware-software codesign, Embedded software generation.

1. INTRODUCTION

Technological evolution -particularly shrinking of silicon fabrication geometries- enables the integration of complex platforms in a single System on Chip (SoC). In addition to specific hardware subsystems, a modern SoC can also include sophisticated interconnects and one or several CPU subsystems to execute software. New design flows for SoC design have become essential in order to manage the system complexity in a short time-to-market. These flows include hardware/software (HW/SW) generation tools, the reuse of pre-designed Intellectual Property (IP), and interface synthesis methodologies which are still open problems requiring further research activities [1].

EDA tools propose their own solutions to HW/SW generation. Some use SystemC as a starting point for the hardware design, like Cynthesizer from ForteDesign [2] or Agility Compiler from Celoxica [3]. Several tools use the C language as a starting point for both hardware and software with a custom Application Programming Interface (API) for HW/SW interfaces. It is the case of DK Design Suite from Celoxica [3] with its DSM API and CatapultC from Mentor [4]. In SiliconC [5], structural VHDL is generated for the C functions. Prototypes of the functions become the entities. There are other variants which start from Matlab to produce both hardware and software like SPW from CoWare [6]. Many design methodologies exist for the design of embedded software [7-9]. Some are based on code generated from an abstract model (UML [10]), graphical finite state machine design environments (e.g StateCharts [11]), DSP graphical programming environments (e.g. Ptolemy [8]), or from synchronous programming languages (e.g Esterel [12]). A software generation from a high level model of operating system is proposed by several authors [13-16]. In [15], a software generation from SystemC is based on the redefinition and overloading of SystemC class library elements. In [13], a software-software communication synthesis approach by substituting each SystemC module with an equivalent C struct

is proposed. It requires special SystemC modeling styles (i.e. with macro definitions and preprocessing switches in addition to the original specification code). In [16], software is generated from SpecC with no restrictions on the description of the system model.

Several approaches have been developed to deal with IPs integration. Fast prototyping enables the productive reuse of IPs [17]. It describes how to use an innovative system design flow, that combines different technologies, such as C modelling, emulation, hard Virtual Component reuse and CoWare tools [6]. Prosilog's IP creator, as part of Magillem, aims to improve the integration and re-use of non-VCI compliant IPs by wrapping them into a compatible structure. This tool allows the generation of wrappers from a RTL VHDL description of the IP interface [18]. The Cosy approach is based on the infrastructure and concepts developed in the VCC framework [19]; it defines interfaces at multiple levels of abstraction. Most of those approaches deal with low level protocol adaptation in order to integrate RTL level IPs. A few approaches provide a ready Network on Chip (NoC) to allow easy integration of communication. But these approaches require that the IPs have to be compliant with the NoC interface. Consequently, the designers have to modify the IPs codes.

All these approaches deals with system level synthesis which is widely considered as the solution for closing the productivity gap in system design. System level models are developed for early design exploration. The system specification of an embedded system is made of a hierarchical set of modules (or processes) interconnected by channels. They are described in a system level language as a set of behaviours, channel and interface declarations. Those behaviours mapped onto general or application-specific microprocessors are then implemented as embedded software and hardware. The predominant system level languages are C/C++ extensions [13, 20]. We consider here the SystemC language but an other language can be used. SystemC is mainly used to model and to simulate designs at

system level. However, dedicated powerful hardware description languages like VHDL and Verilog are used for RTL. Embedded software languages like C with static scheduling or POSIX RTOS are used for embedded processors. This leads to a decoupling of behavioral descriptions and implementable descriptions. This decoupling usually requires the recoding of the design from its specification simulation, in order to meet the very different requirements of the final generated code. The recoding step often results in new coding errors that must be identified and debugged. The derivation of embedded software and hardware from system specifications described in a system level language requires to implement all language elements (e.g. modules, processes, channels and port mappings). It is known that SystemC allows the refinement for hardware synthesis, but up to now, SystemC has not been used as an embedded software language. Considering the limited memory space and execution power of embedded processors, the SystemC overhead makes the direct compilation to produce the binary code for target embedded microprocessors highly inefficient. Obviously, it is due to the large SystemC kernel included in the compiled code. This kernel introduces an overhead to support the system level features (e.g. hierarchy, concurrency, communication), but these features are not necessary to the target embedded software code. In addition to direct SystemC compilation inefficiency, some cross compilers for embedded processors may only support the C language. Thus, SystemC has to be translated to C code.

To address system level synthesis, we propose in this paper a top-down methodology. Our challenge is to automate the co-design flow generating the final code for both embedded processors and hardware from a unifying high level language (SystemC). In our methodology, we have developed methods to make the co-design flow smooth, efficient and automated. These methods allow two improvements: a rapid integration of communication and a fast software generation for embedded processors with an efficient interface synthesis. The proposed methodology includes several parsing steps and intermediate models. The first main step is the communication integration based on a custom library of interface adapters that uses the Virtual Component Interface (VCI) standard from VSIA consortium [21]. This library aims to perform the interface synthesis. It allows heterogeneous IPs to communicate in a plug-and-play fashion in the same system. The second main step is the generation of embedded C code from the system specification written in SystemC. Our approach proposes the use of static scheduling and POSIX based RTOS models. It enables also an automatic refinement, while [14] requires its own proprietary simulation engine and needs manual refinement to get the software code. Our method also differs from [13-16] in that our high level SystemC code is translated to a C code with optimized interface synthesis. Optimization is performed according to the processors busses and the NoC as well as according to the SystemC parallel programming model (c.f. 3.4). Other recent propositions have been published in that direction [22].

The paper is organized as follows: in Section 2 we describe the main features of our proposed design flow. The main innovative parts of the design flow are detailed in the next two sections. The first one presents our hardware interface library and our integration methodology of functional IPs, with implementation results from a simple design example. The second one describes the translation process of SystemC elements to C code. This C code targets either a RTOS for dynamic scheduling or a standalone solution with a generated static scheduling. This translation process is validated in Section 5 with implementation results of a producer/consumer and a Code Division Multiple Access (CDMA) radio-communication applications. This work is the result of a project started in 2001 [23][24].

2. DESIGN FLOW

SoC design requires the elaboration and the use of radically new design methodologies. The main parts of a typical system level design flow are: the specification model, the partition into HW/SW elements and the implementation of the models for each element. In Figure 1 we describe the proposed top-down methodology of automatic generation of binary files from SystemC to both embedded software and hardware. The design flow starts with a high level model described in a high level programming language (SystemC). The system is described either through direct programming or through IP reuse. We use Celoxica tools to develop, simulate, analyze and validate the SystemC code (step 1). The first SystemC description is at the functional level. The system is a set of functional IPs including functional models of architectural IPs for fast simulation. The communication between IPs uses SystemC channel mechanisms like `sc_signal` or `sc_fifo` with `read()` and `write()` primitive functions. From the Celoxica graphical tool, we select the IPs which are associated with the hardware side (the architectural IPs substituted by their already VCI-compliant version), and the IPs which are associated with the software side (the monitoring IPs, stimulating IPs, host IPs...). The remaining IPs of the system are targeted to the co-design side, as we need to optimise and well-balance hardware and embedded software to meet several stringent design constraints simultaneously: hard real-time performance, low power consumption and low resources.

Considering the software side (step 2), the SystemC IPs are directly compiled to become binary files targeted to the host processor. This set of software tasks communicate with the

FIGURE 1: Top Down design flow

remaining IPs contained in the FPGA platform through the PCI bus. Because software components run on processors, the SystemC abstract communication needed to describe the interconnection between the software and hardware components is totally different from the existing abstraction of wires between hardware components as well as the function calls abstraction that describes the software communication. In this part, the communication is abstracted as an API which calls PCI bus drivers through an operating system layer. The API hides hardware details such as interrupt controllers or memory and input/output subsystems. We have implemented the Message Passing Interface (MPI-2) library on the host processor and on the embedded processors of our platform [25]. MPI-2 is our HW/SW interface API.

The step 3 is the performing of our SCXML parser tool which allows to convert a given SystemC source code into an XML intermediate representation. The XML format is a subset of the standardized SPIRIT 2.0 format [26]. The system is interpreted as a set of XML files. Each XML file contains the most important characteristics of a SystemC IP, such as:

- name, type and size of each in/out ports, name and type of processes declared in the constructor and also the sensitivity list of each process.
- name and type of IPs building a hierarchical IP, the names of connections between the sub-IPs, and the binding with the IP ports.

Both XML files and profiling reports from Celoxica tool are

treated by our HW/SW partitioning tool (step 4) in order to partition IPs as hardware or software according to the architecture parameters and constraints. After this step we use SynDex tool (step 5) to perform an automatic mapping, routing and static scheduling of IPs on the software and hardware architecture based on a predefined NoC topology [27]. The different SynDex inputs are:

- a hierarchical conditioned data-flow graph of computing operations and input/output operations. The operations are just specified by the type and size of input/output data and execution time of the IPs. The XML files and profiling reports are parsed to produce these inputs. We need also to provide manually information on the non-exclusive execution of IPs in order to help SynDex optimize parallelism.
- specification of the heterogeneous architecture as a graph composed of software processors and hardware processors, interconnected through communication medias. Processors characteristics are: supported tasks, their execution duration, worst case transfer duration for each type of data on the interconnect. The profiling reports and architecture parameters are parsed to produce these inputs.

SynDex implements the IPs onto the multicomponent architecture through a heuristic mapping, routing and scheduling. After the implementation, a timing diagram gives the mapping of the different IPs on the components and the real time predicted behavior of the system. The communication links are represented in order to show all the exchanges

between processors; they are taken into account in the execution time of IPs. The mapping/routing code generated by Syndex tool is then parsed (step 6) in order to manage the NoC configuration and to switch software IPs to the XML/C parser. This parser translates the XML mark-ups to C code with either RTOS calls or a static scheduling provided by SynDex tool. With our SCXML and XML/C parsers, we obtain an embedded C generation tool (SCEmbed) from SystemC. This SCEmbed tool has about 5000 C++ and JAVA code lines. This tool and its XML format can be easily adapted to a different RTOS.

The embedded C code is then treated in step 7 with the Gcc compiler in order to obtain binary executables for the embedded processors. As the C software IPs are mapped on several heterogeneous processors, they need to use a communication library (MPI-2).

In the communication integration (step 8), the identified SystemC hardware IPs are completed with our SystemC VCI adapter library. This point is detailed in the section below. Then point-to-point communication are established between the new VCI-compliant IPs and the VCI hardware IPs through the VCI NoC. We use SynDex configuration information to initialize the VCI adapters, plug the IPs on the NoC, and load the binary code of the software IPs on their corresponding processor memory. Once all the SystemC architecture is produced, we can either simulate it back in the Celoxica tool for evaluation. After validation, we continue with the implementation step.

The last hardware synthesis step plays a very important role in the methodology described above. There have been various research efforts to come up with a good hardware compiler which can generate a synthesizable HDL from high level C/SystemC specifications. The Agility compiler from Celoxica can help the generation of synthesizable VHDL from SystemC. The final product of the design flow is a set of binary files representing programs for the host processor, LEON and Microblaze (Xilinx) processors and FPGAs. These files can be loaded onto the respective components of the prototyping platform (FPGA boards), to build a prototype with a real-time communication system.

3.HW/SW INTERFACE CODESIGN

3.1. Introduction

A SoC can include specific hardware subsystems and one or several CPU subsystems to execute the software tasks. The SoC architecture includes hardware adapters (bridges or communication coprocessors) to connect the CPU subsystems to other subsystems. The HW/SW interface abstraction must hide the CPU. On the software side, the abstraction hides the CPU under a low level software layer ranging from basic drivers and I/O functionality to sophisticated operating system. On the hardware side, the interface abstraction hides CPU bus details through a hardware adaptation layer generally called the CPU interface. This can range from simple registers to sophisticated I/O peripherals including direct memory access

queues and complex data conversion and buffering systems.

3.2. Hardware to Hardware interface synthesis : VCI Adaptation Methodology

We show in Figure 2 the way to establish a communication between IPs with different abstraction levels. We consider here functional IPs and architectural IPs.

FIGURE 2 : VCI connections of non VCI IPs through VCI adapters

The connection can be through wires or through a NoC. The VCI adapters library aims to simplify the (re)use of functional IPs (non VCI compliant) in any SoC based on the VCI protocol. This adapter library is designed in order to change neither the IP cores nor their interface description.

FIGURE 3 : Layers between heterogeneous interfaces of two sets of IPs

The generic architecture shown in Figure 3 helps to clarify the relationship between two hardware IPs connected through a sophisticated VCI NoC. The communication between heterogeneous component interfaces imposes the existence of a wrapper on each side of the communication media (bus or NoC). This wrapper behaves like a bridge which translates the RTL interface between the media and the component. These wrappers (agents) have to be compatible with VCI interface to build a standard media. Thus, an initiator wrapper is connected to VCI initiator ports of a master IP and a target wrapper is connected to VCI target ports of a slave IP.

Considering that these two VCI wrappers are available, the interface synthesis of SystemC functional IPs is a set of steps to replace a primitive channel with a refined channel in order to connect it to the wrappers. A refined channel will often have a more complex interface (e.g VCI) than the primitive channel previously used. The main step in the refining of the interfaces is to create adapters that connect the original modules to the refined channel. Adapters can help to convert the interfaces of the IPs instances into VCI interfaces. The interface refinement

can be made more manageable if new interfaces are developed without making changes to their associated module. The adapter translates the transaction-oriented interface consisting of methods such as *write(data)* into VCI RTL level interface for hardware IPs. Figure 3 depicts the use of adapters to connect functional IPs to the NoC VCI agents. Hook arrow boxes indicate the interface provided by the adapters while the rightleftarrows square boxes represent ports. Our contribution consists in the design of VCI master adapters and VCI slave adapters which manage the VCI initiator and VCI target interfaces respectively. We have chosen a convention that each SystemC output port is an initiating port of transaction and each input port is a target port. Thus, the release of a transaction results in a non blocking write of data on the output port for a *sc_signal* and in a blocking write for a *sc_fifo*. This corresponds to the semantics of the SystemC *sc_signal* and *sc_fifo* primitive channels. Thus, initiating ports of functional IPs are connected to a master adapter and target ports are connected to a slave adapter. In this cas, several IPs may be connected to the same adapter.

The adaptation methodology approach is implemented using a micro-network stack paradigm, which is an adaptation of the OSI protocol stack. Thus the electrical, logical, and functional properties of the interconnection scheme can be abstracted.

3.2.1. Application layer

This layer describes the functional behaviour of a complex system. A system is a set of functional IPs with behavioural models, not architectural IPs such as processors or memories. The communication mechanism is performed with classical *read(data)* and *write(data)* SystemC primitives without additional parameters and no protocol implementation.

3.2.2. VCI adapter layer

The VCI adapter layer is responsible for converting an IP interface towards a lower level interface. A VCI adapter core can manage different ports of different non VCI-compliant IPs. Functional hardware IP ports are implemented as a memory segment accessed through its VCI adapter. They are directly connected to a VCI adapter dedicated to functional hardware IPs with a DMA inside it. The VCI adapter layer is composed of the following sub-layers:

a)Presentation layer: This layer is responsible for translating an abstract data type port towards a SystemC synthesizable data type port.

b)Session layer: The session layer generates a single VCI address between two ports connected to each other in the system level description. This address is divided in two fields, the most significant bits (MSB) identify the destination wrapper, and the least significant bits (LSB) identify the local offset at destination. Each agent of the NoC needs to be configured in order to know the separation position between MSB and LSB, and thus be able to perform address translation to correctly route the data to be sent.

The LSB field is itself divided first according to the target IP port addressed among the different IP ports connected to the same VCI adapter, and second according to the local address

segment managed by the transport layer. VCI adapter address is finally divided in three fields:

- Field-1: Agent number is the address field decoded/generated by the NoC agents and routed in the NoC. Each VCI adapter is connected to a NoC agent and all the NoC agents are numbered from 0 to N.

- Field-2: Port number is the address field decoded/generated by the VCI adapter to switch data to the corresponding IP port.

- Field-3: Word number is the address field decoded/generated by the transport layer. It represents the address in the memory segment of the selected port.

The address translation of each VCI adapter is configured during its connection to the NoC with its NoC agent number and its port number. Already VCI-compliant IPs have to provide configurability of addresses in order to communicate to any IP on the NoC. This configuration of IP VCI adapters is performed during VCI adapter integration step based on Syndex mapping/routing information. For already VCI-compliant IPs, addresses are provided manually as it is IP dependant. This is the second of the very few non fully automated parts of the flow.

c)Transport layer: The basic function of the transport layer is multiple: it accepts data from the IP ports, splits them into smaller units (segments) according to the VCI master adapter data bus size, passes them to the network layer, and ensures that the pieces all arrive correctly at the other end. In addition, the transport layer is responsible of the generation of the segment number which constitutes the third field of VCI address. This layer also resequences and reassembles the messages at the destination (Slave adapter).

d)Network layer: This layer is responsible for the identification of the initiating port. In the case of a multiport master adapter, the network layer launches an arbiter to solve the conflicts and ensures that only one port can have an access to the resource (media). The second treatment is the operation of transfers multiplexing and de-multiplexing. Multiple connections have to be scheduled in time to use the common physical VCI interface. The priority management of the different connections depends on the application constraints, provided statically or dynamically as quality of service requests (QoS).

e)DataLink layer: The data Link layer defines the format of data on the interface and the communication protocol. It is responsible for VCI transactions.

3.2.3. Physical Layer

The physical layer is the physical way of communication. Wires are used for point-to-point connection between VCs. A NoC is used for sophisticated communications.

We have synthesized an example of a simple producer/consumer on the Xilinx FPGA technology. We have

used the PPCI master/slave adapters with an 8-bit data bus and a 5-bit address bus on both IPs. Each adapter unit allows two IP data bus connections of 64-bit and 32-bit size respectively with a static IP port priority management. This implementation was performed with Xilinx Virtex II xc2v3000-6 technology. We present here the post placed/routed results. We have obtained a master adapter cost of 489 units of 4-entries logic and 136 flipflop units, with a 100 MHz clock frequency. So, it occupies 1.7 % of the FPGA. The slave adapter requires 144 4-entries logic units and 204 flipflop units with the same clock frequency. It needs 0.46% of the FPGA resources. A master adapter is four times larger than a slave adapter.

3.3. Software to Software interface synthesis

For embedded software, the SystemC `read(data)` and `write(data)` are implemented with POSIX elements in the case of dynamic scheduling with a RTOS and Message Passing Interface (MPI) elements in the case of static scheduling. We have used the POSIX compliant Real-Time Embedded Multiprocessor Scheduler (RTEMS) as RTOS.

For RTEMS, the read and write primitive functions are replaced with the `rtems_message_queue_receive()` function and the `rtems_message_queue_send()` function respectively. The `sc_fifo blocking read()` function is implemented with the `RTEMS_WAIT` option set in `rtems_message_queue_receive()`. The non blocking `sc_signal` functions are implemented for RTEMS through message queues which are flushed before each data write. The non blocking read is implemented with the option `RTEMS_NO_WAIT`.

For a RTOS-less solution, the SystemC `read(data)` and `write(data)` are implemented as one-sided Remote Memory Access (RMA) with the MPI `MPI_put(data)` primitive only. The blocking mechanism for `sc_fifo` is implemented with the `MPI_wait()` primitive which waits for an acknowledgment.

3.4. Software to Hardware interface synthesis

For software IP on embedded CPUs, communication with the NoC VCI agent is managed with dual-ported memory buffers and DMA from its VCI adapter (dedicated to the CPUs) directly connected to this dual-ported memory. The DMA is controlled by software driver subroutines overloading MPI or RTEMS message queues.

In the case of host processor, the `read(data)` and `write(data)` SystemC primitives are overloaded in order to call the PCI driver services through MPI calls. This software driver configures the hardware DMA which manages the data transactions between host memory and the NoC on the prototyping board through the VCI/PCI bridge.

Using one-sided RMA is an efficient implementation solution of MPI [25][28] and the SystemC programming model is also very well suited to RMA implementation as `sc_signal` reads and writes are not correlated. In practice, efficiency of HW/SW interfaces is obtained with a direct integration of SystemC high level communication library in hardware, that is by a joint optimisation of the implementation of the SystemC programming model with the `MPI_put()` and `MPI_wait()` primitives (RMA model) as well as with the underlying NoC design. The RMA mechanism is limited to write-only transfers

between IPs allowing the design of a specific NoC optimised for those transfers with DMA. This approach is similar to the joint optimisation of compilers and microarchitectures of microprocessors.

We have designed optimised network interfaces for two custom NoC [29] with write-only communications, connected to Microblazes, LEONs and PowerPCs processors through their dedicated ports. The `MPI_put()` primitive needs two I/O access to configure the DMA of the network interface and to launch the DMA transfer in the NoC. Thus the `MPI_put()` takes only 8 processor clock cycles : 6 clock cycles to prepare the DMA configuration and 2 clock cycles for I/O access. In that case the result for the SystemC `sc_signal write()` primitive is 25 clock cycles of overhead comprising two `MPI_put()` executions (one for the control and one for the data), that is 16 clock cycles, and 9 clock cycles to prepare the data to be transferred. Also there is no overhead for the SystemC `sc_signal read()` which is only a local variable access due to the RMA mechanism.

For comparison, the main differences between MPI RMA subset and DSM API from Celoxica presented in Table 1 is that the `MPI_put` is a non blocking mechanism which in conjunction with `MPI_Wait` can implement a blocking mechanism, compared to the `DsmWrite` and `DsmRead` which are only blocking mechanisms. Also the DSM API is a two-sided communication compared to the one-sided RMA subset.

TABLE 1 : DSM AND RMA MPI SUBSET COMPARISON

DSM	MPI
<code>DsmInit()</code>	<code>MPI_Init()</code>
<code>DsmExit()</code>	<code>MPI_Finalize()</code>
<code>DsmWrite()</code> & <code>DsmRead()</code>	<code>MPI_Put()</code> & <code>MPI_Wait()</code>
<code>DsmPortS2HOpen()</code>	--
--	<code>MPI_Barrier()</code>

4. GENERATION OF EMBEDDED C CODE

In modern complex SoCs, the software as an integral part of the SoC is gaining more and more importance. At the system level, the system is composed of a set of hierarchical behaviors connected together through channels. However, for the implementation, many designers use a task-based approach, where the tasks are scheduled by a real time kernel. A whole system design is composed of a set of globally asynchronous/locally synchronous reactive processes that concurrently perform the system functionalities.

Inside the SystemC process code, only `wait()` primitives are allowed and processes lack a sensitivity list except for one signal which is considered as a clock. Therefore, a process will only block when it reaches a `wait()`. These restrictions that we have required are only for the code involved in the embedded HW/SW partitioning process. They help our SCEmbed tool to generate the embedded C code [30]. These restrictions on SystemC coding are also required by Celoxica tools for the SystemC synthesis.

The XML format used by the XML/C parser is easily adaptable for a new target RTOS. The main idea behind is to redefine the SystemC class library elements for the new target RTOS. The original code of these IPs calls the SystemC kernel

functions to support process concurrency and communication. The new code calls the embedded RTOS functions that implement the equivalent functionality. Thus, SystemC kernel functions are replaced either by typical RTOS functions or through direct generation of a statically scheduled code. The functional behavior is not modified during the hardware, software and interfaces generation.

We illustrate the C generation process for the RTOS target with a Producer/Consumer example. The SystemC main code named `sc_main()` is converted to the RTEMS RTOS main code “init”. The channels are implemented with message queues for blocking `sc_fifo` channels and shared variables for non blocking `sc_signal` channels. The clock in the SystemC code is converted into a task sending an event value broadcasted on a message queue. All the tasks read this clock message queue for there synchronization.

	<i>// RTEMS declaration part</i>
<code>int sc_main () {</code>	<code>{ rtems_task init(rtems_task_argument* unused) {</code>
<code>sc_signal<char> medium;</code>	<code>medium= rtems_build_name('m','e','d','i'); { rtems_message_queue_create(medium,..., &mediumID);</code>
<code>sc_clock clock("clock");</code>	<code>Tclk = rtems_build_name('H','L','G','A'); rtems_task_create(Tclk,..., &TclkID); clock= rtems_build_name('c','l','o','c'); rtems_message_queue_create(clock,..., &clockID); Port_clock[0]=clockID;</code>
<code>producer prod_inst("prod");</code>	<code>Tprod= rtems_build_name('p','r','o','d'); rtems_task_create(Tprod,..., &TprodID);</code>
<code>prod_inst.out(medium);</code>	<code>Port_prod_inst[1]=mediumID;</code>
<code>prod_inst.clk(clock);</code>	<code>Port_prod_inst[0]=clockID;</code>
<code>consumer</code>	<code>Tcons= rtems_build_name('c','o','n','s'); rtems_task_create(Tcons,..., &TconsID);</code>
<code>cons_inst("Consumer");</code>	<code>Port_cons_inst[0]=mediumID;</code>
<code>cons_inst.in(medium);</code>	<code>Port_cons_inst[1]=clockID;</code>
<code>cons_inst.clk(clock);</code>	<code>rtems_task_start(TclkID,clock_task, &Port_clock); rtems_task_start(TprodID,producer,&Port_prod_inst); ; rtems_task_start(TconsID,consumer,&Port_cons_inst);</code>
<code>sc_start(-1);</code>	<code>};</code>
<code>return 0;</code>	<code>rtems_task_delete(RTEMS_SELF);</code>
<code>}</code>	<code>}</code>
(a) before	(b) after

FIGURE 4: from SystemC main code to RTEMS code

SystemC concurrent processes need to be converted into RTOS-based tasks. We instantiate the child tasks in a parent one corresponding to the `SC_MODULE` in the system specification. This step is illustrated by our example in Figure 4. The producer and the consumer instances are converted into `Tprod` and `Tcons` parent tasks. In RTEMS, each parent task (`SC_MODULE` in systemC) launches the child tasks (processes in SystemC) and an additional task which is responsible for inter-process communication. This task is created to manage `sc_out` ports writing delay corresponding to the behavioral delay of the SystemC write function (the data are validated after the wait event). The RTEMS equivalent code of the SystemC Producer is shown in Figure 5.

At the system level, synchronization is implemented using channels or SystemC events. During the generation process, the RTOS model provides routines to replace the SystemC synchronization primitives.

```

rtems_task producer(rtems_task_argument
*port) {

Tsend = rtems_build_name('F','C','o','m');
Tmain = rtems_build_name('m','a','i','n');

rtems_task_create(Tsend[0],..., &TsendID);
rtems_task_create(Tmain[1],..., &TmainID);
rtems_task_start(TsendID, ComTask ,&port);
rtems_task_start(TmainID, main, &port);

rtems_task_delete( RTEMS_SELF );
}

rtems_task main(rtems_task_argument *port){
// main code
rtems_task_delete( RTEMS_SELF );
}
// Communication task
rtems_task ComTask (rtems_task_argument
*port)
{
// task code
rtems_task_delete( RTEMS_SELF );
}

```

FIGURE 5: Producer RTEMS code

In the case of POSIX generation, synchronization between tasks is managed by semaphores for `sc_signal` implementation with global shared variables. A special clock management task is generated which schedules the two-step signal assignment process in order to respect the semantic of `sc_signal`. All the signal assignments are performed simultaneously after all the processes are stopped on a `wait()` instruction. The `wait()` instruction is implemented by a semaphore synchronization. The clock task is waiting for all the tasks which are sensitive to the same clock to stop on a wait instruction. Then the second step is performed by this clock management task, which corresponds to the assignment of all the shared global variables with the temporary variables assigned by the different blocked tasks. These blocked tasks are then freed and can read the shared global variables which are now updated. This mechanism is generated for each independent clock in the whole system. When different tasks are mapped on different processors, we assume that they communicate through asynchronous `sc_fifo` channels. Otherwise, the clock management tasks of the different processors have to be synchronized before the assignment of the shared global variables.

The second approach uses a RTOS-less static scheduling. In this solution, the SystemC scheduler is replaced by our custom simulation engine optimized for embedded applications. This scheduler is called from each `wait()` instruction or from `sc_fifo` blocking `read()` or `write()` functions. This scheduler also manages the synchronization of clock sensitive tasks with barrier primitives.

A channel implementation library is provided for all the solutions. Up to now, only primitive channels are available (`sc_signal`, `sc_fifo`). There are three versions of implementation for each channel: SW/SW, HW/HW and SW/HW. SW/SW channels are direct shared variables or message queues implementation. HW/HW channels are RTL level NoC

wrappers. SW/HW are C drivers for embedded processors connected to the NoC.

5.APPLICATION EXAMPLE

In order to evaluate the proposed technique, two designs have been experimented for the SW part in this section. The first one is a simple consumer/producer case with two SystemC components linked together. The second one, a more realistic case, is a CDMA radiocommunication example. The consumer/producer system description has about 86 SystemC code lines and the CDMA system description has about 976 SystemC code lines. the CDMA includes 7 modules with 8 concurrent processes. Both examples have been implemented in a SPARC-based platform that includes 1MB SDRAM and a LEON2 processor synthesized on one 4Mgate Xilinx FPGA with 128 KBytes of RAM. The open source POSIX-compliant RTEMS operating system has been selected as the target embedded RTOS.

The CDMA system has 7 modules: the top (CDMA), one module that generates samples, three modules that compute the QPSK modulation, the THR and the interleaving, one that models the real environment channel behavior by introducing noise, and the last ones that do the reverse treatment that is de-interleaving, ITHR and demodulation. All the modules work in a pipelined dataflow way. Several channel models have been implemented with our design flow. The CDMA application example uses one of them: a non blocking channel (the `sc_signal` channel). The proposed channel models have different implementations depending on the HW/SW partition. Several experiments have been performed with semaphores, mutex condition variables and signals in order to synchronize threads with RTEMS.

Table 2 shows the code size of the different codes on the different operating systems. Table 3 presents there binary size and Table 4 their average execution time per treatment iteration.

TABLE 2 : LINE NUMBER OF PROD/CONS AND CDMA SOURCE CODE

	SystemC Linux	POSIX Linux	RTEMS LEON	POSIX LEON	Static C Linux	Static C LEON
Prod/Cons	86	130	203	161	-	-
CDMA	976	1350	1479	1387	950	950

TABLE 3 : BINARY CODE SIZE OF PROD/CONS AND CDMA

	SystemC Linux	POSIX Linux	RTEMS LEON	POSIX LEON	Static C Linux	Static C LEON
Prod/Cons	592 K	14K	106K	83K	-	-
CDMA	1.8 M	32K	119K	97K	188K	12K

TABLE 4 : EXECUTION TIME OF PROD/CONS AND CDMA

	SystemC Linux	POSIX Linux	RTEMS LEON	POSIX LEON	Static C Linux	Static C LEON
Prod/Cons	43 μ s	81 μ s	2.43 ms	1.85 ms	-	-
CDMA	170 μ s	310 μ s	12.5 ms	9.2 ms	17 μ s	153 μ s

In Table 2, the number of lines of the generated embedded C code is nearly the double for the first simple case which includes 27% of SystemC primitives. For the CDMA, the

generated code size is nearly half more important with only 13% of SystemC primitives. The size of the embedded C generated code is directly linked to the number of SystemC elements included in the original code. As each SystemC primitive is translated with a set of embedded C instructions, a large proportion of `read()`, `write()`, `wait()` and others primitives can result in an important size. However, the increase of the generated code size remains low. Moreover, this generated C code is entirely "readable" and can be completed or optimized manually.

The Table 3, the size of the statically scheduled code for embedded processor is nearly ten times lower than the RTOS one. Thus it is more interesting to use our RTOS-less for embedded processors. For the Linux implementation, the kernel is not included in the code, thus its size is lower than for the standalone one which include its own kernel.

In Table 4, the code execution time with static scheduling is nearly 60 times faster than the RTOS one. We have to consider here that the CDMA application highly communicates and thus highly requests RTOS services with context switching for each communication. Nevertheless we can conclude that RTOS implementation of SystemC elements is not the best solution when static scheduling can be used. Secondly, the validation of the embedded software on a host computer, through direct POSIX execution, obtains comparable execution times compared to SystemC execution. It is thus possible to validate the produced C code before loading it on the target embedded processor. Also we obtain better execution times for a dedicated static scheduling that is nearly ten times faster than pure SystemC execution times. It is thus possible to evaluate more rapidly the whole SystemC model by parsing it in C and execute it instead of using pure SystemC simulations. This is possible only for the functional embedded code.

The results collected in the tables 2, 3 and 4 show the usability of POSIX and static scheduling as embedded C modeling solutions.

We have also experimented different multiprocessor implementations in order to evaluate the impact of HW/SW interface synthesis in term of time overhead. This overhead includes software delays from device drivers and hardware delays due to the NoC crossing. We have experimented several configurations with 1, 2, 4 and 7 processors connected with a one-dimension linear NoC with two processors per node. The speedup obtained is presented on Figure 6. The reduced overhead of software and hardware interfaces (25 clock cycles for a write) combined with the NoC crossing time of almost one clock per NoC node crossed makes the impact of communication low compared to the execution time of the CDMA functions on the different processors. We obtain a speedup of 1.8 with 2 processors and 5 with 7 processors. These results show the low implications of such a higher level interface approach.

FIGURE 6: speedup of CDMA application

6. CONCLUSION

This paper deals with the idea of unifying the use of SystemC to implement both hardware and embedded software. This technique reduces the embedded system design cost with a platform based HW/SW codesign methodology.

The virtual component interfacing is a key aspect of HW/SW codesign. We have shown that it improves significantly the SoC design process by enabling early verification, reusability, and interoperability. We provide VCI adapters intended to tackle a number of technical challenges confronting SoC designers. Each adapter is divided into several layers and may be connected to several ports and/or several IPs. Only transport, network and datalink layers can be implemented as hardware. The presentation and session layers of the VCI adapters are only used for the functional behaviour of the system.

The proposed methodology uses the redefinition of SystemC class library construction elements to generate the embedded software. A first solution is to replace each SystemC element by typical RTOS functions and MPI primitives. A second solution, which is complementary, is to generate a standalone statically scheduled C code which exhibits better results. This method is independent of the selected RTOS and communication API; any of them which is POSIX compliant can be supported by simply adapting the corresponding library. Future works concern the full support of the SPIRIT standard as well as the improvement of Syndex mapping and routing solution.

REFERENCES

- [1] P. Sánchez: "Embedded SW and RTOS," in *E. Villar (Ed.): "Design of HW/SW embedded systems"*. University of Cantabria. 2001.
- [2] Forte Design Systems, "Cynthesizer 3.0", <http://www.forteds.com/>
- [3] celoxica, "Agility compiler user guide", Celoxica, 2005.
- [4] Mentor Graphics, "CatapultC", <http://www.mentor.com/>
- [5] C. Scott Ananian. "SiliconC: A hardware backend for SUIF," <http://flex-compiler.lcs.mit.edu/SiliconC>.
- [6] CoWare, "SPW and Platform Architect", <http://www.coware.com/>
- [7] R. K. Gupta. "Co-synthesis of Hardware and Software for Digital Embedded Systems," *Kluwer*. August 1995.
- [8] J. L. Pino, S. Ha, E. A. Lee, and J. T. Buck. "Software synthesis for dsp using ptolemy," *Journal of VLSI Signal Processing*, 1995.
- [9] F. Baladin, M. Chiodo, P. Giusto, H. Hsieh, A. Jurecska, L. Lavagno, C. Passerone, A. Sangiovanni-Vicentelli, E. Sentovich, K. Suzuki, B. Tabbara. "Hardware-Software Codesign of Embedded Systems: The POLIS Approach," *Kluwer*. 1997.
- [10] Rational. <http://www.rational.com/uml/index.html>.
- [11] D. Harel et al., "Statemate: a working environment for the development of complex reactive systems," *IEEE Trans. on Software Engineering*, April 1990.
- [12] F. Boussinot and R. de Simone. "The ESTEREL Language," *Proceedings of the IEEE*, September 1991.

- [13] T. Grötter, S. Liao, G. Martin, and S. Swan. "System Design with SystemC," *Kluwer Academic Publishers*, 2002.
- [14] D. Desmet, D. Verkest, and H. D. Man. "Operating system based software generation for system-on-chip," *Proceedings of the Design Automation Conference*, June 2000.
- [15] F. Herrera, H. Posadas, P. Sanchez, and E. Villar. "Systematic embedded software generation from systemc," *Proceedings of Design, Automation and Test in Europe*, March 2003.
- [16] H. Yu, R. Dömer, D. Gajski, "Embedded Software Generation from System Level Design Languages," *Proceedings of the Asia and South Pacific Design Automation Conference*, Yokohama, Japan, Jan. 2004.
- [17] F. Pogodalla, R. Hersemeule and P. Coulomb, "FastPrototyping: a system design flow for fast design, prototyping and efficient IP reuse," *Proceedings of Codes*, 1999.
- [18] OCP Adoption Adds Value to Prosilog. www.prosilog.com/Documents/OCP_Newsletter_04-2003.pdf
- [19] J. Y. Brunel et al., "Cosy communication IP," *Proceedings of the Design Automation Conference*, 2000.
- [20] D. Gajski, J. Zhu, R. Dömer, A. Gerstlauer, and S. Zhao. "SpecC: Specification Language and Methodology," *Kluwer Academic Publishers*, January 2000.
- [21] "Virtual Component Interface Standard (OCB 2 1.0)," *VSIA On-Chip Bus Development Working Group*, March 14, 2000.
- [22] W. Klingauf, "Systematic Transaction Level Modeling of Embedded Systems with SystemC", *IEEE DATE05 conference*, 2005.
- [23] S. Ouadjaout, D. Houzet, "Easy SoC Design with VCI SystemC Adapters", *Euromicro Symposium on Digital System Design*, Sep. 2004.
- [24] S. Ouadjaout, D. Houzet, "VCI Interface cosynthesis", *IEEE DELTA02 conference*, New Zealand, Jan. 2002.
- [25] W. Gropp, E. Lusk and R. Thakur, "Using MPI-2 Advanced Features of the Message Passing Interface", MIT Press, 1999.
- [26] SPIRIT Consortium, "SPIRIT V2.0 Alpha release", 2006.
- [27] C. Sorel and Y. Lavarenne, "From Algorithm and Architecture Specifications to Automatic Generation of Distributed Real-Time Executives : a Seamless Flow of Graphs Transformations", In *Formal Methods and Models for Codesign Conference, France*, June 2003.
- [28] S.G. Ziavras et al., "coprocessor design to support MPI primitives in configurable multiprocessors", *Integration, the VLSI journal*, 2006.
- [29] S. Evain, J. P. Diguët, and D. Houzet. 'Spider : A CAD Tool for Efficient NoC Design. In *IEEE NORCHIP 2004*, 8-9 November 2004.
- [30]

Salim Ouadjaout received the M.S. degree in computer science from the National Institute of Computers (INI), Algeria, in 2000, and the M.S. degree from INP, ENSEEIHT, Toulouse, France, in 2001. He is a Ph.D candidate in electrical and computer engineering at the Institute of Electronics and Telecommunication, Rennes, France. He is also working as a research engineer at M3Systems Inc. He has been an ACM Student Member. His research interests include design methodologies, interface synthesis, and micronetworks for SoC.

Dominique Houzet received the M.S. degree in computer sciences in 1989 from Paul Sabatier University, Toulouse, France, and the Ph.D. degree and HDR degree in computer architecture in 1992 and 1999 both from INPT, ENSEEIHT, Toulouse, France. He worked at IRIT Laboratory and ENSEEIHT Engineering School from 1992 to 2002 as an Assistant Professor and also as a Digital Design Consultant with SME and large companies. He is an Associate Professor in the Department of Telecom, the INSA Engineering School, and IETR Laboratory, Rennes, France, since 2002. He has published a number of research papers in the area of parallel computer architecture and SoC design and a book on VHDL principles. His research interests include codesign and SoC design methodologies applied to image processing and radiocommunications. He is a Member of the IEEE Computer Society.