

Research Article

System-on-Chip Environment: A SpecC-Based Framework for Heterogeneous MPSoC Design

Rainer Dömer, Andreas Gerstlauer, Junyu Peng, Dongwan Shin, Lukai Cai, Haobo Yu, Samar Abdi, and Daniel D. Gajski

Center for Embedded Computer Systems, University of California, Irvine, CA 92697-2625, USA

Correspondence should be addressed to Rainer Dömer, doemer@uci.edu

Received 1 October 2007; Revised 4 March 2008; Accepted 10 June 2008

Recommended by Christoph Grimm

The constantly growing complexity of embedded systems is a challenge that drives the development of novel design automation techniques. C-based system-level design addresses the complexity challenge by raising the level of abstraction and integrating the design processes for the heterogeneous system components. In this article, we present a comprehensive design framework, the system-on-chip environment (SCE) which is based on the influential SpecC language and methodology. SCE implements a top-down system design flow based on a specify-explore-refine paradigm with support for heterogeneous target platforms consisting of custom hardware components, embedded software processors, dedicated IP blocks, and complex communication bus architectures. Starting from an abstract specification of the desired system, models at various levels of abstraction are automatically generated through successive step-wise refinement, resulting in a pin-and cycle-accurate system implementation. The seamless integration of automatic model generation, estimation, and verification tools enables rapid design space exploration and efficient MPSoC implementation. Using a large set of industrial-strength examples with a wide range of target architectures, our experimental results demonstrate the effectiveness of our framework and show significant productivity gains in design time.

Copyright © 2008 Rainer Dömer et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

The rising complexity of embedded systems challenges the established design techniques and processes. Novel, nontraditional design approaches become necessary in order to keep up with the increasing demands of higher productivity.

A well-known technique to address the system design challenge is system-level design which raises the level of abstraction, exploits the reuse of intellectual property (IP), and integrates the traditionally separate design processes of the heterogeneous system components. By combining the design flows of hardware units, software processors, third-party IPs, and the interconnecting bus architectures, system-level design emphasizes the system perspective of the overall design task and enables design space exploration across domains. However, successful system design depends on efficient design automation techniques and, in particular, effective tool support.

In this article, we describe the system-on-chip environment (SCE), a system-level design framework based on the

SpecC language and methodology [1]. SCE realizes a top-down refinement-based system design flow with support of heterogeneous target platforms consisting of custom hardware components, embedded software processors, dedicated IP blocks, and complex communication bus architectures.

1.1. SCE methodology

Figure 1 shows the design flow with SCE in an overview. Starting with an abstract specification model in the system design phase, the designer automatically generates transaction level models (TLM) of the design, successively at lower levels of abstraction. Based on component models from the system database and design decisions made by the user, the generated models carry an increasing amount of implementation details.

SCE follows a *specify-explore-refine* methodology [2]. The design process starts from a model specifying the design functionality (*specify*). At each following step, the designer first explores the design space (*explore*) and makes the

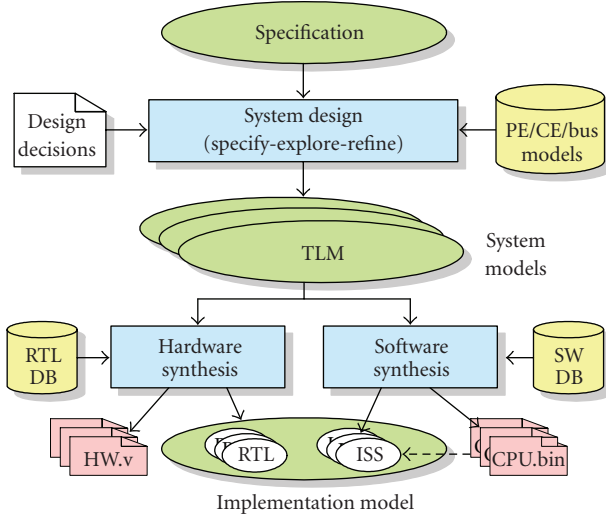


FIGURE 1: System-on-chip environment (SCE) design flow.


necessary design decisions. SCE then automatically generates a new model by integrating the decisions into the previous model (*refine*).

After the system design phase is complete, the hardware and software components in the system model are implemented by the hardware and software synthesis phases, respectively. As a combined result, a pin- and cycle-accurate implementation model is generated. Also, binary images for the software processors, as well as register-transfer level (RTL) descriptions in Verilog for the hardware blocks, are created for further synthesis and manufacturing of the intended Multiprocessor system-on-chip (MPSoC).


Three design models used in the SCE design flow are shown in more detail in Figure 2. In these and later figures that describe design models, we use the graphical notation introduced with the SpecC language [1]. In general, rectangular boxes represent components, and interconnections are indicated by lines (wires) and arrows (busses). Encapsulated computational blocks, called behaviors, are shown as rectangular boxes with round corners, whereas high-level communication is indicated by channels (ellipses) and interfaces (half circles).

Figure 2(a) depicts a simple generic specification model. The model consists of a hierarchy of five behaviors and four communication channels. Except for the system functionality, this model is free of any implementation details. During the system design phase, it will be mapped to a platform architecture (see Section 3.1) and single-threaded processing elements (PEs) will be scheduled (Section 3.2). Communication elements (CEs), such as bus bridges and transducers, and system busses will be added to the model as well (Sections 3.3 and 3.4).

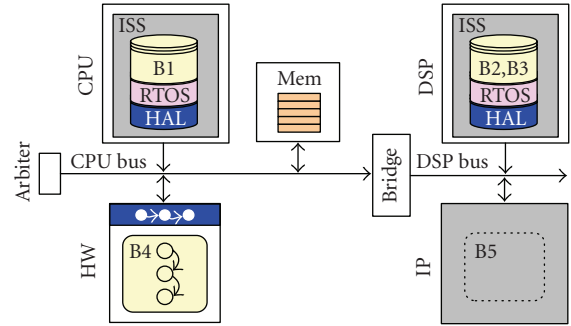
As a result of each of these model refinement steps, a TLM is generated, as shown in Figure 2(b). Depending on the number of implementation decisions taken, the TLM accurately reflects the number and type of PEs in the architecture, the mapping of behaviors to the PEs, and the



(a) Specification model



(b) Transaction-level model (TLM)



(c) Implementation model

FIGURE 2: Generic SCE design models.

mapping of channels to the system busses and CEs. Note that the communication in this model is still at the abstract transaction level.

After hardware and software synthesis (see Sections 3.5 and 3.6, resp.), a cycle-accurate implementation model is generated, as illustrated in Figure 2(c). In this model, embedded software is represented in detailed layers, including the real-time operating system (RTOS) and the hardware abstraction layer (HAL). Custom hardware blocks, on the other hand, are represented accurately by RTL finite state machine (FSM) models. Finally, system communication is also refined down to a pin- and cycle-accurate level.

1.2. Related work

Traditionally, system design is dominated by simulation-centric approaches with horizontal integration of models at specific levels of abstraction. Approaches range from the cosimulation of different low-level languages [3–5] to the combination of heterogeneous models of computation in a common simulation environment [6]. In between, C-based system-level design languages (SLDLs), such as SystemC [7]

and Handel-C [8], emerged as vehicles for transaction-level modeling (TLM) [9]. Most cases, however, are limited to simulation only and lack vertical integration with synthesis flows that provide a path to implementation.

The first attempts at providing system design environments were approaches for hardware/software codesign. Examples of such environments include COSYMA [10], COSMOS [11], and POLIS [12]. These approaches, however, are based on architecture templates consisting of a single microcontroller assisted by a custom hardware coprocessor, and are thus limited to narrow target architectures.

More recently, design environments emerged that provide support for more complex multiprocessor systems. The OCAPI system [13, 14] is based on an object-oriented modeling of designs using a C++ class library and focuses on reconfigurable hardware devices (FPGAs). The OSSS methodology [15] defines an automated system design flow from a cycle-accurate specification written in an object-oriented variant of SystemC. Supporting architecture exploration and automated refinement via intermediate design models, OSSS feeds into the FOSSY synthesis tool for implementation in hardware and software.

Around the TLM concept, several SystemC-based approaches exist that deal with assembly, validation and to some extent automatic generation of communication [16–20]. Metropolis [21, 22] is a modeling and simulation environment based on the platform-based design paradigm. The key idea is to separate function, architecture, and model of computation into separate models. Although Metropolis allows cosimulation of heterogenous PEs as well as different models of computation, a refinement or verification flow between different abstraction levels has not emerged. None of the above frameworks provides a comprehensive, automated approach for the design of complete MPSoCs from abstract specification down to final implementation.

SCE was built on experiences obtained from its predecessor, SpecSyn [2]. While SpecSyn was based on the SpecCharts language, an extension of VHDL, SCE is based on SpecC, which extends ANSI-C for hardware and system modeling.

With respect to our previous publications (previous publications focus on point-tools within the SCE environment and are referenced where applicable), this article is the first comprehensive, cohesive, and complete description of the SCE framework. In other words, for the first time we describe the entire SpecC methodology as implemented by a real working environment. As such, this article focuses on the integration of the tools (including scripting facilities, file formats, and annotations) that realize an efficient top-down system design flow, all the way from an abstract system specification down to a pin- and cycle-accurate implementation. We also list the design decisions taken at each step and thus provide a complete picture of the input the system designer needs to provide based on his application knowledge and design experience. We also demonstrate the effectiveness of the SCE framework and the complete design flow using the combined results of six design experiments using real-world examples. Furthermore, this article describes for the first time the integration of verification tools into the design flow and framework.

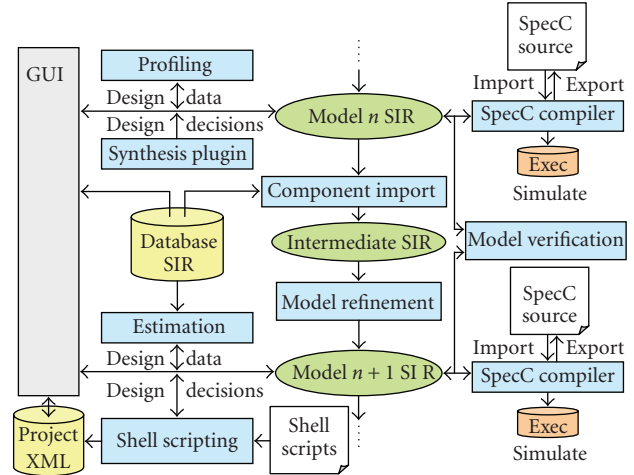


FIGURE 3: SCE software architecture.

2. SCE ARCHITECTURE

SCE is based on the separation of design tasks into two distinct steps: decision making and model refinement. Model refinement takes design decisions and generates a new model of the design reflecting and implementing the decisions.

In SCE, model refinement is automated. Decisions, on the other hand, can be entered manually or through a tool box of automated synthesis algorithms. Together, SCE supports an interactive and automated system design process. Automatic model generation removes the need for error-prone and tedious model rewriting. Instead, designers can focus on design exploration and decision making.

Figure 3 shows the generic software architecture for each task in the SCE design and refinement flow. In each step, design decisions are entered by the user through a graphical user interface (GUI), via a command-line scripting and shell interface, or with the help of automated synthesis plugins implementing optimizing algorithms. Based on the design decisions, a refinement process generates a new design model from the input model automatically.

Overall, the SCE framework is formed by the combination of point tools. These tools exchange information through command line interfaces and design models. In general, all tools operate on a given design model. Design decisions, profiling data, and meta-information about the design are stored as annotations attached to the corresponding objects in the design and database models. All models and databases in SCE are described and captured in the form of SpecC internal representation (SIR) files. Using the SpecC compiler (*scc*), SCE models and databases can be imported from and exported into source files in standard SpecC language format at any time.

2.1. Graphical user interface

The main interface between the designer and the tools is the *sce* GUI [23] which provides various displays and dialogs for browsing of design models and databases, interactive

decision entry, and graphical analysis of profiling and estimation results. Furthermore, it includes menus and tool bars to trigger simulation, profiling, refinement, synthesis, and verification actions. For each action, specific command-line tools are called and executed as needed where the GUI supplies the necessary parameters, captures the output and handles (normal or abnormal) results.

In each session, multiple candidate designs and models can be explored and generated. Information about design models and their relationships, including project-specific compiler and simulator parameters, are tracked by the GUI and can be stored in project files in a custom XML format, allowing for persistent storage, documentation, and exchange of meta-information about the exploration process.

2.2. Simulation and profiling

All design models in the SCE flow are executable for validation through simulation. Using the SpecC compiler and simulator, models can be compiled and executed at any time. SCE also includes profiling tools to obtain feedback about design quality metrics. Based on a combination of static and dynamic analysis, a retargetable profiler (`scprof`) provides a variety of metrics across various levels of abstraction [24]. Initial dynamic profiling derives design characteristics through simulation of the input model. The system designer chooses a set of target PEs, CEs, and busses from the database, and the tool then combines the obtained profiles with the characteristics of the selected components. Thus, SCE profiling is retargetable for static estimation of complete system designs in linear time without the need for time consuming resimulation or reprofiling.

The profiling results can also be back-annotated into the output model through refinement. By simulating the refined model, accurate feedback about implementation effects can then be obtained before entering the next design stage.

Since the system is only simulated once during the exploration process, the approach is fast yet accurate enough to make high-level decisions, since both static and dynamic effects are captured. Furthermore, the profiler supports multilevel, multimetric estimation by providing relevant design quality metrics for each stage of the design process. Therefore, profiling guides the user in the design process and enables rapid and early design space exploration.

2.3. Verification

SCE also integrates a formal verification tool `scver`. Our equivalence verification technology is based on *model algebra* [25], which is a formalism for symbolic representation and transformation of system level models. The formalism itself consists of a set of objects and composition rules. The objects are behaviors, synchronization channels, variables, and ports. The composition rules for control flow, blocking, and nonblocking communication, and hierarchy allow creation of formal models. Functionality preserving transformation rules are also defined on model algebraic expressions. Each of these transformation rules are proven sound with respect to a trace-based notion of functional equivalence.

The incorporation of model algebra-based verification in SCE follows the refinement flow. Well-formed models in SpecC can easily be translated to respective model algebraic expressions. The system designer simply selects an original and a refined model and invokes the verification tool. `scver` then converts the models and applies the transformation rules to derive the refined model from the original model. The two models are equivalent by virtue of the soundness of the transformation rules. The original model is then checked for isomorphism against the derived model and the differences, if any, are reported. It must be noted that the number and order of transformation rules used for the model derivation step depend on the type of refinement. Since the key concept in SCE is the well-defined semantics of models at different abstraction levels, the order of transformation rules can be easily established. Therefore, equivalence verification becomes not only tractable, but straightforward.

2.4. Databases

In the SCE design flow, the system is gradually refined using system components from a set of databases [26]. Specifically, SCE includes databases for processing elements (PEs), communication elements (CEs), operating system models, bus or other communication protocols, RTL units and software components. The database components are described as SpecC objects (behaviors or channels). The SpecC hierarchy for a component object in the database defines its structure and functionality for simulation and synthesis. In addition, metadata, such as attributes, parameters, and general information, is stored in the form of annotations attached to the components.

2.5. Scripting interface

SCE supports scripting of the complete environment from the command line without the need for the GUI. For scripting purposes, a GUI-less command shell, `scsh`, of SCE is available. The SCE shell is based on the same libraries as the SCE GUI (not including the GUI layer itself) and offers interactive command-prompt based- or automatic script-based execution.

The SCE shell is based on an embedded Python interpreter that is extended with an API for low-level access to SCE core functionality and internals. For user-level scripting, a complete set of high-level tools on top of the SCE shell are available. Provided scripts include command-line utilities for component allocation (`sce_allocate`), mapping/partitioning (`sce_map`), scheduling (`sce_schedule`), connectivity definition (`sce_connect`), component import (`sce_import`), and project handling (`sce_project`). These scripts provide a convenient command-line interface for all SCE high-level functionality and decision entry. Together with command-line interfaces to refinement tools and the compiler, a complete scripting of the SCE design flow, through shell scripts or via Makefiles, is available.

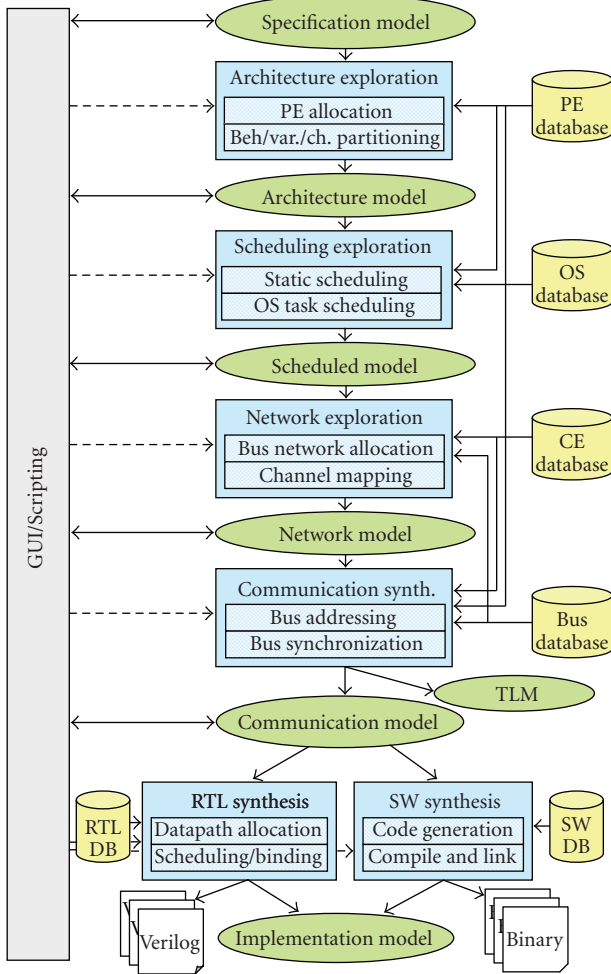


FIGURE 4: Refinement-based tool flow in SCE.

3. SCE DESIGN FLOW

Figure 4 shows the refinement-based tool flow in SCE from the initial abstract specification down to the final implementation model. In particular, the SCE flow consists of six specific tools which we will describe in the following sections.

3.1. Architecture exploration

The first step in the SCE design flow, architecture exploration, defines the target platform and, under a set of design constraints, maps the computational parts of the specification model onto that platform. The target architecture consists of a set of PEs, that is, software processors, custom hardware blocks, and memories. These components are selected by the system designer as part of the decision making. In particular, the designer selects the type and the number of PEs, CEs, and communication busses.

Architecture exploration consists of two tasks: PE allocation and partitioning. PE allocation defines the target architecture by selecting system components (software and hardware processors, memories) from the PE database.

Partitioning then maps behaviors and variables to the allocated PEs and memories, respectively.

Following the design decisions of PE allocation and partitioning, the SCE architecture refinement tool *scar* inserts an additional layer of hierarchy representing the PEs into the model and groups behaviors and variables under these according to the partitioning. Next, it refines given complex channels into a client-server implementation using message-passing communication between the PEs and inserts necessary synchronization to properly preserve the original execution semantics. Finally, *scar* automatically generates the output architecture model [27].

3.2. Scheduling exploration

A key feature in the SCE design flow is the early evaluation of different scheduling strategies for software processors that are sequential and physically can only execute one task at a time. To evaluate different static and dynamic scheduling algorithms, such as round-robin or priority-based scheduling, we utilize a high-level RTOS model on each processor in the system [28]. Our abstract RTOS model is written on top of the SpecC language and does not require any specific language extensions. It supports all the key concepts found in modern RTOS, including task management, real-time scheduling, preemption, task synchronization, and interrupt handling.

After the designer chooses the desired scheduling strategy (e.g., round-robin, priority-based, or first-come-first-served), the SCE scheduling refinement tool *scos* automatically groups the given behaviors in the software PE into tasks and inserts the RTOS model with the user-defined scheduling strategy into the design model. *scar* then wraps all primitives and events that can trigger scheduling, such as task activation and termination, IPC synchronization and communication, and timing wait statements so that the inserted RTOS is called. It finally generates the refined model that can then be simulated for accurate observation and evaluation of dynamic scheduling behavior in the multitasking system. Since our abstract RTOS model requires only minimal overhead in simulation time, this approach enables early and rapid design space exploration.

3.3. Network exploration

Network exploration defines the system communication topology and maps the given communication channels onto a network of busses and communication elements (CEs), that is, bridges and transducers. For this, network refinement inserts the required CEs from the database into the model and implements the end-to-end communication over point-to-point links between PEs and CEs [29].

In the input architecture model, PEs communicate via abstract, typed end-to-end channels, and memory interfaces. During network exploration, the user allocates the actual communication media, bridges, and transducers for the system busses and CEs, respectively. Furthermore, the designer defines the connectivity of PE and CE ports to the busses,

and maps architecture-level end-to-end channels onto the allocated bus network.

Based on the network decisions by the designer, the SCE network refinement tool `scnr` inserts and implements the ISO/OSI presentation, network and transport layers, which implement data conversion, packeting, and routing; and acknowledgements, respectively. `scnr` then generates the new network model such that it reflects the selected network topology including typed end-to-end architecture level communication over untyped point-to-point links between the components in each network segment.

3.4. Communication synthesis

Next, the task of communication synthesis is to implement the point-to-point logical links between stations over the actual bus media, and to select and define the final pin- and bit-accurate parameters of the communication architecture under a set of constraints. Communication refinement then inserts protocols and bus-functional component descriptions from the bus and PE/CE databases, respectively, and generates a refined communication model that implements the communication links in each network segment over the actual, shared bus protocol and bus wires. In addition to this pin-accurate model (PAM), our communication refinement also generates a fast-simulating TLM of the system, which abstracts away the pin-level details of individual bus transactions [29].

In the input network model, communication in each network segment is described as a set of logical links. During communication synthesis, the designer (through the GUI, scripting or using synthesis plugins) defines the bus parameters, such as address and interrupt assignments, for each logical link over each bus. Based on these decisions, the SCE communication refinement tool `sccr` inserts low-level (transaction-level down to pin-accurate) models of busses and components from the databases, and generates a new communication model (PAM or TLM) of the design. In the output model, PE and CE components are refined to implement the lower communication layers (link, stream, media access, and protocol layer) for synchronization, addressing, and media accesses over each bus interface. On top of bus models from the bus database, the generated model hence implements all system communication down to the level of timing-accurate bus transactions (TLM), or cycle-accurate events for sampling and driving of the bus wires (PAM).

3.5. RTL synthesis

The task of RTL synthesis is to generate structural RTL from the behavioral description of the hardware components in the design. Although the designer can freely choose all behavioral synthesis parameters, including scheduling, allocation, and binding decisions, the SCE RTL synthesis tool `scrtl` supports automatic decision making through plugins. The designer can choose an algorithm to apply to all or only parts of their design. Critical parts of the design, on the other hand, can be manually preassigned or postoptimized [30].

Both designers and algorithms can rely on a set of estimates to aid them in the decision making. SCE includes RTL-specific profiling and analysis tools that provide feedback about a variety of metrics including delay, power, and variable lifetimes.

RTL synthesis in SCE takes full advantage of the designers' insight by allowing them to enter, modify, or override their decisions at will. On the other hand, tedious and error-prone tasks including code generation are automated.

3.6. Software synthesis

For implementing the software components in the system model, SCE relies on a layer-based modeling of the programmable processors and the software stack executing on them. Our embedded processor model supports task scheduling and interrupt handling.

Given scheduling priorities defined by the system designer, the SCE software synthesis tool `sc2c` automatically generates embedded software code for each processor from the system model [31]. More specifically, we generate efficient ANSI-C code from the SLDL code of the mapped application, and compile and link it against the selected RTOS. The resulting software binary can then be used for cycle-accurate instruction-set simulation within the system model, as well as for the final implementation.

4. EXPERIMENTS AND RESULTS

We have applied SCE to a large set of industrial-strength examples. In the following, we will first demonstrate the SCE design flow in detail as applied to a case study. Next, we summarize our experiences with different examples and show exploration results. Finally, we will present a set of verification experiments.

4.1. Modeling experiment

In order to demonstrate the overall SCE design flow, we have applied the flow to the example of a mobile phone baseband platform. The specification model of the system is shown in Figure 5. The design combines a JPEG encoder for processing of digital pictures taken by a camera and a voice encoder/decoder (vocoder) for speech processing based on the mobile phone GSM standard. Both JPEG and Vocoder processes are hierarchically composed of subbehaviors implementing the encoding and decoding algorithms in nested and pipelined loops and communicating through abstract message-passing channels. At the top level, a channel `Ctrl` between the two processes is used to send control messages from the JPEG encoder to the vocoder.

For the target platform (for space reasons, we do not show the platform model separately; the model is almost identical to Figure 6, with the exception that the OS layer and OS channel are omitted), we decide to use two software processors assisted by several hardware accelerators. For the JPEG encoder, we select a Motorola Coldfire processor for the main execution, assisted by a special IP component `DCT_IP` which performs the needed discrete cosine transformation

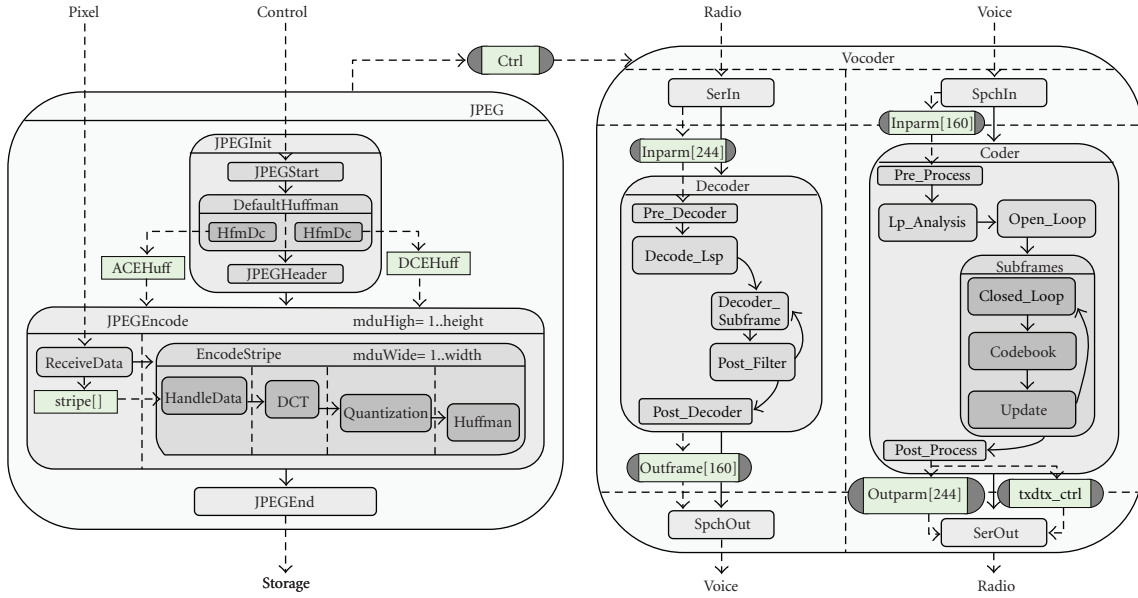


FIGURE 5: Baseband example: specification model.

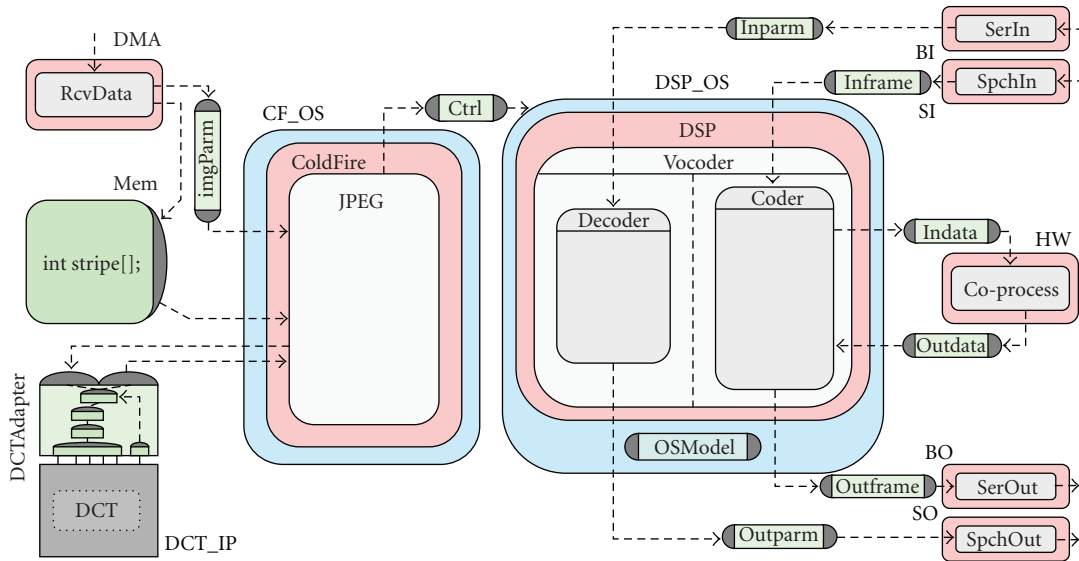


FIGURE 6: Baseband example: scheduled architecture model.

(DCT) in hardware. We also choose a direct memory access component *DMA* that receives pixel stripes from the camera and puts them into a shared memory *Mem*. On the other hand, we select a digital signal processor *DSP* to perform the majority of the voice encoding and decoding tasks. To reach the required performance, the *DSP* is assisted by four hardware blocks dedicated to input and output of the data streams, and one custom coprocessor in charge of the codebook search, the most time-critical function in the vocoder.

In the scheduled model obtained after architecture partitioning and scheduling (Figure 6), the *ColdFire* processor runs the JPEG encoder in software assisted by the hardware

DCT_IP. Since this processor only executes this one task, no operating system is needed and the OS layer *CF_OS* is empty. On the other hand, the *DSP* performs two concurrent speech encoding and decoding tasks. These tasks are dynamically scheduled under the control of a priority-based operating system model that sits in an additional OS layer *DSP_OS* around the *DSP*. The encoder on the *DSP* is assisted by a custom hardware coprocessor (*HW*) for the codebook search. Furthermore, four custom hardware I/O processors perform buffering and framing of the vocoder speech and bit streams.

Table 1 summarizes the design decisions made for implementing the communication channels in the example. As a

TABLE 1: Communication design parameters for baseband example.

Channel	Network	Addr.	Link	Medium
	Routing		Intr.	
imgParm	linkDMA	0x00010000	int7	cfBus
stripe[]	Mem	0x0010xxxx	—	
hData dData	linkDCT	0x00010010	int1	
Ctrl	linkTx1	0x00010020	int2	dspBus
Ctrl	linkTx2	0xB000	intA	
inframe	linkSI	0x800x	intB	
outparm	linkBO	0x950x		
indata	linkHW	0xA000	intD	
outdata				
inparm	linkBI	0x850x	intC	
outframe	linkSO	0x900x		

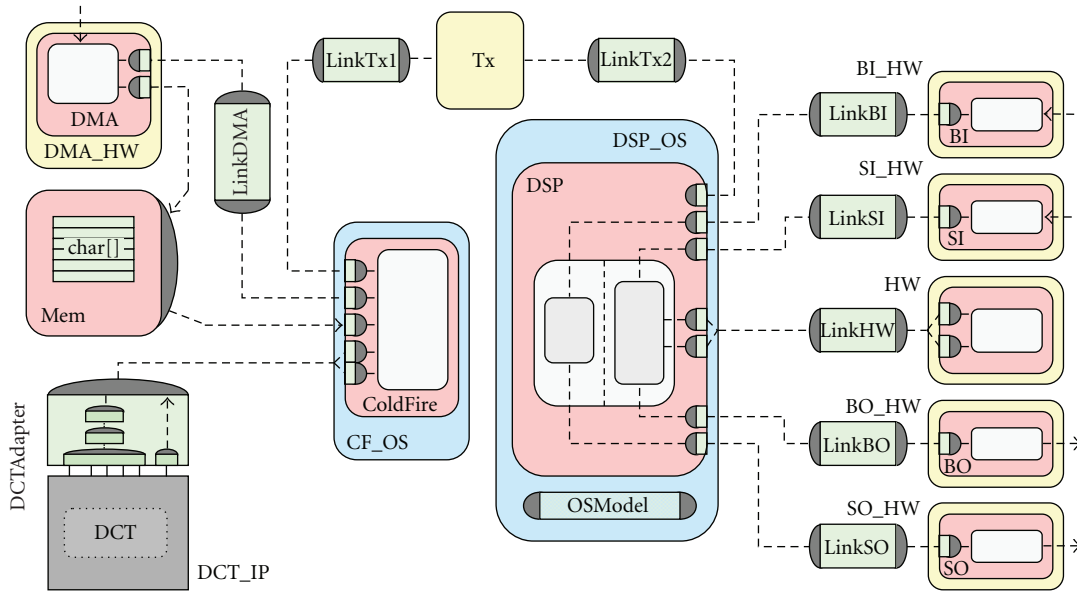


FIGURE 7: Baseband example: network model.

result of the network exploration, the network is partitioned into one segment per subsystem with a transducer T_x connecting the two segments (Figure 7). Individual point-to-point logical links connect each pair of stations in the resulting network model. Application channels are routed statically over these links where the $Ctrl$ channel spanning the two subsystems is routed over two links via the intermediate transducer.

During communication synthesis, all links within each subsystem are implemented over a single shared medium. In both cases, the native ColdFire and DSP processor busses are selected as communication media. Within the segments, unique bus addresses and interrupts for synchronization are

assigned to each link. On the ColdFire side, the memory is assigned a range of addresses with a base address plus offsets for each stored variable. On the DSP side, two of the four available interrupts are shared among the four I/O processors. In those cases, additional bus addresses for slave polling are assigned to each link (base address plus one). Finally, a bridge DCT_Br is inserted to translate between the DCT_IP and ColdFire bus protocols.

As a result, SCE communication synthesis generates two models, a fast-simulating TLM (Figure 8), and a pin-accurate model (PAM, Figure 9) for further implementation. In the TLM, link, stream, and media access layers are instantiated inside the OS and hardware layers of each station. Inside the

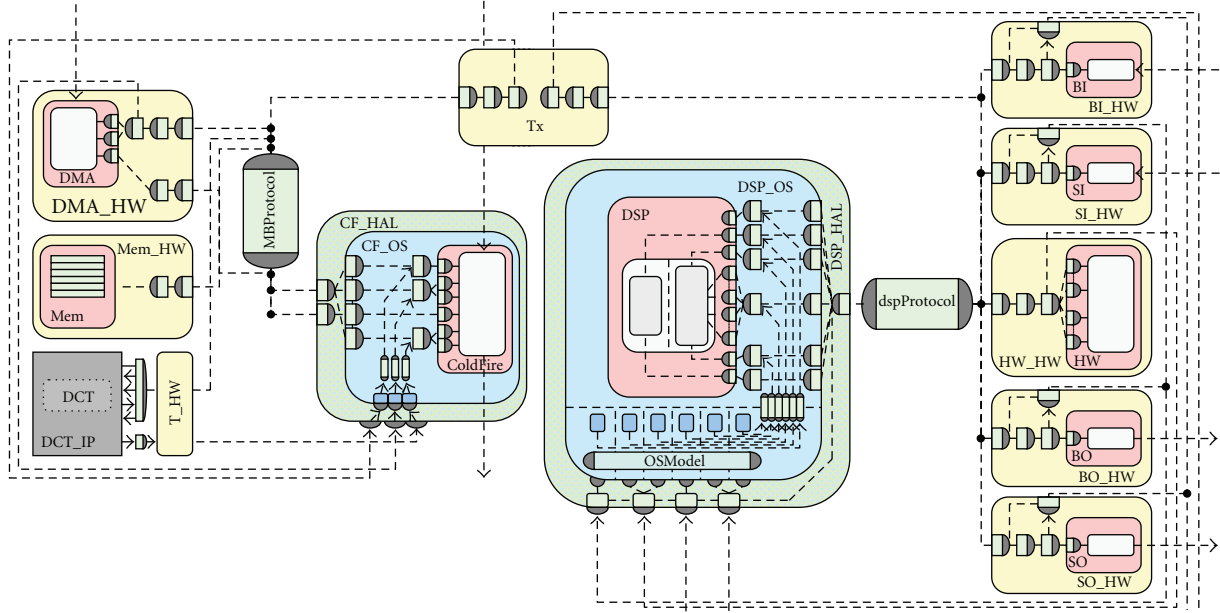


FIGURE 8: Baseband example: transaction-level model (TLM).




FIGURE 9: Baseband example: pin-accurate model (PAM).

processors, interrupt handlers that communicate with link layer adapters through semaphores are created. Interrupt service routines (*ISR*) together with models of programmable interrupt controllers (*PIC*) model the processor's interrupt behavior and invoke the corresponding handlers when triggered.

In the PAM, additionally the communication protocol layers are instantiated. Components are connected via pins and wires driven by the protocol layer adapters. On the ColdFire side, an additional arbiter component regulates bus accesses between the two masters, *DMA_BF* and *CF_BF*.

Table 2 summarizes the results for the example design. Using the refinement tools, models of the example design were automatically generated within seconds. A testbench common to all models was created which exercises the design by simultaneously encoding and decoding 163 frames of

speech on the vocoder side while performing JPEG encoding of 30 pictures with 116×96 pixels. We created and refined both models of the whole system and models of each subsystem separately. Note that code sizes (lines of code, LOC) in each case include the testbenches. Since testbench code is shared, the size of the system model is less than the sum of the subsystem model sizes. All models were simulated on a 2.7 GHz Linux workstation using the QuickThreads version of the SpecC simulator.

Figure 10 plots simulation times on a logarithmic scale, that is, the graph shows that simulation times generally grow exponentially with each new model at the next lower level of abstraction. On the other hand, results of simulated overall frame transcoding (back-to-back encoding and decoding) and picture encoding delays in the vocoder and JPEG encoder, respectively, are shown in Figure 11. As can be seen,

TABLE 2: Modeling and simulation results for baseband example.

Model	ColdFire subsystem			DSP subsystem			System	
	LOC	Simul. time	JPEG delay	LOC	Simul. time	Vocoder delay	LOC	Simul. time
Specification	1,819	0.02 s	0.00 ms	9,736	1.31 s	0.00 ms	11,481	2.25 s
Architecture	2,779	0.03 s	9.66 ms	11,121	1.21 s	8.39 ms	13,866	2.56 s
Scheduled	3,098	0.02 s	22.63 ms	13,981	1.20 s	12.02 ms	17,020	2.00 s
Network	3,419	0.02 s	22.63 ms	14,319	1.22 s	12.02 ms	17,658	2.03 s
TLM	5,765	1.04 s	24.03 ms	15,668	27.4 s	13.00 ms	21,446	92.3 s
PAM	5,916	14.3 s	24.02 ms	15,746	34.8 s	13.00 ms	21,711	2,349 s
RTL-C	7,991	14.9 s	23.48 ms	23,661	147 s	12.88 ms	33,511	2,590 s

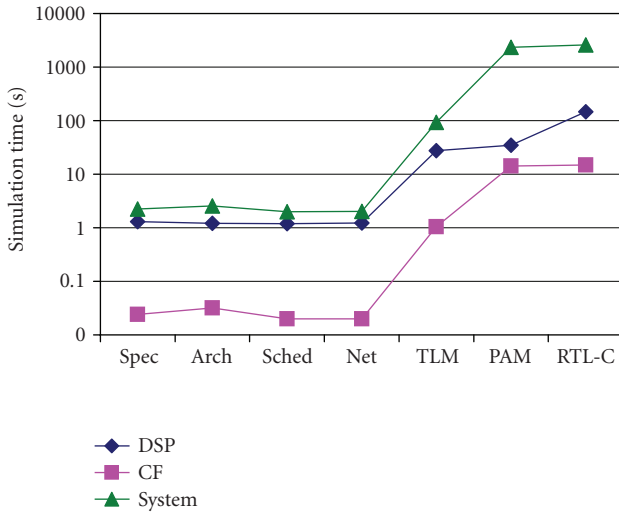


FIGURE 10: Simulation speeds for the baseband example.




FIGURE 11: Simulated delays in the baseband example.

with each new model, measured delays linearly converge towards the final result.

Note that initial specification models are untimed and hence do not provide any delay measurements at all. Beginning with the architecture level, estimated execution

TABLE 3: Design examples and target architectures.

Examples		Buses (masters → slaves)
JPEG	A1	CF → HW
	A2	DSP → HW
	A3	DSP → HW1, HW2, HW3
Vocoder	A1	CF → HW1 CF → HW1, HW2, HW3
	A2	HW1 → HW3 HW2 → HW3
	A3	CF → HW1, HW2, HW3, HW4 HW1 → HW3 ↔ HW5 HW2 ↔ HW4 ↔ HW5
MP3float	A1	ARM → 2 I/O
	A2	ARM → 2 I/O, LDCT, RDCT
	A3	ARM → 2 I/O, LDCT, RDCT LDCT ↔ I/O RDCT ↔ I/O
MP3fix	A1	DSP → HW, 4 I/O, T
	A2	CF, DMA → Mem, BR, T, DMA BR → DCTIP
Baseband	A1	ARM → 4 I/O, 2 DCT, T LDCT, RDCT → I/O
	A2	DSP → HW, 4 I/O, T
Cellphone	A1	LDCT, RDCT → I/O DSP → HW, 4 I/O, T
	A2	ARM → 4 I/O, 2 DCT, T

delays are back-annotated into the computation blocks. As expected, scheduling has a large effect on simulation accuracy where abstract OS modeling enables evaluation of scheduling decisions at native simulation speeds (note that since the amount of simulated parallelism decreases, simulation is potentially even faster than at the specification level). Depending on the relation of communication versus computation, introducing bus models and communication delays at the transaction-level further increases accuracy, potentially at the cost of significantly longer simulation times. On the other hand, TLMs allow for accurate modeling of communication close or equivalent to pin-accurate models but at higher speed.

TABLE 4: Results for exploration experiments.

Examples		Model size (LOC)					Refinement time				Total
		Spec	Arch	Sched	Net	PAM	scar	scos	scnr	sccr	
JPEG	A1	1806	2409	2732	2780	4642	0.27 s	0.37 s	0.16 s	0.21 s	1.01 s
	A1		8449	9594	9775	10679	2.29 s	1.30 s	0.62 s	0.56 s	4.77 s
Vocoder	A2	7385	8508	9632	9913	10989	2.41 s	1.36 s	0.75 s	0.69 s	5.21 s
	A3		8535	9659	9949	11041	2.64 s	1.69 s	0.79 s	0.64 s	5.76 s
MP3float	A1		6963	28190	28204	29807	0.82 s	3.24 s	0.90 s	0.90 s	5.86 s
	A2	6900	7181	28275	28633	31172	0.93 s	2.66 s	1.11 s	1.48 s	6.18 s
	A3		11069	28736	30202	32795	4.66 s	4.05 s	7.10 s	1.88 s	17.69 s
MP3fix	A1		13724	17131	17270	21593	0.95 s	1.37 s	0.58 s	0.95 s	3.85 s
	A2	13363	16040	18300	18564	23228	3.28 s	1.68 s	0.85 s	1.20 s	7.01 s
	A3		16023	18748	19079	24471	2.72 s	1.76 s	1.97 s	0.95 s	7.40 s
Baseband	A1	11481	13866	17020	17658	21711	4.27 s	2.46 s	1.24 s	1.02 s	8.99 s
Cellphone	A1	16441	18653	21936	22570	30072	3.86 s	3.10 s	1.31 s	1.22 s	9.49 s

TABLE 5: Results for equivalence verification.

Examples	Refinement	Model 1		Model 2		No. of transformations	Verification time		
		Type	No. of nodes	No. of edges	Type			No. of nodes	No. of edges
JPEG	Architecture	spec.	148	219	arch.	180	257	1602	1.6 s
	Scheduling	arch.	180	257	sched.	180	287	2740	2.1 s
	Network	sched.	180	287	net.	201	253	2852	2.1 s
Vocoder	Architecture	spec.	436	761	arch.	528	882	6131	3.3 s
	Scheduling	arch.	528	882	sched.	528	881	7065	3.7 s
	Network	sched.	528	881	net.	569	933	7229	3.8 s

Our results show that with increasing implementation detail at lower levels of abstraction, accuracy (as measured by the simulated delays) improves linearly while model complexities (as measured by code sizes and simulation times) grow exponentially. All in all, our results support the choice of intermediate models in the design flow that allows for fast validation of critical design aspects at early stages of the design process.

4.2. Exploration experiments

In order to demonstrate our approach in terms of design space exploration for a wide variety of designs, we applied SCE to the design of six industrial-strength examples: stand-alone versions of the JPEG encoder (*JPEG*) and the GSM voice codec (*Vocoder*), floating- and fixed-point versions of an MP3 decoder (*MP3float* and *MP3fix*), the previously introduced baseband example (*Baseband*), and a *Cellphone* example combining the JPEG encoder, the MP3 decoder, and the GSM vocoder in a platform mimicking the one used in the RAZR cellphone. For each example, we generated different architectures using Motorola DSP56600 (*DSP*), Motorola ColdFire (*CF*), and ARM7TDMI (*ARM*) processors together with custom hardware coprocessors (*HW*, *DCT*) and I/O units. We used various communication architectures with *DSP*, *CF*, *ARM* (AMBA AHB), and simple handshake busses.

Table 3 summarizes the features and parameters of the different design examples we tested. For each example, the target architectures are specified as a list of masters plus slaves for each bus in the system where the bus type is implicitly determined to be the protocol of the primary master on the bus. For example, in the case of the *MP3float* design, the ColdFire processor communicates with dedicated hardware units over its *CF* bus whereas the *HW* units communicate with each other through separate handshake busses. For simplicity, routing, address, and interrupt assignment decisions are not shown in this table.

Table 4 shows the results of exploration of the design space for the different examples. Overall model complexities are given in terms of code size using lines of code (LOC) as a metric. Results show significant differences in complexity between input and generated output models due to extra implementation detail added between abstraction levels.

Note that manual refinement would require tremendous effort (in the order of days). Automatic refinement, on the other hand, completes in the order of seconds. Our results therefore show that a significant productivity gain can be achieved using SCE with automatic model refinement.

4.3. Verification experiments

We implemented the SCE equivalence verification tool *scver* to verify the refinements above network level. Since the

lowest abstraction level of communication in model algebra is the channel, models below network level in the SCE flow could not be directly translated into model algebraic representation.

The results for verification of architecture, scheduling, and network refinements are presented in Table 5. We used two benchmarks, namely, the JPEG encoder and Vocoder as shown in column 1. The model algebraic representation was stored in a graph data structure, with nodes being the objects and edges being the composition rules. Column 5 shows the total transformations applied to derive model 1 from model 2 using the transformation rules of model algebra. As we can see, since the order of transformation is decided, it only took a few seconds to apply them even for representations with hundreds of nodes and edges. The verification time also includes the time it took to parse the SpecC models into model algebraic representation and to perform isomorphism checking between the derived and original model graphs.

The results demonstrate that the SCE tool flow based on well-defined model abstractions and semantics enables fast equivalence verification.

5. SUMMARY AND CONCLUSION

In this work, we have presented SCE, a comprehensive system design framework based on the SpecC language. SCE supports a wide range of heterogeneous target platforms consisting of custom hardware components, embedded software processors, dedicated IP blocks, and complex communication bus architectures.

The SCE design flow is based on a series of automated model refinement steps where the system designer makes the decisions and SCE quickly provides estimation feedback, generates new models automatically, and validates them through simulation and formal verification. The effective design automation tools integrated in SCE allow rapid and extensive design space exploration. The fast exploration capabilities, in turn, enable the designer to optimize the system architecture, the scheduling policies, the communication network, and the hardware and software components, so that an optimal implementation is reached quickly.

We have demonstrated the benefits of SCE by use of six industrial-size examples with varying target architectures, which have been designed and verified top-to-bottom. Compared to manual coding and model refinement, SCE achieves productivity gains by orders of magnitude.

SCE has been successfully transferred to and applied in industrial settings. SER, a commercial derivative of SCE, has been developed and integrated into ELEGANT, an environment for electronic system-level (ESL) design of space and satellite electronics that was commissioned by the Japanese Aerospace Exploration Agency (JAXA). ELEGANT and SER have been successfully delivered to JAXA's suppliers and are currently being introduced into the general market [32].

ACKNOWLEDGMENTS

The authors would like to thank all members of the CECS SpecC group who have contributed to SCE over the years.

Special thanks go to David Berner, Pramod Chandraiah, Quoc-Viet Dang, Alexander Gluhak, Eric Johnson, Raphael Lopez, Gunar Schirner, Ines Viskic, Shuqing Zhao, and Jianwen Zhu.

REFERENCES

- [1] D. D. Gajski, J. Zhu, R. Dömer, A. Gerstlauer, and S. Zhao, *SpecC: Specification Language and Design Methodology*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2000.
- [2] D. D. Gajski, F. Vahid, S. Narayan, and J. Gong, *Specification and Design of Embedded Systems*, Prentice Hall, Upper Saddle River, NJ, USA, 1994.
- [3] P. Coste, F. Hessel, Ph. Le Marrec, et al., "Multilanguage design of heterogeneous systems," in *Proceedings of the 7th International Workshop on Hardware/Software Codesign (CODES '99)*, pp. 54–58, Rome, Italy, May 1999.
- [4] P. Gerin, S. Yoo, G. Nicolescu, and A. A. Jerraya, "Scalable and flexible cosimulation of SoC designs with heterogeneous multi-processor target architectures," in *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC '01)*, pp. 63–68, Yokohama, Japan, January–February 2001.
- [5] *ModelSim SE User's Manual*, Mentor Graphics Corp.
- [6] J. Buck, S. Ha, E. A. Lee, and D. G. Messerschmitt, "Ptolemy: a framework for simulating and prototyping heterogeneous systems," *International Journal of Computer Simulation*, vol. 4, no. 2, pp. 155–182, 1994.
- [7] T. Grötter, S. Liao, G. Martin, and S. Swan, *System Design with SystemC*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2002.
- [8] M. Aubury, I. Page, G. Randall, J. Saul, and R. Watts, *Handel-C language reference guide*, Oxford University Computing Laboratory, Oxford, UK, August 1996.
- [9] F. Ghenassia, *Transaction-Level Modeling with SystemC: TLM Concepts and Applications for Embedded Systems*, Springer, New York, NY, USA, 2005.
- [10] A. Österling, T. Brenner, R. Ernst, D. Herrmann, T. Scholz, and W. Ye, "The COSYMA system," in *Hardware/Software Co-Design: Principles and Practice*, J. Staunstrup and W. Wolf, Eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.
- [11] C. A. Valderrama, M. Romdhani, J.-M. Daveau, G. F. Marchioro, A. Changuel, and A. A. Jerraya, "Cosmos: a transformational co-design tool for multiprocessor architectures," in *Hardware/Software Co-Design: Principles and Practice*, J. Staunstrup and W. Wolf, Eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.
- [12] F. Balarin, M. Chiodo, P. Giusto, et al., *Hardware-Software Co-Design of Embedded Systems: The POLIS Approach*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.
- [13] G. Vanmeerbeeck, P. Schaumont, S. Vernalde, M. Engels, and I. Bolsens, "Hardware/software partitioning for embedded systems in OCAPI-xl," in *Proceedings of the International Symposium on Hardware-Software Codesign (CODES '01)*, Copenhagen, Denmark, April 2001.
- [14] P. Schaumont, S. Vernalde, L. Rijnders, M. Engels, and I. Bolsens, "A programming environment for the design of complex high speed ASICs," in *Proceedings of the 35th Annual Conference on Design Automation (DAC '98)*, pp. 315–320, San Francisco, Calif, USA, June 1998.

- [15] K. Grüttner, F. Oppenheimer, W. Nebel, A.-M. Fouiliart, and F. Colas-Bigey, "SystemC-based modelling, seamless refinement, and synthesis of a JPEG 2000 decoder," in *Proceedings of the Design, Automation and Test in Europe Conference (DATE '08)*, pp. 128–133, Munich, Germany, March 2008.
- [16] W. O. Cesário, D. Lyonard, G. Nicolescu, et al., "Multiprocessor SoC platforms: a component-based design approach," *IEEE Design and Test of Computers*, vol. 19, no. 6, pp. 52–63, 2002.
- [17] D. Lyonard, S. Yoo, A. Baghdadi, and A. A. Jerraya, "Automatic generation of application-specific architectures for heterogeneous multiprocessor system-on-chip," in *Proceedings of the 38th Annual Conference on Design Automation (DAC '01)*, pp. 518–523, Las Vegas, Nev, USA, June 2001.
- [18] K. van Rompaey, I. Bolsens, H. De Man, and D. Verkest, "CoWare—a design environment for heterogeneous hardware/software systems," in *Proceedings of the European Design Automation Conference (EURO-DAC '96)*, pp. 252–257, Geneva, Switzerland, September 1996.
- [19] W. Klingauf, H. Gädke, and R. Günzel, "TRAIN: a virtual transaction layer architecture for TLM-based HW/SW code-sign of synthesizable MPSoC," in *Proceedings of the Design, Automation and Test in Europe Conference (DATE '06)*, vol. 1, Munich, Germany, March 2006.
- [20] T. Kempf, M. Doerper, R. Leupers, et al., "A modular simulation framework for spatial and temporal task mapping onto multi-processor SoC platforms," in *Proceedings of the Design, Automation and Test in Europe Conference (DATE '05)*, vol. 2, pp. 876–881, Munich, Germany, March 2005.
- [21] F. Balarin, Y. Watanabe, H. Hsieh, L. Lavagno, C. Passerone, and A. Sangiovanni-Vincentelli, "Metropolis: an integrated electronic system design environment," *Computer*, vol. 36, no. 4, pp. 45–52, 2003.
- [22] A. L. Sangiovanni-Vincentelli, "Quo vadis SLD: reasoning about trends and challenges of system-level design," *Proceedings of the IEEE*, vol. 95, no. 3, pp. 467–506, 2007.
- [23] S. Abdi, J. Peng, H. Yu, et al., "System-on-chip environment (SCE version 2.2.0 beta): tutorial," Tech. Rep. CECS-TR-03-41, Center for Embedded Computer Systems, University of California, Irvine, Calif, USA, July 2003.
- [24] L. Cai, A. Gerstlauer, and D. Gajski, "Retargetable profiling for rapid, early system-level design space exploration," in *Proceedings of the 41st Annual Conference on Design Automation (DAC '04)*, pp. 281–286, San Diego, Calif, USA, June 2004.
- [25] S. Abdi and D. Gajski, "Verification of system level model transformations," *International Journal of Parallel Programming*, vol. 34, no. 1, pp. 29–59, 2006.
- [26] A. Gerstlauer, L. Cai, D. Shin, H. Yu, J. Peng, and R. Dömer, *SCE database reference manual, version 2.2.0 beta*, Center for Embedded Computer Systems, University of California, Irvine, Calif, USA, July 2003.
- [27] J. Peng and D. Gajski, "Optimal message-passing for data coherency in distributed architecture," in *Proceedings of the 15th International Symposium on System Synthesis*, pp. 20–25, Kyoto, Japan, October 2002.
- [28] A. Gerstlauer, H. Yu, and D. D. Gajski, "Rtos modeling for system level design," in *Proceedings of the Design, Automation and Test in Europe Conference (DATE '03)*, Munich, Germany, March 2003.
- [29] A. Gerstlauer, D. Shin, J. Peng, R. Dömer, and D. D. Gajski, "Automatic layer-based generation of system-on-chip bus communication models," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 26, no. 9, pp. 1676–1687, 2007.
- [30] D. Shin, A. Gerstlauer, R. Dömer, and D. D. Gajski, "An interactive design environment for C-based high-level synthesis of RTL processors," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 16, no. 4, pp. 466–475, 2008.
- [31] H. Yu, R. Dömer, and D. Gajski, "Embedded software generation from system level design languages," in *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC '04)*, pp. 463–468, Yokohama, Japan, January 2004.
- [32] *CECS eNews Volume 7, Issue 3*, Center for Embedded Computer Systems, University of California, Irvine, Calif, USA, July 2007, <http://www.cecs.uci.edu/enews/CECSeNewsJul07.pdf>.