CONCURRENCY AND COMPUTATION: PRACTICE AND EXPERIENCE
Concurrency Computat.: Pract. Exper. 2000; **00**:1–7 *Prepared using* cpeauth.cls *[Version: 2002/09/19 v2.02]*

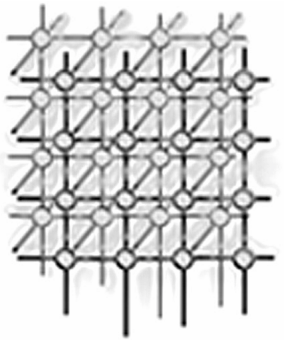

Formal System-level Design Space Exploration

Daniel Knorreck, Ludovic Apvrille and Renaud Pacalet*,†

System-on-Chip Laboratory (LabSoC)
Institut Telecom, Telecom ParisTech, LTCI CNRS
2229, routes des Crêtes, B.P. 193
F-06904 Sophia-Antipolis Cedex

SUMMARY

DIPLODOCUS is a UML profile intended for the modeling and the formal verification of real-time and embedded applications commonly executed on complex Systems-on-Chip. DIPLODOCUS implements the Y-Chart Approach, i.e., application and HW architecture (e.g., CPUs, bus, memories) are first described independently and are subsequently related to each other in a mapping stage. Abstract tasks and communication primitives are therefore mapped onto platform elements like buses and CPUs. DIPLODOCUS endows all models with a formal semantics, thereby paving the way for formal proofs both before and after mapping. More concretely, application, architecture and mapping models can be edited in TTool - an open-source toolkit - using UML diagrams. Then, pre or post mapping UML models may be automatically transformed into a LOTOS-based representation. This specification is in turn amenable to model-checking techniques to evaluate properties of the system, e.g., safety, schedulability, and performance properties. A smart card system serves as case study to illustrate the formal verification capabilities of DIPLODOCUS.

KEY WORDS: Design Space Exploration, Systems-on-Chip, UML, Formal Specification, Model Checking, TTool

*Correspondence to: System-on-Chip Laboratory (LabSoC)
Institut Telecom, Telecom ParisTech, LTCI CNRS
2229, routes des Crêtes, B.P. 193
F-06904 Sophia-Antipolis Cedex
†E-mail: {daniel.knorreck, ludovic.apvrille, renaud.pacalet}@telecom-paristech.fr

Received 10 December 2010
Revised 10 January 2011

1. Introduction

A System-on-Chip (SoC) is a set of functions distributed over hardware computation elements (CPUs, hardware accelerators) which are interconnected by a complex communication infrastructure (e.g., Network-on-Chip). The high complexity of applications executed on SoC - the smart card example provided in this paper being a prominent example - along with shortened time-to-market have pushed to their limits usual SoC designs methodologies. The analysis of systems at low abstraction levels yields a high degree of accuracy but comes with the downside of being demanding and slow. Indeed, traditional simulation techniques operating at register transfer level (RTL), instruction or transaction level are not appropriate for early design stages for two reasons: Only a very limited number of implementation alternatives can be examined due to the high modeling effort and extensive simulation runtime. Moreover, an incomplete specification early in the design flow may prohibit the construction of detailed models even if the effort were acceptable. Thus, abstractions are the key to success and furthermore make the models amenable to formal methods by reducing the state space.
Design Space Exploration (DSE) at system level is a major step in SoC design: it consists in selecting a software / hardware architecture complying to a set of constraints (performance, power consumption, functional properties, etc.). DIPLODOCUS[†] is the environment we advocate for tackling DSE. It forces the user to adhere to the Y-Chart approach [1] which has been extensively discussed in literature. In short it states that if application (functional view) and architecture (platform view) are handled in an orthogonal fashion, the burden of experimenting with several design points is considerably alleviated. The Y-Chart approach is thus very efficient for DSE, where an application is to be assessed with different architecture constraints. In DIPLODOCUS, for the time being, the interpretation of performance results and corresponding modifications of the architecture are carried out manually. Indeed, DIPLODOCUS does not incorporate any algorithm to prune the design space automatically. However, a huge body of work is concerned with multivariate optimization, genetic algorithms, etc. which could be tailored to DIPLODOCUS with reasonable effort.
DIPLODOCUS explicitly takes into account the hardware platform on which application tasks are executed. All system elements (application, architecture, mapping) can be efficiently represented in an abstract way using non-deterministic operators, complexity operators, and abstract hardware components. While modeling [2] and simulation [3] capabilities of DIPLODOCUS, as well as the toolkit supporting DIPLODOCUS [4] (i.e., TTool [5]) were already described in previous publications, the semantics of mapping models and the abstractions applied to architecture components have not been addressed so far. More precisely, the formal semantics of all DIPLODOCUS diagrams (applications, hardware architectures, mapping of applications onto hardware architectures) is defined in terms of a transformation function to LOTOS [6]. LOTOS is a process algebra supported with model-checkers. The paper is more particularly focused on abstractions applied to applications and hardware platforms, and on how formal methods take advantage of these abstractions. Moreover, we showcase

[†]DIPLODOCUS stands for design space exploration based on formal description techniques, UML and SystemC

the associated toolkit (TTool) which masks the underlying formal techniques to the designer (press-button approach).
The paper is organized as follows. Section 2 reviews related contributions. Section 3 recalls the DIPLODOCUS environment. Section 4 focuses on the formal semantics of DIPLODOCUS, and more precisely on the abstractions offered by DIPLODOCUS to allow formal analysis with limited combinatory explosion. Section 5 presents the implemented support toolkit. Section 6 illustrates our approach with a smart card system. Finally, section 7 concludes the article.

2. Related Work

Design Space Exploration (DSE) of Systems-on-Chip is the process of analyzing various functionally equivalent implementation alternatives to select an optimal solution [7]. The most suitable design is commonly chosen based on metrics such as *functionality*, *performance*, *cost*, *power*, *reliability*, and *flexibility*. DSE is usually very challenging because the system design space is extremely large and so usual simulation-based analysis techniques fail to efficiently observe the above mentioned metrics.
A first key factor is to choose an adequate abstraction level for the targeted design task. At early design stages, abstraction is a natural way to account for incomplete specifications and potentially unavailable low level models. Due to their accuracy, some approaches are rather suited for later design stages, where time consuming simulations can be tolerated. For example, [8] relies on an Instruction Set Simulator which executes the real code of the application. In [9], hardware components are considered at micro-architecture level, hence leading to long simulation times. Other environments offer formal exploration, but are sometimes limited to sub-elements of the platform [10].
Marculescu et al. [11] present a framework for computation and communication refinement for multiprocessor Soc Design. Stochastic automata networks represent the application behavior and the authors claim that this formalism allows for fast analytical performance evaluations. When it comes to binding an application to an architecture, transitions and states have to be added to the application model. Hence, application and architecture matters and not strictly handled in an orthogonal fashion. Due to a lack of data abstraction, the modeling of memory elements can quickly lead to state space explosion problem.
Especially when applying formal techniques, it is crucial to make the methodology accessible to practitioners. Intuitive high level languages which are automatically translated into formal languages serve this goal. The work of Hendriks and Verhoef [12] relies on timed automata to analyze timeliness properties of embedded systems. The UPPAAL model checker is used to evaluate the automata which must be created manually. There is no automated translation routine from a high level language (UML,...) and thus the creation of the automata turns out to be quite error prone, and not reusable.
Contributions on DSE environments such as [8,9,13–18] generally rely on a high-level language to describe application functions and architectures. For example, [16–18] rely on UML or MARTE diagrams.
Unfortunately, in many of these environments, architecture and application concerns are not independent [14], thus burdening the study of many alternative solutions. Other approaches

primarily target the performance evaluation of software alone.
The PUMA [19] framework is a unified approach to software modeling. It provides an interface between high level input models (such as UML diagrams) and performance oriented models. For that purpose, input models are first translated into an intermediate format called CSM so as to filter out irrelevant information for performance evaluations. In a second step, CSM can be converted to Petri Nets, Markov models, etc., and the resulting performance figures and design advice is fed back to the initial model. However, this framework concentrates on the modeling of software and thus does not yield a mapping where functionality is associated to software or hardware elements.
Some environments make simplifying assumptions on the control flow of the application. This may entail some nice properties (such as decidability, upper bounds on the worst case execution time, etc.) of the application model which can be evaluated with e.g. static methods. Viehl et al. [20] provide means for formal and simulation based evaluation of UML/SysML models for performance analysis of SoC. UML Sequence diagrams constitute the starting point for the functional description. They are subsequently transformed into so-called communication dependency graphs (CDGs) which thus capture the control flow, synchronization dependencies and timing information. CDGs are in turn amenable to static analysis in order to determine key performance parameters like best case response times, worst case response times and I/O data rates. A drawback of this approach is that data flow independence has to be kept, thus preventing case distinctions and loops with variable bounds to be part of the application model.
In stateless static approaches control flow cannot be represented at all. Thus the application is always assumed to be in a steady state and so upper and lower bounds on consumed and produced data - as well as the execution time - can be determined. One way of doing is for instance to characterize applications merely with their execution time on given resources [21].
SymTA/S [22] [23] and Real Time Calculus (RTC) rely on formal methods such as the real-time scheduling theory and deterministic queuing systems to determine characteristics of distributed systems. In SymTA/S the behavior of the environment is modeled by means of standard event arrival patterns including periodic and sporadic events with jitters or bursts. RTC imposes less restrictions by allowing deterministic event streams to be modeled with the aid of arrival curves denoting lower and upper bounds for event occurrences. Event streams are propagated among resources of distributed systems in a way that each resource may be analyzed separately with classical algorithms. However, the applicability of scheduling theories requires the task model to be simplistic and thus it merely reflects best case and worst case execution times. The so defined interval tends to be too large for tasks with a data dependent or irregular behavior: Stateless techniques are therefore probably inappropriate in many application domains.
DIPLODOCUS [2] forces the developer to follow the Y-Chart approach to ease the process of examining multiple design points. The use of an intuitive subset of UML masks the underlying theory and thus makes the approach accessible to developers having no skills in formal techniques. Also, the recommended but not compulsory graphical interface encourages developers to apply abstractions. DIPLODOCUS is focused on control rather than on data. Data is untyped and carries no value: only the amount of data is a relevant attribute. In turn, data abstraction calls for functional abstractions when it comes to data dependencies of the control flow. These abstractions greatly reduce simulation complexity and are a necessary prerequisite to make models amenable to formal proof techniques. In this paper, we apply the

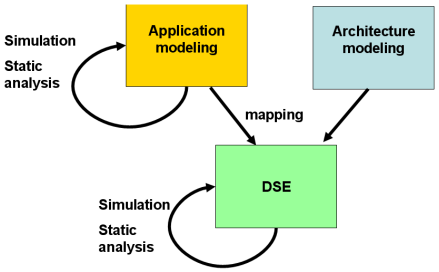

Figure 1. Methodology for Design Space Exploration

latter to conduct safety, performance and schedulability analysis based on the LOTOS process algebra. While LOTOS has already proven its value for hardware verification [24], we suggest its use in the context of more generic platforms (SoCs).

3. The DIPLODOCUS UML Profile

A UML profile customizes the UML language [25] for a given domain of systems. It extends the UML meta-model, according to *semantic variation points*, and is usually accompanied by a methodology. The DIPLODOCUS UML profile targets the modeling and Design Space Exploration of Systems-on-Chip at a high level of abstraction [2]. The DIPLODOCUS methodology, depicted in Figure 1, complies with the Y-Chart approach and includes three main steps, that are further reviewed in next subsections:

1. **Applications** are first described as a network of abstract communicating tasks using a UML class diagram. The latter represents the static view of the application. Each task behavior is expressed in terms of one UML activity diagram.
2. Targeted **architectures** are modeled independently from applications as a set of interconnected generic hardware nodes in UML: execution nodes (e.g. CPUs, hardware accelerators), communication nodes (e.g., buses, bridges), and storage nodes (e.g., memories). These nodes are parametrized for calibration purposes.
3. A **mapping** phase defines how application tasks are bound to execution nodes and also how abstract communications between tasks are assigned to communication and storage nodes.

3.1. Application model

At first, the application is modeled using UML class and activity diagrams. Tasks are modeled as classes interconnected with *channels*, *events*, or *requests* to communicate. Data abstraction is a key point: *channels* convey only unvalued samples i.e. a given amount of data. Also, *events*

are used for synchronization purpose, and *requests* are used to activate tasks.
As stated before, functions are modeled as a set of abstract tasks described with UML class diagrams. Task behavior is described using UML activity diagrams which are built upon the following operators: control flow and variable manipulation operators (loops, tests, assignments, etc.), communication operators (reading/writing abstract data in/to channels, sending/receiving events and requests), and computational cost operators (refer to 4.2.1 and 4.4.1 for detailed semantics and abstractions). Operators have been defined to encourage designers to follow these abstractions:

- **Data abstraction**: Only the amount of data exchanged between functional entities is modeled. Data dependent decisions are abstracted and thus expressed in terms of non-deterministic operators, such as non-deterministic choices, complexity and time operators.
- **Functional abstraction**: Algorithms are described using abstract cost operators. The complexity of computations is thus taken into account without actually having to carry them out.

Three communication and synchronization primitives have been defined:

- **Channels** are characterized by a point-to-point unidirectional communication between two tasks. Channel types are: Blocking Read/Blocking Write (BR-BW), Blocking Read/Non Blocking Write (BR-NBW) and Non Blocking Read/Non Blocking Write (NBR-NBW)
- **Events** are characterized by a point-to-point unidirectional asynchronous communication between two tasks. Events are stored in an intermediate FIFO between the sender and the receiver. This FIFO may be finite or infinite. In case of an infinite FIFO, incoming events are never lost. Indeed, when adding an event to a finite FIFO, the incoming event may be discarded or the oldest event may be dropped if the FIFO is full. Thus, a FIFO containing at most one element may be used to model asynchronous signals. In tasks, events can be sent (*notify*), received (*wait*) and tested for their presence (*notified*).
- **Requests** are characterized by a multipoint-to-point unidirectional asynchronous communication between tasks. A unique infinite FIFO between senders and the receiver is used to store requests. Consequently, a request cannot be lost.

3.2. Architecture model

A candidate architecture is modeled in terms of interconnected hardware components (or *nodes*) using UML stereotypes placed in UML deployment diagrams. The following nodes types are available in DIPLODOCUS:

- **Computation nodes**. Typically, an abstract **CPU** model merges both the functionality of the hardware component and its operating system. The behavior of a CPU model can be customized with parameters such as: data size, pipeline size, cache miss ratio and scheduling algorithm. For the time being, **DMAs** and **hardware accelerators** are represented by adequately parametrized CPU nodes. See 4.4.2 for comments on abstractions made on CPU nodes.

- **Communication nodes**. A communication node is either a **bus** or a **bridge**. The bus model has the following parameters: data size, latency and arbitration policy. A **Link** connects a hardware node - except for buses - to a bus. A link priority parameter may be considered by bus arbitration policies.
- **Storage nodes**. **Memories** have latency and data size parameters.

3.3. Mapping

At the mapping stage, application tasks and channels are bound to hardware components within a UML deployment diagram. As a rule of thumb, a task is mapped onto exactly one execution node and channels are mapped onto n buses, $n-1$ bridges and one memory. A mapping imposes additional constraints on the application model by associating the latter to shared hardware resources. The objective of the mapping stage is to determine whether an architecture is able to accommodate the load defined by the application whilst complying to constraints.

3.4. Simulation and formal verification

As stated earlier, the DIPLODOCUS design flow is especially suited for early stages. The main DIPLODOCUS objective is to help designers to identify a suitable hardware architecture even if algorithmic details have not yet been stipulated thoroughly. To achieve this, DIPLODOCUS relies (i) on fast simulation and formal proof techniques, both at application and mapping level, and (ii) on application models clearly separated from architecture models. Due to the high abstraction level of both application and architecture models, simulation speed is significantly increased with regards to simulations usually performed at lower abstraction level [3], and formal proofs can be achieved: this article focuses on formal analysis techniques that may be applied before and after mapping. At application level, simulation and formal verification usually yield the verification of functional properties. A mapping assigns a semantics to complexity operators and allows to resolve the physical duration of operations. Therefore, performance properties are commonly investigated after mapping, when constraints due to resource sharing come into play. In this context, our framework yields key figures related to the scheduling on CPUs (compliance to deadlines, end to end delays, execution times,...), bus load (congestion, average contention delay experienced by bus masters, etc.), and also information on power consumption and silicon area.

4. Formal semantics and abstractions

4.1. Formal basis: LOTOS

The semantics of DIPLODOCUS models is established by a transformation function to LOTOS. LOTOS [6] is an ISO-standardized Formal Description Technique for distributed system specification and design. A LOTOS specification, being itself a process, is structured into processes. A LOTOS process is a black box that communicates with its environment

through gates using a multiway rendezvous offer. Values can be exchanged at synchronization time. LOTOS specifications may be formally verified with, for instance, the CADP toolkit [26]. CADP implements reachability graph generation, model-checking and graph minimization.

4.2. Semantics at application level

4.2.1. Tasks operators

As described in the previous section, an application is composed of a set of communicating tasks. Operators used to describe task behavior are of four types:

- **Communication operators**: **read** from a channel, **write** to a channel, **notify** an event, **wait** for an event, know whether an event has been sent (**notified**), **request** a task.
- **Control operators**: direct the control flow, such as **variable modifications**, **loops**, **tests**, random number.
- **Complexity operators**: stand for a number of operations on integers (**EXECI**), floats (**EXECF**) or custom (**EXECC**).
- **Temporal operators**: denote deterministic and non-deterministic physical **delays**.

This set of operators makes it possible to describe the communication behavior of applications and algorithms whilst encouraging the user to abstract data (by means of DIPLODOCUS channels) and data dependent decisions (using non deterministic operators such as choices and delays).
The LOTOS semantics of all task operators is further described in Table I, column "LOTOS Semantics before mapping".

4.2.2. Communications between tasks

Channels denote data dependencies between tasks, which impact the control flow of a task by implicitly imposing conditions on availability of data or storage capacity. The semantics of channels (as explained in 3.1) is naturally expressed in LOTOS: since channels convey no value, but only a number of samples, BR-BW and BRNBW channels can easily be translated into a simple process (see Figure 2) sharing a natural value (which represents the number of elements in the FIFO) between two processes using two gates: one gate to add a sample ($wr\_\_ch$), another one to remove a sample ($rd\_\_ch$). The last channel type (NBR-NBW) is translated into a similar LOTOS process with the only difference that no counter is necessary - since its is always possible to read and write -, and so no guards ($[]$ operator) are used before the actions on gates $wr\_\_ch$ and $rd\_\_ch$.

Events express synchronization dependencies between tasks. They can carry up to three parameters, and support three semantics: *infinite FIFO*, *finite blocking FIFO*, and *non-blocking FIFO*. The two first semantics have been selected because they reflect common synchronization schemes of embedded systems. The last one (Non-blocking finite FIFO) is particularly useful to model signal exchanges between tasks: indeed, software and hardware signals usually override previous occurrences of the same signal (e.g., Programmable Interrupt

Type	Task operators	LOTOS Semantics before mapping	LOTOS Semantics after mapping
Channel	Write n samples to a channel	n Write operations in FIFO, i.e., n times action on gate $wr\_\_ch$, see Figure 2	n cycles, and a request on a bus.
	Read n samples from a channel	n read operations from FIFO, i.e., n times action on gate $rd\_\_ch$, see Figure 2	n cycles and a request on a bus.
Event	Notify an event	Adds an event to the corresponding FIFO, i.e., performs an action on gate $notify\_\_evt$, see Figure 3	Same as before mapping.
	Wait for an event	Tries to get an event from a FIFO, i.e., performs an action on gate $wait\_\_evt$, see Figure 3	Same as before mapping.
	Notified	Returns the number of events in a FIFO using action $notified\_\_evt$, see Figure 3	Same as before mapping.
Request	Send a request (operator is called "request")	FIFO management is similar to the one used for events	Same as before mapping.
Control	loop, variable modifications, tests	Direct translation in LOTOS with corresponding LOTOS operators	Direct translation. Operators are executed in 0-cycle.
Complexity	EXECx n,m i.e., between n and m integer instructions	No semantics before mapping, i.e., this operator is ignored	The task executes between $n * perf$ and $m * perf$ cycles with perf denoting a constant value characterizing the performance of the dedicated HW component
Temporal	Delay $d_{min}d_{max}$ unit	No semantics before mapping, i.e., this operator is ignored	The task is blocked for Between n and m cycles with $n = d_{min} * frequency$ and $m = d_{max} * frequency$.

Table I. Task operators

```
process ChannelBRBW_ch[rd_ch, wr_ch](samples:nat) : exit := (
[samples < 8] -> (wr_ch; ChannelBRBW_ch[rd_ch, wr_ch](samples + 1))
[]
[samples > 0] -> (rd_ch; ChannelBRBW_ch[rd_ch, wr_ch](samples - 1)) )
```

Figure 2. Application-level LOTOS semantics for a BR-BW channel, containing at most 8 samples

Controller, or UNIX signals). A separate LOTOS process accounts for each of the three semantics using the *Queue_nat* algebraic type. Figure 3 illustrates a non-blocking finite FIFO (the most complex case) for an event carrying only one natural parameter. Five cases have been taken into account:

1. The FIFO is not empty, and so, a *wait* action can be performed on the FIFO.
2. The FIFO is not full, and so, an event can be added to the FIFO (*notify*).
3. The FIFO is full, and so, an event can be added to the FIFO (*notify*) after the oldest one has been removed.
4. The FIFO is not empty, the *notified* action returns the value 1.
5. The FIFO is empty, the *notified* action returns the value 0.

Unlike channels and events which are one-to-one communications, **requests are many-to-one communications**. They rely on *n-to-one* infinite FIFO. The translation of requests is similar to the one of FIFOs for events, apart from the fact that notification gates are instantiated n times, e.g., $notify\_i$ with $i \in 1 \ldots n$.

```
process Event_evt[notify_evt, wait_evt, notified_evt]
(fifo_1:Queue_nat, fifo_val_1:nat, nb:nat, maxs:nat) : exit :=
  [not (Empty (fifo_1))] -> wait_evt !First(fifo_1);p_0_Event_evt[notify_evt,
    wait_evt, notified_evt](Dequeue(fifo_1), fifo_val_1, nb-1, maxs)
  [] [nb<maxs] -> notify_evt ?fifo_val_1:nat;p_0_Event_evt[notify_evt,
    wait_evt, notified_evt](Enqueue(fifo_val_1, fifo_1), fifo_val_1, nb+1,
    maxs)
  [] [nb == maxs] -> notify_evt ?fifo_val_1:nat;p_0_Event_evt[notify_evt,
    wait_evt, notified_evt](Enqueue(fifo_val_1, Dequeue(fifo_1)),
    fifo_val_1, nb, maxs)
  [] [not (Empty (fifo_1))] -> notified_evt!1;p_0_Event_evt[notify_evt,
    wait_evt, notified_evt](fifo_1, fifo_val_1, nb, maxs)
  [] [Empty (fifo_1)] -> notified_evt!0;p_0_Event_evt[notify_evt, wait_evt,
    notified_evt](fifo_1, fifo_val_1, nb, maxs)
endproc
```

Figure 3. Application-level LOTOS semantics for a Non-blocking Finite FIFO

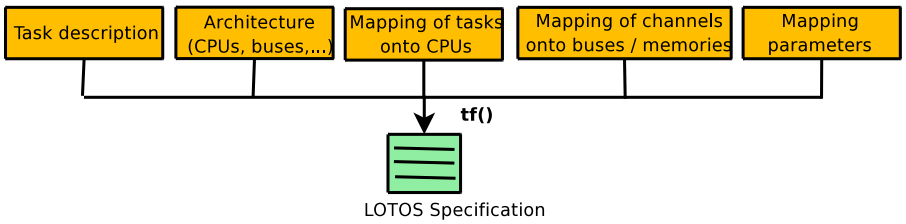

Figure 4. General approach

4.3. Semantics at mapping level

A mapping involves an application (i.e., a network of tasks interconnected by communication primitives), an architecture (i.e., a set of hardware nodes), a distribution of tasks onto hardware nodes (e.g., map the task *task1* onto the CPU *cpu1*), and a mapping of communication channels onto buses / memories. We have therefore defined a transformation function $tf()$ that takes as argument all above mentioned elements and generates a LOTOS specification (see Figure 4).

4.3.1. Basics of the transformation

The mapping phase sheds light on the question whether a system composed of an application and an architecture satisfies a set of constraints (compliance to deadlines,etc.). A mapping resolves arbitration of shared resources (typically, a CPU, a bus, etc.) and its semantics in LOTOS is defined accordingly. As a consequence, the LOTOS specification of a mapping should take into account:

- The arbitration of shared resources, e.g., for tasks: access to CPUs, and for communication: access to buses. To that end, we look at the impact of scheduling policies of operating systems as well as the arbitration policies of buses.
- The time taken by tasks to execute operators, and the time consumed by communication, e.g., bus and memory latencies.

4.3.2. The Mapping-to-LOTOS transformation

All task operators and hardware node parameters are taken into account by the Mapping-to-LOTOS transformation ($tf()$). The transformation does not yet incorporate the latest proposals on resource sharing in DIPLODOCUS (e.g.hierarchical scheduling and virtual nodes [27]). However, we are not aware of major obstacles preventing their integration.
Basically, the LOTOS specification is built upon four functional blocks:

- The **Scheduling manager** schedules tasks on each CPU. $tf()$ transforms each task into a state machine modeled in LOTOS: preemption can only occur when a task is in a state, but not during state transitions.
- The **Communication manager** handles channel-based communication between tasks running on the same CPU, or on different CPUs. Events and requests are assumed not to consume communication resources. Indeed, the amount of data represented by those two synchronization features are assumed to be negligible with regards to channel-based communications. Similar assumptions were made for the simulation semantics [3] (which is less abstract and more tailored to simulation runtime issues).
- The **Task execution manager** handles operators to execute in each task, that is transitions between various task states.
- The **Clock manager** handles clock cycles on hardware nodes, i.e., it activates necessary hardware nodes when a new cycle begins.

The main process of the LOTOS specification works as follows:

1. At first, an initialization phase is used to settle various data structures, for each CPU (e.g., all tasks of a CPU are put in "ready" state), and for the communication manager: data structures related to channels, queues related to events, and so on.
2. A main loop on clock cycles is started: The system waits for the next tick (*tick* bas been defined as a LOTOS action). Then, each CPU plus its operating system are considered one after another. Basically, a CPU is meant to interpret DIPLODOCUS application-level operators of the task selected by the scheduler. More precisely:
   (a) Depending on its clock rate, the CPU is activated or not by the Clock manager.
   (b) If it is activated, then a first test is performed to see whether one task is in *running* scheduling state, or not.
   (c) If one task is *running*, then the task is activated from its former state. The task executes until either (i) it blocks (for example, it tries to receive one given event, and that event is not available): in that case, the scheduler is called, or (ii) it can perform an instruction consuming cycles (e.g., writing a sample to a non-full channel).
   (d) When the scheduler is called, it first checks whether at least one task is *runnable*. If no task is runnable, the CPU goes *idle*. Otherwise, a scheduling algorithm - implemented in LOTOS - is called to select another task. Then, the state machine of that task may be called, and so on.
3. Once all CPUs have been served, a communication manager resolves inter-CPUs communication. Communications requests set-up by tasks in the previous cycle (i.e., all *read*, *write*, *notify events*, etc.) are actually granted in the current cycle. This ensures (i) that a sample written on a CPU during a cycle may not be read by another CPU in the same cycle, and (ii) that the order of CPU evaluation has no impact on results.

The $tf()$ function may also place debug information in the form of LOTOS actions performed at well-chosen points: actions to show scheduler data structures (e.g., list of *runnable* or *blocked* tasks), actions to monitor tasks states, actions to monitor the communication manager, etc.

Finally, $tf()$ has been defined with combinatory explosion in mind. Hence, $tf()$ tries to precompute potential synchronizations between LOTOS processes: if possible, these synchronizations are removed, and resulting processes put in sequence. Unfortunately, combinatory explosion may also be due to (i) non-deterministic elements: for example random, choice and temporal operators of tasks; (ii) Non-determinism in scheduling models: for example, in the round-robin scheduling policy, the possible indexes of tasks, in the tasks list. Abstractions are a key factor to reducing combinatory explosion. The basic intuition is to disregard software and hardware-related details irrelevant to properties to be evaluated (e.g., load on CPUs and buses). The next subsection is dedicated to abstractions.

4.4. Abstractions

4.4.1. Task abstraction (see Table I, column "LOTOS Semantics after mapping")

- **Communication operators**. These operations are given a cost (in clock cycles), and are executed by the execution manager along with the communication manager, to make requests on related buses. The cost in cycles depends of the hardware platform. For example, writing an 8-byte sample on a 32-bit processor takes two cycles. Also, the communication manager is involved for storing output samples, and for providing data to input operations. Note that these operations may be blocking, and so, the scheduling manager may also be involved.
- **Cost operators** are defined by the number of cycles depending on the hardware platform.
- **Other operators**: choice, loop, variable manipulation, etc. These operations are evaluated by the task execution manager. DIPLODOCUS control flow operators are assumed to be executed instantaneously, e.g. without consuming time. Extensive computations have to be modeled with dedicated *EXECx* operators.
- **Temporal operators**: They are defined by a number of time units.

4.4.2. CPU abstractions

- **Parameters of CPU**: Data size (used for communication in channels), size of default integer and floating point data (used for EXEC operations), cost for each EXECx instructions, pipeline size (used for calculating the penalty induced by miss branching), miss branching rate, data cache-miss ratio and penalty, time to enter/leave the idle mode, clock ratio.
- **The Operating System** is taken into account with scheduling algorithms (e.g., Preemptive priority-based, round-robin), switching time, synchronization management (events, requests) and communication delay (buffering for handling channels).

4.4.3. Communication abstractions (buses, memories)

Concurrent requests for data transfers on buses are processed according to an arbitration policy. The time a given transfer takes depends on the width of the bus. Bus arbitration is

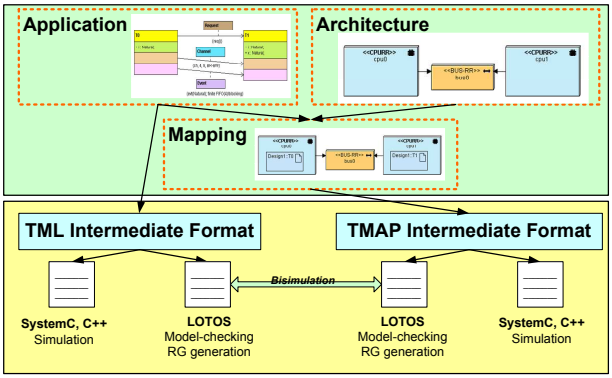

Figure 5. TTool for DIPLODOCUS: code generation capabilities

done on each cycle. Memory delays are modeled throughout bus latencies and cache-miss rates at CPU level, as proposed for the simulation semantics [3].

4.5. Formal verification

LOTOS specifications may be derived either from an application model, or from a mapping of applications onto a given architecture (Figure 1). Our methodology targets the proof of functional properties such as liveness and reachability of task operators. Cycle counts of different execution paths can also be computed, therefore enabling accurate schedulability analysis.

At application level, all tasks are defined to be concurrent. This concurrency is constrained during the mapping phase: buses and CPUs are shared resources. For example, two tasks mapped on the same CPU do not execute in parallel any more. An interesting property would be that formal traces obtained after mapping are a subset of formal traces obtained before mapping. That way, a mapping only eliminates application-level traces, without violating application-level safety properties (e.g., absence of deadlock). One of our ongoing efforts is to prove that scheduling and arbitration policies we have defined for CPUs and buses, respectively, preserve safety properties proved at application level. Optimally, a mapping is always *correct-by-construction* with respect to safety properties.

5. Toolkit

5.1. General overview

TTool [5] is an open-source toolkit initially developed for the TURTLE UML profile [28]. It now supports several other UML/SysML environments such as the *CTTool* profile [29] dedicated to large component-based software systems, the *DIPLODOCUS* profile [2] and the AVATAR profile [30] recently introduced for supporting usual methodological development phases of embedded systems software. TTool comprises diagramming facilities, formal code generators (LOTOS, etc.), analysis tools, diagram animation during simulations and prototyping code generation (C-POSIX).
In particular, TTool includes a Graphical User Interface for drawing DIPLODOCUS UML diagrams. From those diagrams, simulation or formal analysis (see Figure 5) may be performed. Underlying simulation and validation languages (e.g., LOTOS) are totally hidden to DIPLODOCUS users. From LOTOS specification, TTool relies on CADP [26] to generate reachability graphs that can be analyzed directly in TTool (in particular, to detect deadlock situations), to minimize them, and to compare them (bisimulations). Furthermore, TTool is endowed with very fast simulation capabilities [3]. Moreover, at simulation time, UML models can be animated [31] according to the simulation progress.

5.2. Property analysis with TTool and CADP

1. At first, an application is modeled (e.g., Smart Card functionalities: its main application, its TCP protocol stack, etc.). From that model a reachability graph is generated (denoted by rga), and model-checking techniques are used to prove a set P of properties on the application itself.
2. A hardware architecture is described in terms of CPUs, buses, etc..
3. From the mapping (tasks onto CPUs, etc.), a LOTOS specification is generated, and from that specification, CADP is used to obtain a reachability graph rgb. The following verification features are supported:

    - Minimizing the reachability graph to *tick* actions. From that minimization, the longest path of ticks is calculated, therefore resulting in performance figures for the application (e.g., *Worst Case Execution Time*).
    - Minimizing the reachability graph to *tick* and *transferOnBusX*. From that minimization, loads on buses can be deduced. Similar techniques can be used to compute CPU loads.
    - Comparing rgb and the reachability graph rga generated at application level. To that end, a toolkit integrated in TTool first modifies rgb so as to make rgb action names compatible with the one of rga, then CADP minimizes the resulting graphs: if it is proved that $rgb \subset rga$, then safety properties proved at application level are preserved.
    - Of course, all usual model-checking techniques can be directly applied to rga and rgb (e.g., using CADP).

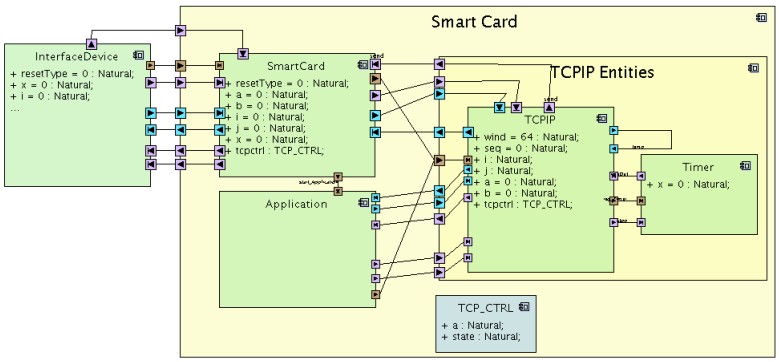

Figure 6. Component based Diagram of the Smart Card Application

6. Case Study

6.1. TCP/IP Protocol Implementation of a Smart Card

A smart card has the size of a credit card and is equipped with a microchip that securely stores data mainly used for identification purposes. The data may be periodically refreshed in order to maintain or enhance the functionality of the card. Smart cards are commonly used for telephone calling, electronic cash payments, establishing identity when logging on to some online account or when demanding public health services, paying small amounts of money (bus, parking, subway fees, etc.). Smart Cards comprise several hardware components like a microprocessor and different kinds of (non-)volatile memories (ROM, EEPROM, RAM Flash). Most smart card systems adhere to the ISO-7816 standard [32] which includes multiple parts defining for example physical characteristics, dimensions, involved protocols and other system properties. For the creation of our model, we mainly relied on the third part of the standard dealing with electronic signals and transmission protocols (ISO 7816-3).

6.2. Application model

The communication application has been decomposed into five DIPLODOCUS tasks (compare Figure 6) corresponding to the main functional blocks. A task called *InterfaceDevice* represents the terminal the smart card communicates with, for instance the card reader at a cash desk. Another task (*SmartCard*) models the transmission protocol defined in ISO 7816-3. The *Application* task models a basic exemplary application which merely makes use of the basic TCP services: establishing a session, sending some application data and finally tearing down the connection. The *TCP* task is structured according to the different phases of the TCP protocol like connection establishment, data transfer, connection termination. Last but not least, a *Timer* task may trigger time outs in the main *TCP* task.

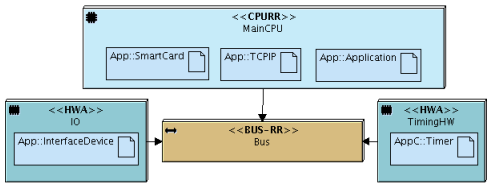

Figure 7. The Second Mapping for the Smart Card Application

Mapping	Min. Cycles	Max. Cycles
1	2474	3500
2	198	456

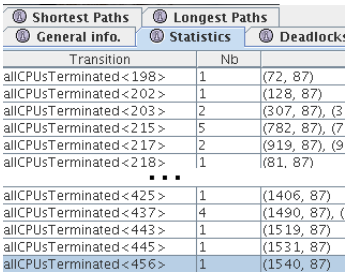

Table II. Verification Results: Performance and Reachability Graph Statistics

6.3. Architecture and mapping

To demonstrate the applicability of our methodology, two candidate architectures are experimented with. A first basic mapping consists of one single CPU onto which all tasks are mapped. A second option is to map the *SmartCard*, *TCP* and *Application* tasks on one CPU named *MainCPU*, and to provide a dedicated Hardware Accelerator to the *Timer* and *InterfaceDevice* task respectively (see Figure 7). The three CPUs are connected via an on chip bus. For space reasons, the mapping of channels, events and requests is omitted in the figure. In this second mapping, up to three tasks may execute concurrently and thus application level parallelism can be better exploited. Due to data and synchronization dependencies, further increasing the number of processing elements would not yield considerable performance improvements.

6.4. Property analysis

At first, performance measurements are carried out by transforming the UML model into its LOTOS equivalent. Subsequently, we rely on CADP to construct a reachability graph comprising all relevant system transitions. TTool permits the user to define state transitions of interest and to minimize the reachability graph accordingly. Depending on the respective verification objective, the user may select synchronization or execution related transitions or add special transitions every x clock cycles. By counting the latter for every possible system

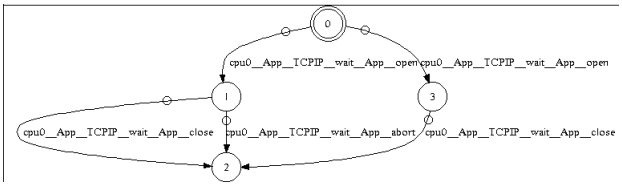

Figure 8. Minimized Reachability graph according to property

execution, upper and lower bounds on the execution time can be obtained. An analysis of the first mapping revealed a minimum of 2474 cycles and a maximum of 3500 cycles, for the second mapping we obtained a minimum of 198 cycles and a maximum of 456 cycles (cf. Table II). TTool provides means for analyzing the reachability graph generated by CADP so that these results can be easily deduced (see statistics depicted in Table II). In addition to performance measurements, we are interested in proving the property that every established connection is correctly relinquished (hence either closed or aborted). To that end, we minimize the reachability graph to transitions corresponding to *open*, *close* and *abort* events generated by the TCP task. That way, we do not even have to express the property in temporal logics. The proof can be simply conducted by observing that every branch of the reachability graph depicted in Figure 8 comprises either a *close* or an *abort* transition.

7. Conclusions and future work

The paper presents an environment - named DIPLODOCUS - for formal functional and performance analysis of complex embedded and distributed systems. A system is described with communicating tasks, hardware architectures and a mapping of tasks and channels onto hardware architectures. A formal semantics is provided to tasks, communication between tasks and hardware architectures, making it possible to perform formal analysis before and after mapping. Moreover, DIPLODOCUS has been implemented in a UML-based and open-source toolkit named TTool. Formal analysis can be performed with absolutely no expertise in formal techniques. DIPLODOCUS has been successfully used within numerous academic and industrial studies, e.g. for LTE-based systems [27] and for automotive systems [33]. As stated in section 4, abstractions are a key factor of our models in order to limit combinatory explosion (and to greatly increase simulation speed [3]): Data and control flow abstraction at application level (see 3.1) and generic hardware components at architecture level (see 3.2).
Trading off accuracy against model complexity of hardware components will remain subject to our research. For example, instruction cache-misses and data cache-misses have been accounted for by static probabilities so far. Indeed, as algorithmic details are represented by symbolic instructions, the real code of the application is not available. That is why state of the art cache models are not appropriate. Furthermore, the accuracy of bus and memory models shall be validated against a real embedded system. A fair comparison with a real implementation

shall therefore reveal whether a set of parameters can be found to limit the inaccuracy to a reasonable percentage. To simplify the modeling of systems making extensive use of DMA engines, a specific UML stereotype could be introduced. This way, the designer would not have to model DMA transfers explicitly using a dedicated execution unit.
While being already operational, our environment will be enhanced with three main features. First, we will define a refinement process from the application modeling step to the after-mapping step, in order to preserve properties proved at application level. Second, we intend to assess and adapt post-mapping hardware abstractions - e.g., the ones used for memories and buses - by confronting mapping results with real implementations. Third, effort will be dedicated to finding adequate trade-offs between the two extreme cases of formal verification and conventional simulation. Thus, simulation could cover several alternative system executions by exploiting indeterminsism inherent to the application model. In case the state space cannot be explored exhaustively, simulation could be guided by heuristics based on non-functional (CPU usage, bus loads, power consumption,...) or functional properties (for instance expressed in TEPE language [30]).

Acknowledgment

The authors would like to thank Chafic Jaber, who kindly granted us the permission to experiment with his smart card model.

REFERENCES

1. Kienhuis B, Deprettere EF, Wolf Pvd, Vissers KA. A methodology to design programmable embedded systems - the y-chart approach. *Embedded Processor Design Challenges: Systems, Architectures, Modeling, and Simulation - SAMOS*, Springer-Verlag: London, UK, UK, 2002; 18–37.
2. Apvrille L, *et al.*. A UML-based environment for system design space exploration. *13th IEEE International Conference on Electronics, Circuits and Systems (ICECS'2006)*, Nice, France, 2006.
3. Knorreck D, Apvrille L, Pacalet R. Fast simulation techniques for design space exploration. *Objects, Components, Models and Patterns*, *Lecture Notes in Business Information Processing*, vol. 33, Springer Berlin Heidelberg, 2009; 308–327, doi:10.1007/978-3-642-02571-6_18.
4. Apvrille L. TTool for DIPLODOCUS: An Environment for Design Space Exploration. *8th annual international conference on New Technologies of Distributed Systems (NOTERE'2008)*, Lyon, France, 2008.
5. TTool. *http://labsoc.comelec.enst.fr/turtle/ttool.html*.
6. ISO-LOTOS. A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. *Draft International Standard 8807, International Organization for Standardization - Information Processing Systems - Open Systems Interconnection*, Geneva, 1987.
7. Muhammad W, *et al.*. Abstract application modeling for system design space exploration. *Euromicro Conference on Digital System Design (DSD'06)*, Dubrovnik, Croatia, 2006.
8. Balarin F, *et al.*. *Hardware-Software Co-Design of Embedded Systems, The POLIS Approach*. 5 edn., KLUWER ACADEMIC PUBLISHERS, 2003.
9. Assayad I, Yovine S. A framework for modelling and performance analysis of multiprocessor embedded systems: Models and benefits. *Proceedings of the 8th conference on Nouvelles Technologies de la Distribution (NOTERE'2007)*, Marrakech, Marocco, 2007.
10. Avnit K, Sowmya A. A formal approach to design space exploration of protocol converters. *Design, Automation and Test in Europe Conference and Exhibition, 2009. DATE'09*, 2009; 129–134.

11. Marculescu R, Ogras UY, Zamora NH. Computation and communication refinement for multiprocessor soc design: A system-level perspective. *ACM Trans. Des. Autom. Electron. Syst.* 2006; **11**(3):564–592, doi:http://doi.acm.org/10.1145/1142980.1142983.
12. Hendriks M, Verhoef M. Timed automata based analysis of embedded system architectures. *Parallel and Distributed Processing Symposium, 2006. IPDPS 2006. 20th International*, 2006; 8 pp.–, doi: 10.1109/IPDPS.2006.1639422.
13. Watanabe Y. Metropolis : An integrated environment for electronic system design. Cadence Berkeley labs, 2001.
14. Wolf PVD, *et al.*. A methodology for architecture exploration of heterogeneous signal processing systems. *1999 IEEE Workshop on Signal Processing Systems (SiPS99)*, 1999.
15. Chatelain A, *et al.*. High-level architectural co-simulation using Esterel and C. *Proc. of IEEE/ACM symposium on Hardware/software codesign*, 2001.
16. Schattkowsky T, *et al.*. A model-based approach for executable specifications on recon figurable hardware. *Design, Automation and Test in Europe Conference and Exhibition, 2005. DATE'05*, 2005; 692–697.
17. Kukkala P, *et al.*. Performance Modeling and Reporting for the UML 2.0 Design of Embedded Systems. *Proc. of the 2005 International Symposium on System-on-Chip*, 2005; 50–53.
18. Vidal J, de Lamotte F, Gogniat G, Soulard P, Diguet JP. A co-design approach for embedded system modeling and code generation with uml and marte. *Design, Automation and Test in Europe Conference and Exhibition, 2009. DATE'09*, 2009; 226–231.
19. Woodside M, Petriu DC, Petriu DB, Shen H, Israr T, Merseguer J. Performance by unified model analysis (puma). *WOSP '05: Proceedings of the 5th international workshop on Software and performance*, ACM: New York, NY, USA, 2005; 1–12, doi:http://doi.acm.org/10.1145/1071021.1071022.
20. Viehl A, Schonwald T, Bringmann O, Rosenstiel W. Formal performance analysis and simulation of UML/sysml models for esl design. *Design, Automation and Test in Europe Conference and Exhibition, 2006. DATE'06* March 2006; **1**:1–6.
21. Ristau B, Limberg T, Fettweis G. A mapping framework for guided design space exploration of heterogeneous mp-socs. *Design, Automation and Test in Europe Conference and Exhibition, 2008. DATE'08* March 2008; :780–783doi:10.1109/DATE.2008.4484910.
22. Hamann A, Jersak M, Richter K, Ernst R. A framework for modular analysis and exploration of heterogeneous embedded systems. *Real-Time Syst.* 2006; **33**(1-3):101–137, doi:http://dx.doi.org/10.1007/s11241-006-6884-x.
23. Henia R, Hamann A, Jersak M, Racu R, Richter K, Ernst R. System level performance analysis - the symta/s approach. *Computers and Digital Techniques, IEE Proceedings* - Mar 2005; **152**(2):148–166, doi:10.1049/ip-cdt:20045088.
24. Wodey P, Camarroque G, Baray F, Hersemeule R, Cousin JP. LOTOS code generation for model checking of stbus based soc: the stbus interconnection. *This paper appears in: Formal Methods and Models for Co-Design, 2003. MEMOCODE '03. Proceedings. First ACM and IEEE International Conference on* June 2003; :204–213.
25. OMG. UML 2.0 Superstructure Specification. *http://www.omg.org/docs/ptc/03-08-02.pdf*, Geneva, 2003.
26. Garavel H, Lang F, Mateescu R, Serwe W. CADP 2006: A Toolbox for the Construction and Analysis of Distributed Processes. *Computer Aided Verification (CAV'2007)*, vol. 4590, Berlin Germany, 2007; 158–163.
27. Jaber C, Kanstein A, Apvrille L, Baghdadi A, Moenner PL, Pacalet R. High-level system modeling for rapid hw/sw architecture exploration. *Proc. of the 20th IEEE/IFIP International Symposium on Rapid System Prototyping (RSP'2009)*, 2009.
28. Apvrille L, *et al.*. TURTLE: A Real-Time UML Profile Supported by a Formal Validation Toolkit. *IEEE transactions on Software Engineering*, vol. 30, 2004; 473–487.
29. Ahumada S, *et al.*. Specifying Fractal and GCM components with UML. *XXVI International Conference of the Chilean Computer Science Society (SCCC'07)*, Iquique, Chile, 2007.
30. Knorreck D, Apvrille L, Saqui-Sannes Pd. TEPE: a sysml language for time-constrained property modeling and formal verification. *Proceedings of the third IEEE International workshop UML and Formal Methods - ULM&FM'2010*, IEEE, 2010.
31. Knorreck D, Apvrille L, Pacalet R. An interactive system level simulation environment for Systems on Chip. *ERTSS - Embedded Real Time Software and Systems*, 2010.
32. Iso 7816 smart card standard:. Http://www.cardwerk.com/smartcards/smartcard_standard_ISO7816.aspx.
33. Idrees M, Roudier Y, Apvrille L. A framework towards the efficient identification and modelling of security requirements. *Fifth Conference on the Security of Network Architecture and Information Systems (SAR-SSI 2010)*, Menton, France, 2010.