

# Virtual Execution Platforms for Mixed-Time-Criticality Systems: The CompSOC Architecture and Design Flow\*

Kees Goossens,<sup>1</sup> Arnaldo Azevedo,<sup>2</sup> Karthik Chandrasekar,<sup>2</sup> Manil Dev Gomony,<sup>1</sup>  
Sven Goossens,<sup>1</sup> Martijn Koedam,<sup>1</sup> Yonghui Li,<sup>1</sup> Davit Mirzoyan,<sup>2</sup>  
Anca Molnos,<sup>2</sup> Ashkan Beyranvand Nejad,<sup>2</sup> Andrew Nelson,<sup>2</sup> Shubhendu Sinha<sup>1</sup>  
<sup>1</sup> Eindhoven University of Technology <sup>2</sup> Delft University of Technology

## ABSTRACT

Systems on chip (SOC) contain multiple concurrent applications with different time criticality (firm, soft, non real-time). As a result, they are often developed by different teams or companies, with different models of computation (MOC) such as dataflow, Kahn process networks (KPN), or time-triggered (TT). SOC functionality and (real-time) performance is verified after all applications have been integrated.

In this paper we propose the CompSOC platform and design flows that offers a *virtual execution platform per application, to allow independent design, verification, and execution*. We introduce the composability and predictability concepts, why they help, and how they are implemented in the different resources of the CompSOC architecture. We define a design flow that allows real-time cyclo-static dataflow (CSDF) applications to be automatically mapped, verified, and executed. Mapping and analysis of KPN and TT applications is not automated but they do run composably in their allocated virtual platforms.

Although most of the techniques used here have been published in isolation, this paper is the first comprehensive overview of the CompSOC approach. Moreover, three new case studies illustrate all claimed benefits: 1) An example firm-real-time CSDF H.263 decoder is automatically mapped and verified. 2) Applications with different models of computation (CSDF and TT) run composably. 3) Adaptive soft-real-time applications execute composably and can hence be verified independently by simulation.

## 1. INTRODUCTION

Systems-on-chip (SOC) complexity grows as more applications are integrated in the same system. Often *applications* are dynamically started and stopped, leading to many *use cases*, i.e. sets of concurrently executing applications. Ap-

\*Copyright is held by the authors.

plications have specific *characteristics*, such as being control- or data-oriented, being more or less time critical (firm FRT, soft SRT, or non real-time NRT), and the extent to which they are adaptive or scalable. The system as a whole may have average or peak power or energy requirements depending on whether it is powered by battery, tethered, or energy-scavenging.

In a SOC applications are usually implemented with multiple communicating tasks, and are executed on an often heterogeneous set of processors and accelerators, memory hierarchy, and advanced on-chip interconnect such as a network on a chip (NOC). Since SOCs are always constrained in terms of area, the processors, interconnect, and memory resources must be shared between applications.

### 1.1 Problem Statement

Integrating a set of applications with different characteristics and requirements on a multi-core SOC is challenging for a number of reasons. First, *resource sharing causes interference* between applications, making their temporal behaviours inter-dependent. In case of an adaptive application the functional behaviour may change too. For example, when decoding H263 video, a B frame instead of an I frame may lead to a shorter computation, fewer memory accesses, lower interconnect throughput, and a different processor task schedule. This affects the performance of other applications (e.g. the audio processing) either positively (e.g. by finishing earlier) or negatively (e.g. in the presence of scheduling anomalies).

Second, SOCs applications are *developed by different groups* in a company, or even by different companies. Often executables are delivered and integrated, since source code is proprietary and not shared for intellectual property reasons. Individual applications are verified on a system that is different from the final system, since other applications are usually absent. When all applications have been delivered and integrated, the system as a whole is verified. Unfortunately each application may fail at this point, since interference from other applications often leads to previously untested conditions. Use-case verification then becomes a *circular process* that must be repeated if an application is added, removed, or modified [19].

Third, applications are *designed and programmed using different methodologies, models of computation (MOC) or programming models, and design flows*. FRT applications could

be programmed and analysed using dataflow, time-triggered, or event-based models of computation. Performance verification would be based on *worst-case* behaviour and formal analysis. NRT applications, on the other hand, could be programmed using threads or tasks using distributed shared memory to communicate. Verification could focus on functional correctness, and be based on *average-case* simulation. SRT applications occupy the middle of the spectrum and could use techniques of both.

The performance verification of *adaptive applications* is yet another case. For example, the quality of a SRT adaptive video decoder is measured with the signal-to-noise ratio (SNR) and deadline misses. The SNR depends on how much work (computation cycles at a certain frequency) and time (before a display deadline) the decoder allocated to decode a compressed video frame. If there is interference from other applications in terms of actual number of computation cycles or scheduling times, then the SNR will change. Hence, while FRT applications and SRT applications are verified on the basis of their worst-case and average-case behaviours, SRT adaptive applications are verified on the basis of their *actual-case* behaviour.

Finally, different application characteristics and MOCs require *different resource-management policies*. For example, different scheduling policies (preemptive vs. non-preemptive, work conserving or not, with jitter constraints and fairness properties, etc.) lead to different worst-case, average-case, and actual-case latency and throughput, but also affect buffer requirements. Regarding power/energy management, FRT applications must use conservative policies, where S/NRT policies can be speculative. It is challenging to map the different scheduling requirements of all applications on a few common schedulers (such as static priority or round robin) that many SOC's offer.

## 1.2 Composable and Predictable Virtual Platforms

The CompSOC platform [17] addresses these problems by offering a *virtual execution platform per application, to allow independent design, verification, and execution*. An automatic *design flow* is available for FRT dataflow applications, and S/NRT dataflow and Kahn process network (KPN) applications. Time-triggered applications are supported but not yet automated. CompSOC relies on two complexity-reducing concepts: *composability* and *predictability*. Composable virtual platforms are completely isolated (partitioned) and cannot affect each other by even a single clock cycle. They are hence *virtualised in terms of actual execution time*, which enables independent verification of S/NRT applications. Our use of composability extends [19] to multiple applications. Each virtual platform is also predictable, which means that it can be *virtualised in terms of performance bounds* such as worst-case execution time. This enables independent formal analysis of FRT applications.

In the next section we describe in more detail how composability and predictability solve or mitigate the problems introduced. Following this, we define the platform architecture (Section 3), describe how resource sharing is dealt with (Sections 4-5) and elaborate the design flow (Section 6). In Section 7 we illustrate our approach with a H.263 video de-

coder implemented as a FRT dataflow application, a FRT time-triggered application, and as a SRT adaptive dataflow application. In all cases, it runs together with other applications, each in their own virtual platform. After discussing related work in Section 8 we conclude.

## 2. COMPSOC CONCEPTS

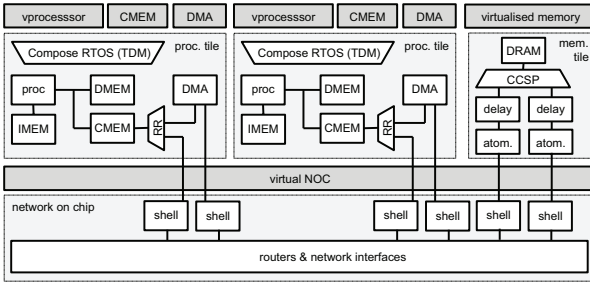
In essence a virtual execution platform is a set of *resource budgets*; intuitively, a percentage of the resource capacity, enforced by a budget scheduler. Resources include processors, NOC, distributed on-chip SRAMs, and off-chip DRAM memory. In other words, each physical resource is divided in smaller virtual resources that are handed out to different applications. Resource budgets are computed at design time (see Section 6), and programmed at run time. Conceptually, composability is ensured on each resource by using *preemptive time-division multiplexing (TDM) between applications*, avoiding all interference between applications. Within a virtual platform, each application can further budget and schedule each virtual resource using whatever scheduling policy is appropriate for the application and resource. CompSOC therefore employs *two-level scheduling* on (selected) resources: *composable and predictable between applications, and application-specific within an application*. However, several different techniques are used to implement these concepts in a concrete platform, as discussed in Sections 3-4. First we discuss how our concepts address the problems raised in the introduction.

### 2.1 How Composability and Predictability Address The Problems

Composability, i.e. virtualisation in terms of actual execution time, has several important consequences. First, the behaviour of an application depends only on its own virtual platform. From the application writer's perspective, sharing with and interference from other applications have disappeared. Verification no longer depends on others, and has become a non-circular process. An application is designed and verified in isolation, and no reverification is required after integration in a larger system.

Second, each application may have its own performance verification methods. Since the *actual-case* behaviour is independent of other applications that may or may not be running at the same time, so is the *average-case* performance of an (adaptive) S/NRT application. Clearly, the *worst-case* behaviour of FRT applications is independent too. As a result, the verification of an application, of whatever criticality, depends only on its own virtual platform.

Third, application-specific scheduling policies are supported through CompSOC's two-level scheduling. Each virtual platform is composable and predictable, and within it, each application defines its own scheduler. Specifically, on a processor, the first-level TDM scheduler multiplexes applications, and a second-level application-specific scheduler multiplexes actors/processes/tasks of an application in a way that fits its characteristics and MOC. For example, a FRT application can use cooperative non-preemptive priority-based scheduling between dataflow actors, a SRT application could use preemptive earliest deadline first (EDF) between tasks, and a NRT application preemptive round robin between KPN



**Figure 1: Example instance of the CompSOC architecture template.**

processes. Similarly, each virtual processor has its own virtual battery, managed by each application’s power/energy manager.

Each virtual platform, i.e. a set of resource budgets, is predictable, which means it can be characterised by performance bounds on throughput and latency. For FRT applications, these bounds may be combined with the bounds given by its own predictable intra-application scheduler. In our FRT dataflow design flow, described in Section 6, an actor budget (intuitively, a percentage of the resource capacity) is the product of the virtual-platform budget and the actor budget within the virtual platform.

We conclude that composability and predictability are complementary concepts that both solve important parts of the verification problem for mixed time-criticality systems, and provide a complete solution when combined.

### 3. NON-SHARED PLATFORM

It is essential that the platform architecture is designed with composability and predictability in mind, as it is not possible to graft these qualities on later. We first explain our user-resource model, and the consequences of composability and predictability on it. We then describe the platform components, and how they are combined into a larger system, and used at run time. A CompSOC platform consists of *resources* that are used by *users*, according to an allocated budget. Specifically, a resource is one of (cf. Figure 1):

1. A processor with its local instruction memory (IMEM) and data memory (DMEM), used by one or more computation tasks.
2. A direct memory access unit (DMA), used by a single DMA task.
3. A communication memory (CMEM), used by one or more storage tasks.
4. A network on a chip (NOC), used by one or more communication tasks.
5. An on-chip SRAM memory, used by one or more storage tasks.
6. An off-chip DRAM memory, used by one or more storage tasks.

As described in more detail in Section 6, for mapping and performance analysis purposes, users represent the resource usage of a “real” application tasks and channels. Appli-

cation tasks are split in computation, DMA, and storage tasks, and channels are split in communication and storage tasks. A processor tile contains a single processor resource including its IMEM and DMEM, and one or more DMAs with their CMEMs. A S/DRAM tile contains the S/DRAM memory. To share the resource between multiple users each tile also contains additional hardware/software, e.g. buffering, multiplexing and scheduling. A platform instance contains any number of processor, SRAM, and DRAM tiles, and a single NOC connecting them all. For physical scalability, CompSOC uses GALS (globally asynchronous locally synchronous), i.e. each tile can operate on its own (scalable) clock frequency, and the NOC has its own single clock. Since tiles operate asynchronously, there is no shared global notion of time, and each resource scheduler operates asynchronously.

For simplicity, we first show how an application executes on a platform, where none of the resources are shared. Each computation task of the application uses a single processor, and its instructions and data must fit in the IMEM and DMEM, respectively, and be loaded before the task executes. The processor has no caches, and the application tasks cannot use interrupts (although its scheduler can). A task communicates with other tasks using distributed shared memory. To communicate with another task on the same tile it uses local DMEM, for another tile it uses the CMEM on the remote tile. A third alternative is to use a remote shared memory, either SRAM or DRAM. A single global address space is used for all memories, except that remote IMEMs and DMEMs are not visible to a processor. The CMEM and DMA are used to access a memory outside of the tile. To write to a remote memory, a task writes the data in its CMEM and then uses a communication library to instruct the DMA to copy the data from the local CMEM to the remote memory. Reading is similar, and copies remote data into the local CMEM. A task and its DMA(s) operate in parallel. The read and write requests of the DMA are transported by a NOC to the remote memories, where they are executed, after which requested responses return over the NOC.

### 4. SHARING IN THE COMPSOC PLATFORM

For composability and predictability, each user is statically bound to one resource. Multiple users can be bound to a single resource, except for DMAs that are not shared. A user cannot use multiple resources, and resources cannot use each other; we discuss this further in Section 6. As a result, we can discuss the sharing of each resource in isolation. Composable sharing requires that the service an application receives is independent of others, and predictable sharing requires a bound on the minimum service an application receives. For each resource, predictable and/or composible scheduling require at least some of the following ingredients: 4.1) bounded scheduling interval, 4.2) fixed scheduling interval, 4.3) a neutral resource state between scheduling intervals, and 4.4/4.5) an appropriate scheduler. The concepts behind these ingredients are detailed in [6]. Next, we discuss each of these in turn.

## 4.1 Bounded Scheduling Interval

Service is given out in *service units*, such as real-time operating system (RTOS) slots or NOC flits, that take a *bounded scheduling interval* to execute. (For simplicity we ignore pipelining here; see [6] for details.) A bounded scheduling interval is ensured either when each user of the resource has a *bounded execution time*, or each resource can be *preempted in a bounded time*. In the former case, cooperative (non-preemptive) scheduling may be used, and the scheduling interval is the worst-case execution time of any user. However, note that in a mixed-criticality system the worst-case execution time for some users may not exist. A NRT task on a processor may never finish (perhaps intentionally), or memory transactions may be infinitely long (as allowed by, e.g., DTL and AHB protocols). This effectively locks up the resource for a single user, which breaks predictability and composability.

In the case of preemption, the scheduling interval is independent of the (worst-case) execution time of users, which is advantageous. However, it may not be possible to preempt a user in a bounded time. Many processors, such as Xilinx Microblaze and simpler ARM processors do not serve interrupts when there is an outstanding read or write transaction to a non-local memory. In the context of CompSOC, IMEM, DMEM, and CMEMs are local (on a local memory bus without handshake / flow control), and serve read/write transactions in a single cycle. Other remote CMEM or S/DRAM memories on the NOC are accessed via a handshaked PLB bus, and this may take an arbitrarily long (possibly infinite) time, depending on the scheduling on the NOC and remote memory. Hence CompSOC uses a DMA with a CMEM to communicate to remote memories, with the processor polling DMA for completion. This results in a bounded scheduling interval because the DMA works in parallel with the processor, and polling to the DMA is local and interruptible.

## 4.2 Fixed Scheduling Interval

Regarding the duration of a scheduling interval, most often it is actually *not* fixed. On a processor, although the interrupt timer may be cycle accurate, the interrupt service latency (in seconds) depends on the instruction that is being executed, and on the frequency the processor runs at. Similarly, reading or writing to a DRAM take a different number of cycles. There are two ways of dealing with this. Either all scheduling intervals are *made the same length* by padding with idle cycles or by disabling the clock of the resource to make each scheduling interval to be as long as the worst-case scheduling interval. In CompSOC, NOC time slots are padded, whereas the processor RTOS time slots are made longer with clock gating. Alternatively, the scheduling interval is *left variable*, and the interference arising from the variable scheduling interval and from variations in timing due to scheduling are removed by delaying to the worst-case interference. This approach is taken for the DRAM. Similar approaches are used in [19, 21, 25]. We introduce the details when discussing the individual resources.

## 4.3 Neutral State

Often resources have internal state that persists between the execution of successive service units. This state often influences the (timing of) execution. When scheduling a resource it is advantageous to return it to a *neutral state* between two

service units. Composability requires that the execution of a user on a resource is independent of others. In particular, when a user starts or resumes execution the state of the resource must be independent of others. Often this is not the case, such as for processor pipeline, branch predictor, and caches, the state of which depends on the instructions of the application(s) that were executed before. The timing behaviour thus depends on other applications. Similarly, for DRAM memories, the open/closed status of pages at the end of a scheduling interval strongly affect the execution time of the following scheduling interval. Since all CompSOC resources are shared compositably, we therefore *reset or restore the resource state between scheduling intervals*. In particular, for processors, branch prediction is turned off, and the pipeline timing and operating frequency are restored to the state when the user was previously interrupted and swapped out, and no caches are used. The DRAM uses close-page policy to enforce a neutral state.

## 4.4 Predictable Scheduling

A predictable scheduler must provide a bound on the service that is delivered to each user of a resource. For this, the scheduling interval must be bounded but not necessarily be fixed. Similarly, a neutral state is not required, as long as the interference of previous users of the resource can be bounded. However, since all resources are shared both compositably and predictably, CompSOC returns every resource to a neutral state between scheduling intervals (except the DMA and CMEM, see below). In terms of scheduling policy, for predictability any *budget scheduler* will do. A budget scheduler guarantees a user a minimum number of service units in a given time frame. Essentially, any scheduler that is free of starvation suffices. In CompSOC we use latency-rate arbiters, such as as TDM, credit-controller static priority (CCSP) [38], round-robin (RR), and static-order (SO) scheduling, depending on the resource.

## 4.5 Composable Scheduling

Finally, the resource scheduler determines in what order the bounded scheduling intervals are given to applications. For composability, this must be done such that the behaviour of an application is not affected by other applications. Scheduling cannot, therefore, take information that depends on (other) applications, such as availability of data or task readiness, into account. The obvious solution, which we use on the processor and NOC, is to use time-division multiplexing (TDM), which is static and not work-conserving. For the same reasons, TTA [19], MERASA [45], and [35] use TDM.

However, TDM inversely couples latency and rate, and couples the scheduling interval and frame size to the latency. A user that requires low-latency service must be overallocated in terms of bandwidth. Overallocation is not acceptable for scarce resources, since a composable, i.e. non-work-conserving, scheduler cannot give unused capacity to other applications that may want it. For this reason, a second technique is also used. *Any predictable scheduler can be made composable* by delaying the result of the execution of a user on a resource until its worst-case interference [4]. In particular, in the case of DRAM, the response of a read transaction is delayed as if it had experienced its worst-case interference from other users whether they are present or



not. Since the worst-case interference of others is constant, the net result is a composable behaviour.

## 4.6 Mixed-Time-Criticality Multi-MOC Applications

Given that each resource can be shared, it is useful to see how different applications, with different time criticality (F/S/NRT) and different models of computation (MOC), are mapped on it and run concurrently. In fact, the CompSOC platform is independent of the MOC of the applications that run on it [9]: a single platform execution model supports multiple models of computation. Processors execute generic tasks, and DMAs and NOC transport generic data.

To run a (FRT) *cyclo-static dataflow (CSDF) application* on CompSOC, its actors are automatically converted to tasks by adding a wrapper that receives tokens before firing and sends them after firing, using the C-HEAP FIFO communication library [34] that uses the DMAs. The user can choose a given (F/S/NRT) scheduler that implements actor firing rules (e.g. preemptive round robin, non-preemptive static order), or supply his/her own. Similarly, a conservative FRT power management policy is provided [30]. As described in Section 6, CSDF applications are automatically mapped on the platform, and are FRT, (given, of course, appropriate scheduler and power manager, and worst-case execution times for actors).

The processes of (SRT) *Kahn process network (KPN) applications*, are normal CompOSe RTOS [16] tasks that use a blocking read/write API on FIFO channels, also implemented with [34]. A mapping of processes and tasks on the platform is automatically generated, along with default schedulers, but without real-time guarantees. The user can use a built-in speculative SRT power manager, or supply another.

Since dataflow and KPN applications work on the basis on availability of data/space in their channels, the entire CompSOC platform uses flow control / backpressure. All resources (processors, NOC, DMA, memories) use handshakes to transfer data, avoiding channel overflow or underflow. Equally important this enables work conserving behaviour within an application, unlike time-triggered architectures [19]: if data arrives early (when there is slack in the virtual platform), the application can work ahead. Conversely if data arrives later than expected, the consumer waits until it arrives.

Hence, dataflow actors and KPN processes tend to be scheduled on the basis of data availability (using firing rules and information regarding blocking on channels) rather than on the basis of time. Tasks of *time-triggered applications*, on the other hand, use a (often static periodic) schedule in time, or use a schedule based on task deadlines. (Non)-preemptive time-triggered task scheduling within an application is supported. Mapping is like for KPN applications.

## 5. SHARING PER RESOURCE TYPE

To recapitulate, predictable sharing requires 4.1) a bounded scheduling interval, and 4.4) a predictable scheduler. Composable sharing is additionally either achieved with a 4.2)

fixed scheduling interval, and 4.3) neutral state, with a 4.5) TDM scheduler, or by 4.5) converting predictable scheduler to be composable. In the following (see Table 1), we describe for each resource what its users are, what its service unit is, and how its scheduling interval is (made) bounded, whether its scheduling interval is fixed, and how its neutral state is enforced. The composable and predictable inter-application scheduling and (optionally predictable) intra-application scheduling policies are also discussed. As we shall see, different resource characteristics (service unit size, size of state, abundant or not) lead to the use of different combinations of techniques.

## 5.1 Processor

A processor serves computation tasks. The CompOSe RTOS [16] runs on Microblaze and ARM processors and multiplexes multiple tasks on the same hardware using preemptive scheduling since tasks may have an unbounded or infinite execution time. Preemption is implemented using an interrupt generated by a timer that runs at the maximum frequency of the tile. Although the interrupt is generated at fixed points in time, the interrupt service routine (ISR) takes a variable number of cycles (e.g. depending on the currently executing instruction). Moreover, the time (rather than cycles) to serve the interrupt varies too, since each task can run at its own frequency. On receiving an interrupt, the ISR first sets the frequency to the maximum, saves the stack, records the ISR service delay due to the processor pipeline emptying, and jumps to the RTOS. Scheduling and housekeeping of the RTOS take a variable amount of time too. Therefore, to enforce a fixed scheduling interval, the processor gates its clock and instructs its clock generator to re-enable the clock at a fixed time at the end of the scheduling interval. In this way, a complete RTOS slot comprising task execution, interrupt handling, and RTOS scheduling, each of which may take a variable amount of time, is served with a constant scheduling interval. In Table 1 this is indicated by the worst-case scheduling interval (WCSI) delay. The next task slot length is reduced by the cycles that it took to empty the pipeline when the interrupt arrived, otherwise the task would receive more cycles than budgeted. The neutral state at the start of a scheduling interval is enforced by restoring the stack, and by enabling the clock at the task's frequency. Since the frequency may be changed several times in each RTOS slot, it must be quite fast (in the order of tens of clock cycles). Technology such as [26, 42] offers this capability for ASIC implementation. The CompSOC FPGA prototype models frequency scaling by clock subsampling [30].

A fixed scheduling interval coupled with a TDM scheduler results in composable scheduling. Since a processor is an expensive resource, its utilisation is important. As we saw before, composable scheduling cannot be work-conserving, and any slack is therefore lost. However, by offering two levels of scheduling, first between applications and then between tasks within the application, only slack between applications is lost. In fact, during unused slots the processor can be switched off, reducing energy/power. Within a single application any application-specific (non)-work-conserving scheduler can be used. The cost of using two levels of scheduling implemented in software is amortised over the relatively large scheduling interval (at least 10,000 cycles at maximum tile frequency). Similarly, each application can have its own

**Table 1: Service units and scheduling levels per resource. SI is scheduling interval.**

resource	user	service unit	bounded SI through	fixed SI size (by)	SI neutral state	inter-app. sched. (comp. & pred.)	intra-app. sched.
processor	proc. task	OS slot	interrupt	WCSI-delay	pipel. delay & restore freq.	TDM	any
DMA	DMA task	transaction	comm. lib.	no	no state	—	—
CMEM	storage task	transaction	DMA	no	—	—	RR
NOC	comm. task	flit	shell	padding	no state	TDM	—
SRAM	storage task	word	atomiser	yes	no state	TDM	—
SRAM	storage task	word	atomiser	yes	no state	any pred. & WCI-delay	—
DRAM	storage task	pattern	atomiser	no	closed pages	any pred. & WCI-delay	—

power manager that computes and sets its frequency, as part of the scheduler or its tasks. Both intra-application scheduling and power management are user-supplied code and run in user time, i.e. a misbehaving scheduler or power manager cannot affect the RTOS or other applications.

## 5.2 DMA and CMEM

A DMA serves DMA tasks. When a computation task wishes to communicate with a remote memory it uses a DMA communication library to instruct the DMA to copy (a finite amount of) data to/from the corresponding local CMEM. A computation task and its DMA task(s) run in parallel, and do not affect each other because they use distinct ports on the CMEM. When needed, the computation task polls for DMA completion, which ensures that a computation task can always be interrupted in a short time, as described previously. The DMA service unit is a single DMA transaction, which results in a variable, but bounded, number of read/write transactions to the NOC, and then a remote memory. When a remote DMA accesses a local CMEM to read or write (synchronisation) data, it uses the same port on the CMEM as the local DMA. Since only tasks from the same application communicate, this sharing needs to be at most predictable and not composable. For this reason, a round-robin scheduler is used here. The DMA has no state after completion, and the round-robin scheduler does not need to have a neutral state.

## 5.3 Network on Chip

The  $\mathcal{A}$ ethereal NOC [14] serves communication tasks. A communication task (usually called a connection) is a virtual wire over which the NOC transports requests from master (DMA) to slave (memory), and optional responses from slave to master. Requests are received from the master using a DTL handshaked protocol. In theory, a transaction may be infinitely long, or may not arrive in a finite amount of time (since each word of the transaction is separately handshaked). A “shell” hardware block serialises the command and data groups of the request to a stream of words. The NOC executes a global TDM schedule of flit service units, with a scheduling interval of 3 cycles. Whenever a network interface port can inject a flit in the NOC, it consumes 3 words from the shell, or else pads with dummy words. This guarantees a fixed service unit. When flits arrive in the receiving network interface over their allocated connection, dummy words are removed, and the result deserialised to a DTL request. Responses are handled similarly.

From a scheduling perspective the NOC behaves as a single deeply-pipelined resource, even though it is made up of many routers and network interfaces. Logically it executes a global static TDM schedule. Flits are injected such that they never wait in the NOC to minimise buffering since this is the main contributor to the area cost. Introducing a second level of intra-application arbitration would require an additional level of buffering (per-application virtual circuits), which is prohibitively expensive. Additionally, two-level scheduling at wire-speed is hard, and in any case, NOC bandwidth is essentially relatively cheap. For these reasons, all NOC communication tasks (connections), whether they belong to the same application or not are scheduled compositably (and predictably).

## 5.4 SRAM

SRAM memory is the simplest resource, and serves transactions that are generated by storage tasks (or requestors). Transactions are chopped in single-word transactions; e.g. a 8-word write is converted in 8 single-word writes. The service unit is a single-word transaction, and the scheduling interval one cycle. Regarding atomicity of transactions, note that this chopping may interleave execution of multi-word transactions. While perhaps unexpected for software engineers, on-chip communication protocols, such as DTL and AXI, explicitly allow this, as they only require byte-level atomicity. Higher-level software synchronisation, such as C-HEAP, must be used to avoid unexpected results. In our current CompSOC platform, a single-level TDM arbiter is used most frequently, since bandwidth overallocation tends not to be a problem. However, all the techniques of the DRAM can also be used for the SRAM.

## 5.5 DRAM

A DRAM memory serves transactions of storage tasks. CompSOC uses the Predator real-time DRAM memory controller [3]. Like for the SRAM, incoming transactions are chopped into fixed-size transactions. However, DRAMs are only efficient with large access granularity, of at least 64 or 128 bytes. We refer to [5] for more information on how to select an appropriate access granularity, and corresponding DRAM configuration. The time to serve a read or write transaction on a DRAM varies greatly depending on the current state of the DRAM, i.e. whether the requested page is open or not, and whether the previous transaction was also a read or write. A simple worst-case approach would yield a maximum guaranteeable bandwidth of about 20%, which is unacceptably low: at least 80% utilisation is required. For this reason, Predator

uses service units based on patterns, which are precomputed sequences of memory commands corresponding to a read, write, or read/write or write/read switch. These patterns leave the DRAM in a neutral state (all pages are closed), and a tight scheduling interval can be computed. Like the NOC, the DRAM is pipelined, and the scheduling interval is less than the service time.

TDM is not an acceptable scheduler for the DRAM for two reasons. First, the service units do not have the same length, since that would reduce utilisation of each service unit too much. Using TDM with variable scheduling intervals is not composable. Second, TDM couples latency and rate, which means that low-latency requestors require overallocation, which when unused cannot be given to other applications through work-conserving scheduling. We therefore use the credit-controller static priority (CCSP) [38] scheduler, which does not have these defects. However, as CCSP is only predictable, we convert it in a composable arbiter by delaying each response until the time it would have had with worst-case interference [4]. In Table 1 this is indicated with WCI-delay. In essence, by simulating that each request always encounters worst-case interference, its behaviour is independent of other requestors whether they are present or not. As a result, without overallocation the DRAM is efficiently used. For S/NRT applications, the actual/average-case DRAM performance has degraded to the worst-case performance. But there are no drawbacks for FRT applications, since the worst-case performance is unaffected.

## 6. MIXED-CRITICALITY DESIGN FLOW

Programming a multi-processor SOC is hard at the best of times, but additionally guaranteeing real-time performance is even more challenging. For this reason, given a FRT cyclo-static dataflow (CSDF) [10] application with worst-case execution time and memory usage per actor, the CompSOC design flow automatically computes the mapping of users to resources and the scheduling budgets of each user on each resource, such that required end-to-end, i.e. first to last actor, throughput and latency are guaranteed. All required C wrappers for actors, drivers that configure the resources and their schedulers at run time, as well as boot code to load the application(s) at run time are automatically generated. As mentioned before, KPN, time-triggered, and applications without a specific MOC must be mapped and compiled by the user, but are loaded automatically at run time. In the remainder of this section we focus on the FRT design flow, since it is the hardest, and is a superset of what is required for other applications. Lacking space for a full formal description, we aim to convey the intuition, and refer to [17, 29, 39] for technical details.

At the basis of the FRT design flow are the notions of user, resource, mapping, and budget. Each user (i.e. computation, DMA, communication, or storage task) is bound to a single resource (i.e. processor, DMA, NOC, or CMEM/-SRAM/DRAM, respectively). Multiple users can be bound to a single resource, except for DMAs that are not shared. A user cannot use multiple resources, and resources cannot use each other. The essence of these restrictions is that the performance of a user (e.g. computation or storage task) depends only on the one resource (a processor or memory, respectively) it is bound to. This is essential for *composi-*

*tional performance analysis*, as we discuss below. Each user has a budget on the resource it is bound to; a budget is a percentage of the resource. A budget is expressed as a minimum number of cycles of service during a given finite time duration. Equivalently it may be stated in terms of service units and scheduling intervals.

*The cyclo-static dataflow (CSDF) paradigm is used both as a programming model and as a model of computation for performance analysis.* To disambiguate, we talk about application actors and application channels versus actors and channels in the CSDF model. As shown in Figure 2(a), a CSDF application is written as a set of application actors that communicate using application channels. Each application actor is implemented by a computation task (bound to a processor), and as many DMA tasks (bound to DMAs), communication tasks (bound to the NOC) and storage tasks (bound to CMEMs) as it has outgoing application channels. In other words, the implementation of an application channel involves DMA, NOC, and CMEM, which must be modelled as three separate resources to comply with the mapping restrictions outlined above. (The distributed SRAM and DRAM memories are not used in the automated binding, but can be analysed when bound manually.)

A CSDF application graph is converted into its corresponding CSDF model for performance analysis in four steps. In Step 1 the application is expanded into its users, as explained above. In Step 2 users are bound to resources, i.e. receive resource budgets. At this point the application has its virtual platform, as shown in Figure 2(b). Steps 1 and 2 are performed for all applications, i.e. including those not written in the CSDF MOC. Later steps are for CSDF applications only. Step 3 adds the worst-case execution time (WCET) of each user on its non-shared resource, i.e. when it has a budget of 100%. Our flow computes the WCET of all users, except computation tasks, which it assumes as given. Step 4 translates each user with its budget on a resource to its corresponding CSDF model. Shared resources contain a scheduler, which is also part of the model. Different resources and schedulers result in different CSDF models.

For example, in Figure 2(c), the first two computation tasks are bound to the same processor 1, which uses a static-order (SO) scheduler. The rates  $R_i$  capture the effect of the budgets ( $2/5$  and  $1/5$ ) on the WCETs of the tasks. (Intuitively,  $R_1 = 150/(2/5)$  and  $R_2 = 100/(1/5)$ .) Note that the third task, which is bound to processor 2 with TDM has a different model [22]. The DMA task is bound to the DMA that is not shared. The NOC model uses a latency-rate (LR) abstraction of the communication performance [17]. CMEM is scheduled with CCSP, which has another model [38]. The self-edge on an actor ensures that the actor cannot restart before it finished.

The resulting CSDF graph with WCETs is used to compute end-to-end throughput and latency of the mapped application using standard dataflow maximum-cycle-mean analysis or similar. Note that this performance analysis is compositional in several senses. First, the WCET of a user bound to a resource is independent of other users and resources (since the resource is not yet shared, users can only use one resource, and resources cannot use each other). This is very

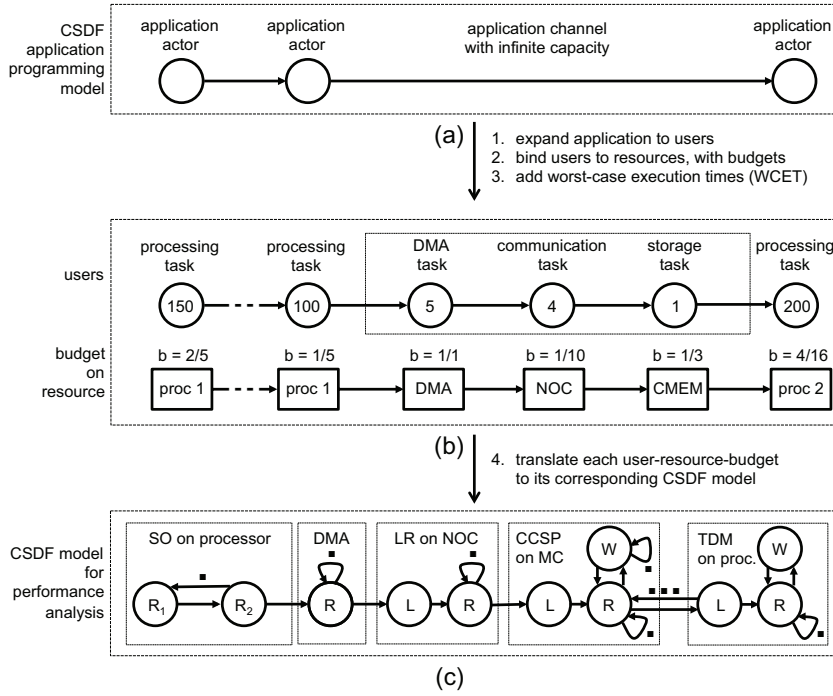


Figure 2: (a) Application written in a dataflow programming model; (b) its user-resource representation; and (c) its dataflow performance analysis model.

advantageous in computing e.g. WCET of application actors on processors: DMA, NOC, and memories do not need to be taken into account. Second, sharing is modelled with budgets that only depend on the user and the resource. An actor that receives e.g. a 3/10 budget does not care if the remaining 7/10 of the resource are used (or not) by other users belonging to the same, or indeed different applications. Finally, a virtual platform, i.e. a set of resource budgets for an application, can be defined and given to an application designer. Each resource budget can be subdivided when multiple users of the application share the resource. However, nowhere is there any dependency on other applications that may or may not co-exist on the same physical platform.

The FRT CSDF design flow is based on SDF3 [40, 41]. It analyses the trade-off between the storage space assigned to the channels (i.e. size of CMEMs) and the throughput of the graph, binds users to resources, and generate static-order schedules for processors. All application C code, and system configurations, drivers and boot code are generated.

## 7. CASE STUDIES

This section contains three case studies illustrating: 1) automatically mapping a FRT CSDF application on a platform instance, using formal performance analysis; 2) running applications with different MOCs and intra-application schedulers at the same time; 3) running multiple SRT adaptive applications at the same time, and using simulation-based performance verification. The CompSOC platform is prototyped on a Xilinx ML605 FPGA board.

In the first case study, a FRT H.263 decoder is implemented as a CSDF graph, and worst-case execution times for all application actors were obtained manually. After this, our design flow automatically finds a mapping on the given two-tile platform instance that satisfies resource (e.g. memory) and timing constraints. Non-preemptive static-order schedulers are used for intra-application task scheduling on the Microblaze processors. The CompSOC instance with the generated C code (application actors with CSDF wrappers) and C drivers and boot code were mapped on FPGA using standard synthesis tools. Experiments on the FPGA board confirmed that our platform provides a composable and predictable behaviour when running the H.263 decoder.

The second case study uses the same decoder, but with a different manual mapping, shown in Figure 3. A second synthetic application with three tasks also runs on tile 1, and uses preemptive time-triggered static-priority scheduling. Figure 4 is a trace measured on the FPGA, and illustrates CompSOC's two-level scheduling on processor tiles. Each composable RTOS slot is delineated by gray lines. In




Figure 3: Binding of a FRT H.263 decoder.






Figure 4: Schedule traces of the synthetic application (App 1) and the H.263 decoder (App 2).

particular, we observe 1) the preemptive TDM schedule between virtual platforms, and 2) the preemptive fixed-priority schedule within the synthetic application’s virtual platform (App 1), and 3) the non-preemptive static-order schedule within the H.263 application’s virtual platform (App 2). Although the exact clock cycles are not shown for legibility, each RTOS slot is exactly the same duration.

The final case study uses an adaptive SRT CSDF dataflow implementation of H.263. At run time it trades the image quality (peak signal to noise ratio) for lower energy usage, depending on the available energy in the application’s *virtual battery*, i.e. energy budget on each processor. In fact, *two H263 decoders run in different virtual platforms*. Both decoders use the same manual mapping as before. The four curves of Figure 5 illustrate that four runs of one the decoders are identical. Situation 1, run 1 & 2 are two successive runs of one decoder, when it runs alone on the platform. In situation 2, run 1 & 2 are two successive runs, when the other decoder runs too. All runs are identical in terms of energy (as shown) and cycles (not shown). Figure 6 shows a screen shot of the system. Each decoder sends its output to shared memory, where it is sampled (composably) by a picture-in-picture application running on a third processor tile. The energy in the virtual battery of each application, and the energy per frame on each processor (C1/C2) per application are also shown at run time. The “%CPU” graph displays the frequency at which each frame is decoded. Decoder 1 decodes and upsamples to the large window, while decoder 2 decodes and downsamples to the smaller picture-in-picture (PIP) window, which results in a lower energy usage on tile 2. This case study shows that multiple adaptive SRT applications can run compositably in the same CompSOC platform, and can hence be verified independently, using simulation or other means.

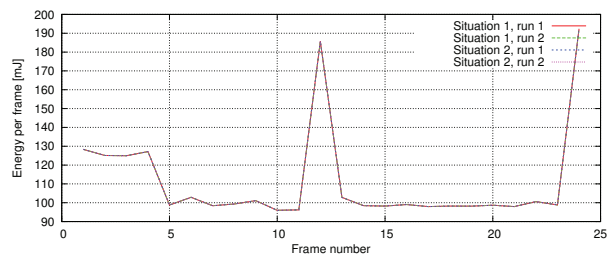


Figure 5: Comparison of energy usage of a H263 decoder in its virtual platform

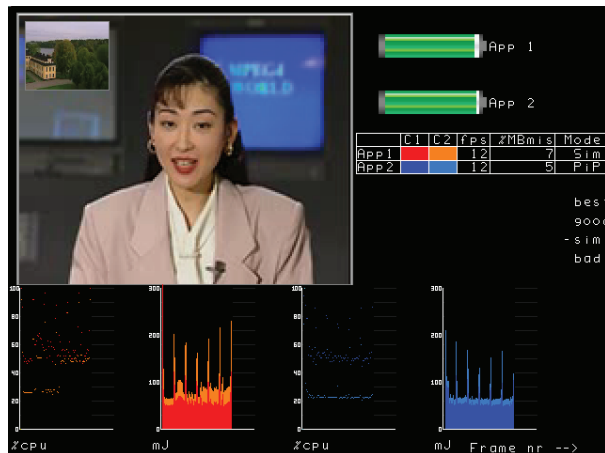


Figure 6: Screenshot of PIP with two adaptive H263 decoders.

In addition, the CompSOC FPGA prototype has been used for several years in to Master lab courses at the TU Eindhoven [15] and TU Delft [31], where students parallelise and map a (motion) JPEG decoder, fractal, and instagram applications.

## 8. RELATED WORK

Specific related work was cited when we discussed particular techniques. Therefore, here we focus on general related work. Wilhelm et al. [43] give an overview on methods to predict the temporal behaviour and the implications of various hardware architecture details on predictability. Moreover, several approaches [23, 44, 45] propose to discard the modern architectural features that aim to improve average performance, but are, however, unpredictable. From this body of work we learn how to achieve predictable hardware.

Conventional real-time systems rely on priority-based scheduling, on top of predictable hardware, and are temporally verified using formal analysis [36]. RTOSs that follow this approach include.  $\mu\text{C}/\text{OS-II}$  [20], RTEMS ([www.rtems.com](http://www.rtems.com)), eCOS [24], QNX Neutrino ([www.qnx.com](http://www.qnx.com)), ERIKA ([www.erika.tuxfamily.org](http://www.erika.tuxfamily.org)), and FreeRTOS ([www.freertos.org](http://www.freertos.org)). Since priority-based scheduling is used across applications, these approaches are not composable, and often not suitable for a mix of F/S/NRT applications.

In the RTOS domain, the techniques that aim at temporal

isolation, but not yet composability, are resource reservation, partitioning, and two-level scheduling. Resource reservation simplifies temporal verification by isolating applications, and ensuring independent temporal bounds. Few examples of such approaches are timing isolation [12, 13, 27, 32, 37], and real-time virtual resources [28]. Other approaches focus only on the interference in the most congested points of a system, such as a central bus [7], or combine reservations and priority-based techniques [1, 2] to reduce, instead of eliminate, the inter-application interference.

[11, 29, 33] are noteworthy because their approaches consider the entire system, with different resources, like our work. In particular, [29] uses the same CSDF approach for real-time performance analysis as advocated here, but for FRT applications only. All these techniques provide the isolation (compositionality) of temporal bounds, but not of actual-case temporal behaviour, which is required for our strict notion of composability.

A resource is partitioned when each user has allocated a (fixed) fraction of the resource (in time and/or space). In two-level scheduling the processor time is first split between applications, and second within an application between its tasks. The RTOSs that implement partitioning are TTA [19] and VxWorks (www.windriver.com). Two-level scheduling is implemented in OKL4 [18] and hypervisors, in general. Only PikeOS (www.sysgo.com), INTEGRITY (www.ghs.com), and LynxOS-178 (www.linuxworks.com) implement both. PikeOs and INTEGRITY have multi-core versions. PikeOs and LynxOS-178 do not provide any power management mechanisms, while INTEGRITY has a user-mode power management support. All the three RTOSs conform to the ARINC standard [8], which is a specification for time and space partitioning in safety-critical avionics real-time operating systems. According to ARINC, processing time is divided into number of time slices known as the partitions. Every partition is assigned an application, however none of the three provide cycle-accurate isolation between partitions, and subsequently, between the applications. Furthermore, none of them implement virtual interrupts for preemptive two-level scheduling.

To the best of our knowledge, CompSOC is the only platform that implements application independence down to the cycle level, power management per application, mixed-criticality and multi-MOC applications, with an automatic design flow for FRT CSDF applications.

## 9. CONCLUSIONS

This overview paper presented the CompSOC platform architecture and its accompanying real-time design flow. By offering a virtual execution platform to each application, it is possible to simultaneously execute applications that have different requirements in terms of timing (firm/soft/non real-time), different models of computation (dataflow, Kahn process networks, time triggered, and C), and different implementations (e.g. task scheduler, power management). Since each application experiences no interference from other applications in its virtual platform, it can be designed, verified, and executed in isolation. This *composability* reduces the cost and complexity of integrating multiple applications in a single platform. Additionally, since each virtual platform is

*predictable*, any mix of firm/soft/non-real-time applications can be supported. The CompSOC design flow automatically maps and verifies cyclo-static dataflow applications on the platform.

## Acknowledgements

We thank Benny Akesson for his contribution to CompSOC. This work was partially funded by projects EU FP7 288008 T-CREST and 288248 Flexiles, Catrene CA104 Cobra, and NL STW 10346 NEST.

## 10. REFERENCES

- [1] L. Abeni and G. Buttazzo. Integrating multimedia applications in hard real-time systems. In *Proc. of the IEEE Real-Time Systems Symposium*, 1998.
- [2] L. Abeni, L. Palopoli, C. Scordino, and G. Lipari. Resource reservations for general purpose applications. *IEEE Trans. Industrial Informatics*, 5(1):12–21, 2009.
- [3] B. Akesson, K. Goossens, and M. Ringhofer. Predator: A predictable SDRAM memory controller. In *Int'l. Conf. on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, pages 251–256, Oct. 2007.
- [4] B. Akesson, A. Hansson, and K. Goossens. Composable resource sharing based on latency-rate servers. In *Proc. Euromicro Symposium on Digital System Design (DSD)*, pages 547–555, Aug. 2009.
- [5] B. Akesson, W. Hayes Jr., and K. Goossens. Automatic generation of efficient predictable memory patterns. In *Proc. Int'l. Conf. on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, volume 1, pages 177–184, Aug. 2011.
- [6] B. Akesson, A. Molnos, A. Hansson, J. Ambrose Angelo, and K. Goossens. Composability and predictability for independent application development, verification, and execution. In M. Hübner and J. Becker, editors, *Multiprocessor System-on-Chip — Hardware Design and Tool Integration*, Circuits and Systems, chapter 2, pages 25–56. Springer, Nov. 2010.
- [7] A. Andrei, P. Eles, Z. Peng, and J. Rosen. Predictable implementation of real-time applications on multiprocessor systems-on-chip. In *Proc. Int'l. Conf. on VLSI Design*, pages 103–110. 2008.
- [8] Avionics Application Software Standard Interface. *ARINC Specification 653*, Jan. 1997.
- [9] A. Beyranvand Nejad, A. Molnos, and K. Goossens. A unified execution model for data-driven applications on a composable MPSoC. In *Proc. Euromicro Symposium on Digital System Design (DSD)*, pages 818–822, Aug. 2011.
- [10] G. Bilsen, M. Engels, R. Lauwereins, and J. PeperStaeete. Cyclo-static data flow. In *Int'l. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 5, pages 3255–3258, 1995.
- [11] E. Bini, G. Buttazzo, J. Eker, S. Schorr, R. Guerra, G. Fohler, K.-E. Arzen, V. Romero Segovia, and C. Scordino. Resource management on multicore systems: The ACTORS approach. *Proc. Microarchitecture (MICRO)*, 31(1):72–81, May-June 2011.
- [12] R. Bril. Towards pragmatic solutions for two-level hierarchical scheduling: A basic approach for

- independent applications. CS-report 07/19, Technische Universiteit Eindhoven, The Netherlands, 2007.
- [13] G. C. Buttazzo, E. Bini, and Y. Wu. Partitioning real-time applications over multicore reservations. *IEEE Trans. Industrial Informatics*, 7(2):302–315, 2011.
- [14] K. Goossens and A. Hansson. The Aethereal network on chip after ten years: Goals, evolution, lessons, and future. In *Proc. Design Automation Conf. (DAC)*, pages 306–311, June 2010.
- [15] A. Hansson, B. Akesson, and J. van Meerbergen. Multi-Processor Programming in the Embedded System Curriculum. In *Workshop on Embedded Systems Education (WESE)*, Oct. 2008.
- [16] A. Hansson, M. Ekerhult, A. Molnos, A. Milutinovic, A. Nelson, J. Ambrose, and K. Goossens. Design and implementation of an operating system for composable processor sharing. *Microprocessors and Microsystems (MICPRO)*, 35(2):246–260, Mar. 2011. Special issue on Network-on-Chip Architectures and Design Methodologies.
- [17] A. Hansson, K. Goossens, M. Bekooij, and J. Huiskens. CoMPSoC: A template for composable and predictable multi-processor system on chips. *ACM Trans. on Design Automation of Electronic Systems (TODAES)*, 14(1):1–24, 2009.
- [18] G. Heiser and B. Leslie. The OKL4 Microvisor: Convergence point of microkernels and hypervisors. In *Proc. Asia-Pacific Workshop on Systems*, pages 19–24, Aug 2010.
- [19] H. Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications*. Springer, 2011.
- [20] J. J. Labrosse. *Microc/OS-II*. R & D Books, 2nd edition, 1998.
- [21] J. W. Lee and K. Asanovic. Meterg: Measurement-based end-to-end performance estimation technique in qos-capable multiprocessors. *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 135–147, 2006.
- [22] A. Lele, O. Moreira, and P. Cuijpers. Dataflow based temporal analysis for TDM. In *Proc. ACM Int'l. Conf. on Embedded software (EMSOFT)*, 2012.
- [23] B. Lickly, I. Liu, S. Kim, H. D. Patel, S. A. Edwards, and E. A. Lee. Predictable programming on a precision timed architecture. In *Proc. Int'l. Conf. on Compilers, architectures and synthesis for embedded systems (CASES)*, pages 137–146, 2008.
- [24] A. Massa. *Embedded Software Development with eCos*. Prentice Hall Professional Technical Reference, 2002.
- [25] S. Matic and T. A. Henzinger. Trading end-to-end latency for composability. *Proc. of the IEEE Int'l. Real-Time Systems Symposium (RTSS)*, 0:99–110, 2005.
- [26] J. Maurice Meijer, de Gyvez and R. Otten. On-chip digital power supply control for system-on-chip applications. In *Proc. Int'l. Symposium on Low Power Electronics and Design (ISLPED)*, Nov. 2005.
- [27] C. W. Mercer, S. Savage, and H. Tokuda. Processor capacity reserves for multimedia operating systems. In *Proc. IEEE Int'l. Conf. on multimedia computing and systems*, May 1994.
- [28] A. K. Mok and A. X. Feng. Real-time virtual resource: A timely abstraction for embedded systems. In *Proc. Int'l. Conf. on Embedded Software (EMSOFT)*, pages 182–196, 2002.
- [29] O. Moreira. *Temporal Analysis of Hard Real-Time Radios on a Multi-Processor*. PhD thesis, Eindhoven university of technology, Jan. 2012.
- [30] A. Nelson, A. Molnos, and K. Goossens. Composable power management with energy and power budgets per application. In *Proc. Int'l. Conf. on Embedded Computer Systems: Architectures, MOdeling and Simulation (SAMOS)*, pages 396–403, July 2011.
- [31] A. Nelson, A. Molnos, A. B. Nejad, D. Mirzoyan, S. Cotofana, and K. Goossens. Embedded computer architecture laboratory: A hands-on experience programming embedded systems with resource and energy constraints. In *Workshop on Embedded Systems Education (WESE)*, Oct. 2012.
- [32] F. Nemati, M. Behnam, and T. Nolte. Independently-developed real-time systems on multi-cores with shared resources. In *Proc. Euromicro Conf. on Real-Time Systems (ECRTS)*, 2011.
- [33] K. J. Nesbit, J. E. Smith, M. Moreto, F. J. Cazorla, A. Ramirez, and M. Valero. Multicore resource management. *Proc. Microarchitecture (MICRO)*, 28(3):6–16, May-June 2008.
- [34] A. Nieuwland, J. Kang, O. P. Gangwal, R. Sethuraman, N. Busá, K. Goossens, R. Peset Llopis, and P. Lippens. C-HEAP: A heterogeneous multi-processor architecture template and scalable and flexible protocol for the design of embedded signal processing systems. *ACM Transactions on Design Automation for Embedded Systems*, 7(3):233–270, 2002.
- [35] A. Schranzhofer, R. Pellizzoni, J.-J. Chen, L. Thiele, and M. Caccamo. Worst-case response time analysis of resource access models in multi-core systems. In *Proc. Design Automation Conf. (DAC)*, pages 332–337, 2010.
- [36] L. Sha, T. Abdelzaher, K.-E. Arzén, A. Cervin, T. Baker, A. Burns, G. Buttazzo, M. Caccamo, J. Lehoczky, and A. K. Mok. Real time scheduling theory: A historical perspective. *Real-Time Systems*, 28:101–155, November 2004.
- [37] I. Shin *et al.* Periodic resource model for compositional real-time guarantees. In *Proc. of the IEEE Int'l. Real-Time Systems Symposium (RTSS)*, 2003.
- [38] F. Siyoum, B. Akesson, S. Stuijk, K. Goossens, and H. Corporaal. Resource-efficient real-time scheduling using credit-controlled static-priority arbitration. In *Proc. Int'l. Conf. on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, volume 1, pages 309–318. Aug. 2011.
- [39] S. Stuijk, T. Basten, M. Geilen, and H. Corporaal. Multiprocessor resource allocation for throughput-constrained synchronous dataflow graphs. In *Proc. Design Automation Conf. (DAC)*, pages 777–782, 2007.
- [40] S. Stuijk, M. Geilen, and T. Basten. SDF<sup>3</sup>: SDF For Free. In *Proc. Int'l. Conf. on Application of Concurrency to System Design (ACSD)*, pages 276–278, 2006.

- [41] S. Stuijk, M. Geilen, and T. Basten. A predictable multiprocessor design flow for streaming applications with dynamic behaviour. In *Proc. Euromicro Symposium on Digital System Design (DSD)*, 2010.
- [42] P. Vivet, E. Beigne, H. Lebreton, and N.-E. Zergainoh. On line power optimization of data flow multi-core architecture based on vdd-hopping for local DVFS. In *Integrated Circuit and System Design. Power and Timing Modeling, Optimization, and Simulation*. Springer, 2011.
- [43] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. Puschner, J. Staschulat, and P. Stenström. The worst-case execution-time problem - Overview of methods and survey of tools. *ACM Trans. Embed. Comput. Sys.*, 7:36:1–36:53, May 2008.
- [44] R. Wilhelm, D. Grund, J. Reineke, M. Schlickling, M. Pister, and C. Ferdinand. Memory hierarchies, pipelines, and buses for future architectures in time-critical embedded systems. *IEEE Journal of Computer Aided Design*, 28(7):966–978, 2009.
- [45] J. Wolf, M. Gerdes, F. Kluge, S. Uhrig, J. Mische, S. Metzloff, C. Rochange, H. Cassé, P. Sainrat, and T. Ungerer. Rtos support for parallel execution of hard real-time applications on the merasa multi-core processor. In *Proc. IEEE Int'l. Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing (ISORC)*, pages 193–201, 2010.