

Manycore simulation for peta-scale system design: Motivation, tools, challenges and prospects



Javad Zarrin^{a,*}, Rui L. Aguiar^{a,b}, João Paulo Barraca^{a,b}

^aATNoG Research Group, Instituto de Telecomunicações (IT), Aveiro, P-3810-193 Aveiro, Portugal

^bDepartamento de Electrónica, Telecomunicações e Informática Universidade de Aveiro, Campus Universitário de Santiago, 3810-193 Aveiro, Portugal

ARTICLE INFO

Article history:

Received 14 July 2016

Revised 18 November 2016

Accepted 14 December 2016

Available online 17 January 2017

Keywords:

Architectural simulations

Manycore simulators

Full-system simulation

Manycore systems

Computer architecture

ABSTRACT

The architecture design of peta-scale computing systems is complex and presents lots of difficulties to designs, as current tools lack support for relevant features of future scenarios. Novel systems must be designed with great care and tools, such as manycore architecture simulators, must be adapted accordingly. However, current simulation tools are very slow, often specific-purpose-oriented, suffer from various issues and are rarely able to simulate thousands of cores. The emergence of peta-scale systems and the upcoming manycore era brings nevertheless new challenges to computing systems and architectures, adding further difficulties and requirements on the development of the corresponding simulators. Furthermore, the design of architecture simulators for manycore systems involve methods and techniques from various interdisciplinary research areas, which in turn brings more challenges in different aspects. As system complexity grows, the growth of the simulation capacity is being outpaced (reaching the so called simulation wall). In this paper, we present the challenges for simulating future large scale manycore environments, and we investigate the adequacy of current modeling and simulation tools, methodologies and techniques. The aim of this work is to highlight how current approaches can best deal with the identified problems, smoothing the challenges of research in future peta-scale systems.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Current trends in computation technology have focused in improving performance by increasing the number of cores per die (parallelism) rather than by increasing the rate of clock frequency of each core, due to the exhaustion of the Moore's law. Many companies and academic communities pushed this trend, designing multicore and manycore systems with capacity of tens to hundreds of cores per single die. These manycore processors are more like data-centers-on-a-chip than previous single processors, as a complex communication network connects the different cores. It is predictable in a near future to consider systems with a very large interconnect network in manycore machines with dimensions from thousands to millions of cores.

A consequence is that both system design and programming concepts must increasingly focus in heterogeneous parallelism. Future parallel and distributed applications, compilers, operating systems and tools must be able to scale well with the hardware nature of manycore and distributed execution. However, increasing the number of cores on a die increases the

* Corresponding author.

E-mail addresses: javad@av.it.pt (J. Zarrin), ruilaa@ua.pt (R.L. Aguiar), jbarraca@ua.pt (J.P. Barraca).

complexity of hardware designs, and has considerable impacts which result in enlargement of the potential design-space. Moreover, it brings serious challenges particularly for memory hierarchies and on-die interconnect bandwidth, both within the die and off die.

Systems must then be designed with great care and tools, such as manycore architecture simulators, must be adapted to address these disruptive challenges. Manycore simulators (i.e. manycore architecture simulators) can assist researchers from different areas. These areas include software (such as programming models, operating systems, compilers, etc.), hardware and computing architectures to model and assess future systems. The upcoming manycore era brings new challenges to computer architects that must be paralleled by the development of adequate architecture simulators. Ideally, to simulate a fully parallel system we can expect that an efficient architecture simulator should (at-least) be able to be parallelizable, and use the benefit of concurrency, enabling faster evaluation of future systems. Moreover, architecture simulators should provide a highly scalable, fast and accurate model to describe, emulate/simulate and measure the hardware details, memory hierarchy and interconnection networks. Furthermore, they must meet the stringent requirements along the lines of productivity, multi-modeling, synchronization, modularity, and event sampling capabilities. However, these properties are not true for the majority of the currently available architecture simulators. As complexity grows, the performance of a single simulated CPU core slows down, and the usage of these sequential simulators (i.e. architecture simulators) will be mainly limited by the performance of simulating a single CPU.

Note that, in this paper, we use the term “simulator” to refer to “architecture simulator” or “manycore architecture simulator” which differs from its more general meaning. We use the term “manycore simulation tools” in a slightly more general way to cover a wider range of contexts (methods and techniques). In other words, a simulator (architecture simulator) is a type of manycore simulation tool. Furthermore, the term “emulator” (or architecture emulator), as used in this paper, also differs from its general meaning in other research fields. In this paper, an emulator means a simulator which lacks support for performance measurements. In fact, we associate an emulator with functional correctness only. This means that the notion of time for an emulator is imprecise and often just a representation of the wall-clock time of the host. We use the term “emulation” to refer both to “functional simulation” and to describe the act of an emulator.

The rest of this paper is organized as follows. In [Section 2](#), we discuss why manycore architectural simulation is needed, particularly for research on peta-scale systems. [Section 3](#) presents in detail our taxonomy to addresses current modeling and simulation tools, as well as the methodologies that could be exploited and enhanced in order to design next generation efficient simulators. In [Section 4](#), we provide an overview of recently proposed simulation tools for architectural analysis, which are able to simulate the entire execution cycle of application for the target systems. Other simulator types, such as those that mainly focus on physical modeling aspects (e.g. power, energy and thermal) or interconnect simulators (e.g. Network on Chip simulators), are out of the scope of paper for conciseness. In [Section 5](#), we extract and elaborate a set of major problems and challenging issues created by manycore simulation. Finally, in [Section 6](#), we present our conclusions, followed by a discussion of future directions for research. This includes possible approaches and solutions which can be used to solve the problems and deal with the challenges identified.

2. Simulation and peta-scale systems

Peta-scale systems are defined as systems which are able to provide peta-FLOPS, millions of billions of Floating Operations per Second, computational power [[1,2](#)]. They can be described as the increasingly massive and dynamic networks of interconnected diverse processors and components (i.e. elements). Such as system, as a whole, exhibits a set of properties and behaviors among the elements, which are not distinguishable from the properties of the individual processors and components. These systems are only on its infancy currently, but in future peta-scale manycore systems, we can expect to have computing nodes with more than 10,000 cores per node. We can also expect to have much more diversity (heterogeneity) of cores, interconnections and architecture designs compared to today systems.

Architectural simulation is a common method for studying and analyzing different architectures, designs and algorithms for various target systems through imitating the operation of real-world processes, processors and systems over time. Architectural simulation acts as a low cost alternative to experimentation on real systems by representation of key characteristics, behaviors and functions of the real systems. The objective of simulation is to provide capability to researchers and designers to flexibly and efficiently explore a design space. This can include analyzing the performance of current systems (e.g. architecture assessment), acquiring and predicting processor/system behaviors and evaluating novel designs. Simulation enables today’s designers to analyze and predict different aspects (such as performance, reliability and efficiency) of future’s machines. This means that simulation is a particularly useful tool when the desired target systems, such as peta-scale systems, do not currently exist in reality. However, current simulation tools are very slow, often specific-purpose-oriented, suffer from various issues and are rarely able to simulate more than 2000 cores (we discuss current simulation tools further in [Section 4](#)). Furthermore, introducing the concept of peta-scale system presents more requirements that must be fulfilled by current simulation tools. We discuss these requirements and their corresponding issues further in [Section 5](#).

[Fig. 1](#) demonstrates a generic structure for (software-based) simulation. As it is shown in the figure, a simulator is an application software which runs on single host or multiple networked host machines (distributed simulation). The target system is the system which needs to be simulated by simulator. Depending on user requirements the target system might be a partial system or a full system including target (simulated) OS. The term “target Instruction Set Architecture (or target ISA)”, as used in this paper, refers to the ISA of the processor architecture simulated. Similarly, the term “host ISA” refers to

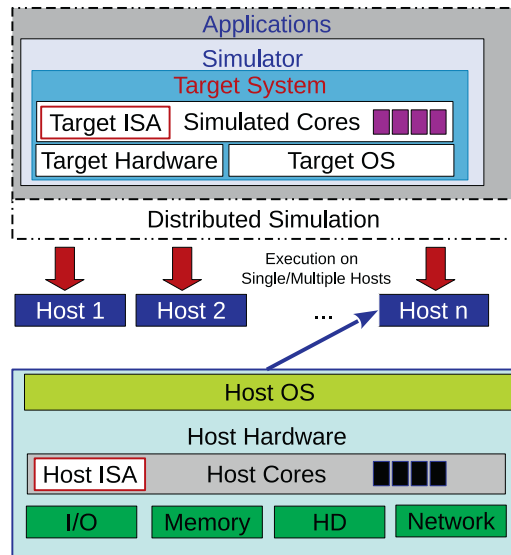


Fig. 1. Overall simulation structure.

the ISA of the host machine which runs the simulation. Note that, it might be possible to have some similarities between a target system and its host system (e.g. target ISA = host ISA or target OS = host OS). In such a case, a simulator can alternatively employ a different simulation approach which might be faster (or easier to implement) (see Section 3). More details of this structure are discussed in the next section.

3. Manycore simulation

Simulators are essential tools for the design and evaluation of computer and system architectures. Taking in consideration the objective of the evaluation, there are a large variety of different simulation tools, techniques and methods in the current literature. Fig. 2 presents a snapshot of the most important classifications for manycore simulation tools and related technologies. Fig. 2 also presents example simulator for each category (note that a simulator can be in multiple categories). We discuss these classifications and their related concepts in detail in Section 3.1.

3.1. Terminology and classifications

There are plentiful tools, methods and techniques for manycore simulation in the current literature. These can be classified along many dimensions, such as: user/application-level vs full-system; functional vs timing; trace-driven vs execution-driven; cycle-driven vs event-driven. In the following, we discuss and briefly explain the terminology and concepts associated with some of the most important classifications.

Emulators vs Simulators: Emulators are tools able to demonstrate the functional behavior of the system, focused on the exact reproduction for an external system. They repeat the function of a target platform on a host platform. Simulators differ from the former because they consider an abstract, simplified model and do not try to replicate all the aspects of a system. A simulator in addition to ensuring the functional correctness, must provide capability to study and analyze the performance of proposed system/hardware designs by using accurate timing information.

Functional vs timing/performance: According to our terminology (see Section 1), architecture emulators include only functional models while simulators contain both functional and performance models. The functional model (functional partition) is in charge for the correct execution of the target Instruction Set Architecture (ISA) (i.e. architectural modeling). In addition, functional models may provide facility to observe the interactions among processors, memory and I/O peripherals without modeling micro-architectural details. As examples of the tasks for a functional model we can mention aspects as decoding instructions, updating simulator memory and verifying the floating point operations.

The performance model (timing partition) is responsible to drive a functional model. This is done by providing accurate timing information in a way that simulating a particular micro-architecture (at least in a specific aspect) would be possible (i.e. micro-architectural modeling). In other words, a timing model is able to determine the time the target architecture takes to execute an instruction. Examples of the tasks for the timing model include making decisions to select the next instruction for execution, tracking branch mispredictions, and predicting the clock cycles to execute instructions.

Creating a new functional model might be complex in terms of implementation, optimization and verification, but after that, it can modularly be coupled and reused across various timing models. On the other hand, implementation of only a timing model with capability to reuse an existing functional model can be much easier and less time consuming than

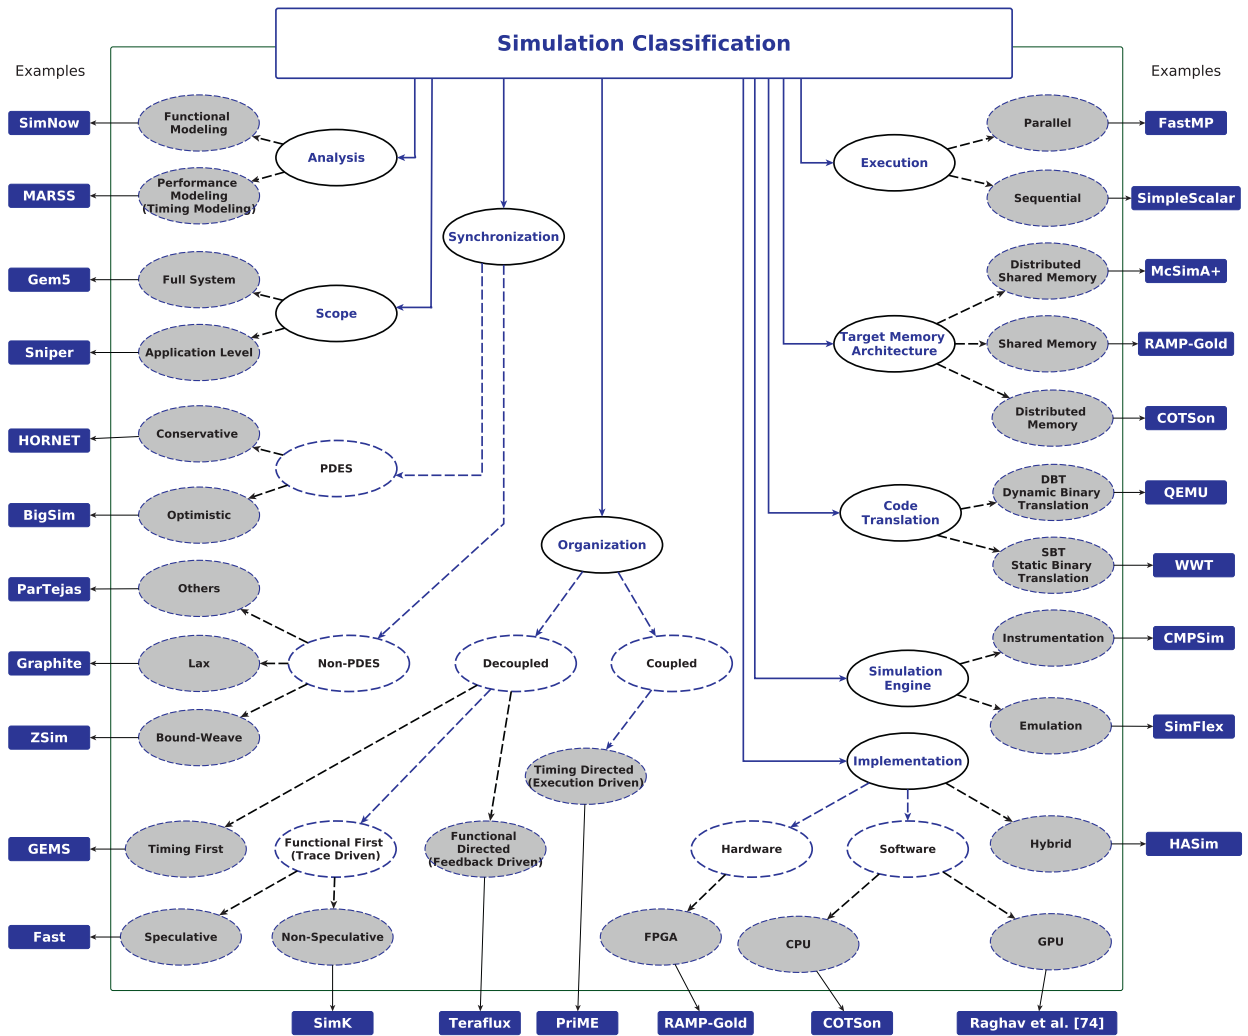


Fig. 2. Simulation classifications.

developing a complete simulator (including both functional and timing models) from scratch. Unlike functional models, the correctness of target architectures (in term of ISA functional) is not of primary concern for timing models. What matters is to track and control the accuracy of micro-architecture-specific timing details [3].

Simulation scope: An architectural simulator is defined as a piece of software which mimics the behavior of a real computer system with ability to estimate performance and outputs for a given input (application). It may model different computer devices and components (i.e. only a single target microprocessor, or an entire computer system including processors, memory system, and I/O devices) with different level of details. The simulation scope specifies the scope of target systems that an architectural simulator can model. According to this, we classify the processor simulators (i.e. architectural simulators) into two categories: full-system simulators and user/application level simulators (see Fig. 3).

A full system simulator provides capability to run a detailed, complete and real software stack on the target (simulated) system without any modification. The software stack might be OS, complex applications such as multithreaded and multi-process workloads, or applications that highly exercise system calls, I/O and networking. The full system model generally includes processor cores, memories, network and interconnection, buses, peripheral devices and privileged modes. However, the supported models and the level of details for each model are various for different simulators. A full system simulator, besides timing/performance modeling might also supports simulation of the physical models such as power, energy and thermal. Gem5 [4], Flexus (SimFlex) [5], and MARSS [6] are some examples of full system simulators.

Furthermore, full system simulators can be very important design tools, particularly for System-on-Chip (SoC) simulation where it is necessary to efficiently cope with the huge hardware/software design space. SoC architecture integrates all components of a (electronic/computing) system (including complex hardware/software) into a single integrated concept. This provides capability to scale computing performance/power efficiency through combining (massively parallel and high

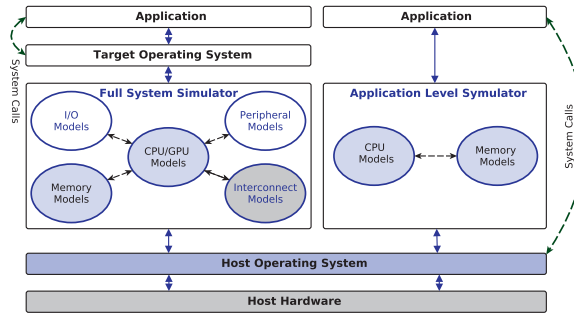


Fig. 3. Full system simulator vs application level simulator.




Fig. 4. Functional first organization – the gray elements (rollback and mismatch detection) are only supported by speculative-functional-first organization.

performance) manycore processors, Network-on-Chip (including several network/interconnect communication protocols) and software (including OS and various computing intensive user applications) [7,8].

On the other hand, user/application level simulators model only the user side of applications without OS support. They rely on the host machine (host OS) to service the system calls and execute a given user code of a benchmark on top of the simulator. These simulators are easier to develop and use, since they do not boot an OS. But they are also limited to only support specific workloads. In other words, they cannot run applications such as multithreaded and JVM workloads that frequently use applications that are sensitive to system time like and client-server workloads. As examples of user/application level simulators we can mention SimpleScalar [9], BigSim [10], Graphite [11], Sniper [12], ZSim [13] and PriME [14].

Simulator organization: Each simulator generally contains functional and timing partitions which interact in order to form a complete simulation. According to whether the functional and the timing partitions are completely separated or not, simulators can be categorized into decoupled and coupled. Coupled simulators potentially can provide flexibility for developing precise simulations (i.e. highly detailed simulation), due to speculative execution modeling (producing all the values and possible side effects) and timing-dependent outcomes. However, this flexibility is reduced when there is need for frequent modifications of functional or performance models. In other words, complexity arises due to new complex devices (functional models), modern performance/timing models and their internal interactions. For coupled simulators, it is challenging to address multiple different conflicting demands (such as simulation precision, accuracy, flexibility and performance) in a single simulator component which integrates both functionality and timing modeling. On the other hand, decoupled simulators aim to reduce this complexity by completely decoupling functional and timing models. Accordingly, they can achieve better flexibility and potentially other advantages such as accuracy (through correctness verification) and modularity (which can lead to faster development and easy modification) [15]. The drawback for this approach is to introduce redundancy which may reduce the simulation performance. Moreover, decoupled simulators typically call an external simulator/emulator to perform functional modeling. This might increase the difficulty for modeling interactions among functional and timing models.

Based on the type of interaction or the relationship between the functional model and the timing model, the simulators can use different organizations. A taxonomy of microarchitectural simulator organizations is introduced in [15] and [16]. We extend those taxonomies and classify the simulator organizations into the following five categories (Table 1 presents the advantages and disadvantages for each of the aforementioned simulator organizations):

1- Functional first (Trace driven): The functional-first is a completely decoupled approach. The functional model is separately executed first, and later, when the functional modeling is finished, the timing model is run. As demonstrated in Fig. 4, the functional model executes instructions and produces traces that are streams of information about the execution. Traces are then fed into the timing partition where the microarchitectural simulation is performed. Depending on whether speculative execution is supported or not, the functional-first organization can be categorized into speculative and non-speculative.

The speculative functional-first assumes that all parts of execution (not just for branch control) are speculative. When the timing partition detects that the execution in functional partition has differed from the timing partition, it asks the

Table 1
Advantages and disadvantages of various simulator organizations.

Simulator Organization	Pros	Cons
Functional-First (Non-Speculative)	Easy implementation and parallelization, Highly-decoupled design, On-way data flow, Highly-Modular, Support for binary translation, direct execution and compiled-code simulation	No support for speculative execution, Not able to model timing-dependent outcomes between threads, No corrections for the timing model, Mismatching functional execution and timing simulation (i.e. execution divergence)
Functional-First (Speculative)	Matching between functional execution and timing simulator, Support for speculative execution and multithreaded applications, Natural parallelism of the functional and timing simulator, On-way data flow	The functional model must support speculation, The timing simulator needs to implement some of the functionality behaviors
Functional-Directed	Efficient trade-off between high speed functional modeling and high accurate timing simulation, Highly-adapted with sampling techniques, Support speculative execution and multithreading	Periodical feedback has impact on the simulation accuracy.
Timing-First	Single timing-functional call per instruction, Fast development, Debuggability of the timing simulator, Matching between functional execution and timing simulator.	No support for accurate modeling of interactions between threads, The timing simulator needs to implement almost all functionality behaviors.
Timing-Directed	Highly-accurate, Natural support for speculative execution and multithreading in timing partition, Potential support for memory and memory consistency modeling, easy validation of the timing simulator, matching between functional execution and timing simulator	Highly-coupled design, complex functional modeling, Highly frequent inter-partition (timing-functional) communication which impacts the simulation speed, Multiple functional-timing calls per instruction.




Fig. 5. Functional directed organization.




Fig. 6. Timing first organization.

functional partition to undo or rollback the effects of wrong path for instruction execution, and to redirect fetch. Examples of this approach include UFast [17], FastSim [18], SimpleScalar [9], Zesto [19], ReSim [20], BigSim [10] and Transformer [21].

Unlike speculative approach, in non-speculative functional-first, rollback mechanism is not supported. This means that the timing partition independently provides the timing model in a highly decoupled fashion and regardless of any potential execution divergence (mismatching functional execution and timing model). Non-speculative approach might be easy for implementation and parallelism through a highly modular and decouple design. However, it potentially suffers from execution divergence issues which results in being inadequate for modeling timing-dependent execution behaviors (such as interactions among threads in a multithreaded application). Examples of this approach include single thread work load simulators (like PTCMP [22]).

2- Functional directed (Feedback driven or event driven): The functional-directed is a completely decoupled approach where the functional model drives the timing model. The functional model is assumed to be accurate, but timing feedback from the timing model is needed in order to correct the timing behavior. As we can see in Fig. 5, using a timing feedback mechanism, the timing model periodically adjusts the speed of functional execution to reflect the timing estimates. This gives the running application a more-precise timing correction. COTSon [23] and Teraflux [24] use this approach.

3- Timing first (instruction driven): The timing-first is a decoupled approach which uses the timing model to drive the functional model. Fig. 6 demonstrates the timing-functional interactions for this organization. The timing model runs ahead of the functional model and simulates the functional behavior (i.e. instruction execution) in addition to the timing

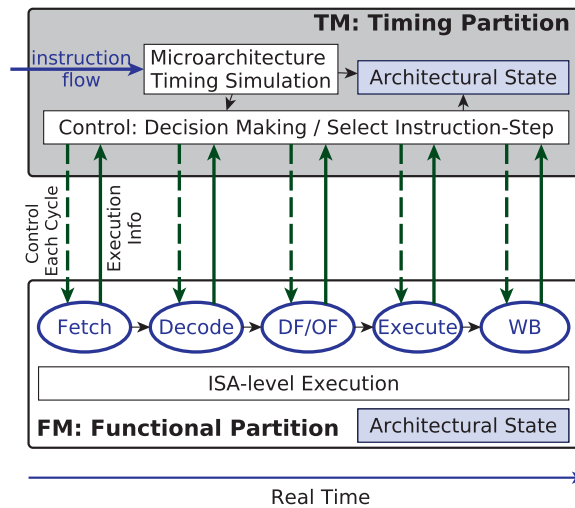


Fig. 7. Timing directed organization.

simulation. But the functional execution (for each instruction) must be checked and verified later by the functional model. This means that the functional model is only used for checking; when a mismatch (execution divergence) is detected, the timing model is notified by the functional model for the correction. In such a case, the architectural state of the timing model is reloaded from the functional partition. Examples of timing-first organization include TFSim [15], FeS2 [25] and GEMS [26].

4- Timing directed (Execution driven or cycled riven): In timing-directed simulators, functional and timing models are tightly coupled. As we can see in Fig. 7, multiple interactions (per instruction) between timing and functional model raise system complexity. The timing model processes the flow of instructions through micro-architectural timing simulation. In each cycle, it directs the functional model with which step (e.g. fetch, decode, operand fetch, memory access, execute, write-back) of which instruction should be executed. Accordingly, the functional model returns the execution information to the timing model for each step. In timing-directed organization, the architectural state in both (timing and functional) partitions are naturally matched with each other, thereby preventing execution divergence. Talisman [27], Graphite [11], ZSim [13] and PriME [14] use this organization.

Simulator execution: The execution of simulations on the host machines can be either sequential or parallel. Sequential simulators are highly accurate but as the complexity of the target architecture increases the simulation speed significantly decreases [28,29]. There are numerous examples of sequential emulators and simulators in the literature including QEMU [30], Embra [31], Mambo [32], SimNow [33], SimpleScalar [9], Simics [34], SimOS [35], Rsim [36], Gem5 [4] and PROTEUS [37]. While most of these tools are able to model multicore architectures and multi-processors (in parallel chips), they are all limited to only utilizing a single processor on a single host machine which leads to significant slowdown in simulation speed. Parallel emulators (e.g. PQEMU [38], COREMU [39], Parallel Mambo [40], and Parallel Embra [41]) and simulators (e.g. FastMP [42], SlackSim [43], Wisconsin Wind Tunnel Simulators (WWT) [44,45], CMPSim [46], ZSim [13], MARSS [6], SimFlex [47], GEMS [26], COTSon [23], Graphite [11], Sniper [12,48] and BigSim [10]) enhance the performance of simulation (in terms of simulation speed) by dividing and distributing the simulation workload across multiple cores on multiple host machines. However, parallelizing a simulator without sacrificing simulation accuracy is challenging. Furthermore, the execution of functional and timing models might be different. For example, COTSon [23] uses SimNow [33], a sequential emulator, for functional modeling, while it benefits of a parallel timing model.

Parallelization and synchronization strategies: Synchronization overhead is one of the major challenges for parallel simulation [49]. Different parallelization strategies and synchronization algorithms such as Parallel Discrete Event Simulation (PDES) [50], Lax synchronization [11] and Bound-Weave parallelization [13] have been proposed for parallel architectural simulators. Along this line we can classify parallel simulators based on their parallelization techniques into PDES-based and non-PDES-based simulators.

1. *PDES-based synchronization:* PDES aims to facilitate fast execution of large simulation programs. It refers to the parallelization strategies for distributed simulation through execution of a single discrete event simulation program on a parallel computer [50]. Many of the conventional parallel micro-architectural simulators in the literature such as HORNET (conservative) [51], BigSim (Optimistic) [10], SimK (conservative) [52], SlackSim (hybrid) [43] and COTSon [23] are based on PDES. Using PDES, the components of the simulator are divided across host threads. The time-stamped simulation events from each component are distributed among multicores in parallel chips of multiple host machines. These events are executed concurrently while maintaining the causality relationship (i.e. cause and effect relationship) between them. This means that

some sequencing order between events executing in two separate processes must be maintained, although maintaining this causality relationship is challenging while exploiting inherent parallelism for faster job scheduling [53].

PDES-based synchronizations can be broadly classified as conservative (pessimistic), optimistic and hybrid. Conservative approaches, using pessimistic estimates, only process events when it is safe to do so. The events are scheduled in time stamp order while avoiding deadlocks efficiently. The synchronization is performed every time an ordering violation (causality error) may happen. On the other hand, in optimistic approaches, all events are executed speculatively and once an ordering violation is detected, the recovery can be performed by invoking a rollback mechanism.

There are several classic algorithms for conservative and optimistic synchronizations including asynchronous-conservative [54], Lower Bound Time Stamp (LBTS) [55] and optimistic time warp [56]. In asynchronous-conservative, there is no need to perform any global synchronization. Instead, various deadlock avoidance algorithms can be applied to ensure that only safe events in the future queue are processed. Using LBTS, in each cycle, the events in the future queue which have the lower bounded time stamp are executed while time progressing is managed by a global reduction and synchronization mechanism. In optimistic time warp, a causality error is detected whenever the time stamp of the received event message by a process is smaller than the process's clock. Many other PDES based synchronization mechanisms have been proposed in the last decade such as conservative null message (or CMB) [57], conservative forecast null message (FNM) [58].

PDES-based parallel simulators might support single or multiple synchronization algorithms. For example, Manifold [59] supports multiple standard PDES algorithms including LBTS [55], CMB [57], FNM [58] and time quantum synchronization [59].

Parallel simulators based on Optimistic-PDES generally suffer from poor scalability due to frequent roll-back and synchronization. On the other hand, simulators based on conservative-PDES provide better scalability due to their simplicity and less synchronization overhead in comparison with optimistic based simulators [51,60]. Nevertheless, both approaches provide reasonable accuracy [13]. The scalability for parallel simulators means that the simulation speedup constantly rises when the number of allocated host processors for simulation is increased [61]. Over the decades of research in the field, many optimization methods such as Lazy Cancellation [62], Lazy Reevaluation [63], Direct Cancellation [50], Early Cancellation [64], Space-time Simulation, Optimistic Time Windows [65], Wolf Calls [66] and Time Warp Straggler Message Identification [67] have been proposed to improve the overall performance of PDES based Simulators. However, PDES based parallel simulators are not yet truly competitive with performant sequential simulators with respect to simulation accuracy and scalability.

II. Non-PDES-based synchronization: Non-PDES based parallel simulators such as Sniper [12,48] and Graphite [11], relax synchronization requirements to obtain scalability by permitting micro-architectural events to occur out of order. However, these approaches provide complexity to model the actual behavior of the target system components such as memory controllers and shared caches while they sacrifice the simulation accuracy. In fact, Graphite leverages lax synchronization models to enable trade-offs between simulation speed and simulation accuracy while Sniper (which is built on Graphite) uses higher levels models by reducing accuracy compromise. ParTejas [68], a more recent work, is a shared memory based parallel simulator written in Java. Unlike Sniper and Graphite, ParTejas doesn't rely on highly relaxed synchronization, but rather it primarily relies on novel concurrent data structures. In fact, it uses a lock free parallel slot scheduler for synchronizing the accesses of multiple threads at a shared resource while uses flexible barriers known as phasers to relax synchronization within bounds. ZSim [13] is one of the latest parallel simulators that provides an alternative approach for synchronization which is called Bound-Weave. ZSim simulation runs in time quanta where each time quanta is defined as a small interval (e.g. 1000 cycles). ZSim divides each time quanta (interval) into two parallel phases: bound phase and weave phase. In the bound phase, similar to the lax synchronization, the cores are simulated without simulation of the interactions among the cores (i.e. unordered simulation), but the core-memory access traces for all the cores are recorded. In the weave phase, parallel event-driven simulation is performed by using the traces to simulate the memory accesses in order. Bound-weave is proposed based on the assumption that path-altering interference is extremely rare. This is a right assumption but only for cores that implicitly communicate through the cache hierarchy. Thus, bound-weave parallelization methods are not applicable to other communication styles (e.g. extremely fine-grained message-passing across whole chip). Moreover, simulating speculation (e.g. transactional memory) and complex workloads (e.g. kernel-intensive applications) would be difficult since ZSim is a user-level simulator. ZSim also provides limitations to model multi-threaded cores, detailed NoC models and virtual memory (TLBs).

In summary, we can claim that non-PDES-based simulators provide better scalability than PDES-based simulators, but they suffer from inaccuracy of simulation.

Simulation engine: The simulation engine specifies the underlying strategy which each simulator uses to perform functional modeling. As we already discussed at the beginning of this section, the main responsibility of the functional model is the correct execution of the simulated ISA (i.e. emulation). But the ISA emulation might not be necessary when simulators directly use the host's ISA, instead, simulators can use instrumentation. This indeed eliminates the need for functional model for such simulators and increases the simulation speed with the potential cost of limiting the ISA of the target machines to only the host's ISA.

Depending on which strategy is used for the simulation engine, we classify the simulators into emulation-based (or interpretation-based) and instrumentation-based. An emulation-based simulator either uses its internal emulator or leverages an external simulator/emulator to model functional behavior. The simulator interprets the instruction and according to the simulator's organization, invokes both functional and timing model to execute the instruction for the simulated ISA.

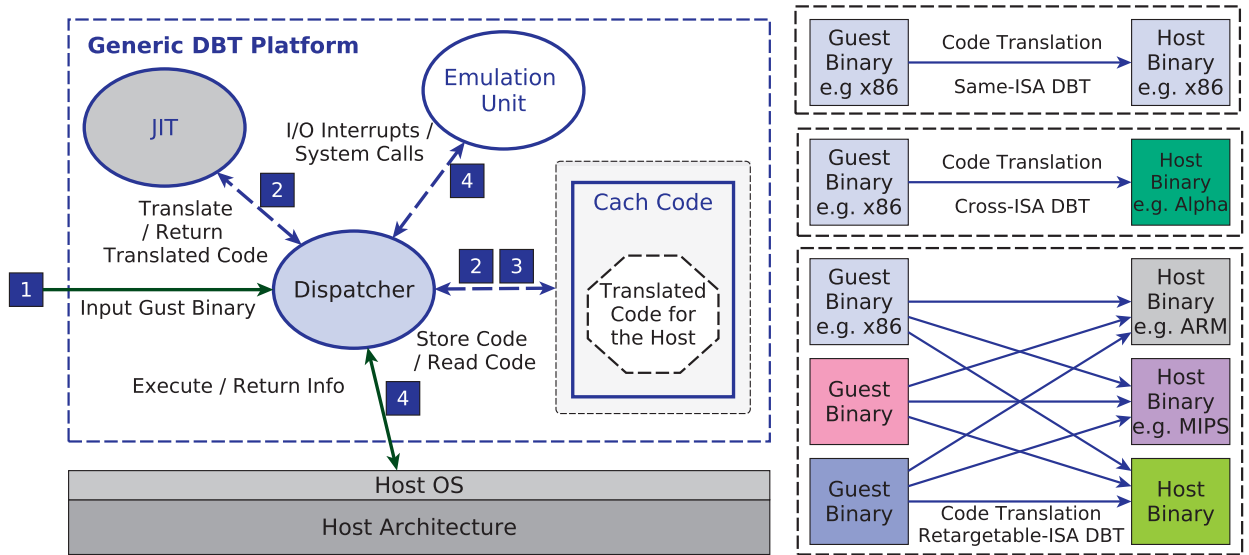


Fig. 8. General DBT framework and DBT types.

Emulation becomes the primary and even the optimal option, particularly when the simulated ISAs are supposed to be different from the host's ISA, or when the simulators are expected to be portable. On the other hand, instrumentation, is faster than emulation, since the instructions are directly executed on the host machine. The simulator adds instrumentation calls to the simulated binary in order to interact with the timing model by calling the timing model before each basic block or memory operation. Instrumentation provides facilities to understand the execution behavior of each instruction and measuring the execution performance on the host machine. It returns a set of useful information about the execution which can be fed the timing model. This can be done by enabling transparent access to the state of host's processor and memory after each instruction execution. Furthermore, instrumentation can be very efficient, particularly when emulation and functional modeling become difficult for complex ISAs. Most of the current instrumentation-based simulators run on x86 hosts, since x86 is common in both desktop and server segments, and the ISA's impact is less relevant. Examples using emulation strategies include MARSS [6], Gem5 [4], SimFlex, COTsOn and BigSim while simulators such as CMPsim [46], Graphite [11], PriME [14], ZSim [13] and McSimA+ [69] benefit from instrumentation strategies.

Binary (Code) translation: Binary translation is the core technology for both emulation and instrumentation strategies through enabling translation techniques to translate binary codes from a simulated architecture to the host architecture. The main difference is that, in emulation, all the functional and execution behavior of the applications for the simulated architecture are modeled through a functional model. In instrumentation, they use instrumentation calls added to the translated binary. Thus, arbitrary statistics about the run-time actions of the executing application can be gathered from the host which can be used to specify the functional behavior of the application for the target architecture without needing a complete functional model.

Two main types of binary translation techniques are static (ahead-of-time translation) and dynamic (translation at run time).

In Static Binary Translation (STB), all the binary code of an executable file is converted into code that can be executed on the host architecture, and after that the translated code is run on the host. This might not be efficient, since discovering some part of the code may depend on the run-time values (e.g. indirect branches, dynamically loaded libraries and self-modifying code). Wisconsin Wind Tunnel (WWT) [70] is a simulator which uses static binary translation.

On the other hand, Dynamic Binary Translation (DBT), relies on on-line code translation which means that that the portions of the binary code (each short sequence of code or single basic block), are translated and executed one after another in the order; the code is only translated as it is discovered. DBT suffers from large amount of overhead during translation which leads to increased execution times. Code cache is a technique which reduces the translation overhead by caching the translated code sequences for later usage when subsequent executions of the same code region can use the already translated code.

Fig. 8 presents the general framework of the DBT systems which includes four main components: dispatcher, just-in-time compiler (JIT), emulation unit and software-based code cache. The dispatcher coordinates translation and execution of the code through directing other components. It gets the address of the next program which is a segment of guest binary code and determines whether a translated copy of that code is available in the code cache. If so, the execution is resumed in the code cache, otherwise, the dispatcher kick-starts the JIT compiler to translate the untranslated guest code segment. JIT fetches the code segment, and then optimizes and translates the code to the host binary in the software code cache.

Optimization can be performed by adding, removing, inserting or replacing instructions to the code before translation. Using this capability, JIT allows to inject various instrumentation instructions into code. Furthermore, the JIT's granularity to fetch and translate (i.e. amount of code which are proceed at a time) can be specified as a basic block, a trace, a tregion, or an entire procedure. The emulation unit is also responsible to handle exceptions and interrupts (such as I/O interrupts for full-system simulation and system calls for user/application level simulation) during execution.

DBTs enable virtualization across ISAs by emulating a guest binary executable code in one ISA on a host machine with a same or different ISA. Modern DBTs also employ dynamic recompilation techniques (e.g. just-in-time compilation). This way, the translated code is instrumented to return information about the execution of each portion of the code which in turn can be used to optimize the rest of the execution (e.g. incremental optimization of hot regions). DBT has been widely used for many different purposes and applications (such as performance optimization, debugging, profiling, performance motoring and application migration). DBT, depending on the simulated ISAs and the host's ISA, can be classified as simulation engine into three categories: same-ISA, cross-ISA and retargetable (see Fig. 8). For the same-ISA, the simulated ISA is identical to the host's ISA. In cross-ISA, guest ISA differs from the host ISA. In retargetable DBT, the guest ISA (simulated ISA) can be retargeted for multiple different ISAs. Both same-ISA and cross-ISA can be considered as dedicated DBTs, since the guest and host architecture are fixed. Dedicated DBTs are limited by assuming that the register set of the host architecture is the same or richer than the guest architecture. This causes lack of translation flexibility and adaptation to highly heterogeneous environments.

Many of the current instrumentation-based simulators such as Graphite [11], PriME [14] and ZSim [13] benefit from DBT tools and libraries which leverage same-ISA DBT on x86 host architecture. Examples of these DBT systems include Pin [71], StarDBT [72], and DynamoRIO [73]. Emulation-based simulator generally leverage cross-ISA and retargetable DBTs to perform virtualization and emulation. IA-32 [74] and FX132 [75] are examples of cross-ISA DBT systems. There are few DBTs in the state of the art which have been designed for retargetability. Among them, QEMU [30] is a well-known emulator which implements retargetability. It enables binary translation from several different guest ISAs such as x86, PowerPC, ARM and SPARC on multiple common host architectures such as ARM, SPARC, Alpha, x86, MIPS and PowerPC. Retargetable DBTs broadly suffers from large emulation overhead ahead of translation, since the guest architecture must be fully virtualized by software in the memory, prior to translation. They also suffer from optimization overhead as well as code optimization overhead. Furthermore, providing near native and high quality translated code is another challenging issue for retargetable DBTs.

Target memory architectures: Depending on the memory architecture of the target simulated systems, manycore simulators can be categorized into shared memory, distributed memory and distributed shared memory.

In shared memory architecture, all the simulated processors simultaneously have access to a common single memory space in order to avoid data redundancy and facilitating inter-processor communications. In this way, the processors do not need to know where data resides, but they might suffer from race conditions. In distributed memory architecture, each simulated processor has its own private memory. Whenever communication is required (e.g. requesting a remote data), the simulated processors can communicate with each other through simulating message passing techniques on top of the simulated network and interconnect. Thus, it is necessary for every distributed memory simulator, to provide capability for network modeling (such as on-chip network). For example, RAMP-Gold [76] is a shared memory simulator which doesn't provide any network model. But distributed memory simulators such as COTSon [23] and BigSim [10] support network modeling.

Distributed memory simulators are more flexible than shared memory simulators for parallel simulation, since every simulated processor (i.e. a core) can independently be run on a different processor of a different host in the system. In distributed shared memory architecture, each node of a cluster has access to a shared memory in addition to its non-shared, private memory. Distributed shared memory simulators (e.g. SimK [52] and McSimA+ [69]) are generally required to support complex models of memory subsystems (i.e. memory and cache hierarchies) as well as communication networks.

Implementation: Depending on the strategy to implement functional and timing models, simulators can be classified into software-assisted simulators (i.e. simulators based on software simulation of the target architectures on CPU/GPU hosts), hardware-assisted simulators (i.e. simulators based on implementation of the target architectures on FPGA hosts) and hybrid simulators (i.e. simulators based on combination of both software and hardware simulation).

Most of the current architectural simulators are software-assisted, since software simulation is low cost, easy to develop, and flexible to explore various target architecture. COTSon, GEMS, SystemC [77], SimFlex and BigSim use software based simulation on CPU for both functional and performance models. In [78], authors introduce a software based simulator which runs on both CPU (by leveraging QEMU for functional modeling) and GPU (by proposing a GPU-based manycore accelerator). The work presented in [79] is another recent example of software based simulators. It leverages QEMU for software based functional modeling on CPU while uses a performance model, written in C and CUDA, designed to execute partly on the host CPU and partly on the host GPU.

On the contrary, hardware-assisted simulators are high cost and not flexible for architectural design exploration. They also take significant amount of time, memory and effort to develop and use, but they provide faster simulation speed compared to software-assisted simulators. Examples include RAMP Gold [76], that uses FPGAs for both functional and timing models. ProtoFlex [80] and FAST [18] are examples of hybrid simulators, where they use FPGAs to accelerate the performance model while using software based simulation for functional models. Furthermore, hybrid simulators like HASim [81] use FPGAs for functional modeling and CPU based software simulation for performance modeling.

Sampling: Sampling and reduction techniques enable faster simulation of large scale execution workloads through simulating representative portions of execution (i.e. samples) [82]. In this way, architectural simulators can be classified based on their sampling strategies. There are various methods for execution sampling in the current literature. Examples include random sampling [10], dynamic sampling [23] and statistical sampling [47,83] such as SimPoint [82] and Sampling Micro-ARchiTecture Simulation (SMARTS) [83]. Sampling techniques might reduce the simulation accuracy and precision in the estimate.

Assuming that a population is a complete set of elements (e.g. execution instructions of a complete workload) which need to be simulated. Sampling means simulation of a sample workload, instead of a complete workload. Each sample includes a subset of population (i.e. a set of sampling units) where a sampling unit is defined as a quantum of population. An ideal sampling approach must create a sample with choosing a minimal but (highly) representative set of sampling units in order to provide a reliable and quantifiable simulation (in terms of accuracy and speed).

A random sampling is performed by selecting and simulating only a random sample from the entire workload (the whole population) through a random selection of a fixed number of sampling units. The number of selected sampling units (or the sample size) must be sufficiently large for obtaining reliable (in terms of accuracy and precision in the estimate) simulation results (close to complete workload simulation). In a systematic (or periodical) sampling, sampling units are chosen from a population at periodical intervals. A uniform sampling, perform this with equal selection probability for (all) sampling units distributed in the entire population. In a representative sampling, sampling units are chosen from weighed regions of the population. In a dynamic sampling, the sample intervals (and the duration of each interval) are determined dynamically through monitoring the variation of execution behavior for the given benchmark [84].

Statistical sampling attempts to estimate (simulate) the execution behavior and characteristics of a given workload (i.e. a population) by selecting and simulating an optimal sample. This can be done by examining the adequacy of potential samples over the entire population using priori profiling (preprocessing), consisting statistical analysis (for example in terms of total, mean, and proportion) of the code. In other words, an statistical sampling prescribes a specified and constructive mechanism to determine an optimal sample (including a large number of tiny sampling units) in order to obtain a desired confidence interval. The sample must be able to capture the inherent variations of the given stream of execution instructions [85]. Statistical sampling is an efficient approach in capturing average behavior. And also it can specially provide an appropriate sampling when the population benefits from a clear internal control and contains a large number of similar transactions. However, it might be more complex and time consuming compared to non-statistical sampling approaches [86,87].

3.2. Virtual machines

Systems based on Virtual Machine (VMs) hypervisors are typical tools to model the hardware. These provide an emulated environment where the guest applications (particularly operating systems) can be executed virtually, very close to the way that they are executed on real hardware. Generally, virtual machines are faster than architecture emulators. This happens because the hypervisors create environments where substantial amount of instructions execute directly on the real underlying hardware without any overhead. VMs also permit execution times near to native speed. However, without support for hardware virtualisation, performance may be severely degraded, and not all currently available processors support virtualisation. They should meet a set of requirements which are elaborated in [88]. Of course, these can only be used when the system to test already exists, precluding the usage of VMs for many design space exploration tasks.

3.3. Architecture emulators

Architecture emulators are software tools which emulate the behavior and characteristics of a given CPU. They do not directly run the instructions of the emulated processor on the underlying real host processor. Rather, they employ methods of interpretation, or of dynamic translation, to translate the emulated instructions to a corresponding set of instructions for the target platform. Then, they execute the translated instructions on the hardware of the host platform. Therefore, architecture emulators generally decrease execution speed, but dynamic translation, a cache intensive technique, is used to enhance the speed.

Definitely architecture emulators are appropriate tools to be embedded in manycore full simulators for the purpose of functional modeling due to their accuracy, and reasonable speed. Qemu [30] and SimNow [33,89] are common examples of architecture emulators which have been employed in several full-system simulators.

Qemu is an open source full system emulator which uses the dynamic translation technique. It provides capabilities to emulate several types of CPU architectures (x86, PowerPC, ARM and Sparc) on several different hosts (x86, PowerPC, ARM, Sparc, Alpha and MIPS) to virtually run a complete and unmodified operating system. It also offers three operation modes: full system emulation mode, where the processor architecture and other peripheral devices are emulated; user mode emulation, where Qemu can launch executables compiled for one processor on another host processor which could have a completely different architecture; and finally Qemu acceleration mode, which executes most of the code directly on the hardware without dynamic translation, resorting to different execution rings.

AMD SimNow is a fast cycle, accurate full system emulator, using caching and dynamic compilation techniques. It can support booting an real operating system and launch complex applications over it. The SimNow emulator supports the x86

and x86-64 instruction sets, with support for other devices of a real system. It performs emulation of the real system with (at least) 10x slowdown in comparison with the native execution. SimNow cores generate a series of event that are stored in the asynchronous queues. COTSon provides timing feedback for SimNow instances, it parses asynchronous queues to create higher level objects such as instructions. COTSon is a complete tool which provides timing information back to the functional emulator in order to affect the behavior of the application. It also uses quantum based simulation [23,59] as synchronization technique. Quantum is the smallest, atomistic timing entity. i.e. the dimension of length (“time”) as a single entity is a quantum. In this way, every time a quantum starts, the timing module will get a notification about the starting time and the quantum length. Similarly once a node ends a quantum, the timing module will let the other modules to know about the network timing information which has been calculated during the past quantum. The functional simulation adds extra latency to all packets submitted, because network packets are sent twice. The source Network Interface Card (NIC) sends packets to the mediator timing module and the mediator will send the packets to the target NIC module. Additional time is required to for the packet processing in the mediator, therefore COTSon uses a Quanta (Q) bigger than real latency time between two nodes [90].

3.4. Network on chip, thermal and energy aspects

Most full-system simulators can simulate the interconnection network (particularly the Network On Chip (NoC)) as well as the entire processor and memory system. However, due to inherent limitations, they are not able to perform NoC simulations, in a very detailed level [91]. ASIM [92] is an example of this case. It is a full-system simulator, used in industrial laboratories, that simulates the processing cores and memory hierarchy. It only models ring interconnects and it cannot model other interconnects such as mesh, hypercube, x-tree, shuffle exchange, fully connected, butterfly, cube connected cycles, etc. RSIM [93] and SESC [94] are other examples of full-system simulators, whose interconnection network is not modeled precisely in these frameworks. On the other hand some of the simulators support modeling of the entire system, including the on-chip network, to a significant degree of detail, such as PharmSim [95].

For the purpose of designing new manycore simulators and particularly with the objective of modeling the future petascale parallel machines, we should consider a very efficient and accurate network-on-chip model. NoC is an integral part of the memory system and not modeling its details leads to an unsatisfactory model. Today, there are several NoC only simulators like NOXIM [96] and SICOSYS [93] which are used by the NoC community for experiments. However, these are not able to be used to perform full-system modeling and simulation.

In the current literature, SystemC [77] (refer to Section 4.3), a simulation framework, has been widely used to design full system simulators and particularly for NoC simulators. SystemC-based NoC simulators [97–103] are powerful tools to evaluate different NoC configurations by means of simulation. Xpipes [103] is an example of this type. It provides capability to simulate both homogeneous/heterogeneous NoC architectures for multiprocessor SoCs through using a set of flexible SystemC-based NoC macros, enabling to act as instance-specific network components at instantiation time.

One of the efficient solutions to enhance the capability of a full-system simulator is leveraging other accurate and high-performance network on chip simulators for NoC modeling instead of using a weak built-in network model. It is required for each of the simulator component and tools to provide modularity. SICOSYS [93] is an example which has been plugged into RSIM [36] for simulating symmetric multiprocessor systems.

Energy consumption in uniprocessor architectures grows linearly with the clock rate frequency and quadratically with voltage. Usually, lowering frequency permits operation at smaller voltages, and this has a cubic effect on energy savings. As for multicore and manycore processors, the power consumption increases linearly with the number of cores while clock frequency increases at a much slower pace [104]. Further, the use of identical processing elements in an homogeneous architecture reduces the overall hardware complexity and verification process in the entire development cycle.

Other than power consumption, thermal issues have a significant impact on the performance of manycore architectures. Simulating the thermal behavior of the manycore processors is an important objective in many different aspects. The architecture designers need to know and analyze the impact of their designs on the temperature of the processors. Similarly, thermal modeling is required by the OS and system designers. The problem is that this kind of modeling is computationally expensive. In processors, thermal time changes in time interval of the order of ten milliseconds. Thus, to simulate a very complex thermal impact of a proposed design, it is required to simulate the system in a high level of details at least for several seconds.

In manycore architectures, thermal status of the individual cores are related to the actual computing workloads of each core. Thus, analytical [105] and mathematical models as Fourier expansion [106] can be used to characterize and model manycore processor workloads and foresee the accurate amount of the processors load.

In the current literature, various tools and simulators have been proposed to investigate power consumption and energy efficiency of microprocessors and manycore architectural designs. Examples include Wattch [107], SimplePower [108], SoftWatt [109], XTREM [110], Orion [111], Orion2 [112], McPAT [113], Sim-PowerCMP [114], PrEsto [115], Sniper/McPAT [116], Manifold [59] and [117], providing capability for power or thermal modeling of computing processors. However, since physical modeling (power, energy and thermal) of manycore systems is out of the scope of our interest in this paper, we do not discuss the details of those approaches here.

Table 2

General specifications of simulators, emulators and virtual machines (Execution info: Information about the execution behavior, functional results: execution results).

Tools	Engine	Code translation	Output information
Virtual machines	Mainly direct execution	Same-ISA DBT	Functional results
Emulators	Emulation (Interpretation)	Cross-ISA & retargetable	Functional results
Emulation-based simulators	Emulation (Interpretation)	Cross-ISA & retargetable	Functional results, execution info
Instrumentation-based simulators	Instrumentation	Same-ISA DBT	Functional results, execution info

4. Architecture simulators

Parallel and distributed simulation tools such as SimFlex [47], GEMS [26,118], FastMP [42], SlackSim [43,60], BigSim [10] and COTSon framework [23,89], have been created in order to enhance the performance of simulation by concurrently distributing the simulator workloads among several parallel hosts. Furthermore, running parallel manycore simulators in accelerated hardware platforms such as FPGAs, and more recently in general purpose GPUs, helps to increase the throughput of the system simulation.

Architectural simulators provide capability to simulate either the target microprocessors, or the full hardware and software functionality of the target machines, as platforms, where platforms may consist of parallel processors, memory hierarchy, storage devices, I/O devices, compiler, OS, etc. Along the line of the growing issues in manycore era architectures, simulators have changed and they used new approaches to solve the problems and improve the system efficiency. They started from a simple approximate uniprocessor simulator, and continually improved to the recent high efficient clustered simulators.

Table 2 provides a generic comparison of architecture simulators and other simulation tools. However, this comparison is very general and depending on each specific tools it might differ. The table shows that simulators (both types) are the more flexible tools to analyze and predict the behavior of future systems. Furthermore, emulation-based simulators are the most powerful tools, since they can potentially simulate any type of architectures. On the other hand, instrumentation-based simulators are faster, since they provide near-native execution speeds.

In the remaining of this section, we discuss some of the most important architecture simulators that have been proposed in the last decade.

4.1. SimpleScalar

SimpleScalar [9] is one of the oldest uniprocessor serialized software simulators, which previously was widely used by the research community focused in processing architectures. As well as other similar simulators, SimpleScalar was limited to only run single-threaded, user-mode workloads. With the advances in manycore micro-architectures, researchers and designers are more interested to use simulators with the capability to run multi-threaded workloads, and to model large number of processing cores along with the memory subsystem and interconnects.

4.2. BigSim

BigSim is a parallel simulator and performance modeling system which is particularly designed to study parallel programming issues [10]. It can predict performance of parallel applications on machines (like IBM Blue-Gene/L systems) [119] with a very large number of processors (i.e. large number of processing nodes). This is done by actual execution of real applications on smaller machines (i.e. small number of processing nodes). Indeed, BigSim can be used for the architecture design of manycore-enabled HPC systems. It has been built on POSE [120] and includes several components. POSE is a general-purpose optimistically synchronized PDES (parallel discrete event simulation) framework which is designed for scalability of fine-grained parallel and distributed large-scale discrete event simulations. However, its load-balancing framework still needs to be improved. BigSim simulates the behaviors of communication and computation separately in two steps. At first, it uses an emulator to execute an application, containing number of virtual processes, on a small number of physical processors to generate trace logs. At the next step, a trace-based simulator uses the log files and simulates activities on a much larger processing system.

BigSim directly executes the application (in small scale) using its emulator and mimics the behavior of the target platform (in large scale). The direct execution creates significant demands on host CPU and memory. For this reason, the simulator allows skipping computations and instead simulates latencies that would be resulted by executing those computations. But the problem is, this works only for data-independent applications, due to the fact that some part of the data are not really computed. The simulator explores the inherent determinacy of several parallel applications. But still it is not an application-independent performance modeling system and its functionality is limited to specific applications. Moreover, tracing in BigSim is specifically designed for its implementation language Charm++, and doesn't support the message passing applications.

BigNetSim, a BigSim's component, uses a simple analytical model (e.g. SimGrid [121]) to simulate interconnection networks by supporting detailed network models of various topologies. However it still doesn't provide capability for packet-level interconnections simulations (e.g. MPI-NetSim [122]) which is the most precise approach (while it is very resource-consuming) for network simulation. This leads to reduction in the overall system accuracy. In fact, the major drawback of BigSim is that it suffers from inaccuracy (with respect to the expected behaviors of the real system) caused by log-based postmortem simulation of the generated traces.

4.3. SystemC

The SystemC language [77] is an extension of C++, providing a cycle-accurate, event-driven, simulation interface for system-level modeling by describing modules of a target architecture as a set of C++ classes. It is a popular framework for SoC architecture simulation, providing a powerful interface to describe HW/SW components as well as interconnections between modules (ports and signals), facilitating description of interconnection between multiple SoC processors. SystemC also provides support for integrating different Instruction Set Simulators (ISS) in a unified system simulation framework, as it is able to plug an independent ISS into the entire simulation framework (as a new system module), where all system modules can be activated and synchronized through a common reference clock. Furthermore, SystemC based simulators (such as [123–127]) can benefit from advantages of C++ language as a hardware description language while bridging the gap between hardware and software description languages [123]. In fact, one of the most important advantages of C/C++ based description languages such as SystemC (or SpecC [128], a similar alternative) is their capability to concurrently specify both hardware and software components in the design (i.e. co-simulation of both hardware and software). This is a necessary requirements for full system simulators (system level).

In the current literature, SystemC has also been extensively used for designing NoC simulators, due to the powerful capability of SystemC to describe various interconnections between hardware components. Examples of this include [97–103].

SystemC based simulators are accurate and sufficient for validating hardware specifications. However, they might fail to adequately support embedded software (in terms of writing or debugging), which is an important requirement for SoC design [129], are often slow compared to the traditional ISSs like SimpleScalar.

MPARM [123,124] is an example of such simulators. It is a full-system SystemC-based architecture simulator, enabling to model functional, performance, and power consumption aspects, as well as a complete OS for a Multi-Processor System-on-a-Chip (MP-SoC), in a cycle accurate manner. Using SystemC, MPARM (or MP-ARM) can provide processor models, memory models, the AMBA bus architecture (for communication between models through ports and signals) and support for parallel programming. However, the MPARM simulator is slow and its models for processing cores are very abstract and relatively simple, lacking detailed core modeling [130].

One of the conventional solutions to overcome the speed limitation of SystemC-based simulators is to integrate SystemC with QEMU, making possible to simultaneously benefit from accuracy of SystemC and speed of QEMU. However, interfacing between SystemC and QEMU might be challenging, since their combination must be capable of accessing all the hardware modeled in QEMU and SystemC for co-simulation of HW/SW. For doing this, SystemC needs the QEMU support to provide I/O operations (initiated by the processor), memory access interface, interrupt handling and also peripherals to access memory directly [131]. Furthermore, there are timing aspects which need to be taken into account for synchronization between SystemC and QEMU models [132], making SystemC-QEMU combination a complex task. There are several research works in the current state of the art, which perform this using different approaches [131,133–136]. Among these types of simulators, Virtualsoc [8,137] is a recent work for many-core-based accelerators, allowing the execution of a full-fledged Linux operating system.

4.4. Graphite

Graphite [11] was created for the exploration of large-scale manycore environments, as well as for research of isolated applications. It can be used as a distributed, high-level parallel simulator. In order to deliver the high performance and scalability needed for useful evaluations, it uses various methods such as direct execution, multi-machine distribution and analytical modeling. In addition, it benefits from lax synchronization schemes like LaxP2P [138–140] (a distributed synchronization technique, in which the progress of one core is periodically checked against another randomly selected core). Graphite has other important capabilities, such as its flexible and extensible architecture, its compatibility with commodity multicores and clusters, its ability to run off-the-shelf p-threads application binaries, and its support for a single shared simulated address space despite running across several physical host machines.

Graphite, unlike BigSim, FastMP and COTSon, allows the analysis of a much wider category of architectures. While it offers the possibility to model distributed memory architectures, it also provides a coherent shared memory between the simulator threads [11]. In addition, Graphite also models compute cores and interconnected networks while operates transparently through providing a single shared address space to off-the-shelf applications.

Graphite can simulate manycore target architectures with hundreds of cores launched on several parallel hosts, but the problem is that the simulator accuracy resides in the application-level. It is not very successful to deal with speed/accuracy

challenges. Graphite has three different methods for synchronization: base model (or Lax synchronization), barrier (LaxBarrier or Lax with quanta-based barrier synchronization) [140] and random-pairs (LaxP2P or Lax with point-to-point synchronization). Lax lets the clocks differ and offers the highest performance and scalability. However, in order to keep the simulated clocks in reasonable agreement, Graphite needs to deploy application events to make them synchronized, otherwise it must let the threads run freely.

LaxBarrier and LaxP2P [58,140] are the mechanisms on top of Lax to improve its accuracy. LaxBarrier is the most accurate synchronization methods of Graphite where all active threads must wait on a barrier after a configurable number of cycles. It has a relatively poor performance and scalability compared to the other two synchronization models. LaxP2P aims to achieve the accuracy of quanta-based LaxBarrier without reducing the scalability and performance of lax synchronization. Using this scheme each tile periodically chooses another random tile and synchronizes with it. If the clocks in the source and target tiles differ by more than the number of configured cycles then the tile which is ahead goes to sleep for a short period of time. LaxP2P is fully distributed therefore it creates less overhead than LaxBarrier. In comparison to the barrier method it provides more scalability and less accuracy [11].

4.5. SimK

SimK [52] is a framework based on the Parallel Discrete-Event Simulator (PDES) synchronization protocol [141] to develop system simulators. PDES [53] is a well known parallel distributed synchronization technique for parallel simulation. SimK provides simulation modules that target system components, such as CPUs and memory modules. All modules communicate through message passing methods, which enables them to run concurrently. A dedicated module maintains the time synchronization of all simulation modules. P-GAS [142], HPPNetSim [143] and G-Cluster [144] are the simulators which have been developed based on the SimK framework.

Since each component has to synchronize the execution state with its peers, continuously at a microsecond rate, SimK employs several optimization strategies and techniques to avoid the severe performance degradation that synchronization would impose. Each node is handled by a single process, which further creates one thread per processor. A user level scheduling scheme is employed where simulation modules are dispatched to each thread. CPU affinity is used to avoid cache related performance penalties. Since the simulation modules run on the same process and share the same memory, SimK employs an asynchronous zero-copy [52] communication mechanism. Further synchronization optimizations are employed at the scheduling level to avoid blocking of the simulation modules. Other optimizations employed include lock-free queues, buffer management and load balance.

SimK requires a host shared-memory multiprocessor system. While it has shown to scale within this system, the lack of cache coherence on manycore systems does not allow the shared-memory dependent approach used by SimK to be effective. Thus, the major bottleneck of this simulation framework, which is synchronization, cannot be solved with the approach taken by SimK on a manycore system. This severely limits the scalability of SimK to multiprocessor systems.

4.6. GEMS

GEMS [26,118] is a full-system simulation platform capable of capturing detailed aspects of processing cores, cache hierarchy, cache coherence, and memory controllers. The simulation platform consists of a set of nodes connected with links allowing for wide variety of topologies, with each link having a particular latency and bandwidth. This has led to the widespread use of GEMS in the computer architecture research community, with a huge amount of contributions for validating research ideas.

A major limitation of GEMS is its simple interconnection model that serves as a communication fabric between various cache and memory controllers. Messages traverse the network hop by hop, which makes GEMS incapable of modeling a detailed router or a network interface. In fact it does not integrate a real interconnection network model [145]. Because of this limitation, GEMS ignores buffer contention, switch and Virtual Channel (VC) arbitration, realistic link contention and pipeline bubbles. The GEMS interconnect SimpleScalar model also assumes perfect hardware multicast support in routers. However, considering on-chip network designs, supporting fast and low power hardware multicast is currently still a challenge.

The limitations in the interconnect model can significantly affect the results reported by the current GEMS implementation. Thus, GEMS has not been adopted by researchers focusing on low-level interconnection network issues. Researchers instead rely on traffic trace-driven simulation with synthetic or actual traces. In a trace-driven approach, the functional simulation is not impacted by the timing simulator. Timing dependent effects are not captured because the trace is generated a priori on a fixed system, and the timing variation caused in the network does not affect the message trace. Trace-driven techniques also do not capture program variability that a full-system evaluation can.

4.7. SimFlex

SimFlex [5,47] is a full system, component-based simulator, inspired by the ASIM simulator [92], which enables creation of timing models for uni and multiprocessor systems running unmodified commercial applications. It integrates SMARTS methodology [83] for simulation sampling and Simics [34] for functional modeling with techniques to avoid runtime overheads that arise from component-based software design. Simics is a high configurable simulator with a basic timing model

(i.e. uniform timing for all instructions and memory accesses). It can provide functional execution of unmodified commercial OSs and applications for wide variety of systems and ISAs (e.g. x86, SPARC, etc.). SimFlex enables full system simulation by augmenting Simics with a framework for rapidly building complex timing models. SimFlex receives a stream of fetched instructions from Simics, and then models the timing behavior of the system while controls the timing advancement in Simics.

SimFlex is a collection of C++ described components connected together in a hierarchical fashion where each component represents a hardware or software component of a real system. Each component is connected to other components by definition of ports (i.e. unspecified C++ template parameters during component development). The connections between components can be configured and specified in a C++ code which is called wiring description. And accordingly, when these wiring descriptions are fed to the compiler, the interconnection between component is created at compile-time. This way, SimFlex framework can produce a custom timing simulator binary which reflects the desired wiring description. Examples of SimFlex timing models include UniFlex (uniprocessor simulation), TraceFlex (no timing), CMPFlex (in-order timing) and CMPFlex.OOO (out-of-order timing).

Most of applications exhibit homogeneous execution phases which can include millions of instructions. Statistical sampling reduces the amount of simulation effort required for performance estimation of such applications and leads to increase the simulation speed. SimFlex, using SMARTS rigorous statistical methods, identifies the minimal sample that assesses application performance with a required confidence level.

SimFlex offers detailed multiprocessor memory systems but lacks detailed I/O models and multiple-system capability. Furthermore, SimFlex suffers from the lack of flexibility for development, particularly for non-academics, since it relies on Simics, which is a commercial functional simulator with licensing restrictions.

4.8. COTSon

COTSon is the HP labs' full system simulator based on AMD SimNow. It allows researchers to trade speed for accuracy, depending on the simulation purpose and the user preferences. It uses a trace driven approach where a single core, full system, simulator generates thread instruction streams. This helps COTSon achieve better simulation speeds, compared to execution-driven simulation, because of the decoupling of functional simulation from detailed simulation. It also is possible to simulate any application that might have been run on different platforms, given the correct tracing infrastructure.

Functional simulation generally adds some extra latency to every transmitted packet. It happens because packets are sent to the mediator and then the mediator sends them to the destination. COTSon performs bandwidth and network simulation in the sender NIC device, but all the network timing characteristics and information is collected at the mediator level. Potentially, COTSon can present reasonable speedups in comparison to some of other simulators, but it assumes an idealized architecture consisting of a perfect memory hierarchy, which is a little far from real architectures.

COTSon leverages some of the other existing simulators and tools for individual sub-components through a robust interface layer, integrated closely with the COTSon timing models to improve simulation accuracy (versus simulation speed). Notable examples are SimNow, for the purpose of functional modeling, and Q-Mediator, for simulation of interconnected network. However, COTSon does suffer from using SimNow, which is a sequential emulator.

SimNow produces a sequential instruction stream as output which is demultiplexed into different threads before timing simulation. It creates a significant drawback for COTSon, limiting parallelism and restricting the simulator to a single host machine for shared-memory simulations. On the other hand, COTSon is able to launch simulations over an overlay network consisting of multiple manycore machines. But, it becomes limited to run applications which are created for distributed memory environments and use a message-passing library (like MPI [146]).

4.9. RAMP-Gold

RAMP-Gold [147] performs system modeling by employing a single timing pipeline, coupled with a single functional pipeline with moderate resource consumption launched on a low-cost mid-size FPGA. The simulator design is constrained in the total amount of cache capacity that we can model by using the BRAM consumption of the timing model [76,148].

FPGAs include Block RAM (BRAM) and Distributed RAMs (DRAM). Block RAMs or BRAMs are dedicated memory blocks and DRAMs are the RAMs that are constructed using Look-Up-Tables (LUT). LUTs are distributed across the FPGA fabric and they can be used as small blocks of RAM by combining DRAMs. RAMP Gold has the capability to simulate 64 SPARC CPUs over 250 times faster than a regular software-based system simulator launched on a Xilinx Virtex-5 board (which is low-cost in comparison to other FPGA hardware). This demonstrates the cost performance benefit of FPGA-based simulation. The design of RAMP Gold also shows that designing FPGA-based architecture simulators is dramatically different from designing multicore processors in either ASICs or in FPGAs.

4.10. Other FPGA-based simulators

FAST [149] is a hybrid FPGA-based simulator, whose functional model runs in software and its timing model runs in FPGAs. It needs a significant amount of communication bandwidth between FPGA and CPU, which may result in limited simulation scalability. ProtoFlex [150] and HASim [81] are other simulators which use FPGAs to implement timing models

for full-cycle-accurate simulations similar to FAST. Unlike FAST and ProtoFlex, HASim implements its functional model in FPGA and timing model in CPU.

In general, FPGA-based solutions are costly and require the user to buy expensive hardware. In addition, FPGA-based approaches are difficult for development and it is not easy to quickly experiment with various designs while implementing new models in FPGA instead of CPU.

4.11. GEM5

GEM5 [4] is a modular, sequential, full system simulator, based on a combination of both M5 [29] and GEMS [26] simulators. It merges the high configurability of M5 with high detailed and flexible memory subsystem modeling of GEMS. This leads to the support of a wide range of simulation features ranging from multiple ISAs (such as Alpha, ARM, MIPS, Power, SPARC, and x86), diverse CPU model (non-pipelined and pipelined CPU models such as AtomicSimple, TimingSimple, In-Order, and Out-Of-Order), detailed cache hierarchies (using Ruby memory model) and multiple cache coherence protocols (using SLICC language) to the instantiation of various interconnection networks (Ruby based network models such as Simple network model and Garnet [151] network model), I/O devices (ranging from simple timers to complex network interface controllers) and multiple systems. Despite the high accuracy and flexibility of GEM5, it suffers from very slow simulation speed due to the lack of parallelization in simulation.

4.12. McSimA+

McSimA+ [69] provides a lightweight, flexible, and detailed microarchitecture-level simulator by offering a middle ground between a full-system simulator and an application-level simulator. It can benefit from the light weight of an application-level simulator, while the simulator is able to fully control the threads and processes (similar to a full-system simulator). McSimA+ can simulate x86-based asymmetric manycore systems (up to more than 1000 cores) for both core and uncore subsystems. The modeling for asymmetric cores can be ranged from single-threaded to multi-threaded workloads and from in-order to out-of-order CPU models. It also supports sophisticated cache hierarchies, coherence hardware, on-chip interconnects, memory controllers, and main memory. Furthermore, using DBT as simulation engine and the ability to support both execution-driven and trace-driven simulations, McSimA+ is able to reasonably improve both simulation speed and accuracy. However, McSimA+ suffers from the inherent limitation of application-level simulators, which is the lack of full support for OSes and applications with complex I/O system calls and extensive system events. Moreover, the simulation workloads for McSimA+ is limited to only Pthread applications due to the frontend Pthread library of McSimA+. Non-Pthread multithreaded applications cannot be executed on McSimA+ without re-targeting the thread interface. In addition, the simulation accuracy of McSimA+ is most likely suboptimal as compared to most of the cycle accurate emulation based simulators. McSimA+ is also limited to support modeling of speculative wrong path executions due to the inherent limitation of its Pin [71] based functional modeling. In fact, Pin does not provide wrong-path instructions, since wrong-path instructions in Pin are invisible beyond the ISA interface and they are not committed in the native hardware.

4.13. ZSim

The ZSim [13] simulator introduces three techniques in order to enable fast, accurate, and scalable simulation for many-core systems (up to 1024 cores). These techniques are: detailed DBT-accelerated core models, in order to increase sequential simulation speed; bound-weave parallelization, in order to achieve scalable and accurate parallelization; and lightweight user-level virtualization, in order to provide support for simulation of complex workloads. ZSim proposes instruction-driven timing models, based on Pin DBT, to perform most of the core's operations in the timing model during instrumentation (i.e. eliminating the need for functional modeling of x86). This reduces most of the FM-TM overheads compared to conventional cycle-driven or event-driven core models. It also leads to fast sequential simulation (between 20 and 90 MIPS) while the core modeling is detailed enough to allow simulation of detailed core models like Out-Of-Order (OOO), including features such as branch prediction, limited fetch and issue widths, and μop fission.

ZSim also introduces the bound-weave algorithm which is a two phase event-driven parallelization technique that scales parallel simulation without increasing overheads or losing accuracy. The main insight behind this algorithm is that, at a small time scale, most concurrent core-memory accesses occur to different unrelated cache lines. This means that out of order simulation of these accesses at first and then simulating their detailed timing in order, can be equivalent as simulating them completely in order. This allows reorderings of instructions only within a small interval (e.g. within 1000 cycles) with assumption that, in such case, path-altering interference is exceedingly rare.

A two core-memory accesses might suffer from path-altering interference, if out of order simulation of those accesses, modifies their paths through the memory hierarchy (e.g. two write for the same cache line from two different cores). ZSim divides the simulation into small intervals of a few thousand cycles each. As we already mentioned in Section 3.1, for each interval the simulator proceeds bound phase and weave phase. In bound phase, cores are simulated in parallel with assumption that all memory accesses have initially zero latencies and for each memory access, the path through memory hierarchy is recorded. The bound phase also puts a lower bound on the cycle of each microarchitectural event. Accordingly per-core event traces are generated through instrumenting all loads, stores and basic blocks of the host thread which simulates a

core. In the weave phase, these traces are used to perform parallel microarchitectural event-driven timing simulation of the aforementioned memory accesses by dividing events among parallel event queues, and simulating them in full order.

While bound-weave parallelization is efficient to provide scalability for parallel simulation, it is only applicable for shared memory target systems. The reason is that the bound-weave approach relies on the assumption that path-altering interference being rare, which is true but only for the cores that communicate implicitly through the cache hierarchy. This means that bound-weave mechanism is not applicable for other target systems like distributed systems with message passing inter-core communications. Moreover, ZSim is a user-level simulator but since it is using a set of lightweight user-level virtualization mechanisms, it can run most modern workloads (such as client-server workloads and multiprocess applications) without any modifications.

4.14. PriME

PriME [14] is an MPI-based, manycore, x86 simulator. It uses Pin as simulation engine and combines both shared memory and message passing techniques in order to achieve high parallelization and distribution of simulation workloads through running simulation across any MPI-enabled cluster of multiprocessor and multicore machines. This allows PriME to support multi-threaded workloads as well as multi-programmed workloads through benefiting from two levels of parallelization; within a host machine and across host machines. The simulation of multi-threaded workloads can be parallelized inside of a host machine while the simulation of multi-programmed workloads can be parallelized across several host machines. This happens by utilizing MPI to communicate among different PriME modules.

PriME demonstrates reasonable performance and speed for manycore simulation (for example, it can simulate more than 2000 cores on 9 machines with a total number of 108 cores). However, similar to other simulators, it also suffers from drawbacks. While the simulator supports detailed modeling of uncore components such as the memory subsystem and NoCs, it can only simulate simple in-order core models with a constant cycles-per-instruction (CPI) for non-memory instructions. Therefore, using PriME, simulation of the detailed core models, like modern out-of-order processors, is not feasible. While PriME is fast enough to explore thousand-core architectures, it is not cycle-accurate. While PriME offers high configurability for uncore models such as memory subsystem, cache coherence and interconnect models, the configurability of the simulator for the core models is poor. For example, the simulated ISA is limited to x86 architectures.

4.15. Comparison

Tables 3 and 4 provide comparative analysis of some of the most important multiprocessor and manycore architecture simulators with respect to the taxonomies presented in Section 3.1. The results of these comparative analysis, presented in the tables, also demonstrate the qualitative ideas expressed on Section 3.1. The discussion of the consequences of these analysis will be deeply used in the Sections 5 and 6 in order to discuss the status of manycore future simulators.

5. Fundamental simulation challenges

Two important features of future peta-scale systems are high-heterogeneity of cores and, very large number of cores. Due to these features, the requirements for simulation of future peta-scale manycore systems include supporting:

- Scalability, in terms of number of cores simulated.
- Heterogeneity, in terms of diversity of cores.
- High speed simulation, due to the complexity and large number of components of such systems.
- Accuracy, in terms of level of detail and similarity to the target systems.
- Flexibility to rapidly explore a very large design space, due to the high diversity of potential architectures for future peta-scale systems.
- And low cost simulation, in terms of hardware required as simulation host.

Efficiently designing future peta-scale manycore architectures is not achievable unless we use powerful simulation tools, fulfilling the abovementioned requirements, and able to drive micro-architecture exploration, evaluate novel systems and design decisions. It is obviously desirable to have a simulator with the highest level of accuracy and scalability, while providing a very fast simulation speed. Along this line, the typical full-cycle-accurate simulators cannot be used anymore, since in some cases, simulating a single second of execution can take between 1 and 12 days. We expect simulators to deal with a lot of design space variables, such as heterogeneity, scalability, modularity, in order or out-of-order core models, cache coherency, memory hierarchy, distributed or shared memory, accelerators, etc. However, there are critical challenges that create difficulties and obstacles to produce such a desirably simulator.

Improving large manycore architectures with large amount of cache memory led simulators to employ highly parallel methods to distribute simulation workloads (as it is the case in current parallel simulators such as PriME and Graphite). However, due to data dependency between parts of the simulated environment, full parallelization is not easy to achieve and we return to the problem of concurrency in current architecture models. Additionally, scale-out from core/tile to die, socket, rack and data-center is really challenging, although this is a necessary requirement for evaluation of future peta-scale machines.

Table 3

A comparative analysis of various many-core architecture simulators.

Characteristics	SimpleScalar (2002)	BigSim (2004)	GEMS (2005)	SimFlex (2006)	SimK (2009)	COTSon (2009)
Simulation Execution	Sequential	Parallel	Parallel (sequential functional simulators)	Parallel (parallel simulation of live points or flex points, sequential functional simulators)	Parallel, distributed	Parallel (sequential functional simulators)
Simulator Organization	Timing directed (execution driven, interpreters) and functional-first (trace driven)	Functional-first (trace-driven)	Timing-first and timing directed (execution driven)	Cycle driven execution model, timing-first for out-of-order multiprocessor simulation	Functional-first (trace driven) and Timing directed (execution driven)	Functional-directed (feedback-driven)
Simulation Scope	User/Application-level simulator	User/Application-level simulator	Full system simulator	Full system simulator	Full system simulator	Full system simulator
Simulation Engine	Emulation	Emulation	Emulation	Emulation	Emulation	Emulation
Target Memory Architectures	Shared memory	Distributed memory (cluster-based simulation)	Shared memory (with cache coherence support)	Shared memory	Distributed shared memory	Shared memory, distributed memory
Scalability	Very Poor	Poor	Poor	Moderate	Moderate	High
Accuracy	Cycle accurate (cycle timers)	Instruction-level/cycle-accurate	Cycle accurate	Cycle accurate	Functional accurate	Cycle accurate
Complexity of Setup	Moderate	Moderate	High	Low	Moderate	Moderate
Cost	Low	Low	Low	Low	Low	Low
Many-Core Support	Uniprocessor	Multiprocessor is supported	Multiprocessor is supported	Multiprocessor and multi-core are supported	Multiprocessor and multi-core are supported	Supported (up to 1024 core)
Implementation Basis	Software simulation Embedded (functional: sim-fast, sim-safe, sim-profile, sim-cache, sim-cheetah, sim-BPred, functional with timing: sim-outorder)	Software simulation BigSim emulator (functional simulation), BigNetSim (performance simulation)	Software simulation Simics (functional simulation), Rubby (timing simulation)	Software simulation Simics (functional simulation), SMART (performance simulation)	Software simulation HppSim (performance simulation), HppSwSim (functional simulation), HppNetSim (performance simulation for interconnection)	Software simulation SimNow
Modeling Features	Single-core CPU models from simple unpipelined processors to detailed dynamically scheduled micro-architectures with multiple-level memory hierarchies	Network models (detailed models of communication networks e.g., network modeling for large message-passing machines such as Blue Waters and BlueGene), uni-processor models	Memory (Ruby model for caches, cache controllers, memory controllers, and banks of main memory), Ruby model for interconnection network, Opal CPU model	CPU models (uni- and multiprocessor systems), CMP and DSM system models, aggressive out-of-order core tuned to produce high memory parallelism	CPU cores, memory without hardware cache coherence, network adapters, routers, switches, interconnection network	Compute cores, networks, memory, the OS and common devices such as disks, video, or network interfaces

(continued on next page)

Table 3 (continued)

Characteristics	SimpleScalar (2002)	BigSim (2004)	GEMS (2005)	SimFlex (2006)	SimK (2009)	COTSon (2009)
Event Sampling	Not Specified	Random sampling of sequential execution blocks, Slicing, Miniaturization	Not Specified	Multi-threaded rigorous statistical sampling using Flex Points and SMARTS	Not Specified	Single-threaded dynamic sampling
Synchronization	Not Specified (only inter-stage latch synchronization for register update and memory load/store in sim-outorder)	Optimistic PDES (taking advantage of the parallel programs' inherent determinacy to reduce synchronization overhead)	Not Specified	SimFlex does not have an explicit synchronization component but synchronization works exactly as it would on a real SPARC machine.	Conservative PDES (fine-grained block/unblock-based synchronization)	PDES
Supported ISAs	ALPHA, PISA, ARM, PowerPC and x86	x86	SPARC-v9, x86, etc.	ALPHA, x86-64, IA-64, ARM, MIPS, MSP430, PowerPC, POWER, SPARC-v8 and v9, etc.	(on CPU)MIPS64, MIPS III	x86-64
Workloads	Single-threaded workloads (realistic applications with relative simplicity and less efforts)	MPI, CHARM++ and Adaptive MPI applications	Not Specified	Not Specified	Not Specified	MPI applications

Table 4

A comparative analysis of various many-core architecture simulators.

Characteristics	Graphite (2010)	RAMP-Gold (2010)	GEM5 (2011)	McSimA+ (2013)	ZSim (2013)	PriME (2014)
Simulation Execution	Parallel, distributed	Parallel	Sequential	Parallel	Parallel (Bound-weave)	Parallel, distributed
Simulator Organization	Timing directed (execution driven)	Multithreaded FAME (FPGA Architecture Model Execution)	Timing directed (execution driven)	Supports both execution-driven and trace-driven, event-driven backend simulator	Timing directed (execution-driven), (instruction-driven core, event-driven uncore)	Timing directed (execution-driven)
Simulation Scope	User/Application-level simulator	Full system simulator	Full system simulator	A middle ground between a user/application-level simulator and a full system simulator	User/Application-level simulator	User/Application-level simulator
Simulation Engine	DBT	Emulation	Emulation	DBT	DBT	DBT
Target Memory Architectures	Coherent shared memory across a cluster of machines, distributed memory	Shared memory (no coherence)	Shared memory (with cache coherence support)	Shared memory (with cache coherence support), distributed memory (cluster-based simulation)	Shared memory (with cache coherence support)	Shared memory (with cache coherence support), distributed memory (across a cluster of machines)
Scalability	High	Moderate	Poor	Moderate	High	High
Speed	Very Fast	Fast	Very Slow	Moderate	Fast	Fast
Accuracy	Not cycle accurate	Cycle accurate	Cycle accurate	Cycle accurate	Not strictly cycle accurate	Not cycle accurate
Complexity of Setup	High	High	Low	Moderate	High	Moderate
Cost	Moderate	High	Low	Low	Low	Low
Many-Core Support	Supported (up to 1024 cores)	Supported (up to 64 cores)	Supported (up to 64 cores)	Supported (up to 1024 cores)	Supported (up to 1024 cores)	Supported (up to 2048 cores)
Implementation Basis	Software simulation Pin	FPGA implementation Embedded	Software simulation M5 (processor simulator), GEMS (network and memory simulator)-> Ruby (cache and simple networks), Garnet (complex networks), SLICC (coherence protocols)	Software simulation Pin based frontend simulator for functional simulations and the event-driven backend simulator for timing simulations	Software simulation Pin	Software simulation Pin based frontend simulator for functional simulations and the event-driven backend simulator for timing simulations
Modeling Features	Core models, memory subsystems (cache hierarchies with full cache coherence), on-chip networks	Cache models, abstract core models, CMP models, no network models	Core models, pipelined model, memory subsystems, networks models, system execution modes	Asymmetric core models, memory subsystems and network models	Detailed DBT-accelerated core models, memory subsystems (cache hierarchies)	One-CPI core models, profiling-based core models, memory subsystem, other uncore models
Event Sampling Synchronization	Statistical sampling Lax, lax with barrier and lax with P2P synchronization	Not Specified Cycle-level synchronization	Not Specified Not Specified	Not Specified Pthread scheduling	Not Specified Bound-Weave parallelization	Not Specified Thread-level and process-level barrier synchronization
Supported ISAs	x86	SPARC-v8	ARM, ALPHA, MIPS, Power, SPARC, and x86	x86	x86	x86
Workloads	Unmodified Pthread applications	Not Specified	Multi-threaded applications in System-call emulation mode and variety of workloads in full-system mode.	Unmodified Pthread applications (single-threaded and multi-threaded workloads)	Most modern and complex workloads (such as multithreaded applications, JVM and client-server workload)	MPI applications

In the remaining of this section, we discuss the major challenges for simulation of future peta-scale systems. We must also note that, in this paper, due to the space limitations, and to avoid losing focus, we avoid the discussion on the issues related to implementation details of the architectural simulators (such as scheduling, synchronization, etc.), since most of these issues are not particularly specific to the area of microprocessor and manycore simulation, but instead generally apply to many different areas in computing.

5.1. System complexity vs simulation capacity

Complexity arises in manycore systems by having more and more diversity and intricacy in the computing architecture and interconnects. For future peta-scale systems, we expect simulation tools to provide functionalities to evaluate the complexity of such nonlinear systems by modeling properties and behavior of each individual component, showing ways where the system behavior cannot be assumed as the sum of the behavior of its parts. For instance, a minor manipulation on the processing/communication properties/behaviors of one system element may cause a significant impact, a proportional impact, or even no impact on other elements or the whole system. However, managing such complexity is a critical challenge in manycore system simulators which are constrained by limited capacity to deal with the system complexity like coupling timing models with functional models. For example, timing models can be multiple *orders of magnitude* slower than real time.

5.2. Performance

Execution performance for manycore simulators is pushed to its limits when a parallel simulator divides its main sequential instruction stream in a set of segmented simulation workloads which have to be executed on a group of individual processors. Basically, it is not just the problem of simulators, but rather it is a challenge in the upcoming manycore architectures, which obligate a paradigm shift in algorithmic design to effectively fully unlock manycore capabilities to achieve maximum performance. Thus, simulators have to deal with challenges such as increasing the level of parallelism, multi-scaling, hardware acceleration, etc., both during development and implementation phases.

5.3. Speed vs accuracy

Interconnection networks and hardware details of manycore machines can be modeled at very different levels of abstraction depending on the target usage of the simulator. So there is a trade-off between simulation detail and simulation speed. The level of detail of the models used for describing the resources employed within the simulator has a direct impact on the final behavior and characteristics of the simulator. Hardware resource Description Languages (HDLs) are essential tools to build an architecture simulator, and to build an efficient simulator we must use a powerful, accurate and appropriate resource description model. While classical HDLs, like Verilog and VHDL [152], are very good at describing detailed hardware characteristics and behaviors (such as timing behavior), they are generally inadequate for expressing the higher-level abstractions required for (today and) future large and complex micro-architectural architecture designs.

The Architecture Description Language (ADL) [153] and the XML-Based ADL [154] are examples of the higher-level resource description model, which have been used to build manycore simulators such as Mhetero [155] and M3C [156]. C λ aSH [157–159] is another example of functional resource description which is based on Haskell. Using an efficient resource description language might help to increase the total performance of the constructed simulators.

Ideally, in order to achieve a proper peta-scale-level simulation platform it is required to describe the hardware details and interconnections as detailed as possible while keeping the total throughput of the system within an acceptable level. This requires the (simulator) designers to build methodologies for constructing calibrated models. Unfortunately, this is not easy to perform, and a common strategy in current simulators is to balance the need for simulation accuracy versus the desire for good simulation performance. Better modeling of target system details by a simulator will produce more accurate results but will result in a slower simulation. This trade-off is particularly important when modeling the interconnection network of a parallel computer.

5.4. Development cost

The cost of developing a new validated and useful simulator refers to two different aspects. On one hand, the developing time of a simulator should not exceed a reasonable duration. For example if we go to design a more detailed cycle-accurate modeling system, this results in consumption of more time on development, which is not desired. On the other hand the designed simulator should be modular, pluggable and able to reuse in future related works. In addition, new methodologies for constructing simulators have to show an acceptable performance while executing on the underlying hardware, and should not be dependent of some specific-purpose hardware infrastructure.

However, sticking to the objective of cost efficient modeling solutions in case of either the cost of the required hardware (e.g. current software-based simulators like GEM5), or the cost of time consuming for development (e.g. SimK), creates substantial constraints on implementation of some efficient solutions in aspects other than cost.

5.5. Design space

Efficient design space exploration and modeling of manycore environments could be desired from different research communities with focusing on different metrics, design aspects, applications and requirements. This is a major research challenge for manycore modeling and simulation. On one hand, a diversity of metrics comes from differentiated market segments and metric emphases, such as power, temperature, latency, or throughput. On the other hand, a diversity of variable multi-processor designs (even designs which are not existing yet) are of interest by users. A comprehensive design exploration model has to locate the optimal amounts of results for different metrics and workloads in a large and high-resolution design space. It should consider all design space parameters simultaneously, while enabling predictions for metrics of interest.

Full design space exploration (e.g. the case in current full system simulators such as SimFlex and GEM5) is constrained by the substantial costs of cycle-accurate simulators, which provide very detailed view into system modeling for a wide range of manycore micro-architectural configurations. We can expose specific trends or interested metrics in design space. However, due to issues as long simulation times, we are limited to constrain design scenarios and consider only small subsets of the full design space (as it is the case in current instrumentation-based simulators such as ZSim and McSimA+). Unfortunately, implying these limitations in design space may lead to results that may not be acceptable for the larger space (e.g. future peta-scale systems with high diversity of processors).

5.6. Simulation time

Designing a high-performance manycore system is extremely time-consuming. It involves exploring and analyzing a huge number of input parameters and configuration elements. Thus, it would be expectable that, in most cases, general simulation methods become infeasible or inefficient. The problem is that architectural simulation is very time consuming. In order to design a microprocessor with optimal characteristics, the simulator must explore and evaluate all possible configurations and find the optimal one. Note that finding the optimum set of configuration values can be different depending on the target and the design criteria. In fact we can say that the total simulation time is directly proportional to the number of configuration parameters which are required to be evaluated. The simulation time is also proportional to the size of the workload space, the average number of instructions in each application-input pair in the workload space, and the simulator slowdown factor [160].

During simulation, the performance evaluation of the system configuration is done through running a software program or a benchmark with a suitable set of inputs. In other words a collection of computer programs (workload) is used for the evaluation process. The simulator must provide capabilities to compose required workloads (the benchmarks with appropriate inputs) which are quite specific depending on the target operation domain of simulated system. Exploration of workload space along with design space has direct impact on the simulation time. Furthermore, the size of workload and extra features of the manycore processor architecture increase the timing cost of the simulation.

It is obvious, that to reduce the total simulation time, we need to reduce the size of any or all of the aforementioned factors (e.g. COTSon has reduced its workload space to only support MPI applications, ZSim has also reduced its design space by only supporting x86 as target ISA). But this might have greater costs in other aspects of simulator performance. Along this line, techniques such as selecting a region of interest through statistical simulation [161,162] (as used in Graphite, LiveSim [163] and BigHouse [164]), choosing a limited but representative set of program-input pairs [165–167], reduced input sets [168], trace sampling [169] (as used in TQSIM [170]), barrier interval time parallelism [171] and simulation optimization [172,173] (as these are partially used in BigSim and COTSon) might be taken into account.

5.7. Multi-model simulations

We must consider modeling manycore systems in different scenarios, as for some studies and analyses we need to have a simulator with capability of multi-dimensional modeling. It means that the simulator must provide an accurate modeling of the system and processor architecture while it offers functionalities to analyze the impact of interconnect networks, thermal and power consumption modeling. It is a challenging task, since we need to put several modeling mechanisms that communicate and work together.

5.8. Scalability

In order to model future peta-scale manycore systems and architectures, the current parallel software simulators are supposed to exploit the highest level of parallelism from the underlying parallel architectural host platforms. The obvious big challenge of such simulation is dealing with the rapid increasing number of cores and threads. We expect modeling and analyzing tools to keep a constant amount of overhead, while coping with the exponentially increasing number of the cores. In other words, overhead must be constant in time and independent from the system size. Handling and managing the intercommunication between threads/cores, resources and threads, which are involved in synchronization and contention points between threads and cores, in both application and architecture level is critical with respect to the synchronization bottlenecks. The problem of synchronization in manycore era will result in large scalability issues.

5.9. Productivity

Researchers and system designers expect simulators to be more productive and come with a set of requirements, features and characteristics such as ease of use, management tools, documentation and deploy-ability, visualization, deployment, debugging, etc. which could help them to perform modeling, designing and evaluating of the future systems in a better way. However, these requirements are changing, and sometimes are unclear. In addition it increases the cost of simulator construction.

6. Discussion and future directions

In order to produce peta-scale manycore simulators with capability to explore a wide range of micro-architectural design space of future parallel systems and architecture we have to use the capabilities of the current parallel hardware. This means that we must be able to fully exploit the current hardware architecture to build future parallel hardware and platforms. This direction, the only that will follow current hardware evolution, brings all of the challenges in application concurrency, many-core and hardware parallelism, to simulator designs, while they face some specific challenges which particularly belongs to the scope of distributed and parallel simulations. We also need to identify the most important capabilities of the current simulators.

The followings summarize the key information on the capabilities of the current architecture simulators (as already discussed in [Sections 3](#) and [4](#)):

- Instrumentation (DBT) based simulators provide faster and more scalable simulation capability compared to emulation based simulators, while the emulation based simulators provide better simulation accuracy to general architectures.
- The configurability (i.e. the ability to explore various ISAs and system architecture) of instrumentation based simulators is limited. Moreover, most of the instrumentation based simulators are user/application level simulators.
- Sampling based simulators are faster and more scalable compared to non-sampling based simulators with the cost of reducing simulation accuracy.
- Hardware-assisted simulators are faster than software-assisted simulators. However, they create complexity for development and they are not flexible for design space exploration. Furthermore, hardware-assisted simulators are more expensive than software-assisted simulators.
- Network/interconnect modeling is a requirement for distributed-memory simulation. Furthermore, distributed memory simulators are more scalable than shared memory simulators and they can provide faster simulation by enabling cluster based simulation.
- Parallel simulators are faster than sequential simulators. However, they generally suffer from drawbacks, such as complicated and inefficient scheduling of simulation segments and high communication overhead. Synchronization techniques improve the performance of the parallel simulators. But still efficient synchronization is a challenging issue for parallel simulation.
- Modularity is a very important feature for simulators, since it provides fast development, ease of use and flexibility for simulators. Accordingly, for most current simulators, efforts have been made to apply modularity through creating component based architectures which enables leveraging of existing tools and components.

We can conclude that, due to the large variety of simulator applications and methodologies trade-offs, it is not feasible to expect a comprehensive simulator to simultaneously satisfy the requirements of researchers in different communities. Therefore simulator designers must clarify the specific target usage of their simulation tools at the first step. Furthermore, due to the aforementioned challenges, we should make efforts towards modular simulation platforms that are able to leverage existing high throughput modules, methods and solutions to build the new tools. Here, base-simulation techniques such as synchronization, sampling, scheduling, etc. are required either to be improved or recreated.

In this paper, we have presented the characteristics and capabilities of the current simulations and modeling tools, aiming to simulate future peta-scale manycore systems and architectures. We have extracted and demonstrated a set of most significant problems and challenging issues which the simulator designers have to deal with. In order to cope with the aforementioned challenges, and to respond to the needs and concerns of different design space exploration and multi/many core architecture simulation requirements, we argue that, several techniques appear to be more promising to develop. Thus, we expect that the future direction of research in modeling and simulation of manycore systems will expand upon the following lines, relying on existing best of breed features currently incipient in different simulators:

- Regression and analytic models
- Statistical simulation
- Acceleration and FPGA-based Prototyping
- Modularity, integrability and aspect-oriented simulation
- Parallel simulation
- Cloud-based simulation
- Raising level of abstraction
- Model-driven simulation

6.1. Regression and analytic models

Regression and analytical modeling is a statistical tool for the investigation of relationships between design space metrics and variables and it is an efficient approach for accurately predicting different metrics and parameters in a large micro-architectural design space. Regularly, in manycore modeling we seek to clarify the causal effect of one metrics upon another metrics and parameters. To explore such issues, we can apply regression techniques to simulators, enabling them to efficiently obtain approximates of the design metrics. Simulators can collect data on the interested underlying metrics and employ regression analytic approaches to foresee the quantitative effect of the causal metrics upon the metric that they influence. For example, the work presented in [174] uses this approach for design space exploration by applying a class of models where, in each model, the response time is modeled as a weighted sum of foreseeing metrics plus a random noise (a noise is defined as the effect of the other metrics which are not considered in the prediction model).

6.2. Statistical simulation

Analyzing all the functional events and creating timing models for each component is costly due to the simulation time. Sampling and statistical simulation techniques are solutions to decrease the time of simulation. These approaches rely on collecting and pre-calculating certain architecture-dependent performance factors. The analysis techniques periodically sample and statistic on some characteristics of the running application in order to accelerate subsequent simulations. This solution proved to be highly accurate while having low overhead. However, the statistical simulation is constrained in the number of architectures that it can support. Furthermore, it is not clear whether these techniques can tolerate significant modifications in the number of cores or the interconnect network topology.

6.3. Acceleration and FPGA-based prototyping

Architecture research already has started to concentrate more on implementation and less on design of instruction sets. Along this line, FPGAs have been employed at different levels of abstraction in the design for the purpose of simulation, implementation and evaluation of current and future computer architectures. Architecture simulators can achieve benefits from increasing the speed of register transfer level simulation by performing logical simulation and evaluation based on FPGA-accelerated hardware platforms. As example of these range of simulators, we can mention Quick-turn [175] (an old product) and CadencePalladium [176] (a more recent one). Another example which integrated FPGA technologies into a tool for exploring and evaluating microarchitectural designs especially for newly proposed architectures is [177].

The simple idea behind FPGA-based simulators is using FPGAs to simulate Register Transfer Level (RTL) (i.e. a design abstraction which models a synchronous digital circuit), because a RTL simulation typically is very slow. In addition, we can map processors directly to FPGAs, which results to enable using higher clock rates. According to this concept the Research Accelerator for Multiple Processors (RAMP) [178] provides a set of research efforts such as RAMP Blue [179], RAMP Red and RAMP Gold [147]. RAMP Blue is a simulated machine, with 1008 cores, which are inter-operated by message passing and distributed memory. RAMP Red is another many-core simulator designed to investigate the issues of transactional memory (TM). Both of these projects used Xilinx FPGAs on Berkeley Emulation Engine 2 (BEE2) boards.

The most important problem of direct-RTL-mapping is that timing models is a property of the functional modeling. RAMP Blue tried to solve this problem by employing a separate simple timing model, which uses clock-gating of various components to get correct functional behavior while faking the desired timing model. This technique implies a tight coupling between the timing and functional simulations because the timing model is a wrapper around each functional component. Direct RTL-mapped systems do not support modularity and they are also suffering from timing wrappers. This leads FPGA-based simulators to completely split the timing and functional modeling of the target system.

6.4. Modularity, integrability and aspect-oriented simulation

There are islands of simulation and modeling tools with various focus on different domains, which are completely isolated from each other. Modularity is an important requirement for simulation tools in order to achieve integrability and reusability. Modularity can be obtained in different levels. Implementation-level (object-level or component-level) modularity is a common approach to design modular simulators, where a simulator composed of a set of interrelated components/objects working together toward a common objective. This even makes simulators more flexible to interoperate with other external simulators or simulator components. In fact, a powerful and comprehensive solution has to benefit from existing tools and individual investments as much as possible. In addition, a proper simulation approach must provide an infrastructure to leverage mature tools via standardized APIs for common simulator services and functionalities (such as time modeling, synchronization, sampling, event modeling, interconnects simulation, bandwidth and latency modeling, thermal and energy consumption modeling). It could support a very large design space exploration in various aspects and objectives. Furthermore, it catalyzes the development process of the new simulator platforms while the resulted tools would be substantially efficient. There are several examples of successful (implementation-level) modular integrations in the current literature, including Qemu-COTSon [30], NOXIM-COTSon (NOXIM is a network on chip simulator), and SystemC-Qemu [131,133–136].

Modularity can also be applied in a more conceptual level by advocating aspect-oriented paradigms into the area of manycore simulation. This type of modularity can be described as model-level (or aspect-level) modularity, where simulation of a target (manycore) system can be performed considering different viewpoints (aspects or simulation aspects). Each viewpoint can be defined as a representation (specification) of a target system (simulated system) in a certain view or aspect (i.e. simulation aspect). In fact, a viewpoint (or an aspect) specifies the simulation focus (e.g. memory and cache, processing cores, network on chip, power and energy, OS and application execution). For example, in NoC simulators, it might be necessary to provide more detailed interconnection while describing cores in more abstract level, since the focus of such simulators is on interconnection networks.

In this context, modularity refers to separating the simulation of different concerns of a target system and decomposing the entire simulation into a set of simulation aspects. Each simulation aspect can simulate the system under study from a particular point of view. The simulation aspects can be recomposed through a weaving mechanism (e.g. [180]), providing facility to simulate a target system from all aspects or only certain aspects of interest. These weaved modular aspects can flexibly communicate to each other in order to provide a detailed model of a target system with respect to desired aspects. This can provide simulation efficiency specially when simulation is required for systematic optimization. In such a case, the optimizer aims to optimize the system under study from a certain point of view, therefore it needs to evaluate the system from a particular view by simulating the certain aspects of the system. There are some works that already introduced aspect-oriented simulators [181–184] in the field of discrete event simulation. However, none of them are specifically designed for the purpose of manycore architecture simulation.

6.5. Parallel simulation

The traditional sequential simulators already began to be converted to parallel simulators. However, there are still major challenges that adversely affect the performance of parallel simulations. These issues include segmentation of simulation workloads, dynamic scheduling, communications between simulator instances, time management and synchronization. Among them, synchronization is a key issue, since lacking an adequate synchronization may result in a time causality error (i.e. violation of the time-stamp order), imposing overhead by periodical synchronization in order to achieve consistency among logical clocks of different host nodes [185]. For example, a high speed node might receive a straggler event with an overdated time-stamp from a low speed node (an straggler event is an event which its time-stamp is less than the local clock value). A potential direction to cope with the synchronization problem might be to use multilevel lax(relax)-based synchronization approaches, as it is getting common for recent parallel (manycore) simulators [185,186] due to their capability to significantly reducing the synchronization frequency and in turn, reducing synchronization overhead and increasing parallelism. These strategies work through efficiently adapting a lax-based technique (e.g. by lengthening of synchronous periods in occurrence of straggler events [11]) in one level, while correcting the simulation accuracy in another level (note that traditional lax-based synchronization suffer from the lack of simulation accuracy).

Overall, regardless of the aforementioned parallelism issues, we propose two future directions for simulation-level parallelism along the thematics of this paper: component-level (object-level) and aspect-level (model-level) parallelism. In component-level parallelism, a simulator contains a set of stand-alone components (a component may include sub-components/objects) and a target system can be simulated by instantiation of components required and definition of inter-component processes (e.g. communication between component instances). Component-level parallelism [187,188] can be achieved by parallel execution of component instances on different host nodes (processors). In aspect-level parallelism, the simulation of a target system includes a set of stand-alone aspects where each aspect provides a detailed representation of a system from the prospective of a certain aspect. More complete (global) simulation of a target system can be extracted through merging a set of aspects desired.

6.6. Cloud-based simulation

The Cloud service models such as Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS) facilitate doing research in all the related fields of distributed/parallel computing. Using Cloud-based systems offer the end users a completely remote, self-organized, load balanced and distributed environment as a single service/single system which is more efficient and easy to work with. Definitely, the Cloud-based software services are able to support the whole range of softwares as services even simulator softwares. Currently there exist some research works that provide system simulation service for Cloud as part of SaaS [189]. However, all the software (such as computation centric and real time applications) and particularly simulators are not the same and they have different operational behaviors and characteristics, so we might consider that using a generic SaaS would not be enough to satisfy all the requirements of the future Peta-scale system simulators. Deploying simulators using SaaS [190] will arise some complexities in different aspects such as configuration, synchronization, workload distribution, etc. In other words it doesn't provide the performance and ease of use as the simulator users are expecting.

A possible approach to improve the performance of the Cloud-based simulators [191–193] is developing the Cloud service model as Simulation-as-a-Service(SiaS), which means that SaaS have to be adapted and customized precisely to run simulators. So by using such a simulation service, the simulator runs on the Cloud servers while leaving the user's local resources free. A large scale manycore simulator explores a very large model/experiment space where the behavior of each model is




Fig. 9. Interval simulation according to miss events (extracted from [194]).

explored during its creation. Furthermore, in runtime each experiment model in the design space will find emergent properties which cannot be deduced from the model and they can only be observed during simulation. Simulation in large scale involves a high number of parameters, values and settings and we can say that it is a resource intensive application which most of its behaviors can be only identified during runtime. Such a Cloud-based simulation service provides a service based on the amount of simulated time for the end user experiment. It takes care of the distribution of the experiment work load on several machines in the Cloud, and finally it will collect all the results from the worker machines.

6.7. Raising level of abstraction

Generally, architecture simulators suffer from two major timing limitations, which are simulation time and development time of simulators. Full-cycle-accurate simulators are examples of this case. They are accurate and fulfill the user requirements for correctness of the experiments, but they are highly time consuming and result from a complex development process which takes time. But this detailed modeling is not necessary for many kinds of design space exploration and just makes the development of new models more difficult. For example, we do not usually need to model the details of cache coherence protocols or many-core interconnection networks when investigating trade-offs in the memory system hierarchy.

Raising the abstraction level of the simulation will help simulators to be faster and easier to use, while remaining accurate. The problem is finding out what level of abstraction is appropriate and how to deal with the tight performance among the co-executing threads and the micro-architectures. Interval simulation [194] is a possible solution along this line. It raises the level of abstraction in the core-level compared with the typical detailed simulation. It uses a mechanistic analytical model [12,194] which drives the timing models of each individual cores with sampling of some set of the instructions through the cores' pipeline stages. Each interval is defined as the distance between two miss events (branch mis-predictions, cache L1 I-cache miss, long-latency load miss or TLB misses) (see Fig. 9) through the cores' pipeline stage which identifies a part of instructions stream. The miss events can be determined by employing branch predictor, memory hierarchy, cache coherence and interconnection network simulators. Thus, the mechanistic analytical model drives the timing models for each intervals instead of the whole instruction set. The integration between miss event simulators and the analytical models facilitate the modeling of the tight performance entanglement between co-executing threads on manycore architectural processors. This approach is a promising technique particularly for system level exploration at the early design stage. Sampling the simulated instruction stream, using host multi-threading and mapping the simulations workloads on FPGAs or GPU are the other alternative approaches that achieve considerable simulation speedups while maintain the performance in a cycle-accurate manner.

6.8. Model-driven simulation

MDE based simulation (simulation based on Model Driven Engineering) can be used as a powerful alternative for many-core simulations in future instead of instrumentation-based simulation. Model-Driven Engineering (MDE) is a software development methodology that evolved as a paradigm shift from the object-oriented paradigm (everything is an object), into the model engineering paradigm (everything is a model) [195]. Currently, this method is used mostly in the domain of software engineering, programming languages and domain specific languages, but not in manycore simulation. Applying MDE concepts in the area of simulation, may lead to a new generation of manycore simulators which can efficiently cope with multiple challenging issues, including interoperability (e.g. ability to simulate various target systems on a fixed host system), multi-model simulation, simulation speed, design space exploration, etc.

In MDE, everything is defined as model and a model basically is a description of a real system (like a manycore system including different components such as memory, processing cores, network and interconnect). In fact, a (real) system is an instantiation of a model. Each model, in turn, is defined by another model, so-called metamodel (i.e. the abstract syntax of a modeling language is specified using another model, describing the syntactic elements and the relationships existing between those elements). The Object Management Group (OMG) proposed a four-level metamodeling framework as a standard to develop modeling languages. In each level, except the bottom level (M0), there is a model that specifies a set of other models at the lower level in a recursive way. The M0 is the bottom level of this hierarchy, specifying various real systems (e.g. a memory component, a processor, different manycore systems). At the higher level (M1), models (e.g. UML class diagrams) represent (abstract) these systems. Each model conforms to its metamodel defined at the upper level (M2). And similarly, metamodels in level M2 conform to another models, defined at the highest level (M3), so-called meta-metamodels.

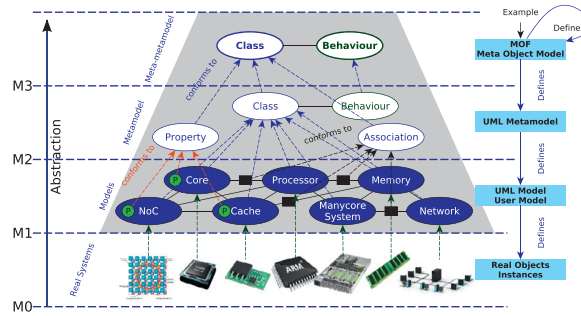


Fig. 10. An example of both structural and behavioral metamodel hierarchy for model-driven manycore simulation (note that “Behavior” entities are specific for behavioral metamodeling).

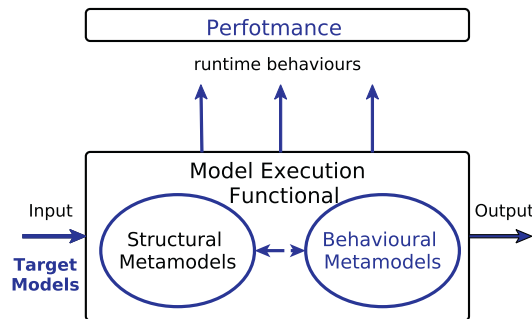


Fig. 11. Extracting behavioral properties of a target model (target system) during run-time.

Meta-metamodels (e.g. OMG’s Meta Object Facility (MOF) [196]) are self-descriptive entities, enabling to conform to themselves. MDE is able to raise the level of abstraction in system description/specification (by using models at the different levels of abstraction) and increase automation in running (execution or simulation of) a system (by using code generation and model transformations mechanism). In other words, the MDE approach promotes the use of models as first-class entities that need to be constructed, maintained, executed, and mapped into other models or artifacts by model transformations.

Overall, MDE covers aspects such as architecture design, code generation, model transformation and model checking. Model checking includes techniques to check and ensure the quality (performance) of the models (e.g. model validation using behavioral properties). Referring to the above-mentioned metamodel hierarchy, a model-driven manycore simulator can flexibly explore the design space through specification of unlimited target systems as different models in different levels of metamodel hierarchy. Since the descriptions of target models in metamodel hierarchy is very abstract, the simulator can swiftly check the functional accuracy of different designed target models, before going to the detailed simulation of a specific target system. This can be done through using model checking/validation mechanism provided by MDE approach. Furthermore, using model transformation languages [197–201], a model-driven simulator can provide interoperability by means of capability to convert any desired target models to codes which can be directly executed on the host machine. This means that a MDE-based simulator can flexibly address platform complexities through simulation of any desired target system on a fixed host machine (similar to retargetable DBT, discussed in Section 3.1).

As we discussed above, current Model-Driven Architectures (MDAs) provide capability for structural metamodeling which can efficiently be used for functional modeling in manycore simulators. A recent work [202] in this area introduces the behavioral metamodeling to complement the structural metamodeling (see Fig. 10). This behavioral metamodeling (including various behaviors described in different levels of metamodel hierarchy) facilitates to extract behavioral properties of a target model (target system) during run-time. This allows performance modeling while functional modeling is performed (see Fig. 11). Model-driven simulation can be used as an alternative for instrumentation-based simulation which extracts the performance behavior of the target system by injecting code into the simulated binary, running on the host machine.

6.9. Summary

Overall, the areas open to innovation are manifold. We do expect next generation manycore simulators to rely on several of the aforementioned techniques to reach the required performance and complexity management of future simulations. Global technology development, and specially the specific development we will see in manycore systems, and its associated market, will determine the relevance of each of these simulation techniques in the future manycore simulations that will handle peta-scale computing systems.

Acknowledgment

The authors acknowledge the support of project FP7-ICT-2009.8.1, Grant Agreement No. 248465, Service-oriented Operating Systems (S[o]OS, 2010–2013) [203–210] and of project Cloud Thinking (2013–2015), CENTRO-07-ST24-FEDER-002031 [211]. We advocate for the research community to develop a collective effort for a community integrated S[o]OS environment, along the lines here discussed.

References

- [1] T. Imamura, S. Yamada, M. Machida, Development of a high performance eigensolver on the petascale next generation supercomputer system, in: *Proceedings of Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2010 (SNA+ MC2010)*, 2010.
- [2] C.-q. Yang, Q. Wu, T. Tang, F. Wang, J.-I. Xue, Programming for scientific computing on peta-scale heterogeneous parallel systems, *J. Cent. South Univ.* 20 (5) (2013) 1189–1203, doi:10.1007/s11771-013-1602-z.
- [3] M. Pellauer, M. Vijayaraghavan, M. Adler, Arvind, J. Emer, Quick performance models quickly: closely-coupled partitioned simulation on fpgas, in: *Proceedings of IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS 2008*, 2008, pp. 1–10, doi:10.1109/ISPASS.2008.4510733.
- [4] N. Binkert, B. Beckmann, G. Black, S.K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D.R. Hower, T. Krishna, S. Sardashti, The Gem5 Simulator 39(2011) 1–7, doi:10.1145/2024716.2024718.
- [5] N. Hardavellas, S. Somogyi, T.F. Wenisch, R.E. Wunderlich, S. Chen, J. Kim, B. Falsafi, J.C. Hoe, A.G. Nowatzky, Simflex: a fast, accurate, flexible full-system simulation framework for performance evaluation of server architecture, *ACM SIGMETRICS Perform. Eval. Rev.* 31 (4) (2004) 31–34.
- [6] A. Patel, F. Afram, S. Chen, K. Ghose, Marss: a full system simulator for multicore x86 cpus, in: *Proceedings of the 48th Design Automation Conference*, ACM, 2011, pp. 1050–1055.
- [7] J. Lu, H. Jia, A. Arias, X. Gong, Z.J. Shen, On-chip bondwire transformers for power soc applications, in: *Proceedings of the Twenty-Third Annual IEEE Applied Power Electronics Conference and Exposition, APEC 2008*, IEEE, 2008, pp. 199–204.
- [8] D. Bortolotti, C. Pinto, A. Marongiu, M. Ruggiero, L. Benini, Virtualsoc: a full-system simulation environment for massively parallel heterogeneous system-on-chip, in: *Proceedings of the 2013 IEEE 27th International Parallel and Distributed Processing Symposium Workshops & Ph.D. Forum (IPDPSW)*, IEEE, 2013, pp. 2182–2187.
- [9] T. Austin, E. Larson, D. Ernst, SimpleScalar: an infrastructure for computer system modeling, *Computer* 35 (2) (2002) 59–67, doi:10.1109/2.982917.
- [10] G. Zheng, G. Kakulapati, L. Kale, Bigsim: a parallel simulator for performance prediction of extremely large parallel machines, in: *Proceedings of the 18th International Parallel and Distributed Processing Symposium*, 2004, 2004, p. 78, doi:10.1109/IPDPS.2004.1303013.
- [11] J. Miller, H. Kasture, G. Kurian, C. Gruenwald, N. Beckmann, C. Celio, J. Eastep, A. Agarwal, Graphite: a distributed parallel simulator for multicores, in: *Proceedings of the 2010 IEEE 16th International Symposium on High Performance Computer Architecture (HPCA)*, 2010, pp. 1–12, doi:10.1109/HPCA.2010.5416635.
- [12] T.E. Carlson, W. Heirman, L. Eeckhout, Sniper: exploring the level of abstraction for scalable and accurate parallel multi-core simulation, in: *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM, 2011, p. 52.
- [13] D. Sanchez, C. Kozyrakis, Zsim: fast and accurate microarchitectural simulation of thousand-core systems, in: *Proceedings of the ACM SIGARCH Computer Architecture News*, 41, ACM, 2013, pp. 475–486.
- [14] Y. Fu, D. Wentzclaff, Prime: A parallel and distributed simulator for thousand-core chips, in: *Proceedings of the 2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2014, pp. 116–125, doi:10.1109/ISPASS.2014.6844467.
- [15] C.J. Mauers, M.D. Hill, D.A. Wood, Full-system timing-first simulation, *SIGMETRICS Perform. Eval. Rev.* 30 (1) (2002) 108–116, doi:10.1145/511399.511349.
- [16] D. Penry, A single-specification principle for functional-to-timing simulator interface design, in: *Proceedings of the 2011 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2011, pp. 186–196, doi:10.1109/ISPASS.2011.5762735.
- [17] D. Chiou, D. Sunwoo, J. Kim, N. Patil, W.H. Reinhart, D.E. Johnson, Z. Xu, The fast methodology for high-speed soc/computer simulation, in: *Proceedings of the 2007 IEEE/ACM International Conference on Computer-aided Design, ICCAD '07*, IEEE Press, Piscataway, NJ, USA, 2007, pp. 295–302.
- [18] E. Schnarr, J.R. Larus, Fast out-of-order processor simulation using memoization, *SIGPLAN Not.* 33 (11) (1998) 283–294, doi:10.1145/291006.291063.
- [19] G. Loh, S. Subramaniam, Y. Xie, Zesto: a cycle-level simulator for highly detailed microarchitecture exploration, in: *Proceedings of IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS 2009*, 2009, pp. 53–64, doi:10.1109/ISPASS.2009.4919638.
- [20] S. Fytraki, D. Pnevmatikatos, Resim, a trace-driven, reconfigurable ILP processor simulator, in: *Proceedings of Design, Automation Test in Europe Conference Exhibition, DATE '09*, 2009, pp. 536–541, doi:10.1109/DATE.2009.5090722.
- [21] Z. Fang, Q. Min, K. Zhou, Y. Lu, Y. Hu, W. Zhang, H. Chen, J. Li, B. Zang, Transformer: a functional-driven cycle-accurate multicore simulator, in: *Proceedings of the 49th Annual Design Automation Conference, DAC '12*, ACM, New York, NY, USA, 2012, pp. 106–114, doi:10.1145/2228360.2228381.
- [22] J. Donald, M. Martonosi, An efficient, practical parallelization methodology for multicore architecture simulation, *IEEE Comput. Archit. Lett.* 5 (2) (2006) 14–14, doi:10.1109/L-CA.2006.14.
- [23] E. Argollo, A. Falcón, P. Faraboschi, M. Monchiero, D. Ortega, Cotsn: infrastructure for full system simulation, *SIGOPS Oper. Syst. Rev.* 43 (1) (2009) 52–61, doi:10.1145/1496909.1496921.
- [24] R. Giorgi, Teraflux: exploiting dataflow parallelism in teradevices, in: *Proceedings of the 9th Conference on Computing Frontiers, CF '12*, ACM, New York, NY, USA, 2012, pp. 303–304, doi:10.1145/2212908.2212959.
- [25] N. Neelakantam, C. Blundell, J. Devietti, M.M. Martin, C. Zilles, Fes2: a full-system execution-driven simulator for x86, *Proceedings of the Poster Presented at ASPLOS 2008(2008)* 6.
- [26] M.M.K. Martin, D.J. Sorin, B.M. Beckmann, M.R. Marty, M. Xu, A.R. Alameldeen, K.E. Moore, M.D. Hill, D.A. Wood, Multifacet's general execution-driven multiprocessor simulator (gems) toolset, *SIGARCH Comput. Archit. News* 33 (4) (2005) 92–99, doi:10.1145/1105734.1105747.
- [27] R.C. Bedichek, Talisman: fast and accurate multicomputer simulation, *SIGMETRICS Perform. Eval. Rev.* 23 (1) (1995) 14–24, doi:10.1145/223586.223589.
- [28] P.M. Ortego, P. Sack, Sesc: Superescalar simulator, in: *Proceedings of the 17th Euro Micro Conference on Real Time Systems (ECRTS'05)*, 2004, pp. 1–4.
- [29] N. Binkert, R. Dreslinski, L. Hsu, K. Lim, A. Saidi, S. Reinhardt, The m5 simulator: modeling networked systems, *IEEE Micro* 26 (4) (2006) 52–60, doi:10.1109/MM.2006.82.
- [30] F. Bellard, Qemu, a fast and portable dynamic translator, in: *Proceedings of the Annual Conference on USENIX Annual Technical Conference, ATEC '05*, USENIX Association, Berkeley, CA, USA, 2005 41–41.
- [31] E. Witchel, M. Rosenblum, Embra: fast and flexible machine simulation, *SIGMETRICS Perform. Eval. Rev.* 24 (1) (1996) 68–79, doi:10.1145/233008.233025.
- [32] P. Bohrer, J. Peterson, M. Elnozahy, R. Rajamony, A. Gheith, R. Rockhold, C. Lefurgy, H. Shafi, T. Nakra, R. Simpson, E. Speight, K. Sudeep, E. Van Hensbergen, L. Zhang, Mambo: a full system simulator for the powerpc architecture, *SIGMETRICS Perform. Eval. Rev.* 31 (4) (2004) 8–12, doi:10.1145/1054907.1054910.
- [33] R. Bedichek, Simnow: fast platform simulation purely in software, 16, *Hot Chips: A Symposium on High Performance Chips Conference*, 2004.
- [34] P. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hallberg, J. Hogberg, F. Larsson, A. Moestedt, B. Werner, Simics: a full system simulation platform, *Computer* 35 (2) (2002) 50–58, doi:10.1109/2.982916.

- [35] M. Rosenblum, E. Bugnion, S. Devine, S.A. Herrod, Using the simos machine simulator to study complex computer systems, *ACM Trans. Model. Comput. Simul.* 7 (1) (1997) 78–103, doi:10.1145/244804.244807.
- [36] C. Hughes, V. Pai, P. Ranganathan, S. Adve, Rsim: simulating shared-memory multiprocessors with ilp processors, *Computer* 35 (2) (2002) 40–49, doi:10.1109/2.982915.
- [37] E.A. Brewer, C.N. Dellarocas, A. Colbrook, W.E. Weihl, Proteus: a high-performance parallel-architecture simulator, *SIGMETRICS Perform. Eval. Rev.* 20 (1) (1992) 247–248, doi:10.1145/149439.133146.
- [38] J.-H. Ding, P.-C. Chang, W.-C. Hsu, Y.-C. Chung, Pqemu: a parallel system emulator based on qemu, in: *Proceedings of the 2011 IEEE 17th International Conference on Parallel and Distributed Systems (ICPADS)*, 2011, pp. 276–283, doi:10.1109/ICPADS.2011.102.
- [39] Z. Wang, R. Liu, Y. Chen, X. Wu, H. Chen, W. Zhang, B. Zang, Coremu: a scalable and portable parallel full-system emulator, *SIGPLAN Not.* 46 (8) (2011) 213–222, doi:10.1145/2038037.1941583.
- [40] K. Wang, Y. Zhang, H. Wang, X. Shen, Parallelization of ibm mambo system simulator in functional modes, *SIGOPS Oper. Syst. Rev.* 42 (1) (2008) 71–76, doi:10.1145/1341312.1341325.
- [41] R. Lantz, Fast functional simulation with parallel embra, in: *Proceedings of the 4th Annual Workshop on Modeling, Benchmarking and Simulation*, Citeseer, 2008.
- [42] S. Kanaujia, I.E. Papazian, J. Chamberlain, J. Baxter, Fastmp: a multi-core simulation methodology, in: *Proceedings of the Workshop on Modeling, Benchmarking and Simulation (MoBS 2006)*, 2006.
- [43] J. Chen, M. Annavaram, M. Dubois, Slacksim: a platform for parallel simulations of cmps on cmps, *SIGARCH Comput. Archit. News* 37 (2) (2009) 20–29, doi:10.1145/1577129.1577134.
- [44] S.S. Mukherjee, S.K. Reinhardt, B. Falsafi, M. Litzkow, S. Huss-Lederman, M.D. Hill, J.R. Larus, D.A. Wood, Fast and portable parallel architecture simulators: Wisconsin wind tunnel ii, *IEEE Concurr.* 8 (4) (2000) 12–20.
- [45] S.S. Mukherjee, S.K. Reinhardt, B. Falsafi, M. Litzkow, S. Huss-Lederman, M.D. Hill, J.R. Larus, D.A. Wood, Wisconsin wind tunnel II: a fast and portable parallel architecture simulator, in: *Proceedings of the Workshop on Performance Analysis and its Impact on Design*, ACM, Denver, Co, 1997.
- [46] A. Jaleel, R.S. Cohn, C.-K. Luk, B. Jacob, Cmp \$ im: a pin-based on-the-fly multi-core cache simulator, in: *Proceedings of the Fourth Annual Workshop on Modeling, Benchmarking and Simulation (MoBS)*, Co-located with ISCA, 2008, pp. 28–36.
- [47] T. Wenisch, R. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, J. Hoe, Simflex: statistical sampling of computer system simulation, *IEEE Micro* 26 (4) (2006) 18–31, doi:10.1109/MM.2006.79.
- [48] W. Heirman, T. Carlson, L. Eeckhout, Sniper: scalable and accurate parallel multi-core simulation, in: *Proceedings of the 8th International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems (ACACES-2012)*, High-Performance and Embedded Architecture and Compilation Network of Excellence (HiPEAC), 2012, pp. 91–94.
- [49] G. D'Angelo, M. Marzolla, New trends in parallel and distributed simulation: from many-cores to cloud computing, *Simul. Model. Pract. Theory* 49 (2014) 320–335. <http://dx.doi.org/10.1016/j.simpat.2014.06.007>.
- [50] R.M. Fujimoto, Parallel discrete event simulation, *Commun. ACM* 33 (10) (1990) 30–53.
- [51] P. Ren, M. Lis, M.H. Cho, K.S. Shim, C. Fletcher, O. Khan, N. Zheng, S. Devadas, Hornet: a cycle-level multicore simulator, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* 31 (6) (2012) 890–903, doi:10.1109/TCAD.2012.2184760.
- [52] J. Xu, M. Chen, G. Zheng, Z. Cao, H. Lv, N. Sun, Simk: a parallel simulation engine towards shared-memory multiprocessor, *J. Comput. Sci. Technol.* 24 (6) (2009) 1048–1060.
- [53] R.M. Fujimoto, Parallel discrete event simulation, *Commun. ACM* 33 (10) (1990) 30–53, doi:10.1145/84537.84545.
- [54] R. Curry, C. Kiddle, R. Simmonds, B. Unger, Sequential performance of asynchronous conservative pdes algorithms, in: *Proceedings of Workshop on Principles of Advanced and Distributed Simulation*, PADS 2005, 2005, pp. 217–226, doi:10.1109/PADS.2005.25.
- [55] R.M. Fujimoto, *Parallel and Distribution Simulation Systems*, 1st ed., John Wiley and Sons, Inc., New York, NY, USA, 1999.
- [56] D.R. Jefferson, Virtual time, *ACM Trans. Program. Lang. Syst.* 7 (3) (1985) 404–425, doi:10.1145/3916.3988.
- [57] K. Chandy, J. Misra, Distributed simulation: a case study in design and verification of distributed programs, *IEEE Trans. Softw. Eng.* SE-5 (5) (1979) 440–452, doi:10.1109/TSE.1979.230182.
- [58] J. Wang, Z. Dong, S. Yalamanchili, G. Riley, Optimizing parallel simulation of multicore systems using domain-specific knowledge, in: *Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, ACM, 2013, pp. 127–136.
- [59] J. Wang, J. Beu, R. Bheda, T. Conte, Z. Dong, C. Kersey, M. Rasquinha, G. Riley, W. Song, H. Xiao, et al., Manifold: a parallel simulation framework for multicore systems, in: *Proceedings of the 2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, IEEE, 2014, pp. 106–115.
- [60] J. Chen, L. Kumar Dabburu, D. Wong, M. Annavaram, M. Dubois, Adaptive and speculative slack simulations of cmps on cmps, in: *Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2010, pp. 523–534, doi:10.1109/MICRO.2010.47.
- [61] Y. Mongkolsin, W. Maruringsith, P-hase: an efficient synchronous pdes tool for creating scalable simulations, in: T. Xiao, L. Zhang, M. Fei (Eds.), *Proceedings of the AsiaSim 2012*, Communications in Computer and Information Science, Springer, Berlin, Heidelberg, 2012, pp. 231–245, doi:10.1007/978-3-642-34387-2_27.
- [62] R. Rajan, P. Wilsey, Dynamically switching between lazy and aggressive cancellation in a time warp parallel simulator, in: *Proceedings of the 28th Annual Simulation Symposium*, 1995, 1995, pp. 22–30, doi:10.1109/SIMSYM.1995.393597.
- [63] A.C. Palaniswamy, S. Aji, P.A. Wilsey, An efficient implementation of lazy reevaluation, in: *Proceedings of the 25th Annual Symposium on Simulation*, IEEE Computer Society Press, 1992, pp. 140–146.
- [64] R. Noronha, N.B. Abu-Ghazaleh, Early cancellation: An active nic optimization for time-warp, in: *Proceedings of the Sixteenth Workshop on Parallel and Distributed Simulation*, PADS '02, IEEE Computer Society, Washington, DC, USA, 2002, pp. 43–50.
- [65] R. Pavel, *Simulation Methodology and Tools for the Development of Novel Program Execution Models and Architectures*, University of Delaware, 2015 Ph.D. thesis.
- [66] D. Ball, S. Hoyt, The adaptive time-warp concurrency control algorithm, in: *Proceedings of the SCS Multiconference on Distributed Simulation*, 22, 1990, pp. 174–177.
- [67] A. Malik, A. Park, R. Fujimoto, Optimistic synchronization of parallel simulations in cloud computing environments, in: *Proceedings of the IEEE International Conference on Cloud Computing*, CLOUD '09, 2009, pp. 49–56, doi:10.1109/CLOUD.2009.79.
- [68] G. Malhotra, P. Aggarwal, A. Sagar, S. Sarangi, Partejas: a parallel simulator for multicore processors, in: *Proceedings of the 2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2014, pp. 130–131, doi:10.1109/ISPASS.2014.6844470.
- [69] J.H. Ahn, S. Li, O. Seongil, N. Jouppi, Mcsima+: a manycore simulator with application-level+ simulation and detailed microarchitecture modeling, in: *Proceedings of the 2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2013, pp. 74–85, doi:10.1109/ISPASS.2013.6557148.
- [70] S.K. Reinhardt, M.D. Hill, J.R. Larus, A.R. Lebeck, J.C. Lewis, D.A. Wood, The Wisconsin wind tunnel: virtual prototyping of parallel computers, *SIGMETRICS Perform. Eval. Rev.* 21 (1) (1993) 48–60, doi:10.1145/166962.166979.
- [71] C.-K. Luk, R. Cohn, R. Muth, H. Patil, A. Klauser, G. Lowney, S. Wallace, V.J. Reddi, K. Hazelwood, Pin: building customized program analysis tools with dynamic instrumentation, *SIGPLAN Not.* 40 (6) (2005) 190–200, doi:10.1145/1064978.1065034.
- [72] C. Wang, S. Hu, H.-s. Kim, S.R. Nair, M. Breternitz, Z. Ying, Y. Wu, Starbdt: an efficient multi-platform dynamic binary translation system, in: *Proceedings of the 12th Asia-Pacific Conference on Advances in Computer Systems Architecture*, ACSAC'07, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 4–15.
- [73] D. Bruening, T. Garnett, S. Amarasinghe, An infrastructure for adaptive dynamic optimization, in: *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*, CGO '03, IEEE Computer Society, Washington, DC, USA, 2003, pp. 265–275.

- [74] L. Baraz, T. Devor, O. Etzion, S. Goldenberg, A. Skaletsky, Y. Wang, Y. Zemach, Ia-32 execution layer: a two-phase dynamic translator designed to support ia-32 applications on itanium®-based systems, in: Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 36, IEEE Computer Society, Washington, DC, USA, 2003, p. 191.
- [75] A. Chernoff, M. Herdeg, R. Hookway, C. Reeve, N. Rubin, T. Tye, S.B. Yadavalli, J. Yates, Fx!32: a profile-directed binary translator, IEEE Micro 18 (2) (1998) 56–64, doi:10.1109/40.671403.
- [76] R. Avizienis, Y. Lee, A. Waterman, Ramp gold: a high-throughput fpga-based manycore simulator, in: Proceedings of the 4th Workshop on Architectural Research Prototyping, Austin, TX, 2009.
- [77] T. Grotker, System Design with SystemC, Kluwer Academic Publishers, Norwell, MA, USA, 2002.
- [78] S. Raghav, A. Marongiu, C. Pinto, D. Atienza, M. Ruggiero, L. Benini, Full system simulation of many-core heterogeneous socs using gpu and qemu semihosting, in: Proceedings of the 5th Annual Workshop on General Purpose Processing with Graphics Processing Units, GPGPU-5, ACM, New York, NY, USA, 2012, pp. 101–109, doi:10.1145/2159430.2159442.
- [79] S. Raghav, M. Ruggiero, A. Marongiu, C. Pinto, D. Atienza, L. Benini, Gpu acceleration for simulating massively parallel many-core platforms, IEEE Trans. Parallel Distrib. Syst. 26 (5) (2015) 1336–1349, doi:10.1109/TPDS.2014.2319092.
- [80] E. Chung, E. Nurvitadhi, J. Hoe, B. Falsafi, K. Mai, Protoflex: Fpga-accelerated hybrid functional simulator, in: Proceedings of the IEEE International Parallel and Distributed Processing Symposium, IPDPS 2007, 2007, pp. 1–6, doi:10.1109/IPDPS.2007.370516.
- [81] M. Pellauer, M. Adler, M. Kinsky, A. Parashar, J. Emer, Hasim: Fpga-based high-detail multicore simulation using time-division multiplexing, in: Proceedings of the 2011 IEEE 17th International Symposium on High Performance Computer Architecture (HPCA), 2011, pp. 406–417, doi:10.1109/HPCA.2011.5749747.
- [82] T. Sherwood, E. Perelman, G. Hamerly, B. Calder, Automatically characterizing large scale program behavior, SIGPLAN Not. 37 (10) (2002) 45–57, doi:10.1145/605432.605403.
- [83] R.E. Wunderlich, T.F. Wenisch, B. Falsafi, J.C. Hoe, Smarts: accelerating microarchitecture simulation via rigorous statistical sampling, SIGARCH Comput. Archit. News 31 (2) (2003) 84–97, doi:10.1145/871656.859629.
- [84] A. Falcón, P. Faraboschi, D. Ortega, Combining simulation and virtualization through dynamic sampling, in: Proceedings of the IEEE International Symposium on Performance Analysis of Systems & Software, ISPASS 2007, IEEE, 2007, pp. 72–83.
- [85] T.F. Wenisch, R.E. Wunderlich, B. Falsafi, J.C. Hoe, Statistical sampling of microarchitecture simulation, in: Proceedings 20th IEEE International Parallel Distributed Processing Symposium, 2006, p. 8, doi:10.1109/IPDPS.2006.1639584.
- [86] J.J. Yi, L. Eeckhout, D.J. Lilja, B. Calder, L.K. John, J.E. Smith, The future of simulation: a field of dreams, Computer 39 (11) (2006) 22–29.
- [87] J.J. Yi, S.V. Kodakara, R. Sendag, D.J. Lilja, D.M. Hawkins, Characterizing and comparing prevailing simulation techniques, in: Proceedings of the 11th International Symposium on High-Performance Computer Architecture, HPCA-11, IEEE, 2005, pp. 266–277.
- [88] J.S. Robin, C.E. Irvine, Analysis of the intel pentium's ability to support a secure virtual machine monitor, in: Proceedings of the 9th Conference on USENIX Security Symposium, SSYM'00, vol. 9, USENIX Association, Berkeley, CA, USA, 2000 10–10.
- [89] D. Lugones, E. Luque, D. Franco, J.C. Moura, D. Rexachs, P. Faraboschi, D. Ortega, G. Gimenez, A. Falcon, Initial studies of networking simulation on cotson, 2009, Accessed: 2013-07-10.
- [90] A. Falcon, P. Faraboschi, D. Ortega, An adaptive synchronization technique for parallel simulation of networked clusters, in: Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS 2008, IEEE, 2008, pp. 22–31.
- [91] M.P. Barcellos, R.S. Antunes, H.H. Muhammad, R.S. Munaretti, Beyond network simulators: fostering novel distributed applications and protocols through extendible design, J. Netw. Comput. Appl. 35 (1) (2012) 328–339. <http://dx.doi.org/10.1016/j.jnca.2011.07.015>. Collaborative Computing and Applications.
- [92] D. Nellans, V.K. Kadaru, E. Brunv, Asim – An Asynchronous Architectural Level Simulator Abstract, 2004.
- [93] V. Puente, J. Gregorio, R. Beivide, Sicosys: an integrated framework for studying interconnection network performance in multiprocessor systems, in: Proceedings of the 10th Euromicro Workshop on Parallel, Distributed and Network-based Processing, 2002, 2002, pp. 15–22, doi:10.1109/EMPDP.2002.994207.
- [94] P.M. Ortega, P. Sack, SESC: SuperEScalar Simulator, The i-acoma Architecture Group at University of Illinois, USA, 2004.
- [95] R.E. Wunderlich, T.F. Wenisch, B. Falsafi, J.C. Hoe, Statistical sampling of microarchitecture simulation, ACM Trans. Model. Comput. Simul. 16 (3) (2006) 197–224, doi:10.1145/1147224.1147225.
- [96] F. Fazzino, M. Palesi, D. Patti, Noxim-noc Simulator, Online:<http://noxim.sourceforge.net> (2010).
- [97] K. Nakajima, T. Hieda, I. Taniguchi, H. Tomiyama, H. Takada, A fast network-on-chip simulator with qemu and systemc, in: Proceedings of the 2012 Third International Conference on Networking and Computing (ICNC), IEEE, 2012, pp. 298–301.
- [98] A. Portero, R. Pla, J. Carrabina, Systemc implementation of a noc, in: Proceedings of the 2005 IEEE International Conference on Industrial Technology, IEEE, 2005, pp. 1132–1135.
- [99] S. Chai, C. Wu, Y. Li, Z. Yang, A noc simulation and verification platform based on systemc, in: Proceedings of the 2008 International Conference on Computer Science and Software Engineering, 3, IEEE, 2008, pp. 423–426.
- [100] Y.-t. Li, L.-p. Liang, A noc modeling and simulating method with systemc, Microelectron. Comput. 3 (2010) 78–82.
- [101] X. Ningyi, L. Xianglun, L. Renfei, Z. Zucheng, A systemc-based noc simulation framework supporting heterogeneous communicators, in: Proceedings of the 2005 6th International Conference on ASIC, 2, IEEE, 2005, pp. 1032–1035.
- [102] J.V. Bruch, M.R. Pizzoni, C.A. Zeferino, Brownpepper: a systemc-based simulator for performance evaluation of networks-on-chip, in: Proceedings of the 2009 17th IFIP International Conference on Very Large Scale Integration (VLSI-SoC), IEEE, 2009, pp. 223–226.
- [103] M. Dall'Osso, G. Biccari, L. Giovannini, D. Bertozzi, L. Benini, Xpipes: a latency insensitive parameterized network-on-chip architecture for multiprocessor socs, in: Proceedings of the 2012 IEEE 30th International Conference on Computer Design (ICCD), 2012, pp. 45–48, doi:10.1109/ICCD.2012.6378615.
- [104] S. Benedict, Energy-aware performance analysis methodologies for {HPC} architectures – an exploratory study, J. Netw. Comput. Appl. 35 (6) (2012) 1709–1719. <http://dx.doi.org/10.1016/j.jnca.2012.08.003>.
- [105] R. Rao, S. Vrudhula, K. Berezowski, Analytical results for design space exploration of multi-core processors employing thread migration, in: Proceedings of the 2008 ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED), 2008, pp. 229–232, doi:10.1145/1393921.1393981.
- [106] B. Wojciechowski, K. Berezowski, P. Patronik, J. Biernat, Fast and accurate thermal simulation and modelling of workloads of many-core processors, in: Proceedings of the 2011 17th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC), 2011, pp. 1–6.
- [107] D. Brooks, V. Tiwari, M. Martonosi, Wattch: a framework for architectural-level power analysis and optimizations, SIGARCH Comput. Archit. News 28 (2) (2000) 83–94, doi:10.1145/342001.339657.
- [108] N. Vijaykrishnan, M. Kandemir, M.J. Irwin, H.S. Kim, W. Ye, Energy-driven integrated hardware-software optimizations using simplepower, SIGARCH Comput. Archit. News 28 (2) (2000) 95–106, doi:10.1145/342001.339659.
- [109] S. Gurumurthi, A. Sivasubramaniam, M.J. Irwin, N. Vijaykrishnan, M. Kandemir, T. Li, L.K. John, Using complete machine simulation for software power estimation: the softwatt approach, in: Proceedings of the 8th International Symposium on High-Performance Computer Architecture, HPCA '02, IEEE Computer Society, Washington, DC, USA, 2002, p. 141.
- [110] G. Contreras, M. Martonosi, J. Peng, G.-Y. Lueh, R. Ju, The xtrem power and performance simulator for the intel xscale core: design and experiences, ACM Trans. Embed. Comput. Syst. 6 (1) (2007), doi:10.1145/1210268.1210272.
- [111] H.-S. Wang, X. Zhu, L.-S. Peh, S. Malik, Orion: a power-performance simulator for interconnection networks, in: Proceedings of the 35th Annual IEEE/ACM International Symposium on Microarchitecture, (MICRO-35), 2002, pp. 294–305, doi:10.1109/MICRO.2002.1176258.

- [112] A. Kahng, B. Li, L.-S. Peh, K. Samadi, Orion 2.0: a power-area simulator for interconnection networks, *IEEE Trans. Very Large Scale Integr. Syst.* 20 (1) (2012) 191–196, doi:10.1109/TVLSI.2010.2091686.
- [113] S. Li, J.H. Ahn, R. Strong, J. Brockman, D. Tullsen, N. Jouppi, Mcpat: an integrated power, area, and timing modeling framework for multicore and manycore architectures, in: *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO-42, 2009*, pp. 469–480.
- [114] A. Flores, J. Aragon, M. Acacio, Sim-powercmp: a detailed simulator for energy consumption analysis in future embedded cmp architectures, in: *Proceedings of the 21st International Conference on Advanced Information Networking and Applications Workshops, AINAW '07, 1, 2007*, pp. 752–757, doi:10.1109/AINAW.2007.334.
- [115] D. Sunwoo, G. Wu, N. Patil, D. Chiou, Presto: an fpga-accelerated power estimation methodology for complex systems, in: *Proceedings of the 2010 International Conference on Field Programmable Logic and Applications (FPL), 2010*, pp. 310–317, doi:10.1109/FPL.2010.69.
- [116] W. Heirman, S. Sarkar, T.E. Carlson, I. Hur, L. Eeckhout, Power-aware multi-core simulation for early design stage hardware/software co-optimization, in: *Proceedings of the 21st International Conference on Parallel Architectures and Compilation Techniques, PACT '12, ACM, New York, NY, USA, 2012*, pp. 3–12, doi:10.1145/2370816.2370820.
- [117] S.M. Faizur Rahman, Q. Yi, H. Homayoun, Just-in-time component-wise power and thermal modeling, in: *Proceedings of the 12th ACM International Conference on Computing Frontiers, CF '15, ACM, New York, NY, USA, 2015*, pp. 2:1–2:8, doi:10.1145/2742854.2742880.
- [118] A. Laganà, A. Costantini, O. Gervasi, N.F. Lago, C. Manuali, S. Rampino, Compchem: progress towards gems a grid empowered molecular simulator and beyond, *J. Grid Comput.* 8 (4) (2010) 571–586.
- [119] N. Saboo, A.K. Singla, J.M. Unger, L.V. Kalé, Emulating petaflops machines and blue gene, in: *Proceedings of the 15th International Parallel & Distributed Processing Symposium, IPDPS '01, IEEE Computer Society, Washington, DC, USA, 2001*, p. 195.
- [120] T.L. Wilmarth, L.V. Kale, Pose: Getting over grainsize in parallel discrete event simulation, in: *Proceedings of the International Conference on Parallel Processing, ICPP 2004, IEEE, 2004*, pp. 12–19.
- [121] H. Casanova, A. Legrand, M. Quinson, Simgrid: a generic framework for large-scale distributed experiments, in: *Proceedings of the Tenth International Conference on Computer Modeling and Simulation, UKSIM 2008, IEEE, 2008*, pp. 126–131.
- [122] B. Penoff, A. Wagner, M. Tuxen, I. Rungele, Mpi-netsim: a network simulation module for mpi, in: *Proceedings of the 2009 15th International Conference on Parallel and Distributed Systems (ICPADS), IEEE, 2009*, pp. 464–471.
- [123] L. Benini, D. Bertozzi, A. Bogliolo, F. Menichelli, M. Olivieri, Mparm: exploring the multi-processor soc design space with systemc, *J. VLSI Signal Process. Syst. Signal Image Video Technol.* 41 (2) (2005) 169–182.
- [124] J. Cong, K. Gururaj, G. Han, A. Kaplan, M. Naik, G. Reinman, Mc-sim: an efficient simulation tool for mp soc designs, in: *Proceedings of the 2008 IEEE/ACM International Conference on Computer-Aided Design, IEEE Press, 2008*, pp. 364–371.
- [125] J. Chevalier, O. Benny, M. Rondonneau, G. Bois, E.M. Aboulhamid, F.-R. Boyer, Space: A Hardware/Software Systemc Modeling Platform Including an RTOS (2003).
- [126] M. Monton, A. Portero, M. Moreno, B. Martinez, J. Carrabina, Mixed sw/systemc SOC emulation framework, in: *Proceedings of the 2007 IEEE International Symposium on Industrial Electronics, IEEE, 2007*, pp. 2338–2341.
- [127] C. Helmstetter, V. Joloboff, Simsoc: a systemc tlm integrated iss for full system simulation, in: *Proceedings of the IEEE Asia Pacific Conference on Circuits and Systems, APCCAS 2008, IEEE, 2008*, pp. 1759–1762.
- [128] D.D. Gajski, J. Zhu, R. Dömer, A. Gerstlauer, S. Zhao, SpecC: Specification Language and Methodology, Springer Science & Business Media, 2012.
- [129] R. Buchmann, F. Pétrot, A. Greiner, Fast cycle accurate simulator to simulate event-driven behavior, in: *Proceedings of the 2004 International Conference on Electrical, Electronic and Computer Engineering, ICEEC'04, IEEE, 2004*, pp. 35–38.
- [130] J. Cong, K. Gururaj, G. Han, A. Kaplan, M. Naik, G. Reinman, Mc-sim: an efficient simulation tool for mp soc designs, in: *Proceedings of the 2008 IEEE/ACM International Conference on Computer-Aided Design, ICCAD '08, IEEE Press, Piscataway, NJ, USA, 2008*, pp. 364–371.
- [131] T.-C. Yeh, M.-C. Chiang, On the interfacing between qemu and systemc for virtual platform construction: using dma as a case, *J. Syst. Archit.* 58 (3) (2012) 99–111.
- [132] D. Quaglia, F. Fummi, M. Macrina, S. Saggini, Timing aspects in qemu/systemc synchronization, in: *Proceedings of the International QEMU Users' Forum, 2011*, pp. 11–14.
- [133] M.-C. Chiang, T.-C. Yeh, G.-F. Tseng, A qemu and systemc-based cycle-accurate iss for performance estimation on soc development, *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* 30 (4) (2011) 593–606.
- [134] T.-C. Yeh, G.-F. Tseng, M.-C. Chiang, A fast cycle-accurate instruction set simulator based on qemu and systemc for soc development, in: *Proceedings of the 15th IEEE Mediterranean Electrotechnical Conference, MELECON 2010–2010, IEEE, 2010*, pp. 1033–1038.
- [135] M. Becker, H. Zabel, W. Müller, Qemu/systemc cosimulation at different abstraction levels, in: *Proceedings of the 1st International QEMU Users Forum, 2011*.
- [136] T.-C. Yeh, M.-C. Chiang, On the interface between qemu and systemc for hardware modeling, in: *Proceedings of the 2010 International Symposium on Next Generation Electronics, IEEE, 2010*, pp. 73–76.
- [137] D. Bortolotti, A. Marongiu, L. Benini, Virtualsoc: a research tool for modern mp socs, *ACM Trans. Embed. Comput. Syst.* 16 (1) (2016) 3:1–3:27, doi:10.1145/2930665.
- [138] J.E. Miller, H. Kasture, G. Kurian, C. Gruenwald, N. Beckmann, C. Celio, J. Eastep, A. Agarwal, Graphite: a distributed parallel simulator for multicores, in: *Proceedings of the 2010 IEEE 16th International Symposium on High Performance Computer Architecture (HPCA), IEEE, 2010*, pp. 1–12.
- [139] S. Jiao, D. Wang, X. Ye, W. Xu, H. Zhang, N. Sun, Partitionsim: A parallel simulator for many-cores, in: *Proceedings of the 2012 IEEE 9th International Conference on Embedded Software and Systems (HPCC-ICES) and 2012 IEEE 14th International Conference on High Performance Computing and Communication, IEEE, 2012*, pp. 119–126.
- [140] G. Kurian, S.M. Neuman, G. Bezerra, A. Giovinazzo, S. Devadas, J.E. Miller, Power modeling and other new features in the graphite simulator, in: *Proceedings of the 2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), IEEE, 2014*, pp. 132–134.
- [141] S. Jafer, Q. Liu, G. Wainer, Synchronization methods in parallel and distributed discrete-event simulation, *Simul. Model. Pract. Theory* 30 (0) (2013) 54–73. <http://dx.doi.org/10.1016/j.simpat.2012.08.003>.
- [142] H. Lv, Y. Cheng, L. Bai, M. Chen, D. Fan, N. Sun, P-gas: parallelizing a cycle-accurate event-driven many-core processor simulator using parallel discrete event simulation, in: *Proceedings of the 2010 IEEE Workshop on Principles of Advanced and Distributed Simulation (PADS), 2010*, pp. 1–8, doi:10.1109/PADS.2010.5471655.
- [143] Z. Cao, J. Xu, M. Chen, G. Zheng, H. Lv, N. Sun, Hppnetsim: a parallel simulation of large-scale interconnection networks, in: *Proceedings of the 2009 Spring Simulation Multiconference, SpringSim '09, Society for Computer Simulation International, San Diego, CA, USA, 2009*, pp. 32:1–32:8.
- [144] E.A. León, R. Riesen, A.B. Maccabe, P.G. Bridges, Instruction-level simulation of a cluster at scale, in: *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09, ACM, New York, NY, USA, 2009*, pp. 3:1–3:12, doi:10.1145/1654059.1654063.
- [145] R. Ubal, J. Sahuquillo, S. Petit, P. López, Multi2sim: a simulation framework to evaluate multicore-multithread processors, in: *Proceedings of the IEEE 19th International Symposium on Computer Architecture and High Performance Computing, 2007*, pp. 62–68.
- [146] W. Gropp, S. Huss-Lederman, A. Lumsdaine, E. Lusk, B. Nitzberg, W. Saphir, M. Snir, MPI – The Complete reference: The MPI-2 extensions, 2, MIT Press, Cambridge, MA, 1998.
- [147] Z. Tan, A. Waterman, R. Avizienis, Y. Lee, H. Cook, D. Patterson, K. Asanović, Ramp gold: an fpga-based architecture simulator for multiprocessors, in: *Proceedings of the 47th Design Automation Conference, DAC '10, ACM, New York, NY, USA, 2010*, pp. 463–468, doi:10.1145/1837274.1837390.
- [148] Z. Tan, A. Waterman, H. Cook, S. Bird, K. Asanović, D. Patterson, A case for fame: fpga architecture model execution, *SIGARCH Comput. Archit. News* 38 (3) (2010) 290–301, doi:10.1145/1816038.1815999.
- [149] D. Chiou, D. Sunwoo, J. Kim, N. Patil, W. Reinhart, D. Johnson, J. Keefe, H. Angepat, Fpga-accelerated simulation technologies (fast): fast, full-system, cycle-accurate simulators, in: *Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2007, 2007*, pp. 249–261, doi:10.1109/MICRO.2007.36.

- [150] E.S. Chung, E. Nurvitadhi, J.C. Hoe, B. Falsafi, K. Mai, PROToFLEX: FPGA-accelerated Hybrid Functional Simulator, 2007 IEEE International Parallel and Distributed Processing Symposium, 2007, pp. 1–6, doi:[10.1109/IPDPS.2007.370516](https://doi.org/10.1109/IPDPS.2007.370516).
- [151] N. Agarwal, T. Krishna, L.-S. Peh, N. Jha, Garnet: a detailed on-chip network model inside a full-system simulator, in: Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS 2009, 2009, pp. 33–42, doi:[10.1109/ISPASS.2009.4919636](https://doi.org/10.1109/ISPASS.2009.4919636).
- [152] J.-M. Berge, A. Fonkoua, S. Maginot, J. Rouillard, Verilog and vhd1, 1992.
- [153] P. Mishra, N. Dutt, Architecture description languages for programmable embedded systems, IEE Proc Comput. Digit. Tech. 152 (3) (2005) 285–297, doi:[10.1049/ip-cdt:20045071](https://doi.org/10.1049/ip-cdt:20045071).
- [154] E.M. Dashofy, A. Van der Hoek, R.N. Taylor, A highly-extensible, xml-based architecture description language, in: Proceedings of Working IEEE/IFIP Conference on Software Architecture, 2001, IEEE, 2001, pp. 103–112.
- [155] B. Christopher, J.J. Lee, A Dynamically Configurable Discrete Event simulation Framework for Many-core Chip Multiprocessors. (2015).
- [156] J.M. Schmidt, Towards Many-core Processor Simulation on Cloud Computing Platforms, Purdue University, USA, 2011 Master's thesis.
- [157] J. Kuper, C. Baaij, M. Kooijman, M. Gerards, Exercises in architecture specification using c.lash, in: Proceedings of Forum on Specification and Design Languages, FDL 2010, ECSI Electronic Chips & Systems design Initiative, Gières, France, 2010, pp. 178–183.
- [158] R. Wester, C. Baaij, J. Kuper, A two step hardware design method using c.lash, in: Proceedings of the 22nd International Conference on Field Programmable Logic and Applications, FPL 2012, IEEE Computer Society, 2012, pp. 181–188.
- [159] C. Baaij, M. Kooijman, J. Kuper, A. Boeijink, M. Gerards, C.lash: Structural descriptions of synchronous hardware using haskell, in: Proceedings of the 13th EUROMICRO Conference on Digital System Design: Architectures, Methods and Tools, IEEE Computer Society, 2010, pp. 714–721.
- [160] L. Eeckhout, K. De Bosschere, Speeding up architectural simulations for high-performance processors, Simul. – Trans. Soc. Model. Simul. Int. 80 (9) (2004) 451–468.
- [161] B. Xypolitidis, R. Shabani, S.V. Khandeparkar, Z. Ul-Abdin, T. Nordstr, et al., Towards architectural design space exploration for heterogeneous many-cores, in: Proceedings of the 2016 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP), IEEE, 2016, pp. 805–810.
- [162] T. Huang, X. Li, T. Zhang, D.-T. Lu, Gpu-accelerated direct sampling method for multiple-point statistical simulation, Comput. Geosci. 57 (2013) 13–23.
- [163] S. Hassani, G. Southern, J. Renau, Livesim: going live with microarchitecture simulation, in: Proceedings of the 2016 IEEE International Symposium on High Performance Computer Architecture (HPCA), 2016, pp. 606–617, doi:[10.1109/HPCA.2016.7446098](https://doi.org/10.1109/HPCA.2016.7446098).
- [164] D. Meisner, J. Wu, T.F. Wenisch, Bighouse: a simulation infrastructure for data center systems, in: Proceedings of the 2012 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), IEEE, 2012, pp. 35–45.
- [165] L. Eeckhout, H. Vandierendonck, K. Bosschere, Workload design: selecting representative program-input pairs, in: Proceedings.Proceedings of the 2002 International Conference on Parallel Architectures and Compilation Techniques, 2002., IEEE, 2002, pp. 83–94.
- [166] M. Van Biesbrouck, L. Eeckhout, B. Calder, Representative multiprogram workloads for multithreaded processor simulation, in: Proceedings of the IEEE 10th International Symposium on Workload Characterization, IISWC 2007, IEEE, 2007, pp. 193–203.
- [167] M.B. Breughe, L. Eeckhout, Selecting representative benchmark inputs for exploring microprocessor design spaces, ACM Trans. Archit. Code Optim. 10 (4) (2013) 37.
- [168] A. KleinOsowski, J. Flynn, N. Meares, D.J. Lilja, Workload Characterization of Emerging Computer Applications, Kluwer Academic Publishers, Norwell, MA, USA, 2001, pp. 83–100.
- [169] I. Uddin, Design Space Exploration in the Microthreaded Many-core Architecture, arXiv preprint arXiv:1309.5551 (2013).
- [170] S.-h. Kang, D. Yoo, S. Ha, TQSIM: a fast cycle-approximate processor simulator based on QEMU, J. Syst. Archit. 66–67 (2016) 33–47, doi:[10.1016/j.sysarc.2016.04.012](https://doi.org/10.1016/j.sysarc.2016.04.012).
- [171] P.D. Bryan, J.A. Poovey, J.G. Beu, T.M. Conte, Accelerating multi-threaded application simulation through barrier-interval time-parallelism, in: Proceedings of the 2012 IEEE 20th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 2012, pp. 117–126, doi:[10.1109/MASCOTS.2012.23](https://doi.org/10.1109/MASCOTS.2012.23).
- [172] L.M. Chan, R. Srinivasan, A hybrid CPU-graphics processing unit (GPU) approach for computationally efficient simulation-optimization, Comput. Chem. Eng. 87 (2016) 49–62, doi:[10.1016/j.compchemeng.2016.01.001](https://doi.org/10.1016/j.compchemeng.2016.01.001).
- [173] M.C. Fu, F.W. Glover, J. April, Simulation optimization: a review, new developments, and applications, in: Proceedings of the 37th Conference on Winter Simulation, 2005, pp. 83–95.
- [174] B.C. Lee, D.M. Brooks, Accurate and efficient regression modeling for microarchitectural performance and power prediction, SIGOPS Oper. Syst. Rev. 40 (5) (2006) 185–194, doi:[10.1145/1168917.1168881](https://doi.org/10.1145/1168917.1168881).
- [175] J.H. Novak, E. Brunvand, Using fpgas to prototype a self-timed floating point co-processor, in: Proceedings of the IEEE 1994 Custom Integrated Circuits Conference, IEEE, 1994, pp. 85–88.
- [176] J. Nunez-Yanez, G. Lore, Enabling accurate modeling of power and energy consumption in an arm-based system-on-chip, Microprocess. Microsyst. 37 (3) (2013) 319–332. <http://dx.doi.org/10.1016/j.micpro.2012.12.004>.
- [177] A. BenA., Y. Tsutomu, S. Masahiro, High-level modeling and fpga prototyping of produced order parallel queue processor core, J. Supercomput. 38 (1) (2006) 3–15, doi:[10.1007/s11227-006-6719-5](https://doi.org/10.1007/s11227-006-6719-5).
- [178] J. Wawrzynek, D. Patterson, M. Oskin, S.-L. Lu, C. Kozyrakis, J. Hoe, D. Chiou, K. Asanovic, Ramp: research accelerator for multiple processors, IEEE Micro 27 (2) (2007) 46–57, doi:[10.1109/MM.2007.39](https://doi.org/10.1109/MM.2007.39).
- [179] A. Krasnov, A. Schultz, J. Wawrzynek, G. Gibeling, P.-Y. Droz, Ramp blue: a message-passing manycore system in fpgas, in: Proceedings of the International Conference on Field Programmable Logic and Applications, FPL 2007, 2007, pp. 54–61, doi:[10.1109/FPL.2007.4380625](https://doi.org/10.1109/FPL.2007.4380625).
- [180] R. Jung, R. Heinrich, W. Hasselbring, Geco: a generator composition approach for aspect-oriented dsls, in: Proceedings of the International Conference on Theory and Practice of Model Transformations, Springer, 2016, pp. 141–156.
- [181] M. Chibani, B. Belattar, A. Bourouis, Toward an aspect-oriented simulation, Int. J. New Comput. Archit. Appl. 3 (1) (2013) 1–10.
- [182] T.B. Ionescu, A. Piater, W. Scheuermann, E. Laurien, An aspect-oriented approach for the development of complex simulation software., J. Object Technol. 9 (1) (2010) 161–181.
- [183] M. Chibani, B. Belattar, A. Bourouis, The use of the aspect oriented programming (aop) paradigm in discrete event simulation domain: overview and perspectives, in: Proceedings of the Third International Conference on Digital Information Processing and Communications, The Society of Digital Information and Wireless Communication, 2013, pp. 653–660.
- [184] M. Chibani, B. Belattar, A. Bourouis, Practical benefits of aspect-oriented programming paradigm in discrete event simulation, Model. Simul. Eng. 2014 (2014) 47.
- [185] J. Wu, X. Zhu, T. Li, X. Sui, Wbsp: a novel synchronization mechanism for architecture parallel simulation, IEEE Trans. Comput. 65 (3) (2016) 992–1005, doi:[10.1109/TC.2015.2439253](https://doi.org/10.1109/TC.2015.2439253).
- [186] M. Eggenberger, M. Strobel, M. Radetzki, Globally asynchronous locally synchronous simulation of nocs on many-core architectures, in: Proceedings of the 2016 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP), IEEE, 2016, pp. 763–770.
- [187] Z. Yu, Y. Chen, Y. Song, S. Huang, Comparison of parallel implementations of controls on gpu for transient simulation of power system, in: Proceedings of the 2016 35th Chinese Control Conference (CCC), 2016, pp. 9996–10001, doi:[10.1109/ChiCC.2016.7554936](https://doi.org/10.1109/ChiCC.2016.7554936).
- [188] N. Belhadj, N. Bahri, Z. Marrakchi, M.A.B. Ayed, N. Masmoudi, H. Mehrez, H.264/avc high definition intra coding implementation on multiprocessor system on chip technology architecture, IET Comput. Digit. Tech. 9 (5) (2015) 259–267, doi:[10.1049/iet-cdt.2014.0151](https://doi.org/10.1049/iet-cdt.2014.0151).
- [189] H. He1, R. Li, X. Dong1, Z. Zhang, H. Han, An efficient and secure cloud-based distributed simulation system, Appl. Math. Inf. Sci. 6 (3) (2012) 729–736.
- [190] G. Sakellari, G. Loukas, A survey of mathematical models, simulation approaches and testbeds used for research in cloud computing, Simul. Model. Pract. Theory 39 (2013) 92–103. S.I. Energy efficiency in grids and clouds. <http://dx.doi.org/10.1016/j.simpat.2013.04.002>.

- [191] F. Caglar, S. Shekhar, A. Gokhale, S. Basu, T. Rafi, J. Kinnebrew, G. Biswas, Cloud-hosted simulation-as-a-service for high school {STEM} education, *Simul. Model. Pract. Theory* 58, Part 2 (2015) 255–273. Special issue on Cloud Simulation. <http://dx.doi.org/10.1016/j.simpat.2015.06.006>.
- [192] W. Tian, M. Xu, A. Chen, G. Li, X. Wang, Y. Chen, Open-source simulators for cloud computing: comparative study and challenging issues, *Simul. Model. Pract. Theory* 58, Part 2 (2015) 239–254. Special issue on Cloud Simulation. <http://dx.doi.org/10.1016/j.simpat.2015.06.002>.
- [193] B. Kamiński, P. Szufel, On optimization of simulation execution on amazon {EC2} spot market, *Simul. Model. Pract. Theory* 58, Part 2 (2015) 172–187. Special issue on Cloud Simulation. <http://dx.doi.org/10.1016/j.simpat.2015.05.008>.
- [194] D. Genbrugge, S. Eyerma, L. Eeckhout, Interval simulation: Raising the level of abstraction in architectural simulation, in: *Proceedings of the 2010 IEEE 16th International Symposium on High Performance Computer Architecture (HPCA)*, 2010, pp. 1–12, doi:10.1109/HPCA.2010.5416636.
- [195] J. Bézin, *Model Driven Engineering: An Emerging Technical Space*, Springer, Berlin, Heidelberg, pp. 36–64. 10.1007/11877028_2.
- [196] R. France, B. Rumpe, Model-driven development of complex software: a research roadmap, in: *Proceedings of the 2007 Future of Software Engineering*, IEEE Computer Society, 2007, pp. 37–54.
- [197] F. Jouault, F. Allilaire, J. Bézin, I. Kurtev, Atl: a model transformation tool, *Sci. Comput. Program.* 72 (1) (2008) 31–39.
- [198] F. Jouault, F. Allilaire, J. Bézin, I. Kurtev, P. Valduriez, Atl: a qvt-like transformation language, in: *Proceedings of the Companion to the 21st ACM SIGPLAN Symposium on Object-oriented Programming Systems, Languages, and Applications*, ACM, 2006, pp. 719–720.
- [199] L. Yonglin, W. Weiping, L. Qun, Z. Yifan, A transformation model from {DEVS} to {SMP2} based on {MDA}, *Simul. Model. Pract. Theory* 17 (10) (2009) 1690–1709. <http://dx.doi.org/10.1016/j.simpat.2009.08.003>.
- [200] A. Alshareef, H.S. Sarjoughian, B. Zarrin, An approach for activity-based devs model specification, in: *Proceedings of the Symposium on Theory of Modeling & Simulation, TMS-DEVS '16*, Society for Computer Simulation International, San Diego, CA, USA, 2016, pp. 25:1–25:8.
- [201] B. Lee, D.-K. Kim, H. Yang, S. Oh, Model transformation between {OPC} {UA} and {UML}, *Comput. Stand. Interfaces* 50 (2017) 236–250. <http://dx.doi.org/10.1016/j.csi.2016.09.004>.
- [202] H.S. Sarjoughian, A. Alshareef, Y. Lei, Behavioral devs metamodeling, in: *Proceedings of the 2015 Winter Simulation Conference, WSC '15*, IEEE Press, Piscataway, NJ, USA, 2015, pp. 2788–2799.
- [203] L. Schubert, A. Kipp, Principles of service oriented operating systems, in: P. Vicat-Blanc Primet, T. Kudoh, J. Mambretti (Eds.), *Networks for Grid Applications*, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, 2, Springer, Berlin, Heidelberg, 2009, pp. 56–69, doi:10.1007/978-3-642-02080-3_6.
- [204] The S(o)OS Consortium, S(o)os (Service-oriented Operating System): Resource-independent Execution Support on Exa-scale Systems, 2010–2013. [Online: accessed 5-September-2014]. <http://www.soos-project.eu/>.
- [205] J. Zarrin, R.L. Aguiar, J.P. Barraca, ElCore: Dynamic elastic resource management and discovery for future large-scale manycore enabled distributed systems, *Microprocess. Microsyst.* 46, Part B (2016) 221–239. <http://dx.doi.org/10.1016/j.micpro.2016.06.007>.
- [206] T. Cucinotta, Challenges in Operating System Design for Future Many-core Systems, 2010, All Hands Meeting (AHM) 2010, Cardiff, UK, Available at <http://retis.sssup.it/~tommaso/presentations/AHM-2010.pdf>. [Online: accessed 15-April-2016].
- [207] J. Zarrin, R.L. Aguiar, J.P. Barraca, Dynamic, scalable and flexible resource discovery for large-dimension many-core systems, *Futur. Gener. Comput. Syst.* 53 (2015) 119–129. <http://dx.doi.org/10.1016/j.future.2014.12.011>.
- [208] J. Zarrin, R.L. Aguiar, J.P. Barraca, A self-organizing and self-configuration algorithm for resource management in service-oriented systems, in: *Proceeding of the 2014 IEEE Symposium on Computers and Communication (ISCC)*, 2014, pp. 1–7, doi:10.1109/ISCC.2014.6912524.
- [209] L. Schubert, A. Kipp, S. Wesner, Above the clouds: from grids to service-oriented operating systems, in: *Future Internet Assembly*, IOS Press, 2009, pp. 238–249, doi:10.3233/978-1-60750-007-0-238.
- [210] S. Kumar, T. Cucinotta, G. Lipari, A latency simulator for many-core systems, in: *Proceedings of the 44th Annual Simulation Symposium, ANSS '11*, Society for Computer Simulation International, San Diego, CA, USA, 2011, pp. 151–158.
- [211] R.L. Aguiar, D. Gomes, J. Barraca, N. Lau, Cloudthinking as an intelligent infrastructure for mobile robotics, *Wirel. Pers. Commun.* 76 (2) (2014) 231–244, doi:10.1007/s11277-014-1687-1.