

A Parallel Dynamic Binary Translator for Efficient Multi-Core Simulation

Oscar Almer · Igor Böhm · Tobias Edler von Koch ·
Björn Franke · Stephen Kyle · Volker Seeker ·
Christopher Thompson · Nigel Topham

Received: 20 March 2012 / Accepted: 11 September 2012 / Published online: 21 September 2012
© Springer Science+Business Media New York 2012

Abstract In recent years multi-core processors have seen broad adoption in application domains ranging from embedded systems through general-purpose computing to large-scale data centres. Simulation technology for multi-core systems, however, lags behind and does not provide the simulation speed required to effectively support design space exploration and parallel software development. While state-of-the-art instruction set simulators (ISS) for single-core machines reach or exceed the performance levels of speed-optimised silicon implementations of embedded processors, the same does not hold for multi-core simulators where large performance penalties are to be paid. In this paper we develop a fast and scalable simulation methodology for multi-core platforms based on parallel and just-in-time (JIT) dynamic binary translation (DBT). Our approach can model large-scale multi-core configurations, does not rely on prior profiling, instrumentation, or compilation, and works for all binaries targeting a state-of-the-art embedded multi-core platform implementing the ARCompact instruction set architecture (ISA). We have evaluated our parallel simulation methodology against the industry standard SPLASH-2 and EEMBC MULTIBENCH benchmarks and demonstrate simulation speeds up to 25,307 MIPS on a 32-core x86 host machine for as many as 2,048 target processors whilst exhibiting minimal and near constant overhead, including memory considerations.

Keywords Instruction set simulators · Just-in-time compilation ·
Multicore processors · Parallel dynamic binary translation ·
Scalable multicore simulation

O. Almer · I. Böhm · T. E. von Koch · B. Franke · S. Kyle ·
V. Seeker (✉) · C. Thompson · N. Topham
Institute for Computing Systems Architecture, University of Edinburgh, Edinburgh, UK
e-mail: s1160444@sms.ed.ac.uk

1 Introduction

With the proliferation of multi-core processor implementations in virtually all computing domains ranging from embedded systems through general-purpose desktop machines to large-scale data centres, the need is growing for scalable and multi-core enabled instruction set simulators (ISS). Software and hardware developers alike rely on efficient ISS technology for design space exploration, convenient software development in the early design stages of a new system, hardware and software performance optimisation, and debugging. Unfortunately, development of scalable multi-core ISS technology has not kept pace with the rapid advances in multi-core architecture design. Ever-increasing numbers of cores contained in future computing systems challenge today's simulation technology and pose serious performance and scalability bottlenecks.

Unlike multi-core simulators, simulation technology for single-processor systems is mature and provides high simulation rates for complex, superscalar, out-of-order architectures [3]. For some embedded processors, simulators even exceed the native execution performance of speed-optimised silicon implementations of the target processor [6] whilst, at the same time, providing architectural observability. These remarkable performance levels for single-core instruction set simulation have been enabled by Just-in-Time (JIT) Dynamic Binary Translators (DBT), which rely on advanced compiling techniques for efficient on-the-fly translation of binary code, and powerful simulation hosts providing the necessary CPU performance for efficient native code execution. However, due to its inherent complexity, up until now JIT DBT technology has not been available for multi-core simulation.

In this paper we develop novel technology for execution-driven, JIT DBT based multi-core simulation enabling faster-than-FPGA simulation speeds [31]. Our approach scales up to 2,048 target processors whilst exhibiting minimal and near constant overhead. We extend single-core JIT DBT to enable multi-core simulation and effectively exploit the parallelism offered by the simulation host for both parallel JIT compilation *and* native execution. In fact, multi-core JIT DBT is not just a single method, but a *combination* of several techniques. It is a widely held view that “we need a 1,000–10,000 MIPS simulator to support effective hardware and software research of a 1,000-way multiprocessor system” [12]. In this paper, we present the first-ever simulator—implemented entirely in software—to reach and exceed this performance level.

Our main contribution is to demonstrate how to effectively apply JIT DBT in the context of multi-core target platforms. The key idea is to model each simulated processor core in a separate thread, each of which feeds work items for native code translation to a parallel JIT compilation task farm shared among all CPU threads. Combined with private first-level caches and a shared second-level cache for recently translated and executed native code, detection and elimination of duplicate work items in the translation work queue, and an efficient low-level implementation for atomic exchange operations we construct a highly scalable multi-core simulator that provides faster-than-FPGA simulation speeds and scales favourably up to 2,048 simulated cores.

We have evaluated our simulation methodology against the MULTIBENCH and SPLASH-2 benchmark suites. Our functional ISS models homogeneous multi-core configurations composed of ENCORE [15] cores, which implement the ARCOMPACT ISA [30]. On a 32-core x86 host machine we demonstrate simulation rates up to 25,307 MIPS for as many as 2,048 target processors. Across all benchmarks our JIT DBT simulation approach achieves an average simulation performance of 11,797 MIPS (for 64 simulated cores) and outperforms an equivalent system implemented in FPGA.

1.1 Motivating Example

Consider the full-system simulation of the SPLASH-2 LU program built for the ENCORE processor. On a 32-core Intel Xeon machine we simulate the application with varying numbers of simulated cores, where each simulated core is modelled as a thread. The LU program is a parallel benchmark that factors a dense matrix into the product of a lower triangular and an upper triangular matrix. The dense $n \times n$ matrix A is divided into an $N \times N$ array of $B \times B$ blocks ($n = NB$) to exploit temporal locality on submatrix elements [37].

Our trace based dynamic binary translator speeds up the simulation by identifying and translating hot traces to native x86 code during simulation. Simulated target cores perform simulation as well as hot trace discovery and dispatch in parallel to each other. A light-weight decoupled JIT DBT task farm runs in parallel to this simulation compiling dispatched hot traces to speed up simulation. To ensure that the most profitable traces are compiled first, we use a work scheduling strategy that dynamically prioritises compilation tasks according to their heat and recency [7].

For the purpose of this motivating example we look at the scalability of our ISS with increasing numbers of simulated target cores. In Chart ① of Fig. 1 we compare the overall simulation rate in MIPS starting with a single-core simulation and scaling up to a multi-core configuration comprising a total of 2,048 target cores. Our simulator scales almost perfectly—*doubling* the simulation rate each time the number of simulated processors is doubled—until the number of simulated target cores equals simulation host cores. Simulating 32 ENCORE cores on a 32-core x86 machine results in an overall simulation rate of 10,500 MIPS for the LU benchmark. Scaling the number of simulated target cores beyond 32 results in a more modest improvement between 10,800 MIPS for 64 simulated cores and the maximum simulation rate of 13,000 MIPS for 1,024 cores. For a total of 2,048 simulated cores we still achieve a simulation rate of 11,982 MIPS.

Application speedup over single-core simulation is shown in Chart ② of Fig. 1. The maximum speedup of $4.4\times$ is reached when simulating 16 target cores and even with 64 simulated cores the application shows a speedup of $1.7\times$.

Chart ③ of Fig. 1 illustrates the ratio of interpreted to natively executed instructions, demonstrating the performance of our parallel JIT DBT. For all configurations up to 2,048 simulated target cores the JIT DBT achieves $>99.1\%$ natively executed instructions.

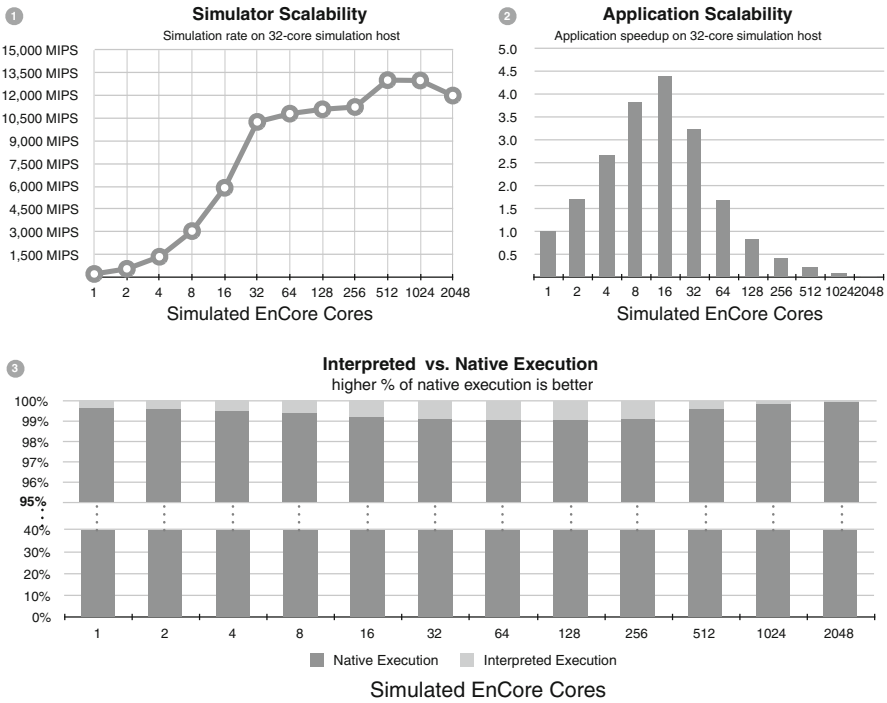


Fig. 1 Full-system simulation of SPLASH- 2 LU benchmark—comparison of ① simulation rate in MIPS, ② speedup over single core simulation, and ③ interpreted vs. natively executed instructions in %, for the simulation of 1–2,048 ENCORE cores on a 32-core x86 simulation host

1.2 Contributions

Among the contributions of this paper are:

1. The development of a scalable multi-core instruction set simulation methodology extending established single-core JIT DBT approaches to effectively exploit the available hardware parallelism of the simulation host,
2. the use of an innovative parallel task farm strategy to achieve truly concurrent JIT compilation of hot traces,
3. the combination of a multi-level cache hierarchy for JIT-compiled code, detection and elimination of duplicate work items in the translation work queue, and an efficient low-level implementation for atomic exchange operations, and
4. an extensive evaluation of the performance and memory consumption of our LLVM-based DBT targeting a multi-core platform implementing the ARCOMPACT ISA and using two benchmark suites: EEMBC MULTIBENCH and SPLASH- 2.

Currently, ARCSIM supports cycle-accurate simulation of each processor core [6]. However, faithful performance modelling of the overall multi-core system including its memory hierarchy and interconnect is beyond the scope of this paper and subject of our future work.

1.3 Overview

The remainder of this paper is structured as follows. In Sect. 2 we present details of the ENCORE processor and provide an overview of the target system architecture. This is followed by an extensive coverage of our proposed parallel simulation methodology in Sect. 3. In Sect. 4 we present the results of our empirical evaluation before we discuss related work in Sect. 5. Finally, we summarise our findings and give an outlook to our future work in Sect. 6.

2 Background

2.1 EnCore Processor

Our simulator models a state-of-the-art embedded processor implementing the ARCOMPACT ISA, namely the ENCORE [15]. Its micro-architecture is based on a 5-stage interlocked pipeline with forwarding logic, supporting zero overhead loops, freely intermixable 16- and 32-bit instruction encodings, static branch prediction, branch delay slots, and predicated instructions.

Although the above configuration was used for this work, the processor is highly configurable. Pipeline depth, cache sizes, associativity, and block replacement policies as well as byte order (i.e. big endian, little endian), bus widths, register-file size, and instruction set specific options such as instruction set extensions are configurable.

2.2 Target System Architecture

As a physical reference system we used a 12-core implementation of the multi-core system, synthesised for a Xilinx X6VLX240T FPGA. The system architecture is shown in Fig. 2. The twelve processor cores (① in Fig. 2) are connected through a 32-bit hierarchical, switched, non-buffered AXI interconnect fabric (② in Fig. 2) to RAM and I/O devices (③ in Fig. 2). An ASIP implementation of the same ENCORE processor, implemented in a generic 90 nm technology node, is currently running in our laboratory at frequencies up to 600 MHz. The processor cores can attain a 50 MHz core clock using this FPGA fabric, while the interconnect is clocked asynchronously to the cores at 75 MHz.

JTAG accessible utility functions and event counters were inserted to be able to record data from the cores. Recorded data for each core includes total clock cycles when not halted, total committed instructions, total I/O operations, and total clock cycles spent on I/O operations. From these counters we calculate the MIPS of each core at 50 MHz (FPGA) and 600 MHz (ASIP), respectively.

2.3 Just-in-Time Dynamic Binary Translation

Efficient DBT relies heavily on JIT compilation for the translation of target machine instructions to host machine instructions. Although JIT compiled code generally runs

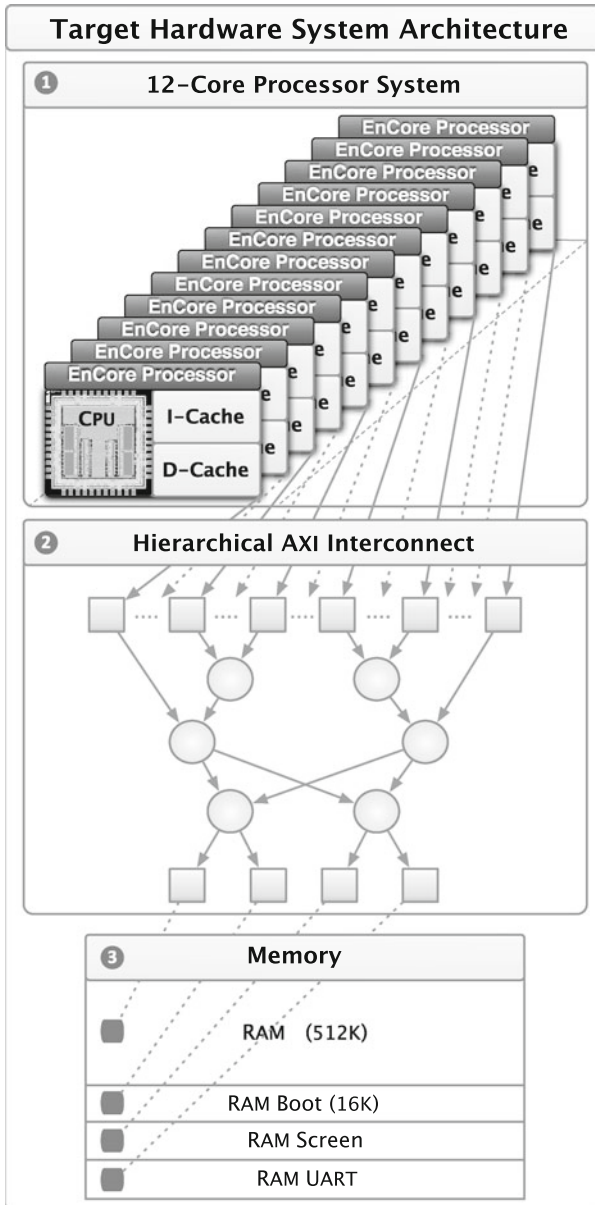


Fig. 2 Multi-core hardware architecture of target system

much faster than interpreted code, JIT compilation incurs an additional overhead. For this reason, only the most frequently executed code regions are translated to native code whereas less frequently executed code is still interpreted (see Box ① in Fig. 3). In a single-threaded execution model, the interpreter pauses until the JIT compiler has translated its assigned code block and the generated native code is executed directly.

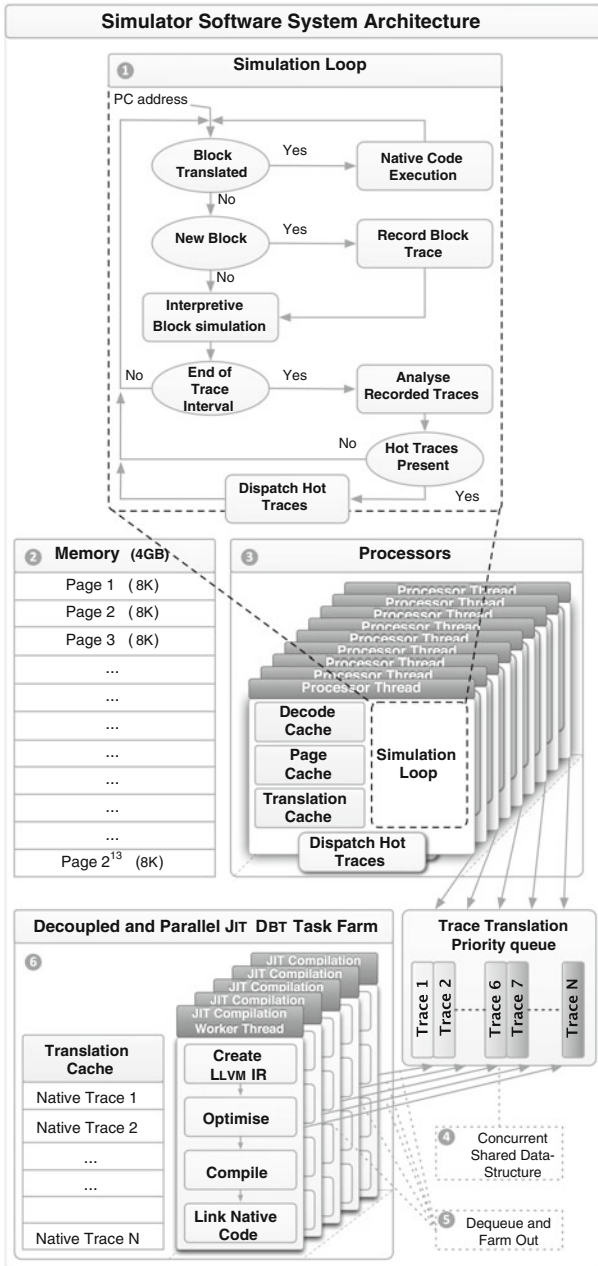


Fig. 3 Software architecture of our multi-core simulation capable ISS using parallel trace-based JIT DBT

However, it has been noted earlier [4] that program execution does not need to be paused to permit compilation, as a JIT compiler can operate in a separate thread while the program executes concurrently. This *decoupled* or *asynchronous* execution of the

JIT compiler increases complexity of the DBT, but is very effective in hiding the compilation latency—especially if the JIT compiler can run on a separate processor. In Sect. 3.2 we extend this concept and introduce a *parallel* JIT task farm (see Box ⑥ in Fig. 3) to further reduce compilation overhead and cope with the increased pressure on the JIT compiler.

3 Methodology

In this section, we discuss the implementation of the ARCSIM simulator, highlighting a number of key contributions which lead to high performance in a multi-core context. These include a carefully designed software architecture, sharing of translations between cores to benefit data-parallel applications, lightweight multi-threading support and an efficient mapping of atomic exchange operations to benefit synchronization.

3.1 Simulator Architecture

Figure 3 shows the architectural design of our simulator. Every simulated core runs within its own interpreter as an individual thread with its own set of CPU state variables. The operation of the interpreter is shown in Box ①, Fig. 3. As this is a functional simulation and is not yet required to be cycle-accurate, these threads are then allowed to simulate the program without restraint. The host operating system is relied upon for scheduling of the cores if more cores are simulated than are available on the host system. Simulated devices, such as screen and terminal I/O, are run in separate threads from the main simulation. Finally, JIT compilation workers also execute in separate threads in a task farm approach where all workers obtain work from a single queue, which all cores dispatch tasks to—see Box ⑥, Fig. 3. The operation of the JIT workers is discussed further in the following paragraph.

3.2 Parallel Just-in-Time Dynamic Binary Translation

Our method for JIT DBT begins with the recording of hot traces as the binary is interpreted. Interpreted simulation time is partitioned into *trace intervals*. After each trace interval every core dispatches the hottest recorded traces, packaged as work units, to a priority queue for JIT compilation before its simulation loop continues. Decoupled from these simulation loops, JIT compilation workers dequeue and compile the dispatched work units, as shown in Box ⑤, Fig. 3. A priority queue, shown in Box ④, Fig. 3, is used to ensure that the most important traces are scheduled for compilation first. Important traces are determined from a combination of a trace's execution frequency (= *heat*) and time since last execution (= *recency*). The main benefit of this approach is that compilation latency is hidden by performing it in parallel with the main simulation loops of the cores. Trace granularity can be adjusted from the level of single basic blocks to as large as an 8 KB page of memory. We use a light-weight

tracing scheme and only record basic block entry points as nodes, and pairs of source and target entry points as edges to construct a CFG.

3.3 Caching of Translated Code

Data parallel applications may run the same code across different cores. To avoid repeated translation of identical traces we employ a multi-level *translation cache* in our simulator.

As translations are generated for traces by the workers, they must be registered with their associated entry points within the simulator's internal representation of the simulated code. When the simulation reaches any basic block entry point, a lookup must be made to see if a translation exists for this physical address, returning a function pointer which can be called to execute this section of code natively. A direct-mapped *translation cache* was implemented which associates physical addresses with their translations, to avoid costly lookup. This cache is checked before consulting the simulator's internal representation, improving performance within the critical path of simulation. Each core has its own private *translation cache*, shown within Box ③, Fig. 3.

This cache also inspired a method for sharing translations between cores. By adding a second level *translation cache*, shown within Box ⑥, Fig. 3, which only the JIT compilation workers can access, the workers can check this cache when handling a work unit. Every time a JIT worker completes a task, it loads the generated translations into the second level *translation cache*. Whereas for the private first level cache the JIT worker had to always generate the code itself to be able to register the translation with the requesting core, the worker can now check the second level cache to see if any other workers have recently produced this translation. If it finds the translation in the cache, it can immediately register it with the requesting core, without having to generate the translation itself, saving time.

Depending on the chosen granularity of the translation units, determining if a translation is in the cache may be more complicated than only checking if the physical addresses match. When our translation units are page-sized, two cores may trace different paths through the page from the same initial physical address, and in this case it would be incorrect to try and use the translation of the first trace to simulate the second trace. For this reason, we also associate a signature with each work unit. The generation of this signature will be discussed in Sect. 3.4. For now, it is enough to say that if the physical address and signature of a JIT task matches those of a translation in the second level cache, it is possible to reuse that translation.

3.4 Detection and Elimination of Duplicate Work Items

The architecture of our multi-core JIT compilation system presents another opportunity to prevent redundant translation work. As multiple cores add tasks to the queue from which JIT workers receive their work, it would be beneficial to identify when a core attempts to dispatch a task which has already been added by another core. Once identified, we can prevent the task from being added to the queue, and instead register

that the requesting core would also like to be updated once the queued task has been handled by a worker.

In order to identify whether two tasks are for the same trace, we generate signatures for each task as it is created. The signature is the result of a hash function applied to the physical addresses of all the basic blocks that are to be translated. This provides us with a means to quickly determine if traces starting at the same address are different. While this may lead to false negatives, the error rate is diminished and is negligible in practice. The hash function is simple in our implementation to maximise performance, but any hash function which accepts 32-bit integers with more robust collision-avoidance could replace it.

A hash table stores lists of all tasks that have a particular key as their signature. Upon attempting to dispatch a new task to the queue, the table is checked for the task's signature. If found, the task is considered a duplicate, and is added to the corresponding list. If not, the task is new, and a new list containing it is added to the hash table, while the task is additionally added to the queue. The JIT compilation workers continue to take tasks from the queue as previously described, and, upon completing a task, check the hash table to see if any other tasks were added to the JIT system while this one was waiting or being processed. If any exist, the cores which created these tasks are also notified about the generated translations.

This technique, in addition to the shared caching of translations described in Sect. 3.3, has the dual effect of reducing the waiting period between the dispatch of a task and the receipt of its translation for many cores. It also reduces the amount of similar tasks in the work queue, resulting in a greater percentage of the simulated code being translated earlier.

3.5 Atomic Exchange Operations

The atomic exchange instructions in the ARCOMPACT instruction set are the way in which the ISA exposes explicit multi-core synchronisation. In ARCSIM, this is implemented using a global lock for all atomic exchange instructions. Whilst potentially expensive, this enables the use of the underlying x86 hardware synchronisation instructions to directly implement the atomic exchange instruction; in practice, it maps to a single x86 instruction.

4 Empirical Evaluation

We have evaluated our parallel JIT DBT multi-core simulator on over 20 benchmarks from the MULTIBENCH and SPLASH-2 benchmark suites. In this section we describe our experimental approach and present and evaluate our results for the speed and scalability of our simulator.

4.1 Experimental Setup and Methodology

We have evaluated our simulator using the MULTIBENCH 1.0 benchmark suite [33], which comprises a total of 14 application kernels from the networking and consumer

domains which can be combined in various ways to reflect typical application workloads of embedded systems. We run each application kernel separately with a varying number of worker threads. Some of the kernels are not multi-threaded, so in these cases we run a separate instance of the kernel on each of the simulated cores.

The SPLASH-2 benchmark suite [37] comprises 12 benchmarks which cover a number of common complex calculations in areas such as linear algebra, complex fluid dynamics, and graphics rendering. Each benchmark is designed to partition its work using threads.¹

Our results are reported in terms of MIPS achieved by the simulator. Each core can calculate its own MIPS rate, where the number of instructions the core has executed is divided by the length of time between when the core itself starts to execute instructions, and when the core is halted. These individual MIPS rates are summed to provide the total simulation rate. Each benchmark and core configuration was run 5 times, with the arithmetic mean taken from these runs to present our results.

The system used to run the simulator was a x86 host with 4 Intel Xeon L7555 1.87 GHz (8-core) processors with hyper-threading disabled, resulting in 32 cores being made available to the openSUSE 11.3 Linux operating system. The system also had 64GB of RAM available, and all experiments were run under conditions of low system load.

4.2 Bare-Metal POSIX Multi-Threading Support

We use a light-weight library that provides the essentials of operating system functionality to run the benchmarks on bare-metal hardware and within the simulator. This library provides features such as startup code, I/O device management, memory management primitives, and in particular basic multi-threading support in the form of a `pthread`s API.

We restrict the number of threads to one per core—which is permitted by the `pthread`s specification—since the implementation of a full preemptive scheduler would have exceeded the scope of this work. Mutexes are implemented using spin locks based around the ARCOMPACT atomic exchange instruction. Our implementation also fully supports condition variables, semaphores, and thread joining.

This approach differs from that taken for instance by UNISIM [2] where threading support is emulated by the simulator. This requires applications to be linked against a `pthread`s emulation library which re-routes `pthread`s API calls to trigger simulator intervention such as suspending or waking up of simulated cores. Our approach, in contrast, produces binaries that can be run both on real hardware and in the simulation environment without modification.

¹ We were unable to build versions of `fmm` and `water-nsquared` which could run on 64 cores, so these are excluded from our results. In the case where contiguous and non-contiguous versions of the same benchmark were available, we have built the contiguous version.

4.3 Summary of Key Results

Our results shown in Figs. 4 and 5 demonstrate that the initial target of “1,000–10,000 MIPS” [12] was easily attained, with seven of our benchmarks exceeding 20,000 MIPS when simulating a 32-core target. This is better than the performance of a theoretical 600 MHz 32-core ASIP. For the SPLASH-2 *fft* benchmark we achieve a maximum overall simulation rate of 25,307 MIPS for a 64-core simulation target, whilst on average we still provide 11,797 MIPS for the same simulation target. For large-scale configurations of up to 2,048 cores our results shown in Fig. 6 demonstrate the ability of our simulator to scale with the number of processors and to sustain its simulation rate beyond the point at which the number of simulated cores exceeds those of the host system.

4.4 Simulation Speed

All of the MULTIBENCH baseline three-core simulations exceed 1,000 MIPS, with *rgbhpg03* reaching 3,100 MIPS. Due to their higher complexity the single-core performance of the SPLASH-2 benchmarks ranges between 100 and 225 MIPS. On the other hand, they exhibit far greater scalability (see Sect. 4.5).

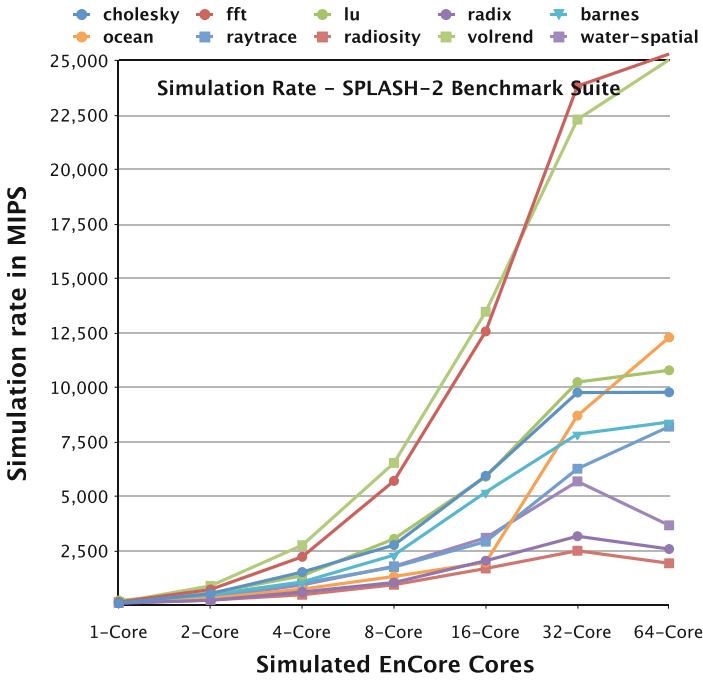
Simulating 64 target cores we achieve simulation rates in excess of 20,000 MIPS for *fft* and *volrend* from SPLASH-2, and for *md5*, *rgbcmk*, *mpeg2*, *rgbyiq03*, and *rgbhpg03* from MULTIBENCH. Only 5 out of 24 applications fail to deliver more than 3,200 MIPS (equivalent to 100 MIPS simulation rate per host core) while the average performance across all benchmarks for this configuration is close to 12,000 MIPS.

Not all benchmarks maintain this simulation rate as the number of cores increases, showing that simulation performance is application-specific. For instance, the MULTIBENCH networking benchmarks (*ippktcheck*, *ipres*, *tcp*) show little, if any, improvement over the baseline for higher numbers of simulated cores. The profile of the instructions executed by these benchmarks indicates a very high rate of memory accesses and memory-dependent branches which quickly saturate the available memory bandwidth of the host system. These findings are in line with the data sheets provided by EEMBC [33].

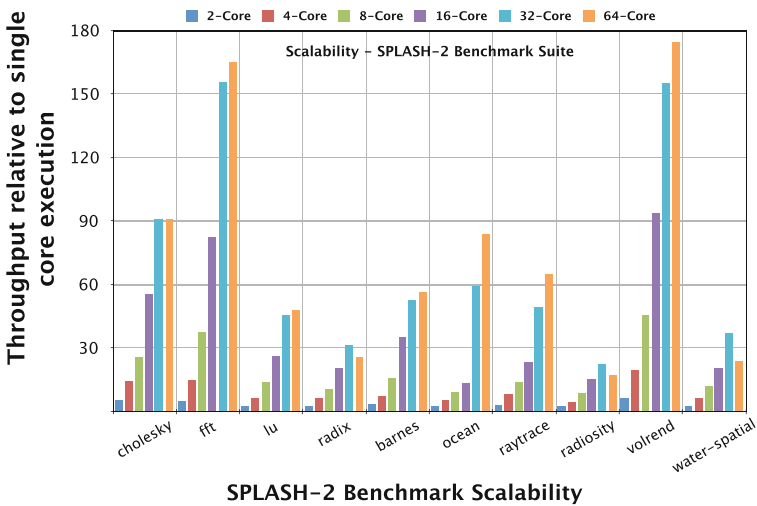
4.5 Scalability

It is important to evaluate how the simulator scales beyond the number of host processor cores for simulating tomorrow’s many-core systems on today’s commodity hardware. Most benchmarks demonstrate that the simulator scales well up to the number of physical cores on the host. Beyond this point we occasionally see modest further improvements (e.g. *cholesky*, *lu*, and *md5*) as shown in Fig. 6.

For the same seven benchmarks that deliver the highest simulation rates, we observe linear scalability as we increase the number of target cores up to 32. Other benchmarks such as *ocean*, *lu*, *cholesky*, *barnes*, *raytrace*, and *x264* do not achieve such high aggregate simulation rates, but still scale favourably.



(a)



(b)

Fig. 4 Running the EEMBC SPLASH- 2 benchmark suite with varying multi-core configurations of our ISS. **a** Simulation rate in MIPS. **b** Throughput relative to *one*-core simulation

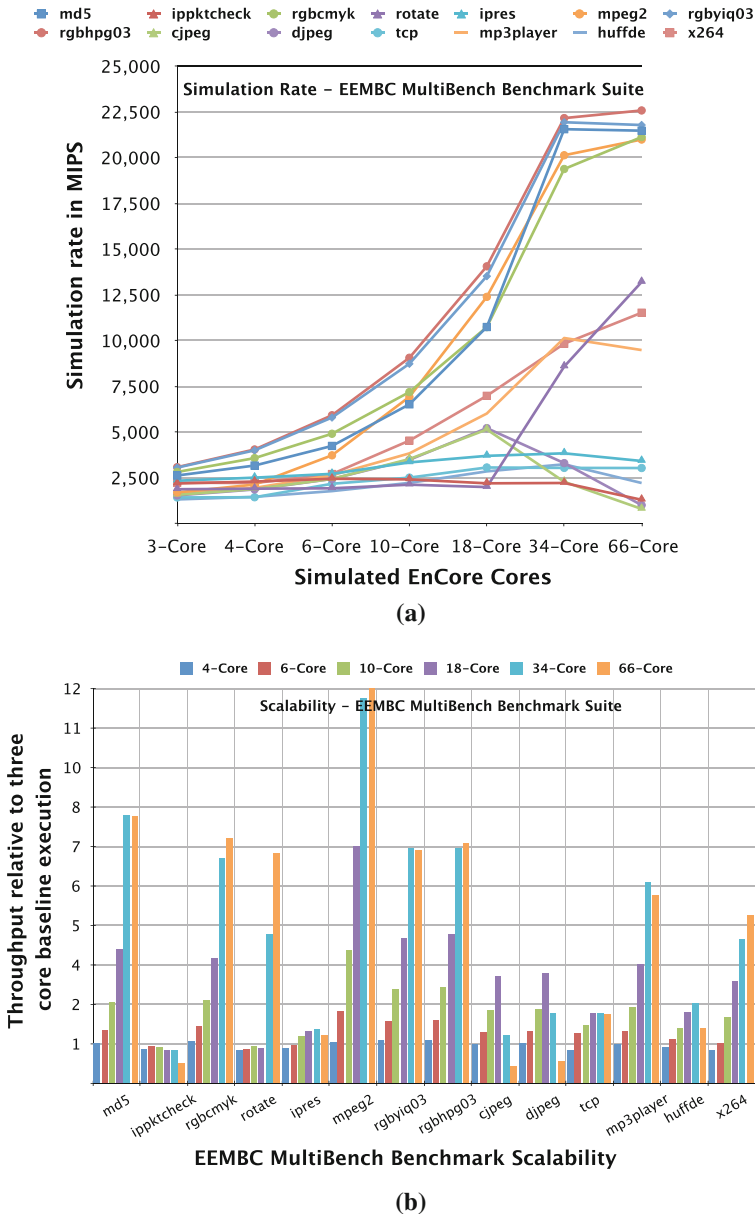


Fig. 5 Running the EEMBC MULTIBENCH benchmark suite with varying multi-core configurations of our ISS. Note that the minimum MULTIBENCH core configuration is three due to the test harness infrastructure. **a** Simulation rate in MIPS. **b** Throughput relative to *three*-core simulation

For six representative benchmarks (three from each benchmark suite) we show scalability up to 2,048 simulated target cores in Fig. 6. Chart ① in Fig. 6 shows the best result, with *cholesky* continuing to scale from 9,767 MIPS for 32 cores, to

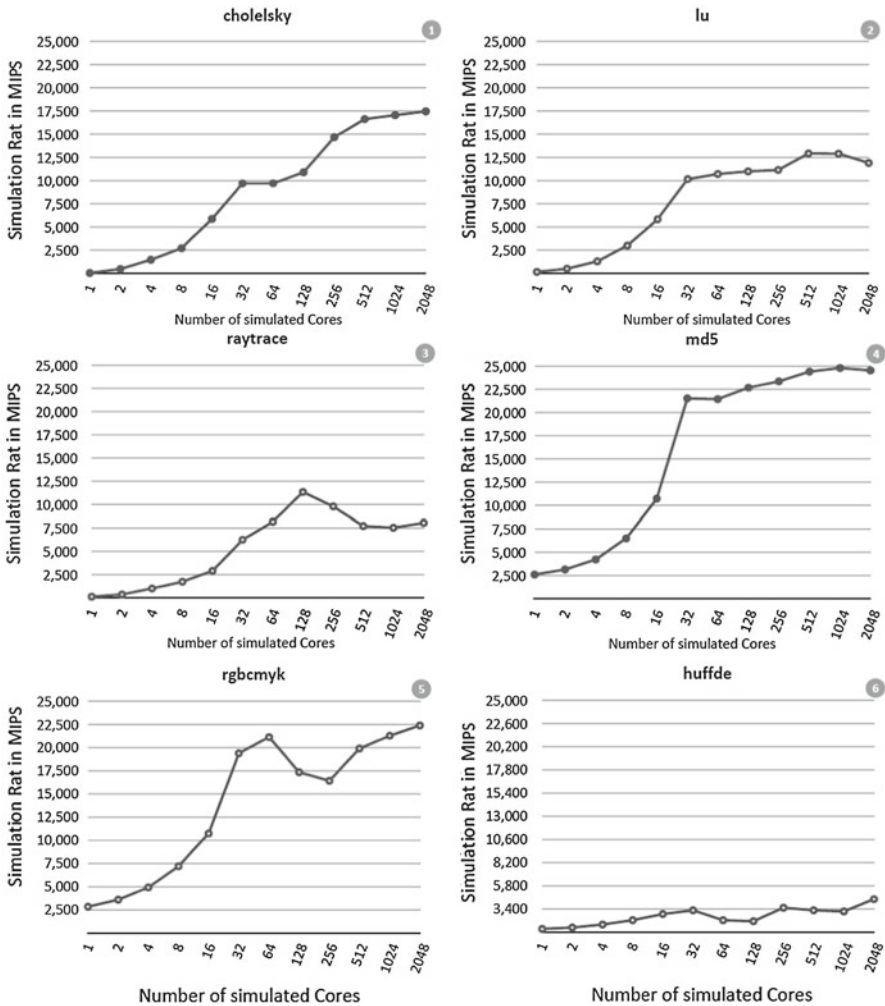


Fig. 6 Results for selected benchmarks from SPLASH-2 ①②③ and EEMBC MULTIBENCH ④⑤⑥ demonstrating the scalability with the number of simulated target cores

17,549 MIPS for 2,048 cores, with the performance always increasing. Chart ④ in Fig. 6 shows a similar result for md5.

In Fig. 4 we see super-linear scalability for a number of benchmarks (e.g. fft, ocean, volrend, cholesky). This is due to excessive synchronisation in the benchmarks under larger core configurations, and the fact that our simulator can execute tight spin-lock loops at near native speed. It is a well-known fact that the SPLASH-2 benchmarks attract high synchronisation costs for large-scale hardware configurations, as shown by other research [24]. The MULTIBENCH results are not affected in the same way due to less synchronisation.

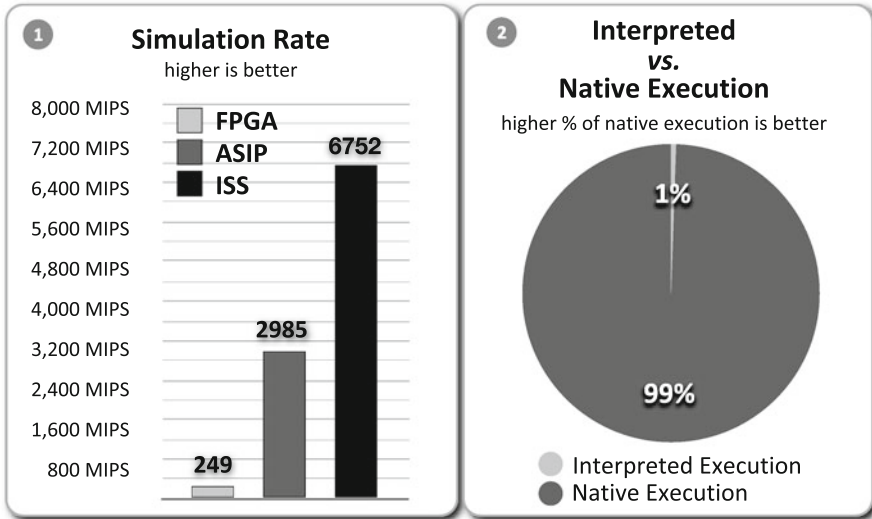


Fig. 7 Comparison of simulator and hardware implementations. ① shows a comparison of maximum achievable simulation rate in MIPS for a 12-core configuration running a parallel Mandelbrot fractal benchmark on an FPGA, ASIP, and ISS platform. ② depicts the ratio of interpreted vs. natively executed instructions on the ISS platform

4.6 Comparison to Native Execution on Real Hardware

Chart ① of Fig. 7 shows a comparison between our simulator and two hardware platforms (FPGA and ASIP, see Sect. 2.2) in terms of MIPS. The application is a parallelised fractal drawing algorithm, executed across 12 cores. We chose this application because of its low memory footprint and its embarrassingly parallel nature, thus avoiding application scalability issues.

Actual FPGA performance is 249 MIPS. We also show the performance of a 600 MHz ASIP implementation which achieves an execution rate of 2,985 MIPS. On the other hand, our instruction set simulator reaches a simulation rate of 6,752 MIPS, thus surpassing a silicon implementation by more than a factor of 2 for this application.

For equivalent configurations, our simulator consistently outperforms the theoretical maximum of the FPGA on a per-core basis. Our 12-core FPGA implementation of the multi-core system is capable of 50 MIPS per core. On the contrary, across all benchmarks, the lowest per-core simulation rate for a 16-core target was 105 MIPS, attained in the SPLASH-2 *radiosity* benchmark. These results show that even in the worst case the simulator maintains more than twice the theoretical maximum execution rate of the FPGA. Compared to the average simulation rate of 11,797 MIPS across all benchmarks, the theoretical maximum of 600 MIPS for a 12-core FPGA implementation is an order of magnitude slower.

4.7 Memory Consumption

In this section, we evaluate our DBT system in terms of its memory consumption. Using the SPLASH-2 benchmark suite, we study the overhead required to simulate a number of multi-core configurations. We chose not to present results for MULTIBENCH as these benchmarks are designed for embedded systems, and therefore are likely to have a much smaller memory footprint than those used in SPLASH-2. Due to the increase in execution time introduced by the infrastructure used for memory profiling, we have gathered results for 2, 4 and 8-core configurations only, and we were unable to obtain memory results for the `radiosity` benchmark. These experiments were executed on a system with a dual-core Intel Core i3-2100 CPU with 4GB of memory available, and with hyper-threading enabled.

In each experiment, we periodically sampled the number of 8KB pages the simulated benchmark had currently requested for allocation. We also periodically recorded the simulator's current heap usage. We use the terms *benchmark memory* and *heap usage* respectively for these measurements in our discussion. As the majority of our data structures used for simulation are allocated dynamically—particularly the data structures involved in the LLVM-based JIT compilation of traces—we feel that tracking the heap is a fair measurement of the memory consumption of our simulator.

For a given run, we use the term *heap overhead* to describe the peak heap usage minus the peak benchmark memory. This overhead accounts for all data in the heap that is required for simulation outside of the data requested by the benchmark itself. In Fig. 8, we present our findings in terms of the increase in this heap overhead as we increase the number of cores over a baseline 2-core simulation. These findings show that the increase in heap overhead grows sub-linearly as we increase the number of cores. As we double the number of cores from 2 to 4, we observe an average overhead increase of 1.35x. Furthermore, when we quadruple the number of cores from 2 to 8, we see an increase of only 1.71x. This indicates that the memory demand for simulating larger core configurations similar to those tested in previous sections will not become overwhelming.

In Fig. 9, we provide some examples of how benchmark memory and heap usage can change over time. We select the `fft` and `volrend` benchmarks as examples of benchmarks with interesting memory allocation patterns; all other benchmarks simply allocated all the pages they needed at startup. We also show how these memory profile traces change as we increase from 2 to 8 cores. We can see that the heap usage closely follows the general pattern of benchmark memory usage, with periodic bursts of increased heap usage. The vertical lines visible on these graphs show points where traces were being handled by the JIT subsystem. From this, we can see that these periods of increased heap usage coincide with JIT activity, indicating that the JIT subsystem accounts for a major part of increased heap usage, although the usage itself is short-lived.

We have already shown that the increase in heap overhead grows sub-linearly as we increase the number of cores, and that any increases in heap usage are usually accounted for by the JIT subsystem. Finally, we present the data used to calculate the heap overhead in Table 1, and show that the amount of extra memory required for each additional core is not significant. The average increase in overhead per core for

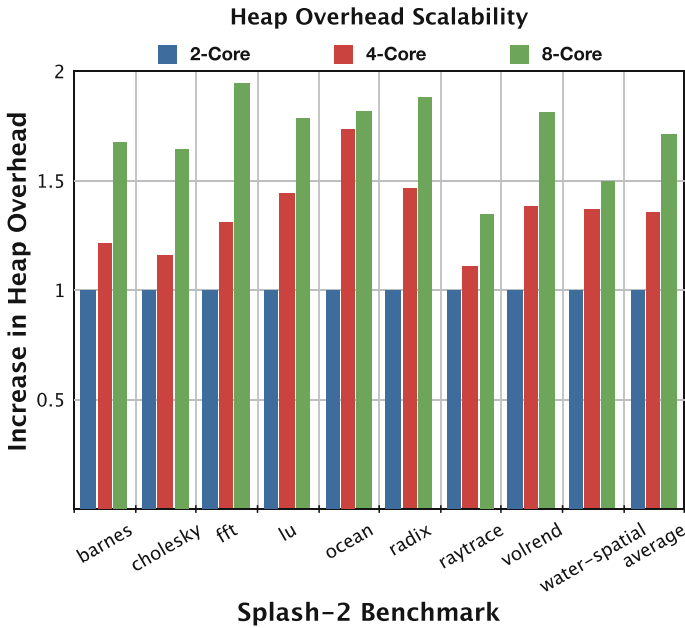


Fig. 8 The increase in *heap overhead* observed as we simulate 4 or 8 cores, in comparison to a 2-core baseline

each benchmark can be seen in the last column of this table. While *ocean* has the largest increase, at 23.16 MB per core, the average increase is only 9.14 MB per core. Although some benchmarks demonstrate a large amount of baseline heap overhead—such as 2-core *barnes*, with a heap overhead of 53.66 MB on top of the 4.82 MB used by the benchmark itself—we can see that the additional overhead introduced by more cores is much smaller. In the case of 2-core *barnes*, this is only 5.92 MB per core. In combination with the results shown in Fig. 8, this reinforces the point that our system has no problems with scaling to a large number of cores in terms of memory consumption.

5 Related Work

Due to the wealth of related research work we restrict our discussion to the most relevant approaches in software and FPGA based multi-core simulation.

5.1 Software Based Multi-Core Simulation Approaches

Early work on efficient simulation of parallel computer systems dates back to the 1990s [14] and examples of conventional sequential simulators/emulators include SIMPLESCALAR [39], RSIM [17], SIMOS [19], SIMICS [21], and QEMU [5]. Some of

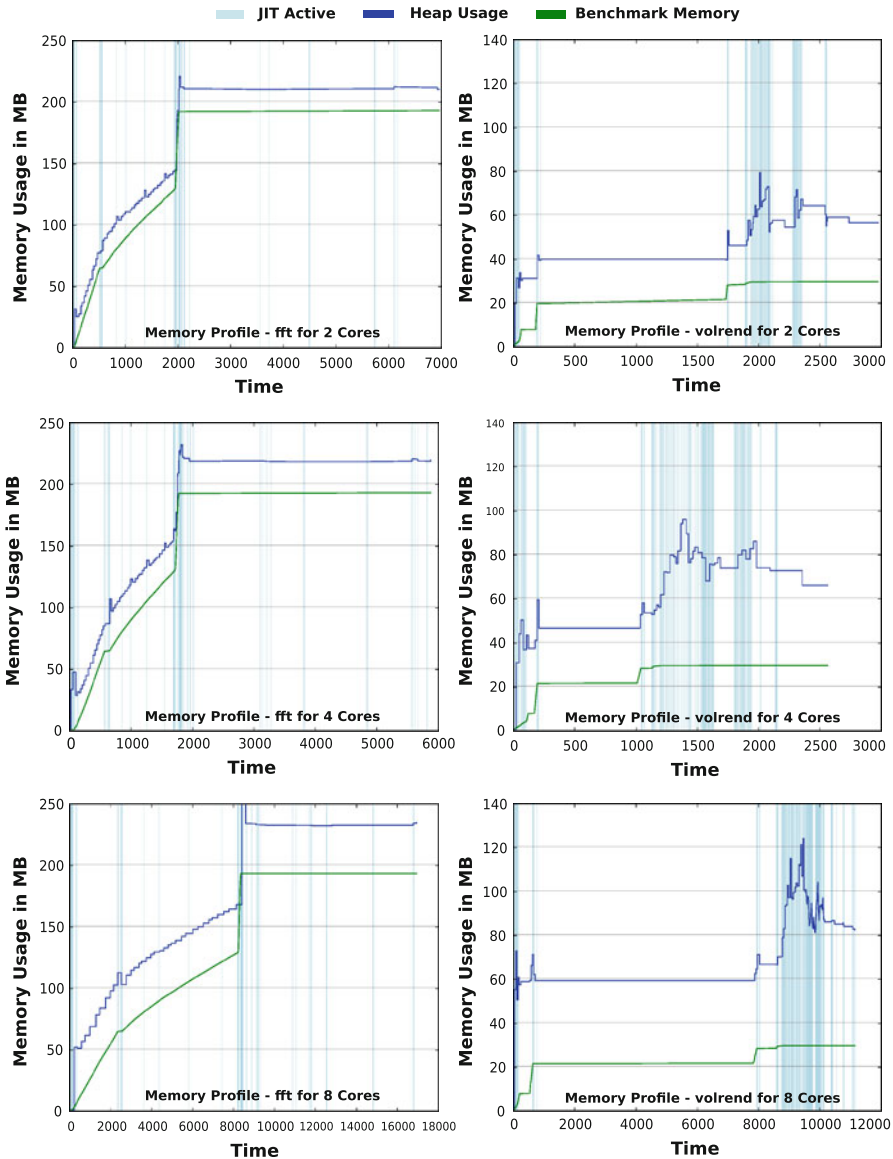


Fig. 9 Memory profile traces for the *fft* and *volrend* benchmarks for 2, 4 and 8 core simulations. The vertical lines indicate points where a dispatched trace was handled by the JIT subsystem

these are capable of simulating parallel target architectures but all of them execute sequentially on the host machine.

Over the last two decades a large number of parallel simulators of parallel target architectures have been developed: SIMFLEX [16], GEMS [22], COTSON [1], BIG-SIM [38], FASTMP [18], SLACKSIM [8], PCASIM [41], Wisconsin Wind Tunnel (WWT) [28], Wisconsin Wind Tunnel II (WWT II) [25], and those described by

Table 1 Memory profiling data for the SPLASH-2 benchmark suite

	Peak heap usage (MB)			Peak benchmark memory (MB)			Heap overhead (MB)			IOC [†]
	2-core	4-core	8-core	2-core	4-core	8-core	2-core	4-core	8-core	
barnes	58.48	69.73	94.55	4.82	4.52	4.57	53.66	65.2	89.98	5.92
chole.	82.2	95.5	134.9	3.24	4.13	5.03	78.95	91.36	129.87	7.35
fft	221.64	230.71	249.53	192.55	192.66	192.94	29.08	38.05	56.59	4.54
lu	48	68.15	83.92	2.38	2.41	2.48	45.62	65.74	81.43	8.01
ocean	106.8	175.03	183.41	15.02	15.61	16.66	91.78	159.42	166.76	23.16
radix	49.56	67.86	85.37	10.5	10.66	11.86	39.06	57.2	73.51	7.41
raytr.	170.35	181.66	206.18	67.59	67.63	67.71	102.77	114.03	138.47	5.79
volrend	81.42	101.38	123.72	29.44	29.45	29.49	51.98	71.94	94.23	8.51
water	89.04	121.19	132.07	2.22	2.24	2.27	86.82	118.95	129.8	11.61

[†] IOC is the increase in overhead per core

Chidester and George [9], and Penry et al. [27]. SIMFLEX and GEMS both use an off-the-shelf sequential emulator (SIMICS) for functional modeling plus their own models for memory systems and core interactions. GEMS uses their timing model to drive SIMICS one instruction at a time, but results in low performance. SIMFLEX avoids this problem by using statistical sampling of the application, but therefore does not observe its entire behaviour.

COTSON [1] uses AMD's SIMNOW! for functional modeling and suffers from some of the same problems as SIMFLEX [16] and GEMS [22].

BIGSIM [38] and FASTMP [18] assume distributed memory in their target architectures and do not provide coherent shared memory between the parallel portions of their simulators. WWT [28] is one of the earliest parallel simulators but requires applications to use an explicit interface for shared memory.

WWT II [25] does not model anything other than the target memory system and requires applications to be modified to explicitly allocate shared memory blocks. ARCSIM also models compute cores and implements a transparent shared memory system.

Like ARCSIM, PARALLEL EMBRA [20] is a fast functional simulator for shared-memory multiprocessors which is part of the PARALLEL SIMOS complete machine simulator. It takes an aggressive approach to parallel simulation; while it runs at user level and does not make use of the MMU hardware, it combines binary translation with loose timing constraints and relies on the underlying shared memory system for event ordering, time synchronisation, and memory synchronisation. While PARALLEL EMBRA shares its use of binary translation with ARCSIM it lacks its scalability and parallel JIT translation facility.

Another effort to parallelise a complete machine software simulator was undertaken with MAMBO [34]. It aims to produce a fast functional simulator by extending a binary translation based emulation mode; published results include a speedup of up to 3.8 for a 4-way parallel simulation. Similarly, the MALSIM [29] parallel functional simulator

has only been evaluated for workloads comprising up to 16 threads. Despite some conceptual similarities with these works our work aims at larger multi-core configurations where scalability is a major concern.

In [36] a DBT based simulator targeting a tiled architecture is presented. It aims at implementing different portions of a superscalar processor across distinct parallel elements thus exploiting the pipeline parallelism inherent in a superscalar microarchitecture. However, this work does not attempt to simulate a multi-core target platform.

ARMN [40] is a cycle-accurate simulator for homogeneous platforms comprising several ARM cores and with support for various interconnect topologies. Whilst this provides flexibility, the performance of ARMN is very low (approx. 10k instructions per second) and, thus, its suitability for both HW design space exploration and SW development is limited.

The GRAPHITE [23] multi-core simulation infrastructure is most relevant to our work. GRAPHITE is a distributed parallel multi-core simulator that combines direct execution, seamless multi-core and multi-machine distribution, and lax synchronisation. GRAPHITE has been demonstrated to simulate target architectures containing up to 1,024 cores on ten 8-core servers. Application threads are executed under a dynamic binary instrumentor (currently PIN) which rewrites instructions to generate events at key points. These events cause traps into GRAPHITE's backend which contains the compute core, memory, and network modelling modules. In contrast, our simulator uses DBT to implement any ISA (currently ARCOMPACT), which can also be different from the target system's ISA. In addition, the primary design goal of our simulator has been highest simulation throughput as showcased by the parallel JIT task farm comprised in ARCSIM. As a result we achieve speedups over native execution for many multi-core configurations, whereas GRAPHITE suffers up to $4007\times$ slowdown.

In [24] a methodology to simulate shared-memory multiprocessors composed of hundreds of cores is proposed. The basic idea is to use thread-level parallelism in the software system and translate it into core-level parallelism in the simulated world. An existing full-system simulator is first augmented to identify and separate the instruction streams belonging to the different software threads. Then, the simulator dynamically maps each instruction flow to the corresponding core of the target multi-core architecture, taking into account the inherent thread synchronisation of the running applications. This approach treats the functional simulator as a monolithic block, thus requiring an intermediate step for de-interleaving instructions belonging to different application threads. ARCSIM does not require this costly preprocessing step, but its functional simulator explicitly maintains parallel threads for the CPUs of the target system.

Outside the academic community the PCXS2 software simulator [26] for the Ps2 games console is a good example of a simulator targeting a multi-core platform which utilises dynamic binary translation.

5.2 FPGA Based Multi-Core Simulation Approaches

We can generally distinguish between two approaches to FPGA based multi-core simulation: The first approach essentially utilises FPGA technology for *rapid prototyping*,

but still relies on a detailed implementation of the target platform (see Sect. 2.1), whereas the second approach seeks to speed up performance modelling through a combined implementation of a functional simulator and a *timing model* on the FPGA fabric. In the following paragraphs we discuss approaches to this latter FPGA architecture simulation.

RAMP GOLD [31] is a state-of-the-art FPGA based “many-core simulator” supporting up to 64 cores. In functional-only mode, RAMP GOLD achieves a full simulator throughput of up to 100 MIPS when the number of target cores can cover the functional pipeline depth of 16. For fewer target cores (and non-synthetic workloads), the fraction of peak performance achieved is proportionally lower. In comparison to the 25,307 MIPS peak performance of our software-only simulation approach (based on an ISA of comparable complexity and similar functional-only simulation) the performance of the FPGA architecture simulation is more than disappointing. Other approaches to FPGA architecture model execution such as RAMP [35], FAME [32], FAST [11, 10] and PROTOFLEX [13] suffer from the same performance issues and for none of the mentioned systems has scalability beyond 64 cores been demonstrated.

6 Summary, Conclusions and Future Work

In this paper we have developed an innovative methodology for high-speed multi-core instruction set simulation based on JIT Dynamic Binary Translation. We have integrated a JIT task farm for parallel translation of hot code traces with a combination of performance-enhancing techniques such as private and shared caching of translated code, detection and elimination of identical translation units in the JIT work queue, and efficient low-level code generation for atomic exchange operations. Through this unique combination of techniques we achieve unprecedented simulator throughput of up to 25,307 MIPS and near-optimal scalability of up to 2,048 simulated cores for the SPLASH-2 and EEMBC MULTIBENCH benchmarks on a standard 32-core x86 simulation host, while maintaining a small increase in memory overhead per additional simulated core. Through our empirical results we demonstrate a simulation performance advantage by two orders of magnitude over leading and state-of-the-art FPGA architecture simulation technology [31] for a comparable level of simulation detail.

In our future work we will extend our cycle-accurate single-core performance model [6] to efficiently capture the detailed behaviour of the shared second level cache, processor interconnect and external memory of the simulated multi-core platform. In addition, a significant amount of work would be required to port our system to another ISA, requiring extensive modification of the interpreter and JIT system. Work is currently being conducted into producing a simulator that can be generated from an ISA description language, and we plan to integrate techniques from this paper’s work into such generated simulators.

References

1. Argollo, E., Falcón, A., Faraboschi, P., Monchiero, M., Ortega, D.: COTSon: infrastructure for full system simulation. *SIGOPS Oper. Syst. Rev.* **43**, 52–61 (2009). doi:[10.1145/1496909.1496921](https://doi.org/10.1145/1496909.1496921)

2. August, D., Chang, J., Girbal, S., Gracia-Perez, D., Mouchard, G., Penry, D.A., Temam, O., Vachharajani, N.: Unisim: an open simulation environment and library for complex architecture design and collaborative development. *IEEE Comput. Archit. Lett.* **6**, 45–48 (2007). doi:[10.1109/L-CA.2007.12](https://doi.org/10.1109/L-CA.2007.12)
3. Austin, T., Larson, E., Ernst, D.: SimpleScalar: an infrastructure for computer system modeling. *Computer* **35**, 59–67 (2002). doi:[10.1109/2.982917](https://doi.org/10.1109/2.982917)
4. Aycock, J.: A brief history of just-in-time. *ACM Comput. Surv.* **35**, 97–113 (2003)
5. Bellard, F.: QEMU, a fast and portable dynamic translator. In: Proceedings of the 2005 USENIX Annual Technical Conference, ATEC '05, pp. 41–41. USENIX Association, Berkeley, CA, USA (2005)
6. Böhm, I., Franke, B., Topham, N.P.: Cycle-accurate performance modelling in an ultra-fast just-in-time dynamic binary translation instruction set simulator. In: Kurdahi, F.J., Takala J. (eds.) ICSAMOS, pp. 1–10. IEEE (2010)
7. Böhm, I., Edler von Koch, T.J., Kyle, S., Franke, B., Topham, N.: Generalized just-in-time trace compilation using a parallel task farm in a dynamic binary translator. In: Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'11), ACM (2011)
8. Chen, J., Annavaram, M., Dubois, M.: SlackSim: a platform for parallel simulations of CMPs on CMPs. *SIGARCH Comput. Archit. News* **37**, 20–29 (2009). doi:[10.1145/1577129.1577134](https://doi.org/10.1145/1577129.1577134)
9. Chidester, M., George, A.: Parallel simulation of chip-multiprocessor architectures. *ACM Trans. Model. Comput. Simul.* **12**, 176–200 (2002). doi:[10.1145/643114.643116](https://doi.org/10.1145/643114.643116)
10. Chiou, D., Angepat, H., Patil, N., Sunwoo, D.: Accurate functional-first multicore simulators. *IEEE Comput. Archit. Lett.* **8**, 64–67 (2009). doi:[10.1109/L-CA.2009.44](https://doi.org/10.1109/L-CA.2009.44)
11. Chiou, D., Sunwoo, D., Angepat, H., Kim, J., Patil, N., Reinhart, W., Johnson, D.: Parallelizing computer system simulators. In: Parallel and Distributed Processing, 2008, IPDPS 2008. IEEE International Symposium on, pp. 1–5 (2008). doi:[10.1109/IPDPS.2008.4536407](https://doi.org/10.1109/IPDPS.2008.4536407)
12. Chung, E.S., Nurvitadhi, E., Hoe, J.C., Falsafi, B., Mai, K.: PROTOFLEX: FPGA-accelerated hybrid functional simulator. *Parallel Distrib. Process. Symp. Int.* **0**, 326 (2007). doi:[10.1109/IPDPS.2007.370516](https://doi.org/10.1109/IPDPS.2007.370516)
13. Chung, E.S., Papamichael, M.K., Nurvitadhi, E., Hoe, J.C., Mai, K., Falsafi, B.: ProtoFlex: towards scalable, full-system multiprocessor simulations using FPGAs. *ACM Trans. Reconfigurable Technol. Syst.* **2**, 15–11532 (2009). doi:[10.1145/1534916.1534925](https://doi.org/10.1145/1534916.1534925)
14. Covington, R., Dwarkada, S., Jump, J.R., Sinclair, J.B., Madala, S.: The efficient simulation of parallel computer systems. *Int. J. Comput. Simul.* **1**(1), 31–58 (1991)
15. ENCORE embedded processor. URL: http://groups.inf.ed.ac.uk/pasta/hw_encore.html
16. Hardavellas, N., Somogyi, S., Wenisch, T.F., Wunderlich, R.E., Chen, S., Kim, J., Falsafi, B., Hoe, J.C., Nowatzyk, A.G.: SimFlex: a fast, accurate, flexible full-system simulation framework for performance evaluation of server architecture. *SIGMETRICS Perform. Eval. Rev.* **31**, 31–34 (2004). doi:[10.1145/1054907.1054914](https://doi.org/10.1145/1054907.1054914)
17. Hughes, C., Pai, V., Ranganathan, P., Adve, S.: RSIM: simulating shared-memory multiprocessors with ILP processors. *Computer* (2002)
18. Kanaujia, S., Papazian, I.E., Chamberlain, J., Baxter, J.: FastMP: a multi-core simulation methodology. In: Proceedings of the Workshop on Modeling, Benchmarking and Simulation (MoBS 2006), Boston, Massachusetts (2006)
19. Lantz, R.: Parallel SimOS: scalability and performance for large system simulation (2007). <http://www-cs.stanford.edu>
20. Lantz, R.: Fast functional simulation with parallel Embra. In: Proceedings of the 4th Annual Workshop on Modeling, Benchmarking and Simulation (2008)
21. Magnusson, P.S., Christensson, M., Eskilson, J., Forsgren, D., Hällberg, G., Högberg, J., Larsson, F., Moestedt, A., Werner, B.: Simics: a full system simulation platform. *Computer* **35**, 50–58 (2002). doi:[10.1109/2.982916](https://doi.org/10.1109/2.982916)
22. Martin, M.M.K., Sorin, D.J., Beckmann, B.M., Marty, M.R., Xu, M., Alameldeen, A.R., Moore, K.E., Hill, M.D., Wood, D.A.: Multifacet's general execution-driven multiprocessor simulator (GEMS) tool-set. *SIGARCH Comput. Archit. News* **33**, 92–99 (2005). doi:[10.1145/1105734.1105747](https://doi.org/10.1145/1105734.1105747)
23. Miller, J.E.M., Kasture, H., Kurian, G., Gruenwald III, C., Beckmann, N., Celio, C., Eastep, J., Agarwal, A.: Graphite: a distributed parallel simulator for multicores. In: The 16th IEEE International Symposium on High-Performance Computer Architecture (HPCA) (2010)

24. Monchiero, M., Ahn, J.H., Falcón, A., Ortega, D., Faraboschi, P.: How to simulate 1000 cores. *SIGARCH Comput. Archit. News* **37**, 10–19 (2009). doi:[10.1145/1577129.1577133](https://doi.org/10.1145/1577129.1577133)
25. Mukherjee, S.S., Reinhardt, S.K., Falsafi, B., Litzkow, M., Hill, M.D., Wood, D.A., Huss-Lederman, S., Larus, J.R.: Wisconsin Wind Tunnel II: a fast, portable parallel architecture simulator. *IEEE Concurr.* **8**, 12–20 (2000). doi:[10.1109/4434.895100](https://doi.org/10.1109/4434.895100)
26. PCSX2. URL: <http://pcsx2.net/>
27. Penry, D.A., Fay, D., Hodgdon, D., Wells, R., Schelle, G., August, D.I., Connors, D.: Exploiting parallelism and structure to accelerate the simulation of chip multi-processors. In: in Proceedings of the Twelfth International Symposium on High-Performance Computer Architecture, pp. 29–40 (2006)
28. Reinhardt, S.K., Hill, M.D., Larus, J.R., Lebeck, A.R., Lewis, J.C., Wood, D.A.: The wisconsin wind tunnel: virtual prototyping of parallel computers. In: Proceedings of the 1993 ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, SIGMETRICS '93, pp. 48–60. ACM, New York, NY, USA (1993). doi:[10.1145/166955.166979](https://doi.org/10.1145/166955.166979)
29. Sui, X., Wu, J., Yin, W., Zhou, D., Gong, Z.: MALsim: a functional-level parallel simulation platform for CMPs. In: 2nd International Conference on Computer Engineering and Technology (ICCET) 2010, vol. 2, p. V2, IEEE (2010)
30. Synopsys Inc.: ARCompact instruction set architecture. URL: <http://www.synopsys.com>
31. Tan, Z., Waterman, A., Avizienis, R., Lee, Y., Cook, H., Patterson, D., Asanović K.: RAMP gold: an FPGA-based architecture simulator for multiprocessors. In: Proceedings of the 47th Design Automation Conference, DAC '10, pp. 463–468. ACM, New York, NY, USA (2010). doi:[10.1145/1837274.1837390](https://doi.org/10.1145/1837274.1837390)
32. Tan, Z., Waterman, A., Cook, H., Bird, S., Asanović, K., Patterson, D.: A case for FAME: FPGA architecture model execution. In: Proceedings of the 37th Annual International Symposium on Computer Architecture, ISCA '10, pp. 290–301. ACM, New York, NY, USA (2010). doi:[10.1145/1815961.1815999](https://doi.org/10.1145/1815961.1815999)
33. The Embedded Microprocessor Benchmark Consortium: MultiBench 1.0 Multicore Benchmark Software (02 February 2010)
34. Wang, K., Zhang, Y., Wang, H., Shen, X.: Parallelization of IBM mambo system simulator in functional modes. *ACM SIGOPS Oper. Syst. Rev.* **42**(1), 71–76 (2008)
35. Wawrzyniek, J., Patterson, D., Oskin, M., Lu, S.L., Kozyrakis, C., Hoe, J.C., Chiou, D., Asanovic, K.: RAMP: research accelerator for multiple processors. *IEEE Micro* **27**, 46–57 (2007). doi:[10.1109/MM.2007.39](https://doi.org/10.1109/MM.2007.39)
36. Wentzlaff, D., Agarwal, A.: Constructing virtual architectures on a tiled processor. In: Proceedings of the International Symposium on Code Generation and Optimization, CGO '06, pp. 173–184. IEEE Computer Society, Washington, DC, USA (2006). doi:[10.1109/CGO.2006.11](https://doi.org/10.1109/CGO.2006.11)
37. Woo, S.C., Ohara, M., Torrie, E., Singh, J.P., Gupta, A.: The SPLASH-2 programs: characterization and methodological considerations. In: Proceedings of the 22nd Annual International Symposium on Computer Architecture, ISCA '95, pp. 24–36. ACM, New York, NY, USA (1995). doi:[10.1145/223982.223990](https://doi.org/10.1145/223982.223990)
38. Zheng, G., Kakulapati, G., Kalé, L.V.: BigSim: a parallel simulator for performance prediction of extremely large parallel machines. In: Parallel and Distributed Processing Symposium, International, vol. 1, p. 78b (2004). doi:[10.1109/IPDPS.2004.1303013](https://doi.org/10.1109/IPDPS.2004.1303013)
39. Zhong, R., Zhu, Y., Chen, W., Lin, M., Wong, W.F.: An inter-core communication enabled multi-core simulator based on simplescalar. *Advanced Information Networking and Applications Workshops, International Conference* **1**, 758–763 (2007). doi:[10.1109/AINAW.2007.87](https://doi.org/10.1109/AINAW.2007.87)
40. Zhu, X., Malik, S.: Using a communication architecture specification in an application-driven retargetable prototyping platform for multiprocessing. In: Proceedings of the Conference on Design, Automation and Test in Europe, DATE '04, vol. 2, pp. 21–244. IEEE Computer Society, Washington, DC, USA (2004)
41. Zhu, X., Wu, J., Sui, X., Yin, W., Wang, Q., Gong, Z.: PCAsim: a parallel cycle accurate simulation platform for CMPs. In: Proceedings of the 2010 International Conference on Computer Design and Applications (ICCD), pp. V1-597–V1-601 (2010)