Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2010-058 December 8, 2010

Heracles: Fully Synthesizable Parameterized MIPS-Based Multicore System

Michel Kinsy and Michael Pellauer

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Heracles: Fully Synthesizable Parameterized MIPS-Based Multicore System

Michel Kinsy
Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139
mkinsy@mit.edu

Michael Pellauer
Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139
pellauer@csail.mit.edu

ABSTRACT

Heracles is an open-source complete multicore system written in Verilog. It is fully parameterized and can be reconfigured and synthesized into different topologies and sizes. Each processing node has a 7-stage pipeline, fully bypassed, microprocessor running the MIPS-III ISA, a 4-stage input-buffer, virtual-channel router, and a local variable-size shared memory. Our design is highly modular with clear interfaces between the core, the memory hierarchy, and the on-chip network. In the baseline design, the microprocessor is attached to two caches, one instruction cache and one data cache, which are oblivious to the global memory organization. The memory system in *Heracles* can be configured as one single global shared memory (SM), or distributed shared memory (DSM), or any combination thereof. Each core is connected to the rest of the network of processors by a parameterized, realistic, wormhole router. We show different topology configurations of the system, and their synthesis results on the Xilinx Virtex-5 LX330T FPGA board. We also provide a small MIPS cross-compiler toolchain to assist in developing software for *Heracles*.

Categories and Subject Descriptors

C.1.2 [**Computer Systems Organization**]: Processor Architecture - Single-instruction-stream, multiple-data-stream processors (SIMD); B.5.1 [**Hardware**]: Register-Transfer-Level Implementation- Design.

General Terms

Design, Experimentation, Performance

Keywords

Multicore Architecture Design, FPGA, Shared-Memory, Distributed Shared Memory, Network-on-Chip, RISC, MIPS, Virtual Channel, Wormhole Router, NoC Routing Algorithm.

1. INTRODUCTION

Multicore architectures have become mainstream computing platforms. These systems typically consist of processing elements (PEs or cores), a memory subsystem, and an infrastructure for inter-core communications. Traditionally, buses have been used in establishing communications between cores, but because of the increasing complexity of these designs and the lack of scalability of wired connections between cores, network-on-chip (NoC) architectures have been introduced as an effective data communication infrastructure [7, 11]. It has been shown that the overall performance of multicore systems is often defined by their communication limits in terms of bandwidth, speed and concurrency [15, 4, 27], and not by the individual computation power of the cores. Therefore, simple reduced instruction set computing (RISC) cores are often used in these architectures.

In this paper, we present a new open-source FPGA-based system for designing multicore architectures called *Heracles*. A complete multicore system written in Verilog, fully parameterized, that can be reconfigured into different topologies and sizes. The main contribution of our work is the fact that *Heracles* is designed with a high degree of modularity to support exploration of future multicore processors of different topologies, routing schemes, processing elements or cores, and memory system organizations. Figure 1 shows the top level view of *Heracles* multicore system arranged in 2D-mesh topology, and Figure 2 shows two different views of the network switch local to a node.

There has been a large body of work on implementing multicore architectures on FPGAs. In contrast, there seems to be very little on complete, modular, multicore system, with reconfigurable network topology, where processing core, memory system, and on-chip network are fully self-containing.

Heracles presents designers with a global and complete view of the inner workings of a multiprocessor machine cycle-by-cycle from instruction fetches at the microprocessor core at each node to the *flit* arbitration at the routers, with RTL level correctness. A flit is the smallest unit of information recognized by the flow control method [8]. This enables the designer to explore different implementation approaches: core micro-architecture, levels of caches, cache sizes, routing algorithm, router micro-architecture, distributed or shared memory, network interface, and to quickly evaluate their impact on the overall system performance.

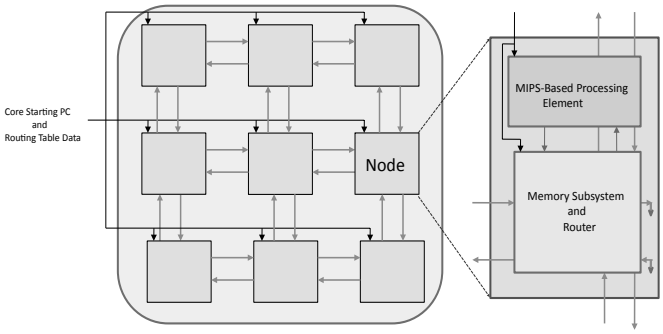

Figure 1: 2D-Mesh Topology Heracles Architecture.

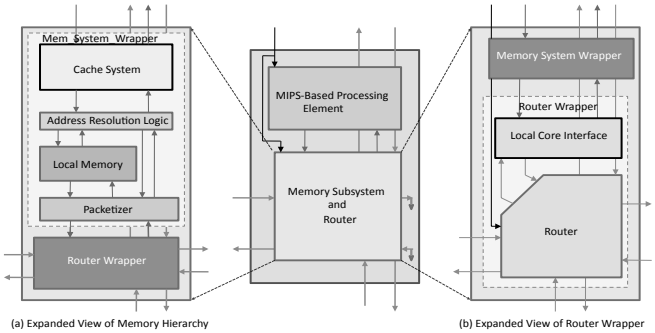

Figure 2: Network Switch Expanded Views.

Section 2 describes an integer-based 7-stage MIPS processing element (PE), and its usage in forming a node in the network. Section 3 presents our structure for supporting an arbitrary memory organization, and details about the network interface. Section 4 deals with the router micro-architecture and supports for various routing algorithms. Section 5 shows different *Heracles* topologies, their FPGA utilization, and performance analysis results. Related work is summarized in Section 6. Section 7 concludes the paper.

2. PROCESSING ELEMENT MODULE

The processing element in *Heracles* consists of an integer-based 7-stage MIPS Core. MIPS (Microprocessor without Interlocked Pipeline Stages) is a register based RISC architecture widely used in commercial products and for teaching purposes [20]. Our implementation is a standard Verilog implementation of the micro-architecture described by Patterson and Hennessy [20], with some modifications for FPGAs. For example, the adoption of a 7-stage pipeline, due to block RAM access time on the FPGA.

Figure 3 shows the core architecture, a 7-stage pipeline architecture, fully bypassed, with no branch predictor or branch delay slot, running MIPS-III instruction set architecture (ISA) without floating point. Instruction and data caches are implemented using block RAMs, when instruction fetch and data memory access take two cycles. Instruction address is issued at *I-Fetch 1* stage and on cache hit, the actual instruction appears in the *I-Fetch 2* stage. Instruction decode and register read stage and the execution stage remains functionally the same as described [20]. Stall and bypass signals are modified to support the extended pipeline. Data memory (cache) is done over *D-Memory 1* and *D-Memory 2* stages. For a read, the memory address is presented to the cache in the *D-Memory 1* stage and the data on a cache hit appears in the *D-Memory 2* stage. On memory write, we also check in the *D-Memory 2* stage that there is a cache hit before continuing execution. Instructions are issued and executed in-order, the data memory accesses are also in-order.

	Used	Available	Utilization
Registers	1,635	207,360	under 1%
Lookup Tables	2,529	207,360	1%

Critical Path (ns)	6.151
Clock Rate (MHz)	162.564

Table 1: Processing Element Synthesis Results

The core is synthesized using Xilinx ISE Design Suite 11.5, with Virtex-5 LX330T as the targeted board. Table 1 gives a brief summary of the synthesis results and the clocking speed of the design. As shown in Table 1, our 7-stage pipeline core architecture runs at 162.5 MHz, and has an FPGA resource utilization of around 1% on the Virtex-5 LX330T. Due to the modular design of *Heracles*, any core with the same memory interface can be plugged into the system.

3. MEMORY SYSTEM ORGANIZATION

The memory system in *Heracles* is completely parameterized, and can be setup in various ways, independent of the rest of the system. Figure 4 shows the block composition of the memory system at a node. The key components are the

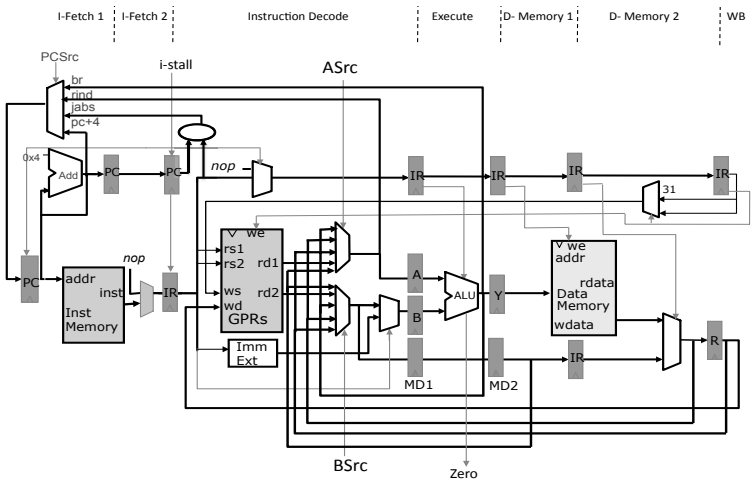

Figure 3: Integer-based 7-stage MIPS processing element (PE).

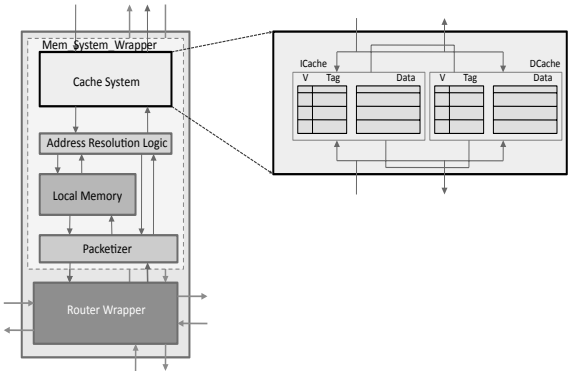

Figure 4: Local View of Memory Subsystem at a Core.

cache system, the local memory, and the network interface.

3.1 Cache System

In *Heracles*, we implement 1-level cache system composed of a direct-mapped instruction cache and a direct-mapped data cache, which can be extended to more cache levels. Each cache can be independently configured. The *INDEX_BITS* parameter controls the number of blocks or cache-line in the cache. The *OFFSET_BITS* parameter determines the cache block size.

The direct-mapped cache is implemented using block RAM, where on a hit, the data appears on the output port in the following cycle. Since block RAMs on the FPGA are constrained resources, we also implement a direct-mapped cache using registers and lookup tables, but at a high FPGA resource cost.

The cache system, like the core, is oblivious to the system level memory organization and network topology. This decoupling is archived through the *Address Resolution Logic*, which sits outside the cache system and interacts with the rest of the memory structure, as shown in figure 4.

3.2 Local Memory Distribution

The memory system in *Heracles* is constructed to allow different memory space configurations. The local memory is parameterized and has two very important attributes: its size can be changed on a per core-basis, and it can service a variable number of caches at round-robin.

For a Shared Memory (SM) implementation, where all processors share a single large memory block, the local memory size is simply set to zero at all nodes except one. At the nodes with no local memory, the *Address Resolution Logic* directly sends all cache system traffic into the network, where it crosses to the target node.

In Distributed Shared Memory (DSM), where each processing element has its own private memory. In this scheme, local memory size can be set the same across all the nodes or with different values. For example, in a mesh network, our experiments show that for a large class of routing algorithms locating larger blocks of memory at the center nodes, can improve network congestion. The *LOCAL_ADDR_BITS* parameter is used to set the size of the local memory size.

The fact that the local memory is parameterized to handle requests from a variable number of caches allows us to present to the local memory the traffic coming into the node from other cores through the network , just as another cache communication. This illusion is created through the network packetizer. Local memory can also be viewed as a memory controller. *Heracles*, at the moment, provides no cache coherence protocol. However the design of the system is set up to support a cache coherence scheme, if one is implemented.

3.3 Network Interface

The *Address Resolution Logic* works with the *Packetizer* module to get the caches and the local memory to interact with the rest of the system. All cache traffic goes through the *Address Resolution Logic*, which determines if a request can be served at the local memory, or if the request needs to be sent over the network. In *Heracles*, an address contains two fields, where the lower order bits represent the real address, and the the higher order bits identify the home core

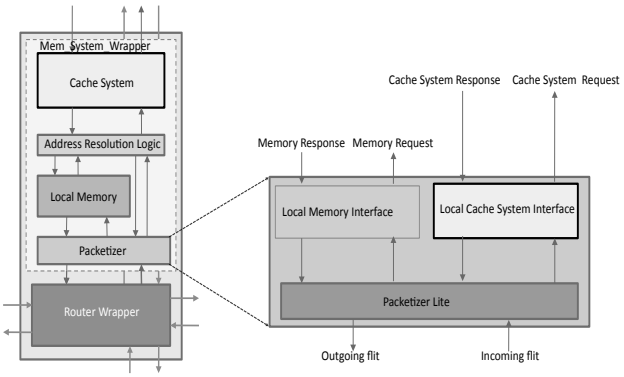

Figure 5: Network Interface Packetizer.

	Used	Available	Utilization
Registers	2,695	207,360	under 1%
Lookup Tables	5,562	207,360	2%
Block RAM/FIFO	75	324	23%

Critical Path (ns)	6.471
Clock Rate (MHz)	155.825

Table 2: Processing Element with Caches and Memory Synthesis Results

for that particular address. These two fields are automatically identified based on the *LOCAL_ADDR_BITS* and the *ADDRESS_BITS* parameters. If the home core of an address is not the core that generated the address, the *Address Resolution Logic* forwards the request to the network, through the *Packetizer* .

Inside the *Packetizer* module, there are three submodules as shown in Figure 5. The *Local Memory Interface* uses a cache-like protocol to interact with the local memory. In our baseline design, the *Local Memory Interface* simply acts as a third cache on the local memory side. The *Local Cache System Interface* uses a memory-like protocol to interact with the cache system like a second larger memory lock. The *Packetizer Lite* is responsible for converting data traffic, such as a load, coming from the local memory and the cache system into packets or flits that can be routed inside the Network-on-chip (NoC), and for reconstructing packets or flits into data traffic at the opposite side when exiting the NoC. The *Packetizer Lite* directly connects to the network router.

Table 2 gives a brief summary of the synthesis results and the clocking speed of our 7-stage pipeline MIPS core with 2 caches (I-Cache and D-Cache), 2KB each, and 262KB of local memory.

4. ROUTER ARCHITECTURE

To provide scalability, *Heracles* uses network-on-chip (NoC) architecture for its data communication infrastructure. An NoC architecture is defined by its topology (the physical organization of nodes in the network), its flow control mechanism (which establishes the data formating, the switching protocol and the buffer allocation), and its routing algorithm (which determines the path selected by a packet to reach its destination under a given application). This section discusses the router micro-architecture, and its support for different network topologies and routing algorithms.

4.1 Router Micro-Architecture

Figure 6 illustrates the virtual-channel router used in *Heracles*. The router fairly conforms, in its architecture and operation, to conventional virtual-channel routers [8, 18, 21]. It has some input buffers to store flits while they are waiting to be routed to the next hop in the network. The routing operation takes four steps or phases, namely routing (RC), virtual-channel allocation (VA), switch allocation (SA), and switch traversal (ST), where each phase corresponds to a pipeline stage in our router. When a head flit (the first flit of a packet) arrives at an input channel, the router stores the flit in the buffer for the allocated virtual channel and determines the next hop for the packet (RC phase). Given the next hop, the router then allocates a virtual channel in the next hop (VA phase). Finally, the flit competes for a switch (SA phase), if the next hop can accept the flit, and moves to the output port (ST phase).

The switch allocation (SA) stage is the critical path in our router design, due to the complexity of the arbiter. During the SA stage, the arbiter grants switch traversal to all input ports requesting output ports for which they have priority. If an input port is requesting an output port, and the priority holder on that outgoing port is either idle or requesting a different output port, it has to compete with all other input ports requesting the same output port. The arbiter is also responsible for adjusting priorities to promote fairness and avoid starvation. The synthesis of the router shows a delay of 14.016 nanosecond, with 24 levels of logic on the critical path.

4.2 Route Computation and Virtual Channel Allocation

Algorithms used to compute routes in network-on-chip (NoC) architectures, generally fall under two categories: *oblivious* and *dynamic* [19]. The router implemented in *Heracles* primarily supports *oblivious* routing algorithms using either fixed logic or routing table. Fixed logic is provided for dimension order routing (DOR) algorithms [24], which are vastly popular and have many desirable properties. For example, they generate deadlock-free routes in mesh or hypercube topologies [6, 23]. Either using *XY-Ordered Routing* or *YX-Ordered Routing*, each packet is routed along one dimension in its first phase followed by the other dimension. On the other hand, table-based routing provides greater programmability and flexibility, since routes can be pre-computed and stored in the routing tables before execution. Table-based routing supports for both *minimal* and *non-minimal* routing algorithms. In this routing scheme, at the beginning of the program, routing tables are updated; and during execution each packet has a flow ID, which is used to address the routing table to determine the packet's outgoing port. The *RT_ALG* parameter is used to select the proper routing algorithm for a given application and topology.

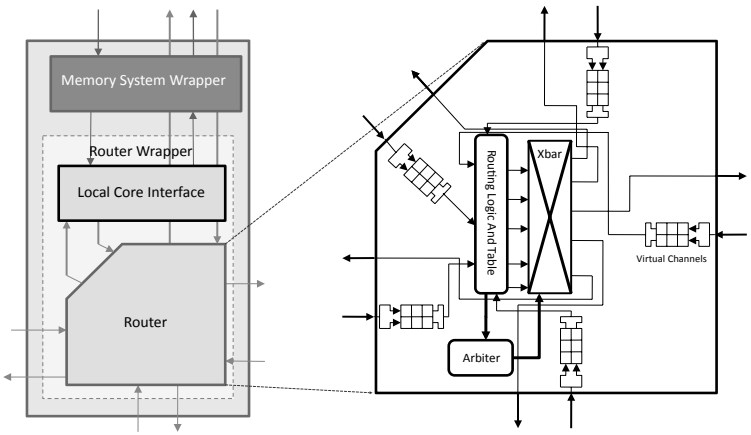

Figure 6: Router Micro-Architecture.

	Used	Available	Utilization
Registers	2,806	207,360	1%
Lookup Tables	2,058	207,360	1%

Critical Path (ns)	14.016
Clock Rate (MHz)	71.345

Table 3: Virtual-Channel Router Synthesis Results

	Used	Available	Utilization
Registers	166,726	207,360	80%
Lookup Tables	165284	207,360	80%

Critical Path (ns)	33.246
Clock Rate (MHz)	30.079

Table 4: Unbalanced Fat-Tree Topology Synthesis Results

Heracles provides support for both static and dynamic virtual channel allocation. When static allocation is used, the routing table stores the outgoing port of the packet along with the virtual channel to be used in the next node. There is no additional hardware cost for supporting static virtual channel allocation, since the entry into the table is also used during dynamic allocation. The number of virtual channels per port and their size are variable parameters (*VC_PER_PORT* and *VC_DEPTH*).

4.3 Network Topology Configuration

The parameterization of the number of input ports and output ports on the router and the table-based routing capability give *Heracles* a great amount of flexibility and the ability to metamorphose into different network topologies; for example, *k*-ary *n*-cube, 2D-mesh, 3D-mesh, hypercube, ring, or tree. A new topology is constructed by changing the *IN_PORTS*, *OUT_PORTS*, and *SWITCH_TO_SWITCH* parameters and reconnecting the routers. In the case of the 3D-mesh, the *IN_PORTS* and *OUT_PORTS* parameters are set to 2, one to connect the router to the local core and a second one to connect the router to the third dimension. Table 3 gives a brief summary of the synthesis results and the clock frequency of our virtual-channel router in 2D-meshes. It runs at 71 MHz, with the limiting factor being logic complexity of the arbiter.

For a *fat-tree* [16] topology, routers at different levels of the tree have different sizes, in terms of crossbar and arbitration logic. The root node contains the largest router, and controls the clock frequency of the system. Figure 7 shows an unbalanced *fat-tree* topology, and Table 4 shows the summary of the synthesis results and the clock frequency of the *fat-tree* multicore system.

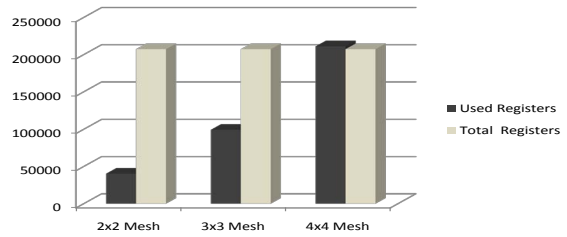

Figure 8: FPGA Register Resource Used Per Mesh Size.

5. FPGA 2D-MESH TOPOLOGY SYSTEMS

A large number of cores can be implemented on a modern FPGA. Moreover, having a simple RISC core, MIPS in our case, for the processing element (PE) allows for a good size multicore system. This section presents three different sizes, 2×2, 3×3, and 4×4, of the complete *Heracles* multicore architecture arranged in 2D-mesh topology. Figure 1 shows the 3×3 mesh topology. Table 5 summarizes the key architectural characterastics of the multicore system. The system is running at 71 MHz, which is the clock frequency of the router, regardless of the size of the mesh. The whole system speed will increase if a less complex arbitration scheme is adopted.

In our experiments, using a mesh topology, we are unable to fit more than 16 cores on the Virtex-5 LX330T FPGA board. In the 2×2 configuration, the local shared memory is set to 260 KB per core. Where for the 3×3 configuration the size of the local shared memory is reduced to 65KB per core, due to limited FPGA block RAM. The local memory in the 4×4 configuration is set to 32KB. Figure 8 shows the

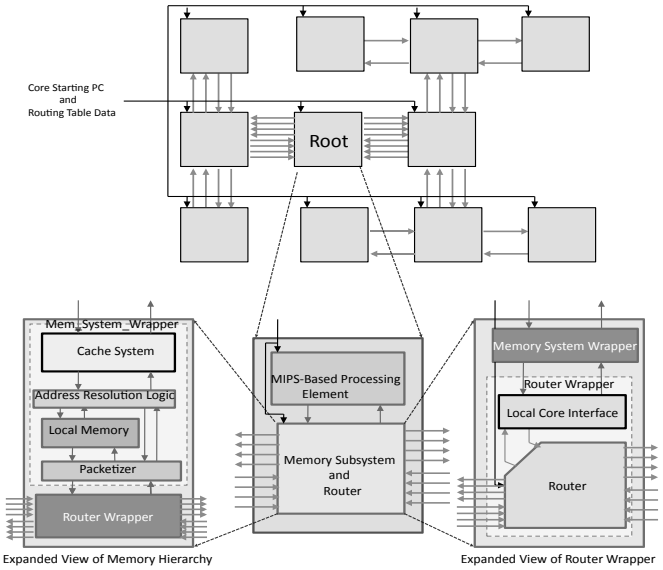

Figure 7: Unbalanced Fat-Tree Topology with Expanded Views of the Root.

	Heracles
Core	
ISA	32-Bit MIPS
Multiply/Divide	Software
Floating Point	Software
Pipeline Stages	7
Bypassing	Full
Branch policy	Always non-Taken
Outstanding memory requests	1
Address Translation	None
Level 1 Instruction/Data Caches	
Associativity	Direct
Size	16KB
Outstanding Misses	1
On-Chip Network	
Topology	2D-Mesh
Routing Policy	DOR and Table-based
Virtual Channels	2
Buffers per channel	8

Table 5: 2D-mesh *Heracles* Multicore Architecture Specification

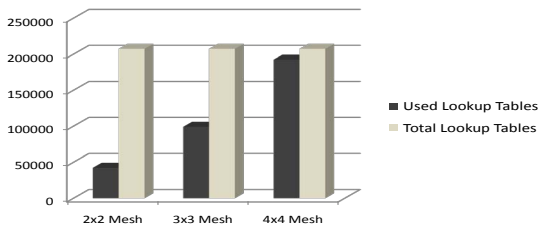

Figure 9: FPGA Lookup Table Resource Used Per Mesh Size.

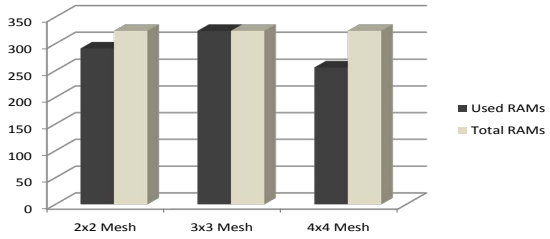

Figure 10: FPGA RAM/FIFO Resource Used Per Mesh Size.

usage of FPGA registers per mesh size. Figure 9 shows the usage of FPGA lookup tables per mesh size. Finally, Figure 10 shows the usage of FPGA RAMs per mesh size, based on the local memory configuration above.

We are releasing *Heracles* with a small open-source software toolchain to assist in developing software for the system. Figure 11 shows the basic tool flow for compiling a C program into the compatible MIPS instruction code that can be executed on the system. This toolchain is built around a MIPS cross-compiler. When *mips-gcc* is executed on a C program an assembly file is generated. GNU C version 3.2.1. is used. The assembly code is then run through isa-checker, the checker's role is to:

- remove all memory space primitives
- replace all pseudo-instructions
- check for floating point instructions

Also at this stage, a small kernel-like assembly code is added to the application assembly code for workload distribution. The assembler is then called by executing *mips-as* on the checked assembly code. The object file is then disassembled using the *mips-objdump* command. Finally, the constructor script is called to transform the dump file into Verilog memory file format.

Here we present a brief analysis of a simple *Fibonacci* number calculation program, to highlight the impact of number of

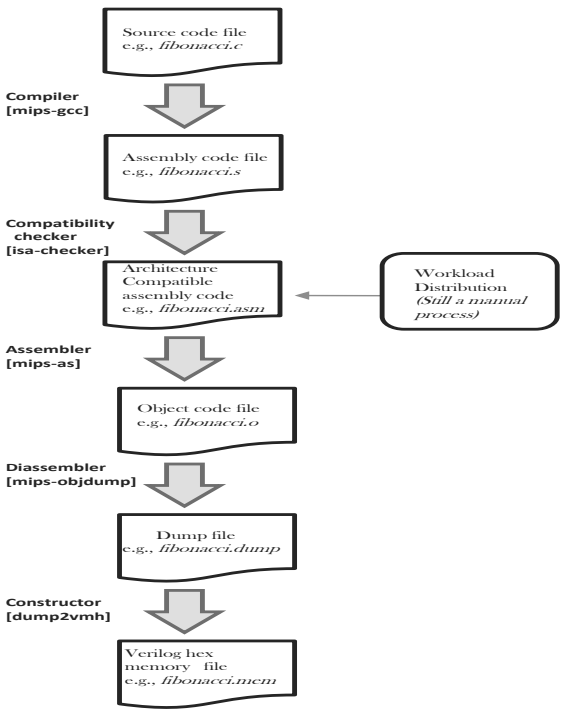

Figure 11: Software Toolchain Flow

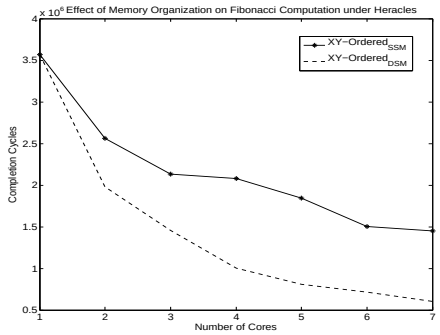

Figure 12: Fibonacci: Effect of Memory Organization on Performance in 2D-Mesh *Heracles*

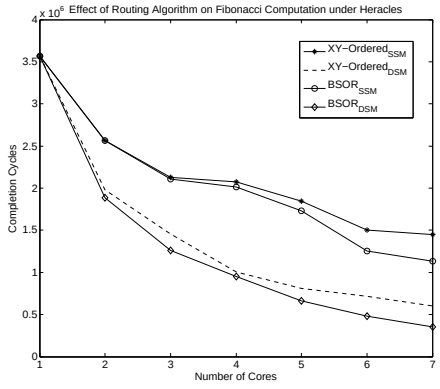

Figure 13: Fibonacci: Effect of Routing Algorithm on Performance in 2D-Mesh *Heracles*

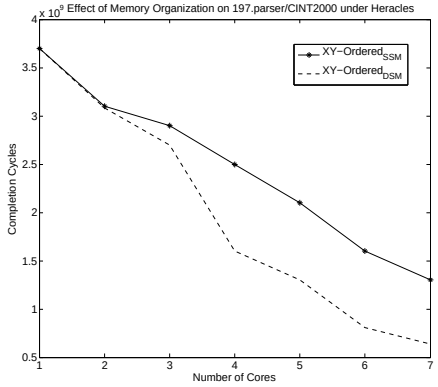

Figure 14: 197.parser: Effect of Memory Organization on Performance in 2D-Mesh *Heracles*

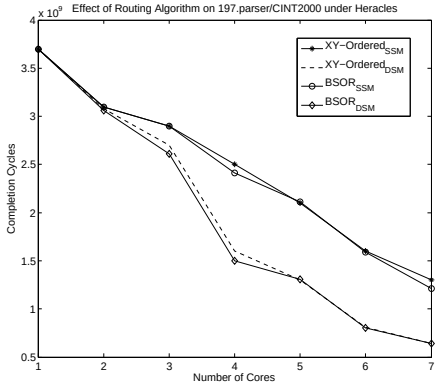

Figure 15: 256.bzip2: Effect of Memory Organization on Performance in 2D-Mesh *Heracles*

cores, memory system organization, and routing scheme on performance in *Heracles*. We also examine the performance of two SPEC CINT2000 benchmarks, namely, 197.parser and 256.bzip2 on *Heracles*. We modify and parallelize these benchmarks to fit into our evaluation framework. For the 197.parser benchmark, we identify three functional units: file reading and parameters setting as one unit, actual parsing another unit, and error reporting as the third unit. When there are more than three cores, all additional cores are used in the parsing unit. Similarly, 256.bzip2 is divided into three functional units: file reading and cyclic redundancy check, compression, and output file writing. The compression unit exhibits a high degree of data-parallelism, therefore we apply all additional cores to this unit for core count greater than three.

Figure 12 shows that with single shared-memory, using *XY-Ordered* routing, increasing the number of cores does very little to improve performance. Because the single shared-memory is the main bottleneck in the system. Figure 13 shows the direct impact of the routing algorithm on the overall system performance, by comparing completion cycles of *XY-Ordered* routing and BSOR [13]. BSOR, which stands for Bandwidth-Sensitive Oblivious Routing, is a table-based routing algorithm that minimizes the maximum channel load (MCL) or maximum traffic across all network links in the effort to maximize application throughput.

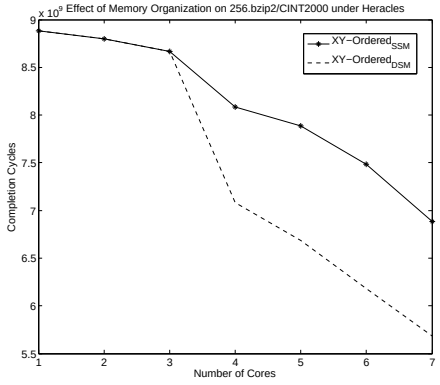

Figure 16: 197.parser: Effect of Routing Algorithm on Performance in 2D-Mesh *Heracles*

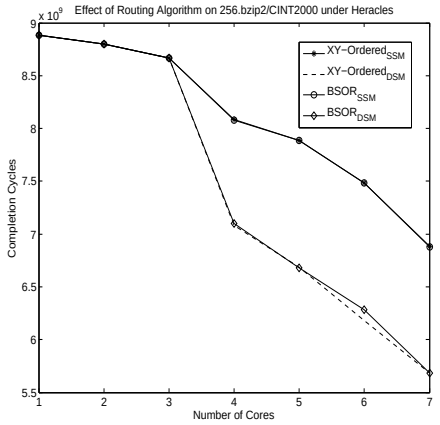

Figure 17: 256.bzip2: Effect of Routing Algorithm on Performance in 2D-Mesh *Heracles*

Figures 14 and 15 show 197.parser and 256.bzip2 benchmarks under single shared-memory (SSM) and distributed shared-memory (DSM), using *XY-Ordered* routing. Increasing the number of cores does improve performance for both benchmarks; it also exposes the memory bottleneck encountered in the single shared-memory scheme. Routing algorithm has little or no effect on the performance of these benchmarks, as shown in Figures 16 and 17, because of the traffic patterns in these applications.

Heracles Verilog files and software toolchain for building MIPS code to run on the system can be found at

`http://web.mit.edu/mkinsy/Public/Heracles`

6. RELATED WORK

Implementation of multicore architecture on FPGA has been the subject of several research projects. In [9] Del Valle *et al* present an FPGA-based emulation framework for multiprocessor system-on-chip (MPSoC) architectures. LEON3 [1], a synthesizable VHDL model of a 32-bit processor compliant with the SPARC V8 architecture, has been used in implementing multiprocessor systems on FPGA. Andersson *et al* [2], for example, use the LEON4FT microprocessor to build their Next Generation Multipurpose Microprocessor (NGMP) architecture, which is prototyped on the Xilinx XC5VFX130T FPGA board. However, the LEON architecture is fairly complex, and difficult to instantiate more than two or three on a medium size FPGA. Clack *et al* [5] investigate the use of FPGA, as a prototyping platform for developing multicore system applications. They use Xilinx MicroBlaze processor for the core, and a bus protocol for the inter-core communication. James-Roxby *et al* [12] shows similar FPGA design in their proposed architecture for supporting a single program multiple data model of parallel processing.

Other FPGA-based multicore architectures are more application specific. Ravindran *et al* [22] demonstrate the feasibility of FPGA-based multicore systems for high performance applications, through the implementation of IPv4 packet forwarding using Xilinx Virtex-II Pro FPGA. Wang *et al* [26] propose a multicore architecture on FPGA for large dictionary string matching. Similarly, Tumeo *et al* [25] present FPGA-based multicore shared memory for dual priority scheduling algorithm for real-time embedded systems.

Some designs focus primarly on the Network-on-chip (NoC). Lusala *et al* [17], for example, propose a scalable implementation of NoC on FPGA using torus topology. Genko *et al* [10] also present an FPGA-based flexible emulation environment for exploring different NoC features. A VHDL-based cycle accurate RTL model for evaluating power and performance of NoC architecture is presented in Banerjee *et al* [3]. Other designs make use of multiple FPGAs. The RAMP Blue project [14] has developed a set of reusable design blocks to emulate multicore architectures on FPGAs. The system consists of 768-1008 MicroBlaze cores in 64-84 Virtex-II Pro 70 FPGAs on 16-21 BEE2 boards.

7. CONCLUSION

We have presented a complete, realistic, fully parameterized, synthesizable, modular, multicore architecture. The system, called *Heracles*, uses a component-based design approach, where the processing element or core, the router and the network-on-chip, and the memory subsystem are independent building blocks, and can be used in other designs. The baseline system has a 7-stage integer-based MIPS core, a virtual-channel wormhole router, with support for both shared memory and distributed shared memory, implemented on the Xilinx Virtex-5 LX330T FPGA board. We have introduce a small software toolchain for compiling C programs onto the system.

We have shown a 2D-Mesh topology and an unbalanced *fat-tree* topology implementation of *Heracles*, to demonstrate the flexibility and the robustness of the system. *Heracles* can serve as a simulator in testing routing algorithms, flow controls, topologies, memory controller organizations, or it can be used as an accelerator when simulating a network-on-chip (NoC) by removing the MIPS cores from the design and placing only the NoC on the FPGA.

Future work will involve adding a small kernel binary code to each core on start up for handling exceptions and proper interrupts for peripheral communications. Multi-threading will also be added, and dynamic runtime workload management among the cores will be explored.

8. ACKNOWLEDGMENTS

We thank Srini Devadas, Joel Emer, Li-Shiuan Peh, Omer Kan, Myong Hyon Cho, and Noah Keegan for interesting discussions throughout the course of this work.

9. REFERENCES

[1] A. G. AB. Leon3 processor. *Available at*: http://www.gaisler.com.

[2] J. Andersson, J. Gaisler, and R. Weigand. Next generation multipurpose microprocessor. *Available at*: http://microelectronics.esa.int/ngmp/NGMP-DASIA10-Paper.pdf, 2010.

[3] N. Banerjee, P. Vellanki, and K. Chatha. A power and performance model for network-on-chip architectures. volume 2, pages 1250 – 1255 Vol.2, feb. 2004.

[4] L. Benini and G. De Micheli. Networks on chips: a new soc paradigm. *Computer*, 35(1):70–78, Jan 2002.

[5] C. R. Clack, R. Nathuji, and H.-H. S. Lee. Using an fpga as a prototyping platform for multi-core processor applications. In *WARFP-2005: Workshop on Architecture Research using FPGA Platforms*, Cambridge, MA, USA, feb. 2005.

[6] W. J. Dally and C. L. Seitz. Deadlock-Free Message Routing in Multiprocessor Interconnection Networks. *IEEE Trans. Computers*, 36(5):547–553, 1987.

[7] W. J. Dally and B. Towles. Route Packets, Not Wires: On-Chip Interconnection Networks. In *Proc. of the 38th Design Automation Conference (DAC)*, June 2001.

[8] W. J. Dally and B. Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, 2003.

[9] P. Del valle, D. Atienza, I. Magan, J. Flores, E. Perez, J. Mendias, L. Benini, and G. Micheli. A complete multi-processor system-on-chip fpga-based emulation framework. pages 140 –145, oct. 2006.

[10] N. Genko, D. Atienza, G. D. Micheli, J. M. Mendias, R. Hermida, and F. Catthoor. A complete network-on-chip emulation framework. In *DATE '05: Proceedings of the conference on Design, Automation and Test in Europe*, pages 246–251, Washington, DC, USA, 2005. IEEE Computer Society.

[11] A. Ivanov and G. D. Micheli. The Network-on-Chip Paradigm in Practice and Research. *Design & Test of Computers*, 22(5):399–403, 2005.

[12] P. James-Roxby, P. Schumacher, and C. Ross. A single program multiple data parallel processing platform for fpgas. In *FCCM '04: Proceedings of the 12th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, pages 302–303, Washington, DC, USA, 2004. IEEE Computer Society.

[13] M. Kinsy, M. H. Cho, T. Wen, E. Suh, M. van Dijk, and S. Devadas. Application-Aware Deadlock-Free Oblivious Routing. In *Proceedings of the Int'l Symposium on Computer Architecture*, June 2009.

[14] A. Krasnov, A. Schultz, J. Wawrzynek, G. Gibeling, and P.-Y. Droz. Ramp blue: A message-passing manycore system in fpgas. pages 54 –61, aug. 2007.

[15] K. Lahiri, A. Raghunathan, and S. Dey. Evaluation of the traffic-performance characteristics of system-on-chip communication architectures. pages 29–35, 2001.

[16] C. E. Leiserson. Fat-trees: universal networks for hardware-efficient supercomputing. *IEEE Trans. Comput.*, 34(10):892–901, 1985.

[17] A. Lusala, P. Manet, B. Rousseau, and J.-D. Legat. Noc implementation in fpga using torus topology. pages 778 –781, aug. 2007.

[18] R. D. Mullins, A. F. West, and S. W. Moore. Low-latency virtual-channel routers for on-chip networks. In *Proc. of the 31st Annual Intl. Symp. on Computer Architecture (ISCA)*, pages 188–197, 2004.

[19] L. M. Ni and P. K. McKinley. A survey of wormhole routing techniques in direct networks. *Computer*, 26(2):62–76, 1993.

[20] D. Patterson and J. Hennessy. *Computer Organization and Design: The Hardware/software Interface*. Morgan Kaufmann, 2005.

[21] L.-S. Peh and W. J. Dally. A Delay Model and Speculative Architecture for Pipelined Routers. In *Proc. International Symposium on High-Performance Computer Architecture (HPCA)*, pages 255–266, Jan. 2001.

[22] K. Ravindran, N. Satish, Y. Jin, and K. Keutzer. An fpga-based soft multiprocessor system for ipv4 packet forwarding. pages 487 – 492, aug. 2005.

[23] D. Seo, A. Ali, W.-T. Lim, N. Rafique, and M. Thottethodi. Near-optimal worst-case throughput routing for two-dimensional mesh networks. In *Proc. of the 32nd Annual International Symposium on Computer Architecture (ISCA)*, pages 432–443, 2005.

[24] H. Sullivan and T. R. Bashkow. A large scale, homogeneous, fully distributed parallel machine, i. *SIGARCH Comput. Archit. News*, 5(7):105–117, 1977.

[25] A. Tumeo, M. Branca, L. Camerini, M. Ceriani, M. Monchiero, G. Palermo, F. Ferrandi, and D. Sciuto. Prototyping pipelined applications on a heterogeneous fpga multiprocessor virtual platform. In *ASP-DAC '09: Proceedings of the 2009 Asia and South Pacific Design Automation Conference*, pages 317–322, Piscataway, NJ, USA, 2009. IEEE Press.

[26] Q. Wang and V. K. Prasanna. Multi-core architecture on fpga for large dictionary string matching. In *FCCM '09: Proceedings of the 2009 17th IEEE Symposium on Field Programmable Custom Computing Machines*, pages 96–103, Washington, DC, USA, 2009. IEEE Computer Society.

[27] C. A. Zeferino, M. E. Kreutz, L. Carro, and A. A. Susin. A study on communication issues for systems-on-chip. In *SBCCI '02: Proceedings of the 15th symposium on Integrated circuits and systems design*, page 121, Washington, DC, USA, 2002. IEEE Computer Society.