OpenPiton: An Open Source Hardware Platform For Your Research

By Jonathan Balkind, Michael McKeown, Yaosheng Fu, Tri Nguyen, Yanqi Zhou, Alexey Lavrov, Mohammad Shahrad, Adi Fuchs, Samuel Payne, Xiaohua Liang, Matthew Matl, and David Wentzlaff

Abstract

Industry is building larger, more complex, manycore processors on the back of strong institutional knowledge, but academic projects face difficulties in replicating that scale. To alleviate these difficulties and to develop and share knowledge, the community needs open architecture frameworks for simulation, chip design, and software exploration that support extensibility, scalability, and configurability, alongside an established base of verification tools and supported software. In this article, we present OpenPiton, an open source framework for building scalable architecture research prototypes from one core to 500 million cores. OpenPiton is the world's first open source, general-purpose, multithreaded manycore processor, and framework. OpenPiton is highly configurable, providing a rich design space spanning a variety of hardware parameters that researchers can change. OpenPiton designs can be emulated on FPGAs, where they can run full-stack multiuser Debian Linux. OpenPiton is designed to scale to very large core fabrics, enabling researchers to measure operating system, compiler, and software scalability. The mature codebase reflects the complexity of an industrial-grade design and provides the necessary scripts to build new chips, making OpenPiton a natural choice for computer-aided design (CAD) research. OpenPiton has been validated with a 25-core chip prototype, named Piton, and is bolstered by a validation suite that has thousands of tests, providing an environment to test new hardware designs while verifying the correctness of the whole system. OpenPiton is being actively used in research both internally to Princeton and in the wider community, as well as being adopted in education, industry, and government settings.

1. INTRODUCTION

Building processors for academic research purposes can be a risky proposition. Particularly as processors have grown in size, and with the focus on multicore and manycore processors,^{17, 19, 20, 21, 14, 22, 6} the number of potential points of failure in chip fabrication has increased drastically. To combat this, the community needs well-tested, open-source, scalable frameworks that they can rely on as baselines to work from and compare against. To reduce “academic time-to-publication”, these frameworks must provide robust software tools, mature full-system software stacks, rely on industry-standard languages, and provide thorough test suites. Additionally, to support research in a broad variety of fields,

these frameworks must be highly configurable, be synthesizable to FPGA and ASIC for prototyping purposes, and provide the basis for others to tape-out (manufacture) their own, modified academic chips. Building and supporting such an infrastructure is a major undertaking which has prevented such prior designs. Our framework, OpenPiton, attacks this challenge and provides all of these features and more.

OpenPiton is the world's first open source, general-purpose, multithreaded manycore processor. OpenPiton is scalable and portable; the architecture supports addressing for up to 500-million cores, supports shared memory both within a chip and across multiple chips, and has been designed to easily enable high performance 1000+ core microprocessors and beyond. The design is implemented in **industry-standard Verilog HDL** and does not require the use of any new languages. OpenPiton enables research from the small to the large with demonstrated implementations from the slimmed-down, single-core PicoPiton, which is emulated on a \$160 Xilinx Artix 7 at 29.5MHz, up to the 25-core Piton processor which targeted a 1GHz operating point and was recently validated and thoroughly characterized.^{12, 13}

The OpenPiton platform shown in Figure 1 is a modern, tiled, manycore design consisting of a 64-bit architecture using the mature SPARC v9 ISA with P-Mesh: our scalable cache coherence protocol and network on chip (NoC). OpenPiton builds upon the industry-hardened, open-source OpenSPARC T1^{15, 1, 18} core, but sports a completely scratch-built uncore (caches, cache-coherence protocol, NoCs, NoC-based I/O bridges, etc), a new and modern simulation framework, configurable and portable FPGA scripts, a complete set of scripts enabling synthesis and implementation of ready-to-manufacture chips, and full-stack multiuser Debian Linux support. OpenPiton is available for download at <http://www.openpiton.org>.

OpenPiton has been designed as a platform to enable at-scale research. An explicit design goal of OpenPiton is that it should be easy to use by other researchers. To support this, OpenPiton provides a high degree of integration and configurability as shown in Table 1. Unlike many other designs where the pieces are provided, but it is up to the

The original version of this paper is entitled “*OpenPiton: An Open Source Manycore Research Framework*” and was published in *Proceedings of ASPLOS 2016*, Atlanta, GA, April 2–6, 2016, ACM.

Figure 1. OpenPiton Architecture. Multiple manycore chips are connected together with chipset logic and networks to build large scalable manycore systems. OpenPiton's cache coherence protocol extends off chip.

Table 1. Supported OpenPiton configuration options. Bold indicates default values. (*Associativity reduced to 2-ways at smallest size).

Component	Configurability options
Cores (per chip)	Up to 65,536
Cores (per system)	Up to 500 million
Threads per core	1/2/4
Floating-point unit	Present /Absent
Stream-processing unit	Present/ Absent
TLBs	8/ 16 /32/64 entries
L1 I-cache	8*/ 16 /32KB
L1 D-cache	4*/ 8 /16KB
L1.5 cache	Number of sets, ways (8KB, 4-way)
L2 cache (per tile)	Number of sets, ways (64KB, 4-way)
Intra-chip topologies	2D mesh , crossbar
Inter-chip topologies	2D mesh, 3D mesh, crossbar, butterfly network
Bootloading	SD/SDHC card, UART

user to compose them together, OpenPiton is designed with all of the components integrated into the same, easy-to-use, build infrastructure providing **push-button scalability**. Researchers can easily deploy OpenPiton's source code, add in modifications, and explore their novel research ideas in the setting of a fully working system. Thousands of targeted, high-coverage test cases are provided to enable researchers to innovate with a safety net that ensures functionality is maintained. OpenPiton's open source nature also makes it easy to release modifications and reproduce previous work for comparison or reuse.

Rather than simply being a platform designed by computer architects for use by computer architects, OpenPiton enables researchers in other fields including operating systems (OS), security, compilers, runtime tools, systems, and computer-aided design (CAD) tools to conduct research at-scale. In order to enable such a wide range of applications, OpenPiton is configurable and extensible. The number of cores, attached I/O, size of caches, in-core parameters, and network topology are all configurable from a command-line

option or configuration file. OpenPiton is easy to extend; the presence of a well documented core, a well documented coherence protocol, and an easy-to-interface NoC make adding research features straightforward. Research extensions to OpenPiton that have already been built include several novel memory system explorations, an Oblivious RAM controller, and a new in-core thread scheduler. The validated and mature ISA and software ecosystem support OS and compiler research. The release of OpenPiton's scripts for FPGA emulation and chip manufacture make it easy for others to port to new FPGAs or semiconductor process technologies. In particular, this enables CAD researchers who need large netlists to evaluate their algorithms at-scale.

2. THE OPENPITON PLATFORM

OpenPiton is a tiled, manycore architecture, as shown in Figure 1. It is designed to be scalable, both intra-chip and inter-chip, using the P-Mesh cache coherence system.

Intra-chip, tiles are connected via three P-Mesh networks on-chip (NoCs) in a scalable 2D mesh topology (by default). The NoC router address space supports scaling up to 256 tiles in each dimension within a single OpenPiton chip (64K cores/chip).

For inter-chip communication, the chip bridge extends the three NoCs off-chip, connecting the tile array (through the tile in the upper-left) to off-chip logic (chipset). The chipset may be implemented on an FPGA, as a standalone chip, or integrated into an OpenPiton chip.

The extension of the P-Mesh NoCs off-chip allows the seamless connection of multiple OpenPiton chips to create a larger system, as shown in Figure 1. OpenPiton's cache-coherence extends off-chip as well, enabling shared-memory across multiple chips, for the study of even larger shared-memory manycore systems.

2.1. Tile

The architecture of a tile is shown in Figure 2a. A tile consists of a core, an L1.5 cache, an L2 cache, a floating-point unit (FPU), a CPU-Cache Crossbar (CCX) arbiter, a Memory Inter-arrival Time Traffic Shaper (MITTS), and three P-Mesh NoC routers.

The L2 and L1.5 caches connect directly to all three NoC routers, and all messages entering and leaving the tile traverse these interfaces. The CCX is the crossbar interface used

Figure 2. Architecture of (a) a tile and (b) chipset.

in the OpenSPARC T1 to connect the cores, L2 cache, FPU, I/O, etc.¹ In OpenPiton, the L1.5 and FPU are connected to the core by CCX.

2.2. Core

OpenPiton uses the open-source OpenSPARC T1¹⁵ core with modifications. This core was chosen because of its industry-hardened design, multi-threaded capability, simplicity, and modest silicon area requirements. Equally important, the OpenSPARC framework has a stable code base, implements a mature ISA with compiler and OS support, and comes with a large test suite.

In the default configuration for OpenPiton, as used in Piton, the number of threads is reduced from four to two and the stream processing unit (SPU) is removed from the core to save area. The default Translation Lookaside Buffer (TLB) size is 16 entries but can be increased to 32 or 64, or decreased down to 8 entries.

Additional configuration registers were added to enable extensibility within the core. They are useful for adding additional functionality to the core which can be configured from software, for example enabling/disabling functionality, configuring different modes of operation, etc.

2.3. Cache hierarchy

OpenPiton's cache hierarchy is composed of three cache levels. Each tile in OpenPiton contains private L1 and L1.5 caches and a slice of the distributed, shared L2 cache. The data path of the cache hierarchy is shown in Figure 3.

The memory subsystem maintains cache coherence using our coherence protocol, called P-Mesh. It adheres to the memory consistency model used by the OpenSPARC T1. Coherent messages between L1.5 caches and L2 caches communicate through three NoCs, carefully designed to ensure deadlock-free operation.

L1 caches. The L1 caches are reused from the OpenSPARC T1 design with extensions for configurability. They are composed of separate L1 instruction and L1 data caches, both of which are write-through and 4-way set-associative. By default, the L1 data cache is an 8KB cache and its line size is 16-bytes. The 16KB L1 instruction cache has a 32-byte line size.

L1.5 data cache. The L1.5 (comparable to L2 caches in other processors) both transduces the OpenSPARC T1's CCX protocol to P-Mesh's NoC coherence packet formats, and acts as a write-back layer, caching stores from the write-through L1 data cache. Its parameters match the L1 data cache by default.

The L1.5 communicates to and from the core through the CCX bus, preserved from the OpenSPARC T1. When a

memory request results in a miss, the L1.5 translates and forwards it to the L2 through the NoC channels. Generally, the L1.5 issues requests on NoC1, receives data on NoC2, and writes back modified cache lines on NoC3, as shown in Figure 3.

The L1.5 is inclusive of the L1 data cache; each can be independently sized with independent eviction policies. For space and performance, the L1.5 does not cache instructions—these cache lines are bypassed directly to the L2 cache.

L2 cache. The L2 cache (comparable to a last-level L3 cache in other processors) is a distributed, write-back cache shared by all tiles. The default cache configuration is 64KB per tile and 4-way set associativity, but both the cache size and associativity are configurable. The cache line size is 64 bytes, larger than the line sizes of caches lower in the hierarchy. The integrated directory cache has 64 bits per entry, so it can precisely keep track of up to 64 sharers by default.

The L2 cache is inclusive of the private caches (L1 and L1.5). Cache line way mapping between the L1.5 and the L2 is independent and is entirely subject to the replacement policy of each cache. Since the L2 is distributed, cache lines consecutively mapped in the L1.5 are likely to be distributed across multiple L2 tiles (L2 tile referring to a portion of the distributed L2 cache in a single tile).

The L2 is the point of coherence for all cacheable memory requests. All cacheable memory operations (including atomic operations such as compare-and-swap) are ordered, and the L2 strictly follows this order when servicing requests. The L2 also keeps the instruction and data caches coherent, per the OpenSPARC T1's original design. When a line is present in a core's L1 instruction cache and is loaded as data, the L2 sends invalidations to the relevant instruction caches before servicing the load.

2.4. P-Mesh network on-chip

There are three P-Mesh NoCs in an OpenPiton chip. The NoCs provide communication between the tiles for cache coherence, I/O, memory traffic, and inter-core interrupts. They also route traffic destined for off-chip to the chip bridge. The packet format contains 29 bits of core addressability, making it scalable up to 500 million cores.

To ensure deadlock-free operation, the L1.5 cache, L2 cache, and memory controller give different priorities to different NoC channels; NoC3 has the highest priority, next is NoC2, and NoC1 has the lowest priority. Thus, NoC3 will never be blocked. In addition, all hardware components are designed such that consuming a high priority packet is never dependent on lower priority traffic.

Classes of coherence operations are mapped to NoCs based on the following rules, as depicted in Figure 3:

- NoC1 messages are initiated by requests from the private cache (L1.5) to the shared cache (L2).
- NoC2 messages are initiated by the shared cache (L2) to the private cache (L1.5) or memory controller.
- NoC3 messages are responses from the private cache (L1.5) or memory controller to the shared cache (L2).

Figure 3. OpenPiton's memory hierarchy datapath.

2.5. Chipset

The chipset, shown in Figure 2b, houses the I/O, DRAM controllers, chip bridge, P-Mesh chipset crossbar, and P-Mesh inter-chip network routers. The chip bridge demultiplexes traffic from the attached chip back into the three physical NoCs. The traffic then passes through a Packet Filter (not shown), which modifies packet destination addresses based on the memory address in the request and the set of devices on the chipset. The chipset crossbar (a modified network router) then routes the packets to their correct destination device. If the traffic is not destined for this chipset, it is passed to the inter-chip network routers, which route the traffic to another chipset according to the inter-chip routing protocol. Traffic destined for the attached chip is directed back through similar paths to the chip bridge.

Inter-chip routing. The inter-chip network router is configurable in terms of router degree, routing algorithm, buffer size, etc. This enables flexible exploration of different router configurations and network topologies. Currently, we have implemented and verified crossbar, 2D mesh, 3D mesh, and butterfly networks. Customized topologies can be explored by reconfiguring the network routers.

2.6. Configurability

OpenPiton was designed to be a configurable platform, making it useful for many applications. Table 1 shows OpenPiton's configurability options, highlighting the large design space that it offers.

PyHP for Verilog. In order to provide low effort configurability of our Verilog RTL, we make use of a Python pre-processor, the Python Hypertext Processor (PyHP).¹⁶ PyHP was originally designed for Python dynamic webpage generation and is akin to PHP. We have adapted it for use with Verilog code. Parameters can be passed into PyHP, and arbitrary Python code can be used to generate testbenches or modules. PyHP enables extensive configurability beyond what is possible with Verilog generate statements alone.

Core and cache configurability. OpenPiton's core configurability parameters are shown in Table 1. The default parameters are shown in bold. OpenPiton preserves the OpenSPARC T1's ability to modify TLB sizes (from 8 to 64, in powers of two), thread counts (from 1 to 4), and the presence or absence of the FPU and SPU. Additionally, OpenPiton's L1 data and instruction caches can be doubled or halved in size (associativity drops to 2 when reducing size).

Leveraging PyHP, OpenPiton provides parameterizable memories for simulation or FPGA emulation. In addition, custom or proprietary memories can easily be used for chip development. This parameterization enables the configurability of cache parameters. The size and associativity of the L1.5 and L2 caches are configurable, though the line size remains static.

Manycore scalability. PyHP also enables the creation of scalable meshes of cores, drastically reducing the code size and complexity in some areas adopted from the original OpenSPARC T1. OpenPiton automatically generates all

core instances and wires for connecting them from a single template instance. This reduces code complexity, improves readability, saves time when modifying the design, and makes the creation of large meshes straightforward. The creation of large two-dimensional mesh interconnects of up to 256×256 tiles is reduced to a single instantiation. The mesh can be any rectangular configuration, and the dimensions do not need to be powers of two. This was a necessary feature for the 5×5 (25-core) Piton processor.

NoC topology configurability. P-Mesh provides other NoC connection topologies than the default two-dimensional mesh used in OpenPiton. The coherence protocol only requires that messages are delivered in-order from one point to another point. Since there are no inter-node ordering requirements, the NoC can easily be swapped out for a crossbar, higher dimension router, or higher radix design. Our configurable P-Mesh router can be reconfigured to a number of topologies shown in Table 1. For intra-chip use, OpenPiton can be configured to use a crossbar, which has been tested with four and eight cores with no test regressions. Other NoC research prototypes can easily be integrated and their performance, energy, and other characteristics can be determined through RTL, gate-level simulation, or FPGA emulation.

Chipset configurability. The P-Mesh chipset crossbar is configurable in the number of ports to connect the myriad devices OpenPiton users may have. There is a single XML file where the chipset devices and their address ranges are specified, so connecting a new device needs only a Verilog instantiation and an XML entry. PyHP is used to automatically connect the necessary P-Mesh NoC connections and Packet Filters.

We have so far connected a variety of devices through P-Mesh on the chipset. These include DRAM, Ethernet, UART, SD, SDHC, VGA, PS/2 keyboards, and even the MIAOW open source GPU.² These devices are driven by the OpenPiton core and perform their own DMA where necessary, routed over the chipset crossbar.

Multi-chip scalability. Similar to the on-chip mesh, PyHP enables the generation of a network of chips starting with the instantiation of a single chip. OpenPiton provides an address space for up to 8192 chips, with 65,536 cores per chip. By using the scalable P-Mesh cache coherence mechanism built into OpenPiton, half-billion core systems can be built. This configurability enables the building of large systems to test ideas at scale.

3. VALIDATION

3.1. Platform stability

One of the benefits of OpenPiton is its stability, maturity, and active support. Much of this is inherited from the OpenSPARC T1 core, which has a stable code base and has been studied for years, allowing the code to be reviewed and bugs fixed by many people. In addition, it implements a mature, commercial, and open ISA, SPARC V9. This means there is existing full tool chain support for OpenPiton, including Debian Linux OS support, a compiler, and an assembler. SPARC is supported on a number of OSs including Debian Linux, Oracle's

Linux for SPARC,^a and OpenSolaris (and its successors). Porting the OpenSPARC T1 hypervisor required changes to fewer than 10 instructions, and a newer Debian Linux distribution was modified with open source, readily available, OpenSPARC T1-specific patches written as part of Lockbox.^{3,4}

OpenPiton provides additional stability on top of what is inherited from OpenSPARC T1. The tool flow was updated to modern tools and ported to modern Xilinx FPGAs. OpenPiton is also used extensively for research internal to Princeton. This means there is active support for OpenPiton, and the code is constantly being improved and optimized, with regular releases over the last several years. In addition, the open sourcing of OpenPiton has strengthened its stability as a community has built.

Validation. When designing large scale processors, simulation of the hardware design is a must. OpenPiton supports one open source and multiple commercial Verilog simulators, which can simulate the OpenPiton design at rates up to tens or hundreds of kilohertz. OpenPiton inherited and then extended the OpenSPARC T1's large test suite with thousands of directed assembly tests, randomized assembly test generators, and tests written in C. This includes tests for not only the core, but the memory system, I/O, cache coherence protocol, etc. Additionally, the extensions like Execution Drafting (ExecD) (Section 4.1.1) have their own test suites. When making research modifications to OpenPiton, the researcher can rely on an established test suite to ensure that their modifications did not introduce any regressions. In addition, the OpenPiton documentation details how to add new tests to validate modifications and extend the existing test suite. Researchers can also use our scripts to run large regressions in parallel (to tackle the slower individual execution), automatically produce pass/fail reports and coverage reports (as shown in Figure 4), and run synthesis to verify that synthesis-safe Verilog has been used. Our scripts support the widely used SLURM job scheduler and integrate with Jenkins for continuous integration testing.

3.2. FPGA prototyping

OpenPiton can also be emulated on FPGA, which provides the opportunity to prototype the design, emulated at tens of megahertz, to improve throughput when running our test suite or more complex code, such as an interactive operating system. OpenPiton is actively supported on three Xilinx FPGA platforms: Artix-7 (Digilent

Nexys Video), Kintex-7 (Digilent Genesys 2) and Virtex-7 (VC707 Evaluation Board). An external port is also maintained for the Zynq-7000 (ZC706 Evaluation Board). Figure 5 shows the area breakdown for a minimized “PicoPiton” core, implemented for an Artix-7 FPGA (Digilent Nexys 4 DDR).

OpenPiton designs have the same features as the Piton processor, validating the feasibility of that particular design (multicore functionality, etc.), and can include the chip bridge to connect multiple FPGAs via an FPGA Mezzanine Card (FMC) link. All of the FPGA prototypes feature a full system (chip plus chipset), using the same codebase as the chipset used to test the Piton processor.

OpenPiton on FPGA can load bare-metal programs over a serial port and can boot full stack multiuser Debian Linux from an SD/SDHC card. Booting Debian on the Genesys2 board running at 87.5MHz takes less than 4 minutes (and booting to a bash shell takes just one minute), compared to 45 minutes for the original OpenSPARC T1, which relied on a tethered MicroBlaze for its memory and I/O requests. This boot time improvement combined with our push-button FPGA synthesis and implementation scripts drastically increases productivity when testing operating system or hardware modifications.

3.3. The Princeton Piton Processor

The Piton processor prototype^{12, 13} was manufactured in March 2015 on IBM's 32 nm SOI process with a target clock frequency of 1GHz. It features 25 tiles in a 5×5 mesh on a $6\text{mm} \times 6\text{mm}$ (36mm^2) die. Each tile is two-way threaded and includes three research projects: ExecD,¹¹ CDR,⁸ and MITTS,²³ while an ORAM⁷ controller was included at the chip level. The Piton processor provides validation of OpenPiton as a research platform and shows that ideas can be taken from inception to silicon with OpenPiton.

With Piton, we also produced the first detailed power and energy characterization of an open source manycore design implemented in silicon.¹³ This included characterizing energy per instruction, NoC energy, voltage versus frequency scaling, thermal characterization, and memory system energy, among other properties. All of this was done in our lab, running on the Piton processor with the OpenPiton chipset implemented on FPGA. Performing such a characterization yielded new insights into the balance between recomputation and data movement,

a Linux for SPARC is hosted at <https://oss.oracle.com/projects/linux-sparc/>

Figure 4. Test suite coverage results by module (default OpenPiton configuration).

Figure 5. Tile area breakdown for FPGA PicoPiton.

the energy cost of differing operand values, and a confirmation of earlier results⁹ that showed that NoCs do not dominate manycore processors' power consumption. Our study also produced what we believe is the most detailed area breakdown of an open source manycore, which we reproduce in Figure 6. All characterization data from our study, as well as designs for the chip printed circuit board (PCB), are now open source at <http://www.openpiton.org>.

Figure 6. Detailed area breakdown of Piton at chip, tile, and core levels. Reproduced from McKeown et al.¹³

3.4. Synthesis and back-end support

OpenPiton provides scripts to aid in synthesis and back-end physical design for generating realistic area results or for manufacturing new chips based on OpenPiton. The scripts are identical to the ones used to tape-out the Piton processor, however the scripts have been made process agnostic and references to the specific technology used have been removed due to proprietary foundry intellectual property concerns. Directions are included with OpenPiton which describe how to port to a new foundry kit. This allows the user to download OpenPiton, link to the necessary process development kit files, and run our full tool flow to produce the chip layout for a new instance of OpenPiton. In this sense, OpenPiton is portable across process technologies and provides a complete ecosystem to implement, test, prototype, and tape-out (manufacture) research chips.

4. APPLICATIONS

Table 2 presents a taxonomy of open source processors which highlights important parameters for research. Since OpenPiton's first release in 2015, it has been used across a wide range of applications and research domains, some of which are described in this section.

Table 2. Taxonomy of differences of open source processors (table data last checked in April 2018).

Processor	Architecture	FPU	OS	MMU	HW multi-threaded	Multicore/ manycore/ GPU	Prototype core count	NoC	HDL	Back-end scripts	License
pAVR	8b AVR	×	×	×	×	No	—	×	VHDL	×	GPL v2
openMSP430	16bMSP430	×	✓	×	×	No	—	×	Verilog	×	BSD
CPU86	16b x86	×	✓	×	×	No	—	×	VHDL	×	GPL
Zet	16b x86	×	✓	×	×	No	—	×	Verilog	×	GPL v3
LatticeMico32	32b LatticeMico32	×	✓	×	×	No	—	×	Verilog	×	GPL
ZPU	32b MIPS	×	✓	×	×	No	—	×	VHDL	×	FreeBSD & GPL
SecretBlaze	32b MicroBlaze	×	×	×	×	No	—	×	VHDL	×	GPL v3
AltOr32	32b ORBIS	×	✓	×	✓	No	—	×	Verilog	×	LGPL v3
aeMB	32b MicroBlaze	×	✓	×	×	No	—	×	Verilog	×	LGPL v3
Amber	32b ARM v2a	×	✓	×	×	No	—	×	Verilog	×	LGPL
OpenRISC	32b/64b ORBIS	✓	✓	✓	×	No	—	×	Verilog	×	LGPL
MIPS32 r1	32b MIPS32 r1	×	✓	×	✓	No	—	×	Verilog	×	LGPL v3
LEON 3	32b SPARC V8	✓(\$)	✓	✓	×	SMP/AMP	—	×	VHDL	×	GPL
OpenScale	32b MicroBlaze	×	✓	×	×	Manycore	FPGA/6	✓	VHDL	×	GPL v3
XUM	32b MIPS32 r2	×	✓	×	✓	Manycore	FPGA/8	✓	Verilog	×	LGPL v3
PicoRV32	32b RISC-V	×	×	×	×	No	FPGA/1	×	Verilog	×	ISC
PULP-RI5CY	32b RISC-V	✓	✓	×	×	Manycore	Chip/9	×	SystemVerilog	×	Solderpad 0.51
PULP-Zeroriscy	32b RISC-V	×	✓	×	×	Multicore	Chip/1	×	SystemVerilog	×	Solderpad 0.51
Nyuzi GPGPU	Nyami ISA	✓	✓	✓	✓	GPGPU	FPGA	✓	SystemVerilog	×	Apache 2.0
MIAOW GPGPU	AMD Southern Islands	✓	×	×	✓	GPU	FPGA/1	✓	Verilog	×	BSD 3-Clause
OpTiMSoC	32b/64b ORBIS	✓	✓	✓	×	Manycore	FPGA/4	✓	SystemVerilog	×	MIT
Simply RISC S1	64b SPARC V9	✓	✓	✓	×	No	—	×	Verilog	×	GPL v2
BERI	64b MIPS/CHERI	✓	✓	✓	✓(BERI2)	Multicore	FPGA/4	×	Bluespec	×	BERI HW-SW
OpenSPARC T1/T2	64b SPARC V9	✓	✓	✓	✓	Multicore	Chip/8	×	Verilog	×	GPL v2
Rocket	64b RISC-V	✓	✓	✓	✓	Manycore	Chip/8	✓	Chisel	×	BSD 3-Clause
AnyCore	64b RISC-V	×	×	×	×	No	Chip/1	×	SystemVerilog	×	BSD 3-Clause
PULP-Ariane	64b RISC-V	×	✓	×	×	Manycore	Chip/1	×	SystemVerilog	×	Solderpad 0.51
BOOM	64b RISC-V	✓	✓	✓	×	Manycore	FPGA	✓	Chisel	×	BSD 3-Clause
OpenPiton	64b SPARC V9	✓	✓	✓	✓	Manycore	Chip/25	✓	Verilog	✓	BSD 3-Clause & GPL v2

* (RTOS)

4.1. Internal research case studies

Execution Drafting. Execution Drafting¹¹ (ExecD) is an energy saving microarchitectural technique for multi-threaded processors, which leverages duplicate computation. ExecD takes over the thread selection decision from the OpenSPARC T1 thread selection policy and instruments the front-end to achieve energy savings. ExecD required modifications to the OpenSPARC T1 core and thus was not as simple as plugging a standalone module into the OpenPiton system. The core microarchitecture needed to be understood, and the implementation tightly integrated with the core. Implementing ExecD in OpenPiton revealed several implementation details that had been abstracted away in simulation, such as tricky divergence conditions in the thread synchronization mechanisms. This reiterates the importance of taking research designs to implementation in an infrastructure like OpenPiton.

ExecD must be enabled by an ExecD-aware operating system. Our public Linux kernel and OpenPiton hypervisor repositories contain patches intended to add support for ExecD. These patches were developed as part of a single-semester undergraduate OS research project.

Coherence Domain Restriction. Coherence Domain Restriction⁸ (CDR) is a novel cache coherence framework designed to enable large scale shared memory with low storage and energy overhead. CDR restricts cache coherence of an application or page to a subset of cores, rather than keeping global coherence over potentially millions of cores. In order to implement it in OpenPiton, the TLB is extended with extra fields and both the L1.5 and L2 cache are modified to fit CDR into the existing cache coherence protocol. CDR is specifically designed for large scale shared memory systems such as OpenPiton. In fact, OpenPiton's million-core scalability is not feasible without CDR because of increasing directory storage overhead.

Memory Inter-arrival Time Traffic Shaper. The Memory Inter-arrival Time Traffic Shaper²³ (MITTS) enables a manycore system or an IaaS cloud system to provision memory bandwidth in the form of a memory request interarrival time distribution at a per-core or per-application basis. A runtime system configures MITTS knobs in order to optimize different metrics (e.g., throughput, fairness). MITTS sits at the egress of the L1.5 cache, monitoring the memory requests and stalling the L1.5 when it uses bandwidth outside its allocated distribution. MITTS has been integrated with OpenPiton and works on a per-core granularity, though it could be easily modified to operate per-thread.

MITTS must also be supported by the OS. Our public Linux kernel and OpenPiton hypervisor repositories contain patches for supporting the MITTS hardware. With these patches, developed as an undergraduate thesis project, Linux processes can be assigned memory inter-arrival time distributions, as they would in an IaaS environment where the customer paid for a particular distribution corresponding with their application's behavior. The OS configures the MITTS bins to correspond with each process's

allocated distribution, and MITTS enforces the distribution accordingly.

4.2. External research use

A number of external researchers have already made considerable use of OpenPiton. In a CAD context, Lerner et al.¹⁰ present a development workflow for improving processor lifetime, based on OpenPiton and the gem5 simulator, which is able to improve the design's reliability time by 4.1×.

OpenPiton has also been used in a security context as a testbed for hardware trojan detection. OpenPiton's FPGA emulation enabled Elnaggar et al.⁵ to boot full-stack Debian Linux and extract performance counter information while running SPEC benchmarks. This project moved quickly from adopting OpenPiton to an accepted publication in a matter of months, thanks in part to the full-stack OpenPiton system that can be emulated on FPGA.

Oblivious RAM (ORAM)⁷ is a memory controller designed to eliminate memory side channels. An ORAM controller was integrated into the 25-core Piton processor, providing the opportunity for secure access to off-chip DRAM. The controller was directly connected to OpenPiton's NoC, making the integration straightforward. It only required a handful of files to wrap an existing ORAM implementation, and once it was connected, its integration was verified in simulation using the OpenPiton test suite.

4.3. Educational use

We have been using OpenPiton in coursework at Princeton, in particular our senior undergraduate Computer Architecture and graduate Parallel Computation classes. A few of the resulting student projects are described here.

Core replacement. Internally, we have tested replacements for the OpenSPARC T1 core with two other open source cores. These modifications replaced the CCX interface to the L1.5 cache with shims which translate to the L1.5's interface signals. These shims require very little logic but provide the cores with fully cache-coherent memory access through P-Mesh. We are using these cores to investigate manycore processors with heterogeneous ISAs.

Multichip network topology exploration. A senior undergraduate thesis project investigated the impact of interchip network topologies for large manycore processors. Figure 7 shows multiple FPGAs connected over a high-speed serial interface, carrying standard P-Mesh packets at 9 gigabits per second. The student developed a configurable P-Mesh router for this project which is now integrated as a standard OpenPiton component.

MIAOW. A student project integrated the MIAOW open source GPU² with OpenPiton. An OpenPiton core and a MIAOW core can both fit onto a VC707 FPGA with the OpenPiton core acting as a host, in place of the Microblaze that was used in the original MIAOW release. The students added MIAOW to the chipset crossbar with a single entry in its XML configuration. Once they implemented a native P-Mesh interface to replace the original AXI-Lite interface, MIAOW could

Figure 7. Three OpenPiton FPGAs connected by 9 gigabit per second serial P-Mesh links.

directly access its data and instructions from memory without the core's assistance.

Hardware transactional memory. Another student project was the implementation of a hardware transactional memory system in OpenPiton. The students learned about the P-Mesh cache coherence protocol from the OpenPiton documentation, before modifying it, including adding extra states to the L1.5 cache, and producing a highly functional prototype in only six weeks. The OpenPiton test suite was central to verifying that existing functionality was maintained in the process.

Cache replacement policies. A number of student groups have modified the cache replacement policies of both the L1.5 and L2 caches. OpenPiton enabled them to investigate the performance and area tradeoffs of their replacement policies across multiple cache sizes and associativities in the context of a full-stack system, capable of running complex applications.

4.4. Industrial and governmental use

So far we are aware of multiple CAD vendors making use of OpenPiton internally for testing and educational purposes. These users provide extra confidence that the RTL written for OpenPiton will be well supported by industrial CAD tools, as vendors often lack large scale designs to validate the functionality of their tools. In government use, DARPA has identified OpenPiton as a benchmark for use in the POSH program.

5. FUTURE

OpenPiton has a bright future. It not only has active support from researchers at Princeton but has a vibrant external user base and development community. The OpenPiton team has run four tutorials at major conferences and numerous tutorials at interested universities and will continue to run more tutorials. The future roadmap for OpenPiton includes adding additional configurability, support for more FPGA platforms and vendors, the ability to emulate in the cloud

by using Amazon AWS F1 instances, more core types plugged into the OpenPiton infrastructure, and integration with other emerging open source hardware projects. OpenPiton has demonstrated the ability to enable research at hardware speeds, at scale, and across different areas of computing research. OpenPiton and other emerging open source hardware projects have the potential to have significant impact not only on how we conduct research and educate students, but also design chips for commercial and governmental applications.

Acknowledgments

This material is based on research sponsored by the NSF under Grants No. CNS-1823222, CCF-1217553, CCF-1453112, and CCF-1438980, AFOSR under Grant No. FA9550-14-1-0148, Air Force Research Laboratory (AFRL) and Defense Advanced Research Projects Agency (DARPA) under agreement No. FA8650-18-2-7846 and FA8650-18-2-7852 and DARPA under Grants No. N66001-14-1-4040 and HR0011-13-2-0005. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory (AFRL) and Defense Advanced Research Projects Agency (DARPA), the NSF, AFOSR, DARPA, or the U.S. Government. We thank Paul Jackson, Ting-Jung Chang, Ang Li, Fei Gao, Katie Lim, Felix Madutsa, and Kathleen Feng for their important contributions to OpenPiton. C

References

1. *OpenSPARC T1 Microarchitecture Specification*. Santa Clara, CA, 2006.
2. Balasubramanian, R., Gangadhar, V., Guo, Z., Ho, C.-H., Joseph, C., Menon, J., Drummond, M.P., Paul, R., Prasad, S., Valathol, P., Sankaralingam, K. Enabling GPGPU low-level hardware explorations with MIAOW: An open-source RTL implementation of a GPGPU. *ACM Trans. Archit. Code Optim.* 12, 2 (June 2015).
3. Bittman, D., Capelis, D., Long, D. Introducing SeaOS. In *2014 International Conference on Information Science and Applications (ICISA)*, May 2014, 1–3.
4. Capelis, D.J. *Lockbox: Helping Computers Keep Your Secrets*. Technical Report UCSC-WASP-15-02, University of California, Santa Cruz, Nov. 2015.
5. Elnaggar, R., Chakrabarty, K., Tahoori, M.B. Run-time hardware trojan detection using performance counters. In *2017 IEEE International Test Conference (ITC)*, Oct. 2017, 1–10.
6. Esmaeilzadeh, H., Blem, E., St. Amant, R., Sankaralingam, K., Burger, D. Dark silicon and the end of multicore scaling. In *Proceedings of the 38th Annual International Symposium on Computer Architecture, ISCA '11* (New York, NY, USA, 2011), ACM, 365–376.
7. Fletcher, C.W., Ren, L., Kwon, A., van Dijk, M., Devadas, S. Freecursive ORAM: [nearly] free recursion and integrity verification for position-based oblivious ram. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15* (New York, NY, USA, 2015), ACM, 103–116.
8. Fu, Y., Nguyen, T.M., Wentzlaff, D. Coherence domain restriction on large scale systems. In *Proceedings of the 48th International Symposium on Microarchitecture, MICRO-48* (New York, NY, USA, 2015), ACM, 686–698.
9. Kim, J.S., Taylor, M.B., Miller, J., Wentzlaff, D. Energy characterization of a tiled architecture processor with on-chip networks. In *Proceedings of the 2003 International Symposium on Low Power Electronics and Design, ISLPED '03* (New York, NY, USA, 2003), ACM, 4–427.
10. Lerner, S., Taskin, B. Workload-aware ASIC flow for lifetime improvement of multi-core IoT processors. In *2017 18th International Symposium on Quality Electronic Design (ISQED)*, March 2017, 379–384.
11. McKeown, M., Balkind, J., Wentzlaff, D. Execution drafting: Energy efficiency through computation deduplication. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Dec. 2014, 432–444.

12. McKeown, M., Fu, Y., Nguyen, T., Zhou, Y., Balkind, J., Lavrov, A., Shahrad, M., Payne, S., Wentzlaff, D. Piton: A manycore processor for multitenant clouds. *IEEE Micro* 37, (2) (Mar. 2017), 70–80.
13. McKeown, M., Lavrov, A., Shahrad, M., Jackson, P., Fu, Y., Balkind, J., Nguyen, T., Lim, K., Zhou, Y., Wentzlaff, D. Power and energy characterization of an open source 25-core manycore processor. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2018.
14. Miller, B., Brasili, D., Kiszely, T., Kuhn, R., Mehrotra, R., Salvi, M., Kulkarni, M., Varadharajan, A., Yin, S.-H., Lin, W., Hughes, A., Stysiac, B., Kandadi, V., Pragasathi, I., Hartman, D., Carlson, D., Yalala, V., Xanthopoulos, T., Meninger, S., Crain, E., Spaeth, M., Aina, A., Balasubramanian, S., Vulih, J., Tiwary, P., Lin, D., Kessler, R., Fishbein, B., Jain, A. A 32-core RISC microprocessor with network accelerators, power management and testability features. In *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2012, 58–60.
15. Oracle. OpenSPARC T1. <http://www.oracle.com/technetwork/systems/opensparc/opensparc-t1-page-1444609.html>.
16. PyHP. PyHP Official Home Page. <http://pyhp.sourceforge.net>.
17. Seiler, L., Carmean, D., Sprangle, E., Forsyth, T., Abrash, M., Dubey, P., Junkins, S., Lake, A., Sugerman, J., Cavin, R., Espasa, R., Grochowski, E., Juan, T., Hanrahan, P. Larrabee: A many-core $\times 86$ architecture for visual computing. *ACM Trans. Graph.* 27, (3) (Aug. 2008), 18:1–18:15.
18. Szefer, J., Zhang, W., Chen, Y.-Y., Champagne, D., Chan, K., Li, W., Cheung, R., Lee, R. Rapid single-chip secure processor prototyping on the OpenSPARC FPGA platform. In *2011 22nd IEEE International Symposium on Rapid System Prototyping (RSP)*, May 2011, 38–44.
19. Vangal, S.R., Howard, J., Ruhl, G., Dighe, S., Wilson, H., Tschanz, J., Finan, D., Singh, A., Jacob, T., Jain, S., et al. An 80-tile sub-100-w teraops processor in 65-nm CMOS. *IEEE J. Solid-State Circuits* 43, 1 (2008), 29–41.
20. Wentzlaff, D., Griffin, P., Hoffmann, H., Bao, L., Edwards, B., Ramey, C., Mattina, M., Miao, C.-C., Brown III, J.F., Agarwal, A. On-chip interconnection architecture of the Tile Processor. *IEEE Micro* 27, (5) (Sept. 2007), 15–31.
21. Wentzlaff, D., Jackson, C.J., Griffin, P., Agarwal, A. Configurable fine-grain protection for multicore processor virtualization. In *Proceedings of the Annual International Symposium on Computer Architecture* (Washington, DC, USA, 2012), 464–475.
22. Woo, D.H., Lee, H.-H.S. Extending Amdahl's law for energy-efficient computing in the many-core era. *Computer* 12 (2008), 24–31.
23. Zhou, Y., Wentzlaff, D. Mitts: Memory inter-arrival time traffic shaping. In *Proceedings of the 43rd International Symposium on Computer Architecture, ISCA '16* (Piscataway, NJ, USA, 2016), IEEE Press, 532–544.

Jonathan Balkind, Michael McKeown, Tri Nguyen, Alexey Lavrov, Mohammad Shahrad, Adi Fuchs, and David Wentzlaff ([jbalkind, mmckeown, trin, alavrov, mshahrad, adif, wentzlaff]@princeton.edu), Princeton University, Princeton, NJ, USA.

Yaosheng Fu* and Samuel Payne* ([yfu, spayne]@nvidia.com), NVIDIA, Santa Clara, CA, USA.

*Work was done at Princeton University.

Yanqi Zhou* (zhouyanqi@baidu.com), Baidu SVAIL, Sunnyvale, CA, USA.

Xiaohua Liang* (xialian@microsoft.com), Microsoft, Redmond, WA, USA.

Matthew Matti* (mmatti@eecs.berkeley.edu), University of California, Berkeley, USA.

Copyright held by authors/owners. Publication rights licensed to ACM.

Advertise with ACM!

Reach the innovators and thought leaders working at the cutting edge of computing and information technology through ACM's magazines, websites and newsletters.

Request a media kit with specifications and pricing:

Ilia Rodriguez
+1 212-626-0686
acmm mediasales@acm.org

Department of Electrical and Computer Engineering Graduate School of Engineering and Management Air Force Institute of Technology (AFIT) Dayton, Ohio Faculty Position

The Department of Electrical and Computer Engineering at the Air Force Institute of Technology is seeking applications for a tenured or tenure-track faculty position. All academic ranks will be considered. Applicants must have an earned doctorate in Electrical Engineering, Computer Engineering, Computer Science, or a closely affiliated discipline by the time of their appointment (anticipated 1 September 2020).

We are particularly interested in applicants specializing in one or more of the following areas: autonomy, artificial intelligence / machine learning, navigation with or without GPS, cyber security, and VLSI. Candidates in other areas of specialization are also encouraged to apply. This position requires teaching at the graduate level as well as establishing and sustaining a strong DoD relevant externally funded research program with a sustainable record of related peer-reviewed publications.

The Air Force Institute of Technology (AFIT) is the premier Department of Defense (DoD) institution for graduate education in science, technology, engineering, and management, and has a Carnegie Classification as a High Research Activity Doctoral University. The Department of Electrical and Computer Engineering offers accredited M.S. and Ph.D. degree programs in Electrical Engineering, Computer Engineering, and Computer Science as well as an MS degree program in Cyber Operations.

Applicants must be U.S. citizens. Full details on the position, the department, applicant qualifications, and application procedures can be found at <http://www.afit.edu/ENG/>. Review of applications will begin on January 6, 2020. The United States Air Force is an equal opportunity, affirmative action employer.