Fast and Accurate Activity-based Cycle-level Power Modeling for Arbitrary RTL

Donggyu Kim, Vighnesh Iyer
Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
{dgkim, vighnesh.iyer}@berkeley.edu

***Abstract*—This paper presents a novel methodology to accurately and quickly predict power dissipation of arbitrary RTL designs during performance simulation. First, important signals that are most sensitive to power dissipation are selected by analyzing the circuit graph of the target RTL design. Next, design-specific cycle-level power models are accurately constructed and trained by running microbenchmakrs, random instruction streams, and execution sample snapshots from long-running applications, and using the important signal list, the RTL signal activities, and the power estimates from commercial CAD tools. In addition, the performance simulators generated from RTL designs are automatically instrumented with signal activity counters. Power dissipation is readily predicted during FPGA simulation by reading out the activity counter values. Finally, the power models are validated by replaying random RTL state sample snapshots on detailed gate-level simulation, which provides accurate power estimates with statistically bounded errors. This methodology will be first demonstrated using RISCV-mini, a simple 3-stage pipeline with instruction and data caches, and applied to more complex designs including RocketChip, Hwacha, and BOOM.**

I. Introduction

Power dissipation has been a major design constraint not only for embedded and portable systems but also for high-end servers and datacenters. A plethora of design strategies have been suggested to meet the power constraint while achieving maximum performance. Thus, power estimation methodologies are very important for the evaluation and the validation of novel hardware designs. However, accurate and fast application-specific power estimation for complex hardware design is still a big challenge for system designers and application developers, which will play a more crucial role in the near future as Moore's law is near the end.

Microarchitecture-level power analysis tools calibrated against representative hardware designs are widely used by computer architects [1], [2], [3], [4], [5]. These power models need activity statistics driven from microarchitectural cycle-level software simulators [6], [7], [8]. This approach helps designers study trace-offs with design parameters in early design phase without developing RTL designs. However, this methodology is limited to designs that are similar to what the abstract model is build upon, and requires long simulation times to collect microarchitectural activities from microarchitectural software simulators. Moreover, it is very hard to apply this methodology to non-traditional hardware designs such as application-specific accelerators as their power models should be validated against RTL designs or existing systems.

Once RTL implementations are available, energy efficiency as well as cycle time and area can be evaluated using commercial CAD tools. The existing CAD tools provide very accurate performance and power estimates. However, their simulation time is extremely slow, preventing design space exporation for realistic applications running on complex hardware designs.

This paper describes an activity-based power modeling methodology for arbitrary RTL, which enables fast and accurate cycle-level power estimation for any RTL designs. First, signals, which determines the cycle behavior of the target design, are automatically selected from the circuit graph. Next, a design-specific power model is constructed and trained with the selected signal activities. In addition, the target RTL design is automatically instrumented so that the signal activities can be probed during FPGA performance simulation. Finally, power dissipation is computed in real time by pausing the performance simulation and reading out the signal activities from FPGA.

II. Related Work

Analytic power modeling [1], [2], [3], [4], [5] combined with microarchitectural cycle-level software simulators [6], [7], [8] is widely used for computer architecture research. This enables early microarchitecture-level design space exploration without necessitating RTL development. However, microarchitectural software simulators run much slower than real systems, making full execution of applications infeasible. Moreover, for different microarchitectures with new technologies, a new analytic power model is constructed and strictly validated against real systems or detailed circuit / gate-level simulations, which is very difficult for non-traditional hardware designs.

Power modeling based on performance-monitoring counters is also popular for power estimation [9], [10], [11], [12], [13], [14]. This method provides a quick power estimate by profiling full execution of applications. However, its application is also limited to well-known microarchitectures. Moreover, it is only practical with existing existing physical systems since standard microarchitectural software simulators are extremely slow.

There are significant efforts to accelerate power estimation using an FPGA [15], [16], [17], [18], [19], [20], [21]. Coburn et al. [15] implement detailed power models directly on the FPGA, which suffers from large FPGA resource overhead. Ghodrat et al. [16] expands the ideas of Coburn et al. [15] by translating part of power computation logic on the FPGA

to software to reduce FPGA resource overhead, but introduces communication overhead between FPGA and software. These methods are inherently lack of scalability due to large overhead of power computation logic.

On the other hand, Atienza et al. [17] manually attached activity monitoring units to the target design on the FPGA. Bhattacharjee et al. [18] manually implement event counters on top of RTL designs to estimate power dissipation from FPGA emulation. PrEsto [19] generates power models from signals manually selected by users. In addition, This technique also requires additional manual efforts to instrument existing FPGA simulators with power models. These methods have less FPGA resource overhead, but require significant manual efforts as well as designers' intuition.

Yang et al. [20] proposes to construct cycle-by-cycle power models by pruning probed signals with machine learning without requiring designers' intuition. However, this method simply assumes dynamic power dissipation is a linear function of primary registers' toggle activities, which is not necessarily always true. Also, it instruments the RTL designs with the power computing logic as in PrEsto [19], which may be automated by commercial CAD tools.

Strober [21] automatically transforms RTL designs into FPGA performance simulators and also automatically instrument the simulators with scan chains to capture RTL state snapshots. Then, the RTL state snapshots are randomly sampled from the FPGA simulation and replayed on detailed gate-level simulation to compute the average power. Although, Strober quickly provides accurate power estimates with statistically bounded errors, the average power should be computed after the FPGA simulation finishes, preventing online power computation, which is the motivation of this paper.

III. Methodology

There are five major steps in our methodology(Figure 1). First, the *signal selection* step walks the circuit graph in FIRRTL[22] and picks important signal candidates according to the circuit topology. Next, the *power model construct and training* step builds a power model using selected signals and trains the model using preliminary waveforms from small workloads. The *RTL instrumentation* step automatically instrument FPGA performance simulators generated by Strober [21] with the activity counters for the signals from the previous step. The *power prediction* step computes the power dissipation in real time during FPGA performance simulation. Finally, the *power validation* step validates the power model by comparing against the power estimation from Strober, and further trains the power model if necessary.

A. Signal Selection

To construct and train a power model, we should select a subset of signals because it is infeasible to keep track of all signal activities on the FPGA. Firgure 2 describes how power sensitive signals are selected.

First of all, a random RTL design is fed into the FIRRTL compiler. In this paper, we assume this design is written with

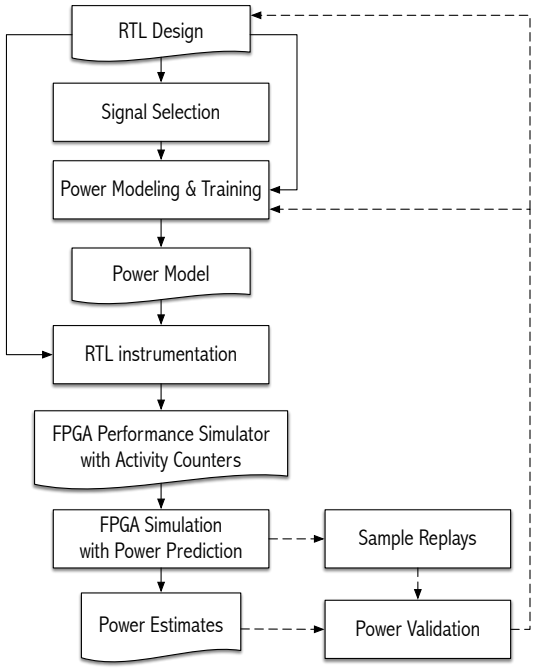

Fig. 1. Overall Tool Flow

Chisel [23]. However, the target designs are not necessarily developed with Chisel because the FIRRTL compiler will support various languages including Verilog in the near future.

In the FIRRTL compiler, a circuit graph is expressed with an intermediate representation, FIRRTL [22], so that a designer can write custom compiler passes for their design. In this step, a custom pass selecting important signals and dumping the signal list is implemented. The output signal list will be passed to the next step to build a design-specific power model.

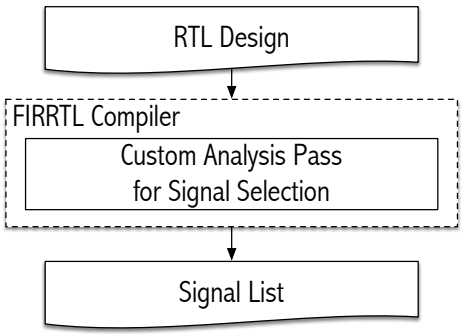

Fig. 2. Signal Selection

B. Power Modeling and Training

Once the signal list is obtained from the previous step, a design-specific power model can be constructed expressed with the signal toggle activities in the list. Figure 4 explains how power models are built and trained. First, an RTL design is fed into the logic synthesis tool(e.g. Synopsys Design

Compiler ®) and the place-and-route tool (e.g. Synopsys IC Compiler ®) to obtain a gate-level design, which is simulated in SDF back-annotated gate-level simulation(e.g. Synopsys VCS ®) for accurate power estimates by the power analysis tool(e.g. Synopsys PrimeTime PX ®).

RTL signal activities are also computed from RTL simulation(e.g. Verilator, Synopsys VCS ®). By providing the signal list, the RTL signal activities, and the detailed power estimates to the power modeling and training algorithm, a design-specific power model can be constructed expressed with the signal toggle activities.

For RTL simulation and gate-level simulation, microbenchmarks or random instruction streams are employed for initial power model training. In addition, there can be random execution sample snapshots from long-running FPGA performance simulation [21], which are used not only for power model validation but also for further training the power model for more accurate power estimates.

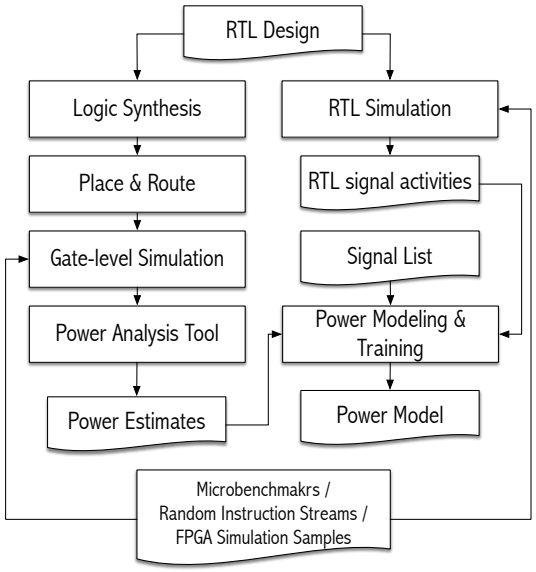

Fig. 3. Power Modeling and Training

C. RTL Instrumentation

The RTL instrumentation is built upon the Chisel3 and FIRRTL port of the Strober's flow to generate FPGA performance simulators. Figure 4 shows how the FPGA performance simulators are instrumented with the signal activity counters.

First, the target RTL design is FAME1-transformed to create the FPGA performance simulator [21] in the FIRRTL compiler. Then, a custom transform is executed to attach activity counters to the FPGA performance simulator by using the signal list from Section III-A. In addition, scan chains are inserted not only to read out the activity counter values but also to capture RTL state snapshots to validate the power model. Finally, simulation mapping and platform mapping are conducted for FPGA simulation [21].

Note that signal selection in Section III-A can be integrated to this flow if signals are not pruned during power modeling and training in Section III-B. In this case, FPGA simulation can run in parallel while training power models, which uses slow CAD tools extensively.

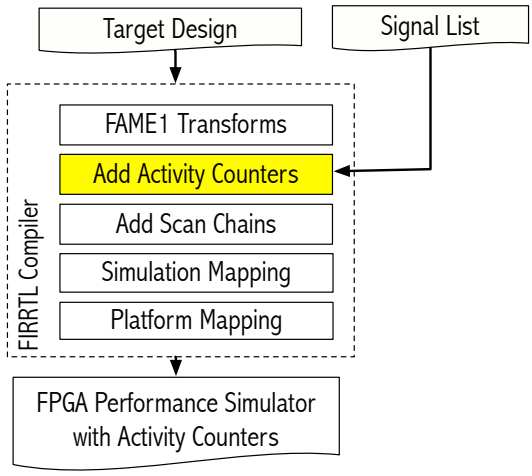

Fig. 4. Power Modeling and Training

D. Power Prediction during FPGA Simulation

Once the FPGA performance simulator instrumented with activity counters is obtained, power estimates are quickly available during FPGA simulation. Since the target RTL is FAME1-transformed, FPGA simulation can be easily paused, and activity counter values are read out from the FPGA in the same way random RTL state sample snapnots are taken in Strober [21]. If the power model is not available during FPGA simulation because it is still trained, the activity counter values are saved to estimate the power dissipation later.

E. Power Validation

Designers may want to validate the power estimates from Section III-D after FPGA simulation. Since this methodology is combined with Strober, an accurate power estimates with statistically bounded errors are available by replaying random RTL state snapshots, which are used to further train the power model for more accurate power estimates (Figure 4).

IV. Experimental Setup

To demonstrate our idea, we will start from a simple design, RISCV-mini [24]. RISCV-mini implements RV32I of the User-level ISA Version 2.0 [25] and the Machine-level ISA of the Privileged Architecture Version 1.7 [26]. RISCV-mini includes a 3-stage pipeline as well as instruction and data caches as shown in Figure 5, which is neither too simple nor too complex. This is designed to validate novel ideas before applying them to more complex designs including RocketChip [27], Hwacha [28], and BOOM [29]. We expect that the methodology proposed in this report will accurately predict the power dissipation of microbenchmarks(median, multiply, qsort, towers, and vvadd) running on RISCV-mini.

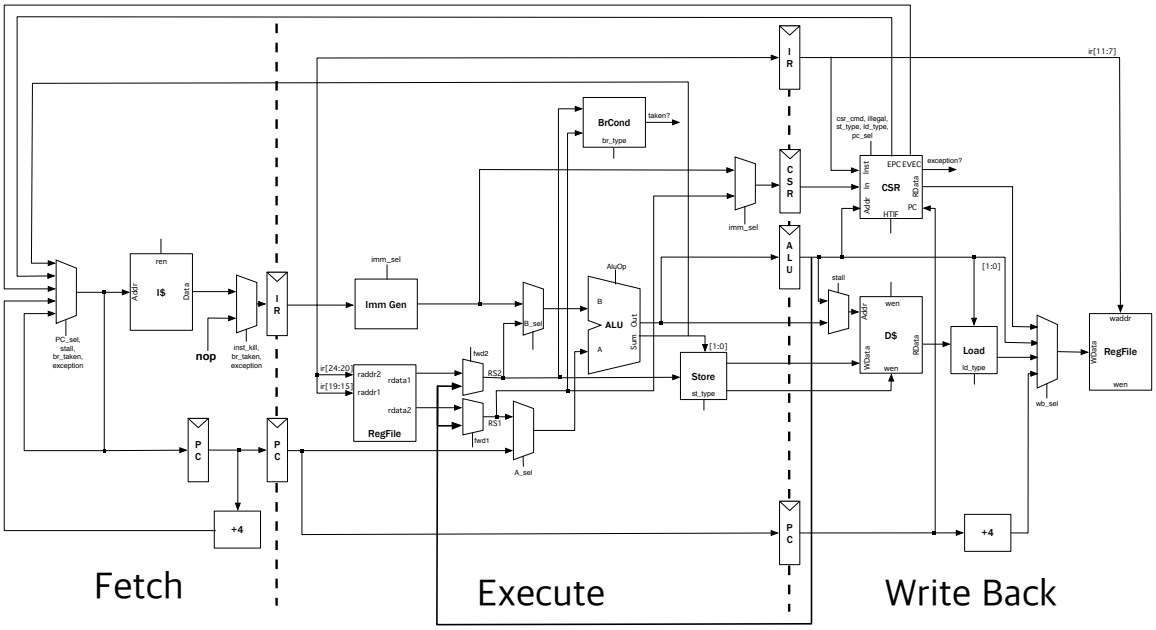

Fig. 5. RISCV-mini Pipeline

V. Conclusion

In this report, a novel approach to accurately and quickly predict cycle-level power estimates for arbitrary RTL was presented. Import signals sensitive to power dissipation were automatically selected using a custom analysis pass in the FIRRTL compiler. Also, power models were built and trained using the signal list, the RTL signal activities, and the accurate power estimates from detailed gate-level simulation. The RTL instrumentation was done on top of the Strober's flow for the FPGA performance simulators to automatically add activity counters using the signal list. During FPGA simulation, online power prediction was readily obtained by pausing the simulation and reading out the activity counters. Finally, the power model was validated and further trained by replaying random RTL state snapshots.

One possible application of this methodology is fast DVFS modeling using FPGA simulation. To enable DVFS in simulation, online power prediction must be available for the power management unit model to control voltage and frequency. By using this methodology, it is expected we will simulate DVFS with FPGA performance simulatros in the near future.

References

[1] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: a framework for architectural-level power analysis and optimizations," in *ISCA*, 2000.

[2] N. Vijaykrishnan, M. Kandemir, M. J. Irwin, H. S. Kim, and W. Ye, "Energy-driven integrated hardware-software optimizations using SimplePower," in *ISCA*, 2000.

[3] S. Li, J. H. Ahn, R. Strong, J. Brockman, D. Tullsen, and N. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *MICRO*, 2009.

[4] J. Leng, T. Hetherington, A. ElTantawy, S. Gilani, N. S. Kim, T. M. Aamodt, and V. J. Reddi, "GPUWattch: enabling energy optimizations in GPGPUs," in *ISCA*, 2013.

[5] Y. S. Shao, B. Reagen, G.-Y. Wei, and D. Brooks, "Aladdin: A pre-RTL, power-performance accelerator simulator enabling large design space exploration of customized architectures," in *ISCA*, 2014.

[6] N. Binkert, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, D. A. Wood, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, and T. Krishna, "The gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol. 39, Aug 2011.

[7] A. Patel, F. Afram, and S. Chen, "MARSSx86: A full system simulator for x86 CPUs," in *DAC*, 2011.

[8] T. F. T. Wenisch, R. R. E. R. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, and J. C. J. Hoe, "SimFlex: Statistical Sampling of Computer System Simulation," *IEEE Micro*, vol. 26, pp. 18–31, Jul 2006.

[9] F. Bellosa, "The benefits of event: driven energy accounting in power-sensitive systems," in *The 9th ACM SIGOPS European workshop*, 2000.

[10] T. Li and L. K. John, "Run-time modeling and estimation of operating system power consumption," in *SIGMETRICS*, 2003.

[11] C. Isci and M. Martonosi, "Runtime power monitoring in high-end processors: Methodology and empirical data," in *MICRO*, 2003.

[12] W. Bircher, M. Valluri, J. Law, and L. John, "Runtime identification of microprocessor energy saving opportunities," in *ISLPED*, 2005.

[13] W. L. Bircher and L. K. John, "Complete System Power Estimation: A Trickle-Down Approach Based on Performance Events," in *ISPASS*, 2007.

[14] R. Bertran, M. Gonzelez, X. Martorell, N. Navarro, and E. Ayguade, "A Systematic Methodology to Generate Decomposable and Responsive Power Models for CMPs," *IEEE Transactions on Computers*, vol. 62, pp. 1289–1302, Jul 2013.

[15] J. Coburn, S. Ravi, and A. Raghunathan, "Power emulation: A new paradigm for power estimation," in *DAC*, 2005.

[16] M. A. M. Ghodrat, K. Lahiri, and A. Raghunathan, "Accelerating system-on-chip power analysis using hybrid power estimation," in *DAC*, 2007.

[17] D. Atienza, P. G. Del Valle, G. Paci, F. Poletti, L. Benini, G. De Micheli, and J. M. Mendias, "A fast HW/SW FPGA-based thermal emulation framework for multi-processor system-on-chip," in *DAC*, 2006.

[18] A. Bhattacharjee, G. Contreras, and M. Martonosi, "Full-system chip multiprocessor power evaluations using FPGA-based emulation," in *ISLPED*, 2008.

[19] D. Sunwoo, G. Y. Wu, N. a. Patil, and D. Chiou, "PrEsto: An FPGA-accelerated Power Estimation Methodology for Complex Systems," in *FPGA*, 2010.

[20] J. Yang, L. Ma, K. Zhao, Y. Cai, and T. F. Ngai, "Early stage real-time SoC power estimation using RTL instrumentation," in *ASP-DAC*, 2015.
[21] D. Kim, A. Izraelevitz, C. Celio, H. Kim, B. Zimmer, Y. Lee, J. Bachrach, and K. Asanović, "Strober : Fast and Accurate Sample-Based Energy Simulation for Arbitrary RTL," in *ISCA*, 2016.
[22] P. S. Li, A. M. Izraelevitz, and J. Bachrach, "Specification for the FIRRTL Language," Tech. Rep. UCB/EECS-2016-9, EECS Department, University of California, Berkeley, 2016.
[23] J. Bachrach, H. Vo, B. Richards, Y. Lee, A. Waterman, R. Avižienis, J. Wawrzynek, and K. Asanović, "Chisel: constructing hardware in a scala embedded language," in *DAC*, 2012.
[24] D. Kim, "RISCV-mini: A Simple RISC-V 3-state Pipeline in Chisel." https://github.com/ucb-bar/riscv-mini.git.
[25] A. Waterman, Y. Lee, K. Asanović, and D. Patterson, "The RISC-V Instruction Set Manual: User-Level ISA Version 2.0," Tech. Rep. UCB/EECS-2014-54, EECS Department, University of California, Berkeley, 2014.
[26] A. Waterman, Y. Lee, R. Avizienis, D. Patterson, and K. Asanovi, "The RISC-V Instruction Set Manual: Privileged Architecture Version 1.7," Tech. Rep. EECS-2015-157, EECS Department, University of California, Berkeley, 2015.
[27] K. Asanović, R. Avizienis, J. Bachrach, S. Beamer, D. Biancolin, C. Celio, H. Cook, D. Dabbelt, J. Hauser, A. Izraelevitz, S. Karandikar, B. Keller, D. Kim, J. Koenig, Y. Lee, E. Love, M. Maas, A. Magyar, H. Mao, M. Moreto, A. Ou, D. A. Patterson, B. Richards, C. Schmidt, S. Twigg, H. Vo, and A. Waterman, "The Rocket Chip Generator," Tech. Rep. UCB/EECS-2016-17, EECS Department, University of California, Berkeley, Apr 2016.
[28] Y. Lee, *Decoupled Vector-Fetch Architecture with a Scalarizing Compiler*. PhD thesis, University of California, Berkeley, May 2016.
[29] C. Celio, D. A. Patterson, and K. Asanović, "The Berkeley Out-of-Order Machine (BOOM): An Industry-Competitive, Synthesizable, Parameterized RISC-V Processor," Tech. Rep. UCB/EECS-2015-167, EECS Department, University of California, Berkeley, Jun 2015.