Our Contribution

Experiment results

Summary 0

Early Stage Real-Time SoC Power Estimation Using RTL Instrumentation

Jianlei Yang^{1,2} Liwei Ma¹ Kang Zhao¹ Tin-Fook Ngai¹ Yici Cai²

¹Intel Labs China, Intel

²Department of Computer Science, Tsinghua University

ASPDAC, January 2015



Motivation	Our Contribution	Experiment results	Summary o
Outline			

- Accurate Real-time Power Estimation
- Previous Work
- 2 Our Machine Learning Approach and Contribution
 - Main Results
 - Machine Learning Approach
 - Synthesizable RTL Instrumentation for Real-Time
- 3 Experiment results on Real IP
 - Experiment results on H.264/AVC
 - Experiment results on AES and AC97

Motivation 00	Our Contribution	Experiment results	Summary o
Outline			

- Accurate Real-time Power Estimation
- Previous Work
- 2 Our Machine Learning Approach and Contribution
 - Main Results
 - Machine Learning Approach
 - Synthesizable RTL Instrumentation for Real-Time
- 3 Experiment results on Real IP
 - Experiment results on H.264/AVC
 - Experiment results on AES and AC97

Early Power Estimation for Architecture Exploration

- SoC architecture explorations need,
- Software/hardware co-designs need,
- Always a target of EDA tools.

Early Power Estimation for Architecture Exploration

- SoC architecture explorations need,
- Software/hardware co-designs need,
- Always a target of EDA tools.
- FPGA co-emulation prototype widely adopted,
- But how about power estimation?

Previous Works Focusing on Module Boundaries.

Capacity



Previous Works Focusing on Module Boundaries.



Solution space

Previous RTL hack

- On boundaries [1][2][3],
- On modules/funcitons
 [2][4],
- Complex with cross-term [1].
- No huge data employed,
- No machine learning,
- Lack automation for radom logic.

Mo	otiv	ati	on
00			

Outline

1 Motivation

- Accurate Real-time Power Estimation
- Previous Work

2 Our Machine Learning Approach and Contribution

- Main Results
- Machine Learning Approach
- Synthesizable RTL Instrumentation for Real-Time
- 3 Experiment results on Real IP
 - Experiment results on H.264/AVC
 - Experiment results on AES and AC97

A New EDA Flow for Early Power Estimation

Our Contribution

- Work for random logic,
- Automatic machine learning aproach,
- Within 5% accuracy loss,
- Sythesizable RTL instrumentation,
- Within 7% extra LUTs.

Merged EDA flow



Summary 0

Key Registers Indicate Power Consumption

Logic Cone

 Key register toggling prorogate with a logic cone,



Summary o

Key Registers Indicate Power Consumption

Synergy

- Key register toggling prorogate with a logic cone,
- Registers flipping synergy,



Summary o

Key Registers Indicate Power Consumption

Invariant

- Key register toggling prorogate with a logic cone,
- Registers flipping synergy,
- Invariant boundary between ASIC and FPGA.



Summary o

Let Machine Learn the Relationship

Title

- Power trace,
- Register toggle,
- SVD machine learning,
- Calibration X,
- Power Prediction.



Summary o

RTL instrumentation with Adder Tree

Group according bits

- XOR to get toggle,
- Same coefficients,
- +|- Group,
- 3-stage adder tree.

Toggle-Coefficients adder tree



Mo	tiv	/at	io	n
00				

Outline

1 Motivation

- Accurate Real-time Power Estimation
- Previous Work

2 Our Machine Learning Approach and Contribution

- Main Results
- Machine Learning Approach
- Synthesizable RTL Instrumentation for Real-Time

3 Experiment results on Real IP

- Experiment results on H.264/AVC
- Experiment results on AES and AC97

H.264/AVC baseline decoder of QCIF [5]

Title

- Total 31K registers,
- 400K cycles.
- ~2K registers remained.
- 7% extra for 1-stage tree,
- 12%, for 2-stage.

Toggle-Coefficients adder tree



Cross-Prediction Error Results

Normalized RMS error of cycle-by-cycle prediction

NRMSD	Akiyo	Carphone	Claire
Akiyo	2.51%	2.68%	3.20%
Carphone	4.35%	2.53%	4.13%
Claire	3.43%	3.62%	2.22%

Relative errors of total power prediction

Relative Error	Akiyo	Carphone	Claire
Akiyo	0.09%	1.07%	1.89%
Carphone	2.58%	0.24%	3.47%
Claire	0.40%	1.19%	0.29%

Mot	iva	tio	

Our Contribution

Experiment results

Power Prediction Waveform



Summary o

Experiment results on AES [6] and AC97 [7]

- AES : 678 registers,
- AC97: 2288 registers.

IP Core	Calibration		Prediction	
	NRMSE	RelErr	NRMSE	RelErr
AES	3.25%	3.39%	3.35%	2.45%
AC97	1.74%	0.27%	0.85%	0.75%

Motivation	Our Contribution	Experiment results	Summary o
Outline			

- Accurate Real-time Power Estimation
- Previous Work

2 Our Machine Learning Approach and Contribution

- Main Results
- Machine Learning Approach
- Synthesizable RTL Instrumentation for Real-Time

3 Experiment results on Real IP

- Experiment results on H.264/AVC
- Experiment results on AES and AC97

- Work for random logic by a machine learning approach to abstract power model,
- Real-time estimation by RTL instrumentation with synthesizable model,
- <5% power estimation accuracy and <7% LUTs resource overhead.

For Further Reading

- Dam Sunwoo, Gene Y. Wu, Nikhil A. Patil, and Derek Chiou. PrEsto: An FPGA-accelerated power estimation methodology for complex systems. In *Proc. FPL*, pages 310–317, 2010.
- [2] Joel Coburn, Srivaths Ravi, and Anand Raghunathan. Power emulation: a new paradigm for power estimation. In *Proc. DAC*, pages 700–705, 2005.
- [3] Sumit Ahuja, Avinash Lakshminarayana, and Sandeep Kumar Shukla. Low Power Design with High-Level Power Estimation and Power-Aware Synthesis, chapter Regression-Based Dynamic Power Estimation for FPGAs. Springer, 2012.
- [4] Abhishek Bhattacharjee, Gilberto Contreras, and Margaret Martonosi.
 Full-system chip multiprocessor power evaluations using FPGA-based emulation. In *Proc. ISLPED*, pages 335–340, 2008.
- [5] H.264/AVC Baseline Decoder.

http://opencores.org/project, nova.

[6] AES IP Core.

http://opencores.org/project,aes_core.

[7] AC 97 Controller IP Core.

http://opencores.org/project,ac97.