Energy-Efficient Processor Design

C2S2 Center for Circuit & System Solutions

Omid Azizi, Mark Horowitz
Stanford University

Microprocessor Design

- **Processors are complex systems with a large design space**
  - There are plenty of design parameters and choices
    - In-order vs out-of-order execution, cache size, pipe latencies, …
- **Would like to design aggressively to maximize performance**
  - But power consumption considerations limit design performance
- **Instead, we must evaluate design choice tradeoffs**
  - A design feature or parameter may increase performance
  - But it usually comes with an associated cost (power, area)
- **Problem: What is the "best" set of design choices?**

The Processor Design Space

- **Each design configuration is a point in the cost-performance space**
- **We want to find the *efficient* designs**
  - Designs which maximize performance for a given power or area budget
  - Designs which make the best use of available resources

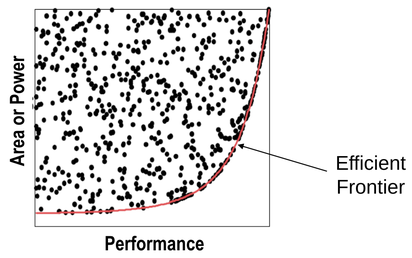

Trade-off Analysis & Marginal Costs

- **Finding efficient designs is a trade-off analysis problem**
  - A design feature usually affects both performance and energy
- **To gauge efficiency of design choices, use marginal costs**
  - Want those choices with the lowest cost per unit performance
- **If we know marginal costs, then we can optimize a design**
  - "Buy" parameters with a low marginal cost, "sell" parameters with high cost

$$\text{Marginal Cost of } x = \frac{\Delta E}{\Delta P}$$

Energy cost of x

Performance benefit of x

An Integrated Hierarchical Optimization

- **Costs of underlying units depend on their implementations**
  - True optimization needs to be aware of the tradeoffs of the underlying units
    - Circuit design and architecture design are not independent

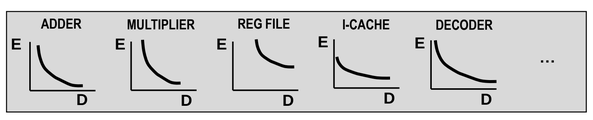

- **Our goal**
  - An integrated optimization framework that simultaneously considers both circuit and architectural parameters to produce a globally optimized processor design

Optimization Framework: Overview

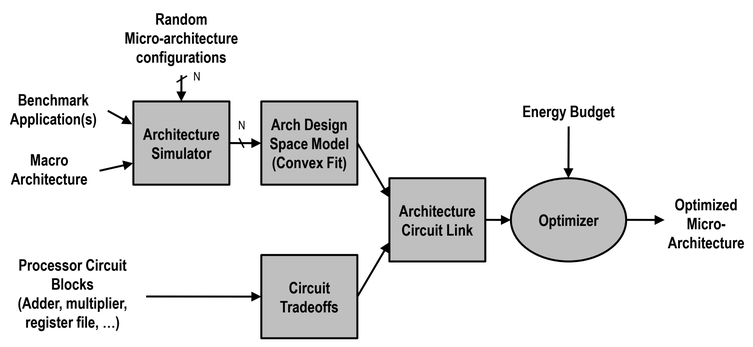

Circuit Tradeoffs Library

- **For each circuit, we need circuit-level energy-delay tradeoffs**
- **Use a synthesis tool to explore circuit design space**
  - Circuit topology, gate sizing, etc.
- **Fit mathematical model to data**
- **Store results in a library**

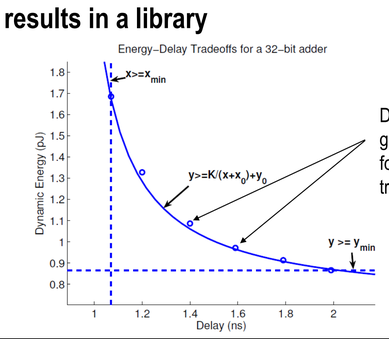

Design parameters (topology, gate sizing, etc.) are annotated for all designs on the area-delay tradeoff curve

Architectural Models

- **The architectural models need to translate the characteristics of the underlying blocks to an effect on the whole system**

$$E_{total} = \sum(\alpha_i E_{dyn\text{-}i}) + T \times \sum(P_{lkg\text{-}i})$$

dynamic energy; leakage energy

Activity factor of component *i*; Dynamic energy of component *i*; Leakage power of component *i*

$$P_{total} = f(delay_{adder}, delay_{L1\text{-}cache}, size_{L1\text{-}cache}, \ldots)$$

Need to characterize this function
step 1: randomly sample design space through simulation
step 2: create a posynomial fit using random samples
step 3: validate design space error

Architectural Performance Modeling and Optimization

- **Performance modeling with simulation becomes prohibitive**
  - Too many design parameters, exponentially growing space
  - Solution: Use statistical techniques
    - Randomly sample design space
    - Then create models by fitting data using *posynomial* fits
- **Final Step: Optimize**
  - We have mathematically expressed circuit tradeoffs, architectural models, and the link between them
  - Then geometric programming allows us to optimize across the design space *quickly* and *globally*

Optimization Results

- **Optimized 4 different architectures running SPEC**
  - Dual-issue out-of-order efficient for high performance
  - Dual-issue in-order efficient for low energy
  - Never use single issue (for this benchmark set)

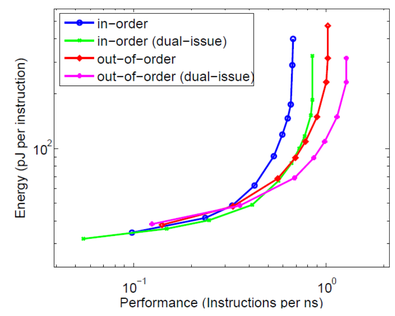

Energy Breakdown of Designs on Efficient Frontier

- **What is the optimal allocation of your energy budget to units?**
  - A lot of energy spent in memory accesses – an energy-efficiency bottleneck

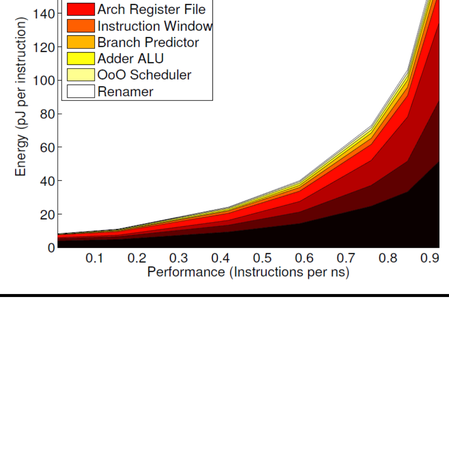

Under the Hood: How parameters change

- **What are the optimal settings for the design knobs?**
  - Optimizer finds these optimal operating points for each parameters
  - As we push for more performance, each of the knobs change too
    - But each parameter has its own profile
  - Optimizer helps guide the design process by telling designers how to tune their designs for efficient operation

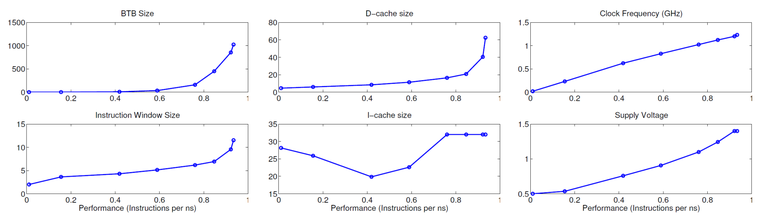

Conclusion & Future Work

- **We have developed a new processor optimization framework**
  - Optimizer performs integrated optimization of circuit and architectural design spaces
  - Produces the optimal architecture and design configuration for efficient operation
- **Future Work**
  - Exploration of new high-level architectures
  - Incorporation of customized code compilation into the optimization design space
- **Acknowledgments**
  - This work was performed in collaboration with Aqeel Mahesri and Sanjay Patel from the University of Illinois at Urbana-Champaign
- **Related References**
  - Omid Azizi; Mahesri, A.; Patel, S.; Horowitz, M., "Area-Efficiency in CMP Core Design: Co-optimization of Circuits and Microarchitecture," *Workshop on Design, Architecture, and Simulation of Chip Multiprocessors (dasCMP), 41st International Symposium on Microarchitecture*, November 2008.
  - Omid Azizi; Collins, J.; Patil, D.; Wang, H.; Horowitz, M., "Processor Performance Modeling using Symbolic Simulation," *IEEE International Symposium on Performance Analysis of Systems & Software, 2007. ISPASS 2008.* 20-22 April 2008.

FCRP

STANFORD UNIVERSITY