

# Wattch, Hotspot, Hotleakage, McPAT

<http://www.eecs.harvard.edu/~dbrooks/wattch-form.html>

<http://lava.cs.virginia.edu/HotSpot>

<http://lava.cs.virginia.edu/HotLeakage>

<http://www.hpl.hp.com/research/mcpat/>

Ramon Canal  
NCD – Master MIRI



# Wattch, architecture power modelling

**Idea:** Perform cycle by cycle tracking of power dissipation  
By estimating unit capacitances and activity factors.

## Methodology:

Parameterized power models of common structures present in modern processors.

## Basic mathematical equation:

Dynamic power consumption in CMOS processors

$$P_d = \alpha C V_{dd}^2 f$$

C : Load capacitance (farads)

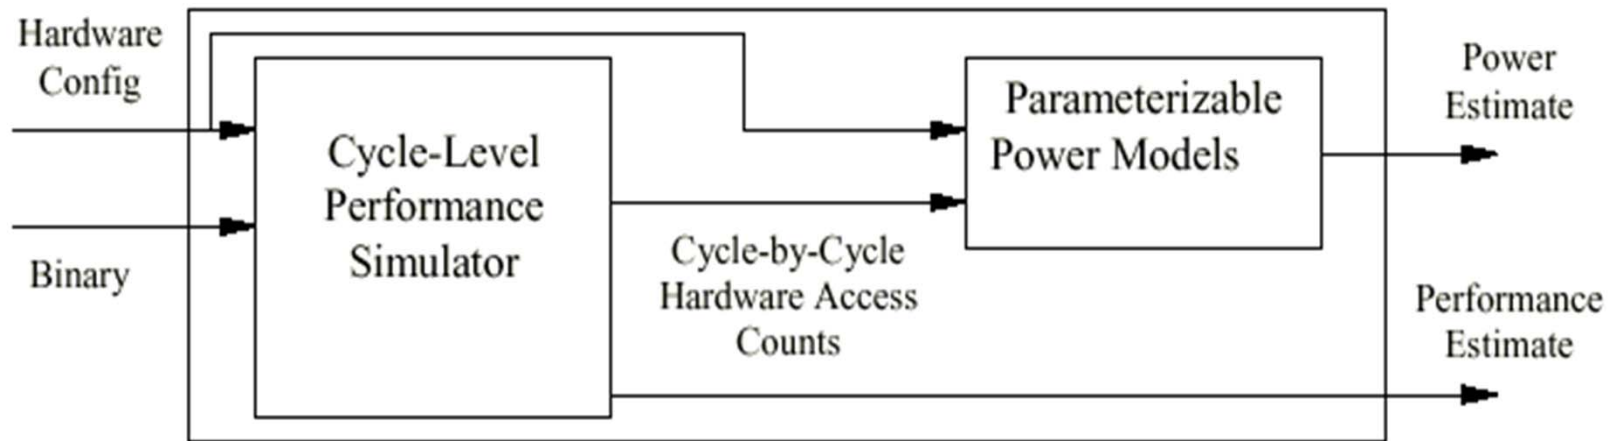
$V_{dd}$ : Supply voltage (V)

f : Clock frequency (Hz)

$\alpha$  : A number between 0 and 1 indicating how often clock ticks lead to activity on an average.



# Power Modeling Methodology



Overall Structure of the Power Simulator

# Classification of the processor units:

## 1. Array structures:

Data and instruction caches, cache tag arrays, all register files, register alias tables etc.

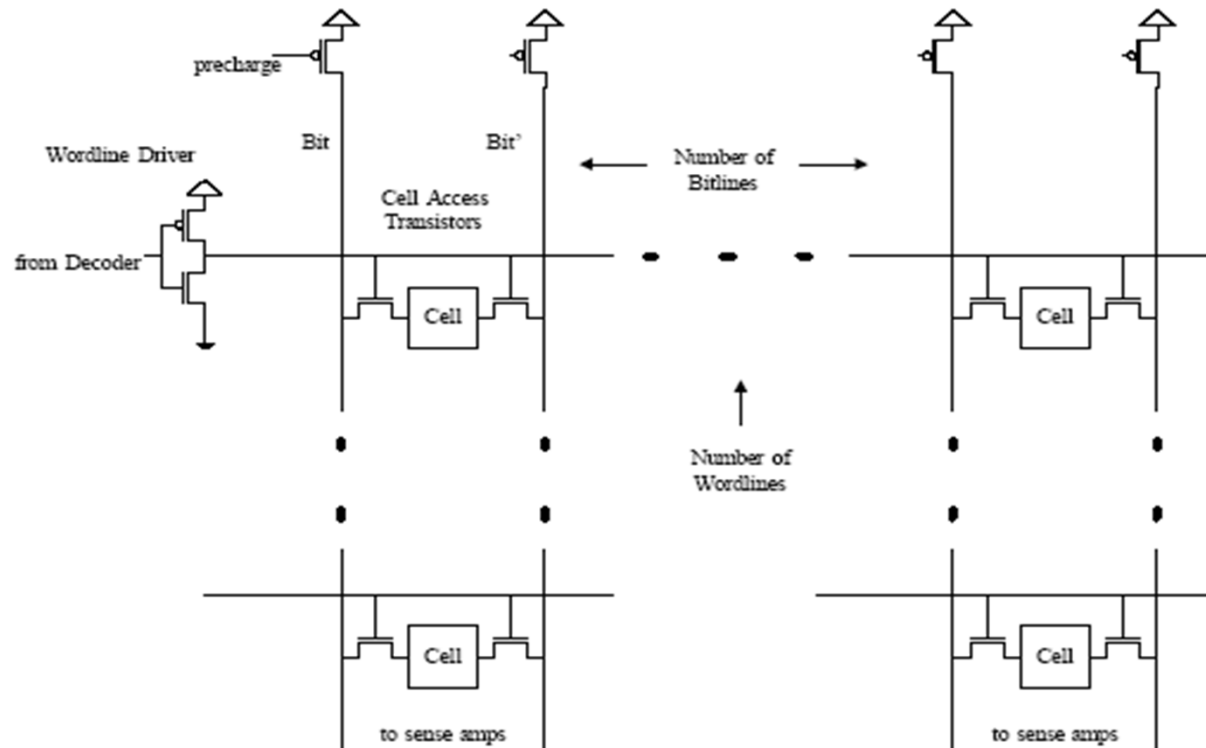


Figure 3: Schematic of wordlines and bitlines in array structure.

# Classification of the processor units:

- 2. Fully associative Content-addressable memories:  
TLBs, Instruction window wake-up logic etc...

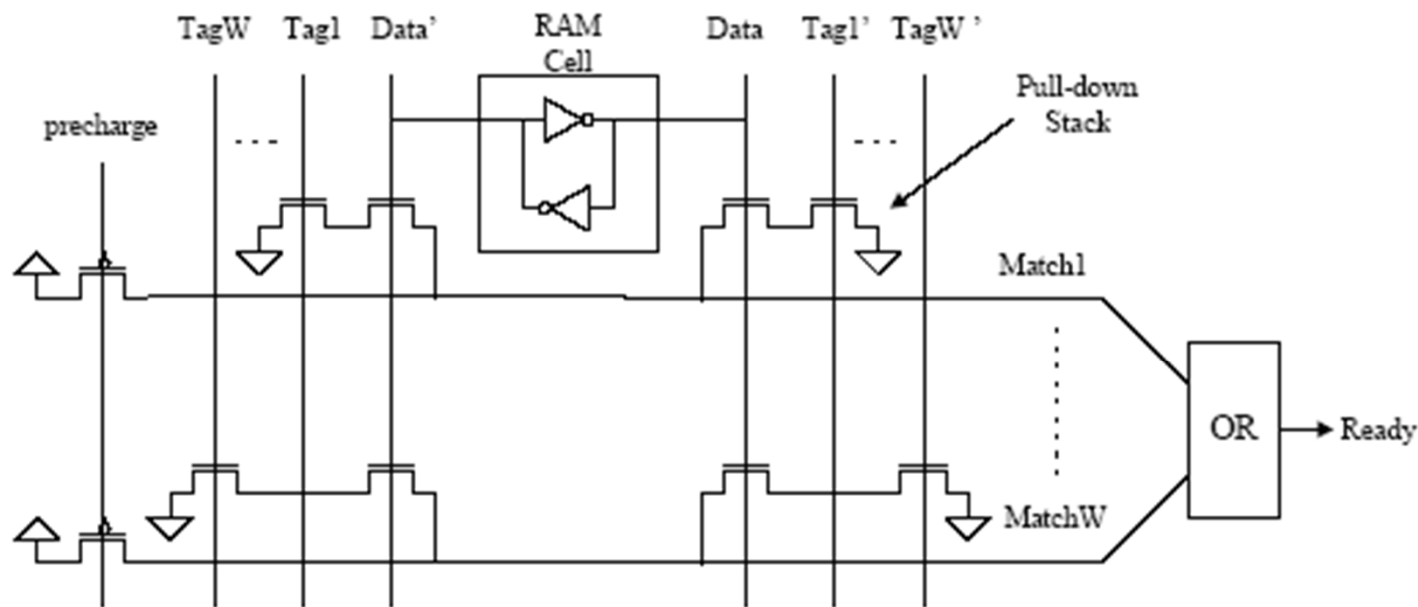


Figure 4: Core cell of wakeup logic modeled as a CAM.

# Classification of the processor units:

## 3. Combinational Logic and Wires:

Instruction selection logic, Functional units, result buses etc...

- Result buses: Find their lengths and multiply by their capacitance per unit length.
- Functional units: based on their own model and further scaling for process technology and frequency

## 4. Clocking Model:

Clock buffers, clock wires and capacitive loads.

- Global clock metal lines : A modified H-tree network in which the global clock signal is routed to all portions of the chip



# SimpleScalar Interface

- The power models are interfaced with SimpleScalar
  - *sim-outorder*
    - modern out-of-order processors with 5-stage pipelines :  
fetch, decode, issue, writeback and commit
    - speculative execution
  - Conditional Clocking Styles
    - All or Nothing Clock Gating
      - full power for use and no power for idle
    - Linear Clock Gating with 10%
      - power is scaled linearly but with minimum background power
    - Linear Clock Gating
      - power is scaled according to the portion of a unit's ports accessed



# SimpleScalar Interface

- The power models are interfaced with SimpleScalar
  - Simulation Speed
    - performance simulator + power model
    - 105K instructions per second → 80K instructions per second
    - PowerMill takes one hour to run 64-bit adder for 100 test vectors.  
→ Wattch can simulate a full CPU running 280M instructions
    - quite tolerable overhead
  - Simulation Method
    - Compute the base power dissipation for each unit at program startup.
    - The access counts per unit are obtained from SimpleScalar simulator.
    - The base power costs are scaled with per-unit access counts.



# How accurate is it?

Makers claim accuracy within 10% of estimates from industry's leading tools.

- ✓ Seen as a complement to low level tools but is 1000 times faster.

| Hardware Structure        | Intel Data | Model |
|---------------------------|------------|-------|
| Instruction Fetch         | 22.2%      | 21.0% |
| Register Alias Table      | 6.3%       | 4.9%  |
| Reservation Stations      | 7.9%       | 8.9%  |
| Reorder Buffer            | 11.1%      | 11.9% |
| Integer Exec. Unit        | 14.3%      | 14.6% |
| Data Cache Unit           | 11.1%      | 11.5% |
| Memory Order Buffer       | 6.3%       | 4.7%  |
| Floating Point Exec. Unit | 7.9%       | 8.0%  |
| Global Clock              | 7.9%       | 10.5% |
| Branch Target Buffer      | 4.7%       | 3.8%  |

Table 4: Comparison between Modeled and Reported Power Breakdowns for the Pentium Pro®.

| Hardware Structure        | Alpha 21264 | Model |
|---------------------------|-------------|-------|
| Caches                    | 16.1%       | 15.3% |
| Out-of-Order Issue Logic  | 19.3%       | 20.6% |
| Memory Management Unit    | 8.6%        | 11.7% |
| Floating Point Exec. Unit | 10.8%       | 11.0% |
| Integer Exec. Unit        | 10.8%       | 11.0% |
| Total Clock Power         | 34.4%       | 30.4% |

Table 5: Comparison between Modeled and Reported Power Breakdowns for the Alpha 21264.



# Shortcomings

## ❑ Fails to model interconnect capacitances

- Difficult to model capacitances of polysilicon wires as they vary with the physical layout.
- There is no way to estimate the length of interconnect. Capacitance estimation errors gives us an overall error of 6-11%
- Physical implementation may be different from the models assumed in Wattch. A critical thing to note!
- Models only one decoder for an array structure.



# McPAT: Multicore Power Area and Timing

Sheng Li, Jung Ho Ahn, Jay B. Brockman,  
Norman P. Jouppi  
HP Labs

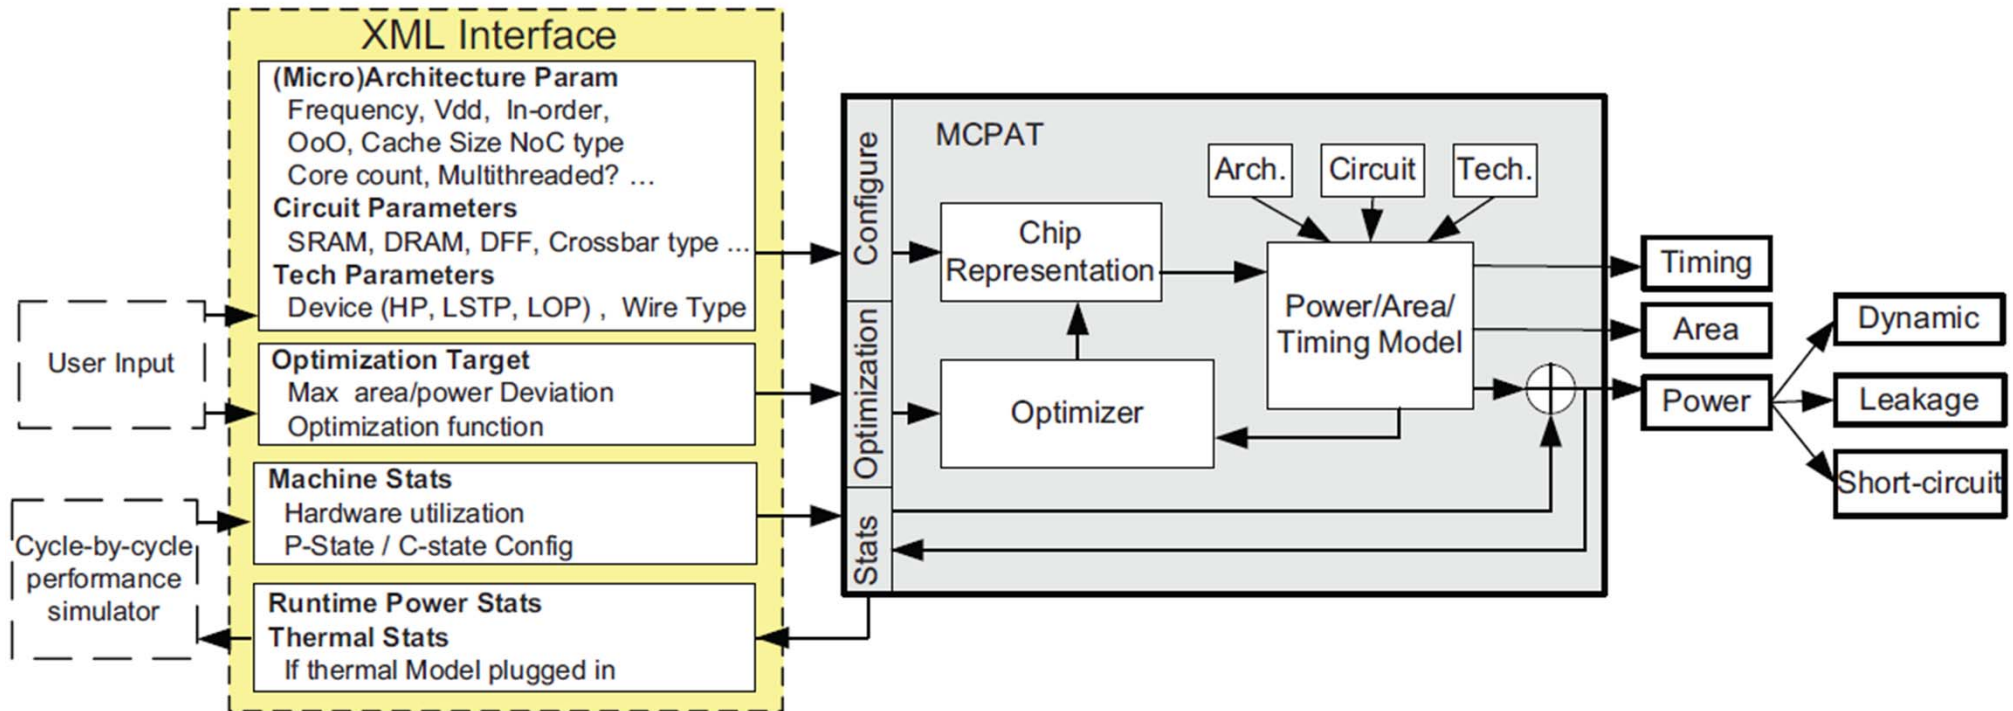


# Overview

- Integration of core/cache/network power models.
- Technology nodes from 90nm to 22nm
- Missing temperature models.




# Overview



# Accuracy

■ Cores ■ Crossbar ■ Leakage ■ L2 Cache ■ Global Clock ■ Interconnect ■ I/Os




Published Niagara2 Data  
@ 1.4GHz/1.1V

McPAT Modeled Data


(b) Validation against Niagara2 processor.

# Accuracy

■ Core Dyn ■ Core Lkg ■ L3 Dyn ■ L3 Lkg ■ Control Dyn ■ Control Lkg



Published Xeon (Tulsa) Data  
@ 3.4GHz/1.25V



McPAT Modeled Data

(d) Validation against Xeon Tulsa processor.

# Summary

- Yet another model, but:
  - *Most complete power model to date*
    - Dynamic and static power models
    - Core, caches and interconnect models
- Main memory not included
- Temperature effects not included



# Hotspot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design

Kevin Skadron, Mircea R. Stan, et al.  
University of Virginia



# Overview

- Introduction to Thermal Models
- Hotspot thermal modeling in detail
- Thermal Model for Interconnects
- Validation of the Model
- Hotspot Applications

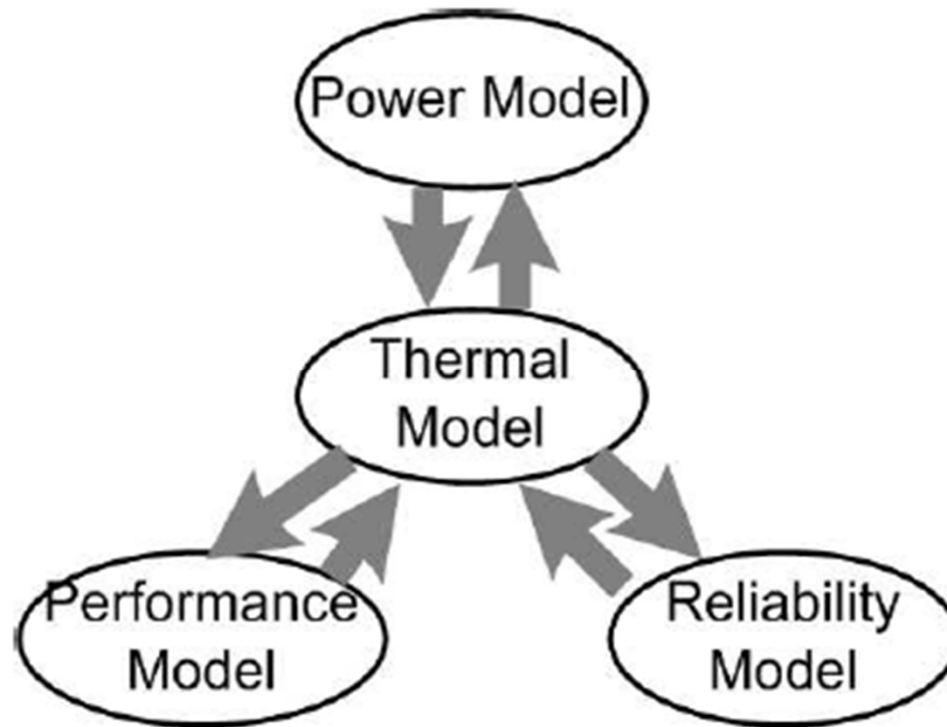


# Impact of high temperatures on Very Large Scale Integrated Systems

- Degradation in carrier mobility
  - A single inverter is 35% slower at 110°C than at 60°C.
- Exponential increase in sub-threshold leakage.
- Increase in interconnect resistivity.
- Decrease in device lifetimes.
- Temperature plays an important role in the early and accurate estimation of power, performance and reliability.



# Interactions among thermal model and power, performance and reliability models



# The Need for Architecture level Thermal Modeling

Ability of architecture domain

- To use runtime knowledge of application behavior and
- Current thermal status of different units of the chip to adjust execution, distribute the workload to control thermal behaviour and exploit instruction-level parallelism.<sup>2</sup>



# Hotspot - A microarchitecture level compact thermal model

- Modeling methodology based on Compact Thermal Models (CTM) and stacked layer packaging configurations.
- Analytical investigation of the relationship between the number of nodes in the CTM and accuracy of the model.
- High-level model for on-chip interconnect self-heating power and temperature.




# Compact Thermal Model in Hotspot

- The model is based on the duality between heat transfer and electrical phenomena.
- Rationale behind this duality – Heat flow and electrical current can be described by similar set of differential equations.

| Thermal quantity               | unit  | Electrical quantity          | unit     |
|--------------------------------|-------|------------------------------|----------|
| $P$ , Heat flow, power         | $W$   | $I$ , Current                | $A$      |
| $T$ , Temperature difference   | $K$   | $V$ , Voltage                | $V$      |
| $R_{th}$ , Thermal resistance  | $K/W$ | $R$ , Electrical resistance  | $\Omega$ |
| $C_{th}$ , Thermal capacitance | $J/K$ | $C$ , Electrical capacitance | $F$      |

# A typical Ceramic Ball Grid Array Package




# Model Overview

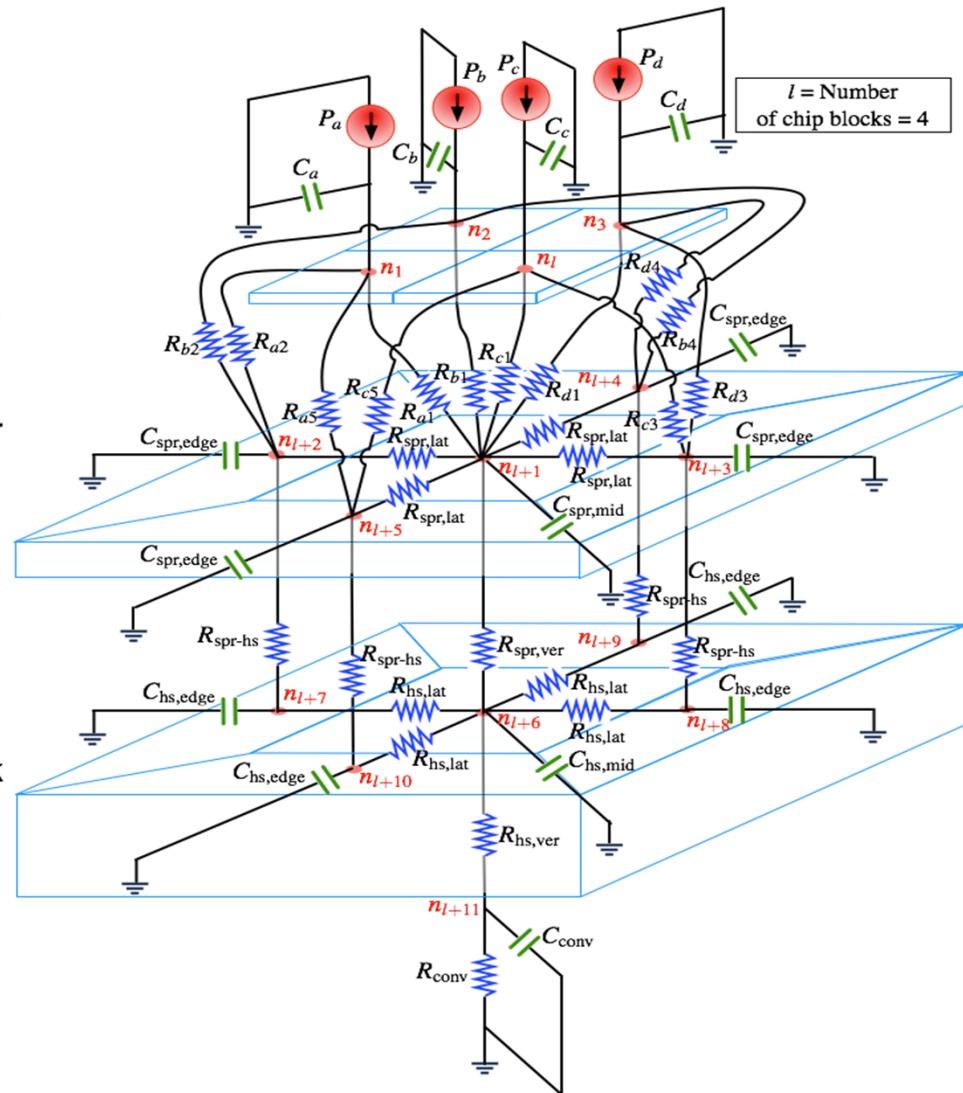
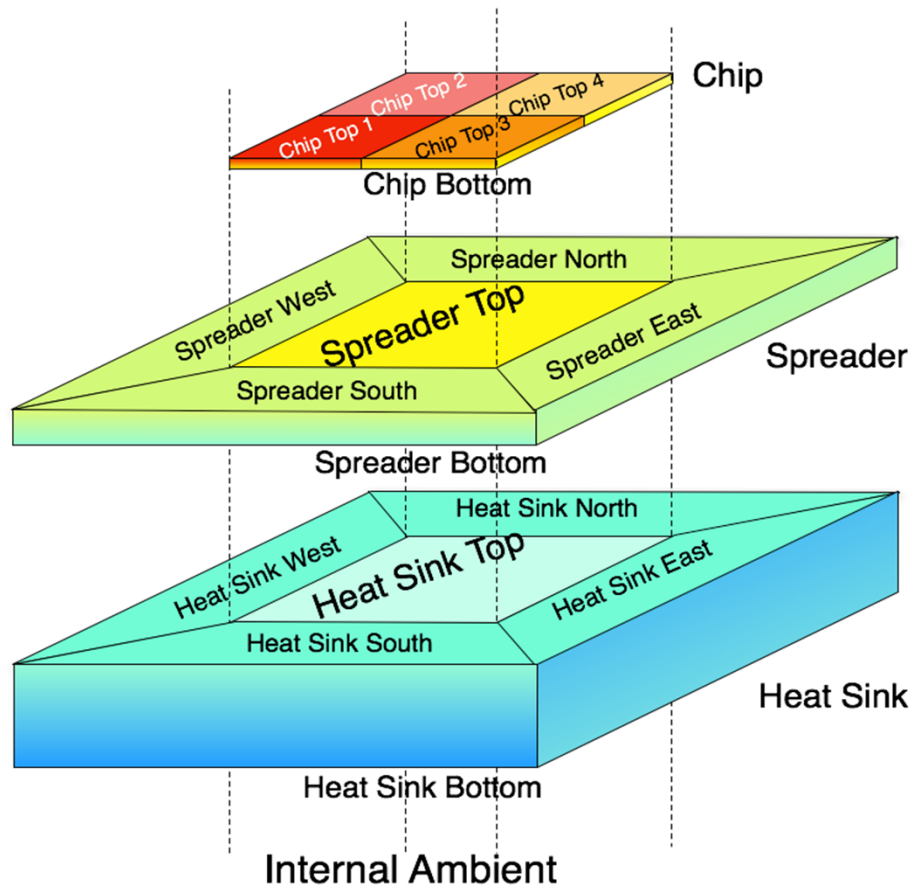
- The CTM proposed is essentially a thermal RC circuit.
- Each node in the circuit corresponds to a block at the desired level of granularity.
- Heat dissipation is modeled using a current source connected to each node.
- Solving the thermal RC circuit gives the temperature at each node.



# Example Hotspot RC model



# A closer view of the RC circuit




# Rs and Cs in detail



- Thermal Resistance and Capacitances

$$R = \frac{t}{k \cdot A}$$


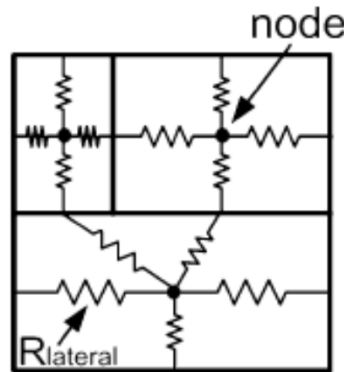
$$C = c \cdot t \cdot A$$



area covered by a smaller neighbor layer



(b)




# Solving the thermal RC circuit

$$P_{i,j} = \frac{T_{i,j} - T_{i-1,j}}{R_1} + \frac{T_{i,j} - T_{i,j+1}}{R_2} + \frac{T_{i,j} - T_{i+1,j}}{R_3} + \frac{T_{i,j} - T_{i,j-1}}{R_4} + \frac{T_{i,j} - T_{bottom}}{R_{bottom}} + \frac{C_{i,j} \Delta T_{i,j}}{\Delta t}$$


$$\Delta T_{i,j} = \frac{P_{i,j} \Delta t}{C_{i,j}} + \frac{\Delta t}{C_{i,j}} \left( \frac{T_{i-1,j}}{R_1} + \frac{T_{i,j+1}}{R_2} + \frac{T_{i+1,j}}{R_3} + \frac{T_{i,j-1}}{R_4} + \frac{T_{bottom}}{R_{bottom}} \right) - \frac{\Delta t}{C_{i,j}} G_{T_{i,j}} T_{i,j}$$

$$G_{T_{i,j}} = \frac{1}{R_1} + \frac{1}{R_1} + \frac{1}{R_2} + \frac{1}{R_3} + \frac{1}{R_4} + \frac{1}{R_{bottom}}$$



# Modeling at different levels of Granularity

Temperature can be modeled at different granularity levels by dividing the die into a number of grid cells.



# Derivation for the minimum grid cell size

*Condition :*

Temperature diff across one grid cell  $\leq$

p % of the max. temperature difference across the die.

$$\Delta T_{\text{grid}} \leq T\left(\frac{-w}{2}\right) - T(0).$$

$$T\left(\frac{-w}{2}\right) - T(0) = \frac{T_{\text{max0}}}{2} (1 - e^{\frac{-w}{2t}}) = p\% \cdot T_{\text{max0}}$$

$$w = 2 \cdot t \cdot \ln \frac{1}{1 - 2 \cdot p\%}.$$



# Interconnect Self-Heating Power and Thermal Modeling

1. The average self-heating power of interconnects in each metal layer,
2. The equivalent thermal resistance for metal wires. (including vias)



# 1. Interconnect self-heating power model

- Self-heating power of a metal wire is given by

$$P_{\text{self}} = I^2 \cdot R = I^2 \cdot \rho_m \cdot l / A_m$$

- Metal resistivity is temperature-dependent.
- The model needs to predict wire temperature before physical layout is available.
- Hence, it has to be able to predict the average wire length and self-heating current.

# a) Average interconnect length in each metal layer

- Different methods of prediction for Signal Interconnects and Power Distribution Network because of the difference in their routing schemes.
- The wire-length distribution model presented by Davis et al.<sup>3</sup> ( involving the concept of Rent's rule) is used in determining the average interconnect length in each layer.

## b) Average Interconnect rms self-heating current in each metal layer

- The average current flow through the interconnect is solved from the equation,

$$I_{\text{RMS}}^2(R_{\text{tr}} + R_{\text{wire}})t_d = \frac{1}{2}\alpha C_L V_{\text{dd}}^2$$

$I_{\text{RMS}}$  – Self-heating current per wire in each metal layer,

$R_{\text{tr}}$  – On-resistance of the transistor,

$R_{\text{wire}}$  – Wire resistance,

$t_d$  – Delay of the switching event.

## c) Interconnect rms self-heating current for Power supply sections

- Method 1:

Modeled and solved using an RC circuit similar to that of the silicon die layer.

- Method 2:

By dividing the total current delivered to a metal layer by the number of power grid sections.



## d) Total Interconnect self-heating power in each metal layer

- Self-heating power of a metal layer  $i$  is given by

$$P_{\text{self}_i} = P_{\text{wire\_sig}_i} \cdot n_{\text{sig}_i} + P_{\text{wire\_pwr}_i} \cdot n_{\text{pwr}_i}$$

$P_{\text{wire\_sig}_i}$  and  $P_{\text{wire\_pwr}_i}$  are the self-heating power of each individual signal interconnect and power supply wire for metal layer  $i$ ,

$n_{\text{sig}_i}$  and  $n_{\text{pwr}_i}$  are the number of signal interconnects and power supply sections in metal layer  $i$ .

# Accuracy concerns about the interconnect power and thermal model

- Usefulness of the Interconnect model
- Accuracy concern of Rent's Rule
- Concern about current loading accuracy.



# Computation Speed of Hotspot CTMs

- In the order of milliseconds to minutes. (AMD MP 1.5 GHz Dual-Processor)

| simulated time interval | CPU time |
|-------------------------|----------|
| 0.1 ms                  | 20 ms    |
| 20 ms                   | 100 ms   |
| 2 sec                   | 7.9 sec  |
| 20 sec                  | 78.5 sec |
| steady-state            | 14.9 sec |

- The small overhead is due to the manageable number of nodes in the lumped RC circuit and
- Because of the use of first-order differential equations to iteratively solve the RC network.

# Validation

- An FPGA based system with 6 functional blocks and temperature sensors is implemented.
- The errors between Hotspot model and the thermal sensor measurements is found to be within 10%.

| Unit     | Power(mW) | Sensor Temperature | HotSpot Temperature |
|----------|-----------|--------------------|---------------------|
| blank1   | 0.1       | 3.4                | 3.37                |
| left_ppc | 75        | 3.5                | 3.69                |
| bott_ppc | 75        | 3.4                | 3.67                |
| ppc      | 45        | 3.5                | 3.66                |
| mb       | 313       | 4.1                | 3.96                |
| blank2   | 0.1       | 3.4                | 3.38                |

# Hotspot Applications

- Temperature estimations are used for temperature-aware design.
- The CTMs have been used to leakage power calculations.
- Used to explore different Dynamic Thermal Management Techniques (DTM).
- For an accurate interconnect lifetime predictions.



# Floorplan

- Use available floorplans in the tool
- Manually adapt the floorplan
- Use a graphic tool such as QUILT
  - <http://www.ece.rochester.edu/research/acal/quilt/>

# Example:

## Power/Performance/Thermal trade-offs in microarchitecture



# Main Goals

- Exploration of the power/thermal behaviour of architectures with value compression
- Compress values in all microarchitectural blocks
  - Size compression
  - Significance compression
  - Zero compression

Single core, 4 wide issue, 64KB L1 cache, 2MB L2 cache

(Benchmark: Crafty -SpecInt2K-)



# Value compression styles


32-bit Value



Size Compression

Sign ext'n byte

Sign extension bit



Zero Compression

zero byte

Zero extension bit



Significance Compression


Sign ext'n byte

Sign extension bit

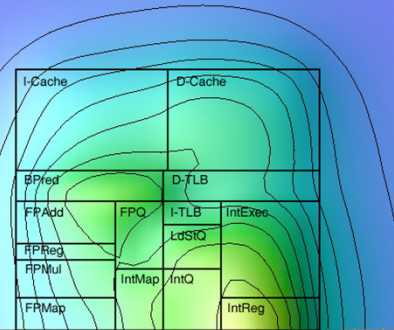


# Thermal behaviour


*No Compression*





*Size Comp.*



*Zero Comp.*



*Sign Comp.*



Configurations used:

Size 8-16-32-64

Zero 8-16-24-32-40-48-56-64

Sign 8-16-24-32-40-48-56-64

# Leakage models

- Hotleakage (University of Virginia)
- Temperature dependant leakage model:
  - Weiping Liao, Lei He, and K.M. Lepak. Temperature and supply Voltage aware performance and power modeling at microarchitecture level. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 24(7):1042, July 2005.

# Conclusions

- Framework for power/temperature/leakage evaluation of microarchitectures:
  - Wattch/McPAT + HotSpot + Leakage model
- But... beware of:
  - Integration: All the tools should use the same technology parameters
  - Floorplan: needs to be updated if the structures change
  - Accuracy is limited (but accepted in the area)