ISLPED 2003
8/26/2003

Reducing Power Density through Activity Migration

Seongmoo Heo, Kenneth Barr,
and Krste Asanović
Computer Architecture Group, MIT CSAIL

Background

- **Hot Spots**
  - Rapid rise of processor power density
  - Uneven distribution of power dissipation
    - Blocks such as issue windows have more than 20x power density of less active block such as L2$
  - Reduced device reliability and speed, increased leakage current
- **Existing Solutions**
  - Packaging/cooling: high cost, not possible at laptop
  - Dynamic thermal management: performance loss
    - Total power dissipation must be reduced until all hot spots have acceptable junction temperature

Introduction

- **Activity Migration (AM) to reduce power density**
  - **With AM, we spread heat by transporting computation to a different location on the die**
  - **If one unit heats past a temperature threshold, the computation is transferred to a second unit allowing the first to cool down**
- **AM for lowering temperature and power or for doubling maximum power dissipation at a given package**

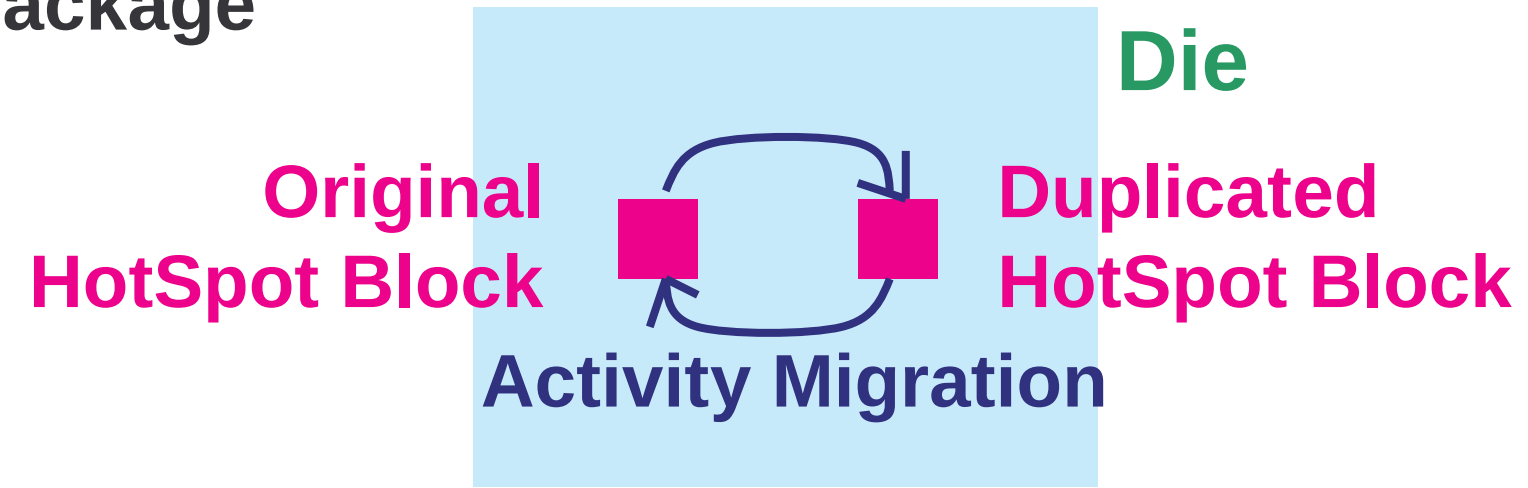

Die Thickness and Power Density

- **Two technology cases**
  - **180nm case: present, based on TSMC process**
  - **70nm case: near future, based on BPTM process**
- **Die thickness**
  - **Most heat is removed through back of die**
  - **Thinning chips: 250um → 100um**
  - **Increasing lateral resistance**
- **Power density**
  - **Ideal scaling → constant power density**
  - **Vdd scale-down slowed, clock frequency increase accelerated due to deep pipelining → power density increase: 5W/mm^2 → 7.5W/mm^2**

Equivalent RC Thermal Model

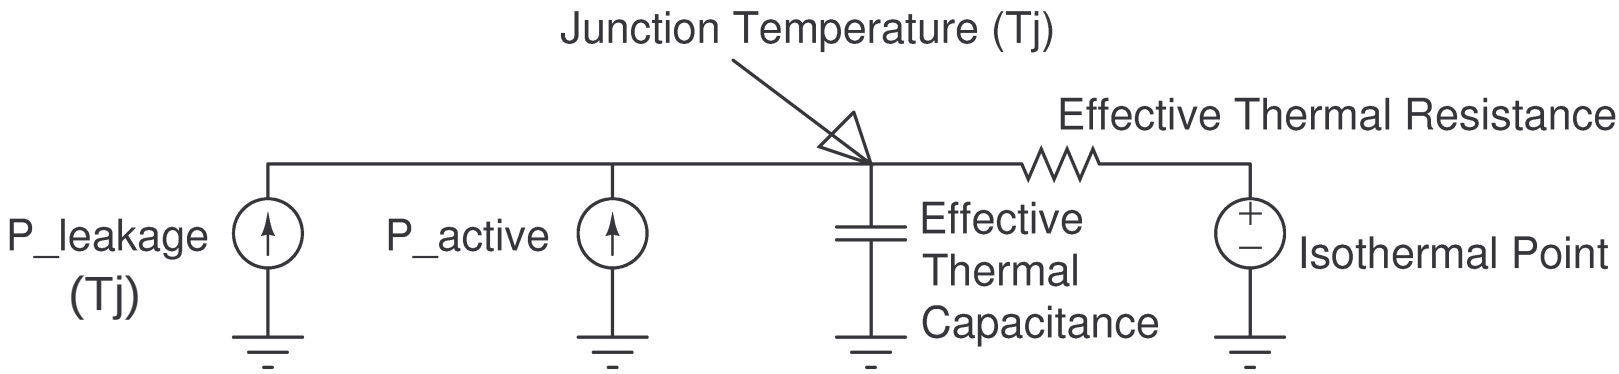

- **Equivalent RC Thermal Model:**
  - **temperature - voltage, power - current**
- **Thermal resistance: lateral resistance ignored**
- **Thermal capacitance: package capacitance modeled as a temperature source (isothermal point)**
- **Exponential dependence of leakage power on temperature modeled as voltage-dependent current source (P_leakage(Tj))**

Benefits of Activity Migration

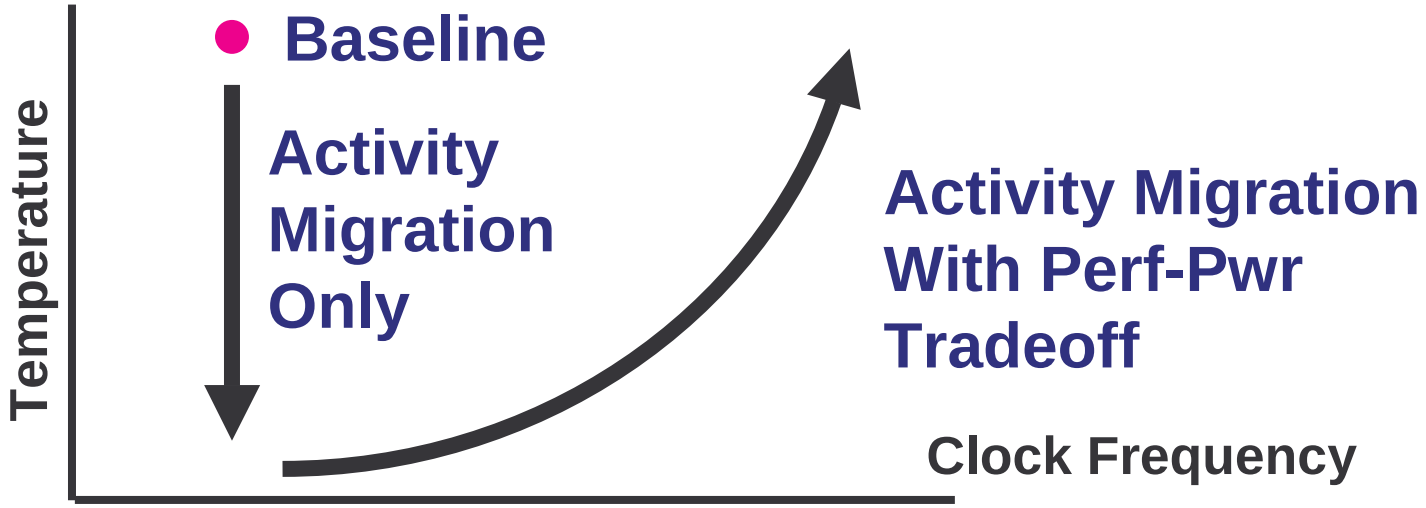

- **AM**: reduced temperature and power
- **AM + Perf-Pwr Tradeoff**: increased frequency and sustainable power
- Example: laptop with limited heat removal
  - Battery mode: AM Only: low temp, low leakage power → energy-efficient execution
  - Plugged mode: AM+Perf-Pwr Tradeoff: more power, more performance → max. performance execution without raising die temperature

Activity Migration Model

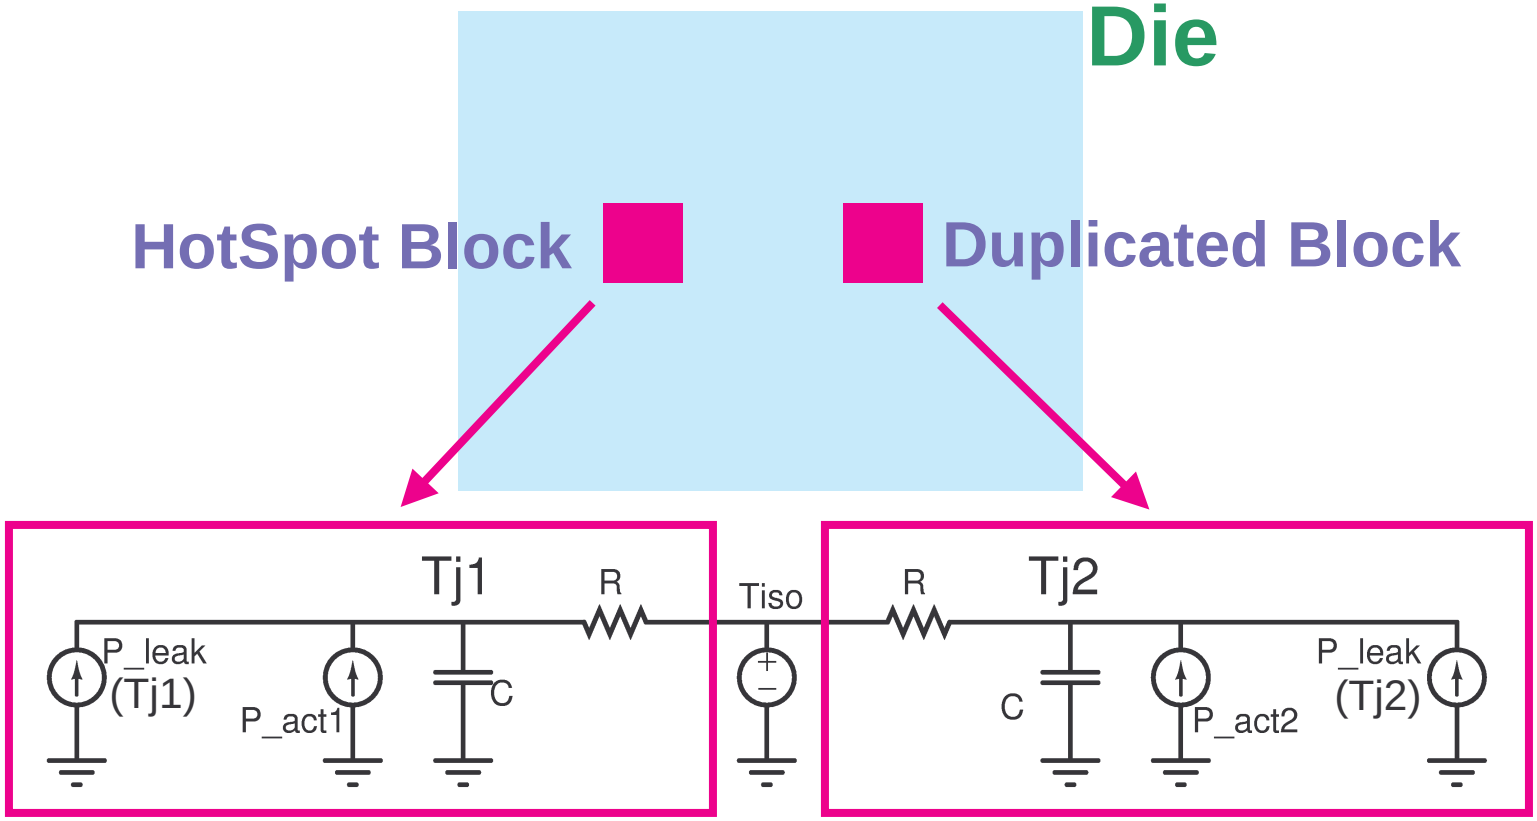

- **Activity Migration by turning on and off active power of hotspot and duplicated blocks (P_act1 and P_act2)**
- **Identical thermal resistance and capacitance**
- **Identical leakage power at same temperature**

AM Only

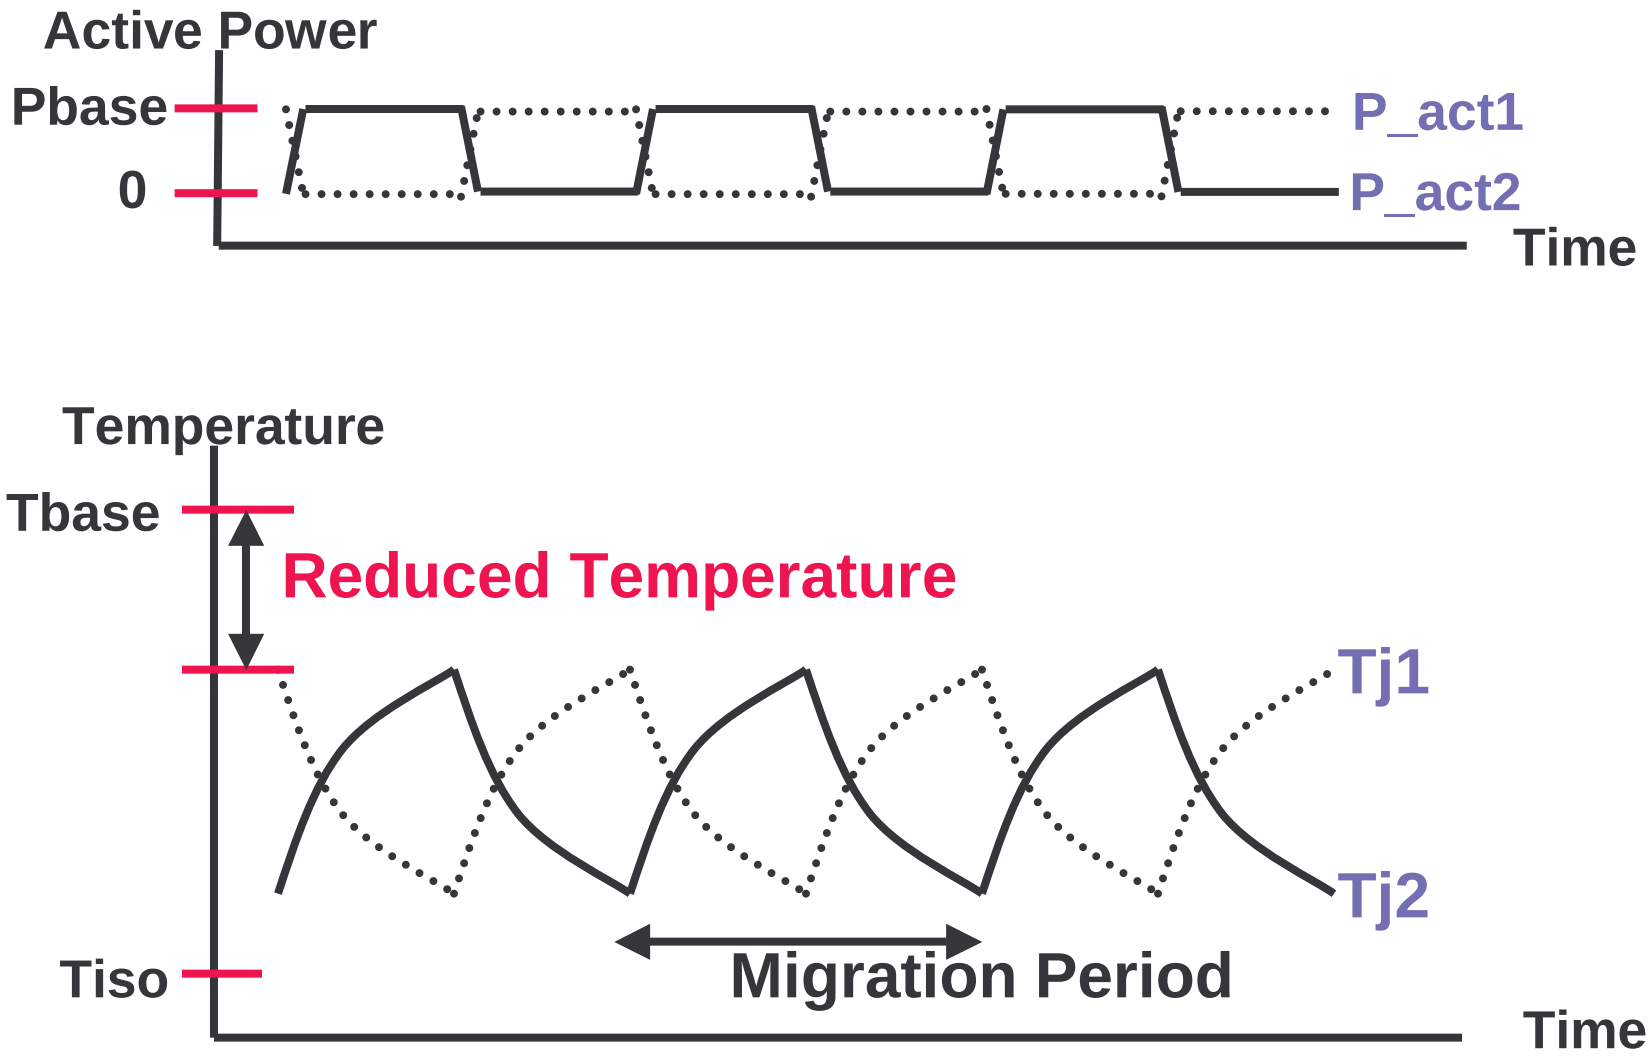
Active Power
Pbase
0
P_act1
P_act2
Time
Temperature
Tbase
Reduced Temperature
Tj1
Tj2
Tiso
Migration Period
Time

AM + Perf-Pwr Tradeoff

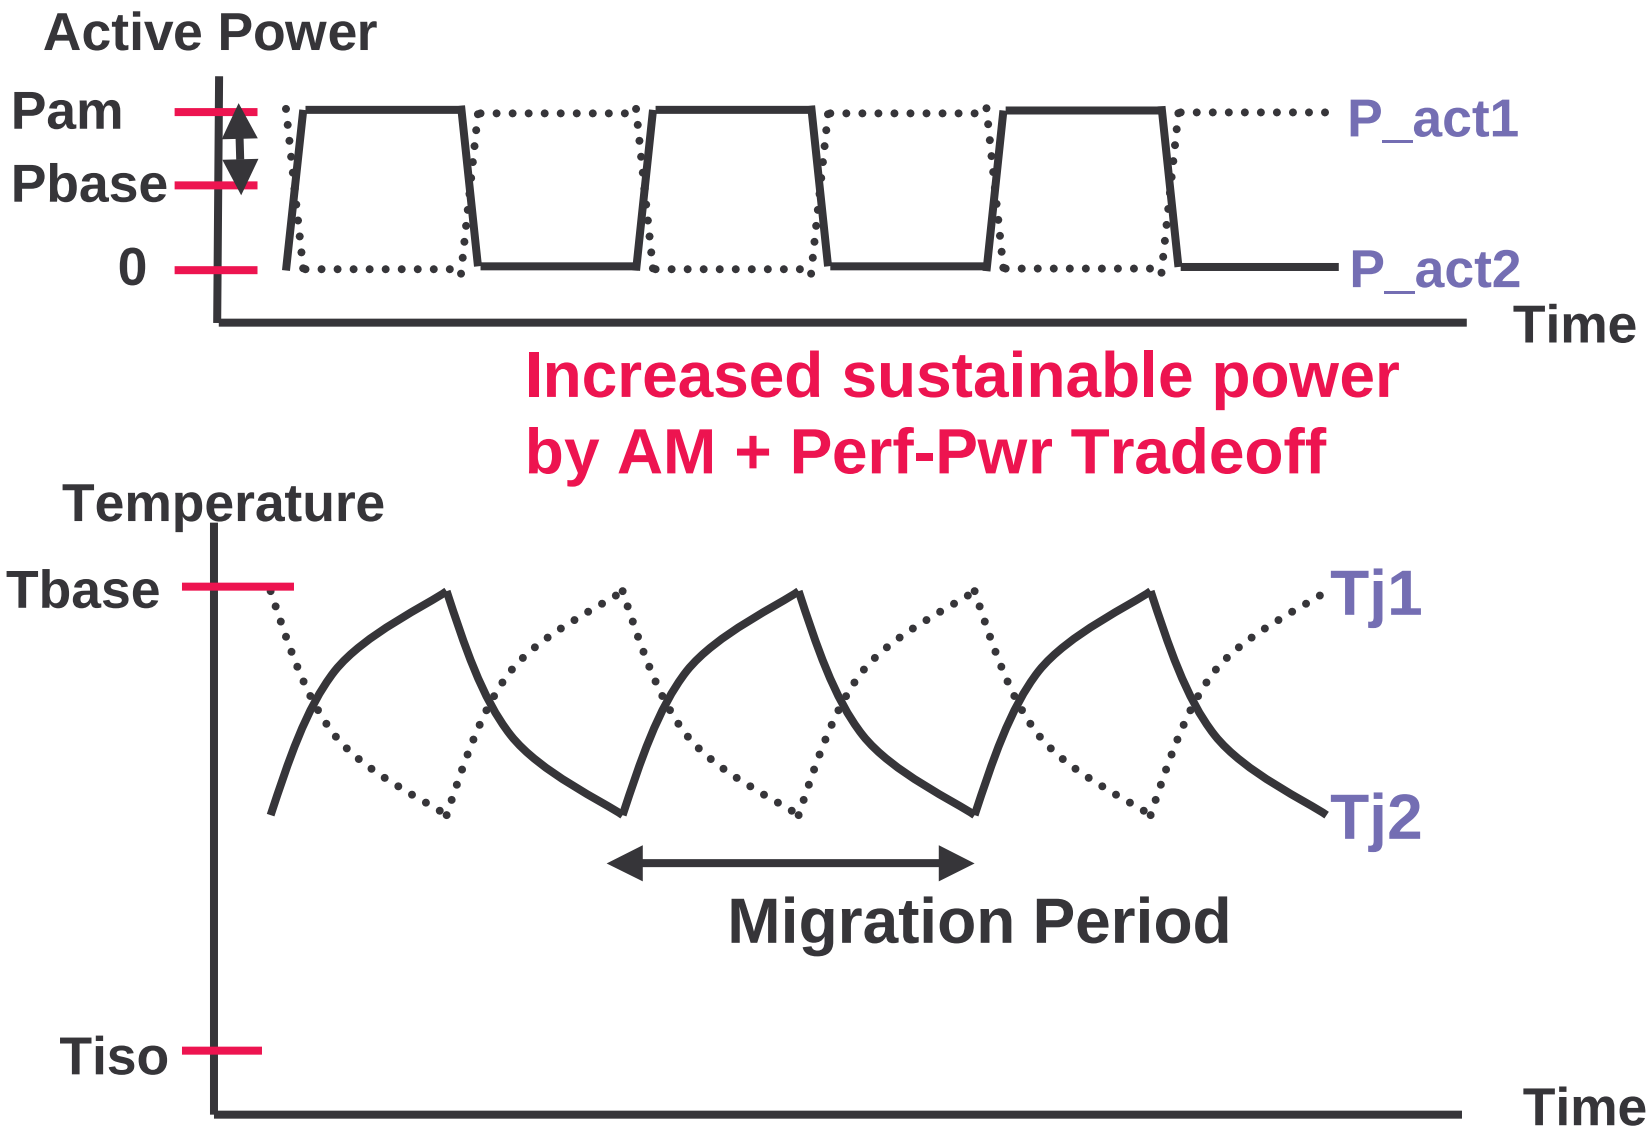

Migration Period: AM Only

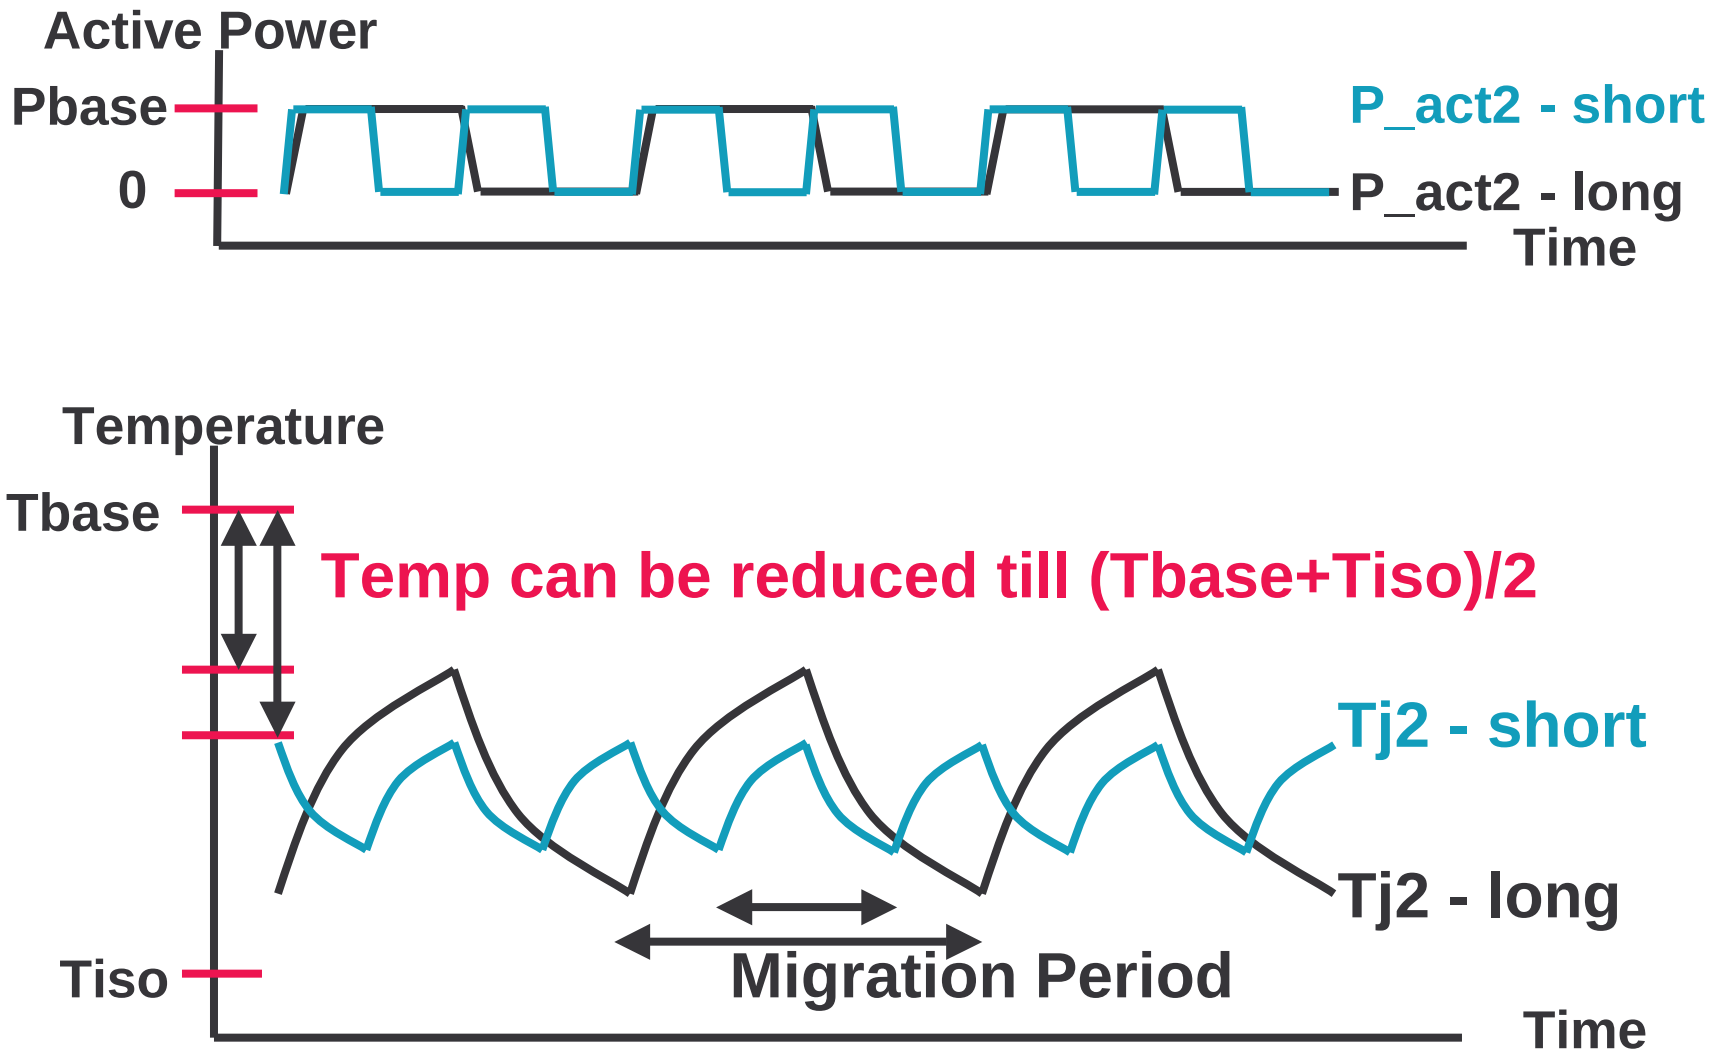

Migration Period: AM + Perf-Pwr Tradeoff

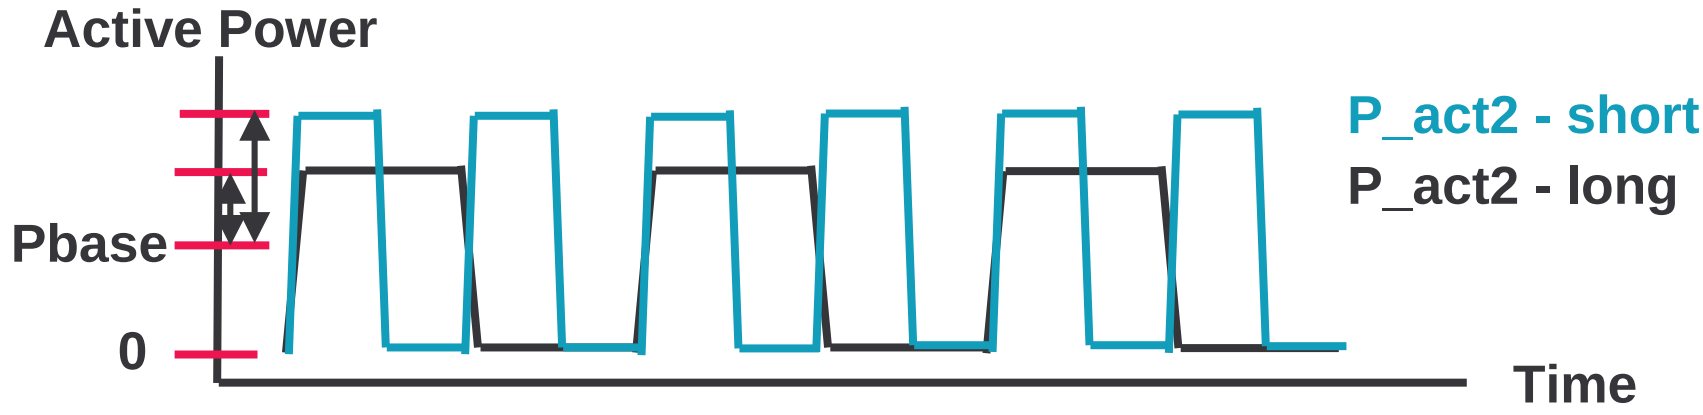

Sustainable power can be increased till 2*Pbase

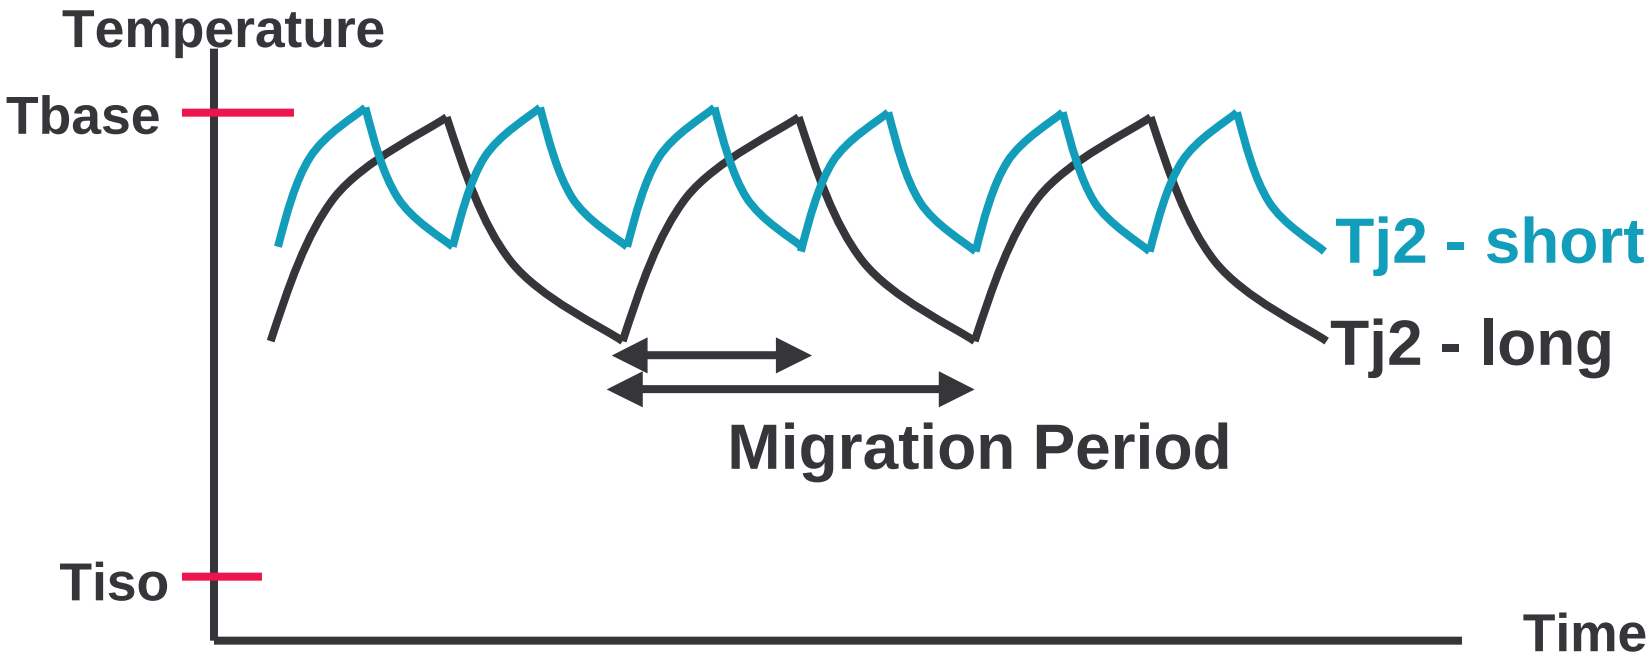

Effect of Migration Period

- **Small migration period**
  - **+ More temperature drop (More power increase)**
  - **- Greater CPI penalty**
  - **- AM in hardware: Hardware overhead**
- **Large migration period**
  - **+ Smaller CPI penalty**
  - **+ AM in software: OS context swap**
  - **- Less temperature drop (Less power increase)**

Simulation Results: AM Only

- Reduced temperature → reduced leakage power
- Reduced latency due to increased drain current at low temperature is exploited by reducing Vdd → reduced active power

	180nm Case			70nm Case		
Migration period (μs)	1800	600	200	600	200	60
Temperature drop (K)	9.2	11.5	12.4	3.4	6.4	7.5
Leak power reduction (%)	29.6	35.3	37.6	5.9	10.8	12.6
Act power reduction (%)	3.7	7.6	9.7	3.3	9.5	9.7

Simulation Results: AM+Perf-Pwr Tradeoff

- Same temperature as baseline
- Perf-Pwr Tradeoffs: DVS, dynamic cache configuration modification, fetch/decode throttling, or speculation control
- DVS chosen for Perf-Pwr Tradeoff due to its simplicity

	180nm Case			70nm Case		
Migration period (μs)	1800	600	200	600	200	60
Freq increase (%)	10.5	14.1	15.9	2.3	5.0	5.9
Power increase (%)	56.8	79.5	90.9	25.0	61.4	79.6

AM Architecture Configuration

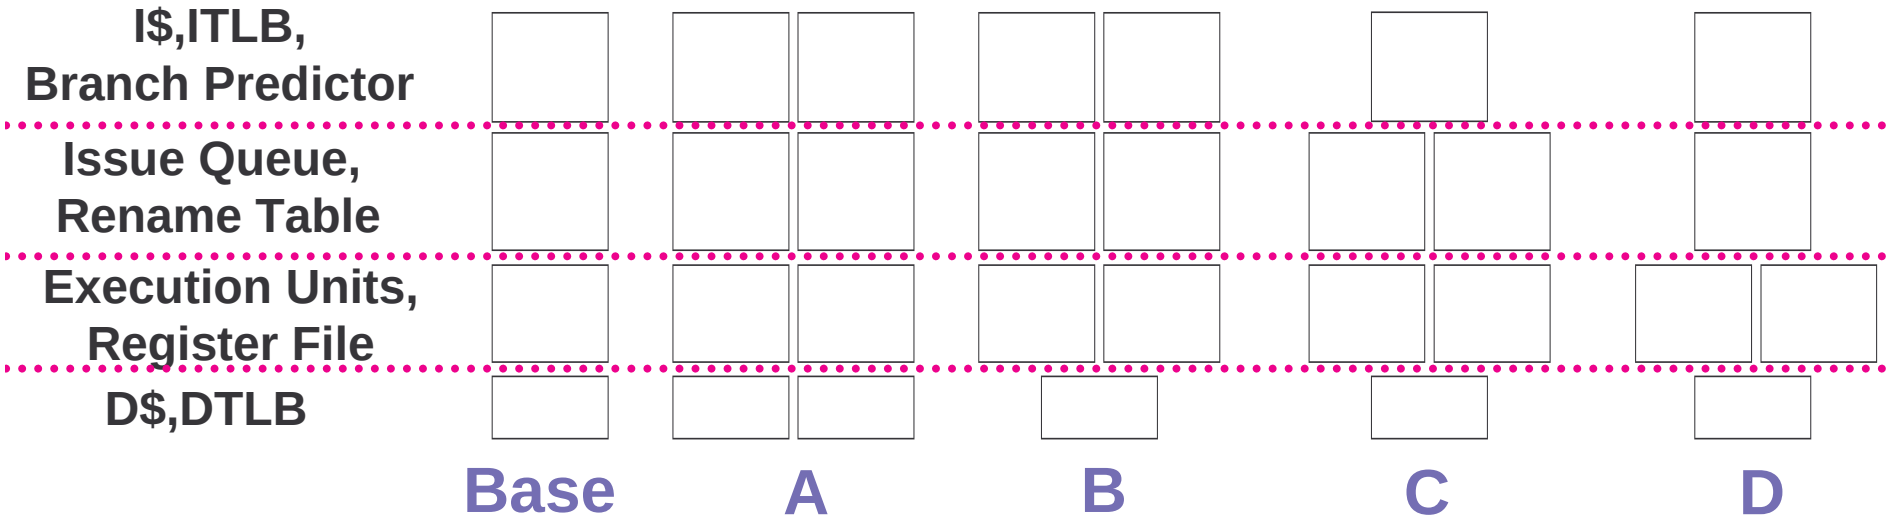

- Base: block areas based on Alpha 21264 floorplan
- Hotspot blocks: execution units and register file
- Pessimistic CPI penalties of AM
  - Cycle penalty due to increased wire latency when sharing a block: e.g. Shared D$ → extra cycle to cache access time
  - Migration penalty: draining and copying

Performance Effects of AM

- **Methodology**
  - **4-wide 32-bit superscalar machine**
  - **SimpleScalar 3.0b**
  - **SPEC2000 benchmarks using SimPoints**
- **Migration Period**
  - **Short migration period chosen: 200K cycles (200μs for 180nm case and 60 μs for 70nm case)**

Only 0~3% CPI penalty on average even at short migration period

Effects of AM for Area and Net Perf

	180nm Case				70nm Case			
Conf	A	B	C	D	A	B	C	D
Area	2.00	1.84	1.56	1.30	2.00	1.84	1.56	1.30
Speed	1.16	1.13	1.12	1.12	1.06	1.04	1.03	1.03

•normalized to baseline, speed = clock freq / CPI

- 180nm Case: conf. D achieves 12% performance gain with 30% area increase
- 70nm Case: performance gain relatively small → AM only to cool down hot spots
- Other issues
  - Extra power for driving increased wire lengths
  - Migration triggering by thermal sensors rather than fixed migration periods

Conclusion

- Activity Migration (AM) was proposed to solve hotspot problem of modern microprocessors
- AM spreads heat by transporting computation to a duplicated block
- AM can be used in two ways
  1. AM only: low temperature, low leakage
  2. AM + Performance-Power Tradeoff: sustainable power and performance increase
- Dynamic fixed-period AM was evaluated on a superscalar machine
  - 12.7 degree temperature reduction
  - 12% clock frequency increase with 3% CPI penalty and 30% area increase

Acknowledgments

- Thanks to Christopher Batten, Ronny Krashinsky, Heidi Pan, and anonymous reviewers

- Funded by DARPA PAC/C award F30602-00-2-0562, NSF CAREER award CCR-0093354, and a donation from Intel Corporation.

BACKUP SLIDES

Thermal and Process Properties

	Symbol	Current Case	Future Case
Die thickness (μm)	T	250	100
Die conductivity (W/K/m)	K	100	100
Die specific heat (J/K/m^3)	C	1e6	1e6
Die area (mm^2)	A_{die}	100	100
Hot spot area (mm^2)	A_{block}	2	2
Hot spot active power density (W/mm^2)	PD_{act}	5	7.5
Hot spot leakage power density (110°C) (W/mm^2)	PD_{leak}	0.015	0.15
Isothermal point (°C)	T_{iso}	70	70
Channel length (nm)	L	180	70
Supply voltage (V)	V_{DD}	1.5	1.0
NMOS threshold voltage (V)	NV_{th0}	0.269	0.120
PMOS threshold voltage (V)	PV_{th0}	-0.228	-0.153

* Transistor models: TSMC 180nm and BPTM 70nm processes

Equivalent RC Thermal Model

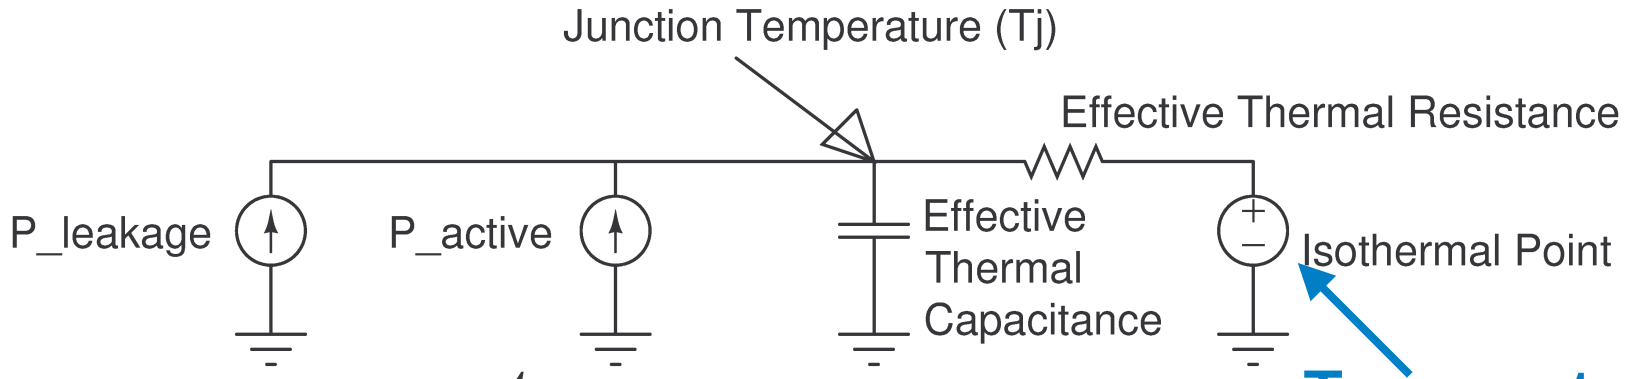

Temperature source in packaging

- $R_{silicon,\,vertical} = \dfrac{t}{k \times A_{block}}$
- $R_{package,\,vertical} = 120 \times \dfrac{t}{k} \times \dfrac{A_{die}}{A_{block}}$
- $R_{total,\,vertical} = (1 + 120 \times A_{die}) \times \dfrac{t}{k \times A_{block}}$
- $C_{silicon} = c \times t \times A_{block}$

***Empirical formula from 3D simulation results [Barcella02]**

Exponential dependence of leakage power upon temperature modeled by voltage-dependent current source

Temperature Dependency of Leakage

- **Leakage power**
  - -Significant part of total power
  - -Exponential dependence upon temperature
  - -Voltage-dependent current source

$$P_{leak} = P_{leak110} \times e^{\beta(Tj-110)}$$

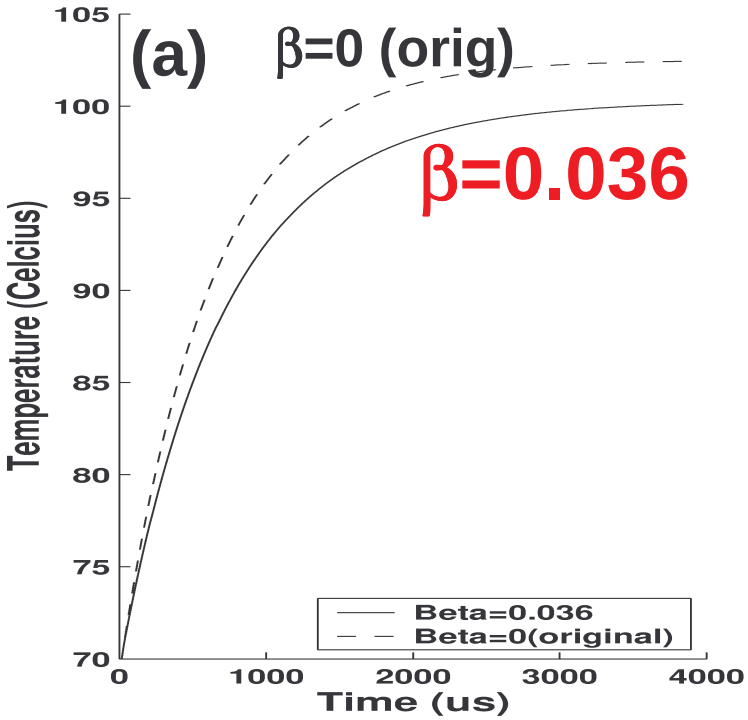

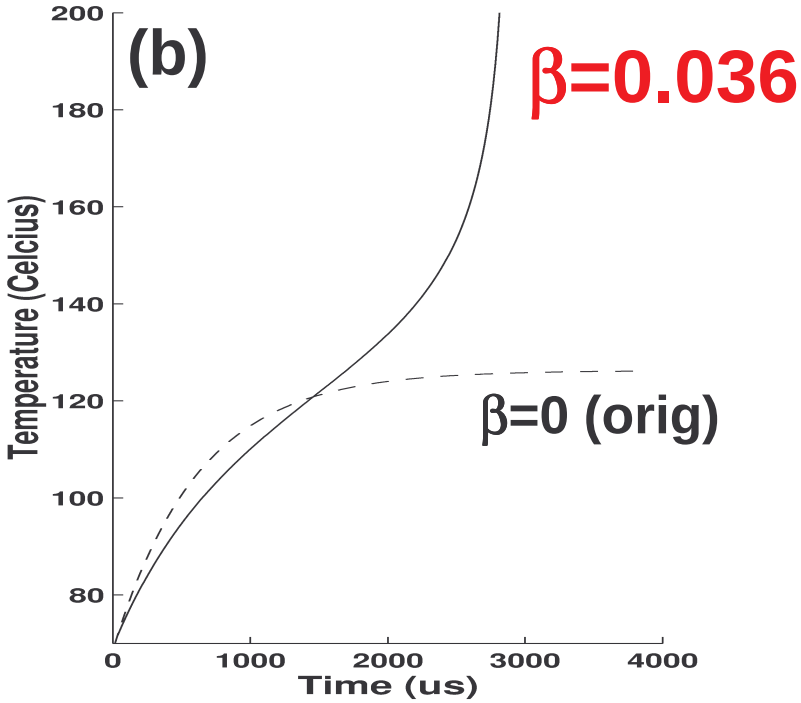

AM Model

HotSpot Block

Duplicated Block

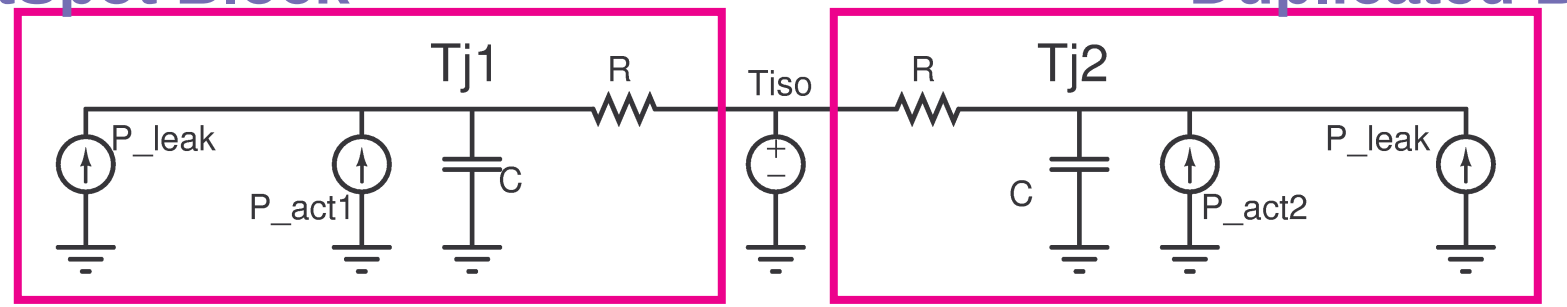

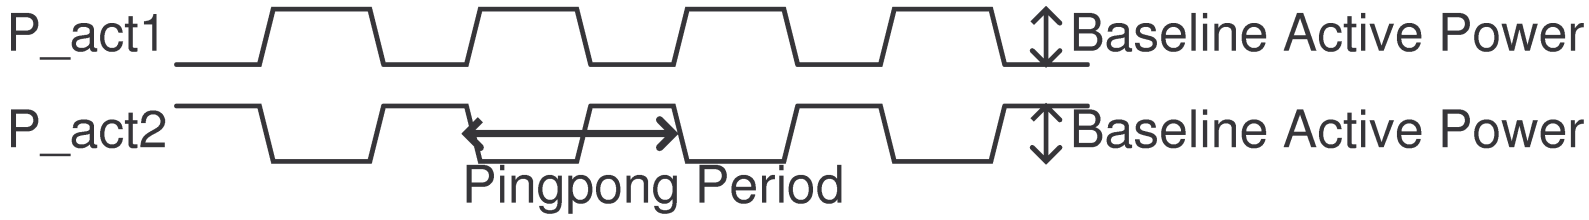

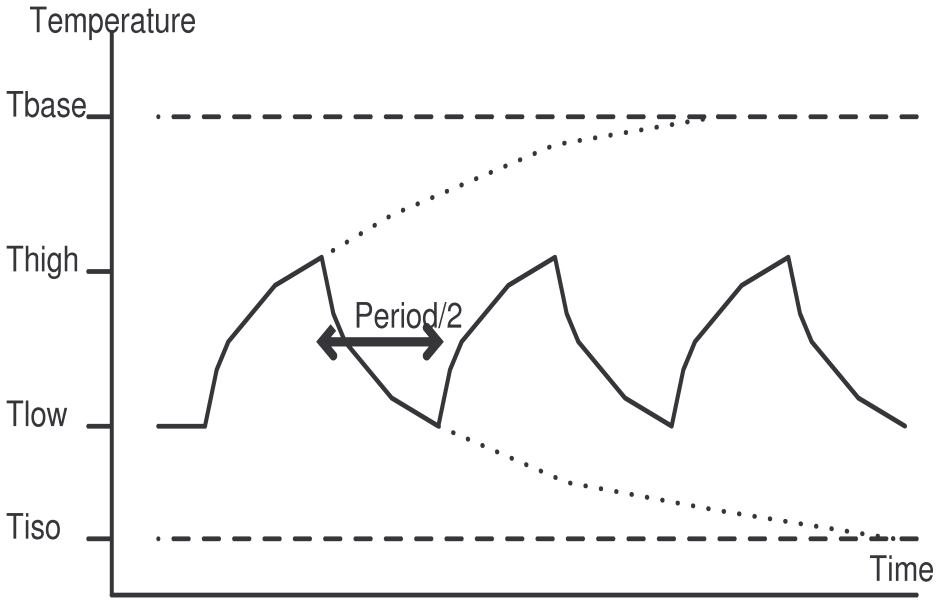

$$T_{high} = \frac{T_{base} - T_{iso}}{1 + e^{-\frac{Period}{2\tau}}} + T_{iso}$$

→ 2

If period is small enough,

- **Halve temp increase**
- **Double sustainable power**

AM Simulation Results: AM + DVS

AM and DVS for various pingpong periods for the hot spot block (Current case)

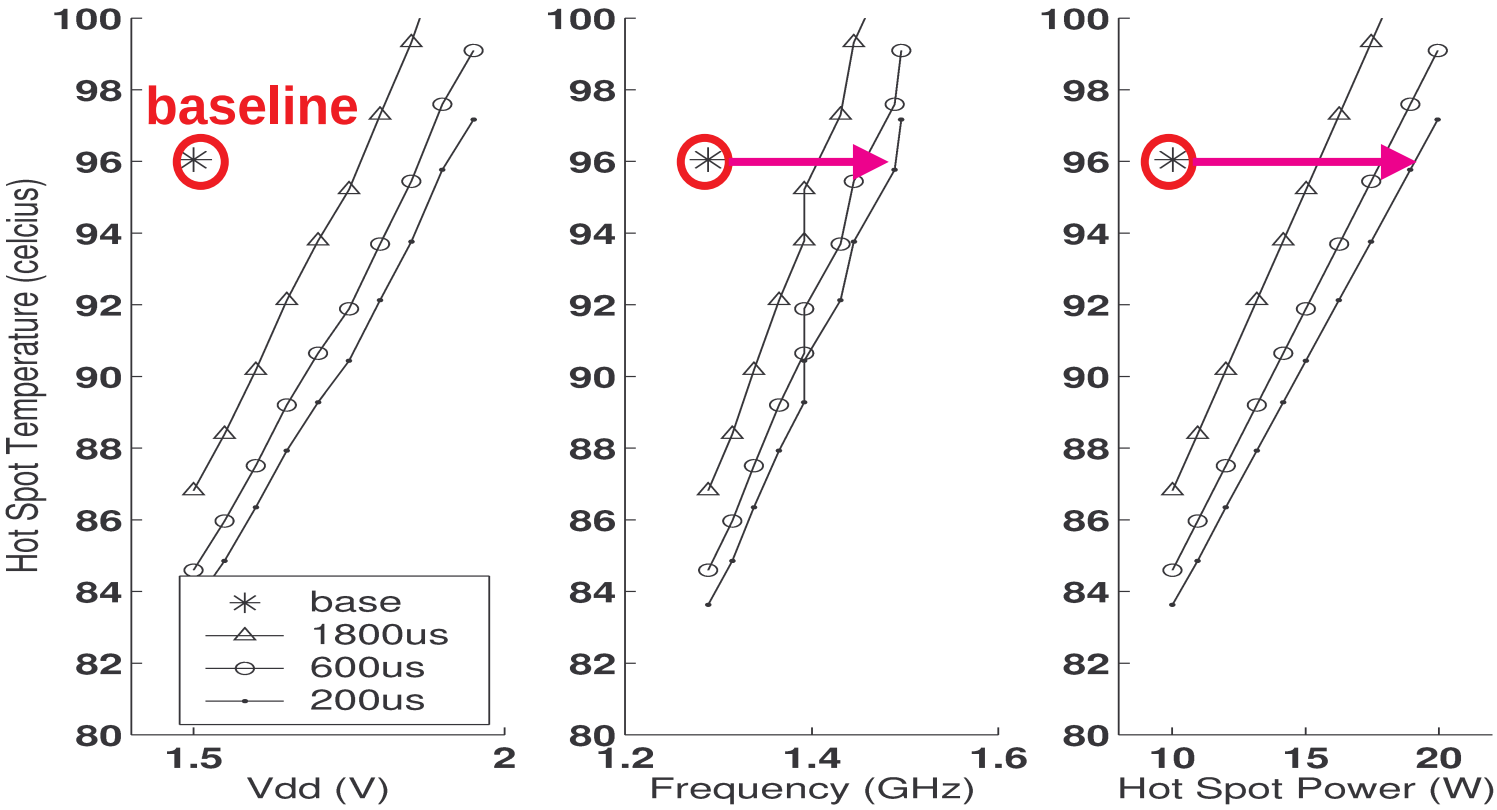

DVS effects were modeled based on Hspice simulation of a 15-stage ring-oscillator

AM Simulation Results: AM + DVS

AM and DVS for various pingpong periods for the hot spot block (Future case)

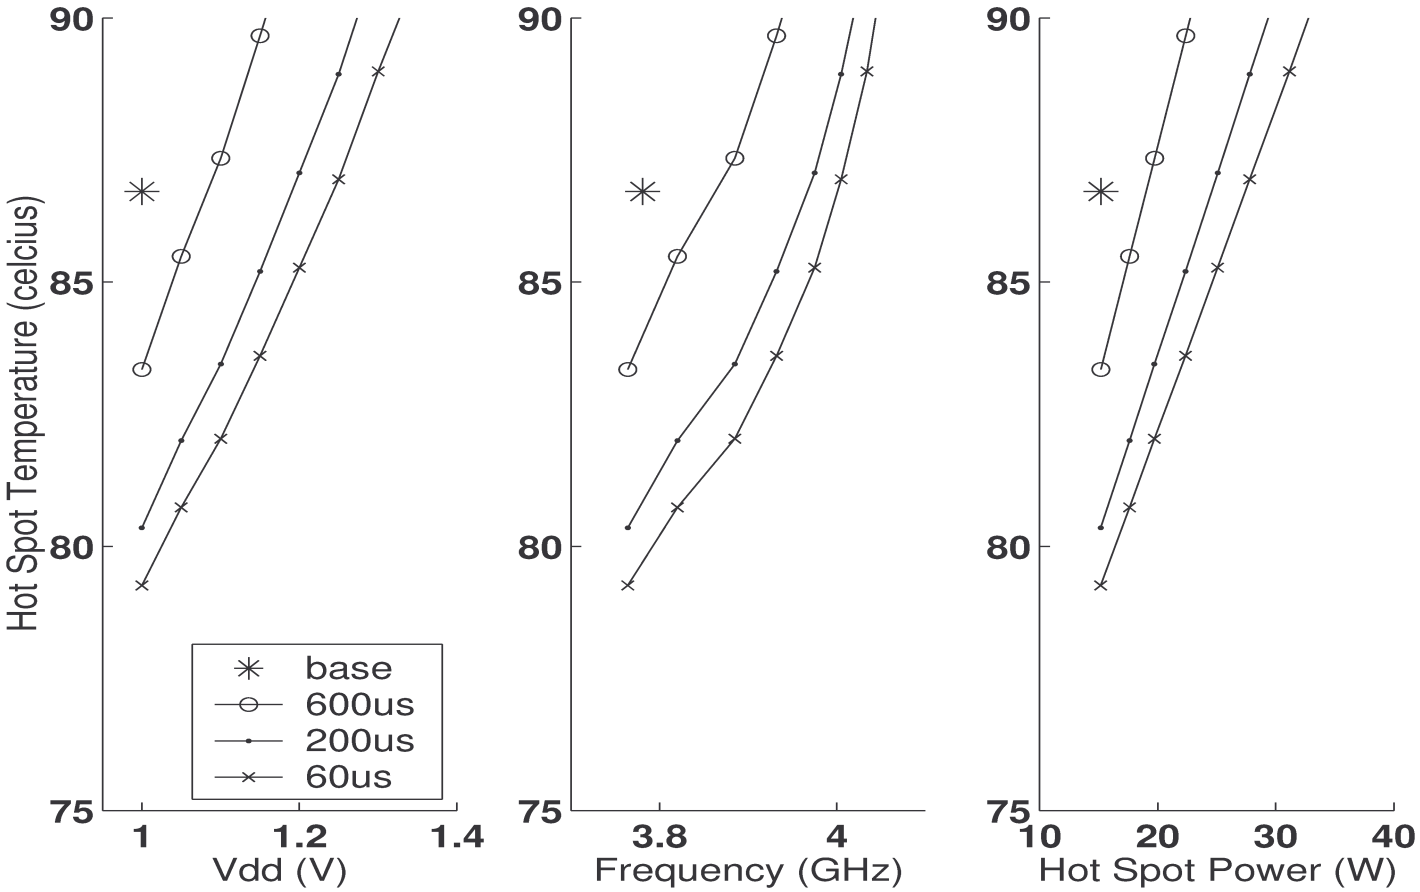

Performance Effects of AM

- 4-wide 32-bit superscalar machine
- SimpleScalar 3.0b
- SPEC2000 benchmarks using SimPoints
- Short migration period chosen: 200K cycles (200μs for 180nm case and 60 μs for 70nm case)