Active Bank Switching for Temperature Control of the Register File in a Microprocessor

Kimish Patel, Wonbok Lee, Massoud Pedram

University of Southern California
Los Angeles CA
March 12th 2007

Introduction

- With increasing power density, the on-chip steady state temperature is increasing
  - Necessitating the use of temperature control mechanisms
- Dynamic thermal management (DTM)
  - Achieve thermally safe state of a microprocessor with minimal performance degradation
- Some previous DTM schemes
  - Reactive:
    - Fetch Toggling
    - Instruction Cache Throttling
    - Dynamic Instruction Window Resizing
    - Activity Migration (AM)
  - Predictive:
    - Workload prediction based dynamic voltage and frequency scaling

Motivation

- **Physical register files**
  - Multi-ported SRAM structure
  - The higher the issue width, the higher the number of ports
- **With technology scaling**
  - Power density in register file increases
  - Temperature becomes a critical problem for register files
- **Register file is one of the hottest on-chip structure**

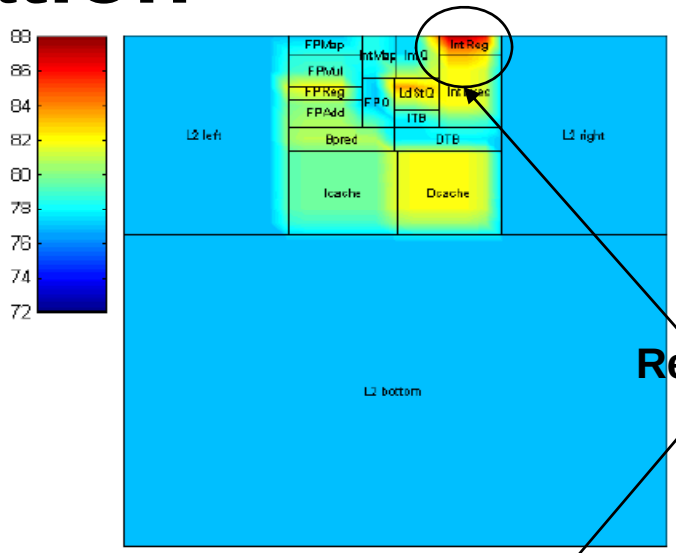

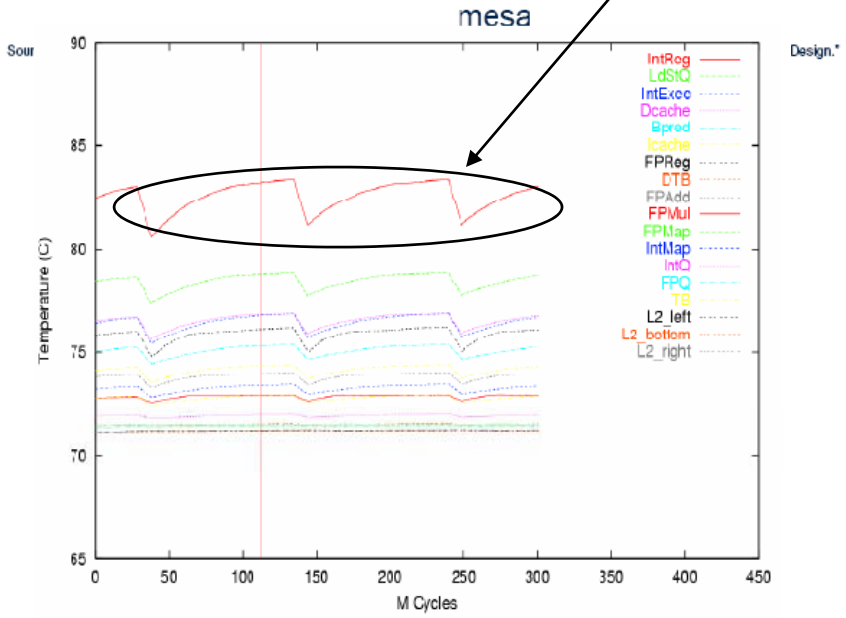

Source : Mircea Stan

Motivation: Register File Utilization

- **Modern Superscalar processors:**
  - **With 32-bit ISA have 32 architectural registers**
  - **To deal with high issue widths:**
    - **More physical registers are available, typically 64, 128, etc.**
- **Physical register file utilization:**
  - **Is however typically low for major portion of the program run**
    - **Due to intrinsic data dependency among instructions**

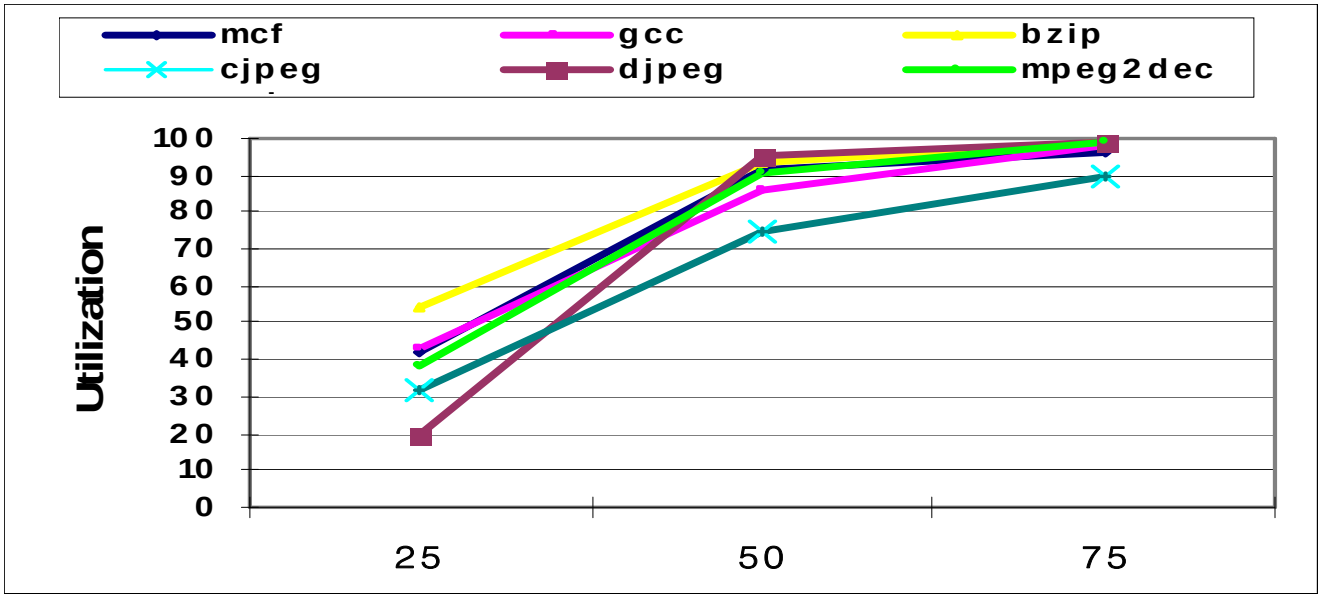

Active Bank Switching: The Idea

- **Based on the observations:**
  - **Only 50% of the physical registers are in use for more than 70% of the program execution time**
- **Idea:**
  - **Split the register file in two banks**
  - **Only one is active at a time (denoted as primary bank) while the other one is inactive (denoted as secondary bank)**
  - **New register allocation is done only from the primary bank**
  - **Once the active banks are switched:**
    - **Some live register references may still point to the previously active bank**
    - **However, number of such references are relatively small**

Active Bank Switching: Thermal Behavior

- **Monolithic bank architecture**
  - Single bank, higher power
  - Continuously active and continuously accessed
  - Relatively high steady state temperature (T_{SS})
- **Active bank switching architecture**
  - Dual banks, lower power
  - Only one bank active at a time
  - Relatively low T_{SS}

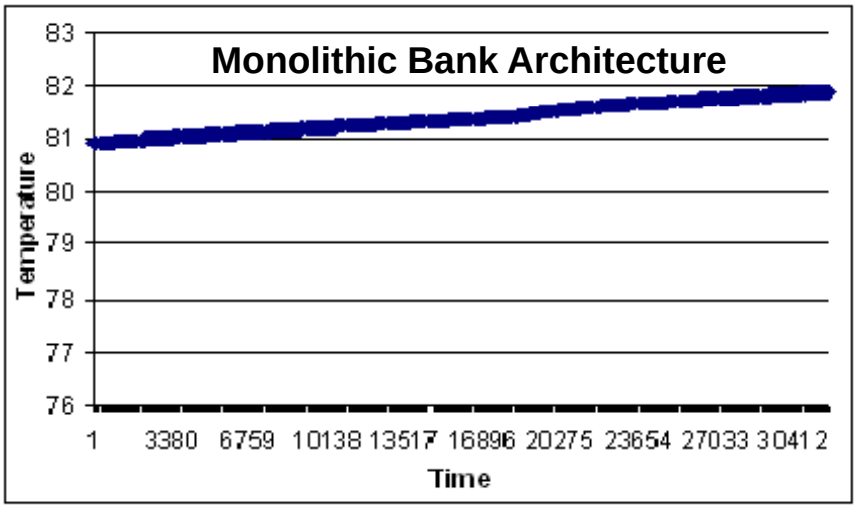

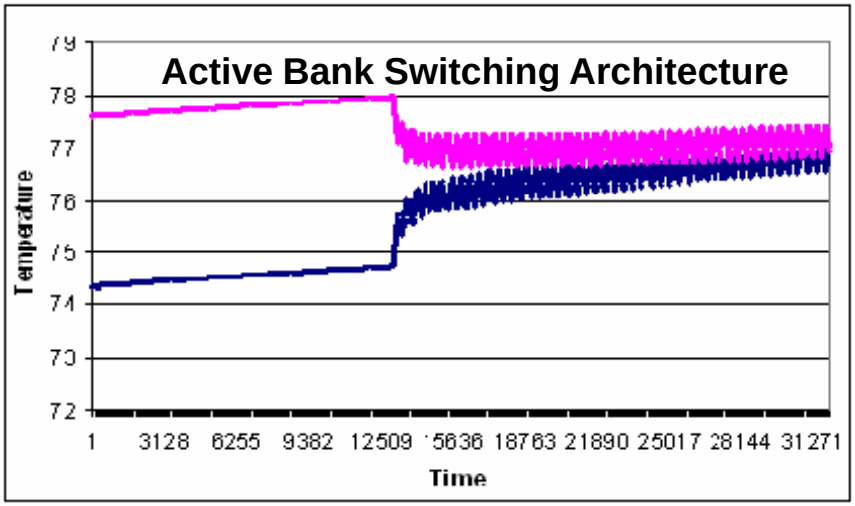

Thermal Model (1/2)

- **Adopt the thermal model proposed by Skadron et al:**

$$\Delta T = \left(\frac{P}{C_{th}} - \frac{T_{old}}{R_{th} \cdot C_{th}}\right) \cdot \Delta t$$

ΔT **: Temperature variation in *Δt* time interval**

P **: Average power**

R_{th}**: Thermal resistance**

C_{th}**: Thermal capacitance**

T_{old}**: Initial temperature at the beginning of the current time interval**

- **While the bank is active:**
  - **Bank's rising thermal gradient is calculated as:**

$$\text{Rising: } \frac{\Delta T_r}{\Delta t} = \left(\frac{P}{C_{th}} - \frac{T_{old}}{R_{th} C_{th}}\right)$$

- **While the bank is inactive:**
  - **Bank's falling thermal gradient is calculated as :**

$$\text{Falling: } \frac{\Delta T_f}{\Delta t} = \left(-\frac{T_{old}}{R_{th} C_{th}}\right)$$

- **The rising and the falling thermal gradients can be very different**
  - **Careful selection of the active period is crucial**

Thermal Model (2/2)

- **Determine the temperature at which the rising and falling thermal gradients are equal**
  - **Point where equal amount of time spent in temperature rise and fall:**
    - **Given the same initial/old temperature (T_{old})**
  - **We denote this point as the break-even temperature (T_{BE})**

$$\frac{\Delta T_f}{\Delta t} = \frac{\Delta T_r}{\Delta t} \Rightarrow \left(\frac{P}{C_{th}} - \frac{T_{BE}}{R_{th} \cdot C_{th}}\right) = \left(\frac{T_{BE}}{R_{th} \cdot C_{th}}\right) \Rightarrow \; T_{BE} = \frac{1}{2} \cdot P \cdot R_{th}$$

- **Break-even temperature, T_{BE}, determines the temperature region (i.e. any temperature $T > T_{BE}$) in which the active bank switching is beneficial**

Temperature Zones and Thermal Gradients

- Thermal behavior of a program can be divided into two regions:

- Fast Temperature Fall (FTF) Zone:
  - Active bank switching is quite effective when applied (i.e. triggered) in this region
  - Less performance penalty since temperature drops faster than it rises

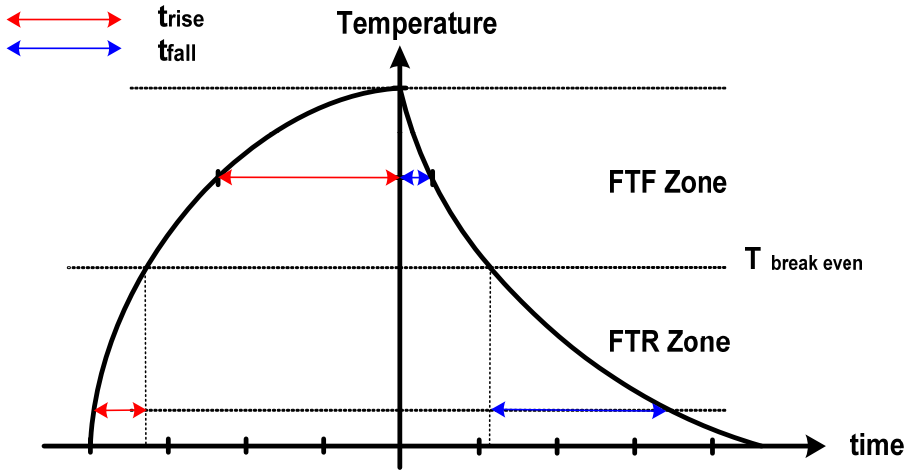

- Fast Temperature Rise (FTR) Zone:
  - Active bank switching is expensive when applied in this region
  - High performance penalty since temperature rises faster than it drops
- Our scheme:
  - Periodic switching (every 10M cycles) which is sufficient to reduce T_{SS} since the original T_{SS} lies in FTF zone

Experimental Results (1/3)

- **Simulator**
  - SimpleScalar/PISA, combined with Wattch and HotSpot, is configured with issue width of 4
  - Thermal measurement every 50K cycles
  - Initial/Ambient temperature: 60/45℃
- **Floor-plan:**
  - Pentium IV 130nm floor-plan (Figure (a)) for with 128 entry of Integer Register File (Figure (b))
  - Our configuration:
    - 64 entry Integer Register File
    - Figure (c)
  - Banked Architecture:
    - Two 32 entry Integer Register Files.
    - Figure (d)

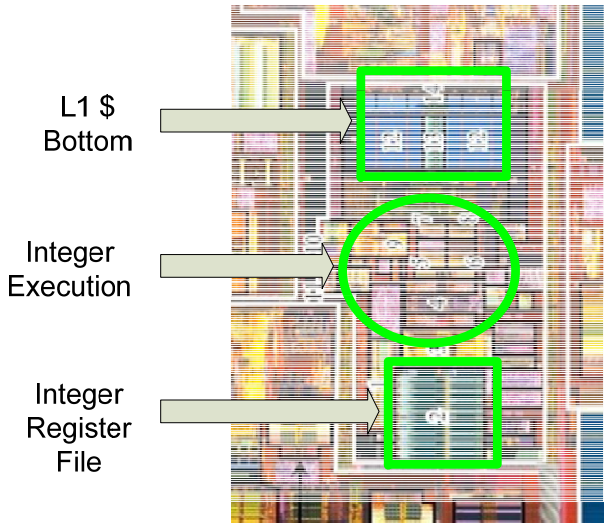

(a)

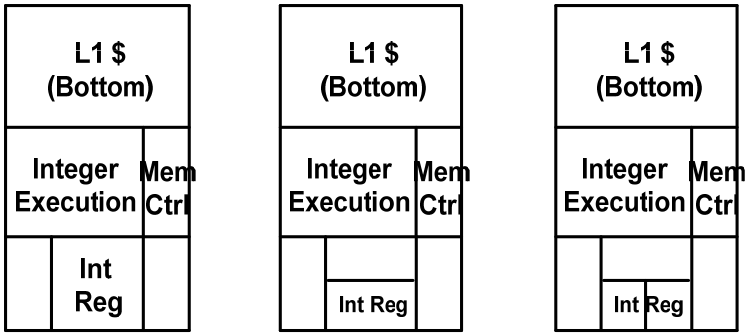

(b) (c) (d)

Experimental Results (2/3)

- Active bank switching reduces steady state temperature by 3.4℃ on average
- Average performance penalty of 0.75%

Program	Steady-state Temp (℃)		Thermal reduction (℃)	IPC
	Monolithic RF (64)	Banked RF (2*32)		
mcf	68.0	66.7	1.3	0.7707
gcc	76.5	73.7	2.8	1.2748
bzip	78.2	75.0	3.2	1.5022
gzip	81.7	77.5	4.2	2.1069
cjpeg	83.0	79.0	4.0	2.2553
mpeg2dec	82.0	77.7	4.3	2.2729
djpeg	82.0	77.0	5.0	2.3825

Table 3 Steady-State Temperature and IPC

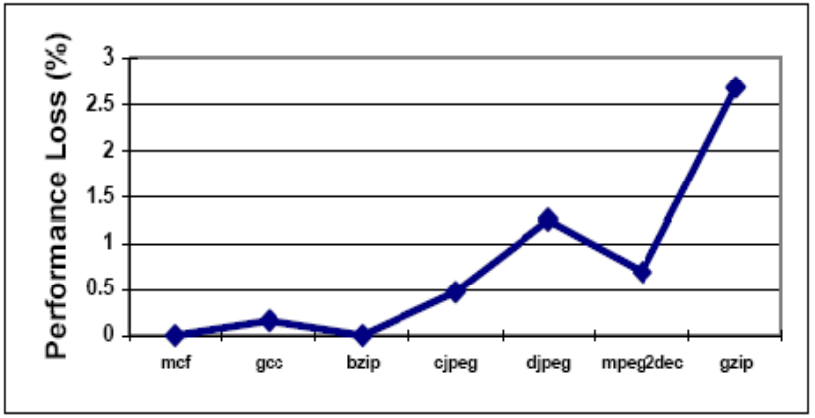

Figure 2 Performance Penalty in Half Sized Register Bank

Experimental Results (3/3)

- Figure shows dynamic thermal behavior of *gcc*
  - Upper curve: For monolithic register file
  - Lower curves: For active bank switching architecture
- Pattern of dynamic behavior remains the same. However the steady state temperature is lowered by active bank switching