Fine-Grain Power-Gating for Run-Time Leakage Power Reduction

Masaaki Kondo

The University of Electro-Communications

Background

- Power efficiency
  - A first class design constraint in today's computer systems
  - Key to achieve good performance/cost ratio even for supercomputers
    - Electrical power costs $$$
    - Size of the footprint
    - Reliability / availability

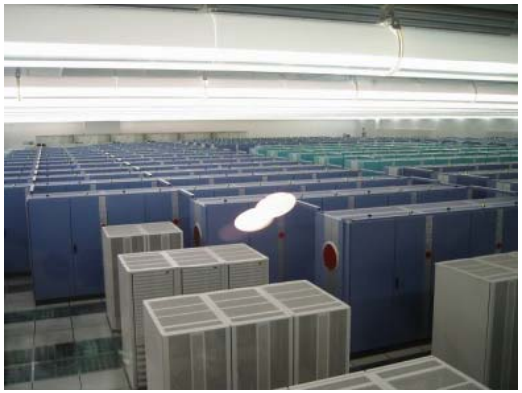

Current Research Project

- Innovative Power Control for Ultra Low-Power and High-Performance System LSIs
- 5 years project which started October, 2006
- Supported by JST (Japan Science and Technology Agency)
- Project Members (faculty members):
  - Prof. H. Nakamura (U. Tokyo): architecture [leader]
  - Prof. M. Namiki (T[illegible]kyo Uni[illegible])
  - Prof. H. Amano (Keio
  - Prof. K. Usami (Shibaura
  - Prof. M. Kondo (Univ. of

CP-PACS: The world fastest supercomputer in 1995

Leakage Power

- Leakage current
  - Flows even when circuits are idle
  - Expected to increase in the future
- Leakage reduction techniques
  - Standby time:
    - power-gating, DTCMOS, VTCMOS, …
  - Runtime:
    - Cache-decay, Drowsy-cache, …
- Leakage for logic parts (ALU, decode etc.) is still a matter

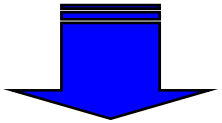

We are trying to reduce runtime leakage power of logic parts

Power-Gating

- Power-gating (PG) technique
  - Sleep transistors between GND and NMOS transistors of logics
    - Cuts power-supply to the logic blocks
    - Active/sleep mode controlled by sleep signals
  - Sleep transistors consist of High-Vth transistors

→ *significant leakage power reduction*

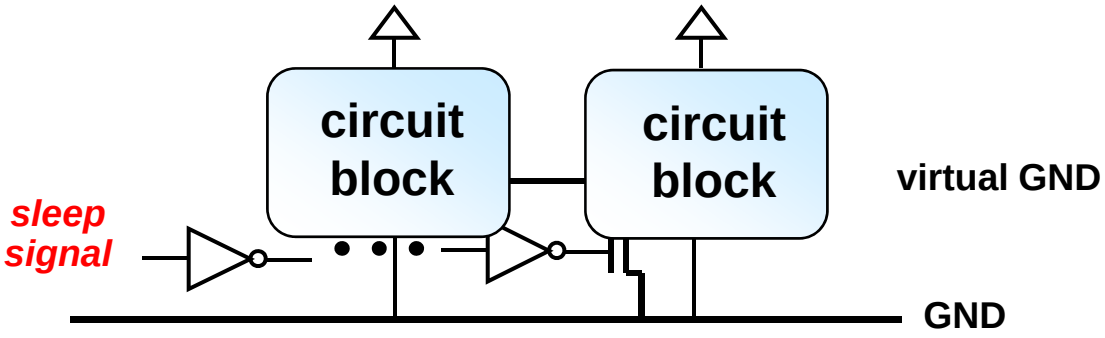

Strategies of Power-Gating

- Currently for a processor with 5-stage pipeline
- Two PG strategies
  1. Put each EX-unit in sleep mode operation-by-operation
  2. Put all the EX-units in sleep mode when cache misses occur

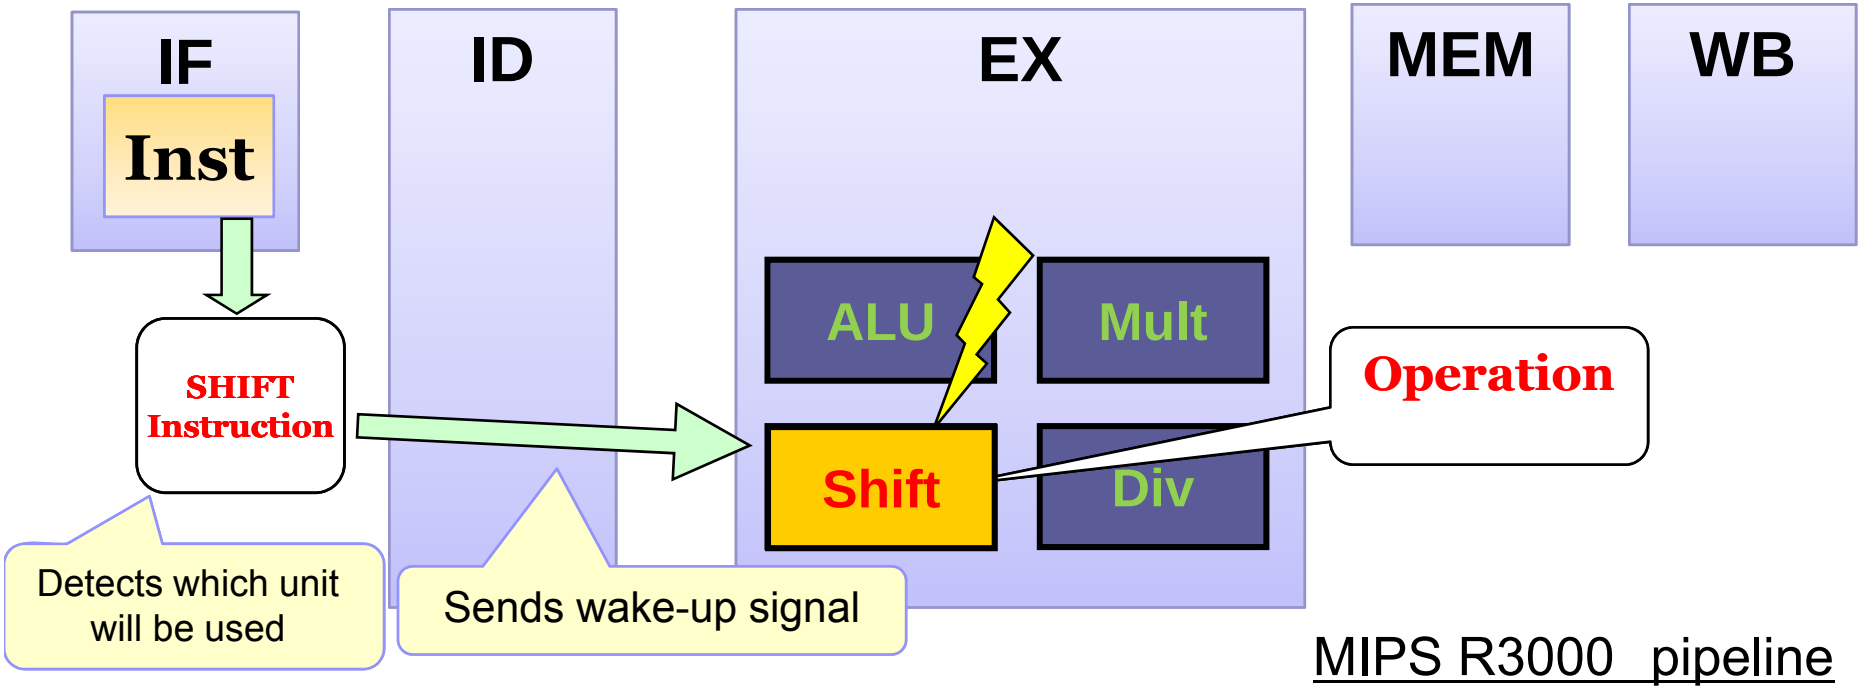

MIPS R3000 pipeline

Challenges of Run-Time Power-Gating

- Wake-up latency when activating a execution-unit
  - The unit cannot be used immediately after waking-up
  - Pre-wakeup before the unit is actually needed to hide the latency
  - In MIPS pipeline, this latency is perfectly hidden
- Energy Overhead for sleep-mode control
  - Caused by dynamic energy to control sleep transistors
  - Important to consider for fine-grain power-gating

Energy Overhead of Run-Time Power-Gating

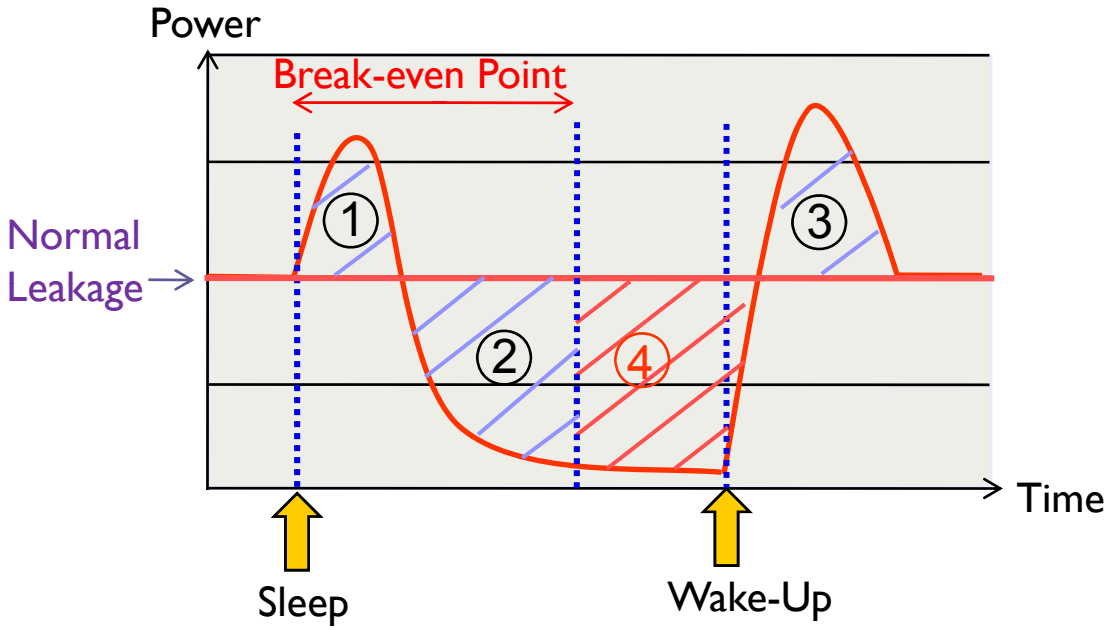

① + ③ : Energy overhead

② : part of leakage saving

① + ③ = ②

Break Even Point (BEP)

④ : Net Energy saving

- Sleep period should be longer than BEP
  - Otherwise, total energy consumption increases

Break Even Point of Each Unit

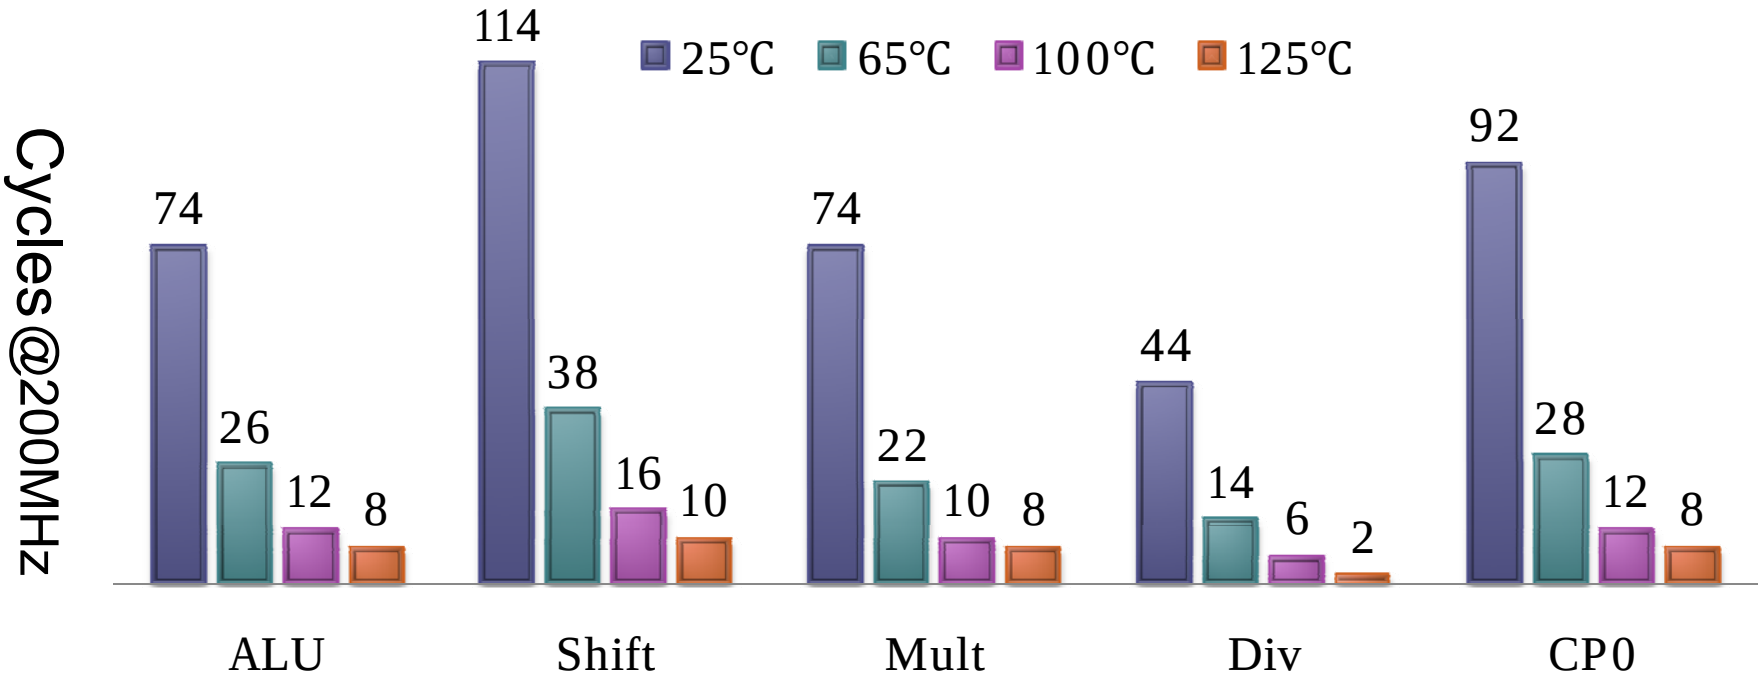

- BEP is short when the chip temperature is high
  - Leakage current is temperature dependent
- We need PG strategies with taking BEP length into account

PG Optimization by Compiler/OS

- Compiler and OS enable/disable Power-gating
  - Depending on expected sleep period and BEP
- Compiler based technique
  - Extended ISA
  - Statically analyzes sleep period from instruction sequence

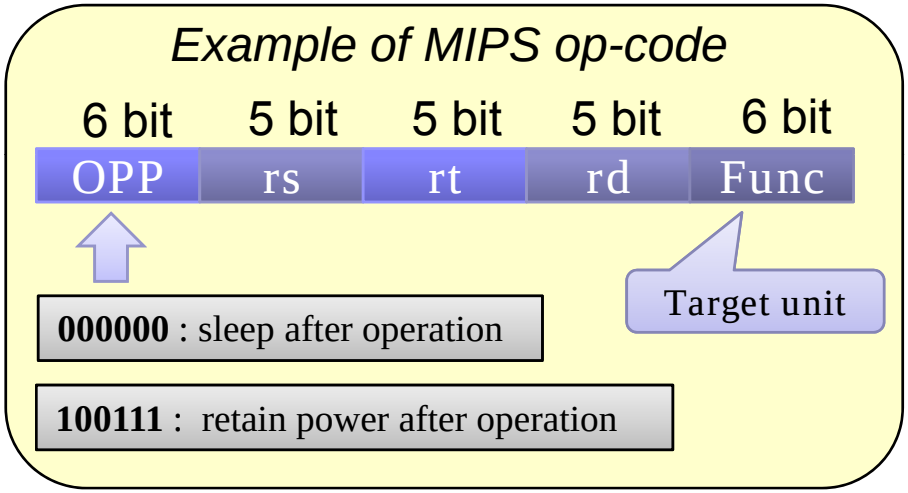

- Operating System based technique
  - A status register for each Ex-unit to enable/disable PG
  - OS controls the status registers depending on the chip temperature

Implementation

- Real implementation
  - MIPS R3000 based CPU
  - 90nm CMOS technology
- 41% area overhead compared with Non-PG one
  - Power Switches
  - Sleep controller

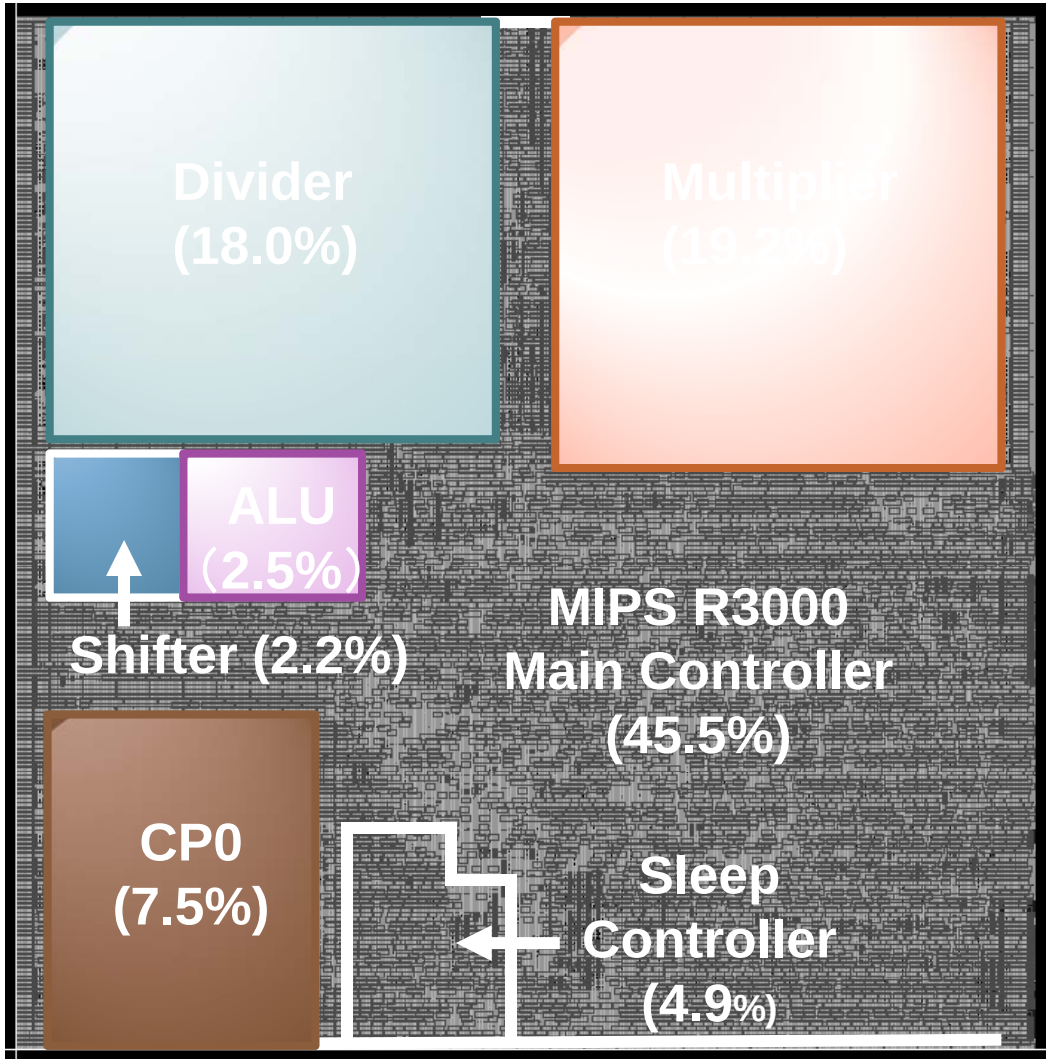

The Utilization of Each Unit

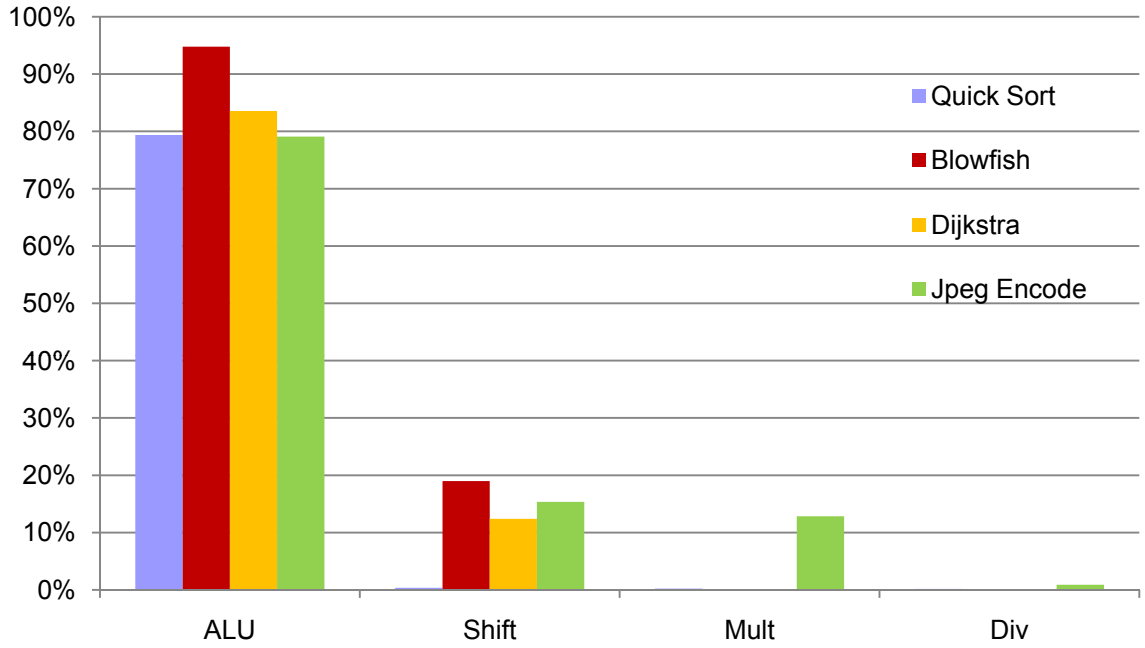

- ALU is frequently used
  - Should be always activated except cache misses happen
- Average sleep period
  - ALU/shift : about 20-40 cycles (mainly caused by cache misses)
  - Mult/Div: about 10000 cycles

Average Leakage Power

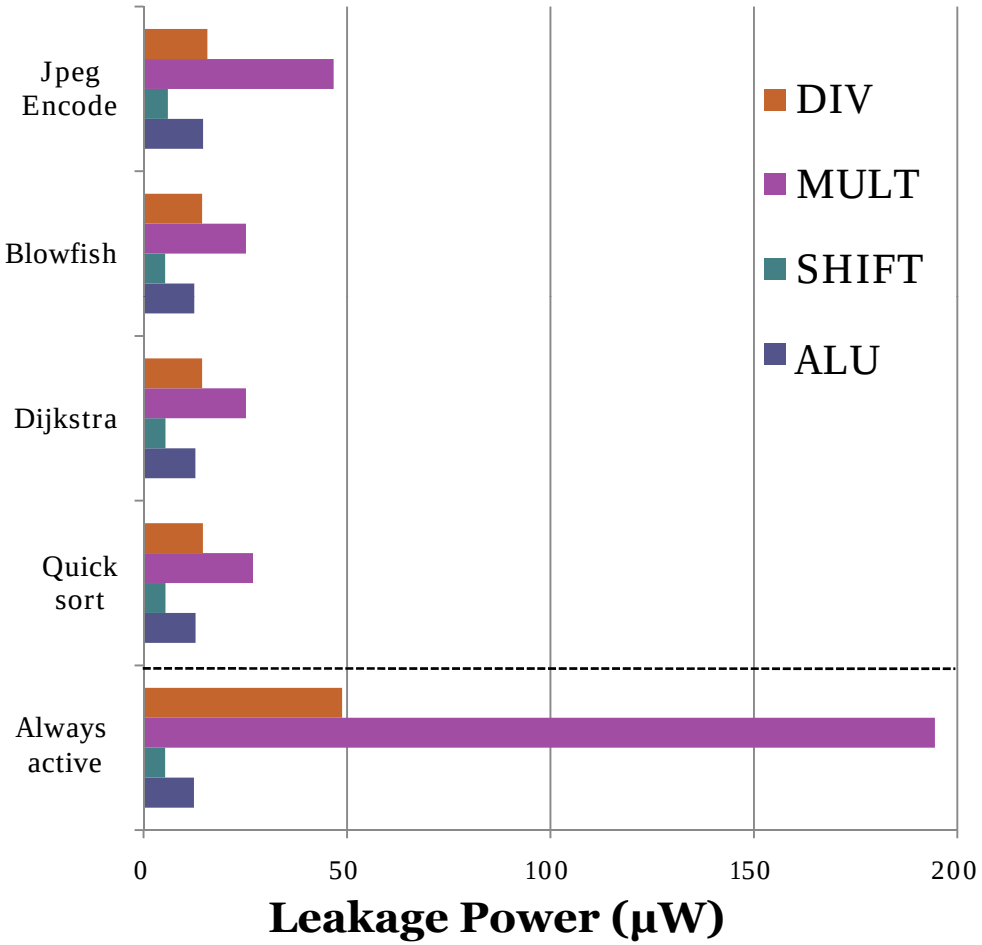

- Leakage Power for each Unit
  - 200MHz, 1V, 25 C°
  - Evaluated by SPICE Simulation
- Always active: no power-gating (normal-processor)
- Huge reduction in MULT/DIV units
  - MULT/DIV: seldom used
- Small or no reduction in ALU/SHIFT
  - ALU/SHIFT: frequently used
- Reduction of core's total leakage power (including other parts)
  - 47% reduction (439µW → 232 µW)

Summary

- Run-time power-gating for leakage-power reduction
- Implemented in a MIPS R3000 like processor
  - Several features for power-gating management
  - 47% leakage power reduction
- Useful for HPC processors
  - Many execution (floating-point) units
- Future work
  - Evaluation with real chip
  - Application to more sophisticated processors

Acknowledgements

- Prof. H. Nakamura
- Prof. M. Namiki
- Prof. H. Amano
- Prof. K. Usami
- Dr. Y. Hasegawa
- Mr. N. Seki
- Mr. L. Zhao
- Ms. J. Kei
- Mr. D. Ikebuchi
- Mr. Y. Kojima
- Mr. T. Kashima
- Mr. S. Takeda
- Mr. M. Nakata,
- Mr. T. Sunata
- Mr. J. Kanai
- Mr. T. Shirai.

Thank you!

Questions / Comments?

Backup Slides

Course-grain vs. Fine-grain

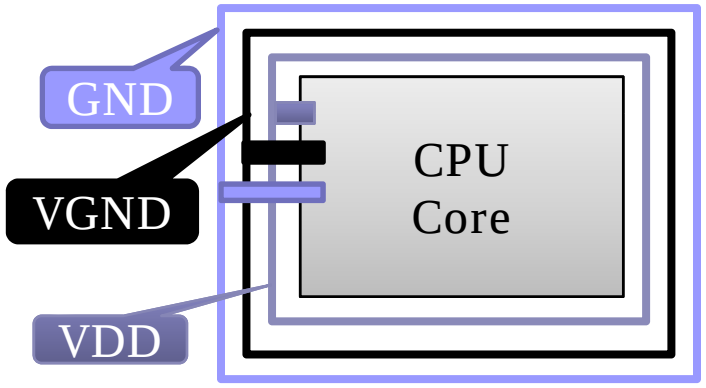

PG used by Power Ring

- Course-grain Power Gating
  - Traditional PG is used Power Ring
    - Added Virtual Ground(VGND) ring and Sleep tr in Core
  - Unit of PG is a Core or Module
  - Appling for a large semiconductor domain and controlled for long time scale (millisecond order)

- Fine-grain Power Gating
  - A few Cell are sharing the VGND (Locally shared)
  - The standard Cells are fixed for PG
    - Include the VGND rail
    - Add New port for the VGND
  - Power Switch Cell, Isolation Cell, Level shifter is added

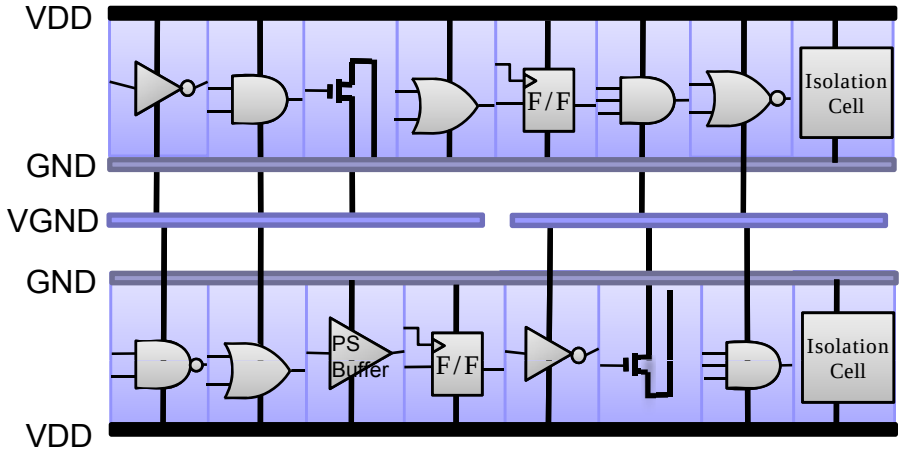

Over review of Geyser-0

- Based MIPS R3000 Core
- For comparing, PG-Core and Non-PG Core(NM-CORE) is existed
  - Only one Core is working
  - PG Cores , NM-Core and other area (Bridge) is divided by the power domain
  - We can measure the current of each power domain
- TLB, Cache, MMU is shared
  - 16 entry TLB
  - 8 KB Instruction Cache and Data Cache
  - MMU: Memory Management Unit
- Mother Board is designed for evaluation
  - Vertex4 FPGA installed
  - 64 MB SDRAM
  - JTAG, USB, Video Output

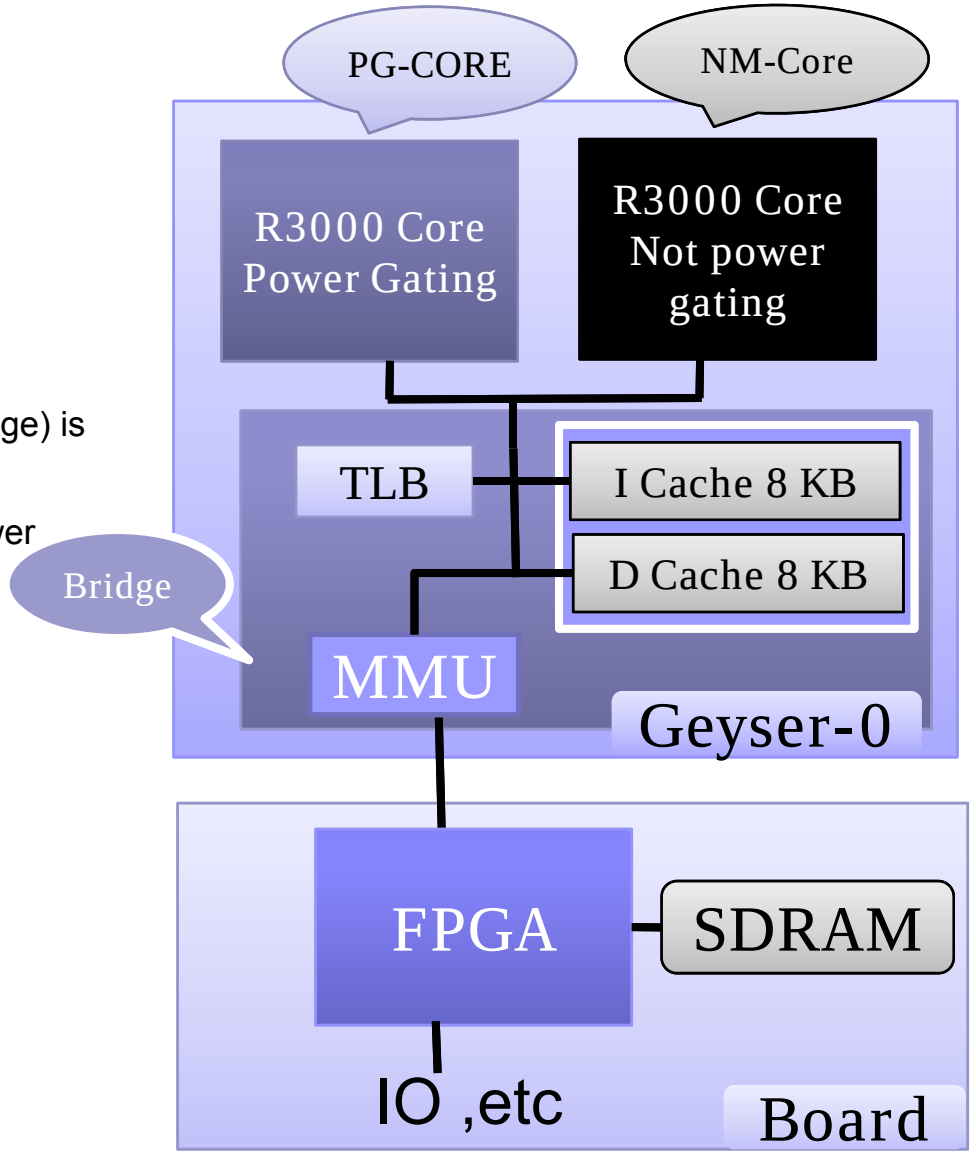

Geyser-0 Design Flow

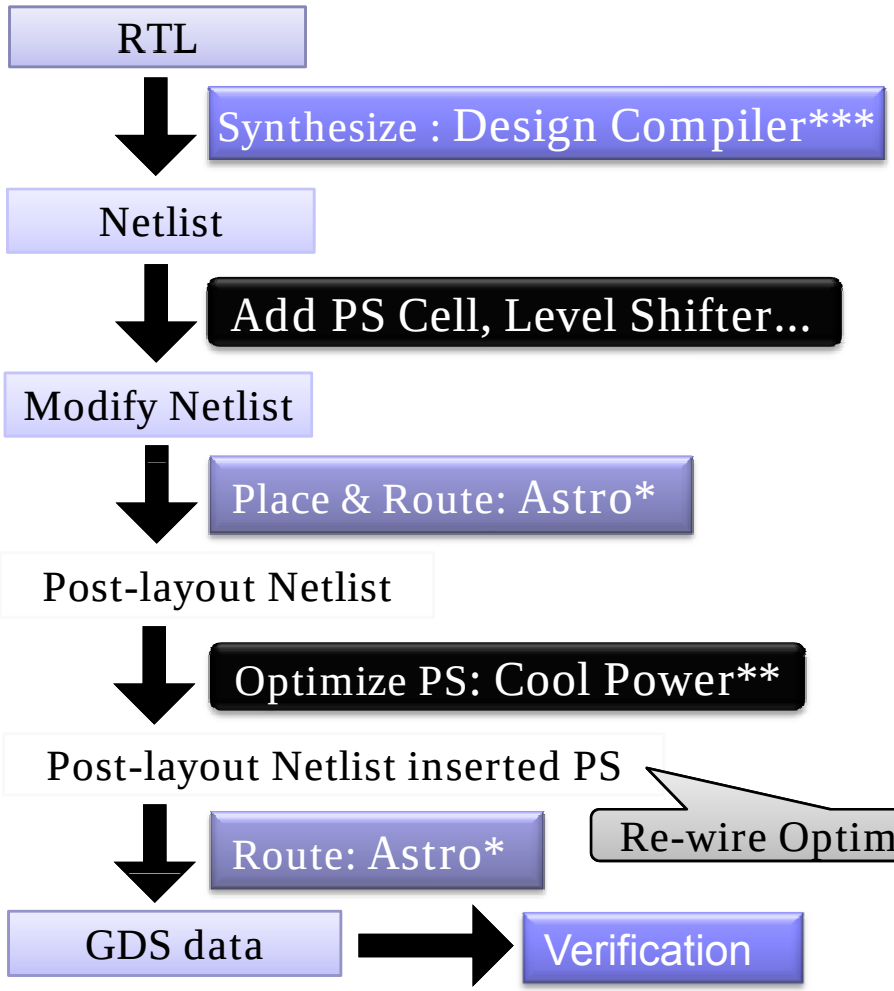

- Geyser-0 Specification
  - ASPLA 90nm technology
  - Chip size: 2.5 mm * 5 mm
  - Core VDD: 1 V
  - Operation Frequency: 200 MHz
- The Common Design flow added new phases
  - Insert Power Switch Cell, Level Shifter and Isolation Cell
  - Optimize these Cells and Re-Route

*Astro 2007.03-SP3 (Synopsys, Inc.)

**Cool Power 2007.3.8.5 (Sequence Design, Inc.)

***Design Compiler 2006.06-SP2 (Synopsys, Inc.)

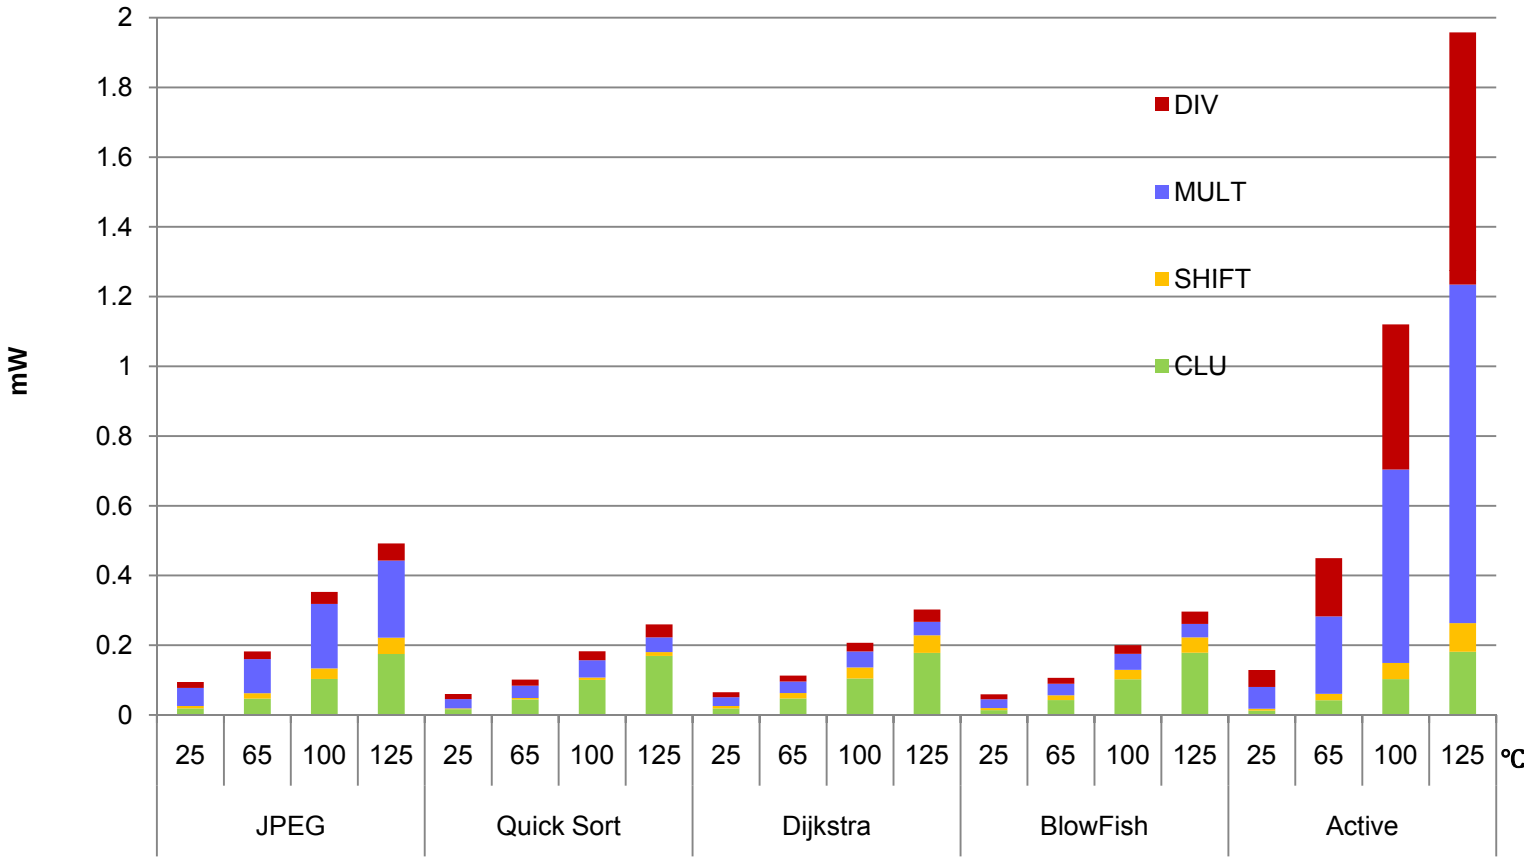
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
mW
DIV
MULT
SHIFT
CLU
25
65
100
125
℃
JPEG
Quick Sort
Dijkstra
BlowFish
Active

Challenges of Run-Time Power-Gating

- Wake-up latency when powering a execution-unit
  - Pre-wakeup before the unit is actually needed to hide the latency
  - In our design, we need to wakeup the unit one cycle before
    - Our target is 200MHz CPU
    - According to our estimation, wake-up latency is about 4 ns at maximum
- Energy Overhead for sleep-mode control
  - Dynamic energy caused by sleep transistors control
  - Important to consider for fine-grain power-gating

Background

- Power efficiency
  - A first class design constraint in today's computer systems
  - Key to achieve good performance/cost ratio even for supercomputers
    - Electrical power costs $$$
    - 50MW
    - Small footprint
    - Reliability / availability
  - The green500 project
    - Power efficiency of top500 supercomputers
    - Maintained by Prof. Chou
    - Importance of power efficiency

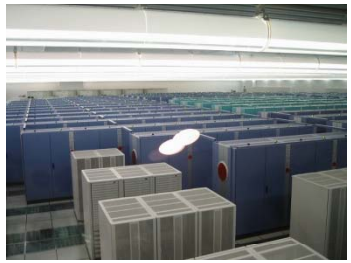

Leakage Power

- Leakage current
  - Flows even when circuits are idle
  - Expected to increase in the future
- Leakage reduction techniques
  - Standby time:
    - power-gating, DTCMOS, VTCMOS, …
  - Runtime:
    - Cache-decay, Drowsy-cache, …
- Leakage for logic parts (ALU, decode etc.) is still a matter

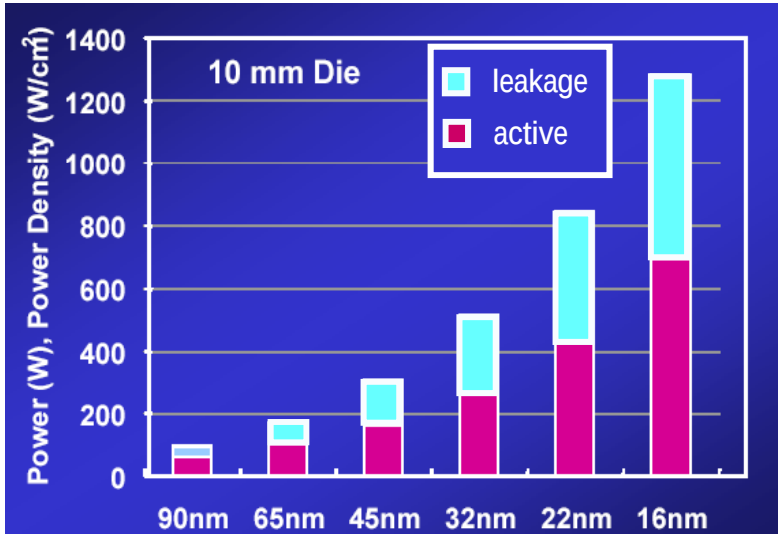

source: Shekhar Borkar, *Intel*

We are trying to reduce runtime leakage power of logic parts