

OSCAR Compiler and API for High Performance Low Power Multicores and Their Application to Smartphones, Automobiles, Medical Systems

Hironori Kasahara

**Professor, Dept. of Computer Science & Engineering
Director, Advanced Multicore Processor Research Institute**

Waseda University, Tokyo, Japan

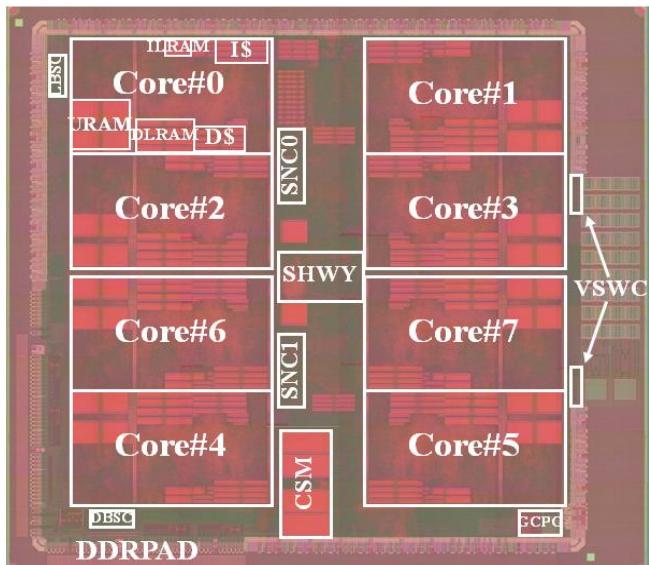
**IEEE Computer Society Board of Governors
IEEE Computer Society Multicore Strategic Technical
Committee (STC) Chair**

URL: <http://www.kasahara.cs.waseda.ac.jp/>

Intel/Kai, 2012.10.18(Thursday)

Multi/Many-core Everywhere

Multi-core from embedded to supercomputers



OSCAR Type Multi-core Chip by Renesas in
METI/NEDO Multicore for Real-time Consumer
Electronics Project (Leader: Prof.Kasahara)



The 37th (June 20,2011) &38th
(Nov.14.2011) Top 500 No.1,
Riken Fujitsu "K" 705,024 cores
Peak 11.28 PFLOPS, (88,128procs)
LINPACK 10.510 PFLOPS (93.2%)

➤ Consumer Electronics (Embedded)

Mobile Phone, Game, TV, Car Navigation, Camera,
IBM/ Sony/ Toshiba Cell, Fujitsu FR1000,
Panasonic Uniphier, NEC/ARM MPCore/MP211/NaviEngine,
Renesas 4 core RP1, 8 core RP2, 15core Hetero RP-X,
Plurality HAL 64(Marvell), Tilera Tile64/ -Gx100(>1000cores),
DARPA UHPC (2017: 80GFLOPS/W)

➤ PCs, Servers

Intel Quad Xeon, Core 2 Quad, Montvale, Nehalem(8cores),
Larrabee(32cores), SCC(48cores), Night Corner(50 core+:22nm),
AMD Quad Core Opteron (8, 12 cores)

➤ WSs, Deskside & Highend Servers

IBM(Power4,5,6,7), Sun (SparcT1,T2), Fujitsu SPARC64fx8

➤ Supercomputers

Earth Simulator: 40TFLOPS, 2002, 5120 vector proc.
BG/Q (A2:16cores) Water Cooled 20PFLOPS, 3-4MW (2011-12),
BlueWaters(HPCS) Power7, 10 PFLOP+(2011.07),
Tianhe-1A (4.7PFLOPS, 6coreX5670+ Nvidia Tesla M2050),
Godson-3B (1GHz40W 8core128GFLOPS) -T (64 core, 192GFLOPS: 2011)
RIKEN Fujitsu "K" 10PFLOPS(8core SPARC64VIIIIfx, 128GGFLOPS)

High quality application software, Productivity, Cost performance, Low power consumption are important
Ex, Mobile phones, Games

Compiler cooperated multi-core processors are promising to realize the above futures

Green Computing Systems R&D Center

Waseda University

Supported by METI (Mar. 2011 Completion)

<R & D Target>

**Hardware, Software, Application
for Super Low-Power Manycore
Processors**

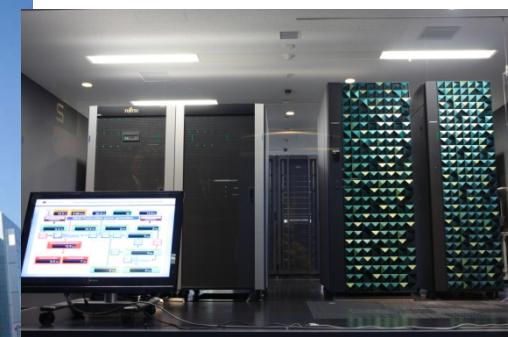
- More than 64 cores
- Natural air cooling (No fan)
Cool, Compact, Clear, Quiet
- Operational by Solar Panel

<Industry, Government, Academia>

Hitachi, Fujitsu, NEC, Renesas, Olympus,
Toyota, Denso, Mitsubishi, Toshiba, etc

<Ripple Effect>

- Low CO₂ (Carbon Dioxide) Emissions
- Creation Value Added Products
 - Consumer Electronics, Automobiles,
Servers



Hitachi SR16000:
Power7 128coreSMP
Fujitsu M9000
SPARC VII 256 core SMP



Beside Subway Waseda Station,
Near Waseda Univ. Main Campus

Research, development and practical utilization through industry-government-academia partnerships (spillover effect)

Environment



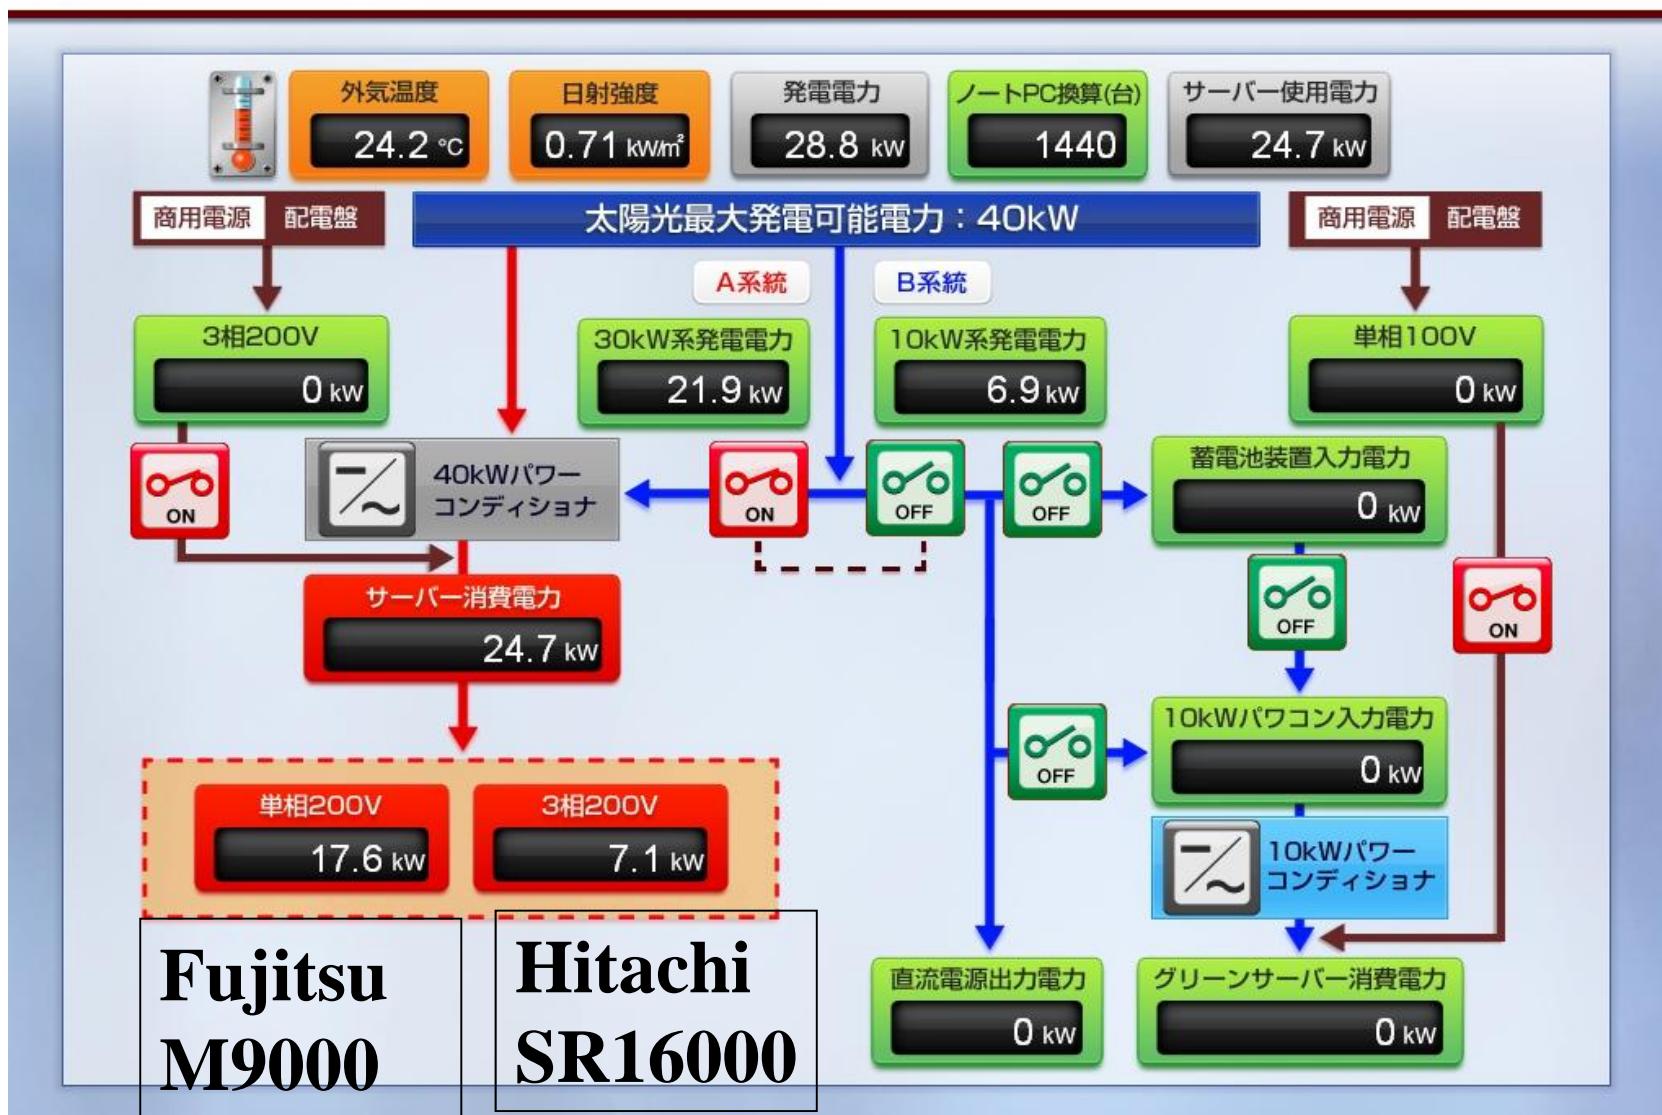
Green Computing Systems R&D Center, 2011.11.1 (Clear)

Solar Power Generation & Server Consumption



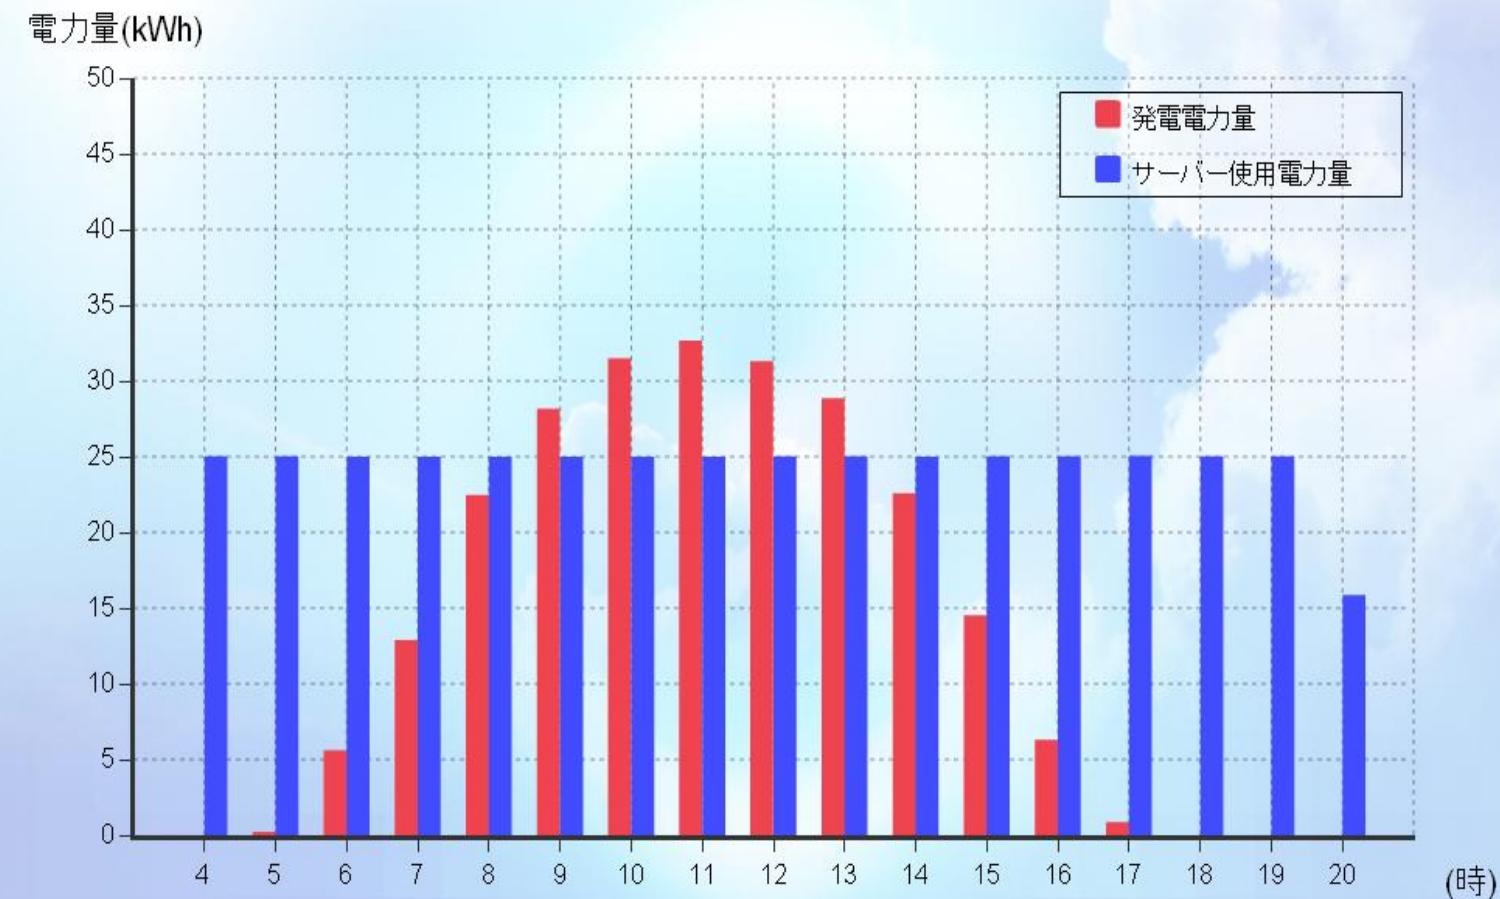
WASEDA University

早稲田大学グリーンコンピューティングシステム研究開発センター太陽光発電システム



2012.4.2(Clear) Power Generation and Server Consumption: One day Trends

電力量の1日の変化



Super Low Power Web Server Using Embedded Multicore Processor RPX

1W with 8 SH4A processor cores



WASEDA UNIVERSITY Computer Science and Engineering

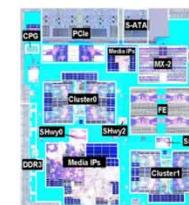
Kasahara Laboratory

Power consumption



0.96 W

by RPX embedded multicore Web-server.



Japanese | English
Private Page



Contents

- ▶ Professor Kasahara
- ▶ Associate Professor
- ▶ Publications
- ▶ Members

News

- 2012.4.25 OSCAR API 2.0 has been released
- 2012.4.2 Low power embedded multicore RPX server started Kasahara & Kimura Laboratory's web service and power consumption indication.
- 2011.10.07 Prof. Hironori Kasahara has been elected to the IEEE Computer Society Board of Governors(2012-2014). Thank you very much for your kind supports.
- 2011.09.06 Information for the 25th Anniversary Workshop LCPC2012 (International Workshop on Languages and Compilers for Parallel Computing) Sep. 11-13, 2012 was upped.

METI/NEDO National Project

Multi-core for Real-time Consumer Electronics

<Goal> R&D of compiler cooperative multi-core processor technology for consumer electronics like Mobile phones, Games, DVD, Digital TV, Car navigation systems.

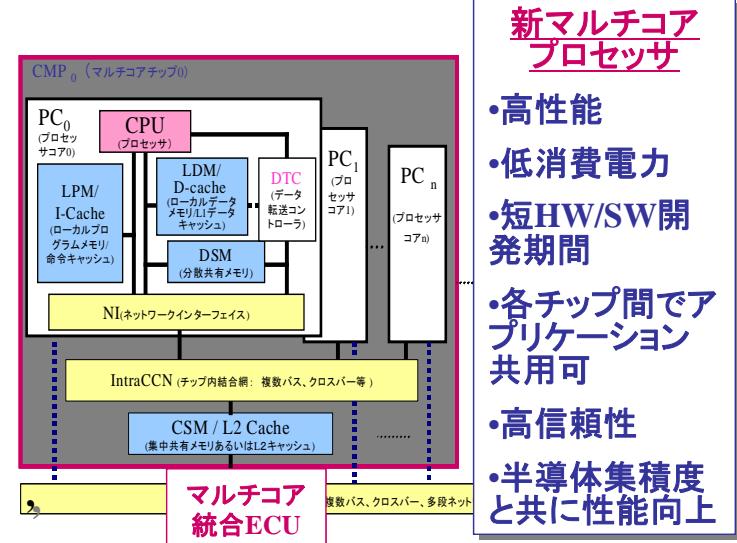
<Period> From July 2005 to March 2008

<Features>

- Good cost performance
- Short hardware and software development periods
- Low power consumption
- Scalable performance improvement with the advancement of semiconductor
- Use of the same parallelizing compiler for multi-cores from different vendors using newly developed API

API: Application Programming Interface

(2005.7~2008.3) **



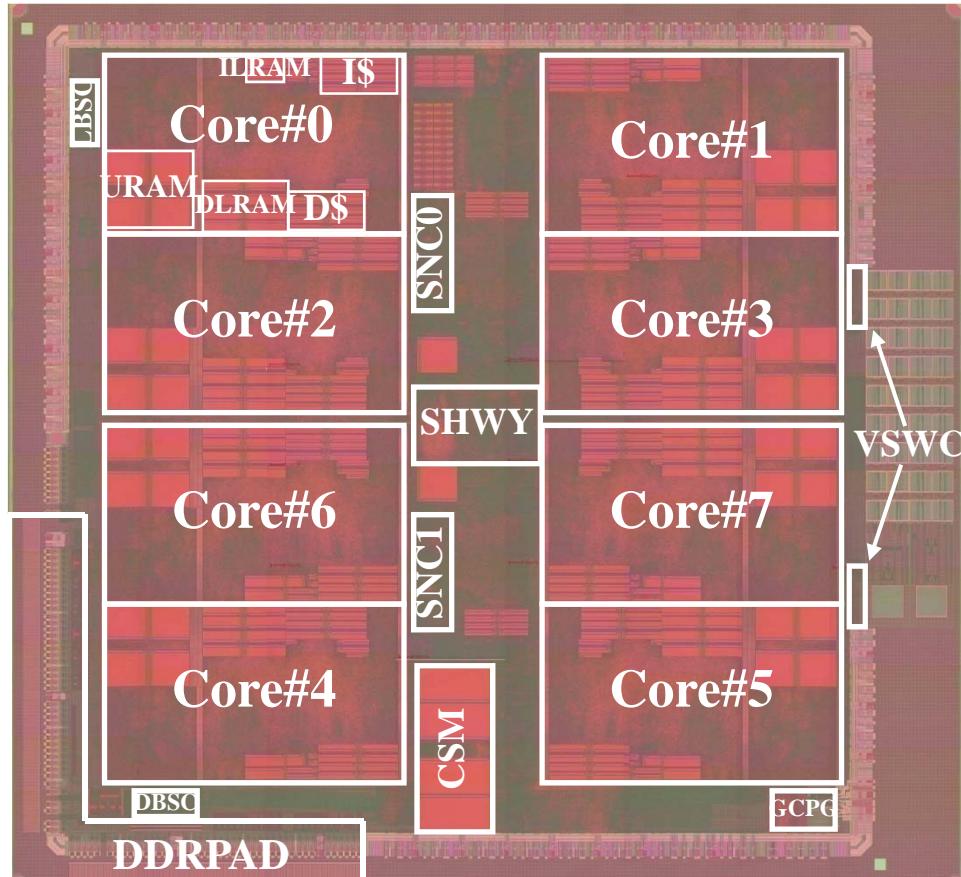
開発マルチコアチップは情報家電へ



**Hitachi, Renesas, Fujitsu, Toshiba, Panasonic, NEC

Renesas-Hitachi-Waseda Low Power 8 core RP2

Developed in 2007 in METI/NEDO project



Process Technology	90nm, 8-layer, triple-Vth, CMOS
Chip Size	104.8mm ² (10.61mm x 9.88mm)
CPU Core Size	6.6mm ² (3.36mm x 1.96mm)
Supply Voltage	1.0V–1.4V (internal), 1.8/3.3V (I/O)
Power Domains	17 (8 CPUs, 8 URAMs, common)

**IEEE ISSCC08: Paper No. 4.5, M.ITO, ... and H. Kasahara,
“An 8640 MIPS SoC with Independent Power-off Control of 8
CPUs and 8 RAMs by an Automatic Parallelizing Compiler”**

Demo of NEDO Multicore for Real Time Consumer Electronics at the Council of Science and Engineering Policy on April 10, 2008

第74回総合科学技術会議【平成20年4月10日】



第74回総合科学技術会議の様子(1)



第74回総合科学技術会議の様子(2)



第74回総合科学技術会議の様子(3)



第74回総合科学技術会議の様子(4)

CSTP Members

Prime Minister:
Mr. Y. FUKUDA

**Minister of State for
Science, Technology
and Innovation
Policy:**

Mr. F. KISHIDA

**Chief Cabinet
Secretary:**

Mr. N. MACHIMURA

**Minister of Internal
Affairs and
Communications :**

Mr. H. MASUDA

Minister of Finance :

Mr. F. NUKAGA

**Minister of
Education, Culture,
Sports, Science and
Technology:**

Mr. K. TOKAI

**Minister of
Economy, Trade and
Industry:**

Mr. A. AMARI

OSCAR Parallelizing Compiler

To improve effective performance, cost-performance and software productivity and reduce power

Multigrain Parallelization

coarse-grain parallelism among loops and subroutines, near fine grain parallelism among statements in addition to loop parallelism

Data Localization

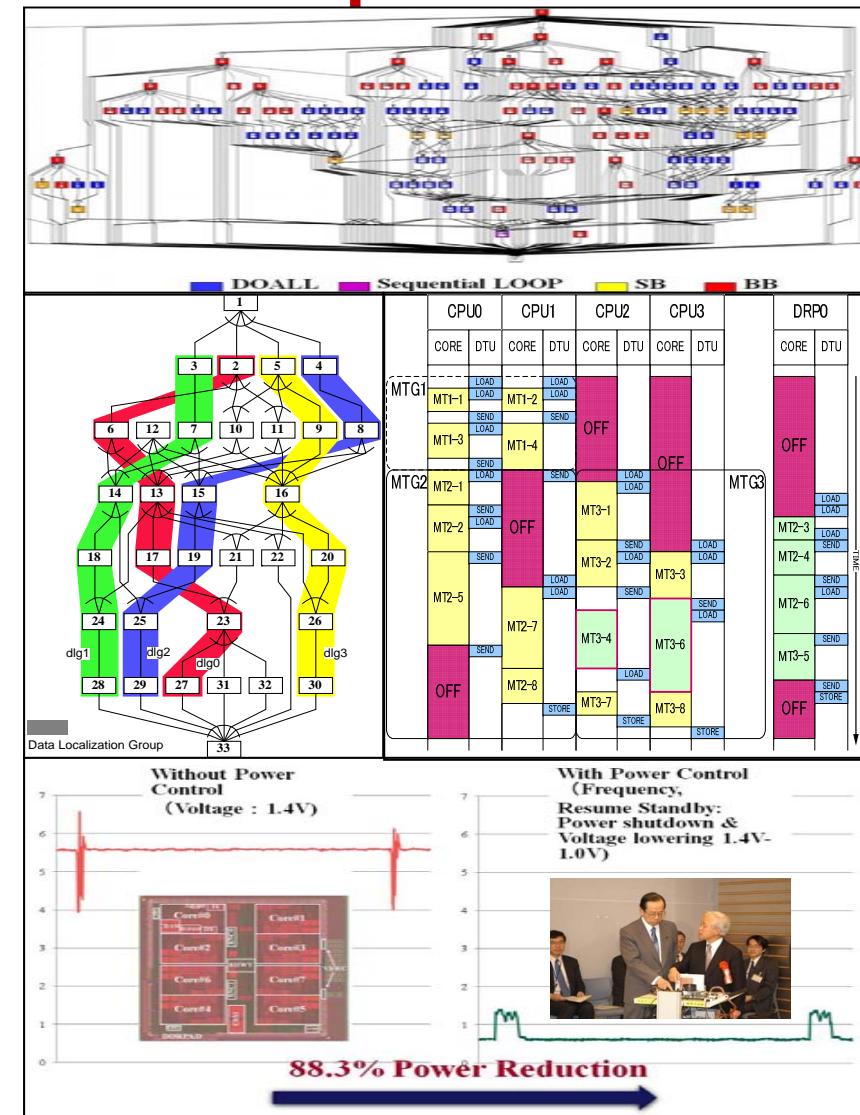
Automatic data management for distributed shared memory, cache and local memory

Data Transfer Overlapping

Data transfer overlapping using Data Transfer Controllers (DMAs)

Power Reduction

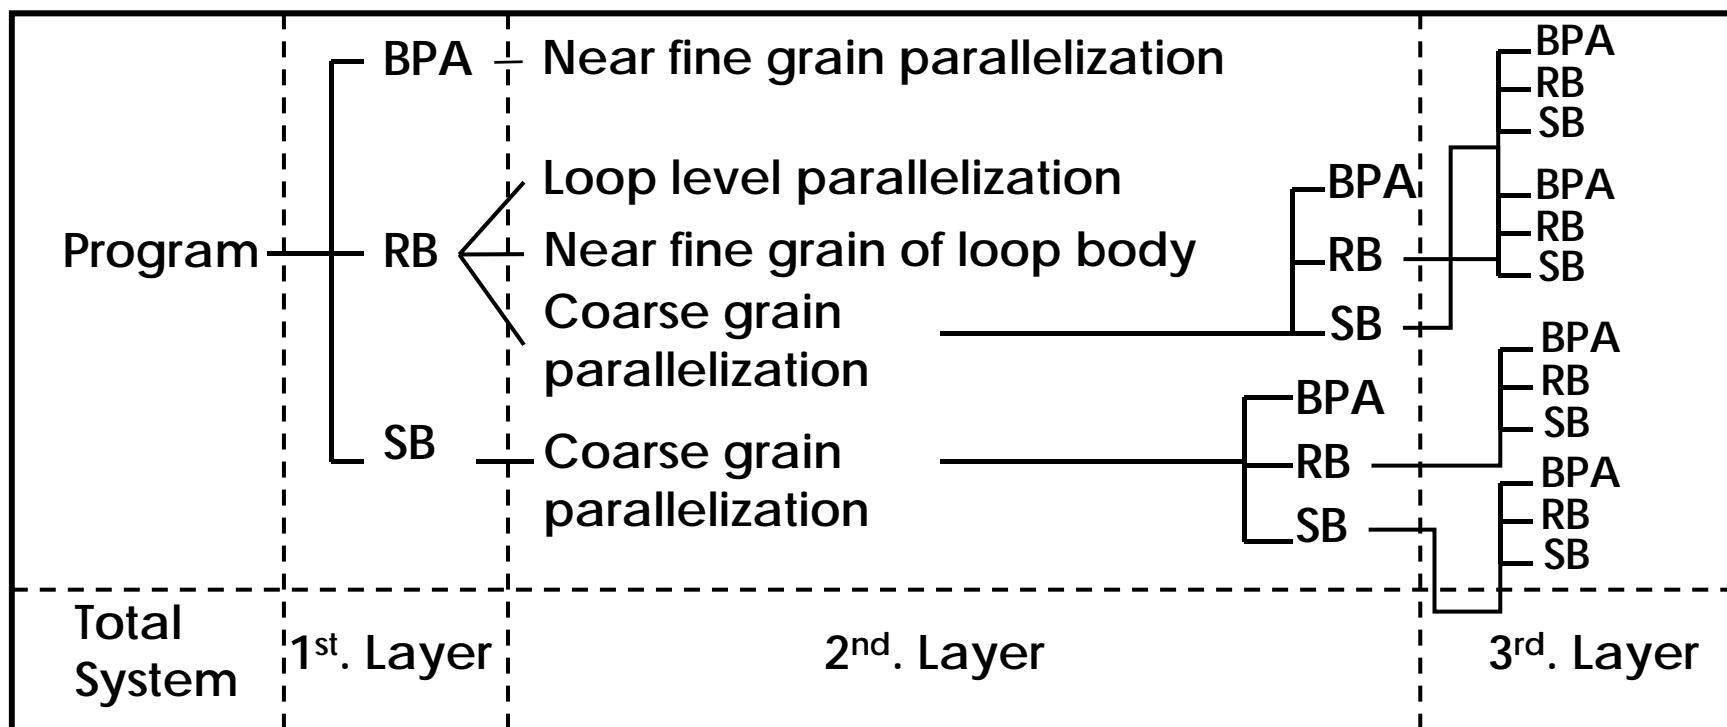
Reduction of consumed power by compiler control DVFS and Power gating with hardware supports.



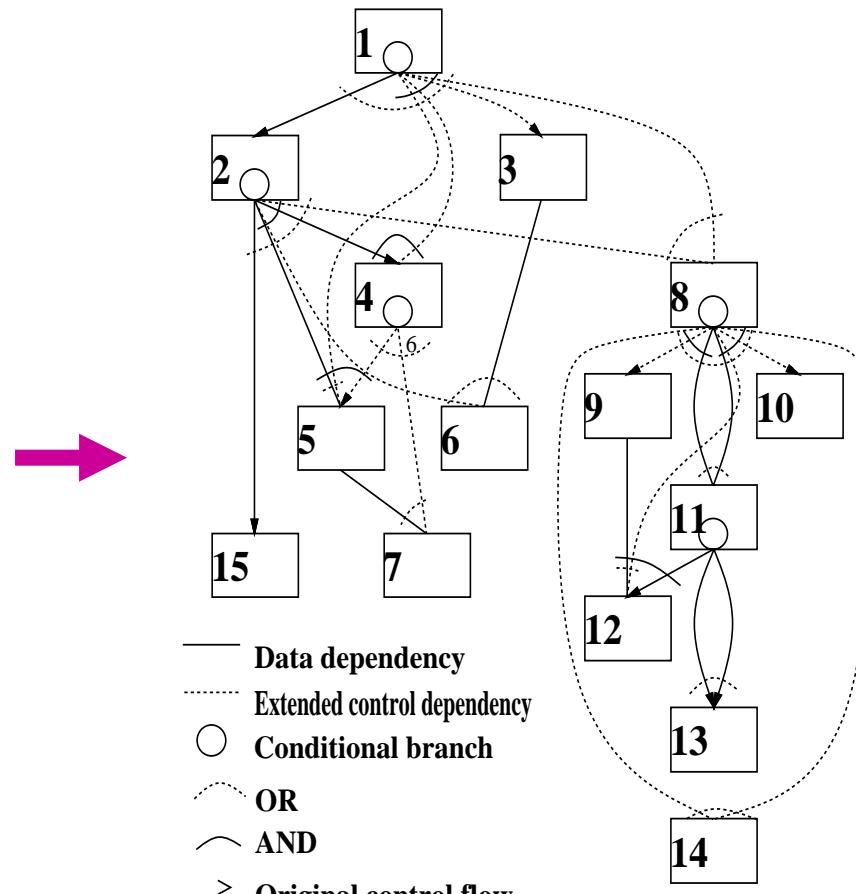
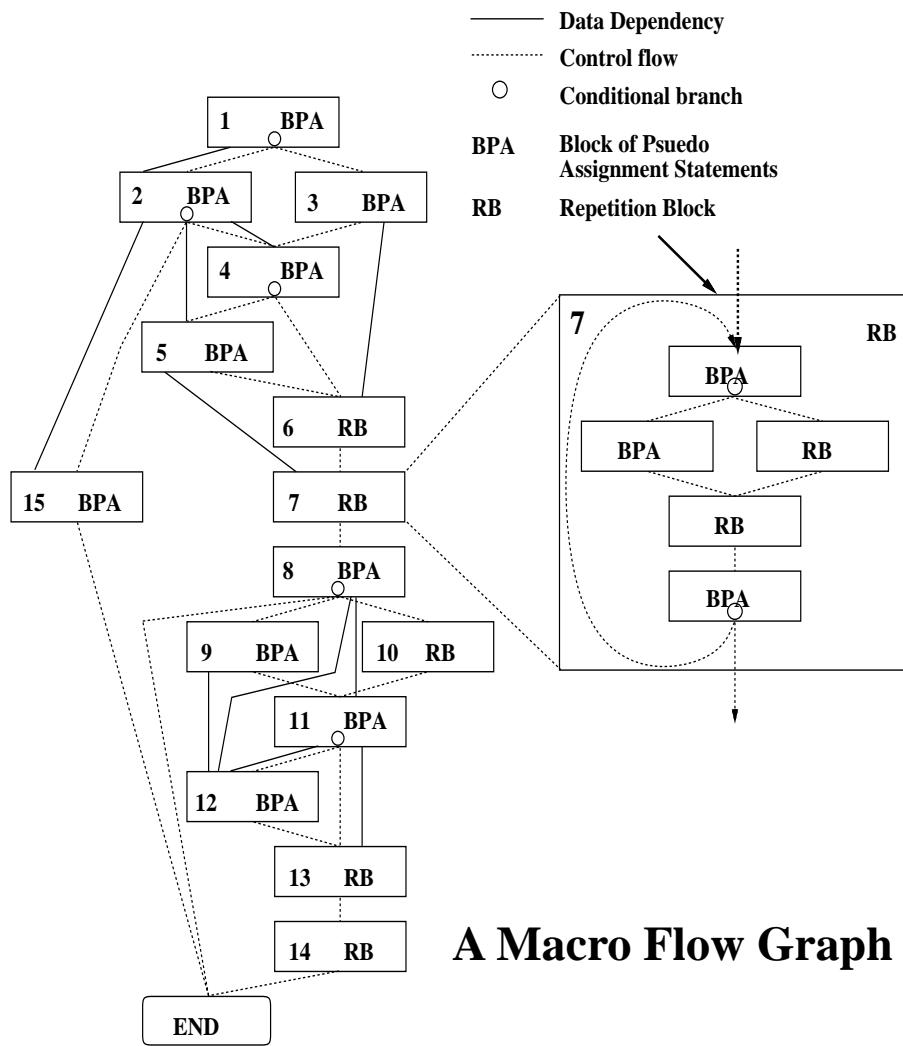
Generation of coarse grain tasks

■ Macro-tasks (MTs)

- **Block of Pseudo Assignments (BPA): Basic Block (BB)**
- **Repetition Block (RB) : natural loop**
- **Subroutine Block (SB): subroutine**

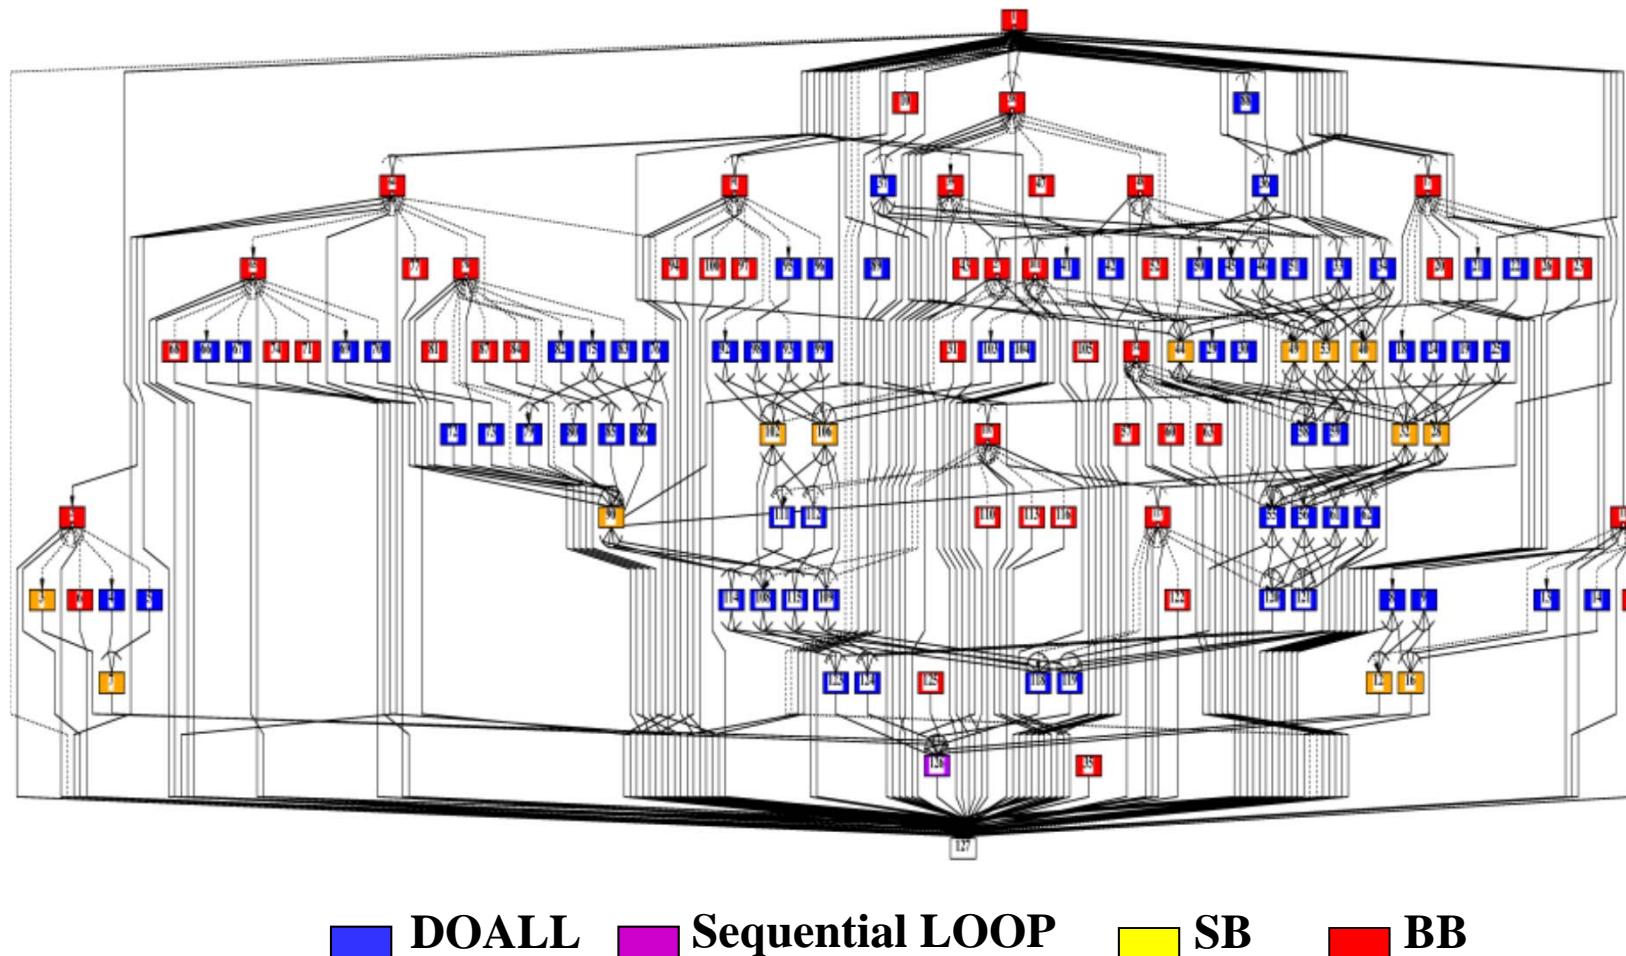


Earliest Executable Condition Analysis for coarse grain tasks (Macro-tasks)

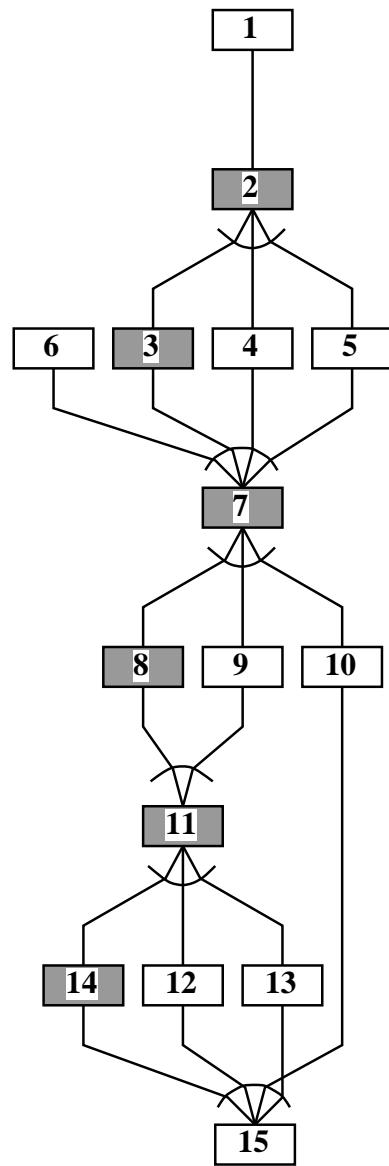


MTG of Su2cor-LOOPS-DO400

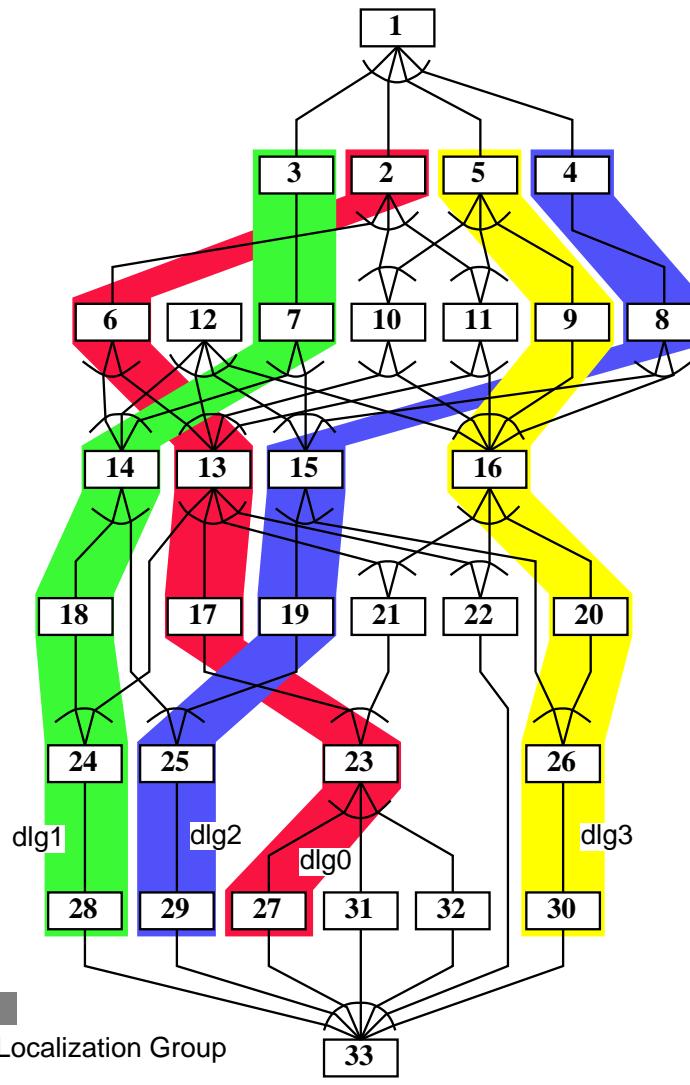
- Coarse grain parallelism PARA_ALD = 4.3



Data Localization



MTG



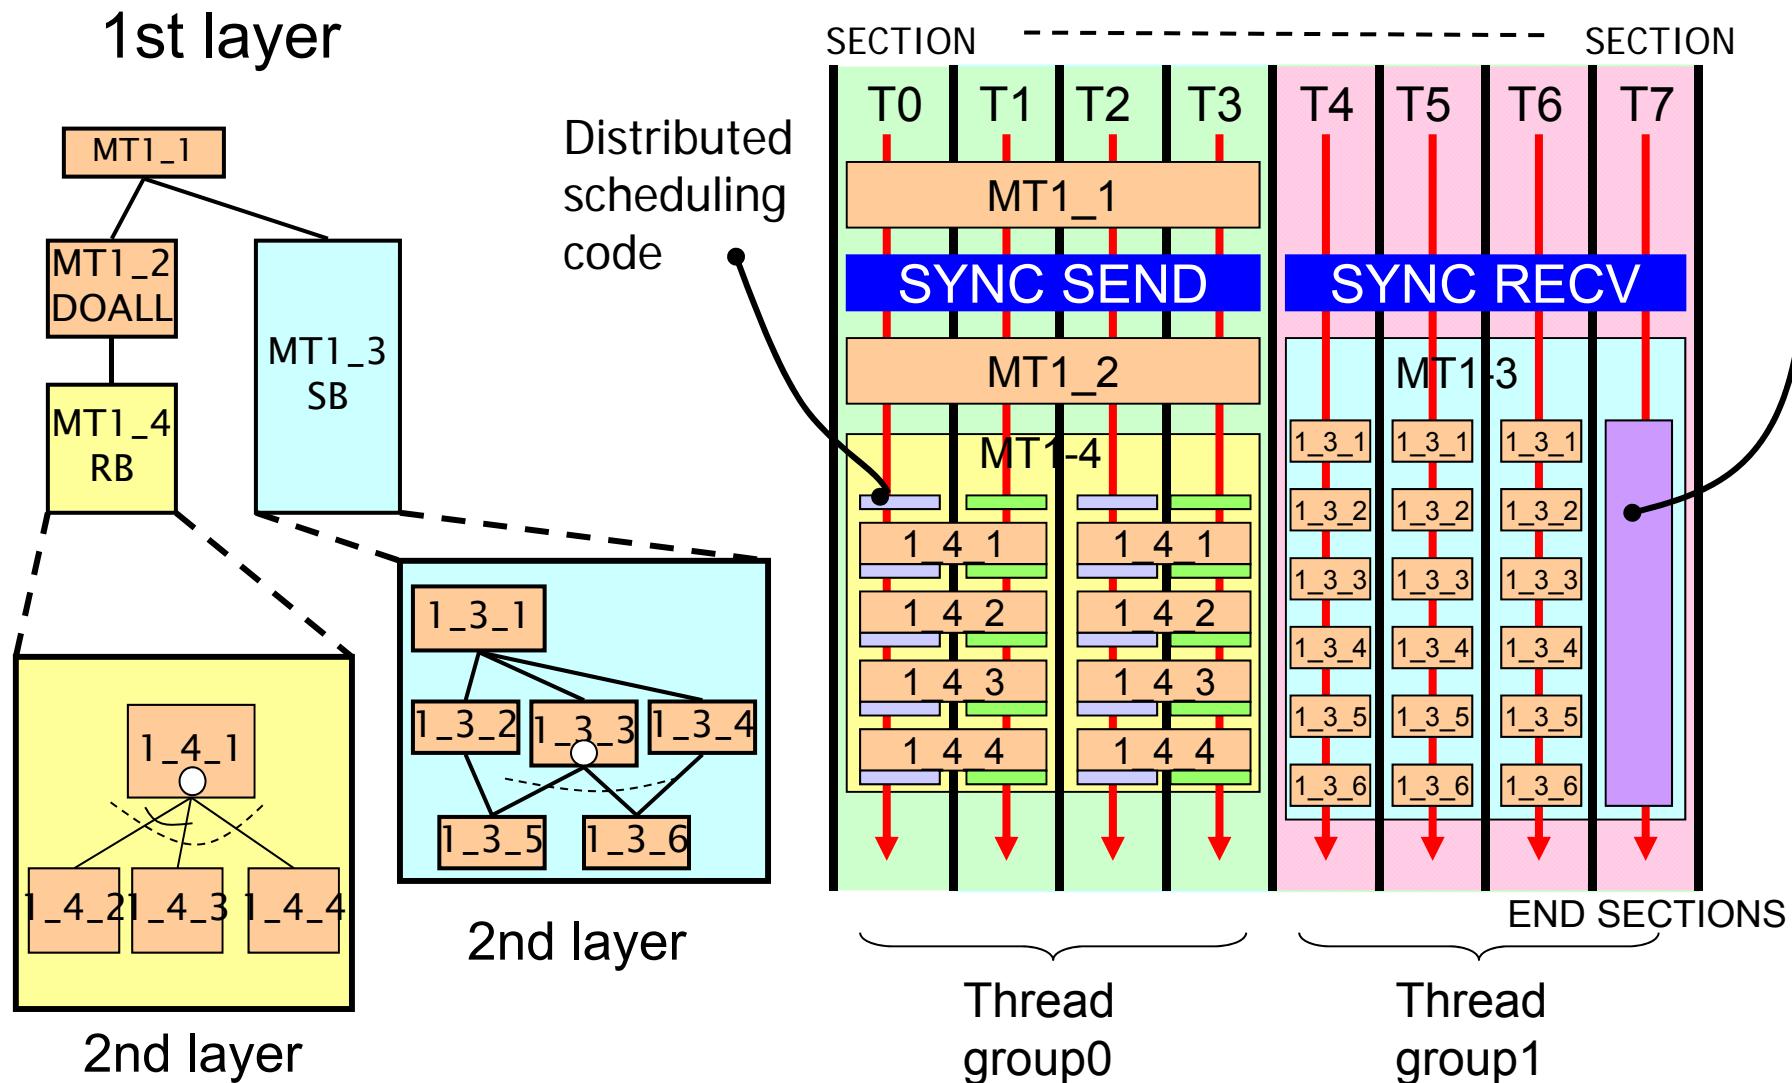
MTG after Division

A schedule for two processors

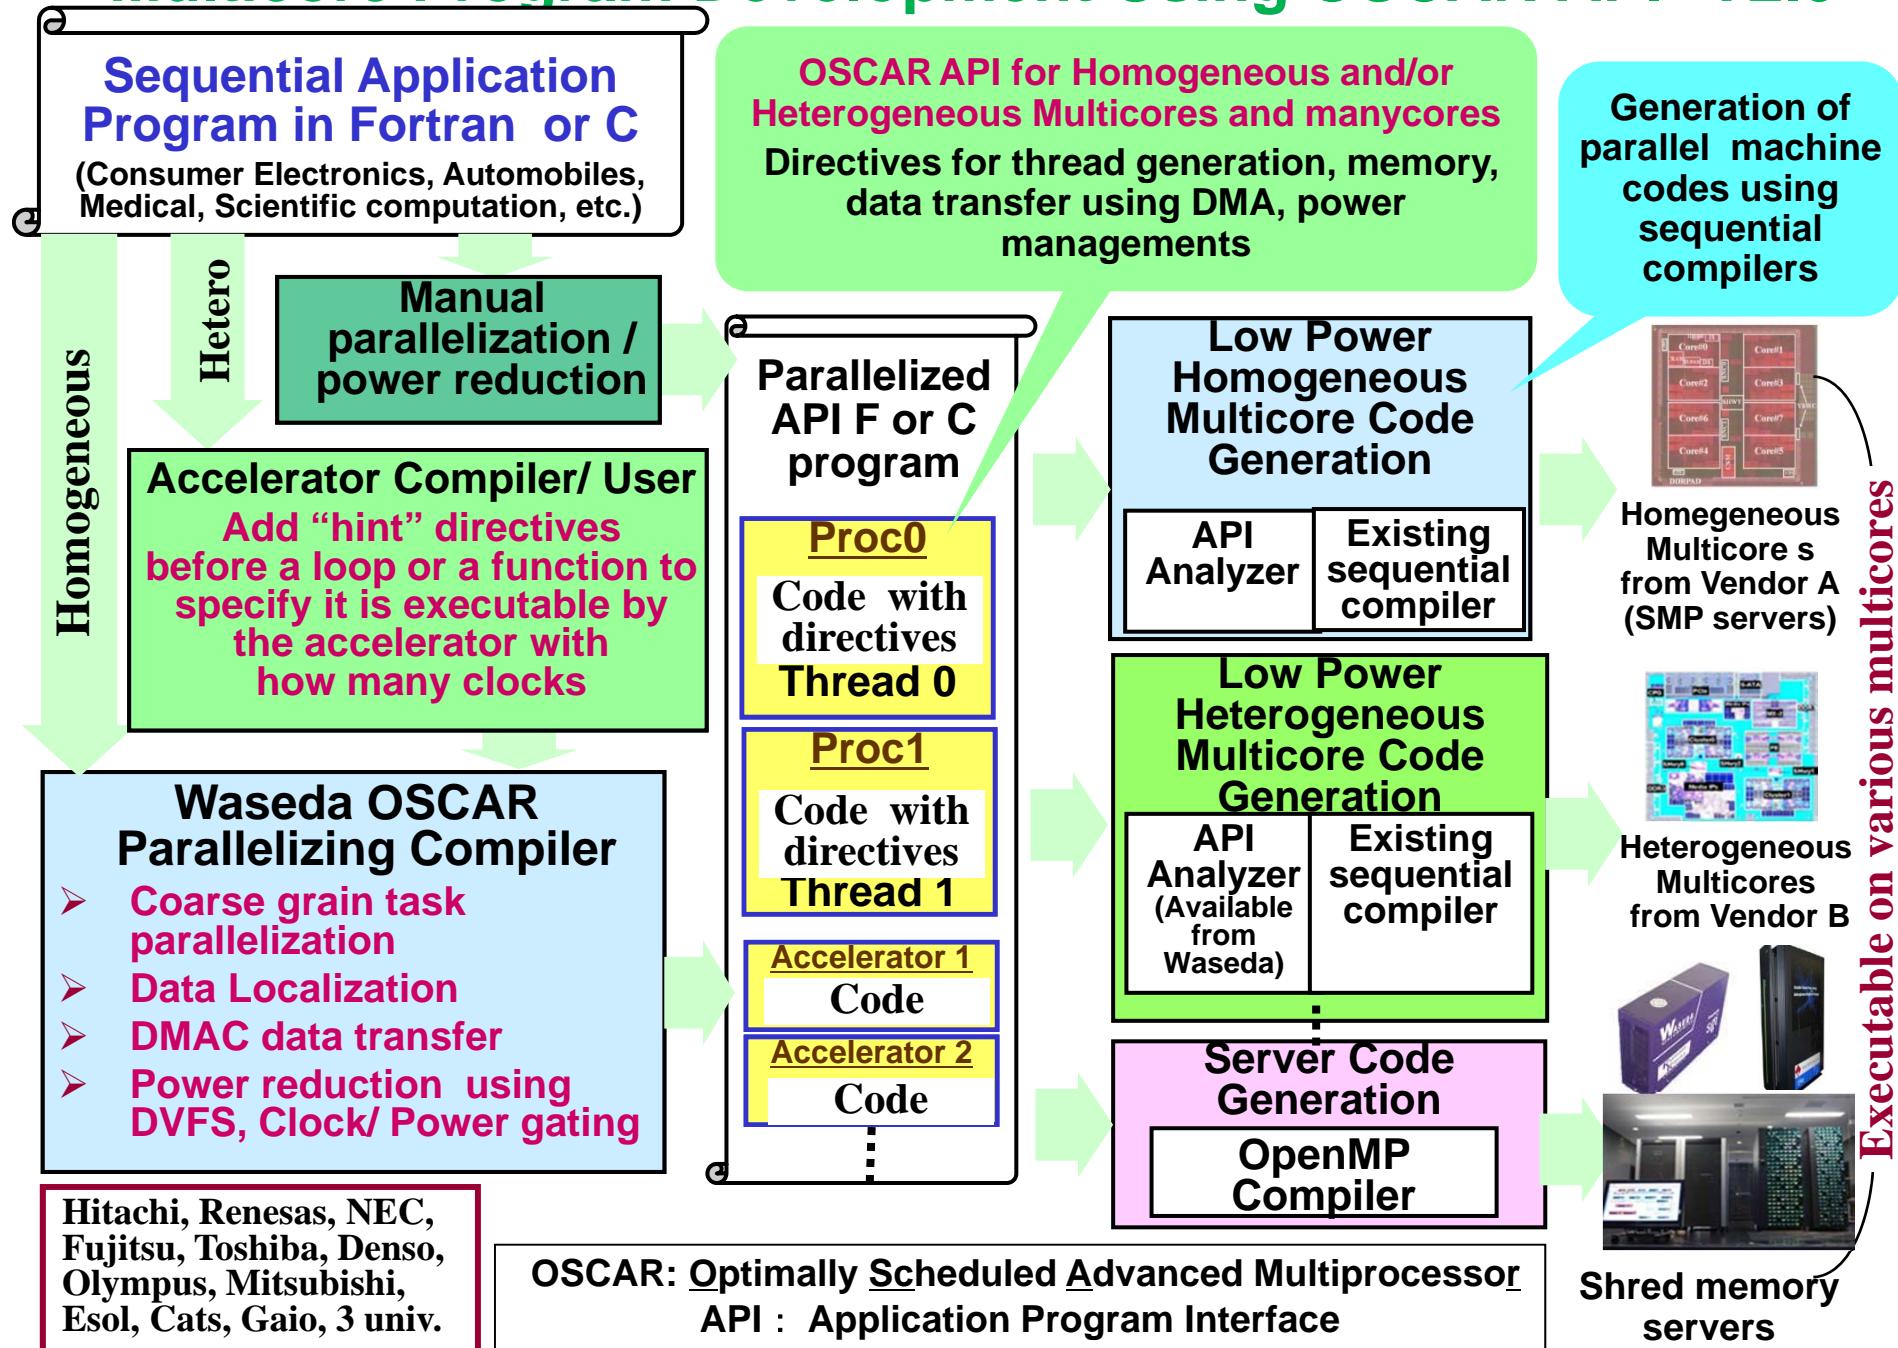
	PE0	PE1
12	1	
2	3	
6	7	
4	14	
8	18	
15	5	
19	9	
25	11	
29	10	
13	16	
17	20	
22	26	
21	30	
23	24	
27	28	
	32	
	31	

Generated Multigrain Parallelized Code

(The nested coarse grain task parallelization is realized by only OpenMP “section”, “Flush” and “Critical” directives.)



Multicore Program Development Using OSCAR API V2.0



OSCAR API Ver. 2.0 for Homogeneous/Heterogeneous Multicores and Manycores

List of Directives (22 directives)

- ▶ Parallel Execution API
 - ▶ parallel sections (*)
 - ▶ flush (*)
 - ▶ critical (*)
 - ▶ execution
- ▶ Memory Mapping API
 - ▶ threadprivate (*)
 - ▶ distributedshared
 - ▶ onchipshared
- ▶ Synchronization API
 - ▶ groupbarrier
- ▶ Data Transfer API
 - ▶ dma_transfer
 - ▶ dma_contiguous_parameter
 - ▶ dma_stride_parameter
 - ▶ dma_flag_check
 - ▶ dma_flag_send

(* from OpenMP)

- ▶ Power Control API
 - ▶ fvcontrol
 - ▶ get_fvstatus
- ▶ Timer API
 - ▶ get_current_time
- ▶ Accelerator
 - ▶ accelerator_task_entry
- ▶ Cache Control
 - ▶ cache_writeback
 - ▶ cache_selfinvalidate
 - ▶ complete_memop
 - ▶ noncacheable
 - ▶ aligncache

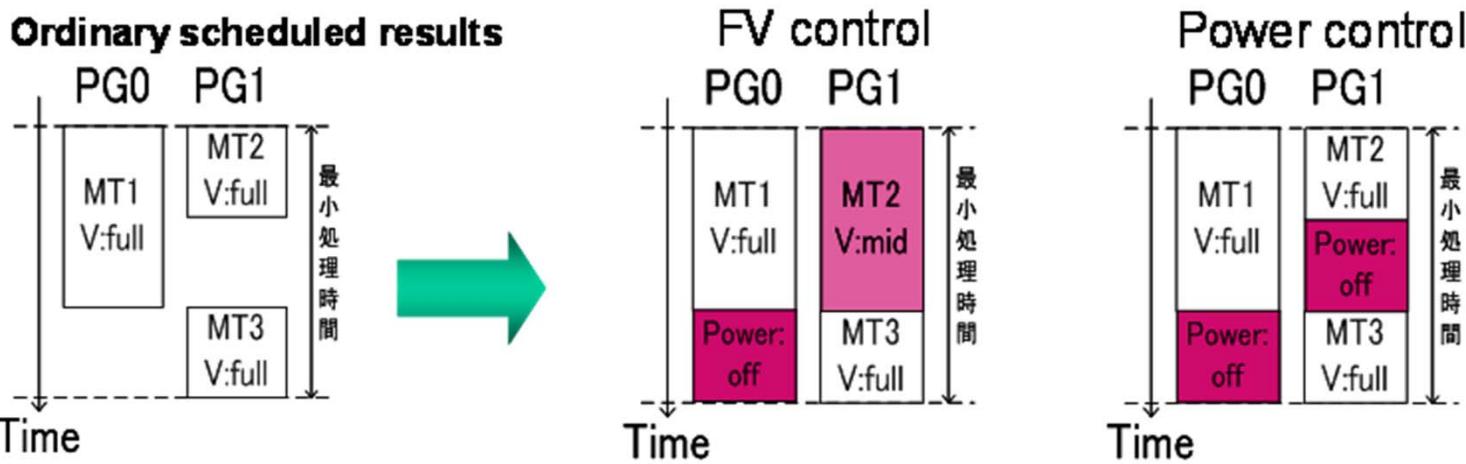
2 hint directives for OSCAR compiler

- accelerator_task
- oscar_comment

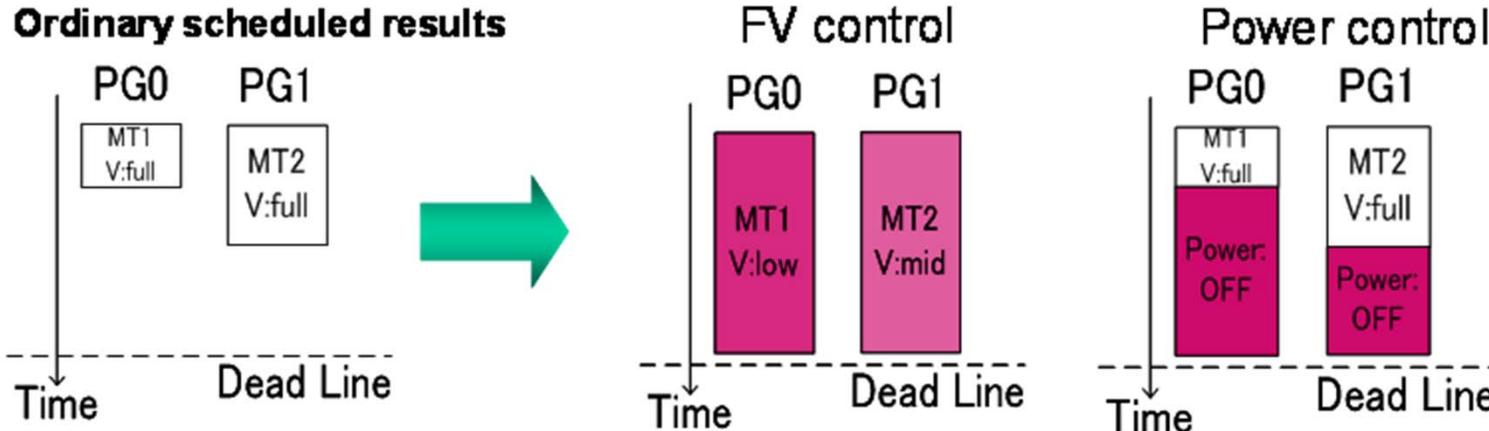
from V2.0

Power Reduction by Power Supply, Clock Frequency and Voltage Control by OSCAR Compiler

- Shortest execution time mode



- Realtime processing mode with dead line constraints



An Example of Machine Parameters for the Power Saving Scheme

- **Functions of the multiprocessor**
 - Frequency of each proc. is changed to several levels
 - Voltage is changed together with frequency
 - Each proc. can be powered on/off

state	FULL	MID	LOW	OFF
frequency	1	1 / 2	1 / 4	0
voltage	1	0.87	0.71	0
dynamic energy	1	3 / 4	1 / 2	0
static power	1	1	1	0

- State transition overhead (Example: not for RP2)

state	FULL	MID	LOW	OFF	state	FULL	MID	LOW	OFF
FULL	0	40k	40k	80k	FULL	0	20	20	40
MID	40k	0	40k	80k	MID	20	0	20	40
LOW	40k	40k	0	80k	LOW	20	20	0	40
OFF	80k	80k	80k	0	OFF	40	40	40	0

Power Reduction Scheduling

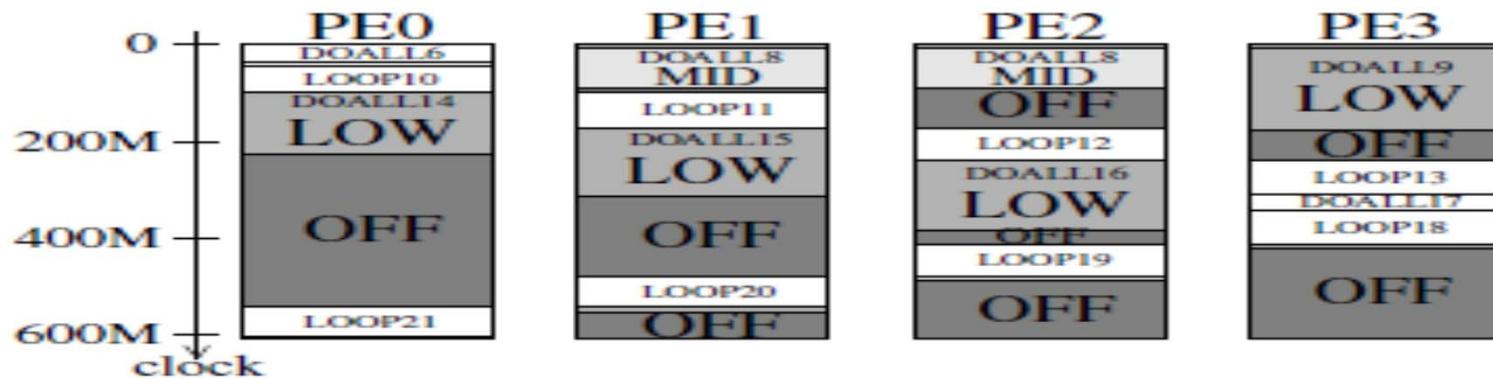
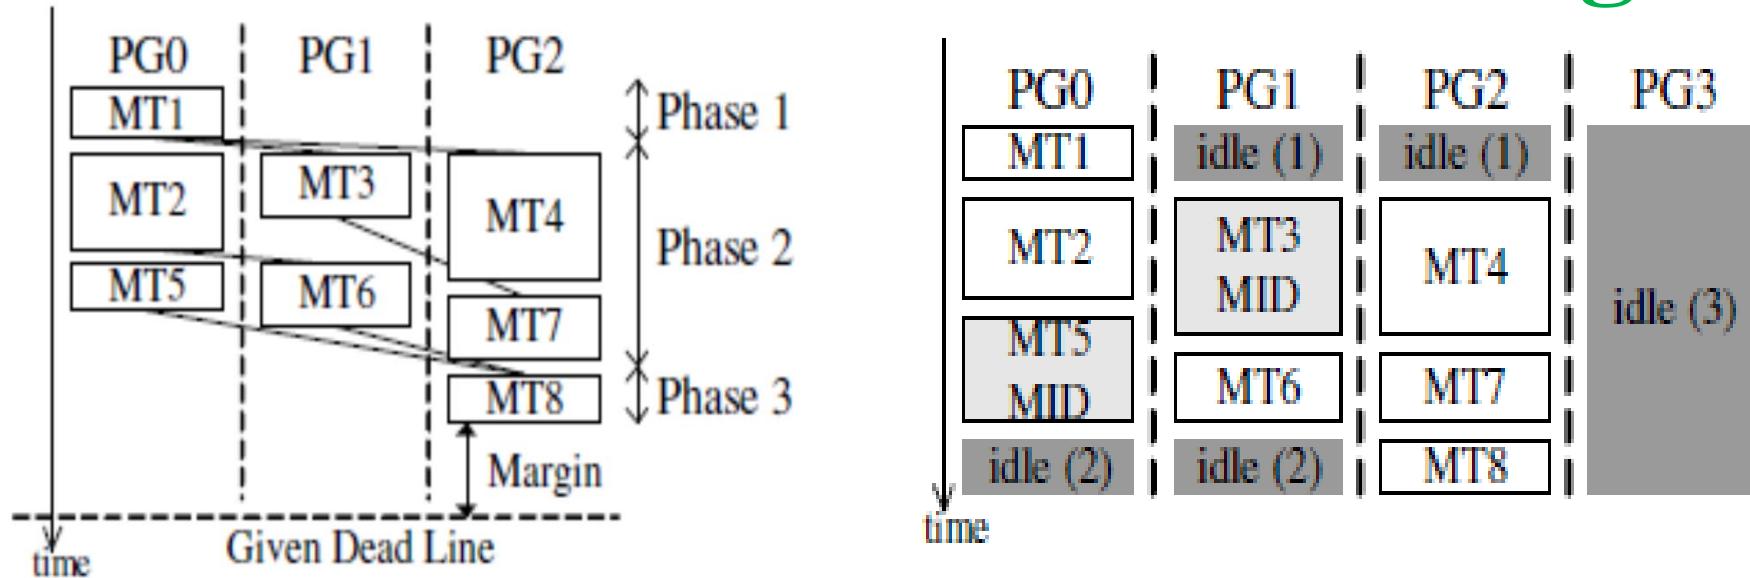
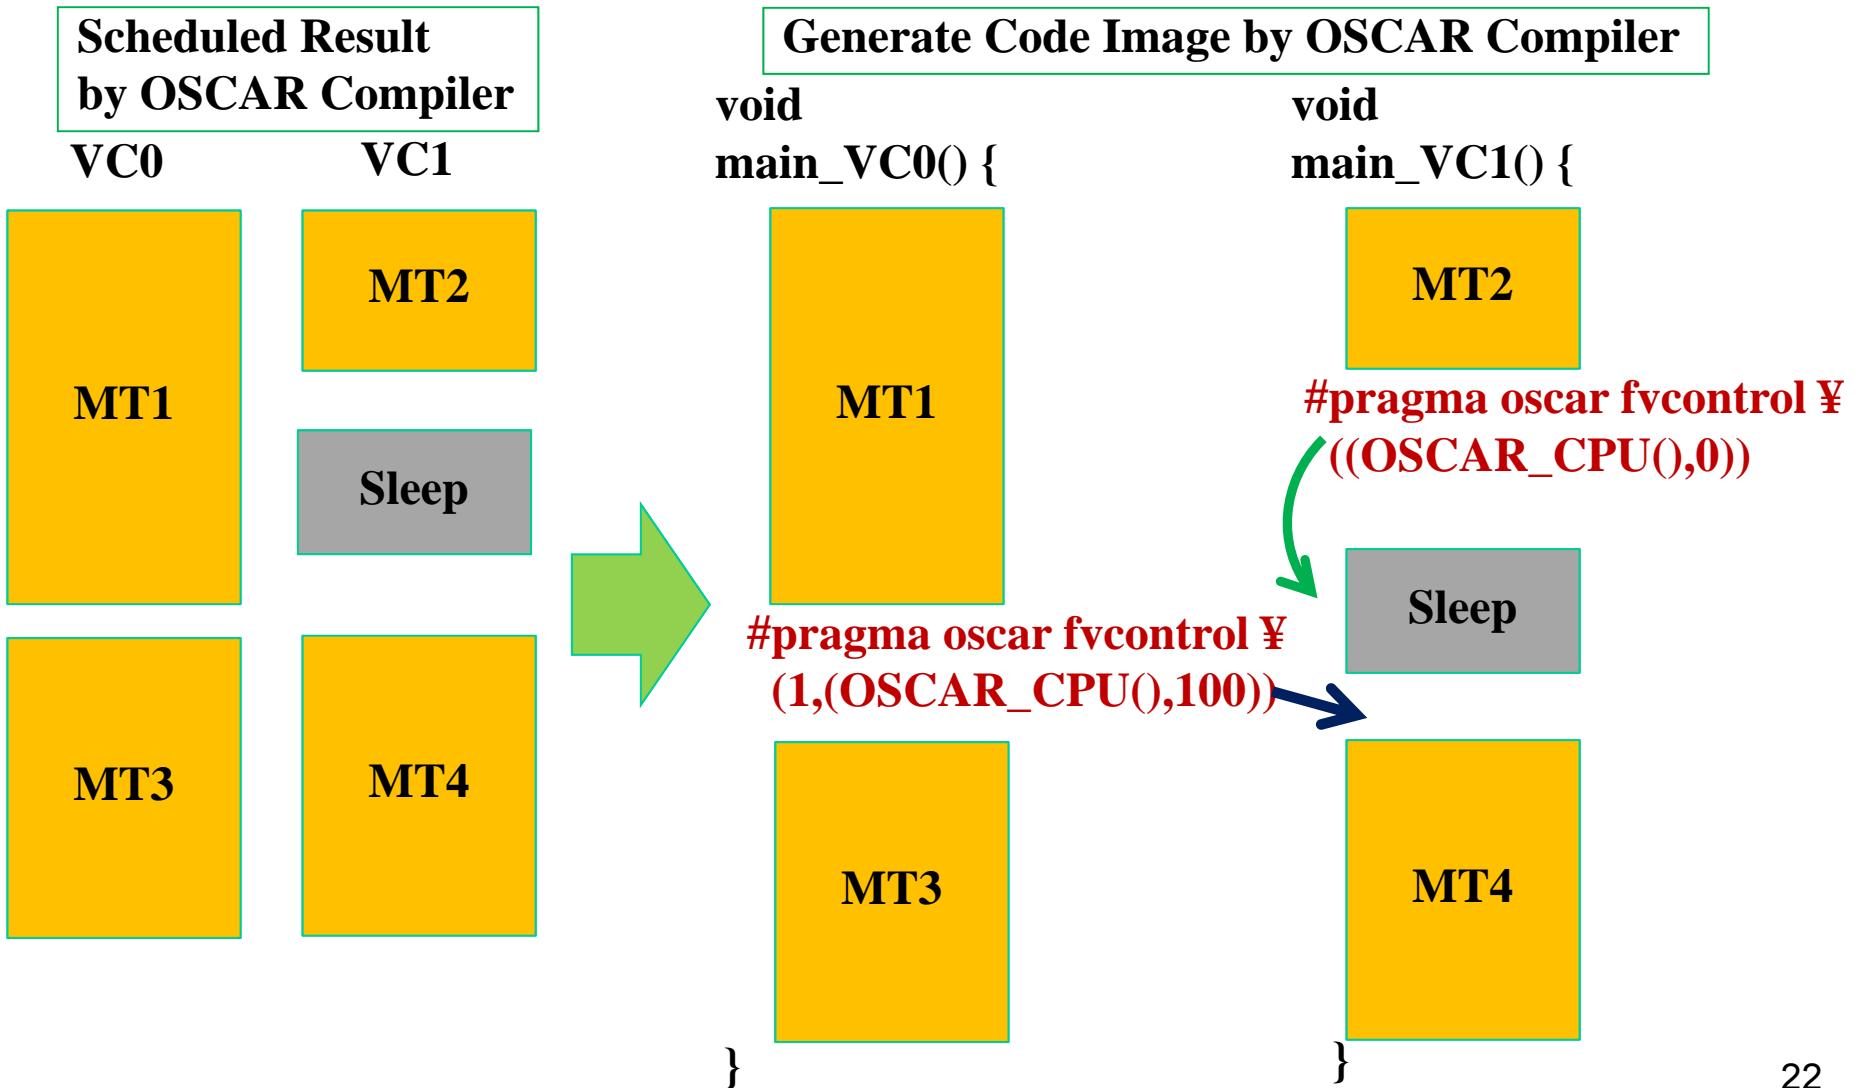
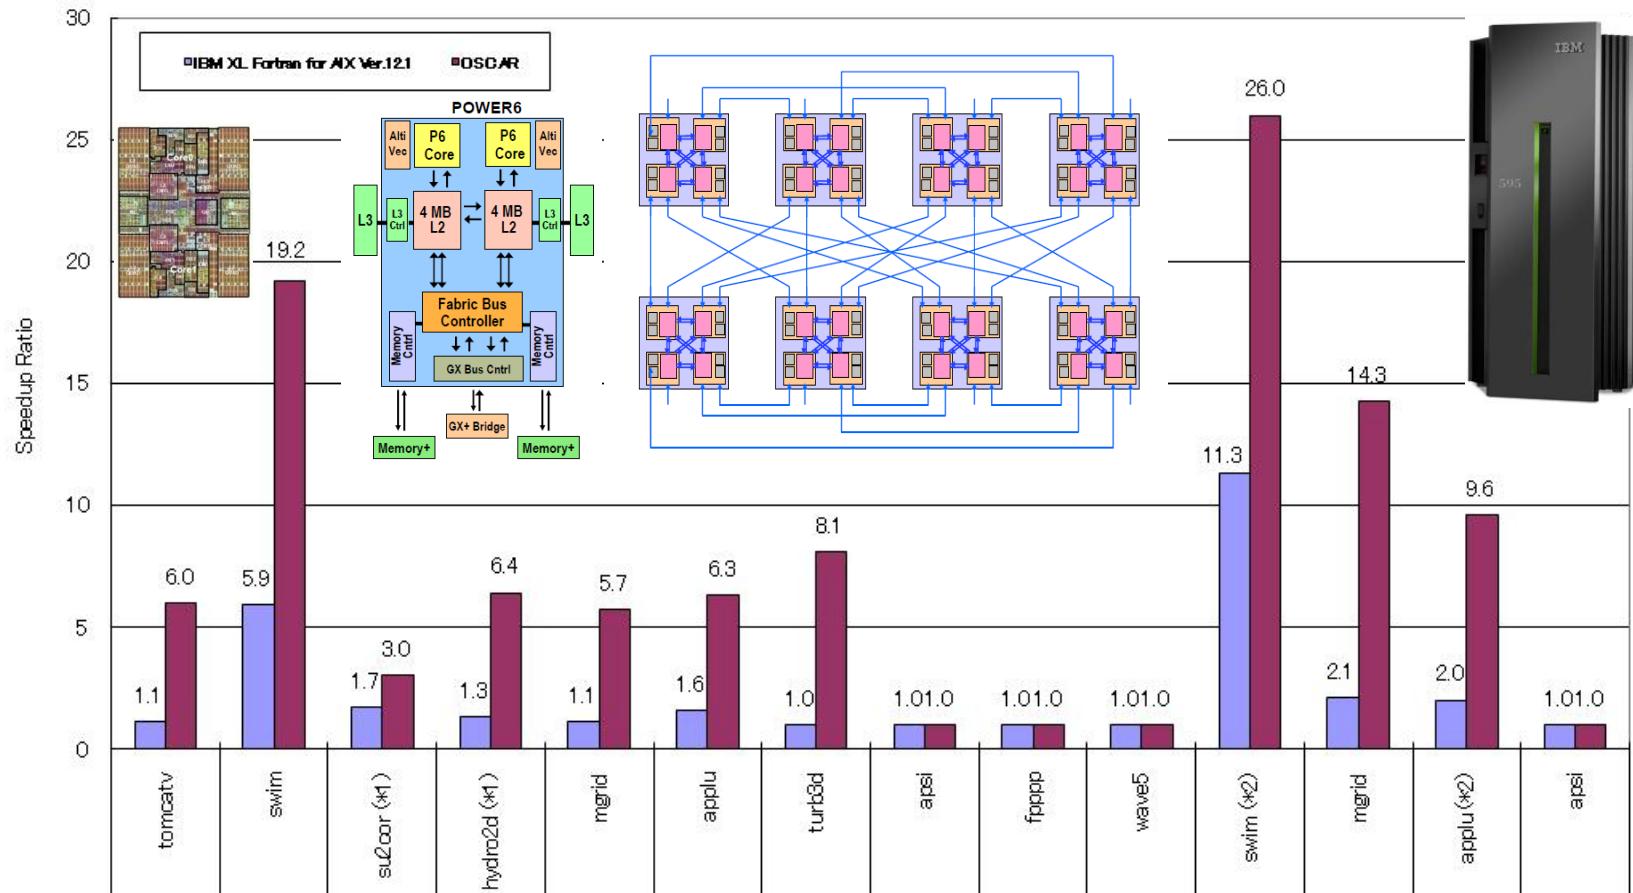


Fig. 6. V/F control of applu(4proc.)

Low-Power Optimization with OSCAR API



Performance of OSCAR Compiler on IBM p6 595 Power6 (4.2GHz) based 32-core SMP Server



OpenMP codes generated by OSCAR compiler accelerate IBM XL Fortran for AIX Ver.12.1 about 3.3 times on the average

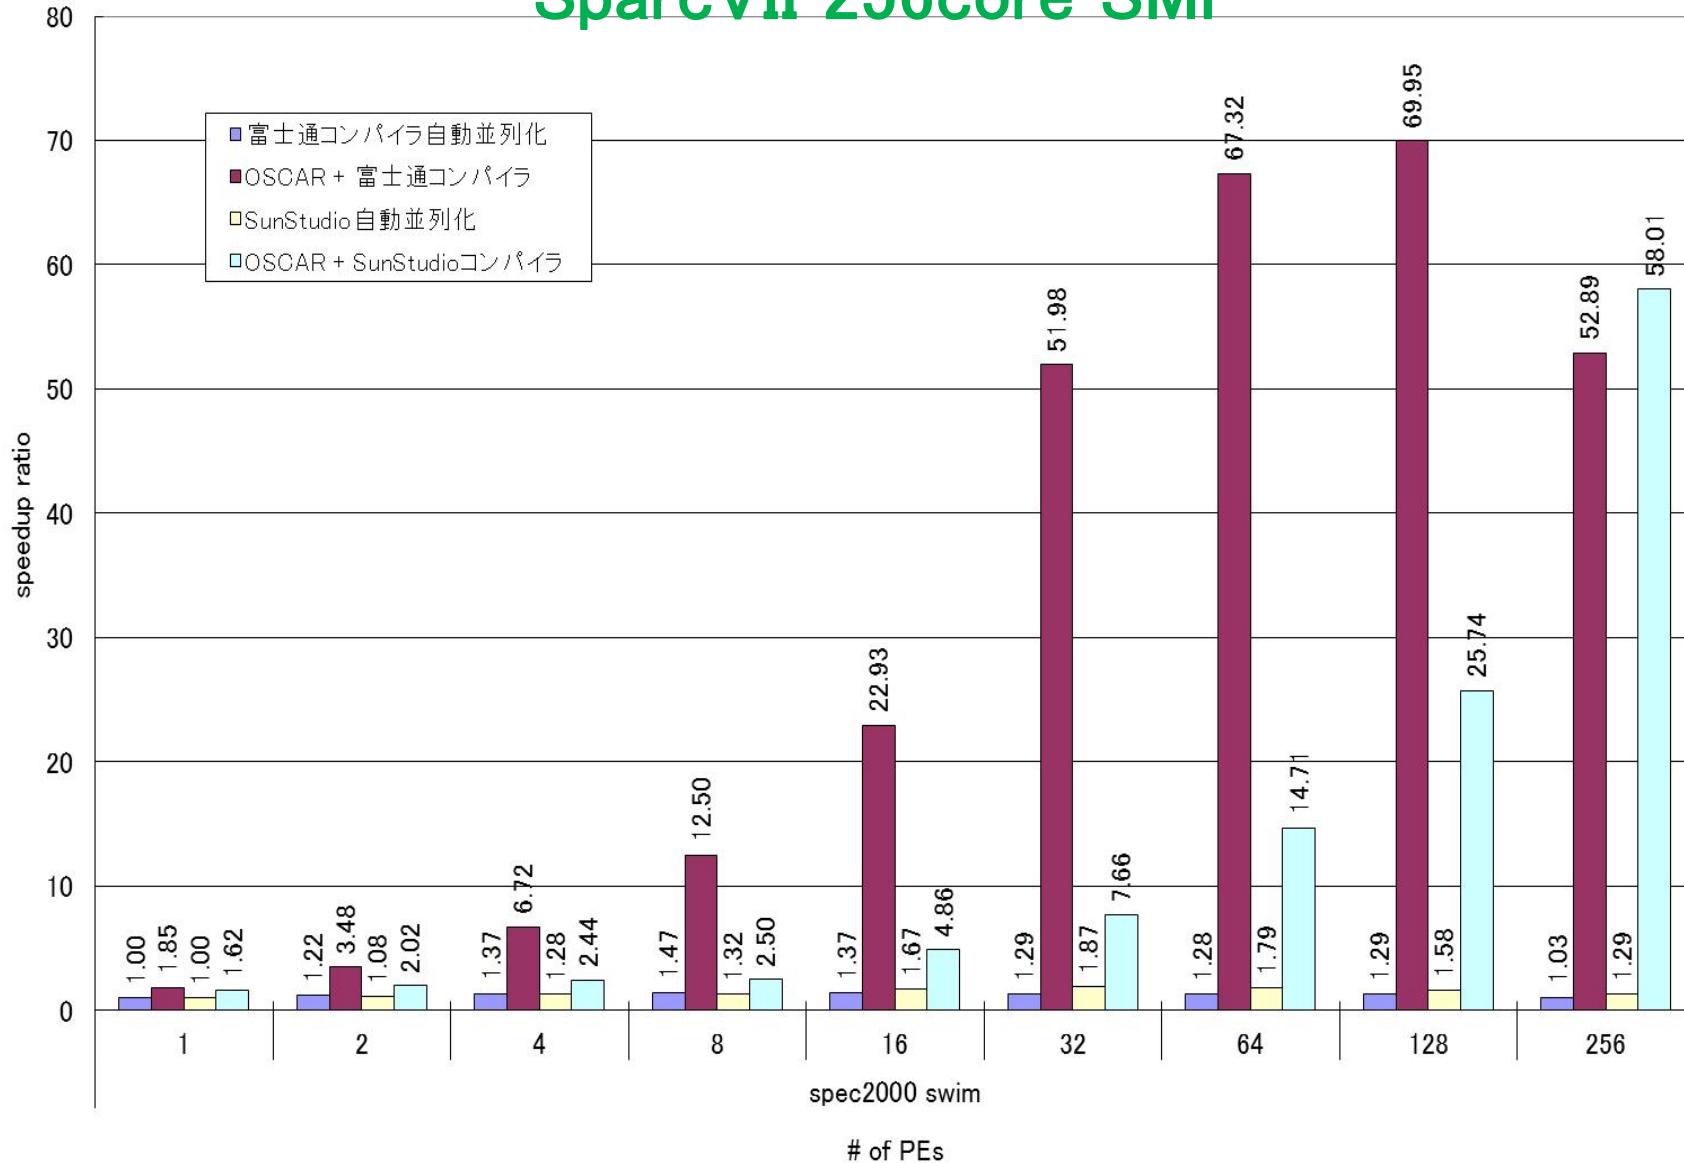
Compile Option:

(*1) Sequential: -O3 -qarch=pwr6, XLF: -O3 -qarch=pwr6 -qsmp=auto, OSCAR: -O3 -qarch=pwr6 -qsmp=noauto

(*2) Sequential: -O5 -q64 -qarch=pwr6, XLF: -O5 -q64 -qarch=pwr6 -qsmp=auto, OSCAR: -O5 -q64 -qarch=pwr6 -qsmp=noauto

(Others) Sequential: -O5 -qarch=pwr6, XLF: -O5 -qarch=pwr6 -qsmp=auto, OSCAR: -O5 -qarch=pwr6 -qsmp=noauto

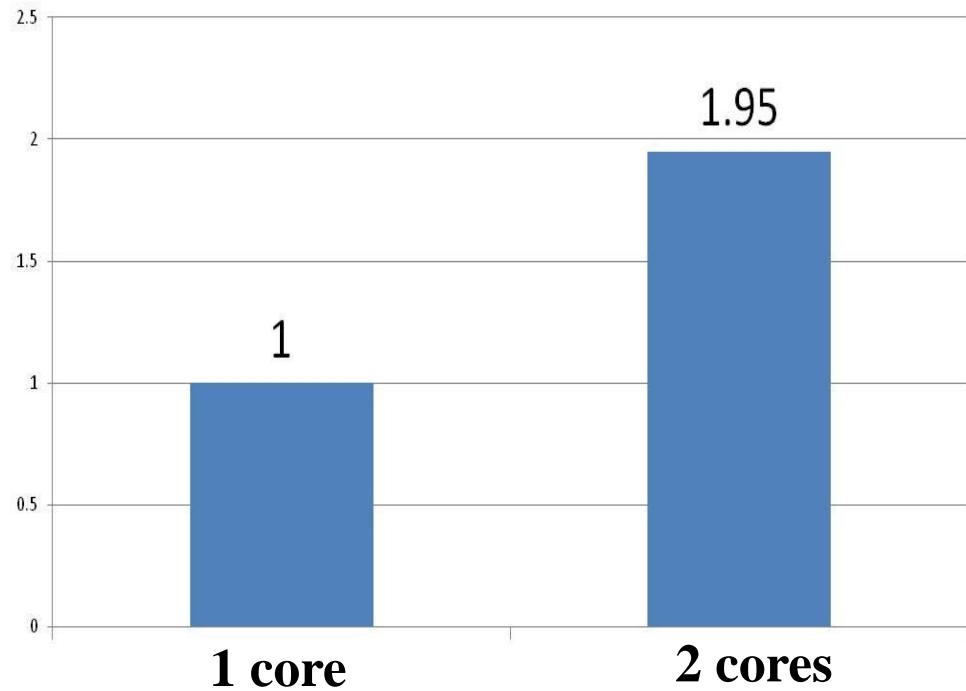
OSCAR Compiler's Performance on Fujitsu9000 SparcVII 256core SMP





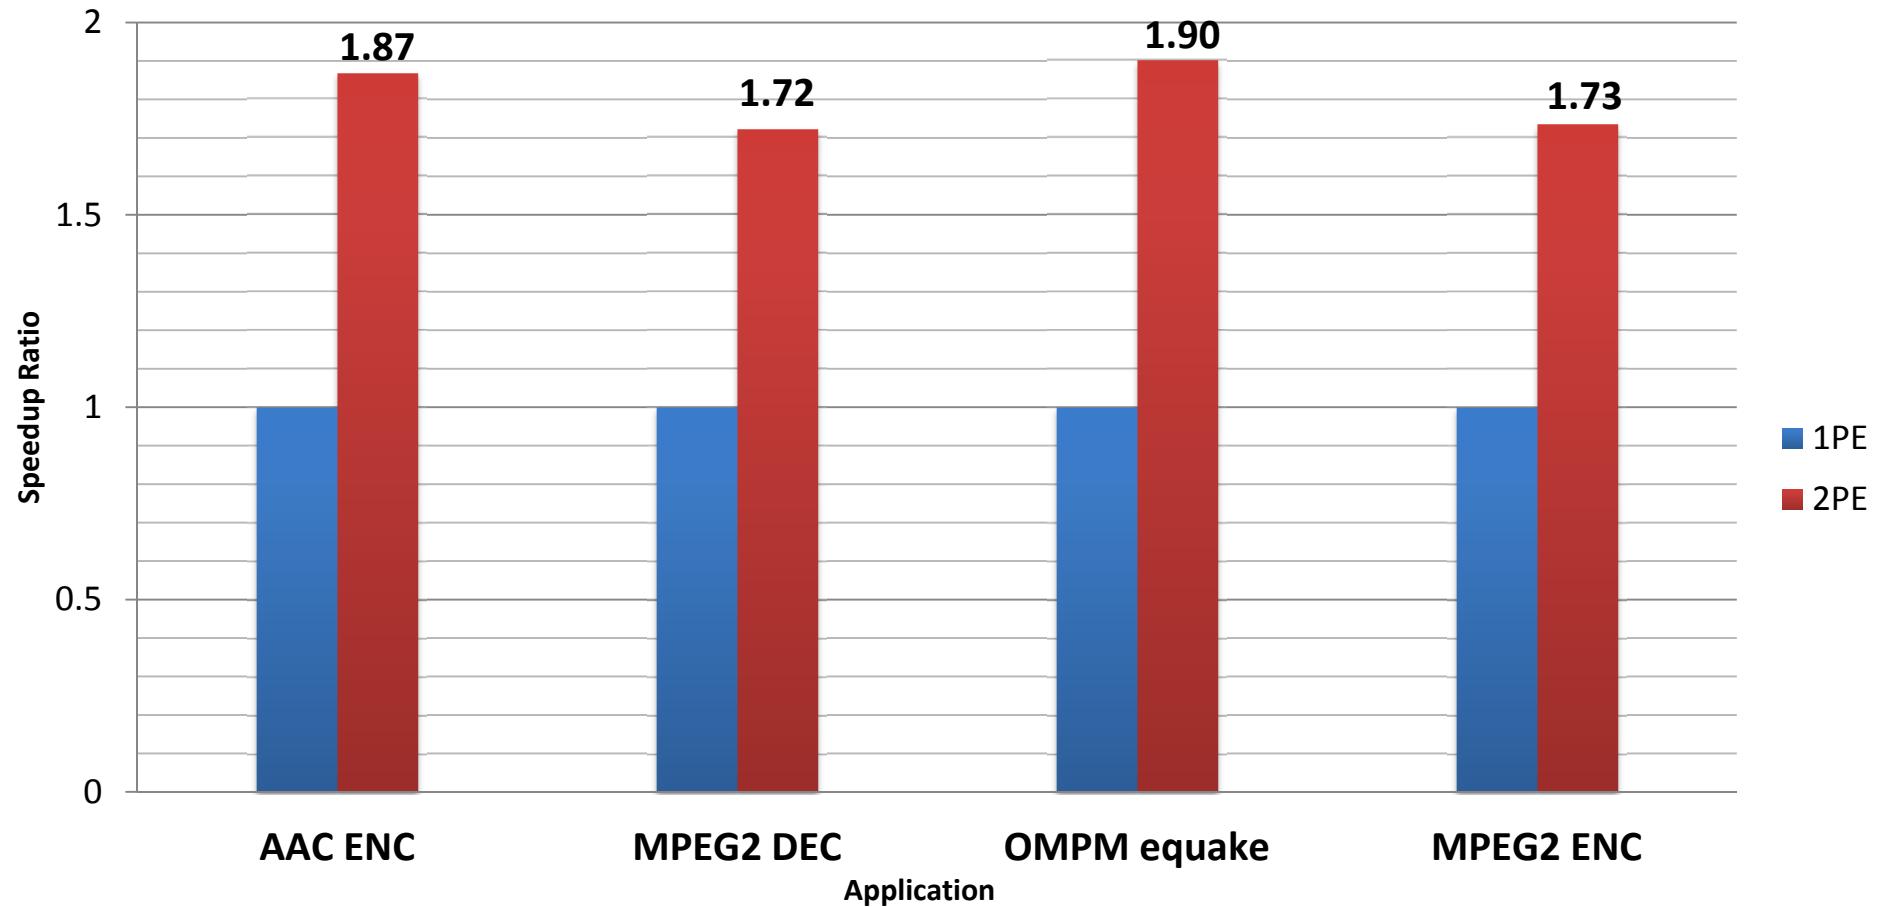
Engine Control by multicore with Denso

Though so far parallel processing of the engine control on multicore has been very difficult, Denso and Waseda succeeded 1.95 times speedup on 2core V850 multicore processor.



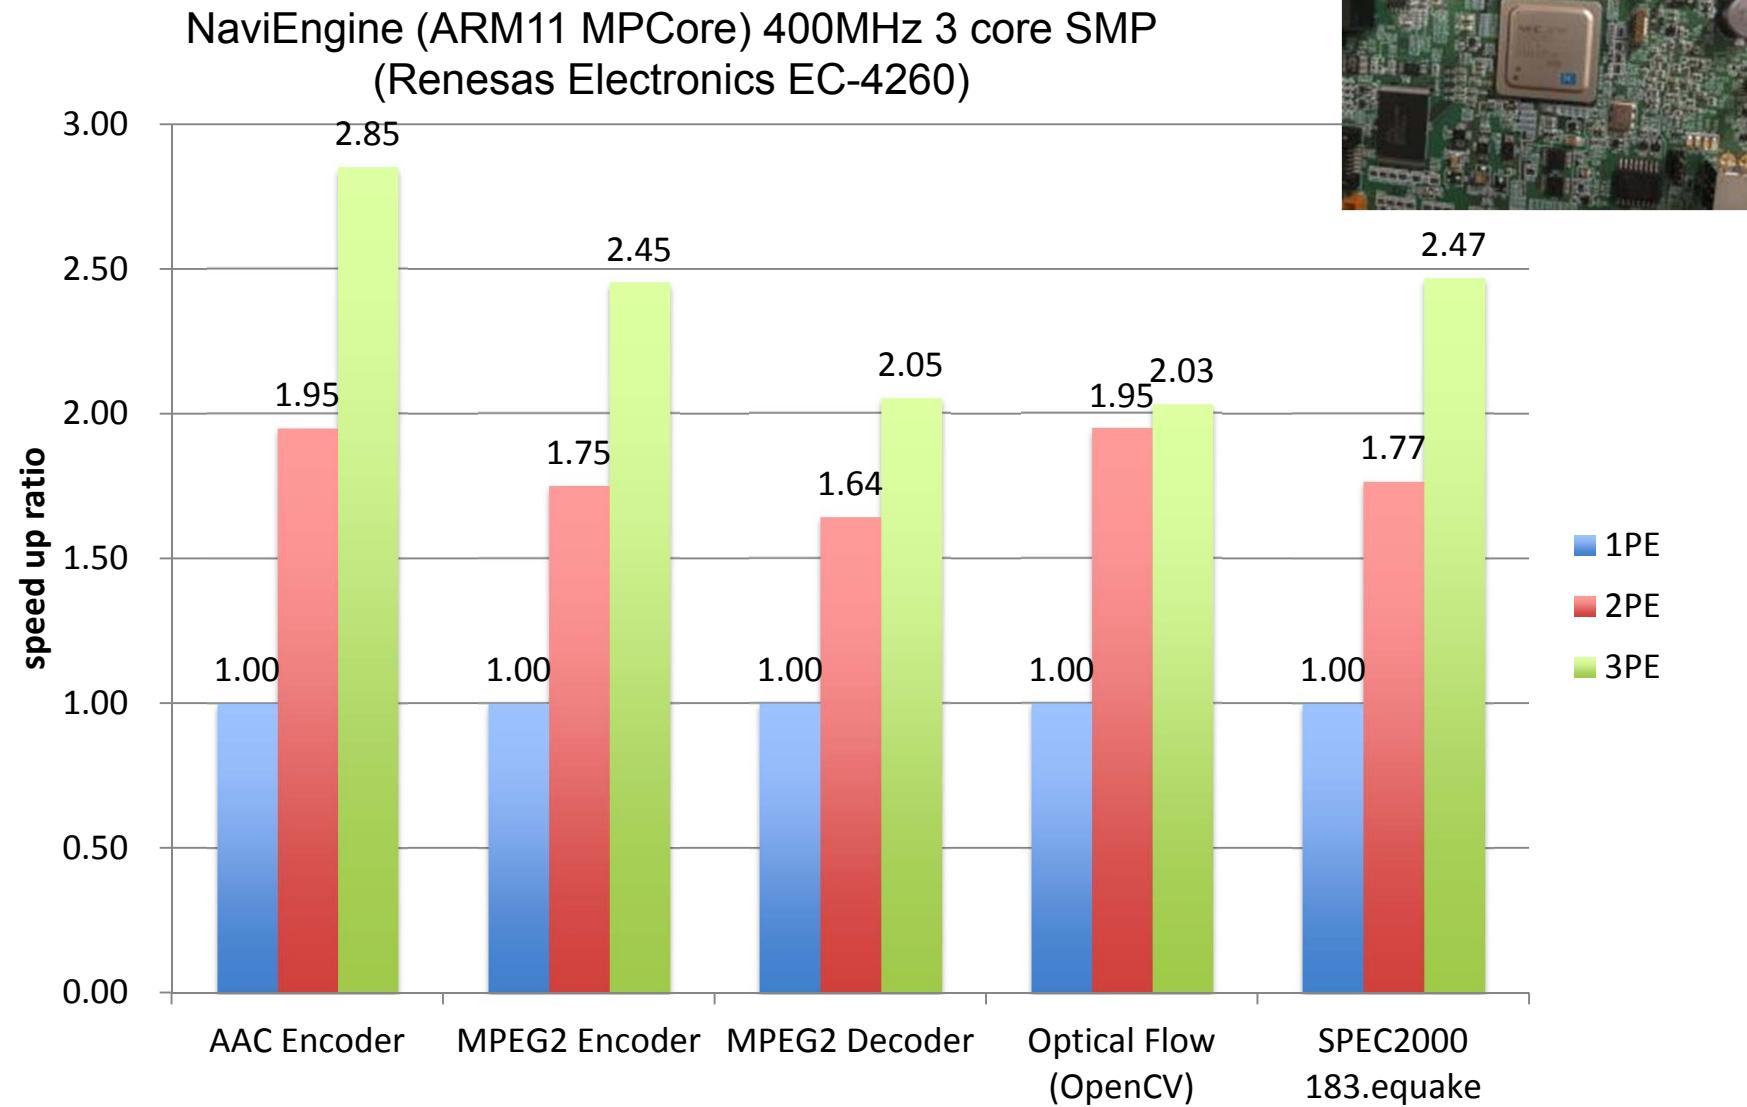
Engine control by multicore
Hard real-time processing

Performance of OSCAR Compiler & API on 2 ARMv7-cores Qualcomm MSM8960 (Snapdragon) Android 4.0 for Smart Phones



1.81 times speedup by 2 cores on the average against 1 core

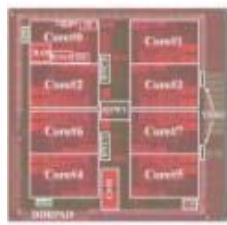
Parallel Processing Performance on 3Cores NaviEngine with Realtime OS eT-Kernel Multi-Core Edition



- 2.37 times speedup on 3ARM cores against 1 core

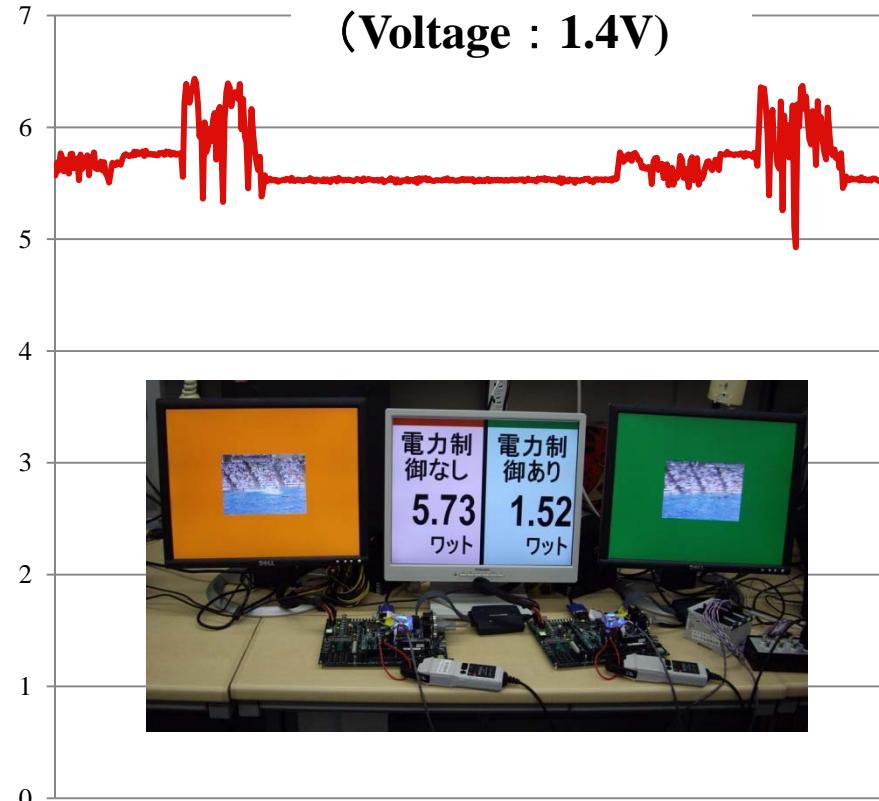
Power Reduction of MPEG2 Decoding to 1/4 on 8 Core Homogeneous Multicore RP-2 by OSCAR Parallelizing Compiler

MPEG2 Decoding with 8 CPU cores



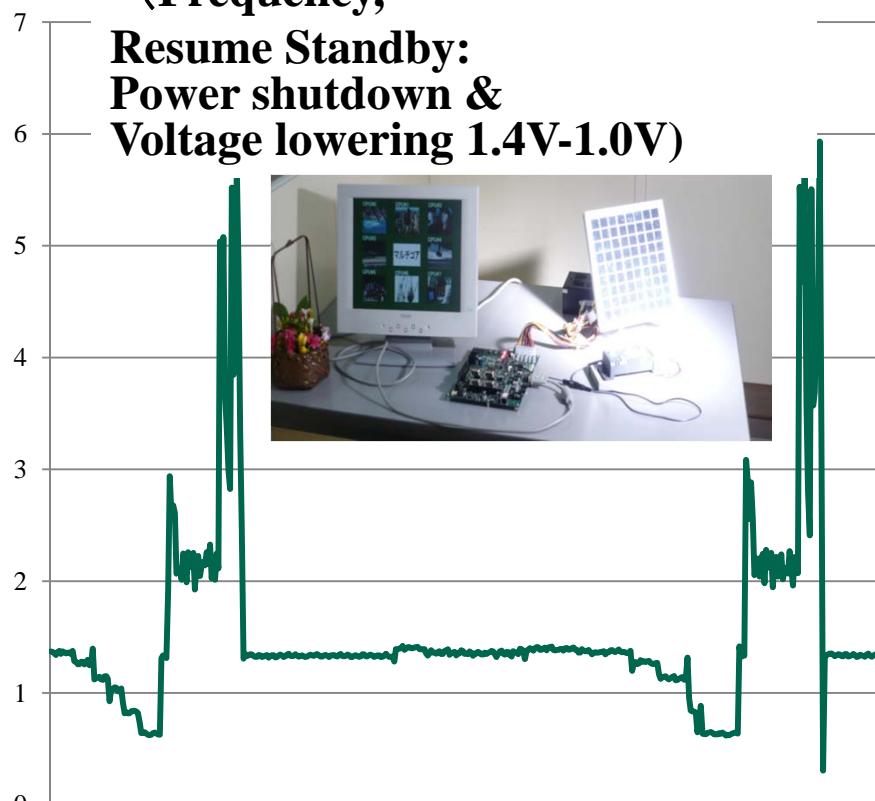
Without Power Control

(Voltage : 1.4V)



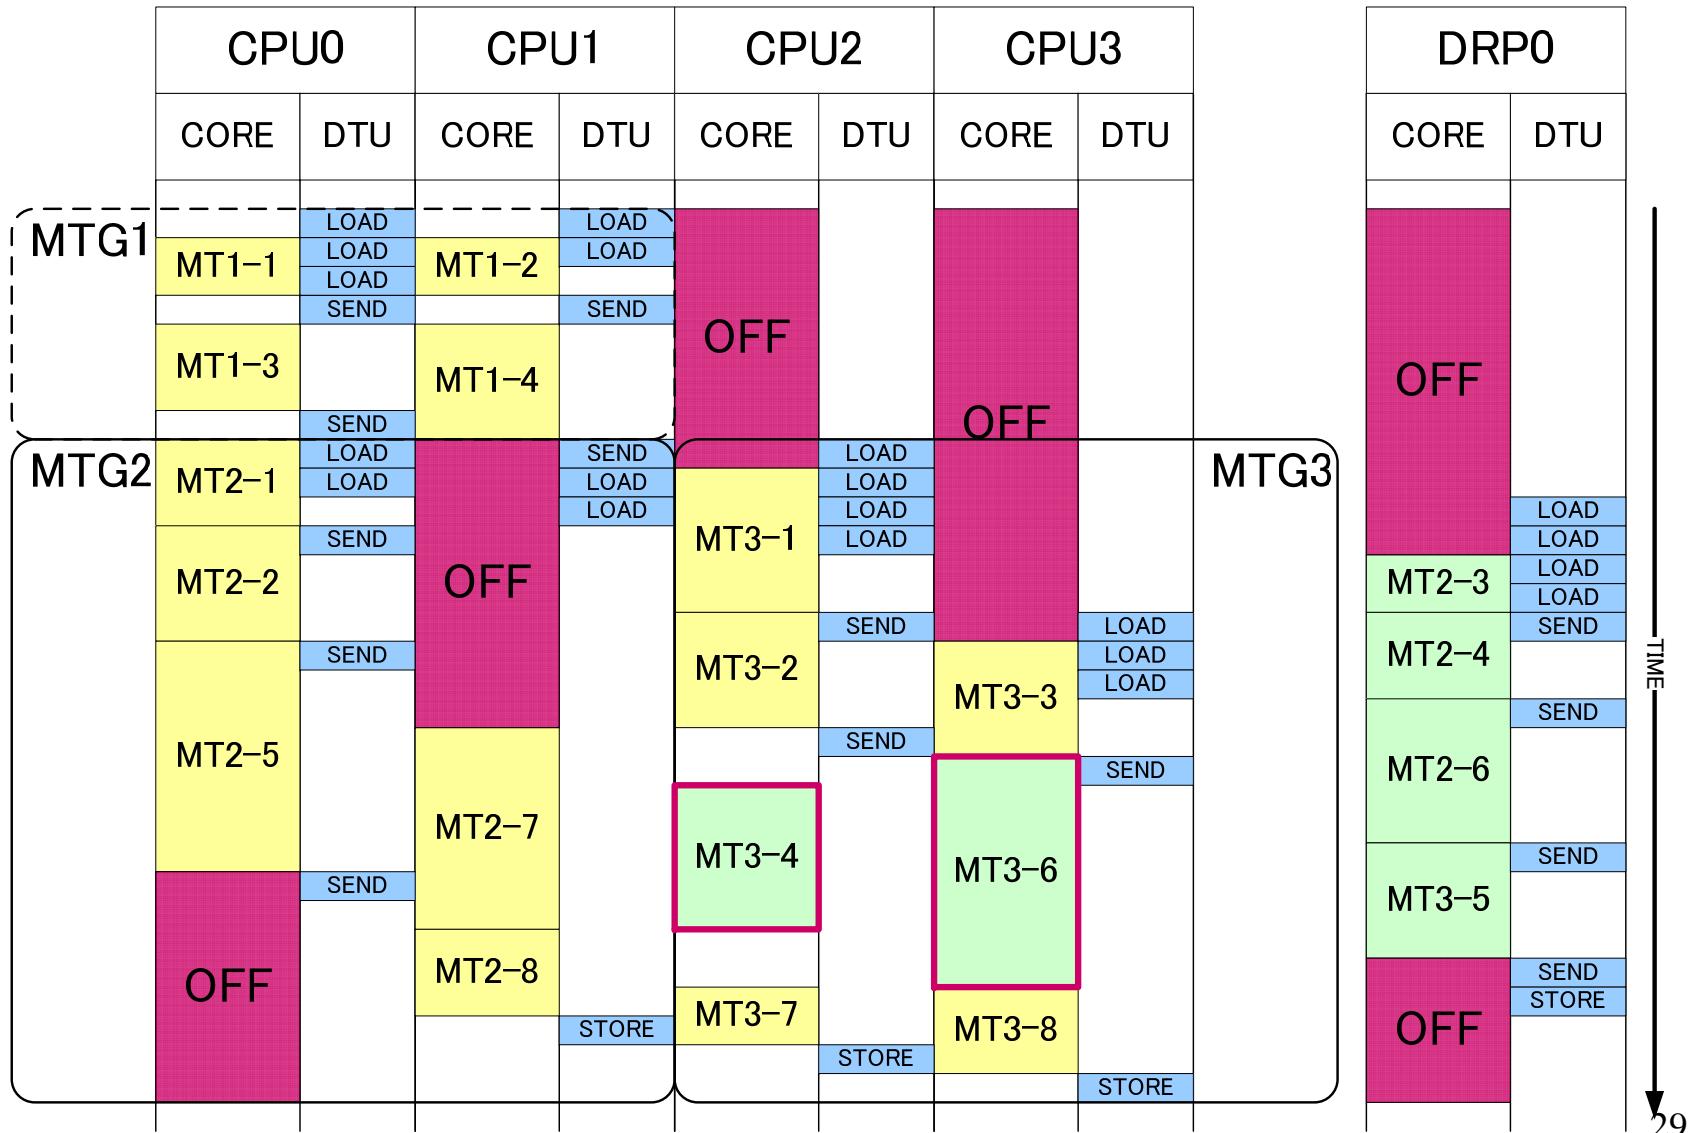
With Power Control (Frequency,

Resume Standby:
Power shutdown &
Voltage lowering 1.4V-1.0V)

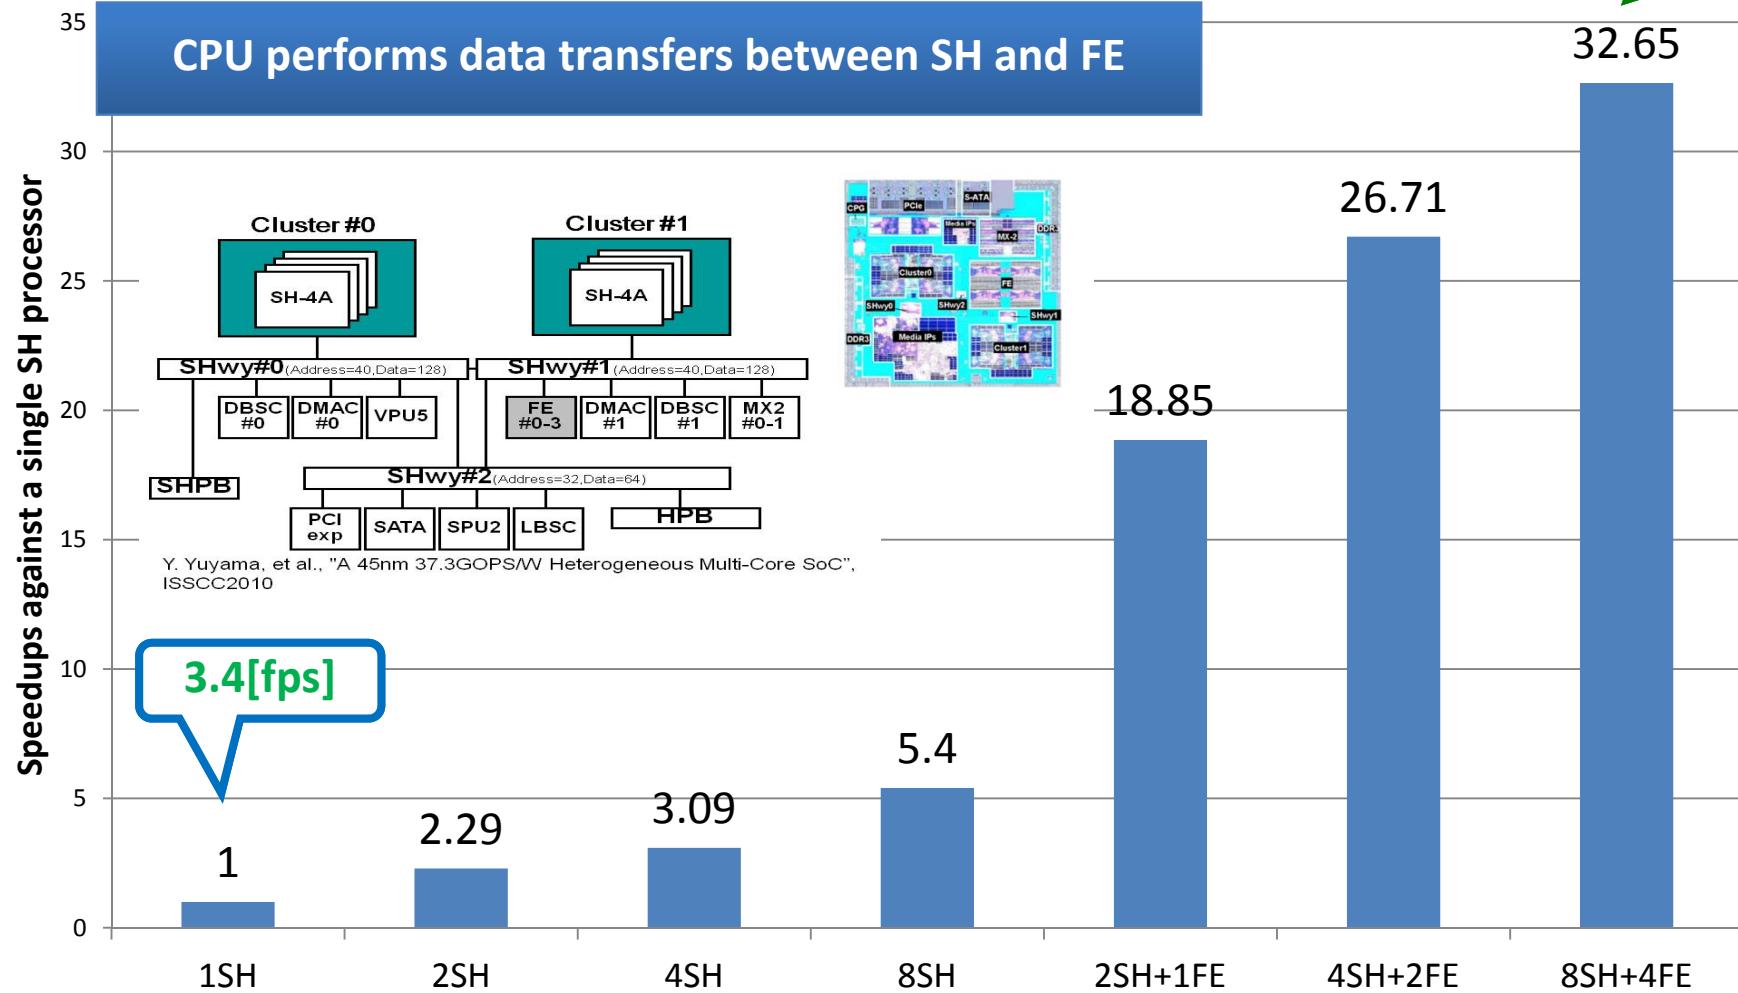


73.5% Power Reduction

An Image of Static Schedule for Heterogeneous Multi-core with Data Transfer Overlapping and Power Control



33 Times Speedup Using OSCAR Compiler and OSCAR API on RP-X (Optical Flow with a hand-tuned library)



Power Reduction in a real-time execution controlled by OSCAR Compiler and OSCAR API on RP-X (Optical Flow with a hand-tuned library)

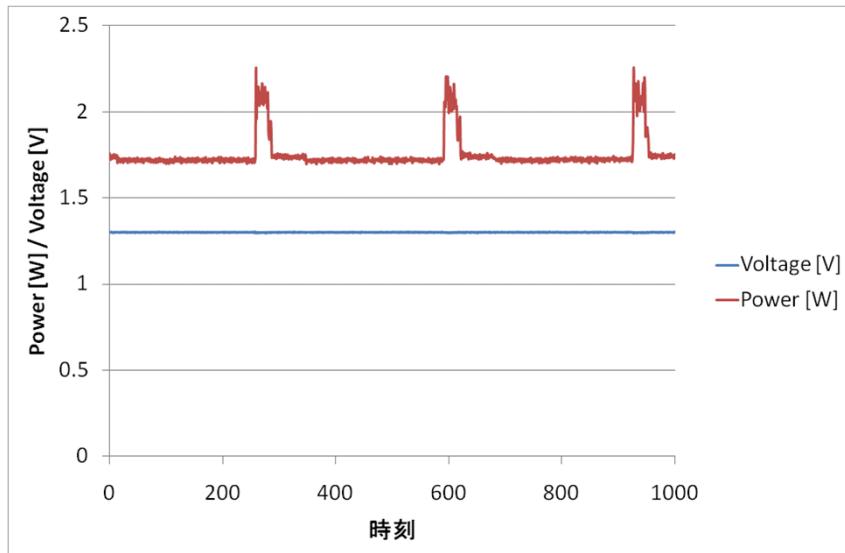
Without Power Reduction

With Power Reduction
by OSCAR Compiler

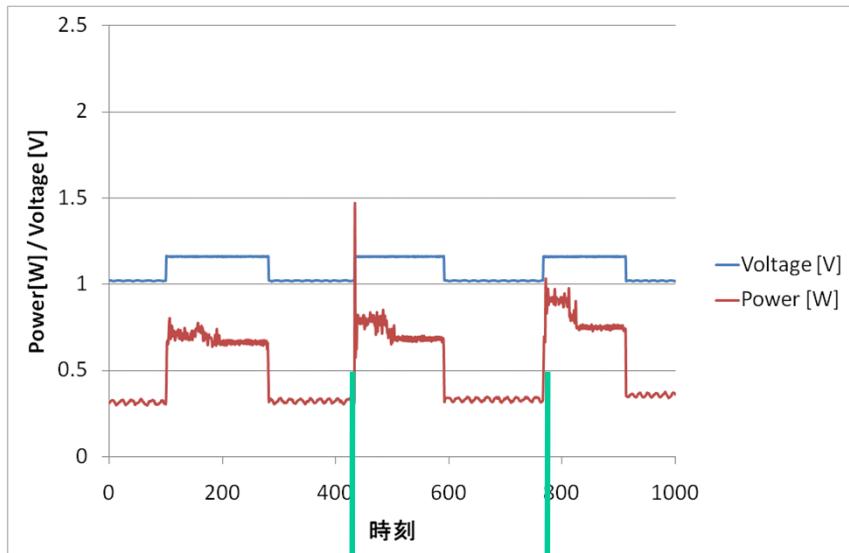
70% of power reduction

Average: 1.76[W]

Average: 0.54[W]

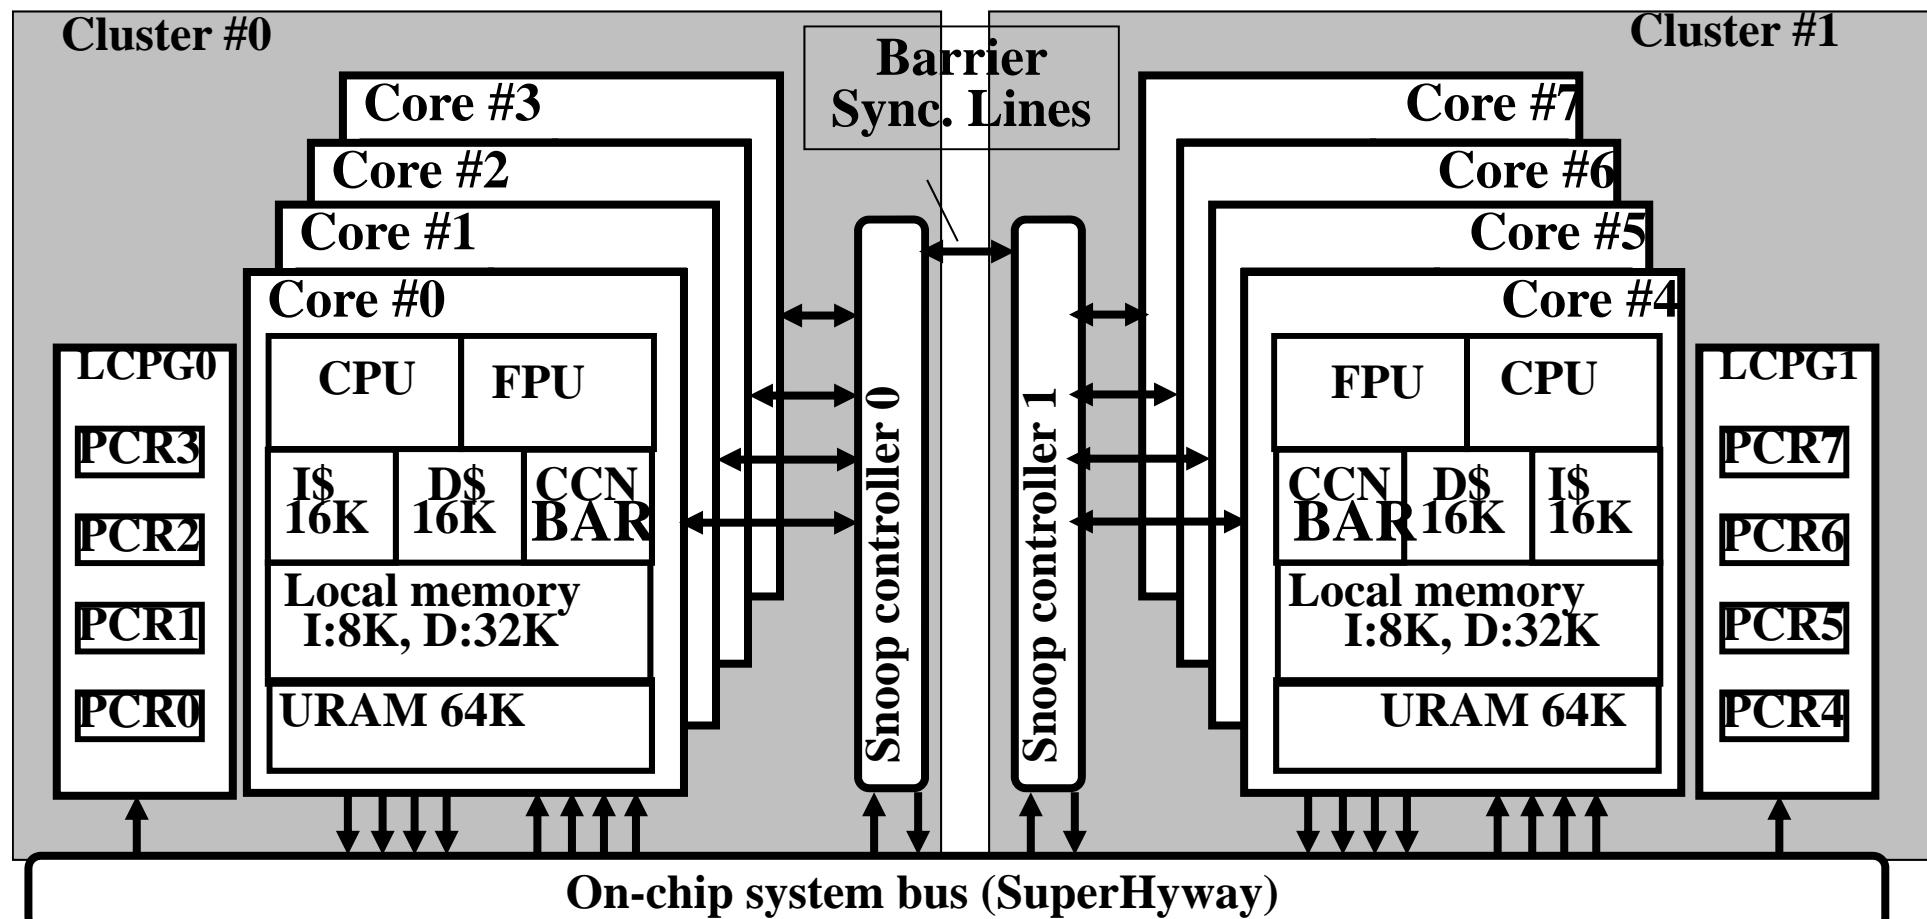


1cycle : 33[ms]
→30[fps]



時刻

8 Core RP2 Chip Block Diagram



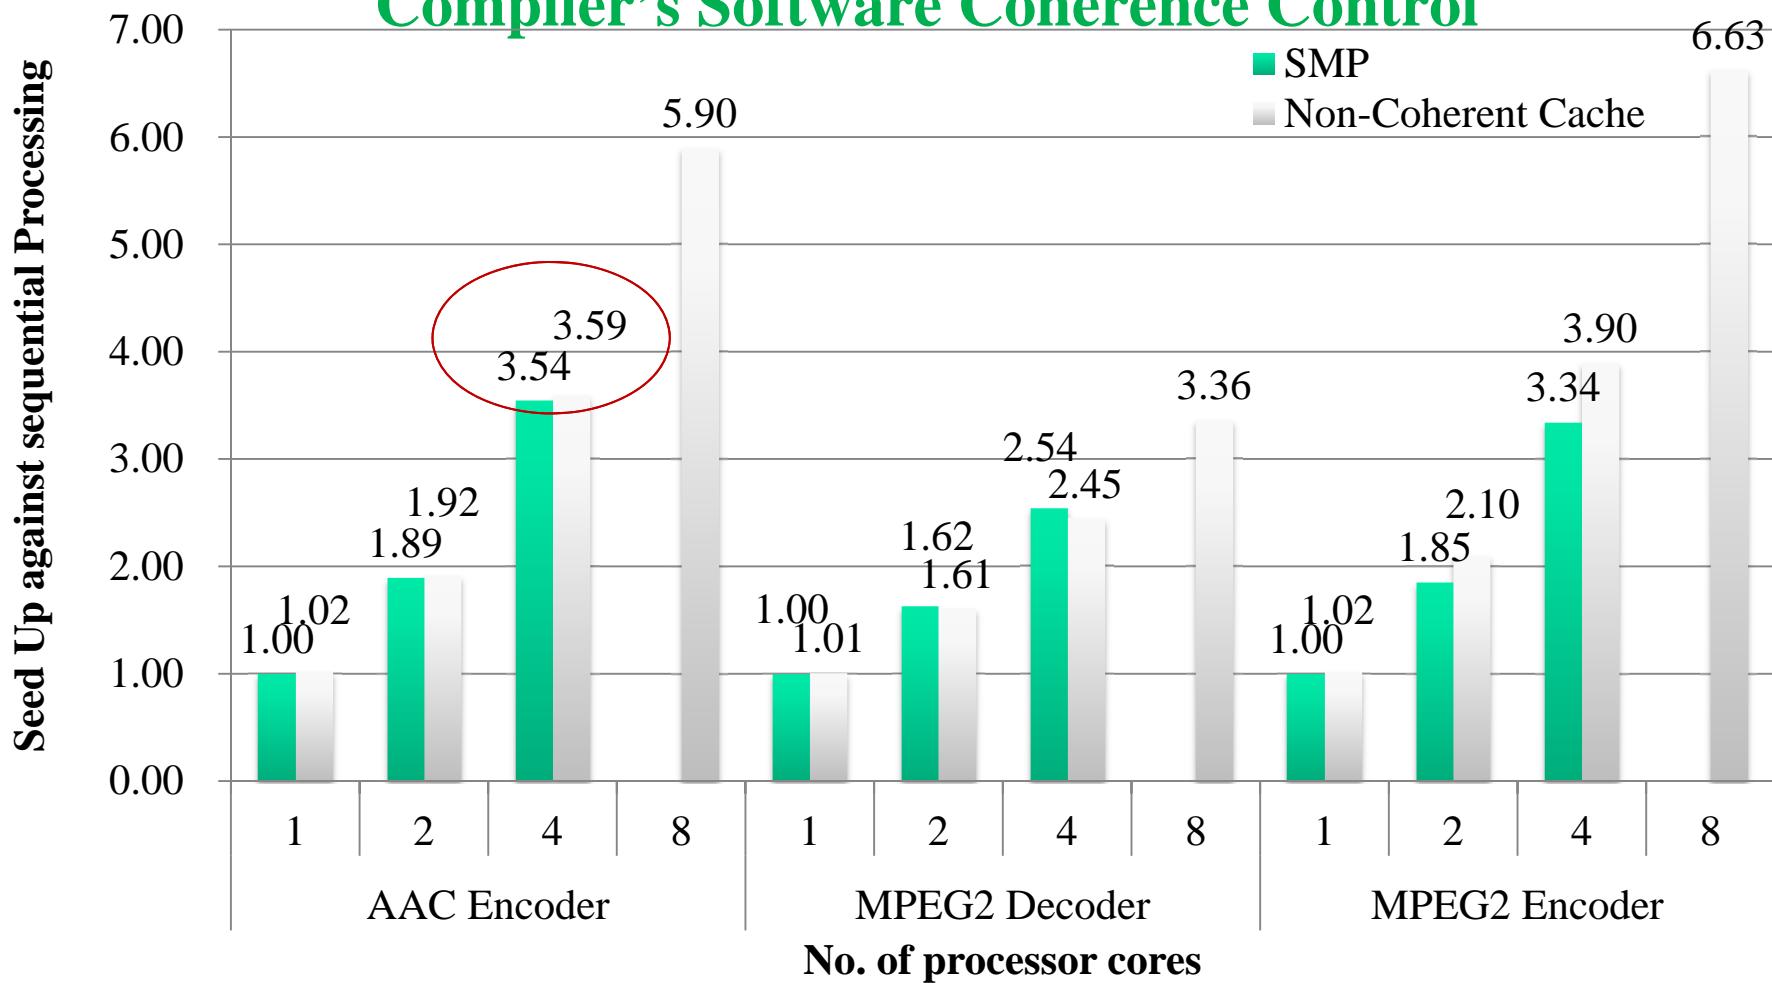
LCPG: Local clock pulse generator

PCR: Power Control Register

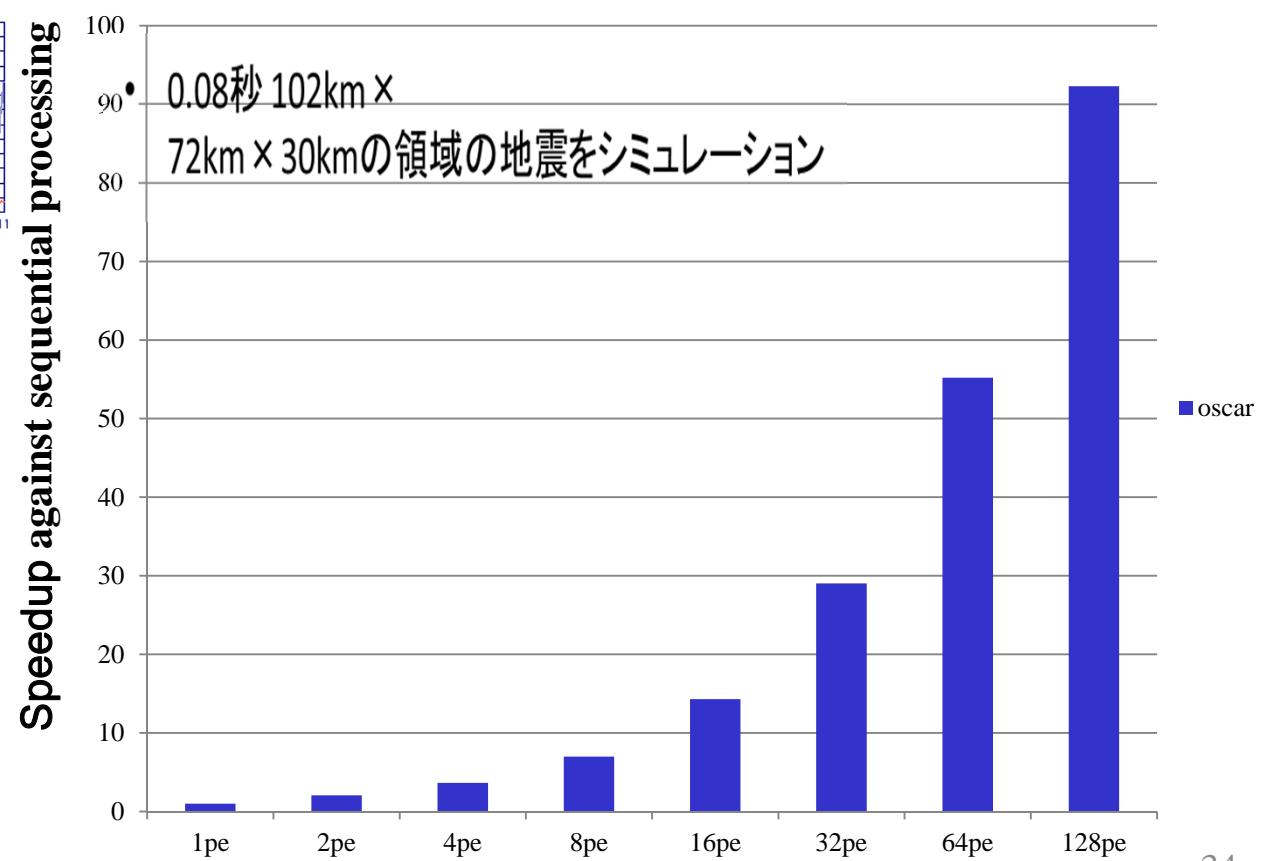
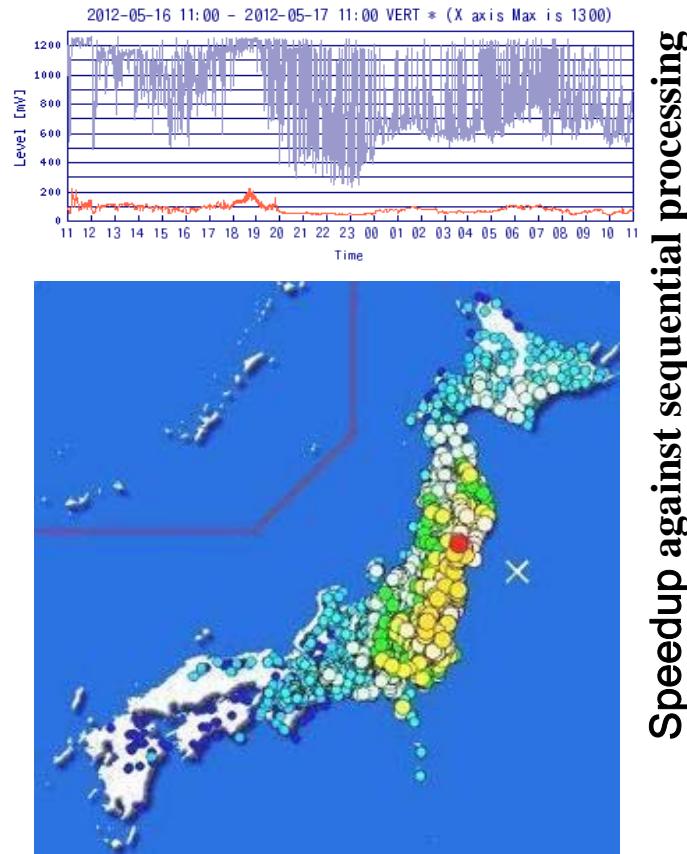
CCN/BAR: Cache controller/Barrier Register

URAM: User RAM (Distributed Shared Memory)

Faster or Equal Processing Performance with Hardware Coherence Control on 8 core RP2 Multicore Precessor Having Hardware Coherent Mechanism Up-to 4 cores by OSCAR Compiler's Software Coherence Control

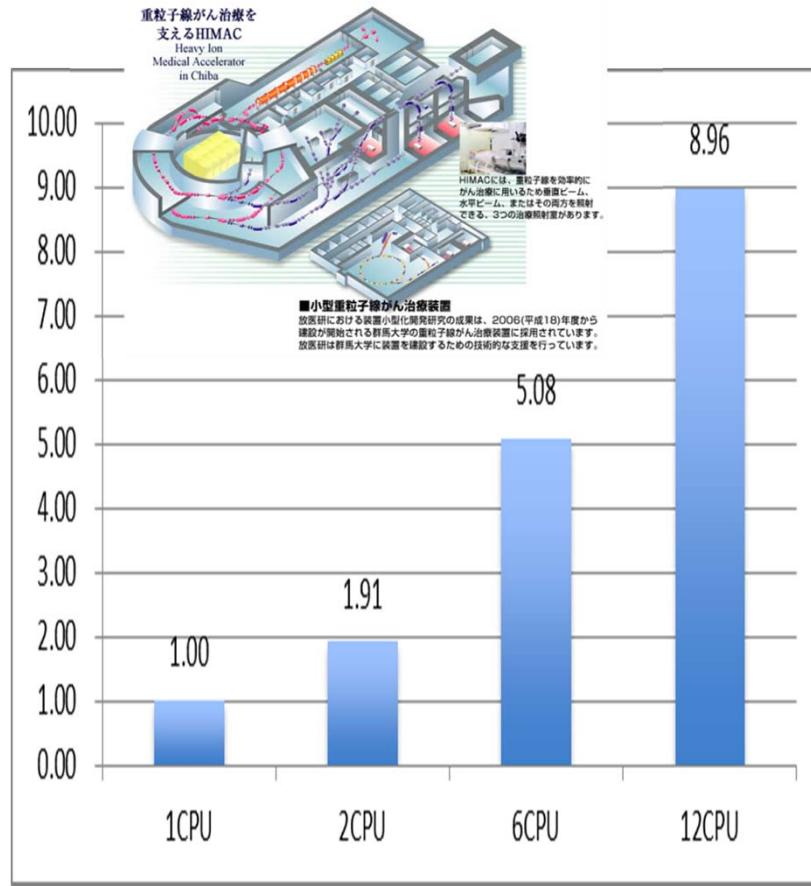


92 Times Speedup against the Sequential Processing for GMS Earthquake Wave Propagation Simulation on Hitachi SR16000 (Power7 Based 128 Core Linux SMP)



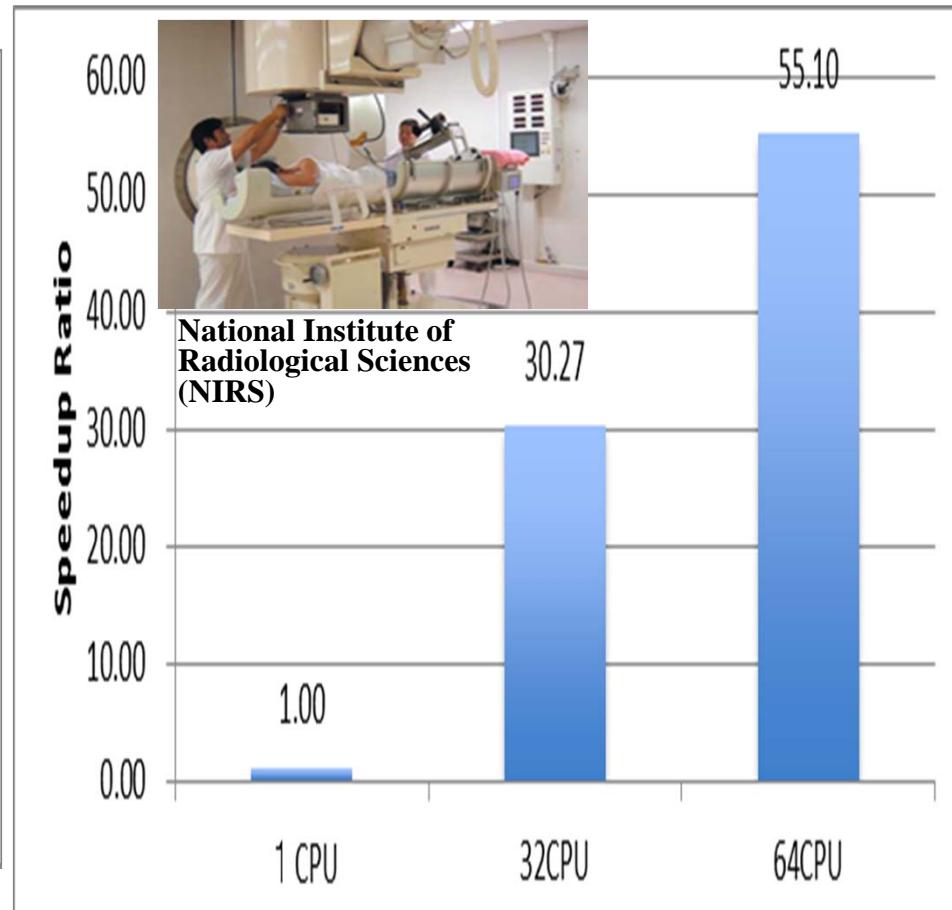
Cancer Treatment Carbon Ion Radiotherapy

(Previous best was 2.5 times speedup on 16 processors with hand optimization)



8.9times speedup by 12 processors

Intel Xeon X5670 2.93GHz 12
core SMP (Hitachi HA8000)



55 times speedup by 64 processors

IBM Power 7 64 core SMP
(Hitachi SR16000)

CS Multicore STC Team

Chair

Hironori Kasahara

FTs Proj. Mgr.

? (to be hired)

BoG "Angel"

Hironori Kasahara



Conferences

- <Hard+Soft+Industrial Applications from Embedded to HPC Systems with Gov., Acad.&Indus.>
- Co-Chair Josep Torrellas (UIUC)
- Co-Chair Hironori Kasahara
- Vivek Sarkar (Rice U.)
- Dr. Ahmed Jerraya, CEA-LETI, MINATEC, Fr
- + Industry: Automobile, Smart Phone, Medical etc
- Carrie Walsh (SE)

Standards

- Architecture Committee
3D Integ., Memory (Non volatile),etc
- Software Committee
API, Development Env. etc
- Industrial Application Comm.
 - Consumer Electronics (Smart Phones): ATT, NTT, Apple,
 - Automobile (GM, Mercedes, Toyota, Bosch, Denso, etc)
 - Medical (Varian,Hitachi,Semens,etc)
 - Anne Marie Kelly (SE)

Publishing

- <Online Publication for quick and low cost>
- Trans. on Multicores
- Multicore Magazine
 - <Start as online publication through Web Portal: Not only written papers and also Online Presentation by especially Industry Leaders>
 - <Think introduction of "Mileage system" for Editorial, Programming Committee members and reviewers
- Lars Jentsch (SE) Alicia Stickley (SE)

Education

- Start from "Online Lecture":
Ask the lecture to the world best researcher for the topics
- David Padua(UIUC)
 - Dorian McClenahan (SE)

Body of Knowledge

- /
- ### Thesaurus
- Based on Parallel Processing Encyclopedia
- David Padua & others
 - Dante David (SE)

Newsletter

- First ,with Conference, Education and Publishing push the latest attractive information to members.
- ?
- Margo McCall (SE)

Conclusions

- OSCAR compiler automatic parallelizes C or Fortran program using multigrain parallelization, data localization for cache and local memory with DMA data transfers and generates C or Fortran parallelized code with OSCAR API version 2.0.
- It supports shared memory homogeneous and heterogeneous multicores and manycores including non-coherent cache architectures.
- In addition to the automatic parallelization, automatic power control using DVFS and Clock and Power gating has been implemented for real-time processing and minimum execution time processing modes.
- The following performance has been attained on various multicores and servers:
 - 55 times speedup by 64 processor cores for Carbon Ion Radiotherapy Cancer treatment on IBM Power 7 64 core SMP (Hitachi SR16000)
 - 92 Times Speedup for GMS Earthquake Wave Propagation Simulation on 128 processor cores SMP (Hitachi SR16000)
 - Faster or Equal Processing Performance with Hardware Coherence Control on 8 core RP2 Multicore Precessor Having Hardware Coherent Mechanism Up-to 4 cores by OSCAR Compiler's Software Coherence Control
 - 33 Times Speedup for Optical Flow on 8 SH4A and 4 DRP accelerators on RP-X heterogeneous multicore.
 - Power Reduction of MPEG2 Decoding to 1/4 on 8 Core Homogeneous Multicore RP-2.
 - 1.95 times speedup on Renesas V850 2 core embedded multicore for automobile engine control program generated by MATLAB/SIMLINK embedded coder.
 - 2.9 Times Speed-up for AAC Encodeing on 3 Core NaviEngine (ARM MPcore) with Realtime OS eT-Kernel Multi-Core Edition