OSCAR Compiler for Automatic Parallelization and Power Reduction for Multicores and Manycores

Hironori Kasahara
Professor, Dept. of Computer Science & Engineering
Director, Advanced Multicore Processor Research Institute
Waseda University, Tokyo, Japan
IEEE Computer Society Board of Governors
URL: http://www.kasahara.cs.waseda.ac.jp/

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, etc

＜Ripple Effect＞

- **Low CO_2 (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)

Cancer Treatment
Carbon Ion
Radiotherapy

Generated Solar Power and Sever Power Consumption on April 2, 2012.4.2（Clear) in Green Computing Center

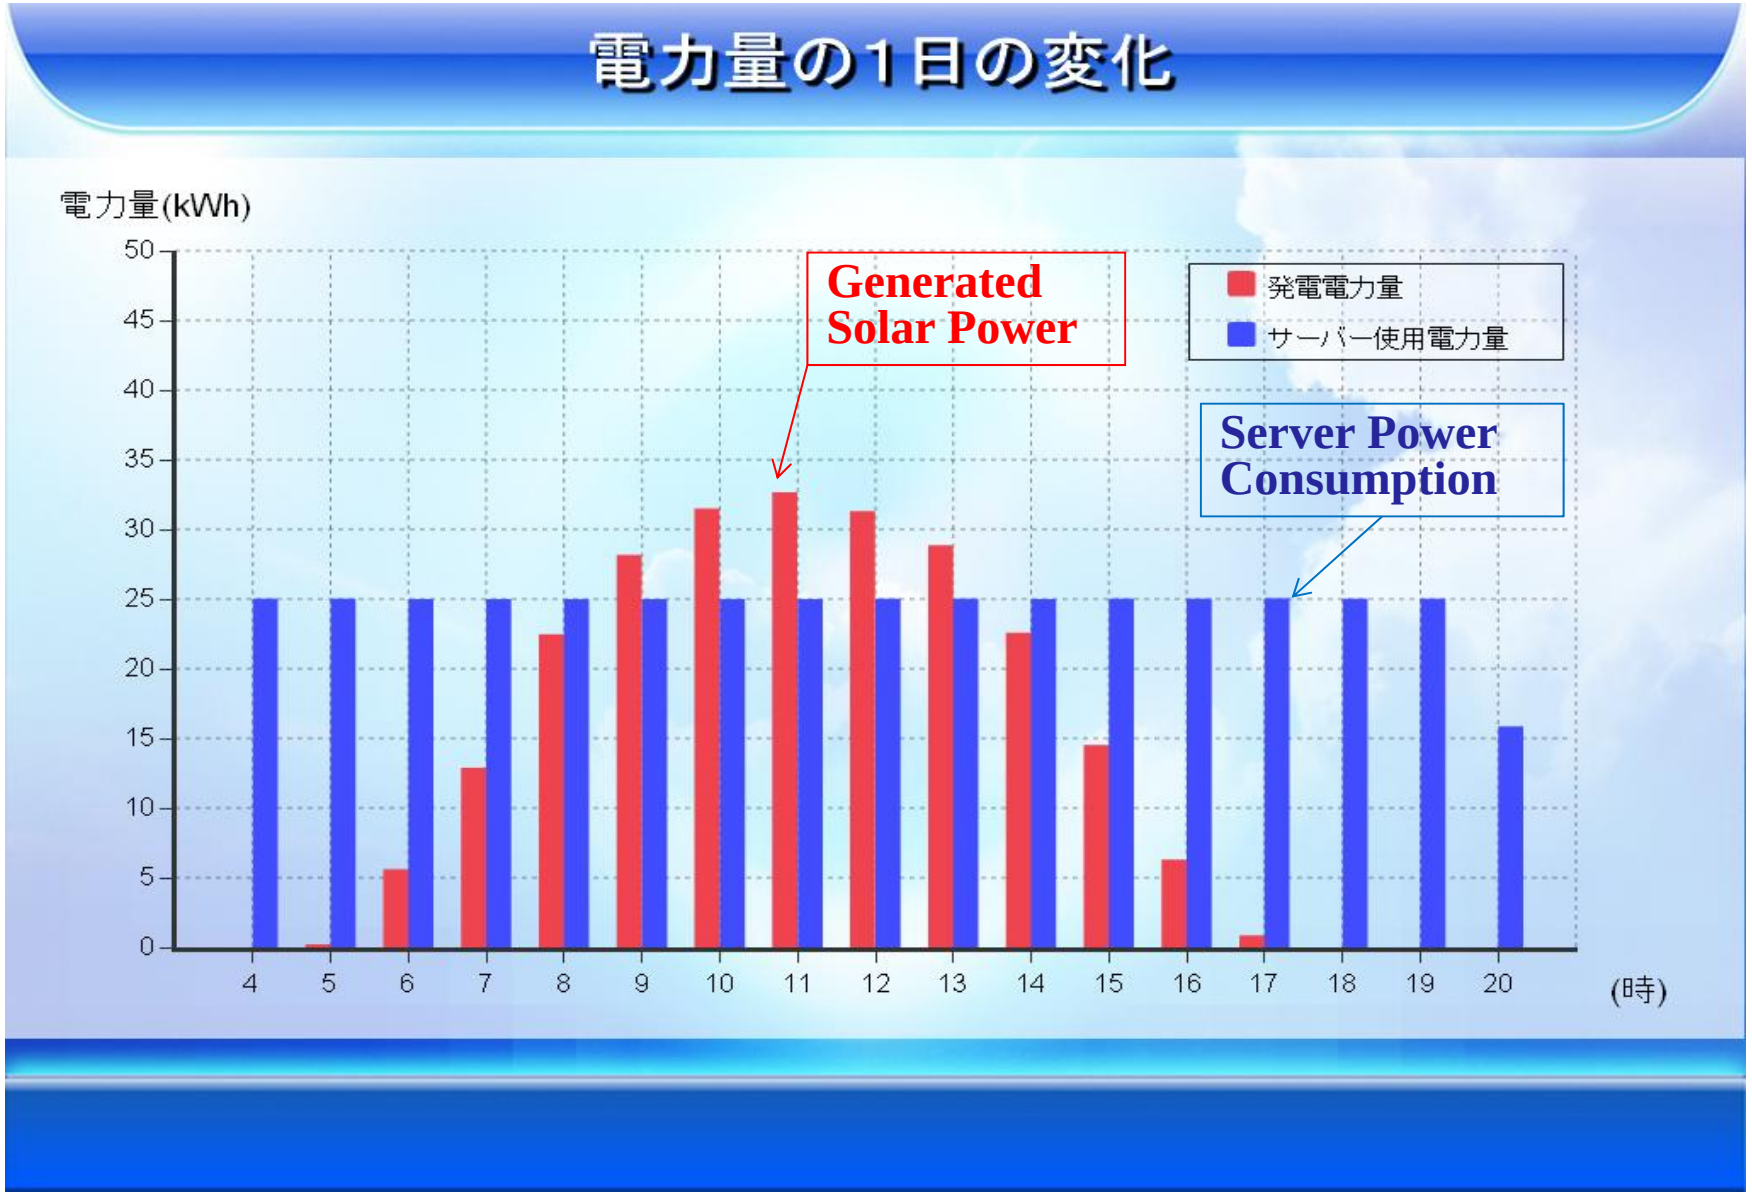

Super Low Power Web Server Using Embedded Multicore Processor RPX

1W with 8 SH4A processor cores

Power consumption

0.96 W

by RPX embedded multicore Web-server.

[Japanese | English]

Private Page

WASEDA UNIVERSITY Computer Science and Engineering

Kasahara Laboratory

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 ParallelComputing) Sep. 11-13, 2012 was upped.

Multi/Many-core Everywhere

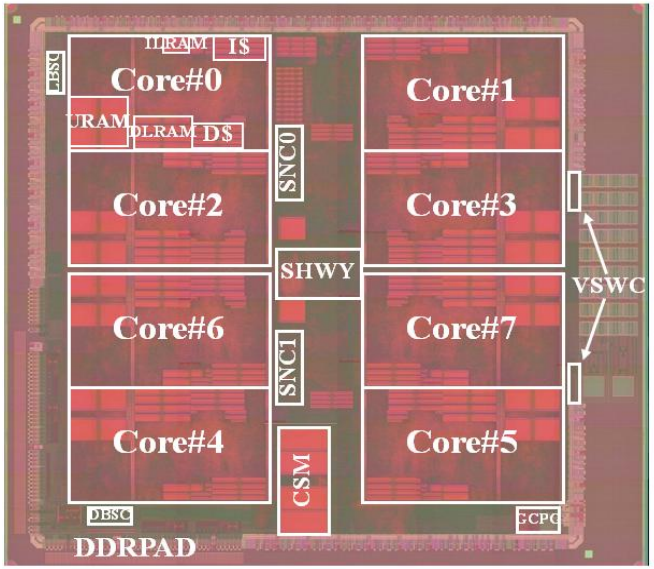

OSCAR Type Multi-core Chip by Renesas in METI/NEDO Multicore for Real-time Consumer Electronics Project (Leader: Prof.Kasahara)

The 37th (Nov. 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%)

Multi-core from embedded to supercomputers

- **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,
  - Plurarity 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 Cooled20PFLOPS, 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 SPARC64VIIIfx, 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

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.

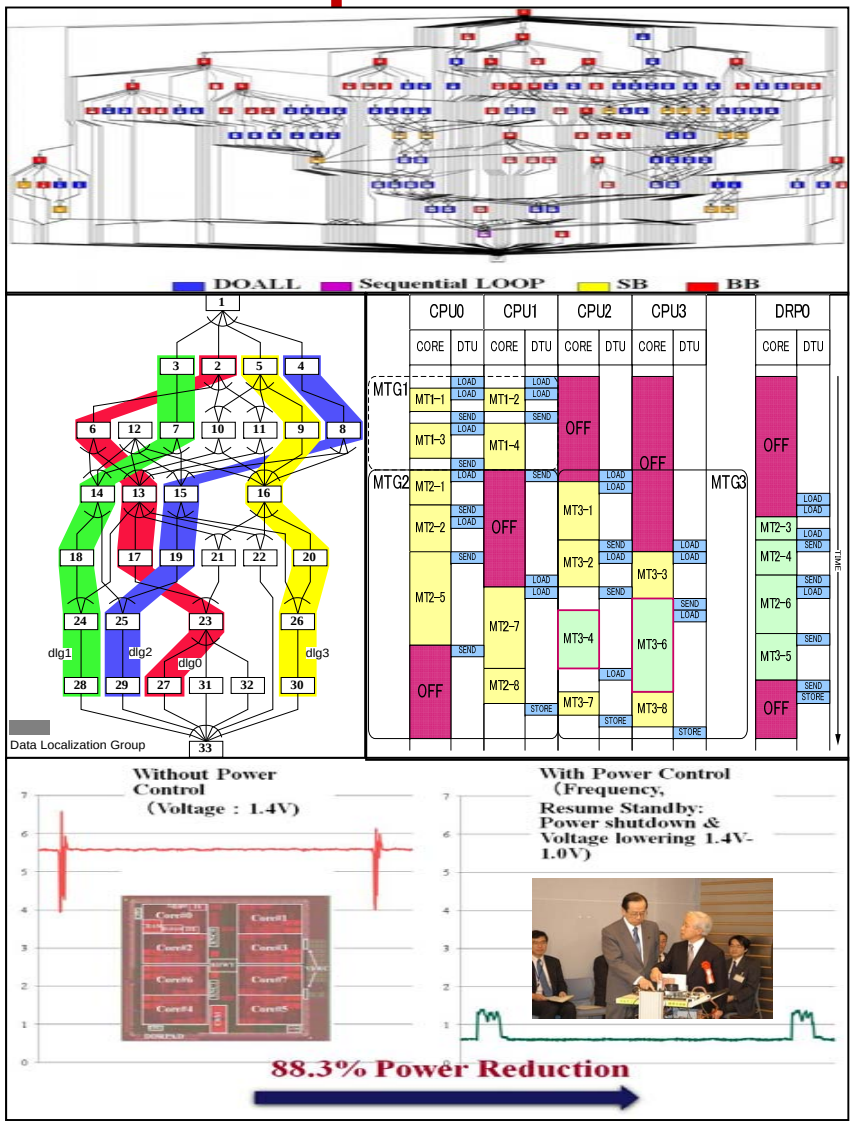

Generation of coarse grain tasks

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

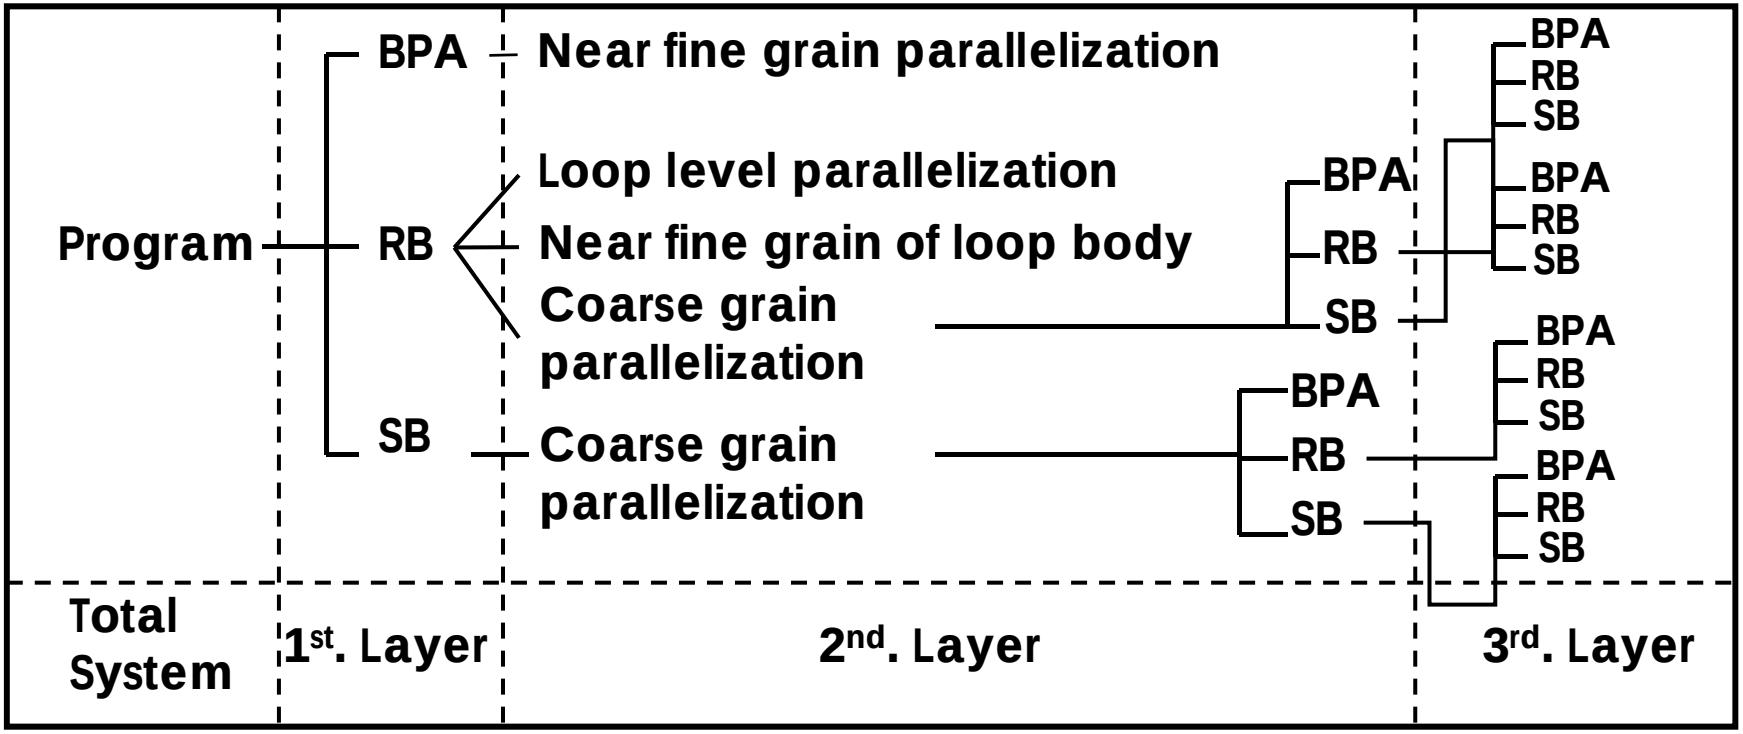

Earliest Executable Condition Analysis for coarse grain tasks (Macro-tasks)

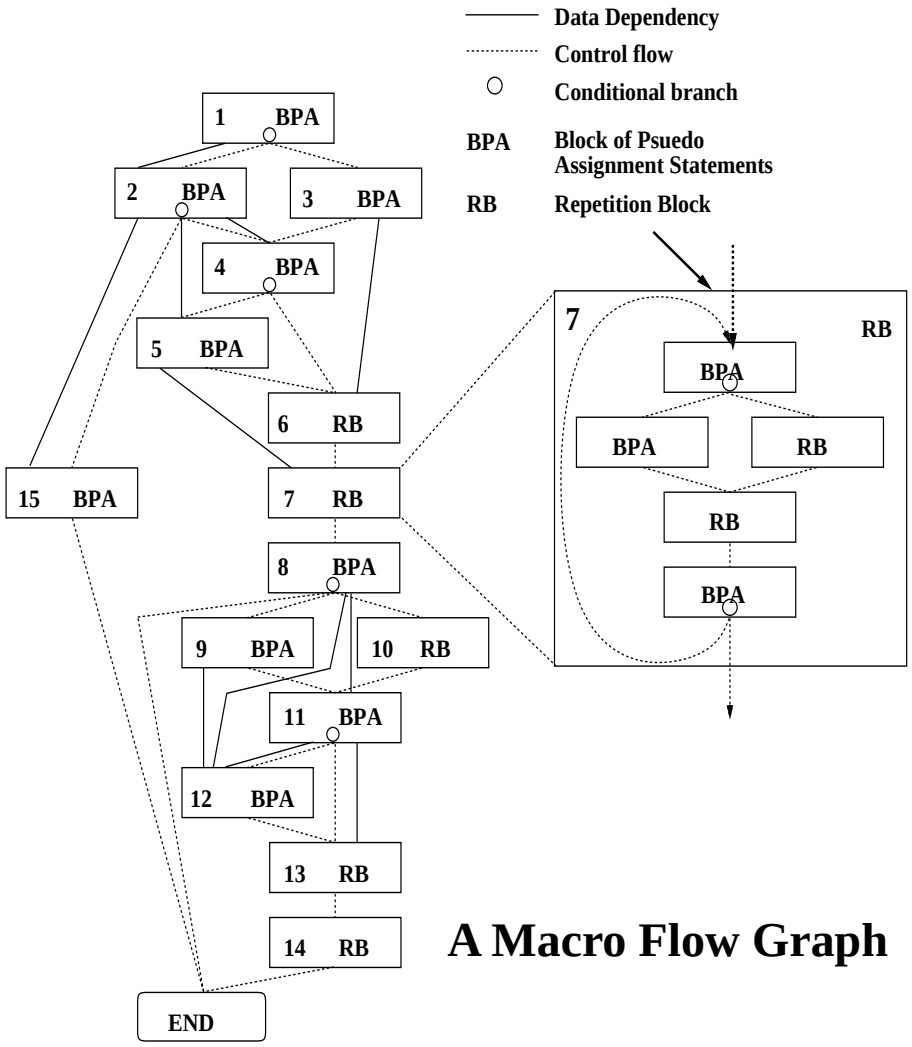

A Macro Flow Graph

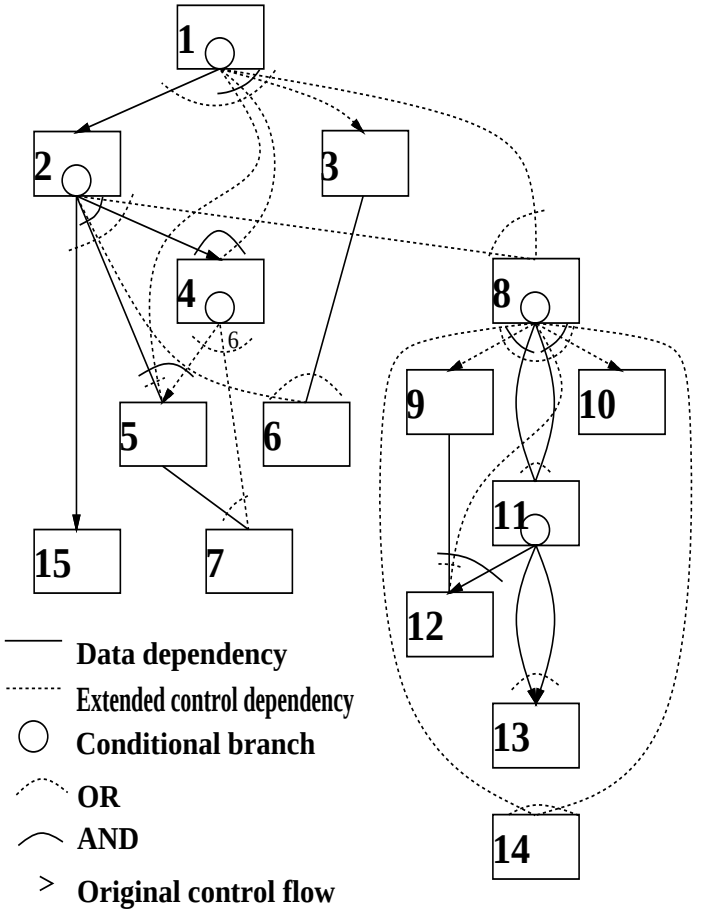

A Macro Task Graph

Automatic processor assignment in su2cor

- **Using 14 processors**
  - **Coarse grain parallelization within DO400 of subroutine**

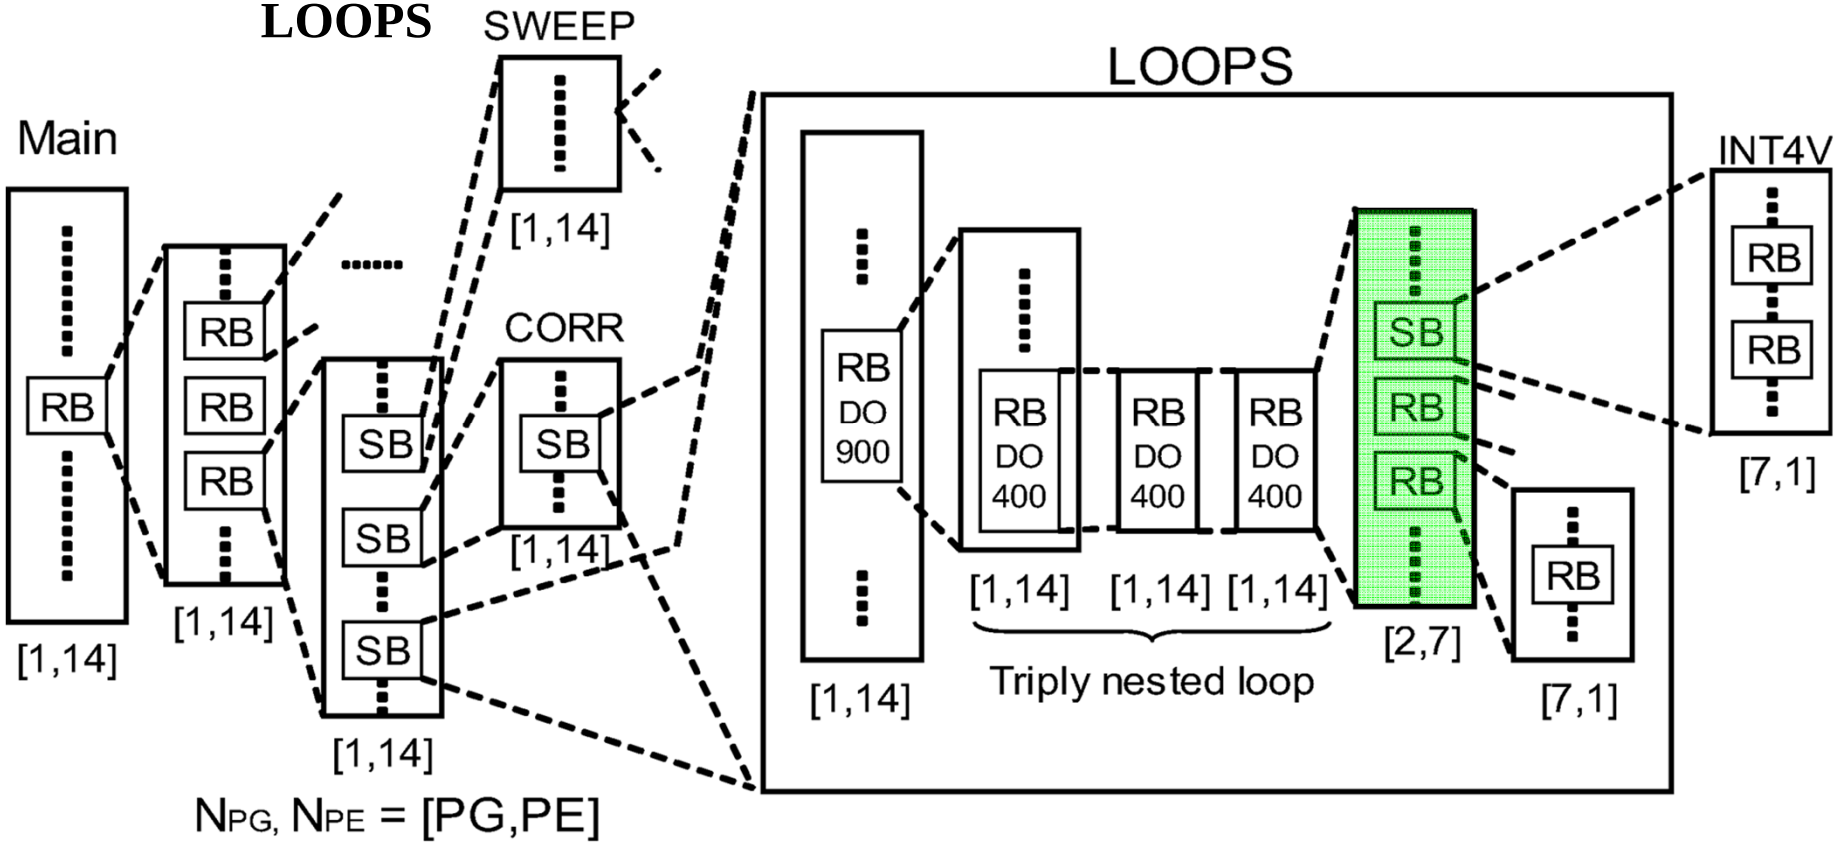

MTG of Su2cor-LOOPS-DO400

- **Coarse grain parallelism PARA_ALD = 4.3**

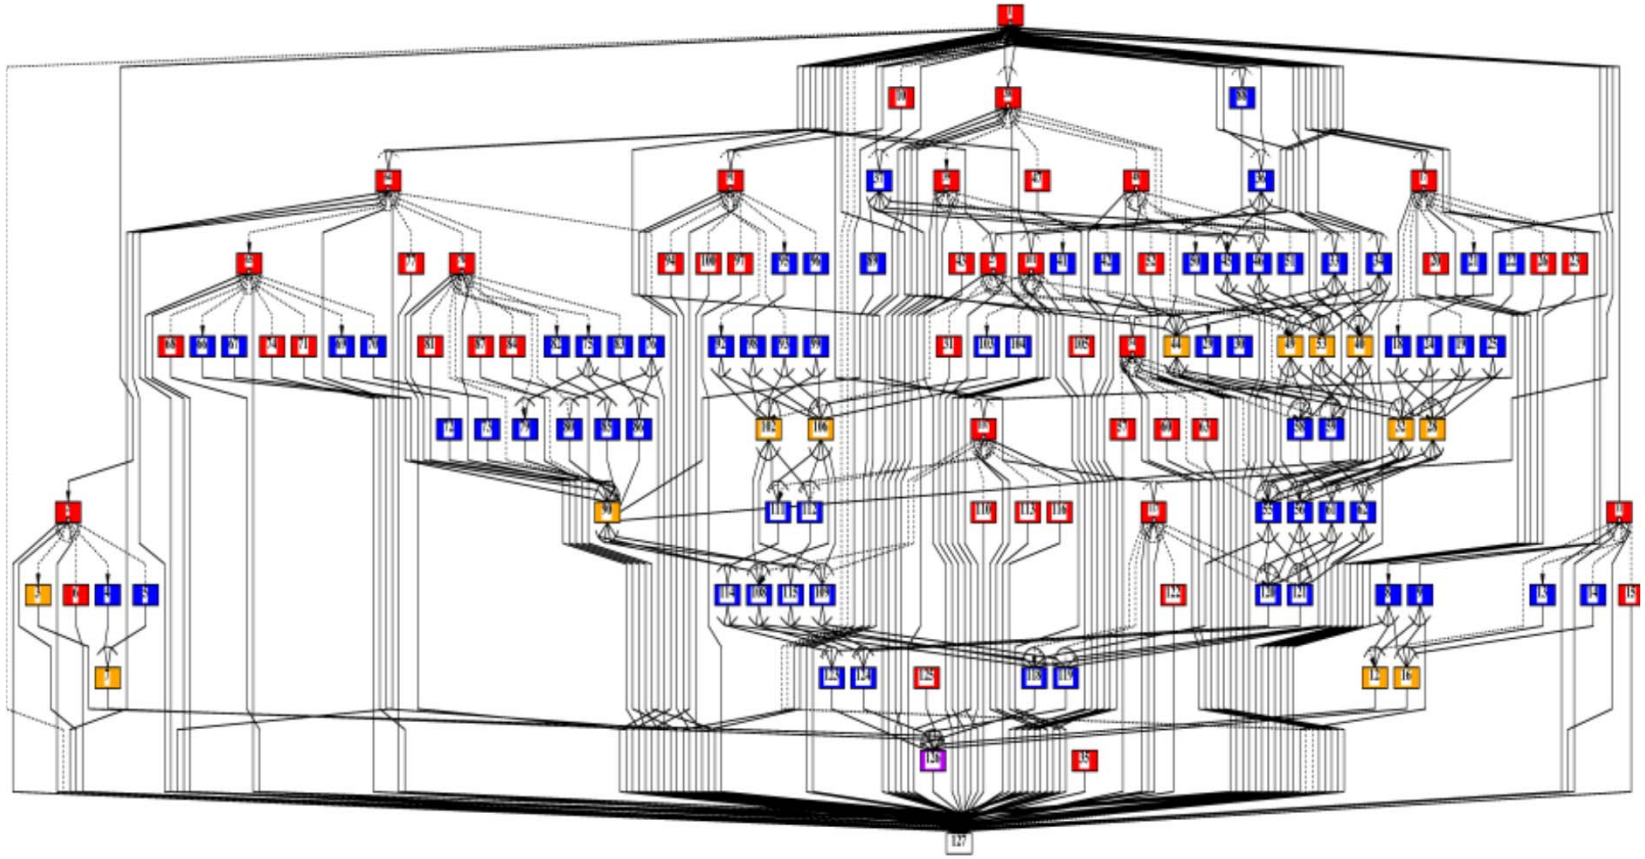

DOALL **Sequential LOOP** **SB** **BB**

Data-Localization
Loop Aligned Decomposition

- **Decompose multiple loop (Doall and Seq) into CARs and LRs considering inter-loop data dependence.**
  - **Most data in LR can be passed through LM.**
  - **LR: Localizable Region, CAR: Commonly Accessed Region**

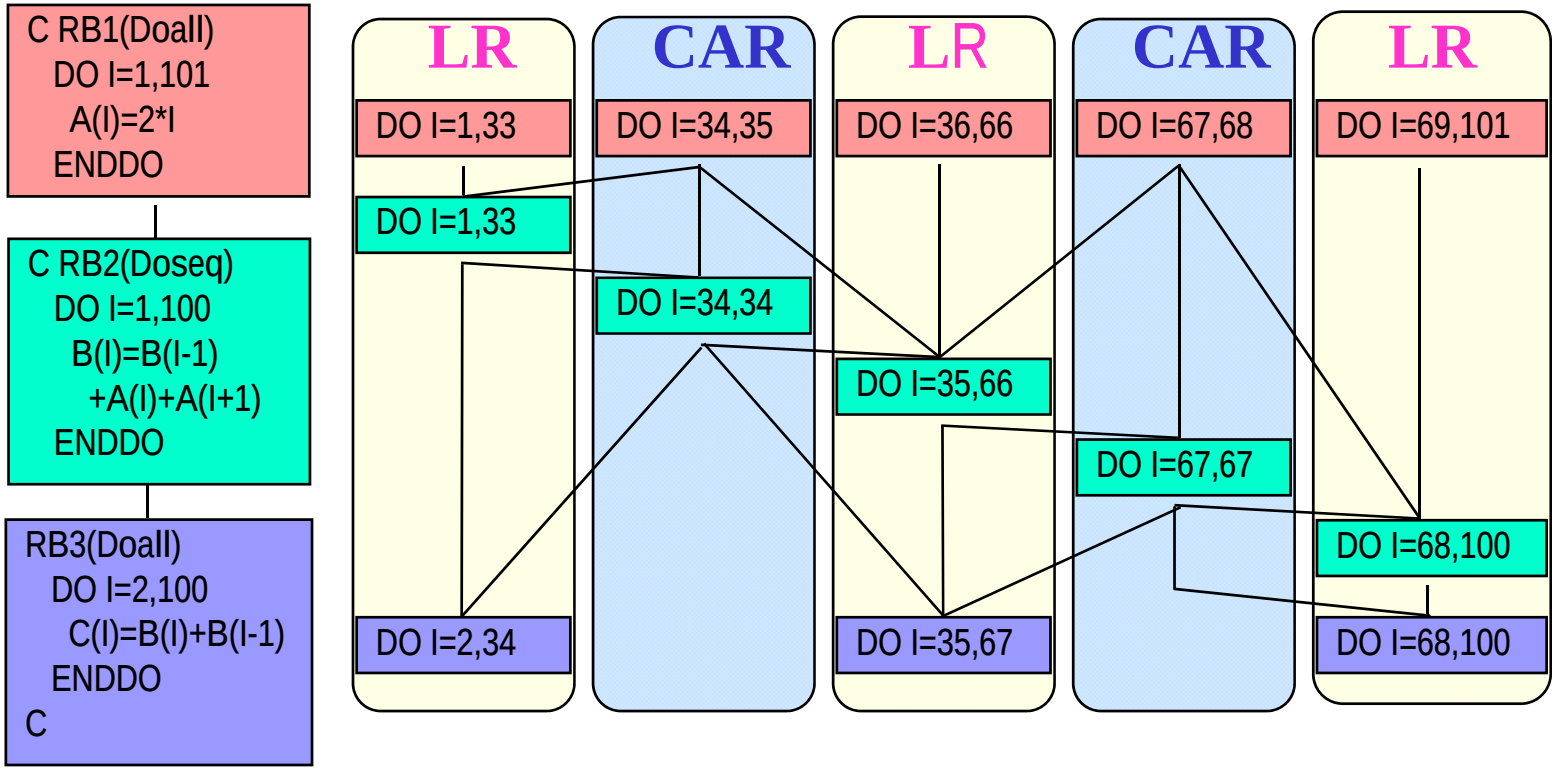

Data Localization

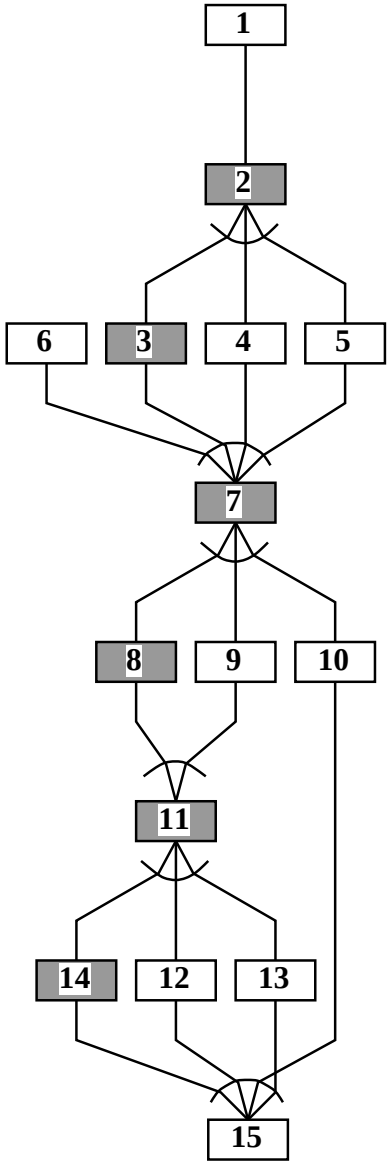

MTG

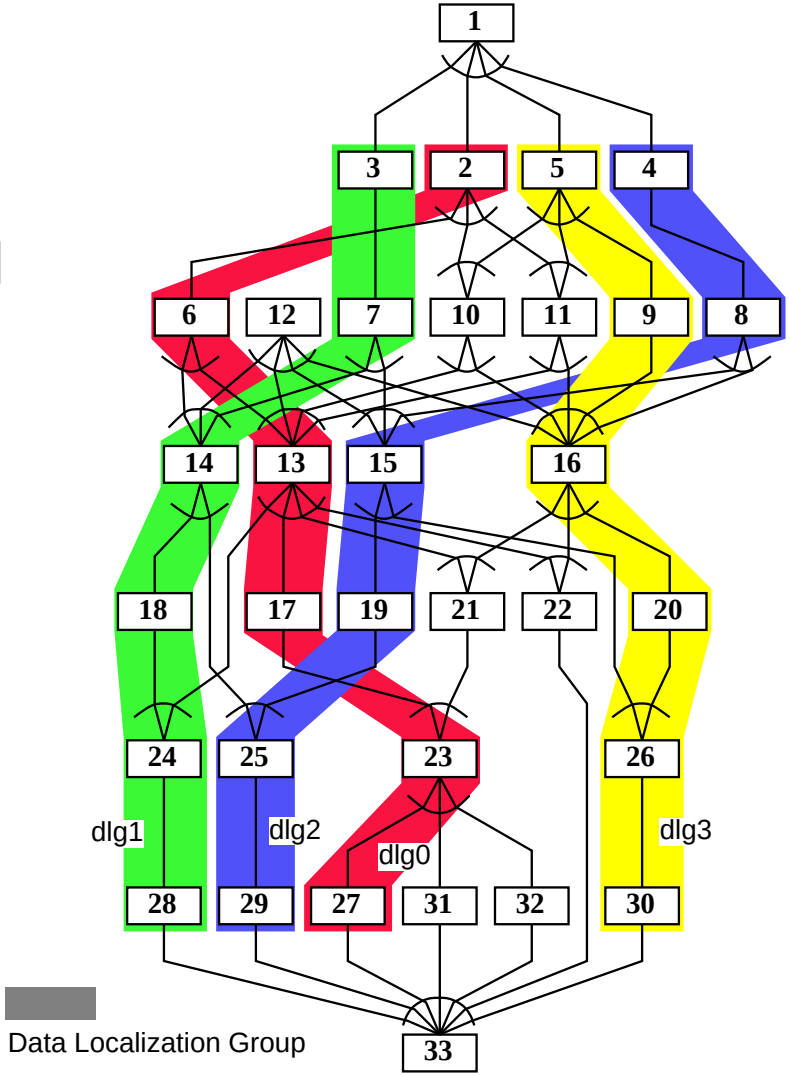

MTG after Division

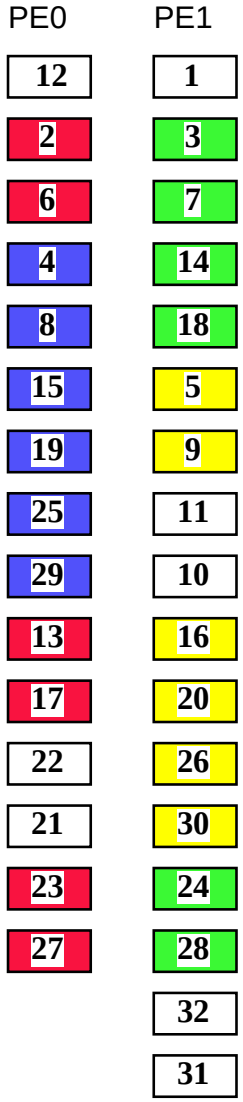

A schedule for two processors

Data Layout for Removing Line Conflict Misses by Array Dimension Padding

Declaration part of arrays in spec95 swim

before padding

```
PARAMETER (N1=513, N2=513)

COMMON  U(N1,N2), V(N1,N2), P(N1,N2),
*       UNEW(N1,N2), VNEW(N1,N2),
1       PNEW(N1,N2), UOLD(N1,N2),
*       VOLD(N1,N2), POLD(N1,N2),
2       CU(N1,N2), CV(N1,N2),
*       Z(N1,N2), H(N1,N2)
```

after padding

```
PARAMETER (N1=513, N2=544)

COMMON  U(N1,N2), V(N1,N2), P(N1,N2),
*       UNEW(N1,N2), VNEW(N1,N2),
1       PNEW(N1,N2), UOLD(N1,N2),
*       VOLD(N1,N2), POLD(N1,N2),
2       CU(N1,N2), CV(N1,N2),
*       Z(N1,N2), H(N1,N2)
```

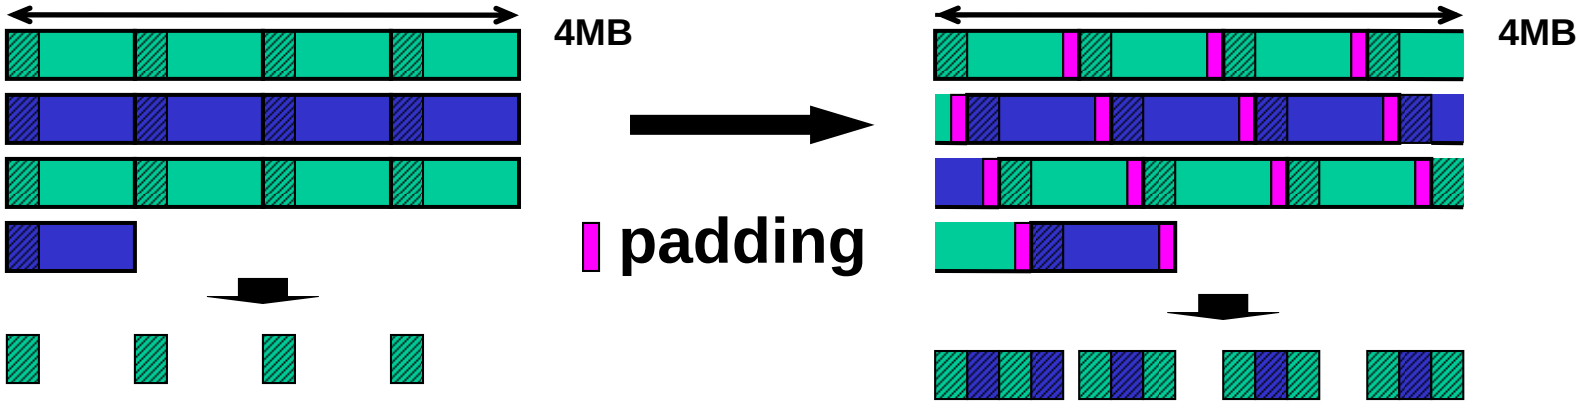

Box: Access range of DLG0

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

- Shortest execution time mode

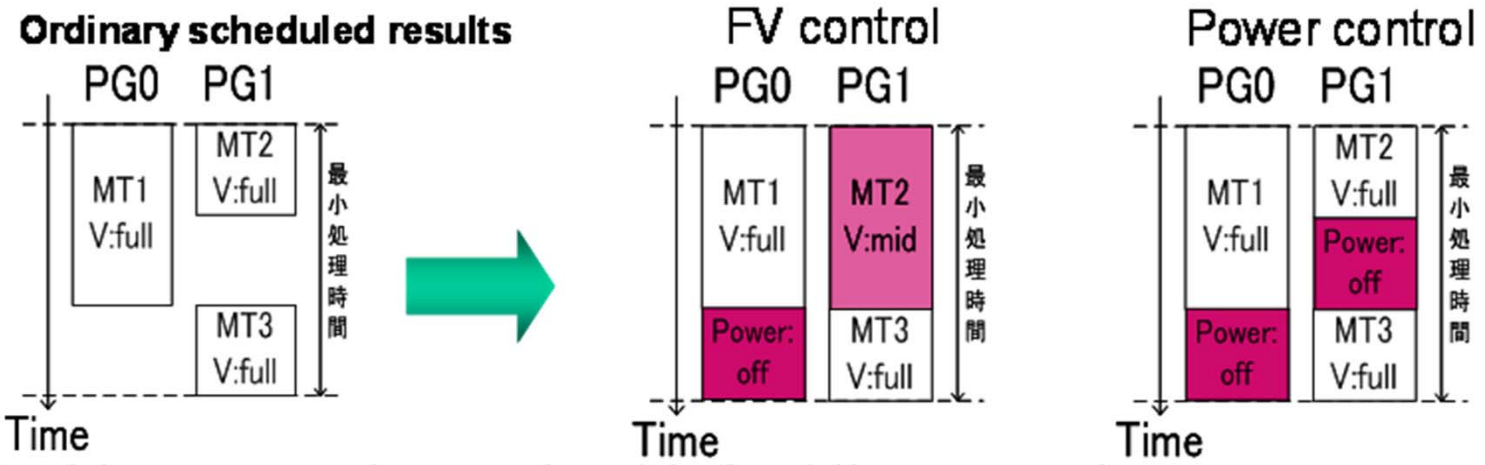

- Realtime processing mode with dead line constraints

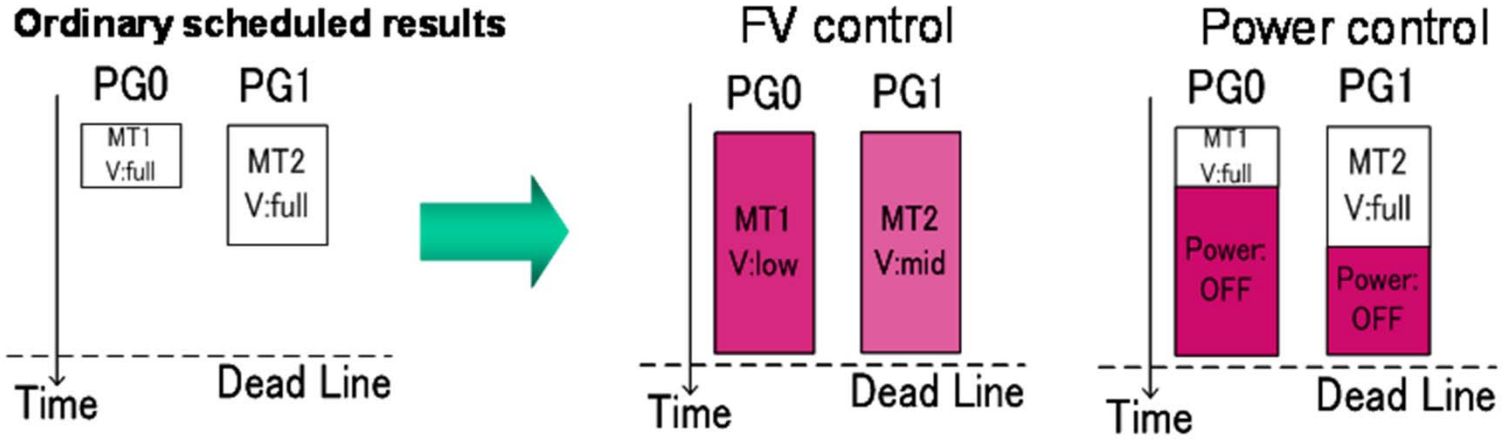

An Example of Machine Parameters for the Power Saving Scheme

- **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
FULL	0	40k	40k	80k
MID	40k	0	40k	80k
LOW	40k	40k	0	80k
OFF	80k	80k	80k	0

delay time [u.t.]

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

energy overhead [μJ]

Power Reduction Scheduling

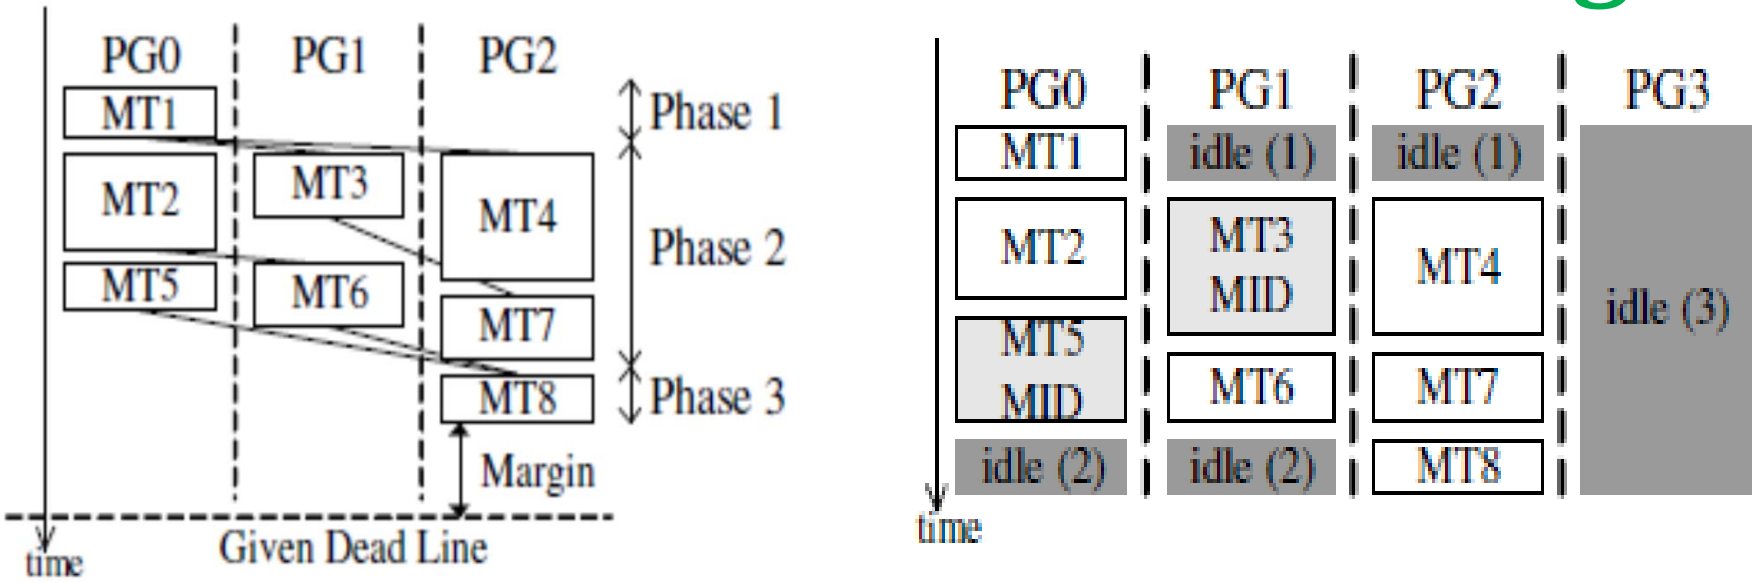

Fig. 6. V/F control of applu(4proc.)

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

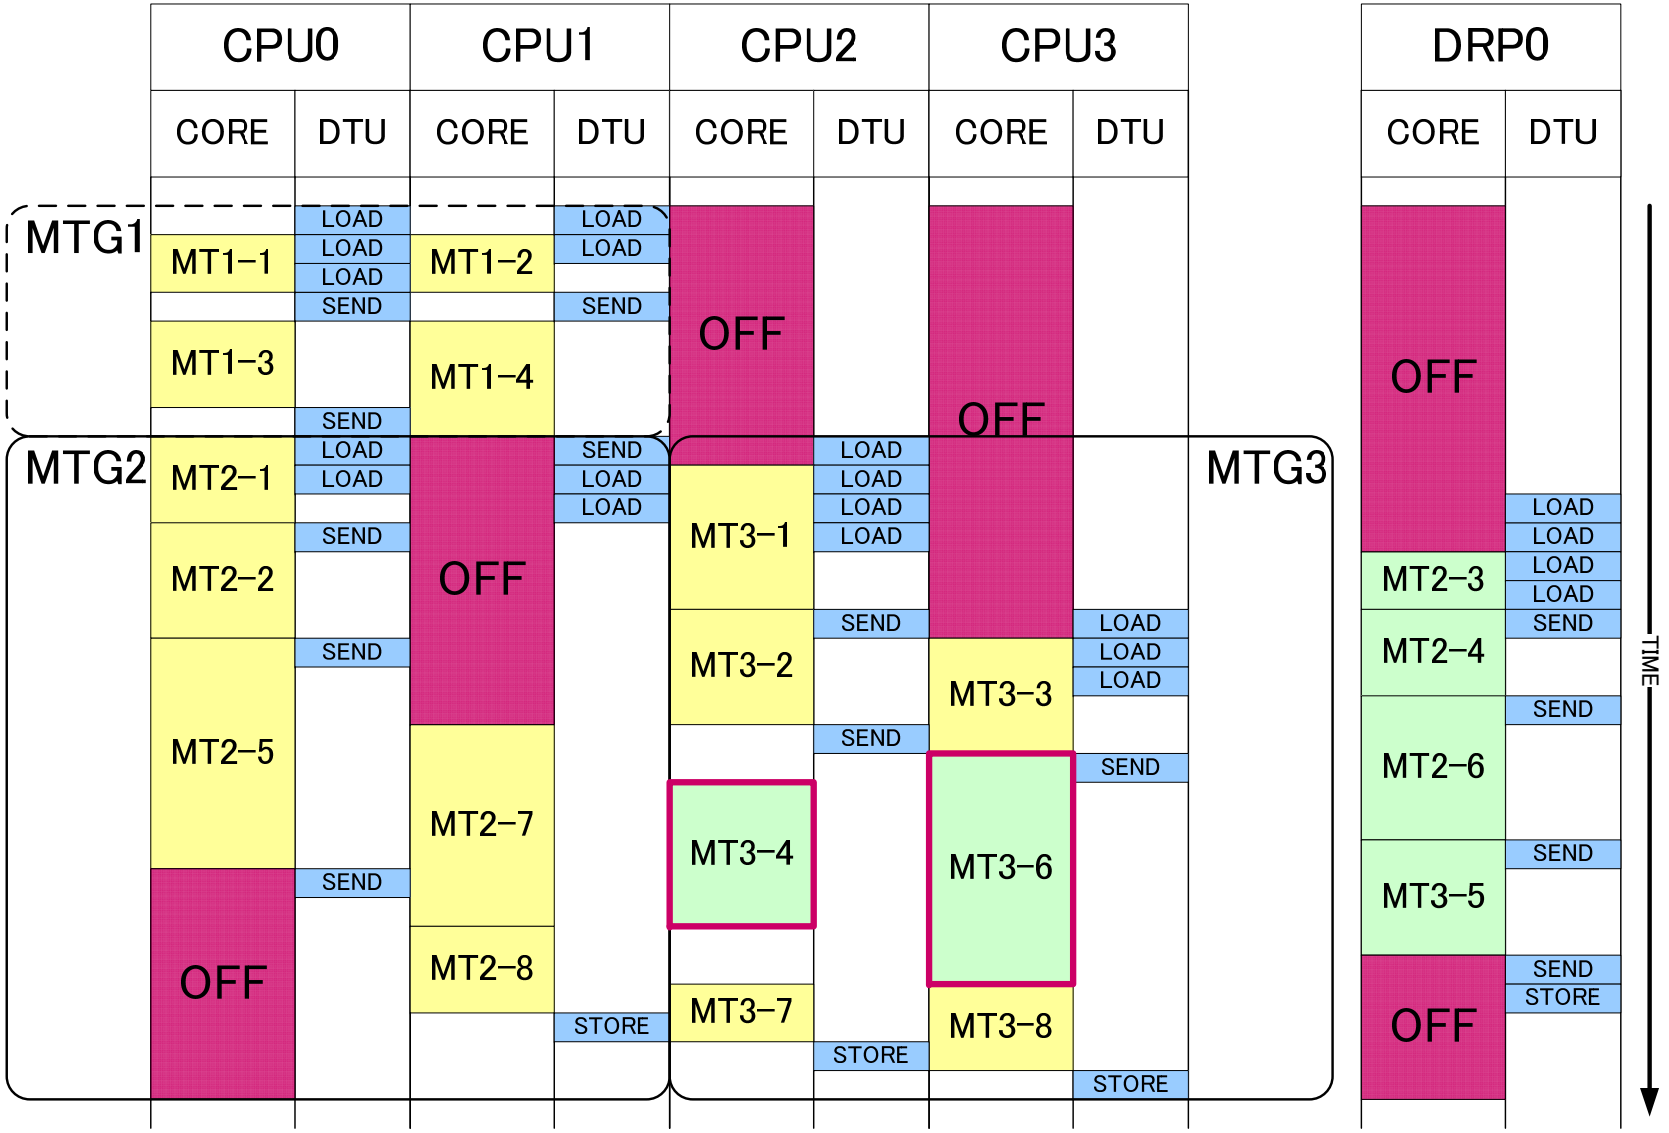

Generated Multigrain Parallelized Code

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

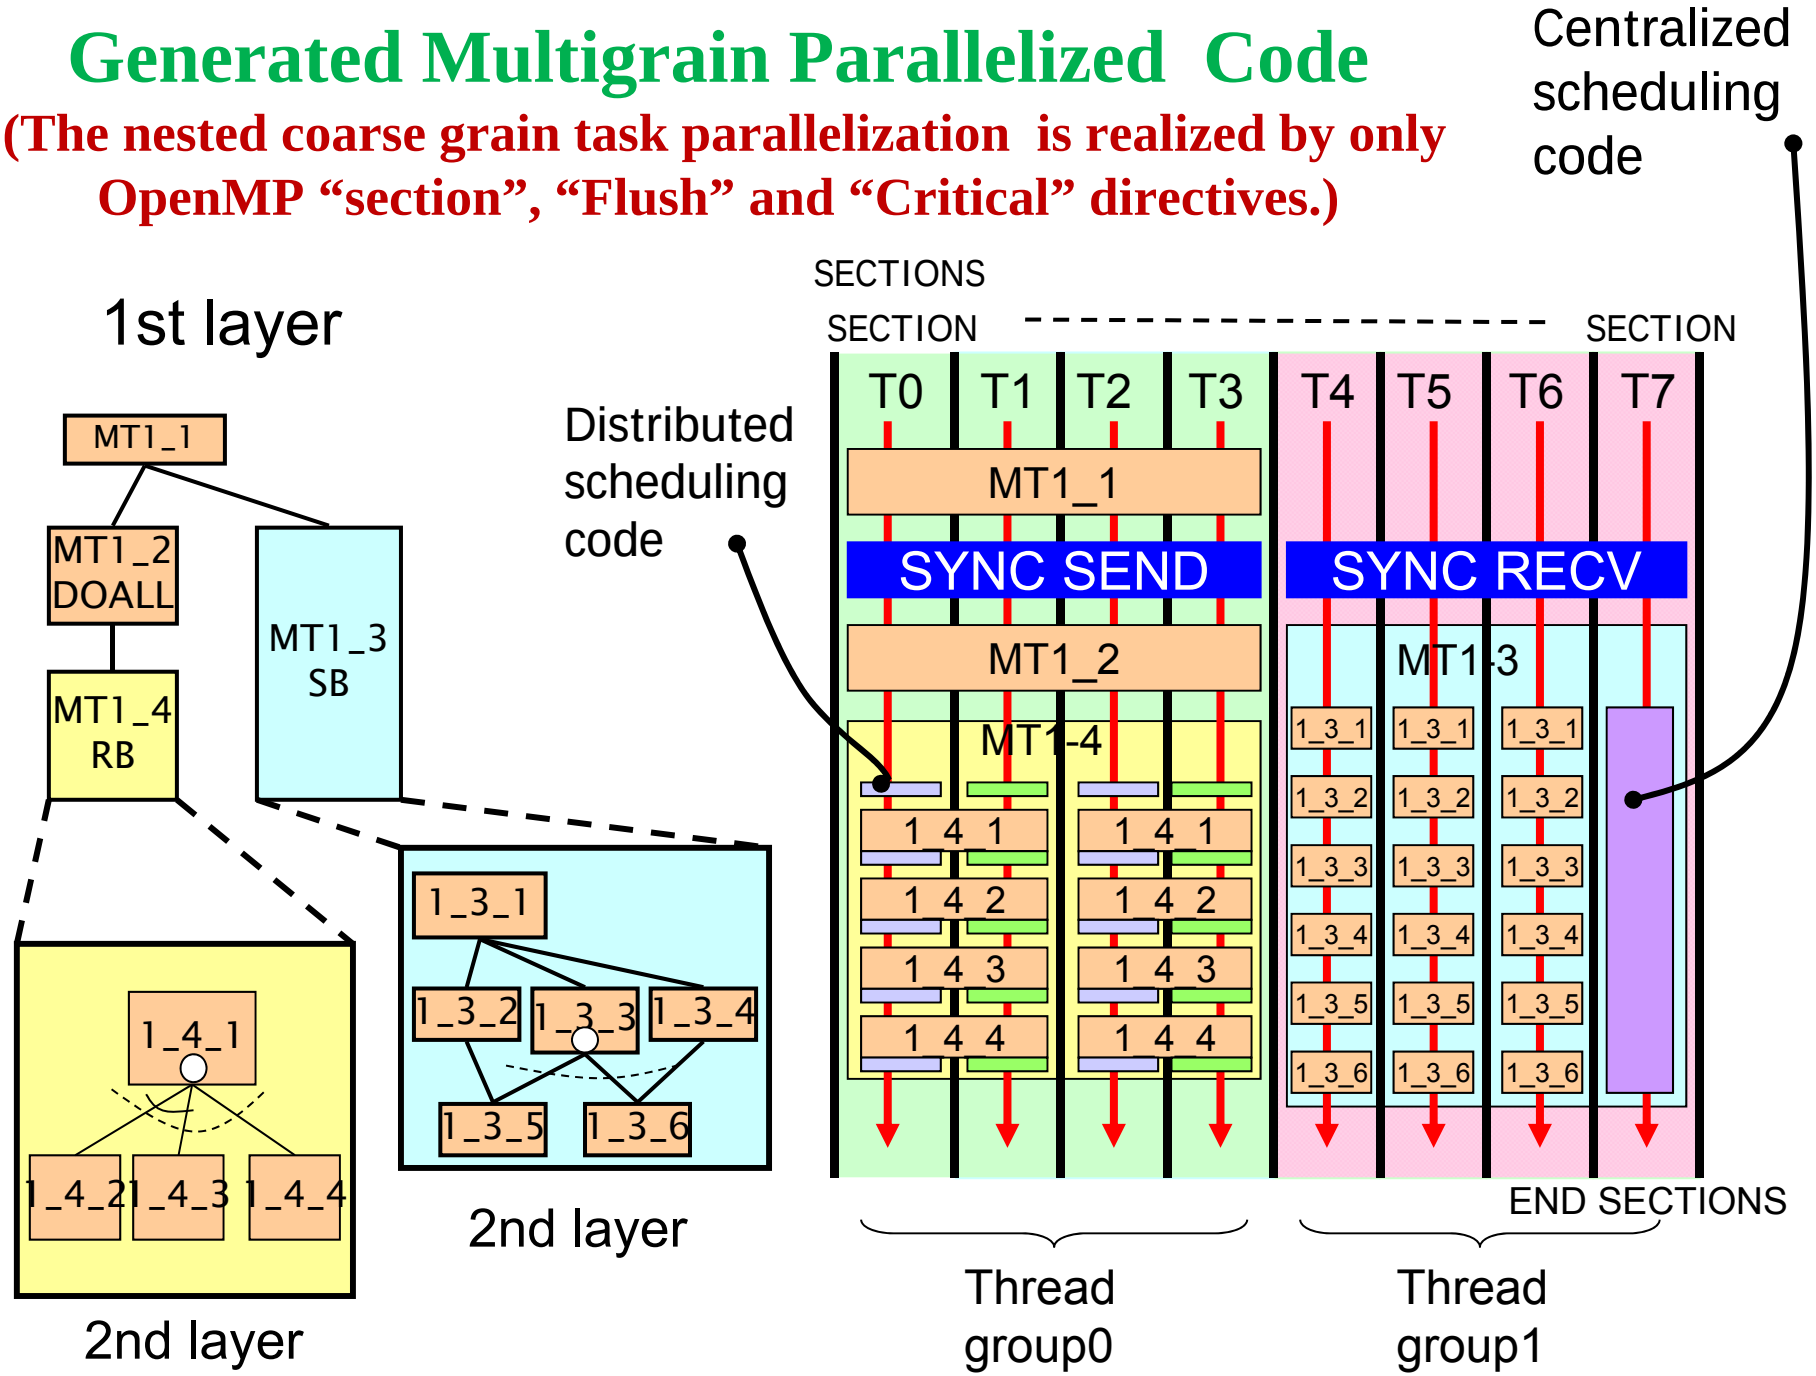

Multicore Program Development Using OSCAR API V2.0

Sequential Application Program in Fortran or C
(Consumer Electronics, Automobiles, Medical, Scientific computation, etc.)

Homogeneous

Hetero

Manual parallelization / power reduction

Accelerator Compiler/ User
Add "hint" directives before a loop or a function to specify it is executable by the accelerator with how many clocks

Waseda OSCAR Parallelizing Compiler
- Coarse grain task parallelization
- Data Localization
- DMAC data transfer
- Power reduction using DVFS, Clock/ Power gating

OSCAR API for Homogeneous and/or Heterogeneous Multicores and manycores
Directives for thread generation, memory, data transfer using DMA, power managements

Parallelized API F or C program

Proc0
Code with directives
Thread 0

Proc1
Code with directives
Thread 1

Accelerator 1
Code

Accelerator 2
Code

Generation of parallel machine codes using sequential compilers

Low Power Homogeneous Multicore Code Generation
API Analyzer | Existing sequential compiler

Low Power Heterogeneous Multicore Code Generation
API Analyzer (Available from Waseda) | Existing sequential compiler

Server Code Generation
OpenMP Compiler

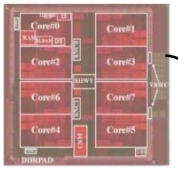

Homegeneous Multicore s from Vendor A (SMP servers)

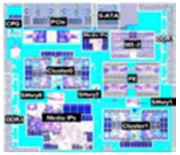

Heterogeneous Multicores from Vendor B

Shred memory servers

Executable on various multicores

Hitachi, Renesas, NEC, Fujitsu, Toshiba, Denso, Olympus, Mitsubishi, Esol, Cats, Gaio, 3 univ.

OSCAR: Optimally Scheduled Advanced Multiprocessor
API : Application Program Interface

Performance of OSCAR compiler on NEC NaviEngine(MPcore) with Linux in 2007

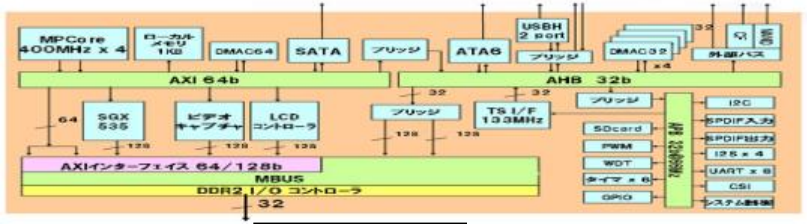

Fortran

Compile Opiion : -O3

C

NaviEngine上での
メディアアプリケーションによる
OSCARコンパイラ評価

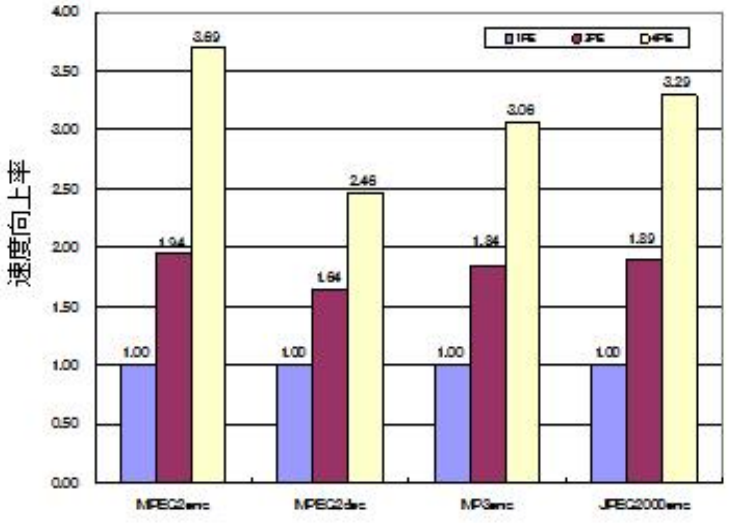

•4プロセッサで、平均3.13倍の速度向上率

- **OSCAR compiler gave us 3.43 times speedup against 1 core for Fortran and 3.13 for C on ARM/NEC MPCore with 4 ARM 400MHz cores**

2.9 Times Speed-up of AAC Encodeing Compared with a Sequential Processing on 3 Core NaviEngine (ARM MPcore) with Realtime OS eT-Kernel Multi-Core Edition

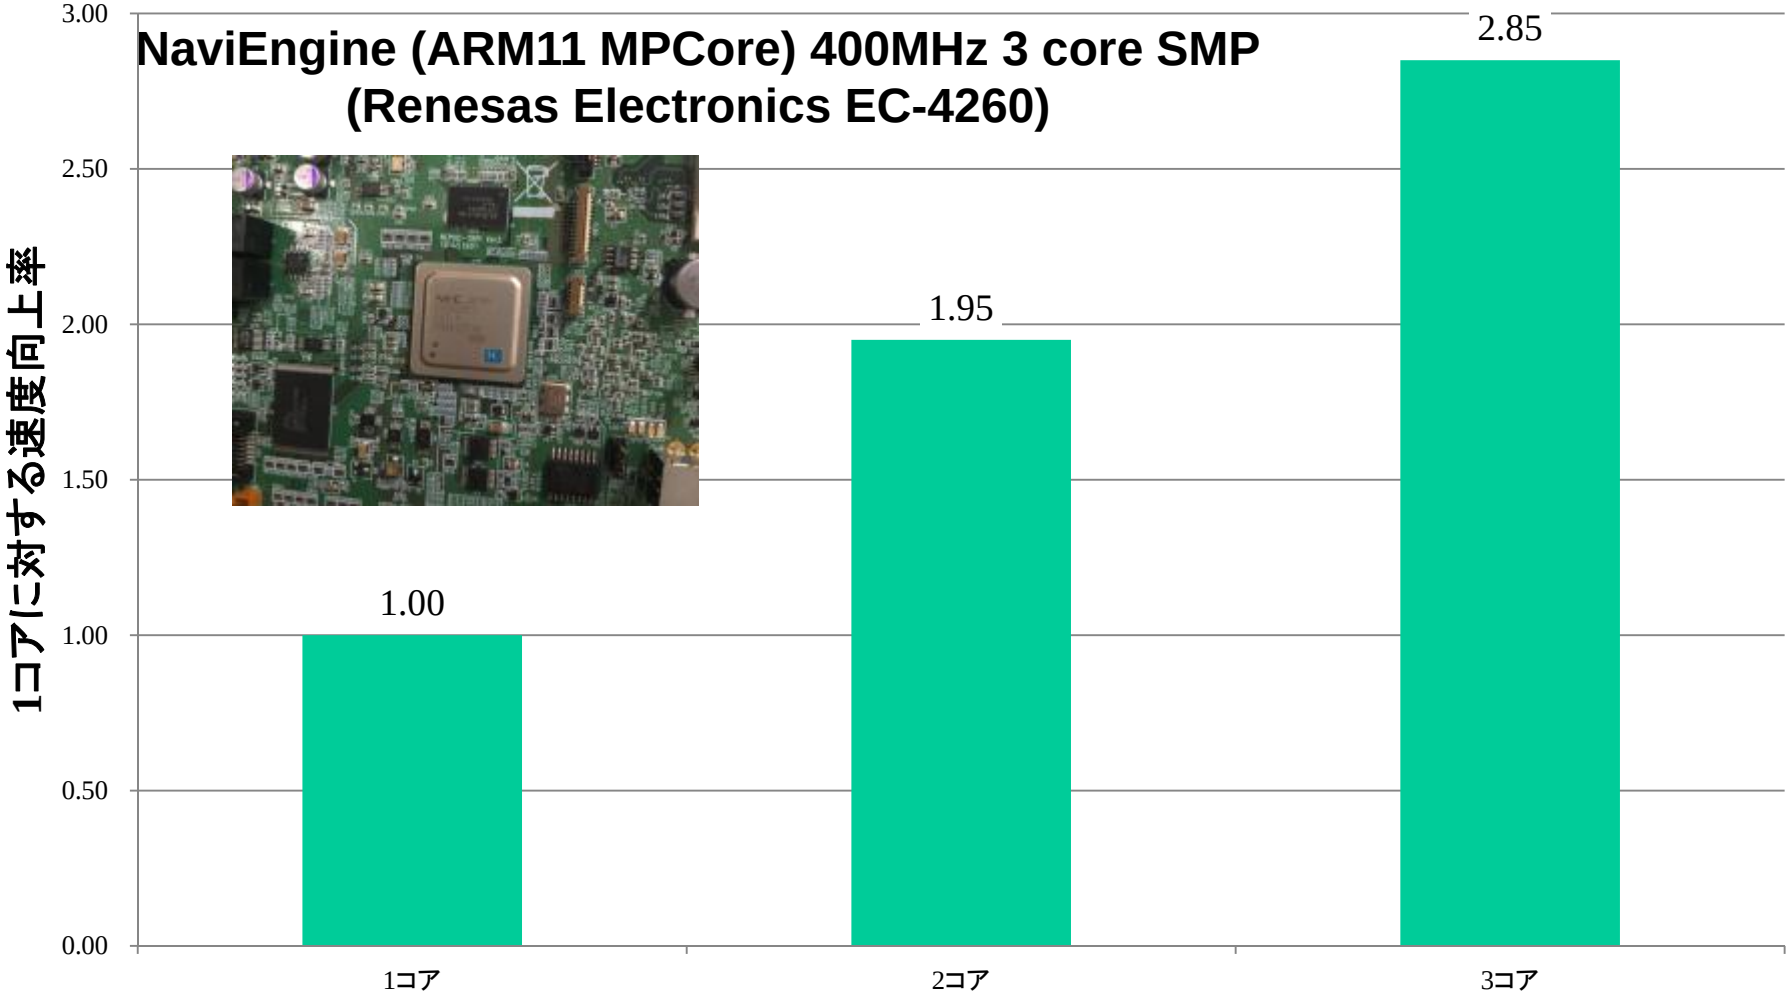

Performance of OSCAR Compiler & API on 2 ARMv7-cores Qualcomm MSM8960 Android 4.0 for Smart Phones

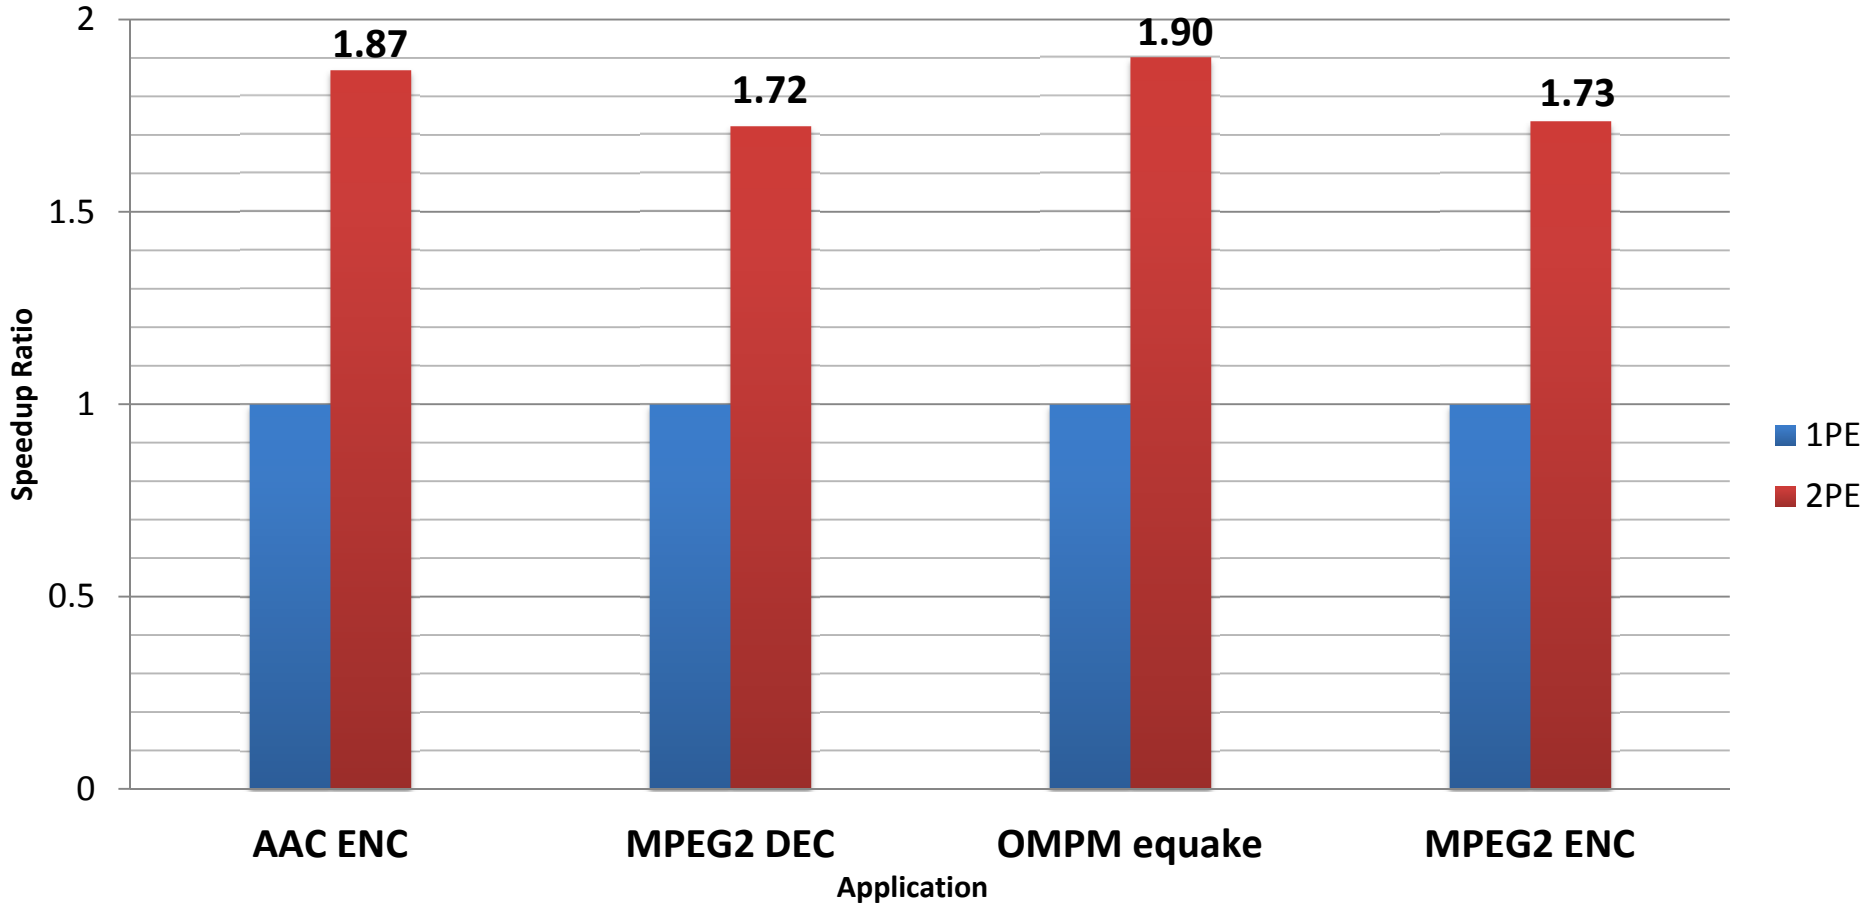

1.81 times speedup by 2 cores on the average against 1 core

Performance of OSCAR Compiler on Intel 12 core SMP based on 6-core Xeon X5670

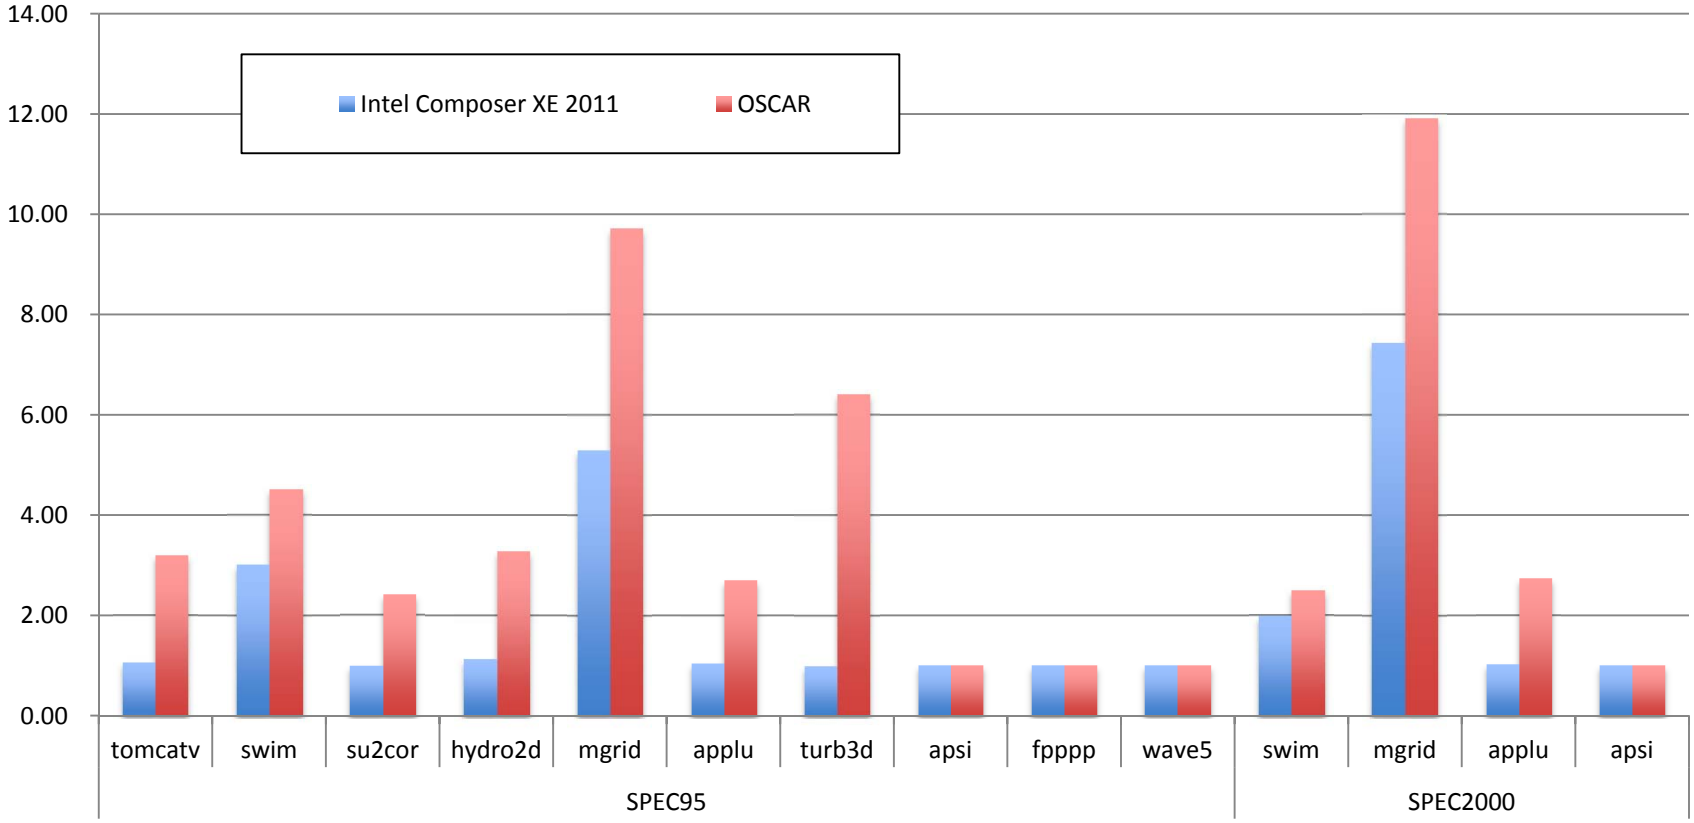

- OSCAR Compiler gives us 1.9 times speedup on the average against Intel Composer XE 2011

Performance of OSCAR Compiler on AMD 12-core SMP Based on Opteron 6174

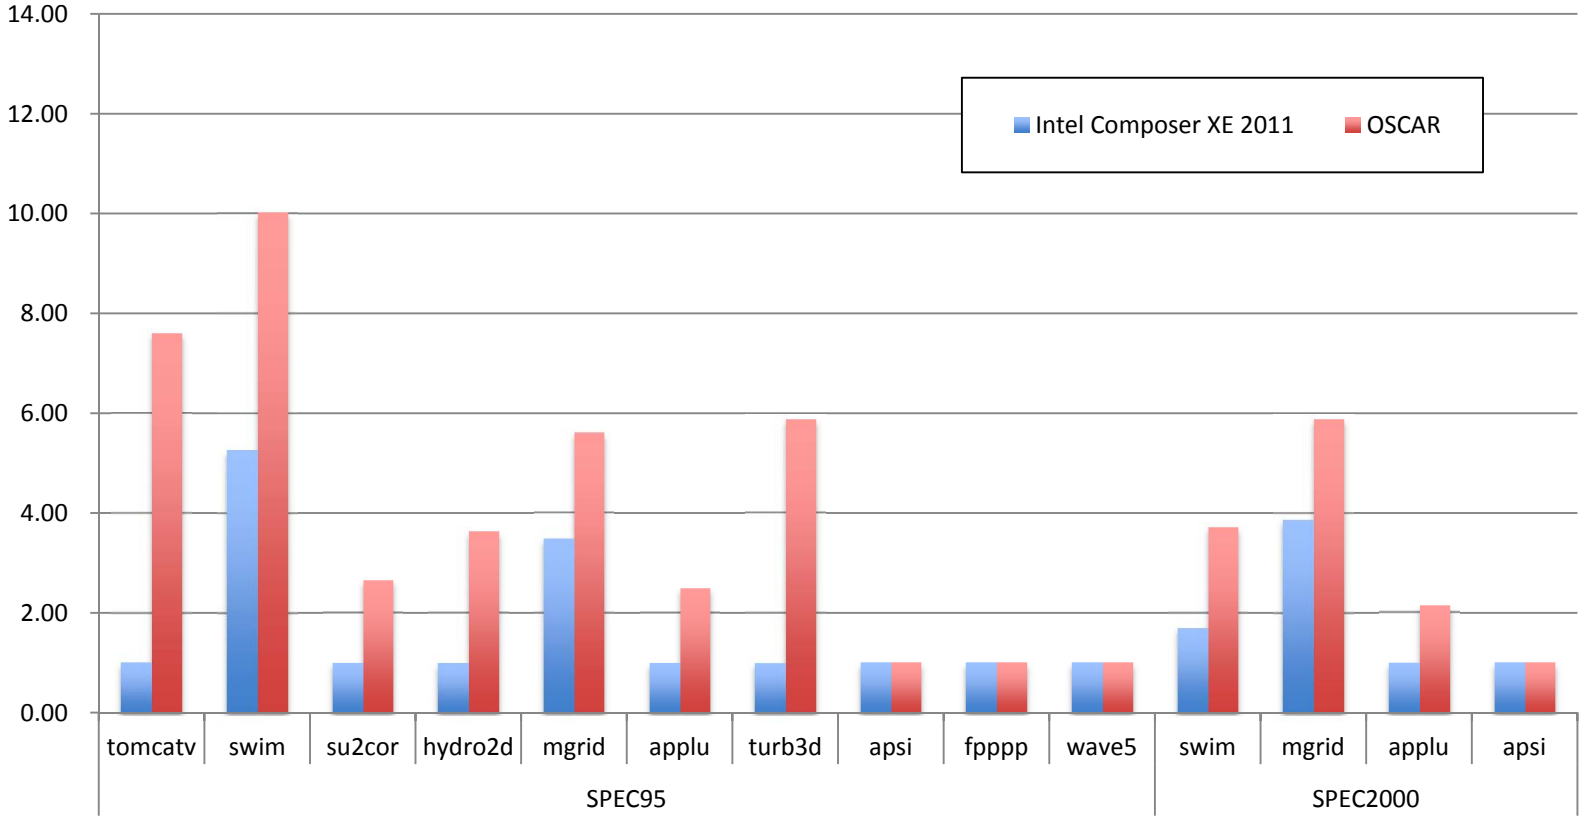

- OSCAR Compiler gives us 2.2 times speedup on the average against Intel Composer XE 2011

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

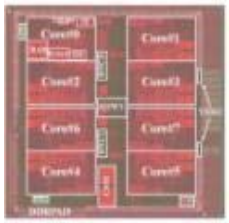

MPEG2 Decoding with 8 CPU cores

Without Power Control （Voltage：1.4V)

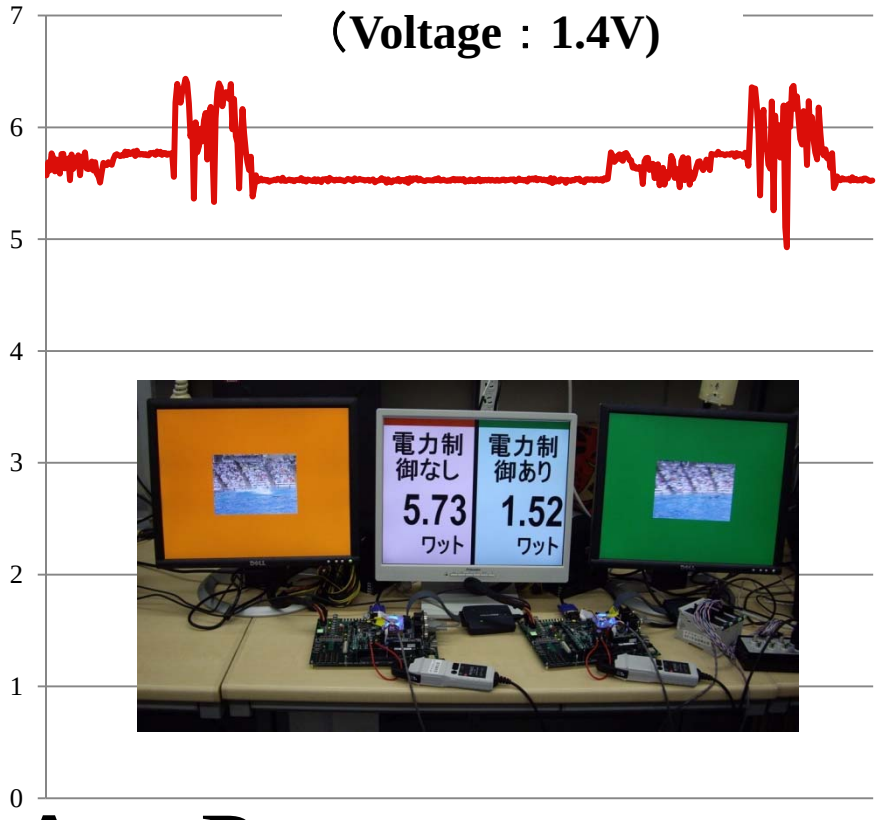

With Power Control （Frequency, Resume Standby: Power shutdown & Voltage lowering 1.4V-1.0V)

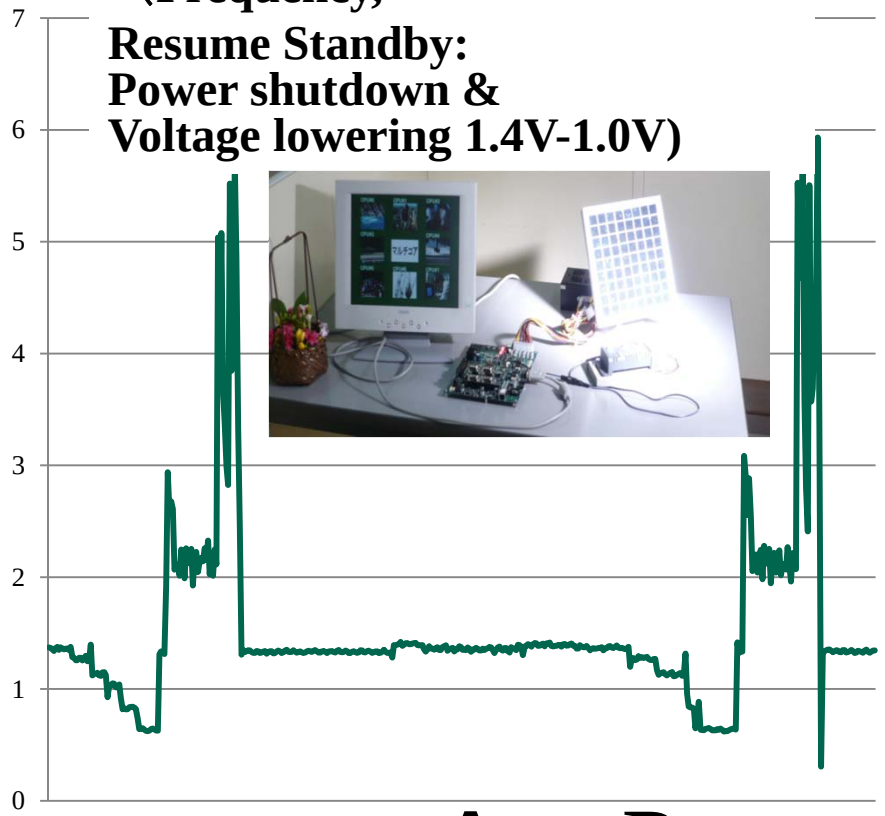

Avg. Power 5.73 [W]

73.5% Power Reduction

Avg. Power 1.52 [W]

33 Times Speedup Using OSCAR Compiler and OSCAR API on RP-X

(Optical Flow with a hand-tuned library)

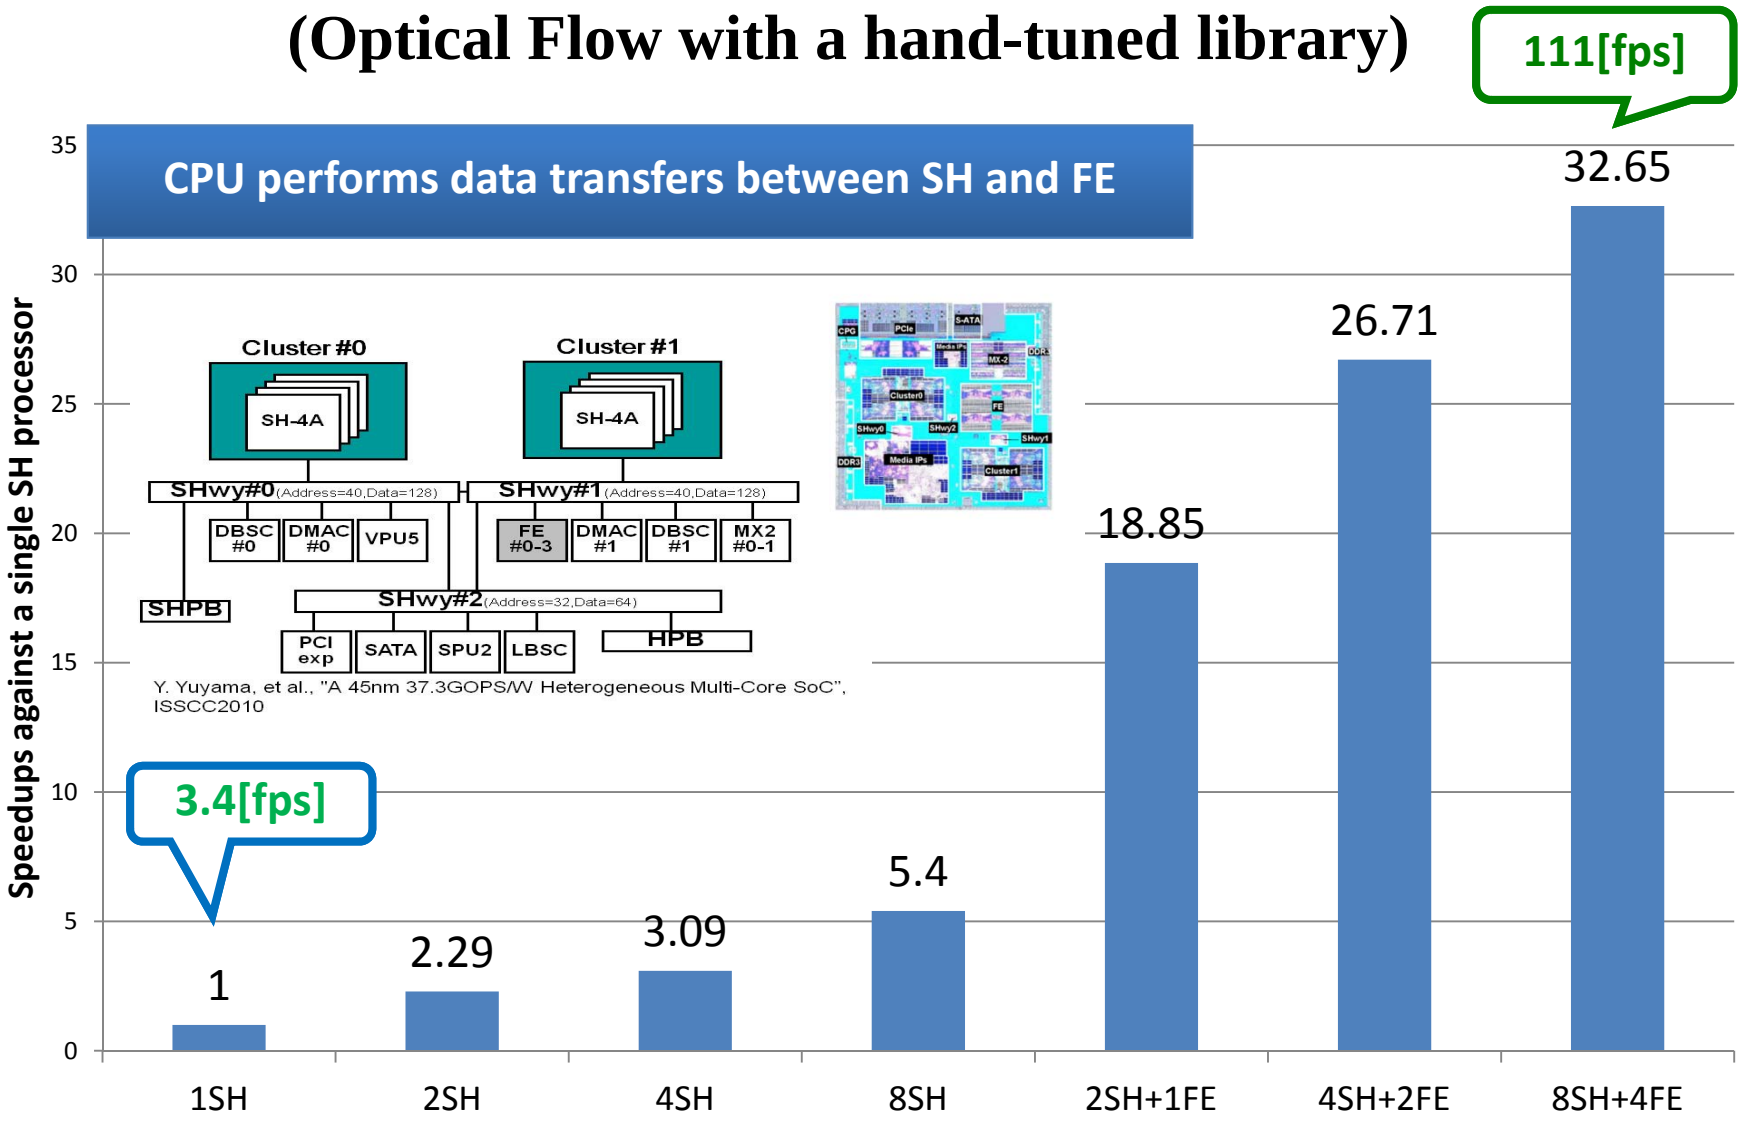

8 Core RP2 Chip Block Diagram

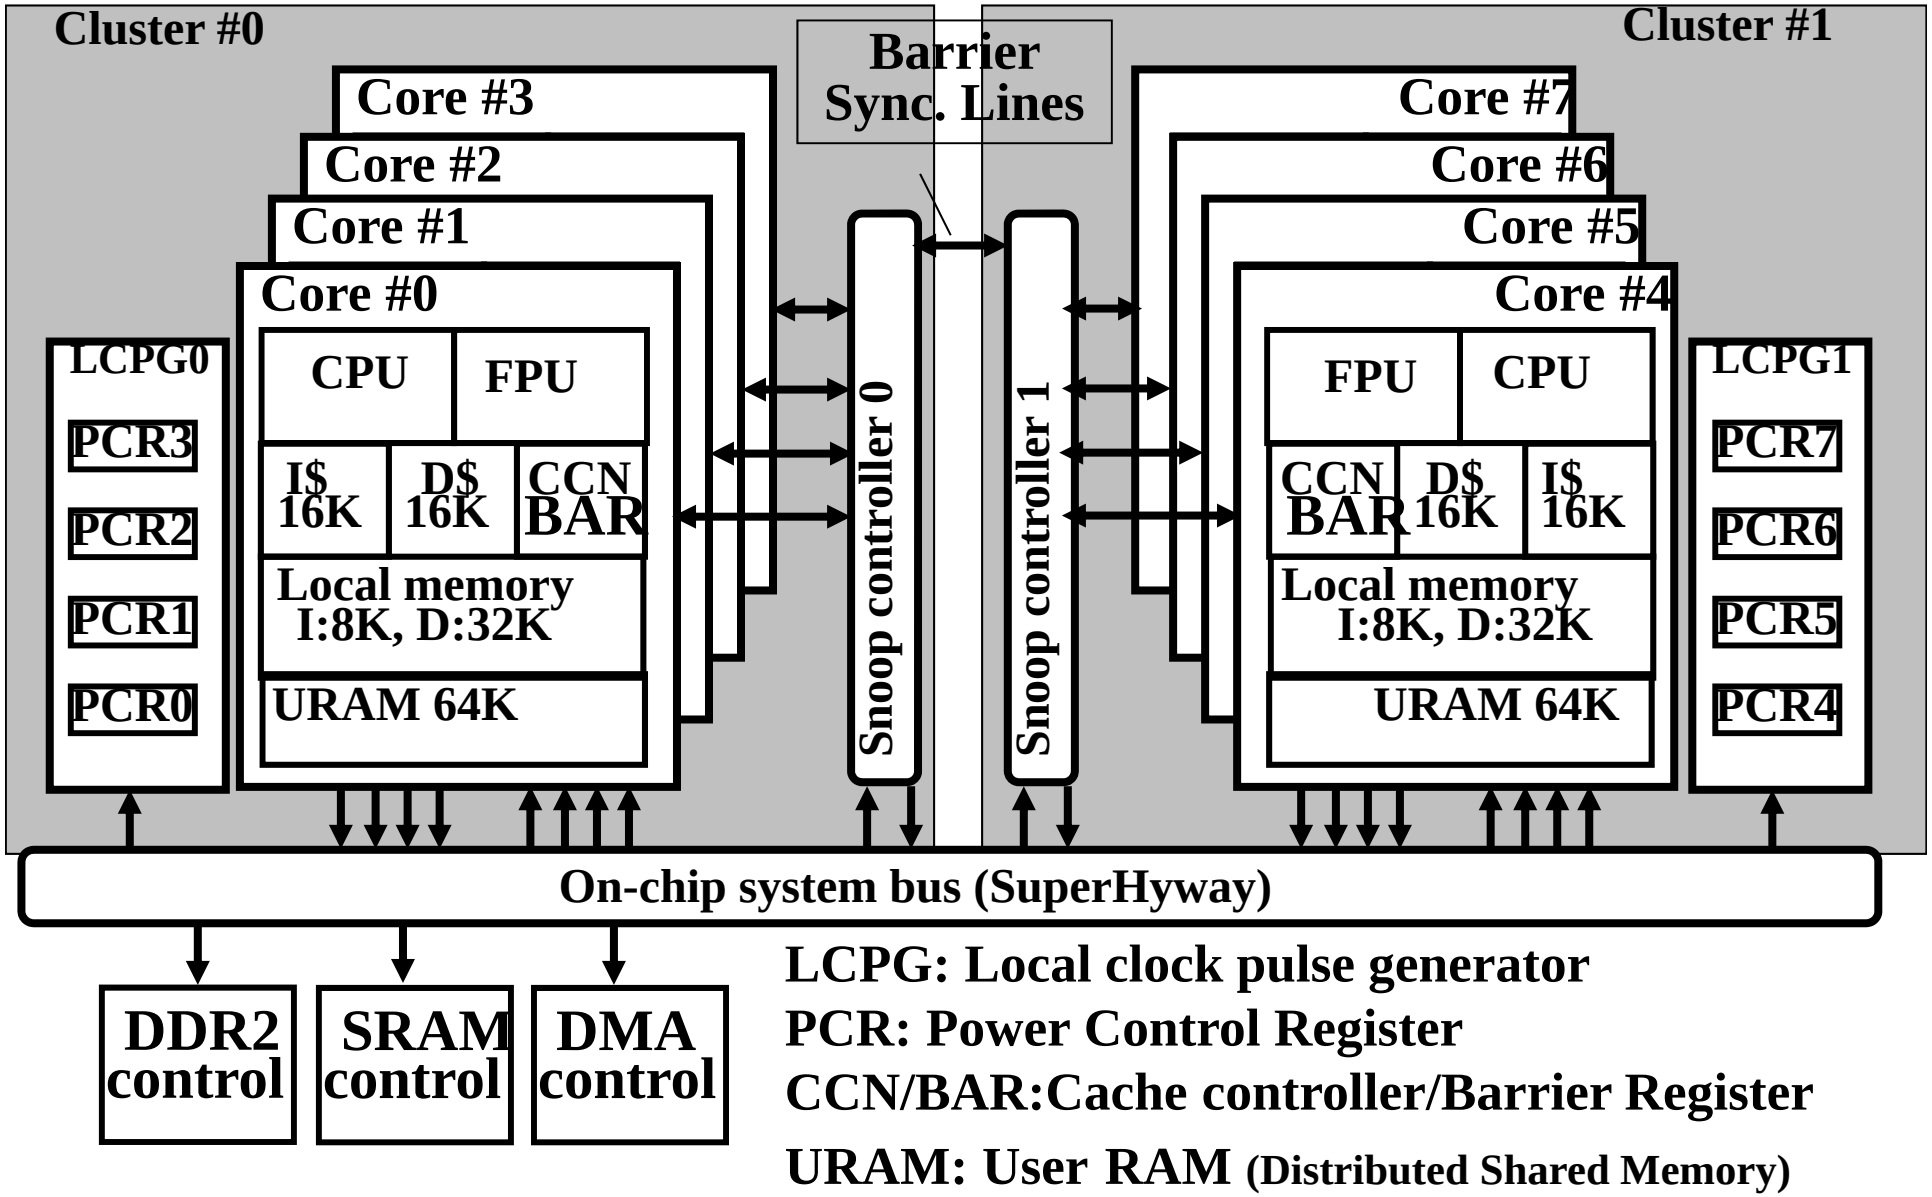

LCPG: Local clock pulse generator

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

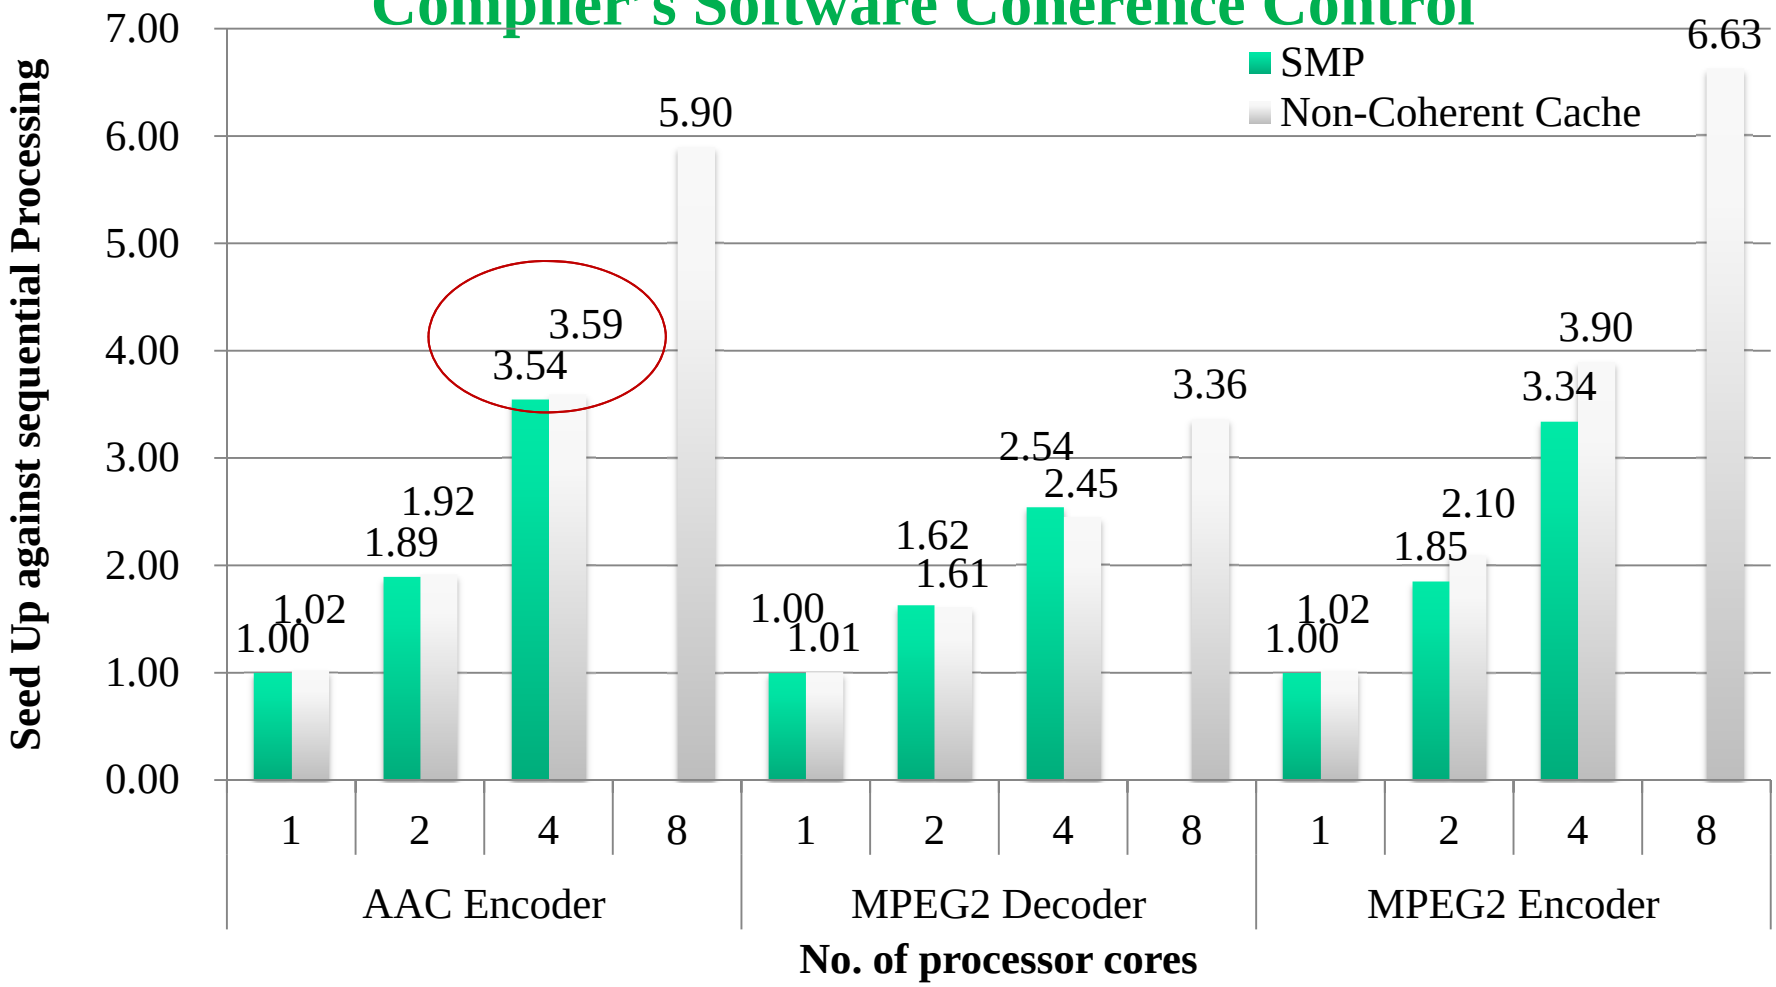

92 Times Speedup against the Sequential Processing for GMS Earthquake Wave Propagation Simulation on Hitachi SR16000

（Power7 Based 128 Core Linux SMP）

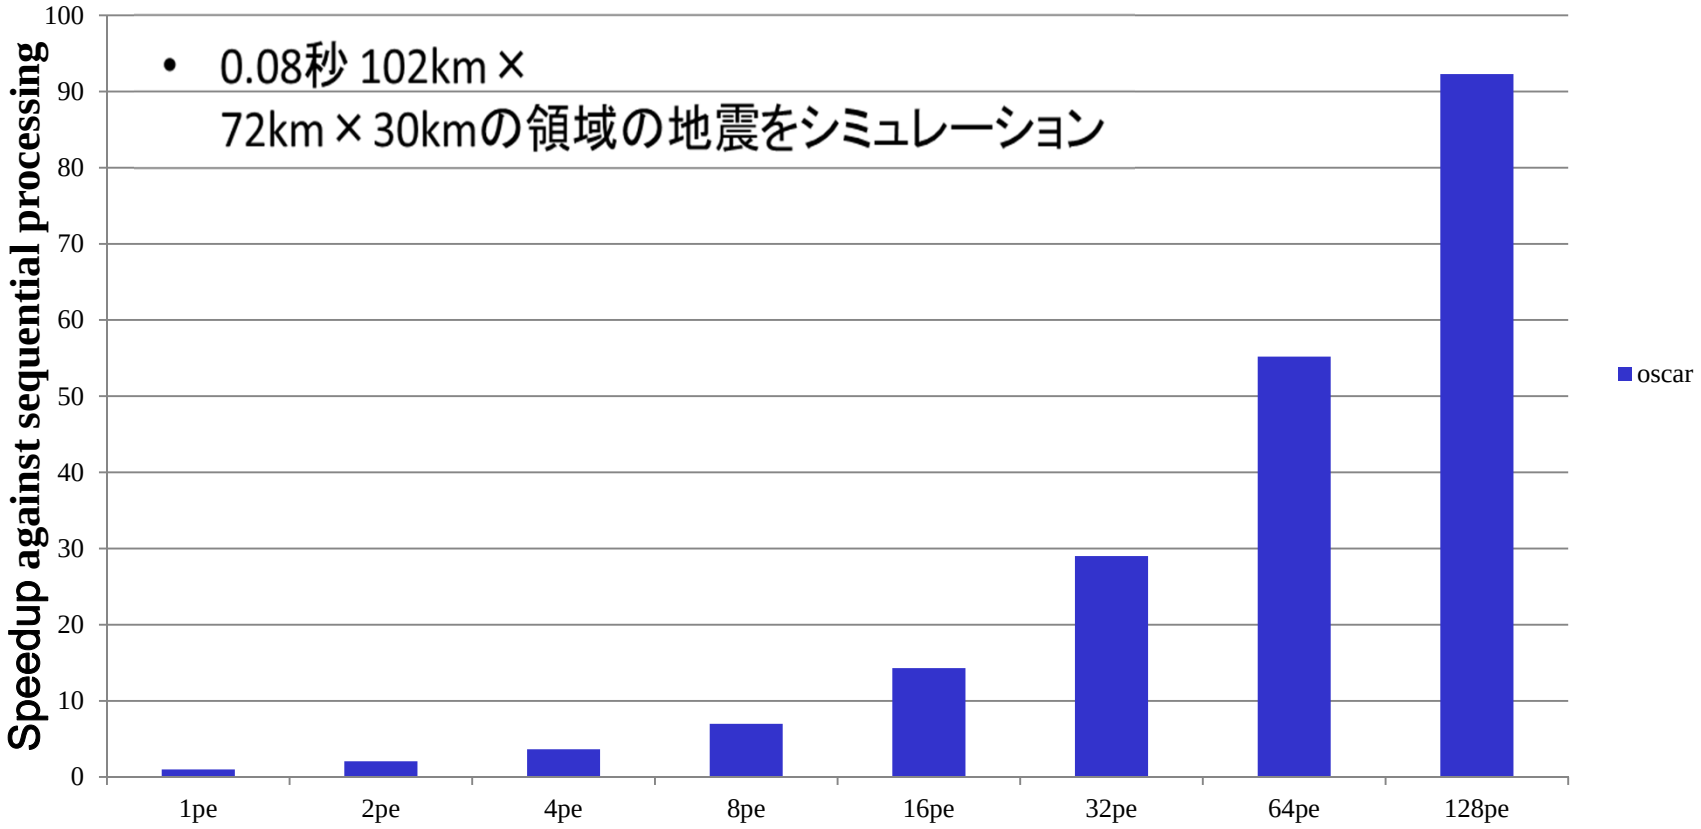

Cancer Treatment
Carbon Ion Radiotherapy

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

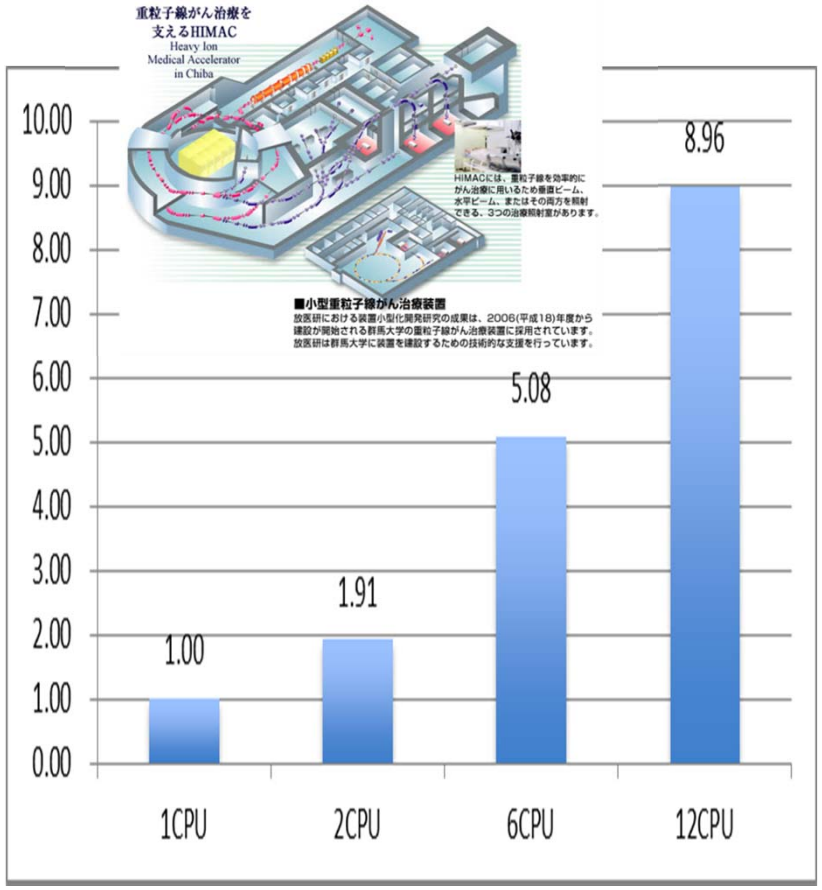

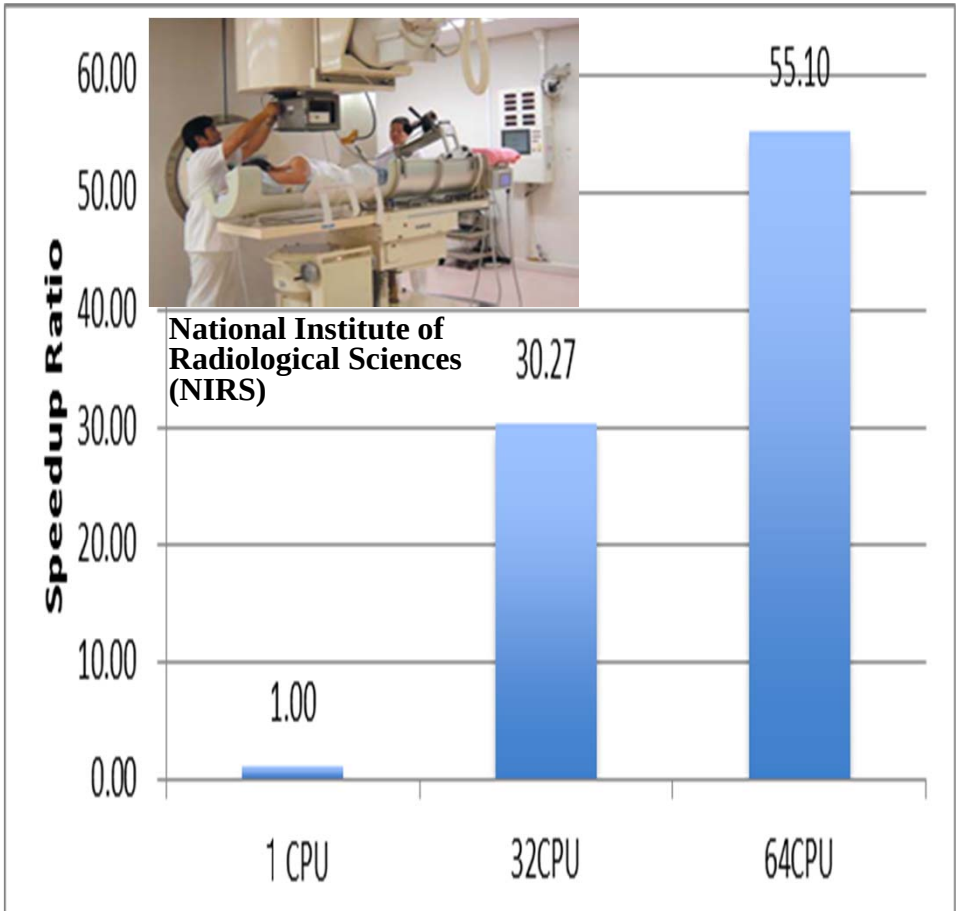

8.9times speedup by 12 processors

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

55 times speedup by 64 processors

IBM Power 7 64 core SMP (Hitachi SR16000)

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 processors for Carbon Ion Radiotherapy Cancer treatment on IBM Power 7 64 core SMP (Hitachi SR16000)
  - 46 Times Speedup for GMS Earthquake Wave Propagation Simulation on 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.
  - 2.2 times speedup on the average against Intel Composer XE 2011on AMD 12-core SMP against Intel Composer XE 2011 Based on Opteron 6174.
  - 1.9 times speedup on the average on Intel 12 core SMP based on 6-core Xeon X5670.
  - 2.9 Times Speed-up for AAC Encodeing on 3 Core NaviEngine (ARM MPcore) with Realtime OS eT-Kernel Multi-Core Edition

OSCAR Multi-Core Architecture

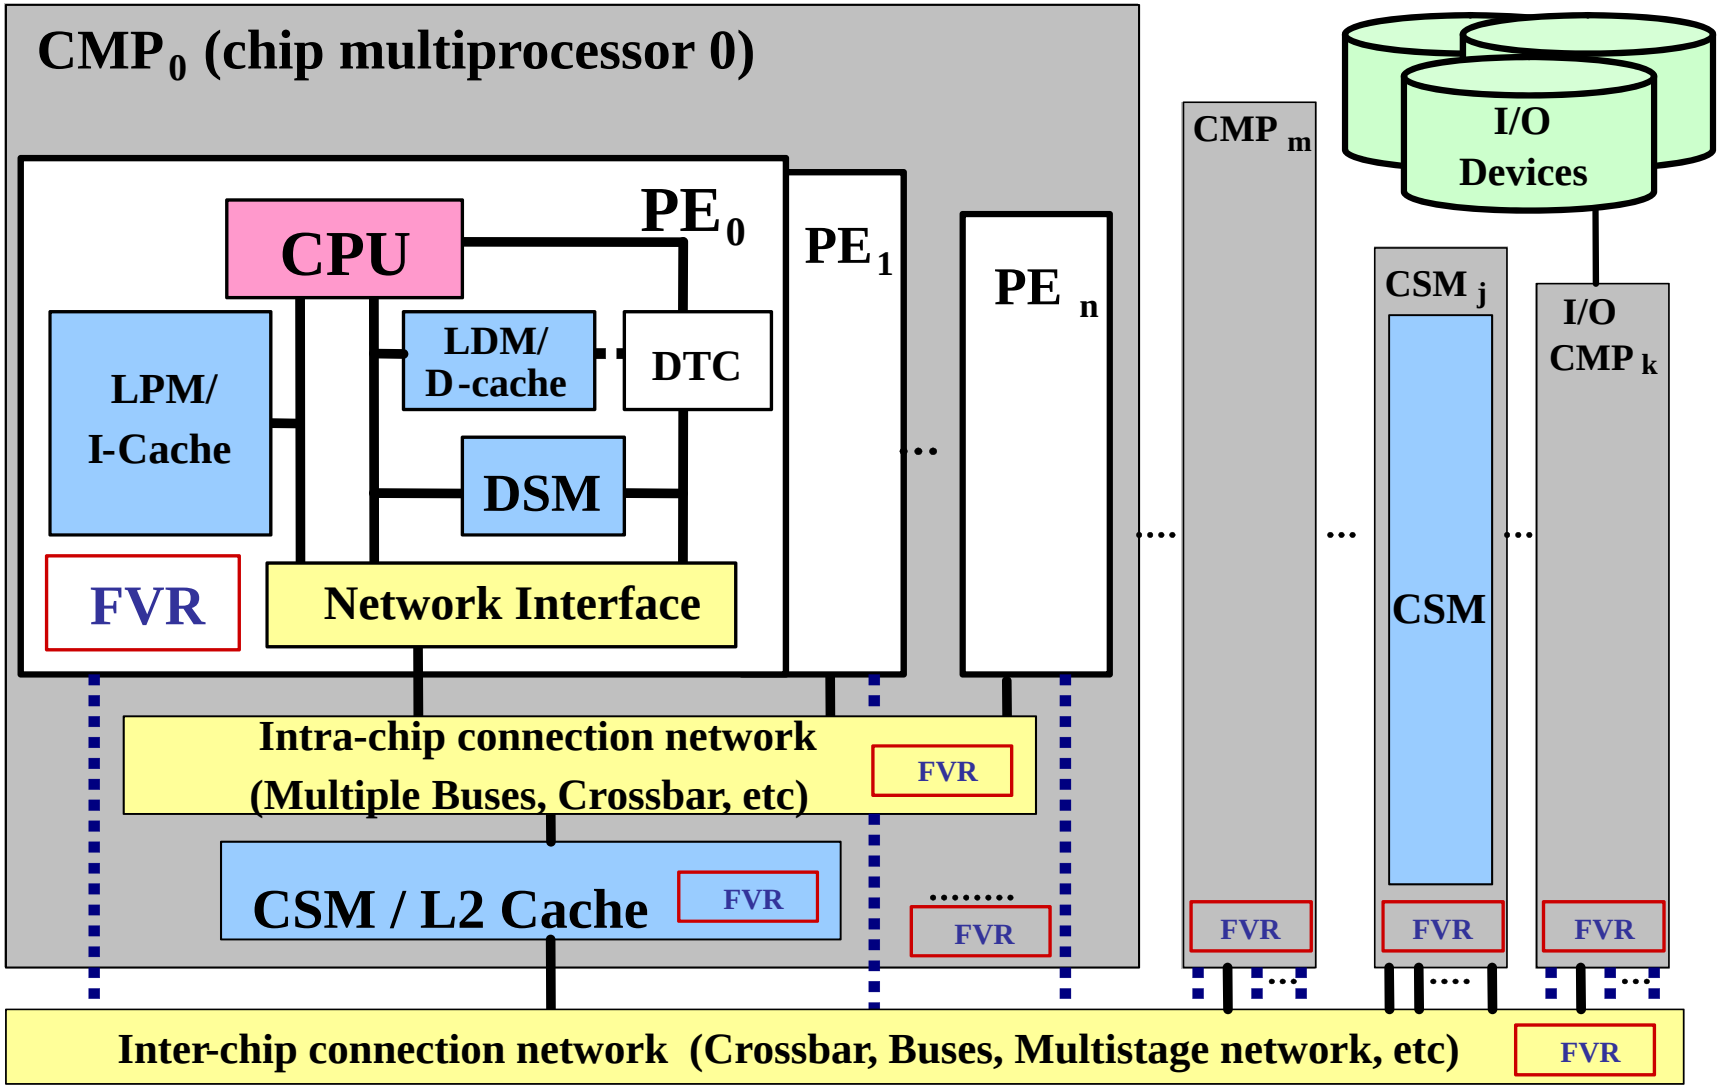

CSM: central shared mem.

DSM: distributed shared mem.

DTC: Data Transfer Controller

LDM : local data mem.

LPM : local program mem.

FVR: frequency / voltage control register

OSCAR Memory Space (Global and Local Address Space)