A Multigrain Parallelizing Compiler with Power Control for Multicore Processors

Hironori Kasahara

Professor, Department of Computer Science

Director, Advanced Chip-Multiprocessor Research Institutes

Waseda University

Tokyo, Japan

http://www.kasahara.cs.waseda.ac.jp

Feb. 5, 2008, 13:30-15:00 at Intel Headquarter

Hironori Kasahara

＜Personal History＞

B.S. (1980,Waseda), M.S.(1982,Waseda), Ph.D.(1985,EE, Waseda). Res.Assoc. (1983,Waseda), Special Research Fellow JSPS (1985) ,Visiting Scholar (1985.Univ.California at Berkeley), . Assist. Prof. (1986.Waseda), Assoc. Prof.(1988,Waseda), Visiting Research Scholar(1989-1990. Center for Supercomputing R&D, Univ.of Illinois at Urbana-Champaign), Prof.(1997-,Dept. CS, Waseda). , IFAC World Congress Young Author Prize (1987), IPSJ Sakai Memorial Special Award (1997), STARC Industry-Academia Cooperative Research Award (2004)

＜Activities for Societies＞

IPSJ：Sig. Computer Architecture(Chair), Trans of IPSJ Editorial Board (HG Chair), Journal of IPSJ Editorial Board (HWG Chair), 2001 Journal of IPSJ Special Issue on Parallel Processing(Chair of Editorial Board: Guest Editor, JSPP2000 (Program Chair) etc.

ACM ：International Conference on Supercomputing(ICS)(Program Committee)
Int'l conf. on Supercomputing (PC, esp. '96 ENIAC 50th Anniversary Co-Prog. Chair).

IEEE: Computer Society Japan Chapter Chair, Tokyo Section Board Member, SC07 PC

OTHER: PCs of many conferences on Supercomputing and Parallel Processing.

<Activities for Governments>

METI：IT Policy Proposal Forum(Architecture/HPC WG Chair),
Super Advanced Electronic Basis Technology Investigation Committee

NEDO:Millennium Project IT21 "Advanced Parallelizing Compiler"(Project Leader),
Computer Strategy WG (Chair).Multicore for Realtime Consumer Electronics Project Leader etc.

MEXT:Earth Simulator project evaluation committee, 10PFLOPS Supercomputer evaluat. comm.

JAERI: Research accomplishment evaluation committee, CCSE 1st class invited researcher.

JST: Scientific Research Fund Sub Committee, COINS Steering Committee ,
Precursory Research for Embryonic Science and Technology (Research Area Adviser)

Cabinet Office: CSTP Expert Panel on Basic Policy, Information & Communication Field Promotion Strategy , R&D Infrastructure WG, Software & Security WG

＜Papers＞

Papers 158, Invited Talks 64, Tech. Reports 114, Symposium 25, News Papers/TV/Web News/Magazine 139, IEEE Trans. Computer, IPSJ Trans., ISSCC, Cool Chips, Supercomputing, ACM ICS, etc.etc.

Multi-core Everywhere

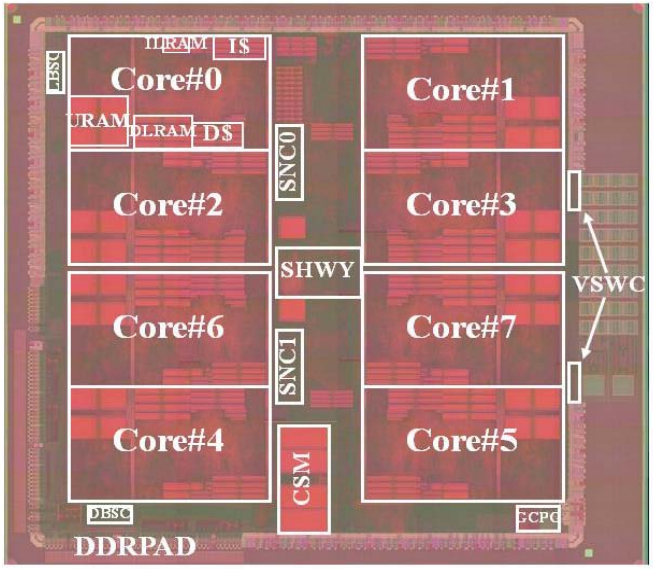

OSCAR Type Multi-core Chip by Renesas in METI/NEDO Multicore for Real-time Consumer Electronics Project (Leader: Prof.Kasahara)

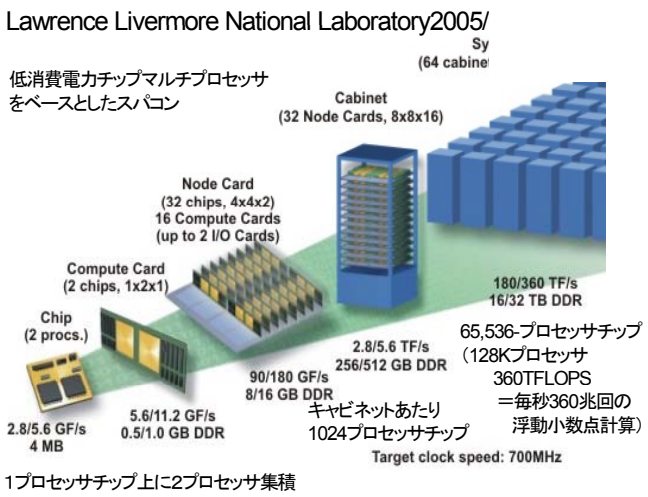

Multi-core from embedded to supercomputers

- **Consumer Electronics (Embedded)**
  - Mobile Phone, Game, Digital TV, Car Navigation, DVD, Camera,
  - IBM/ Sony/ Toshiba Cell, Fuijtsu FR1000,
  - NEC/ARMMPCore&MP211, Panasonic Uniphier,
  - Renesas SH multi-core(4 core RP1, 8 core RP2)
  - Tilera Tile64, SPI Storm-1(16 VLIW cores)
- **PCs, Servers**
  - Intel Quad Xeon, Core 2 Quad, Montvale, Tukwila, 80 core
  - AMD Quad Core Opteron, Phenom
- **WSs, Deskside & Highend Servers**
  - IBM Power4,5,5+,6 Sun Niagara(SparcT1,T2), Rock
- **Supercomputers**
  - Earth Simulator:40TFLOPS, 2002, 5120 vector proc.
  - IBM Blue Gene/L: 360TFLOPS, 2005, Low power CMP based 128K processor chips, BG/P 2008

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

Market of Consumer Electronics 1 Trillion Dollars in 2010 (World Wide)

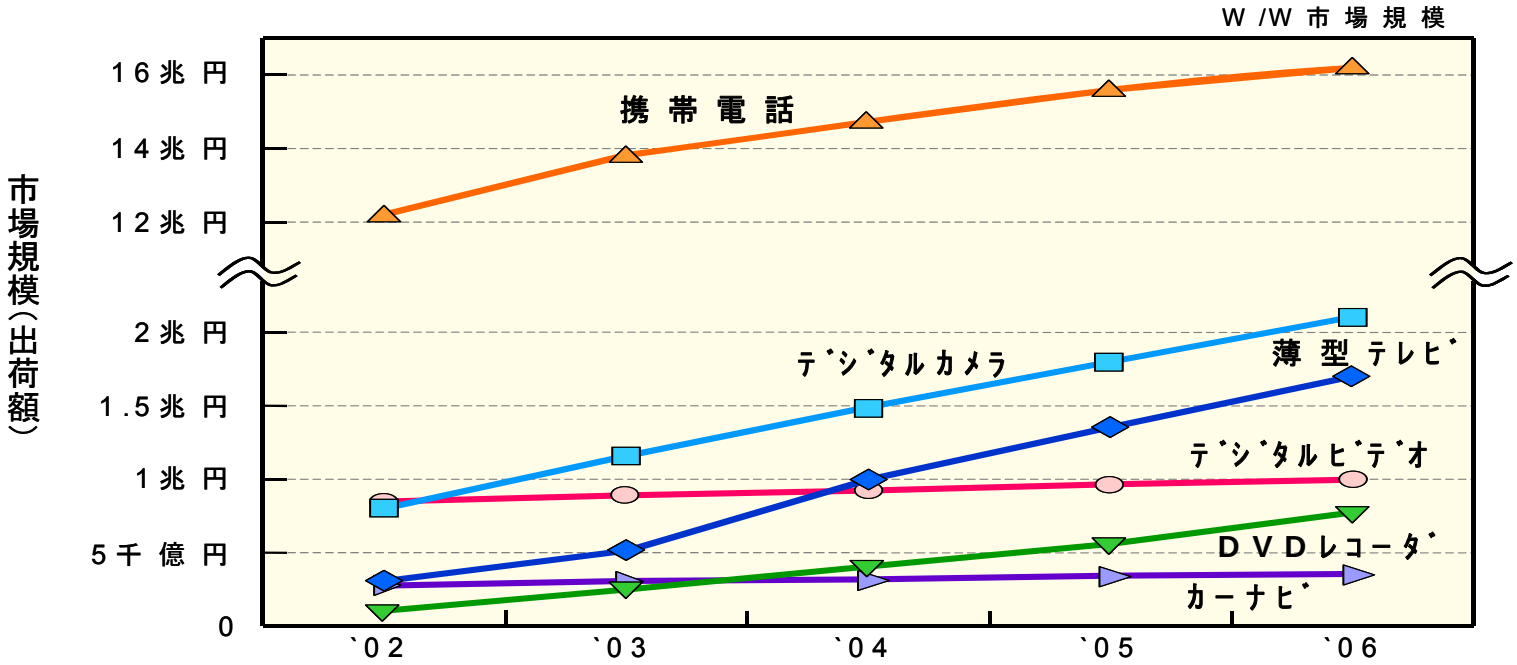

		'03	'07	年平均成長率%
Dig. Camera	デジタルスチルカメラ(M台)	49	76	12
Dig. TV	デジタルTV(M台)	6	27	45
DVD Recorder	DVDレコーダ(M台)	3.6	33	74
DVD for PC	PC用DVD(記録型)(M台)	27	114	43
Mobile Phone	携帯電話(M台)	490	670	8
LSI for Cars	自動車用半導体需要(B$)	14.0	20.9	11

Annual Growth Rates

2005.5.11
NEDOロードマップ報告会
電子・情報技術開発部
「技術開発戦略」より

Roadmap of compiler cooperative multicore project

- Millennium Project IT21 NEDO Advanced Parallelizing Compiler
  （Waseda Univ. Fujitsu,Hitachi, JIPDEC,AIST）
- STARC Compiler Cooperative Chip Multiprocessor
  （ Waseda Univ., Fujitsu, NEC, Toshiba, Panasonic,Sony）
- NEDO (2004.07-2007.06) Heterogeneous Multiprocessor
  （Waseda Univ., Hitachi）
- **NEDO (2005.06-2008.03) Multicore Technology for Realtime Consumer Electronics**
  - **Waseda Univ., Hitachi, Renesas, Fujitsu, NEC, Toshiba, Panasonic**
    - **➢Power Saving Multicore Architecture, Parallelizing Compiler,API**
- **NEDO (2007.02- 2010.03) Heterogeneous Multicore for Consumer Electronics** Waseda

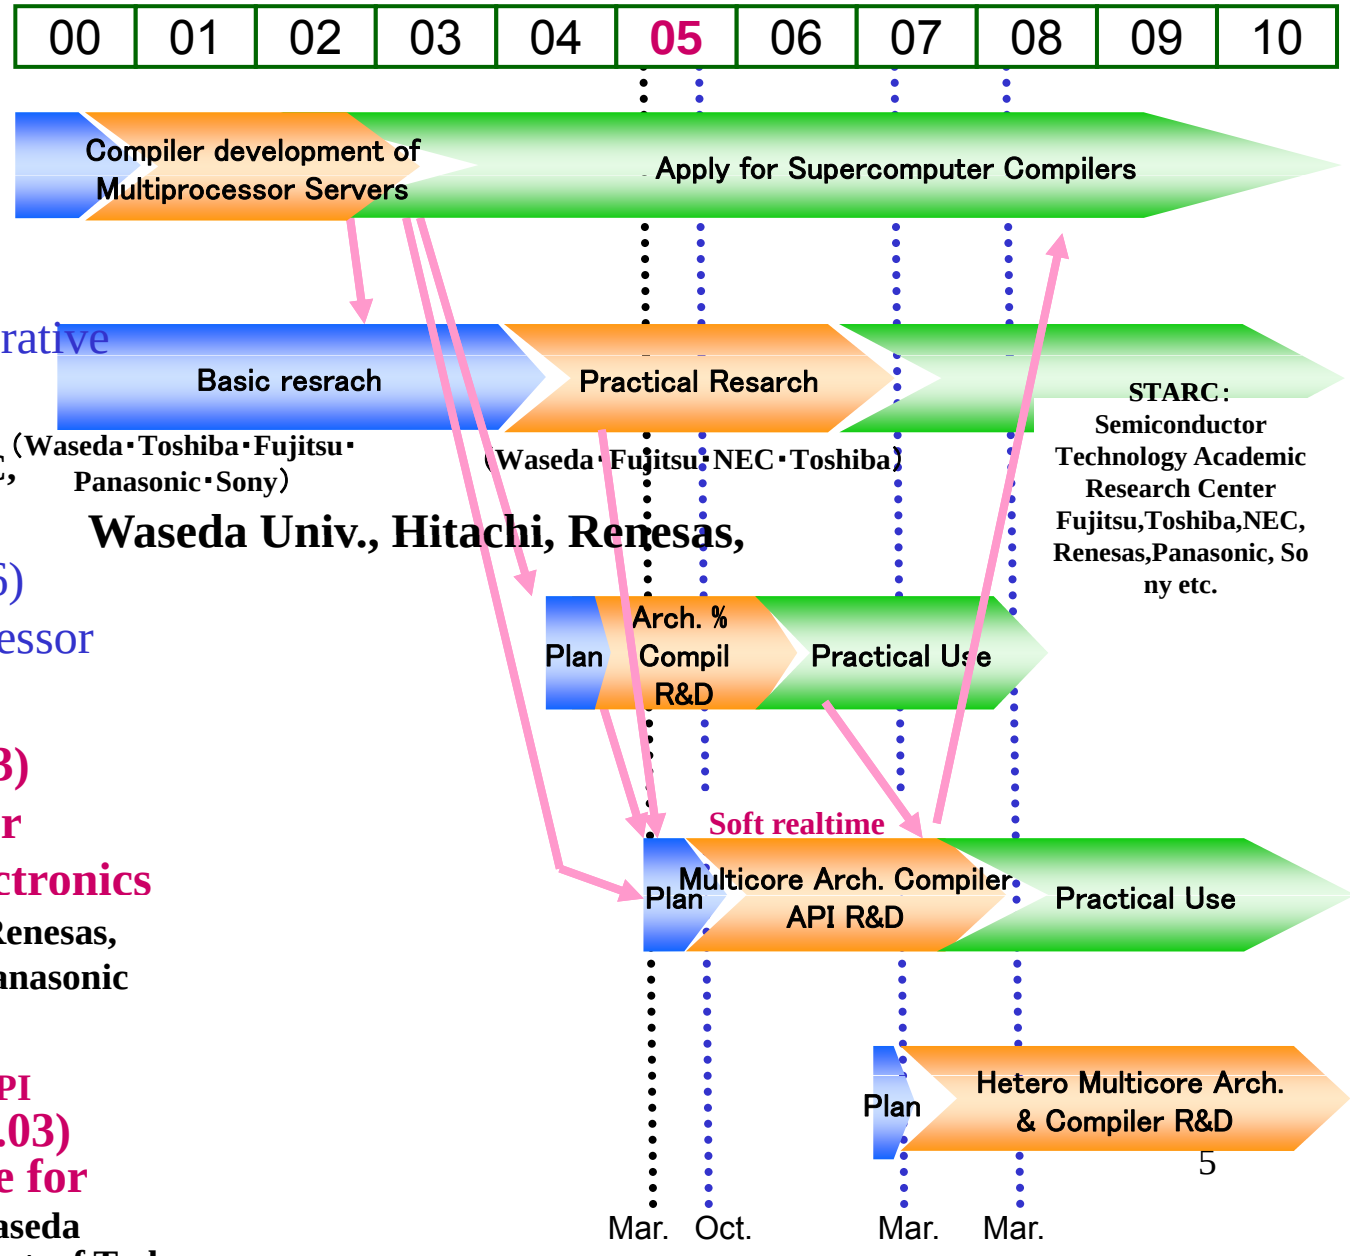

METI/NEDO Advanced Parallelizing Compiler Technology Project

Millenium Project IT21 2000.9.8 –2003.3.31
Waseda Univ., Fujitsu, Hitachi, AIST

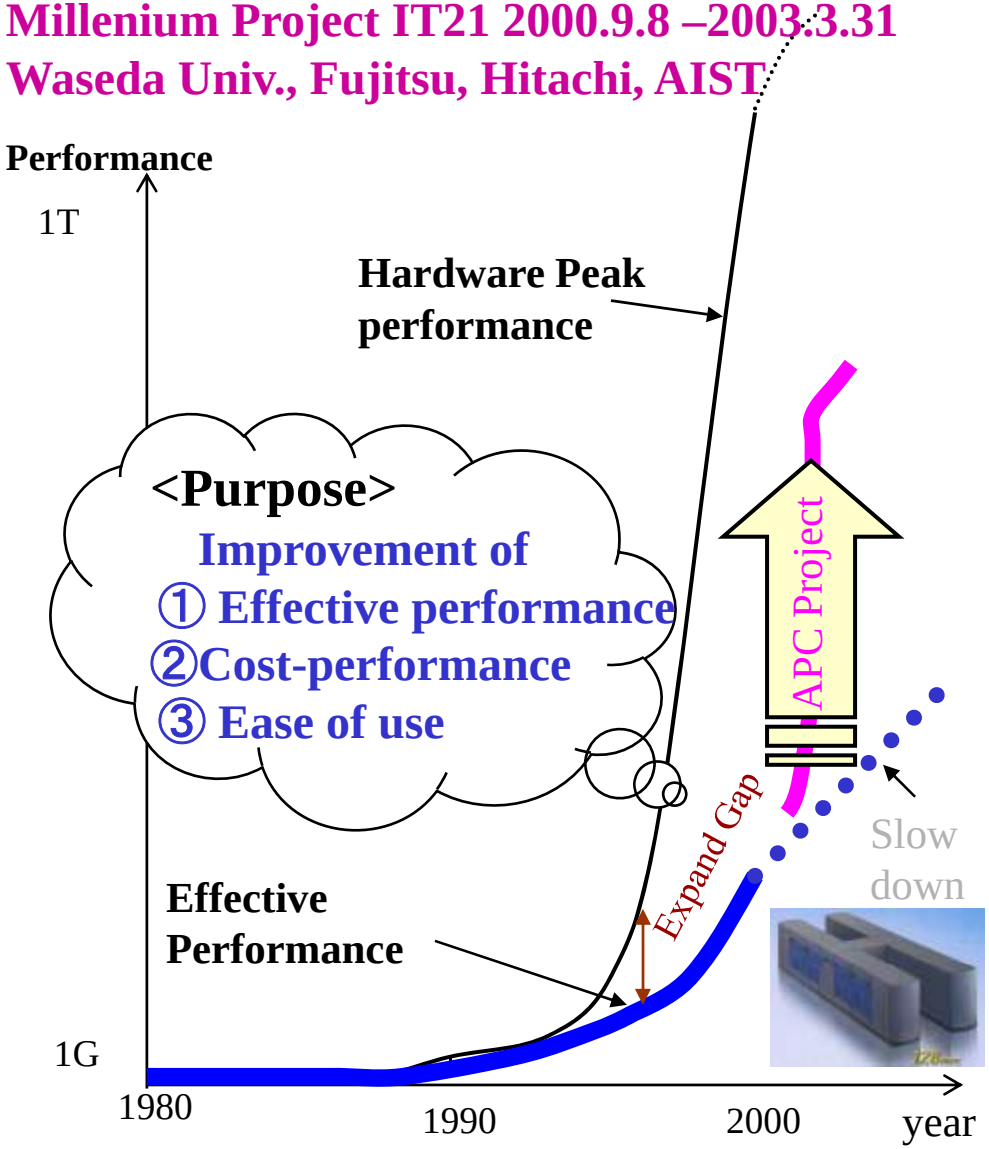

Theoretical maximum performance vs. Effective performance of HPC

Background and Problems

①Adoption of parallel processing as a core technology on PC to HPC
② Increase of importance of software on IT
③ Need for improvement of cost-performance and usability

Contents of Research and Development

① R & D of advanced parallelizing compiler
Multigrain, Data localization, Overhead hiding
② R & D of Performance evaluation technology for parallelizing compilers

Goal: Double the effective performance

Ripple Effect

① Development of competitive next generation PC and HPC
② Putting the innovative automatic parallelizing compiler technology to practical use
③ Development and market acquisition of future single-chip multiprocessors
④ Boosting R&D in the following many fields:
IT, Bio-tech., Device, Earth environment, Next-generation VLSI design, Financial engineering, Weather forecast, New clean energy, Space development, Automobile, Electric Commerce, etc

Performance of APC Compiler on IBM pSeries690 16 Processors High-end Server

- IBM XL Fortran for AIX Version 8.1
  - Sequential execution : -O5 -qarch=pwr4
  - Automatic loop parallelization : -O5 -qsmp=auto -qarch=pwr4
  - OSCAR compiler : -O5 -qsmp=noauto -qarch=pwr4 (su2cor: -O4 -qstrict)

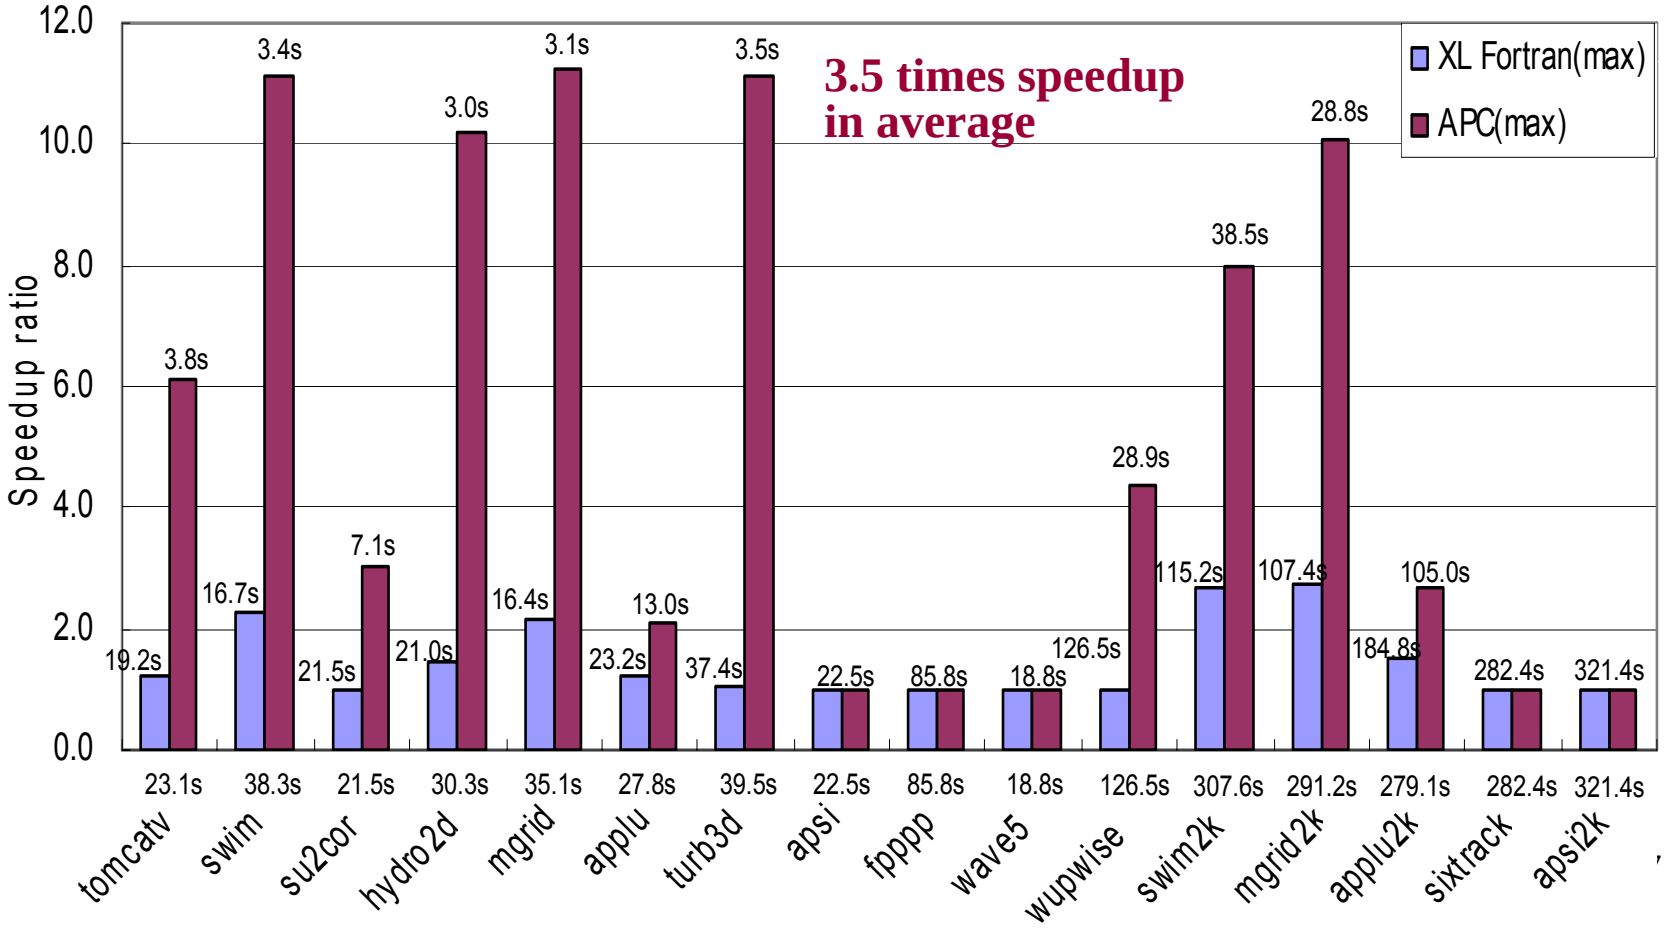

Performance of Multigrain Parallel Processing for 102.swim on IBM pSeries690

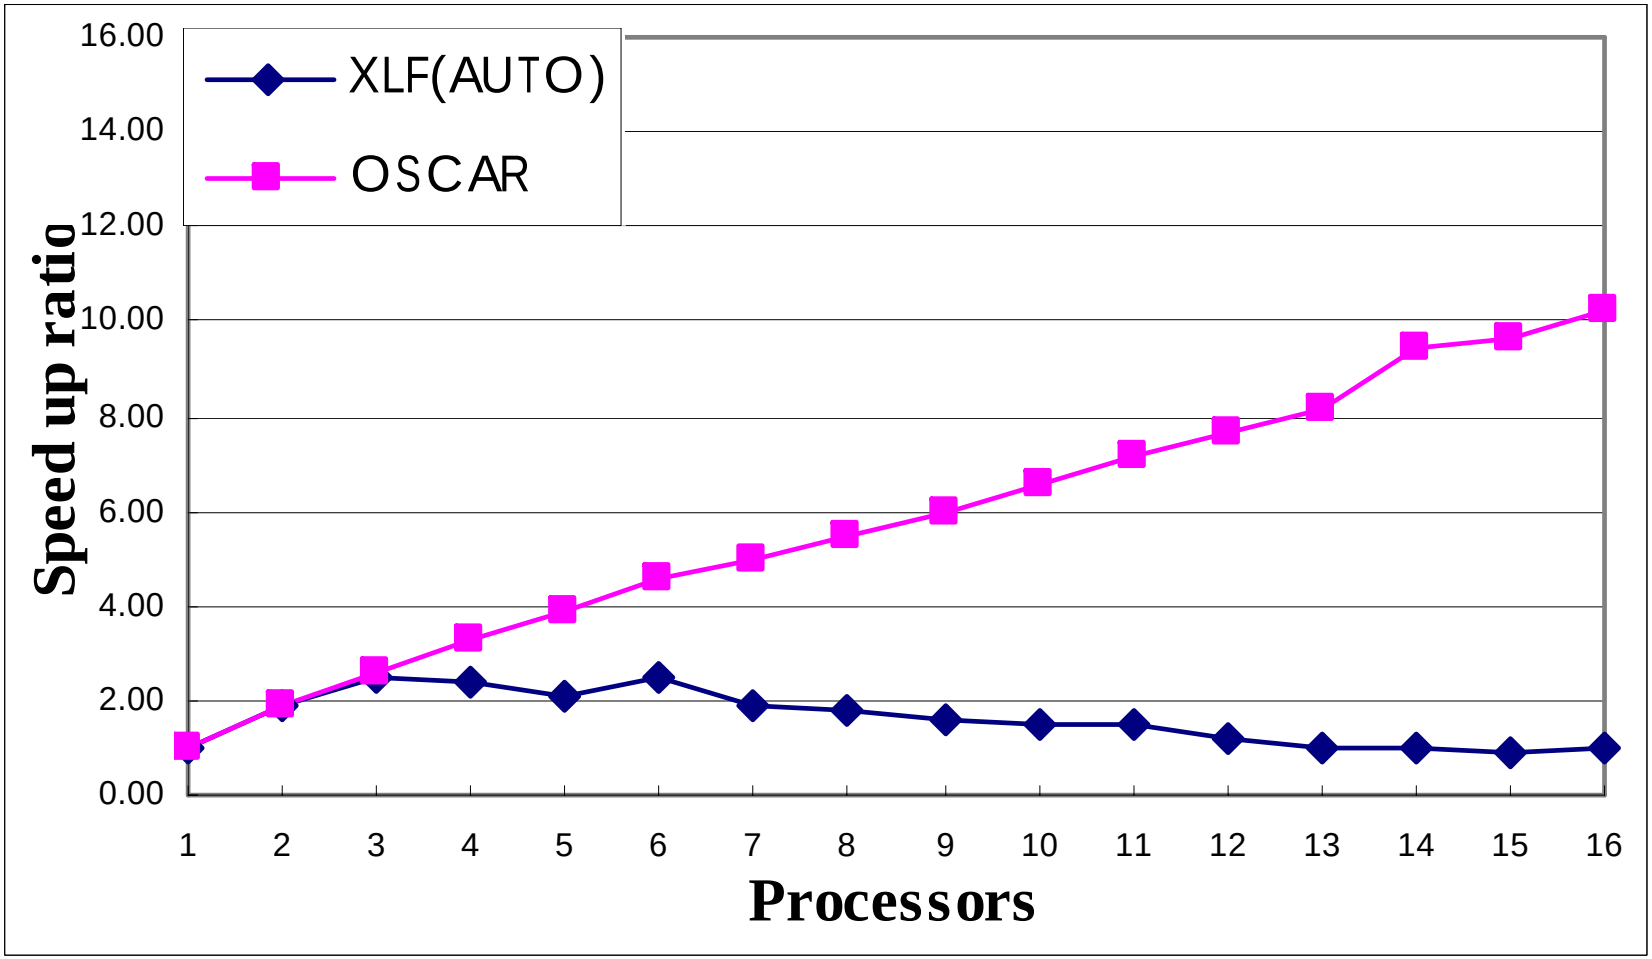

Performance of APC Compiler on Sun Ultra80 4 Processor Workstation

- Sun Forte Developer 6 Update 2
  - Sequential execution : -fast
  - Automatic loop parallelization : -fast -autopar -reduction -stackvar
  - OSCAR compiler : -fast -explicitpar -mp=openmp -stackvar

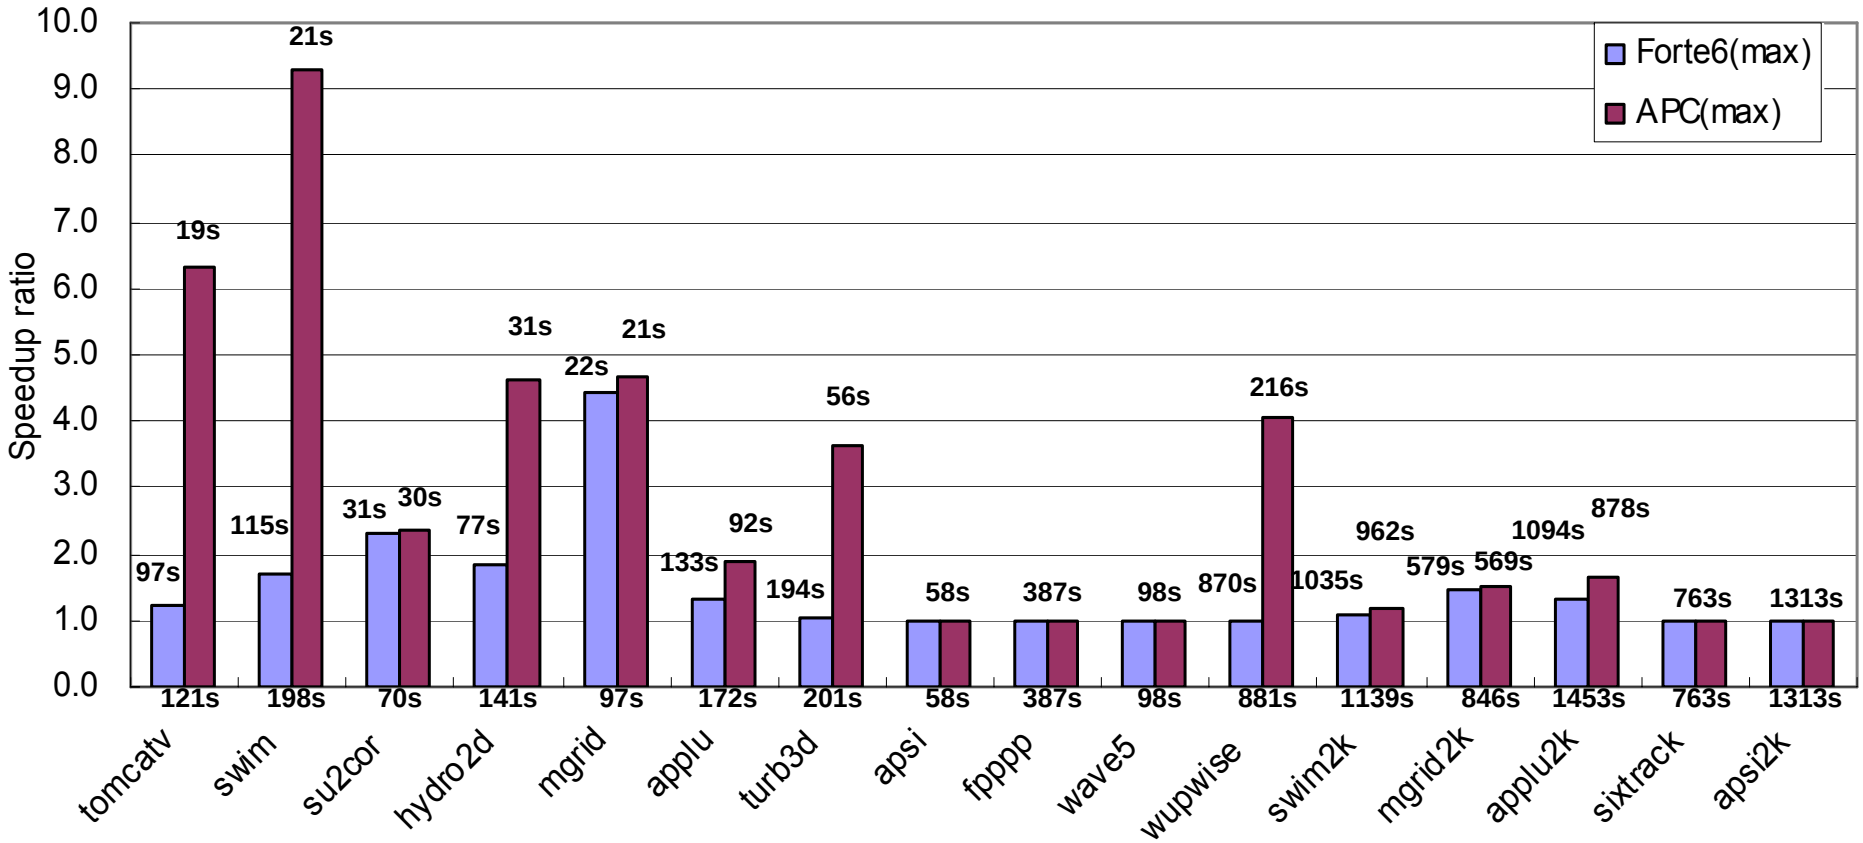

OSCAR Parallelizing Compiler

- **Improve effective performance, cost-performance and productivity and reduce consumed power**
  - **Multigrain Parallelization**
    - **Exploitation of parallelism from the whole program by use of coarse-grain parallelism among loops and subroutines, near fine grain parallelism among statements in addition to loop parallelism**
  - **Data Localization**
    - **Automatic data distribution for distributed shared memory, cache and local memory on multiprocessor systems.**
  - **Data Transfer Overlapping**
    - **Data transfer overhead hiding by overlapping task execution and data transfer using DMA or data pre-fetching**
  - **Power Reduction**
    - **Reduction of consumed power by compiler control of frequency, voltage and power shut down with hardware supports.**

Generation of Coarse Grain Tasks

Macro-tasks (MTs)

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

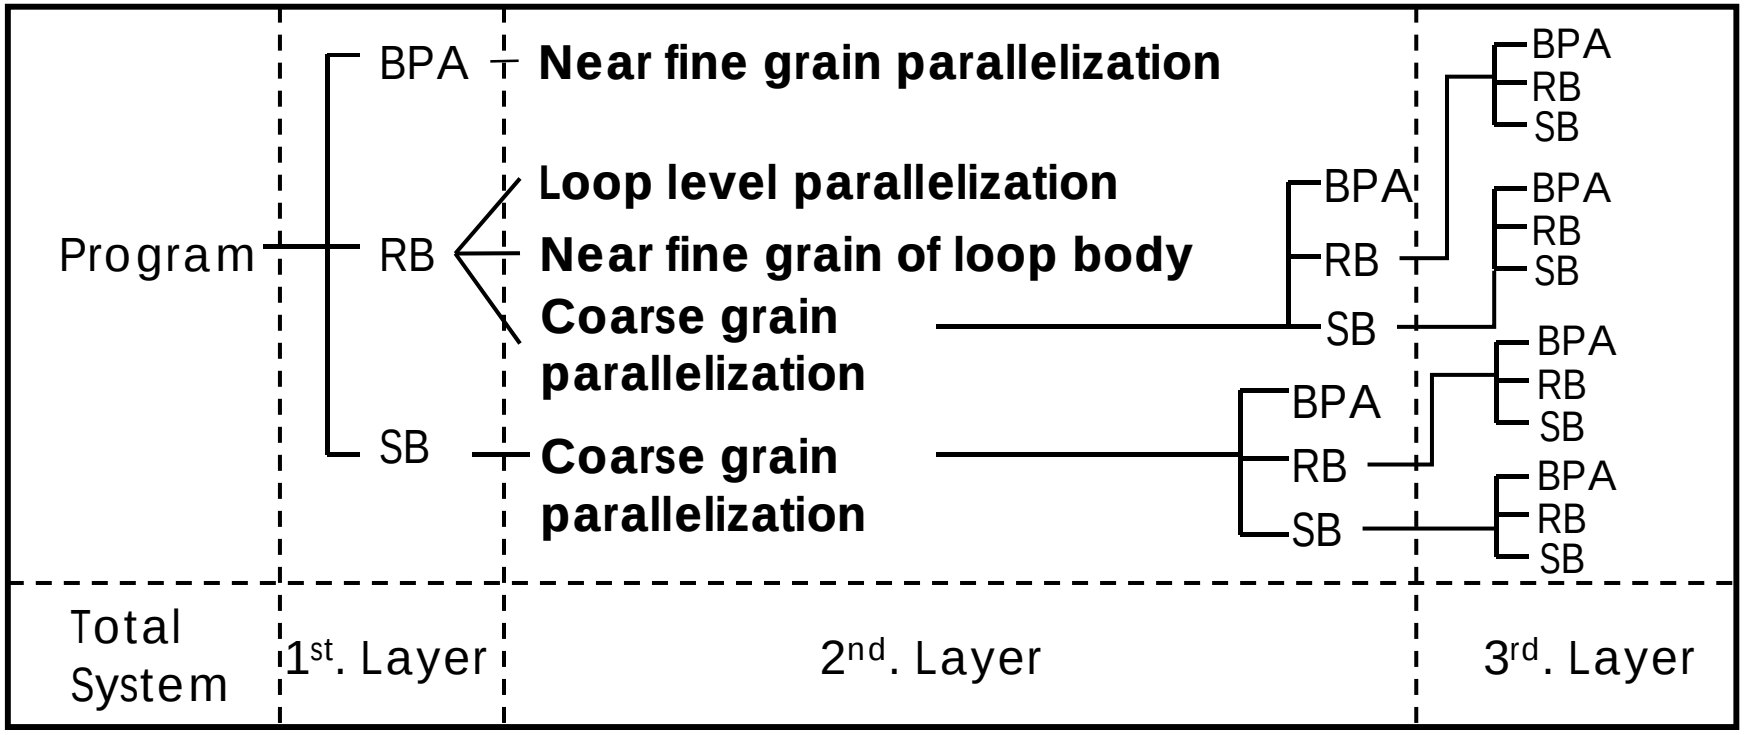

Earliest Executable Condition Analysis for coarse grain tasks (Macro-tasks)

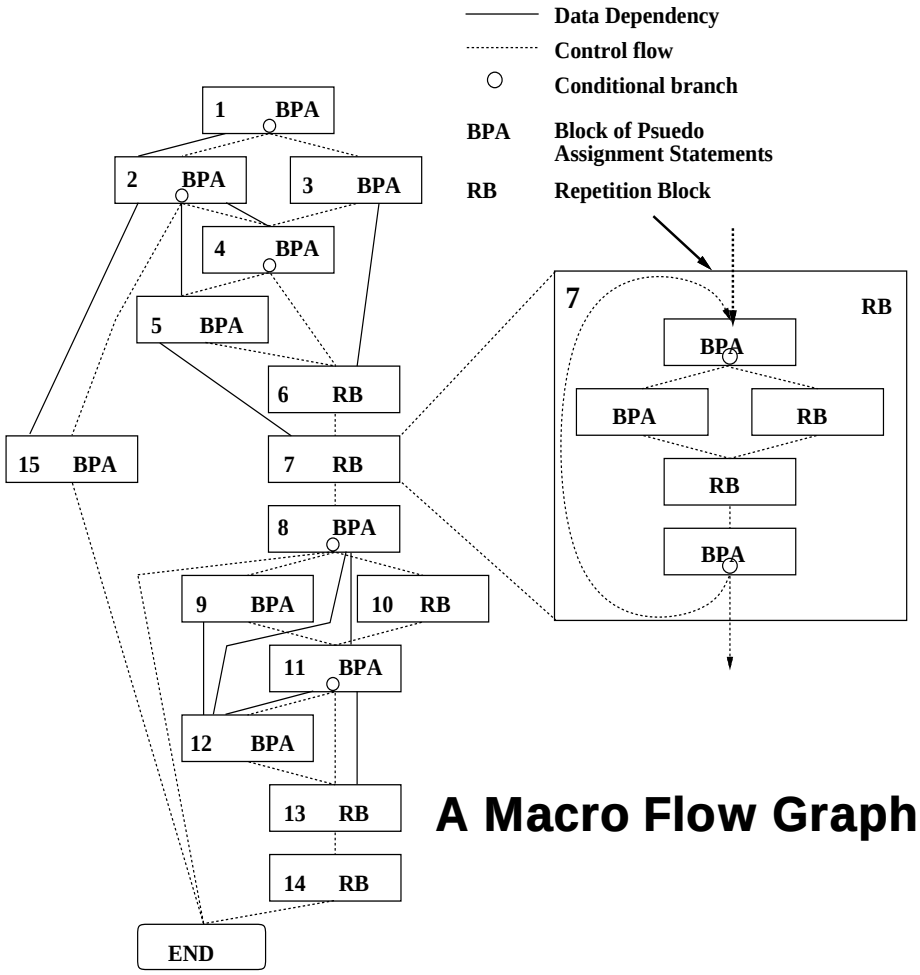

A Macro Flow Graph

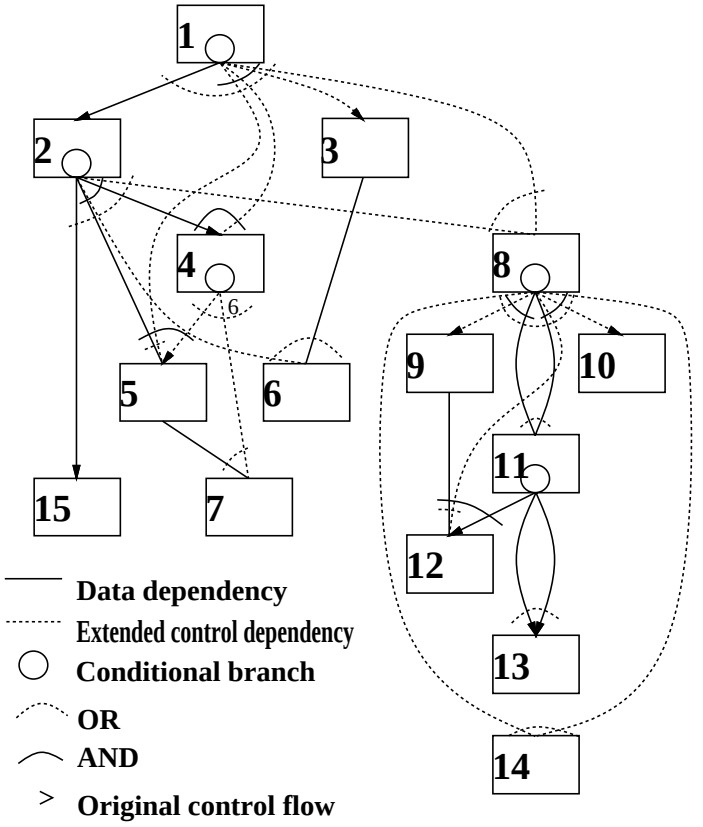

A Macro Task Graph

Automatic processor assignment in 103.su2cor

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

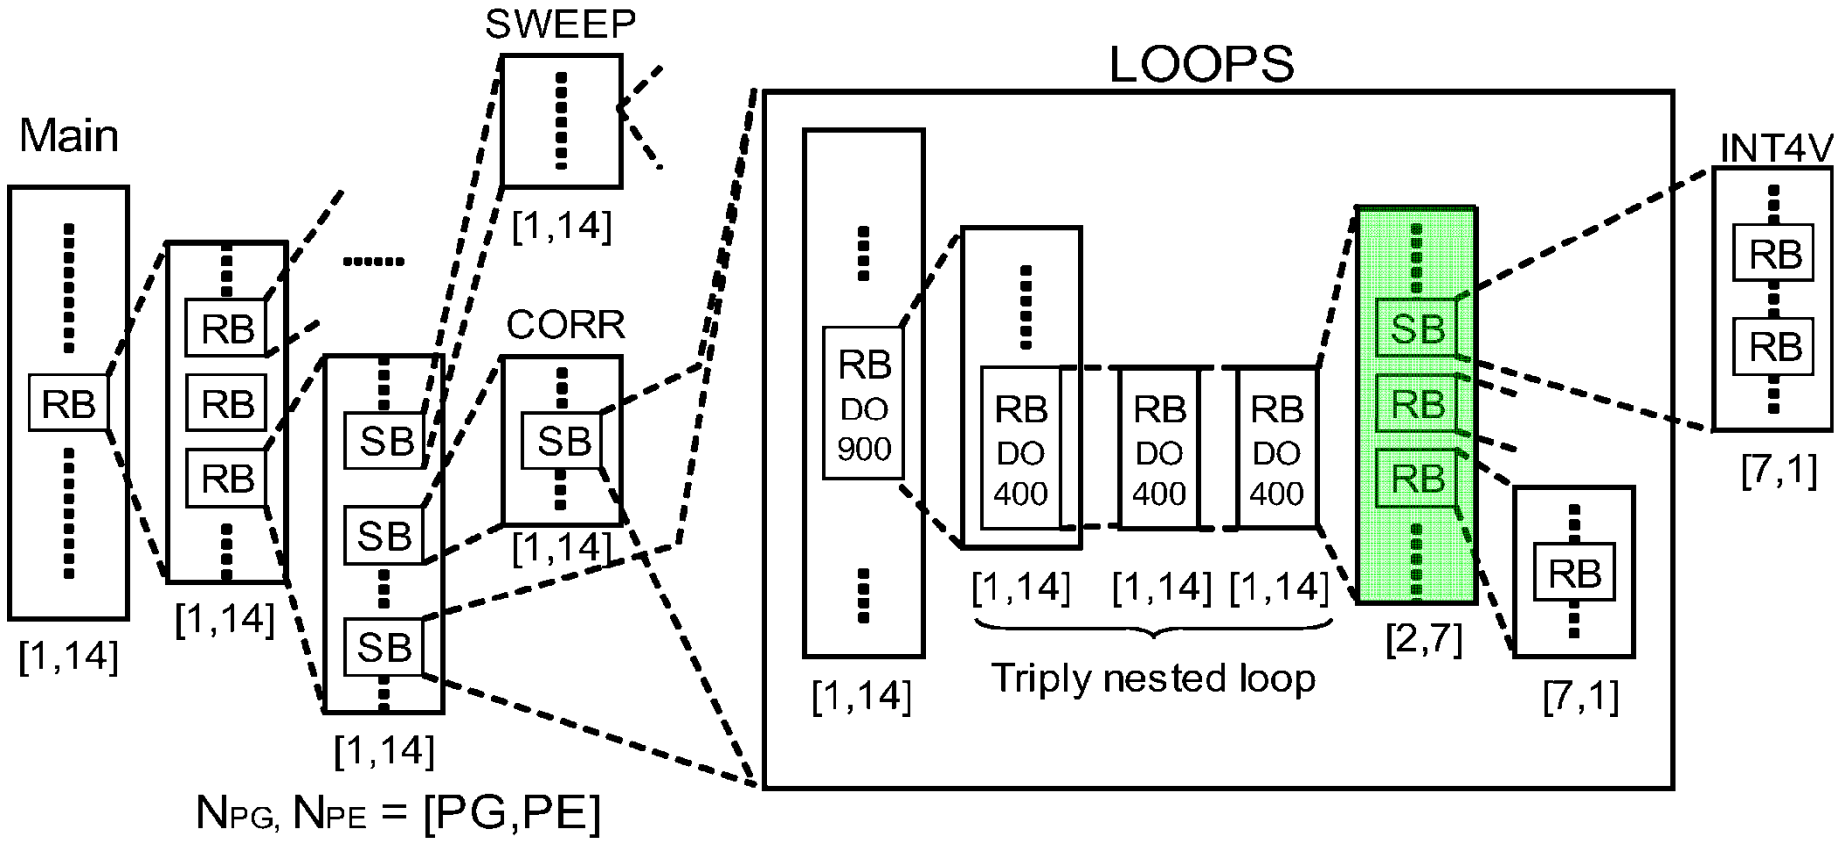

MTG of Su2cor-LOOPS-DO400

- Coarse grain parallelism PARA_ALD = 4.3

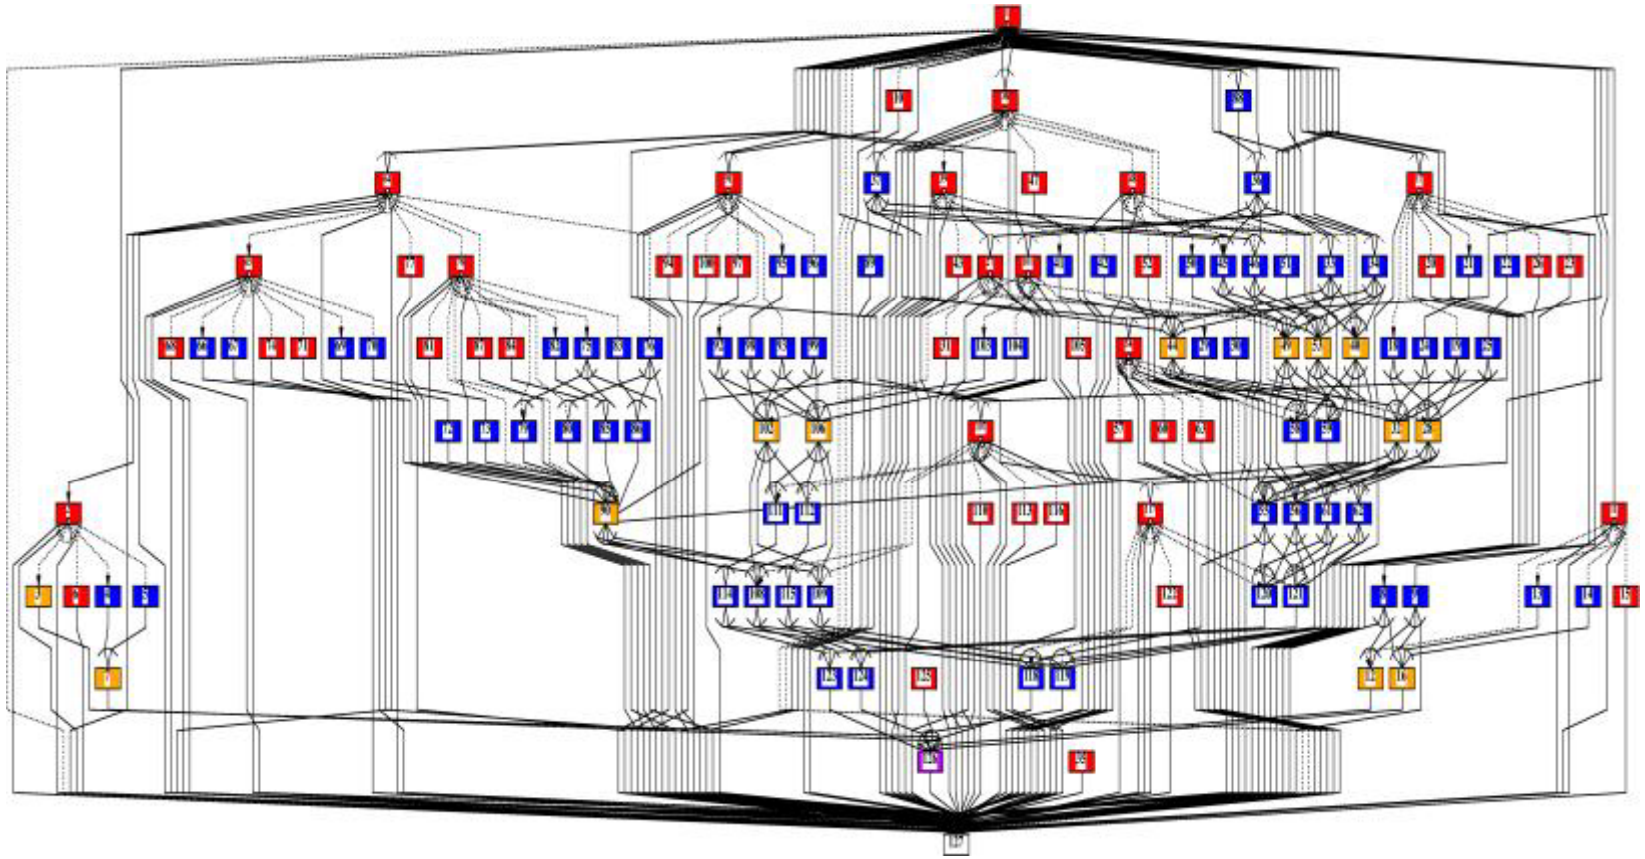

DOALL Sequential LOOP SB BB

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**

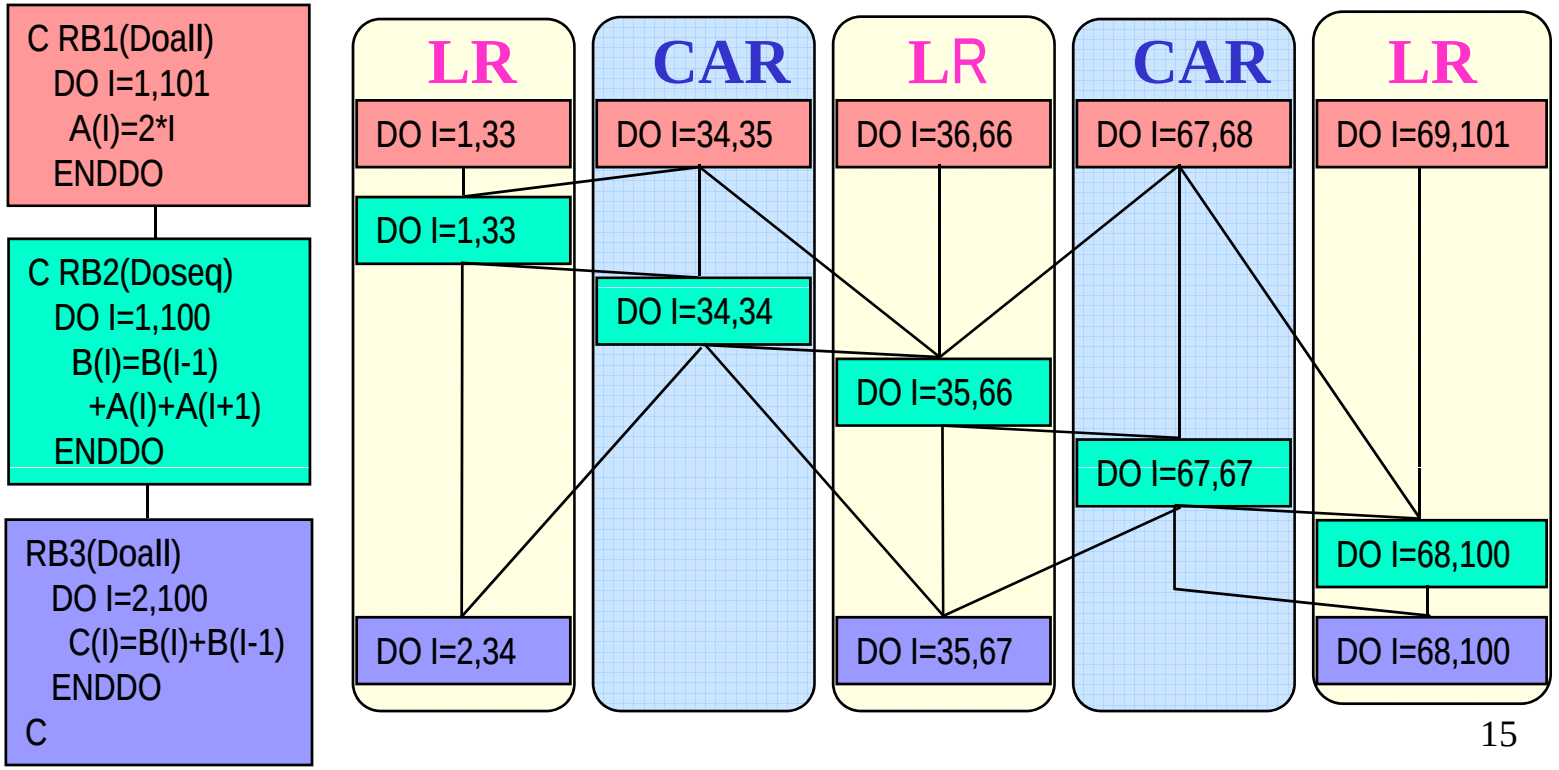

Data Localization

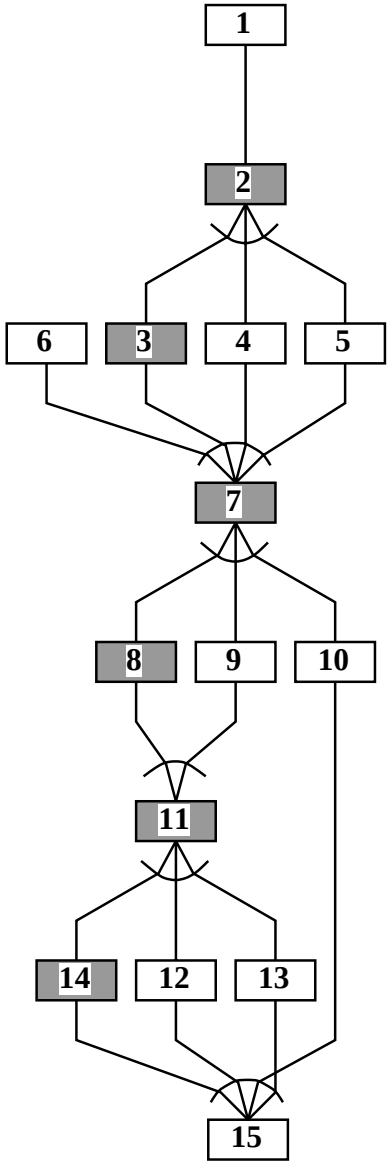

MTG

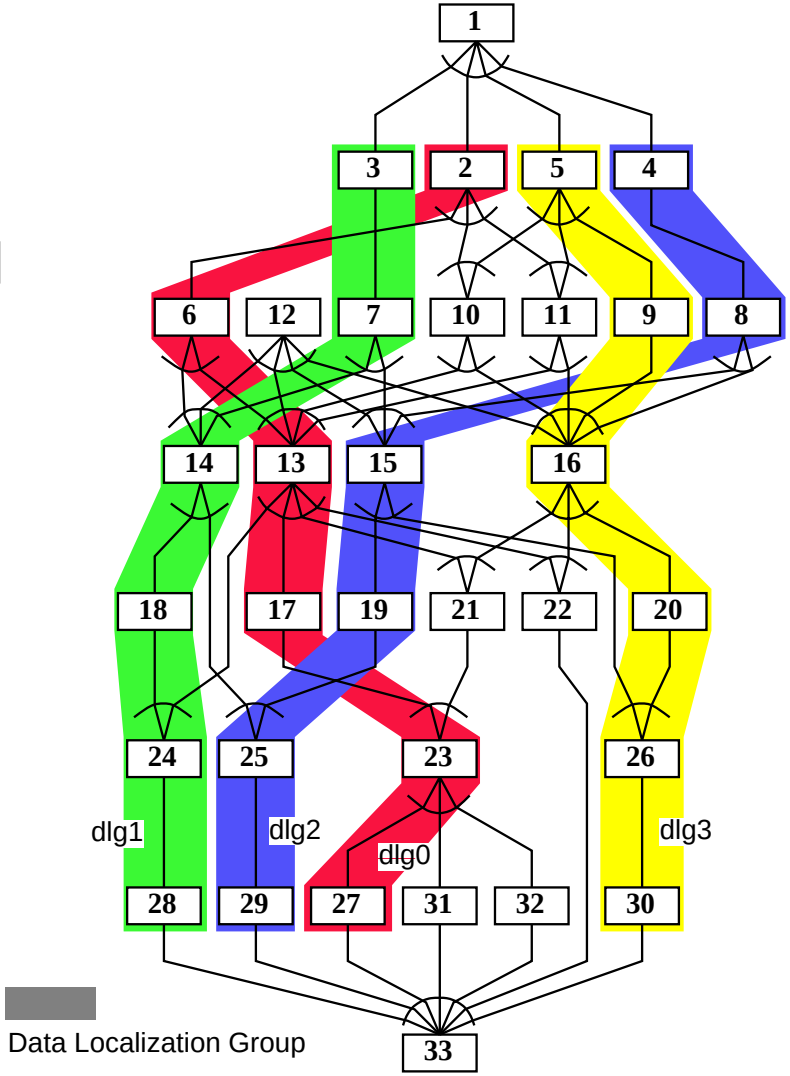

MTG after Division

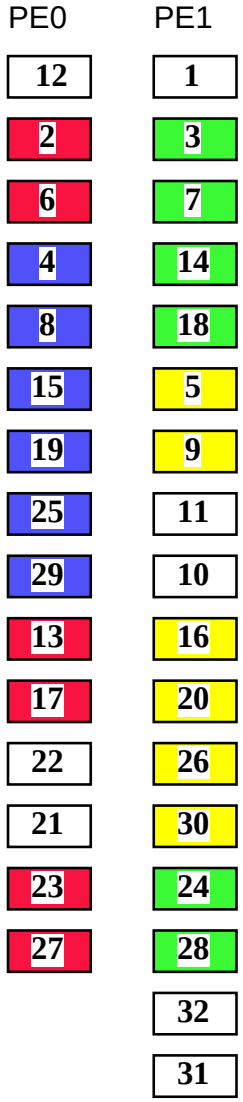

A schedule for two processors

An Example of Data Localization for Spec95 Swim

```
      DO 200 J=1,N
      DO 200 I=1,M
         UNEW(I+1,J) = UOLD(I+1,J)+
     1    TDTS8*(Z(I+1,J+1)+Z(I+1,J))*(CV(I+1,J+1)+CV(I,J+1)+CV(I,J)
     2       +CV(I+1,J))-TDTSDX*(H(I+1,J)-H(I,J))
         VNEW(I,J+1) = VOLD(I,J+1)-TDTS8*(Z(I+1,J+1)+Z(I,J+1))
     1       *(CU(I+1,J+1)+CU(I,J+1)+CU(I,J)+CU(I+1,J))
     2       -TDTSDY*(H(I,J+1)-H(I,J))
         PNEW(I,J) = POLD(I,J)-TDTSDX*(CU(I+1,J)-CU(I,J))
     1       -TDTSDY*(CV(I,J+1)-CV(I,J))
  200 CONTINUE
```

```
      DO 210 J=1,N
         UNEW(1,J) = UNEW(M+1,J)
         VNEW(M+1,J+1) = VNEW(1,J+1)
         PNEW(M+1,J) = PNEW(1,J)
  210 CONTINUE
```

```
      DO 300 J=1,N
      DO 300 I=1,M
         UOLD(I,J) = U(I,J)+ALPHA*(UNEW(I,J)-2.*U(I,J)+UOLD(I,J))
         VOLD(I,J) = V(I,J)+ALPHA*(VNEW(I,J)-2.*V(I,J)+VOLD(I,J))
         POLD(I,J) = P(I,J)+ALPHA*(PNEW(I,J)-2.*P(I,J)+POLD(I,J))
  300  CONTINUE
```

(a) An example of target loop group for data localization

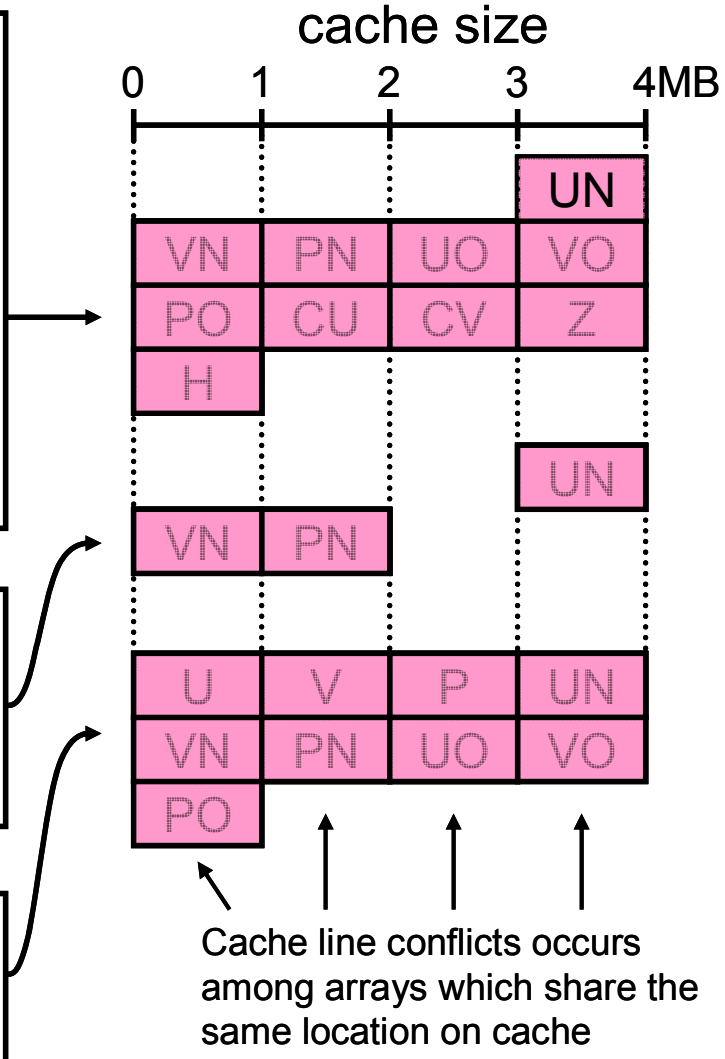

(b) Image of alignment of arrays on cache accessed by target loops

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)
```

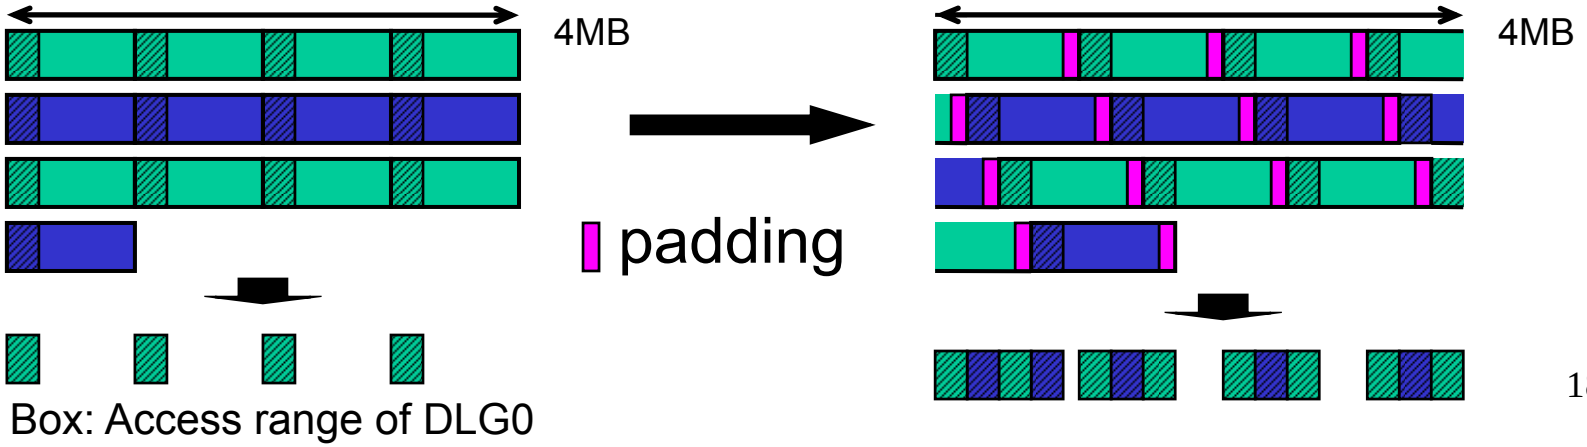

APC Compiler Organization

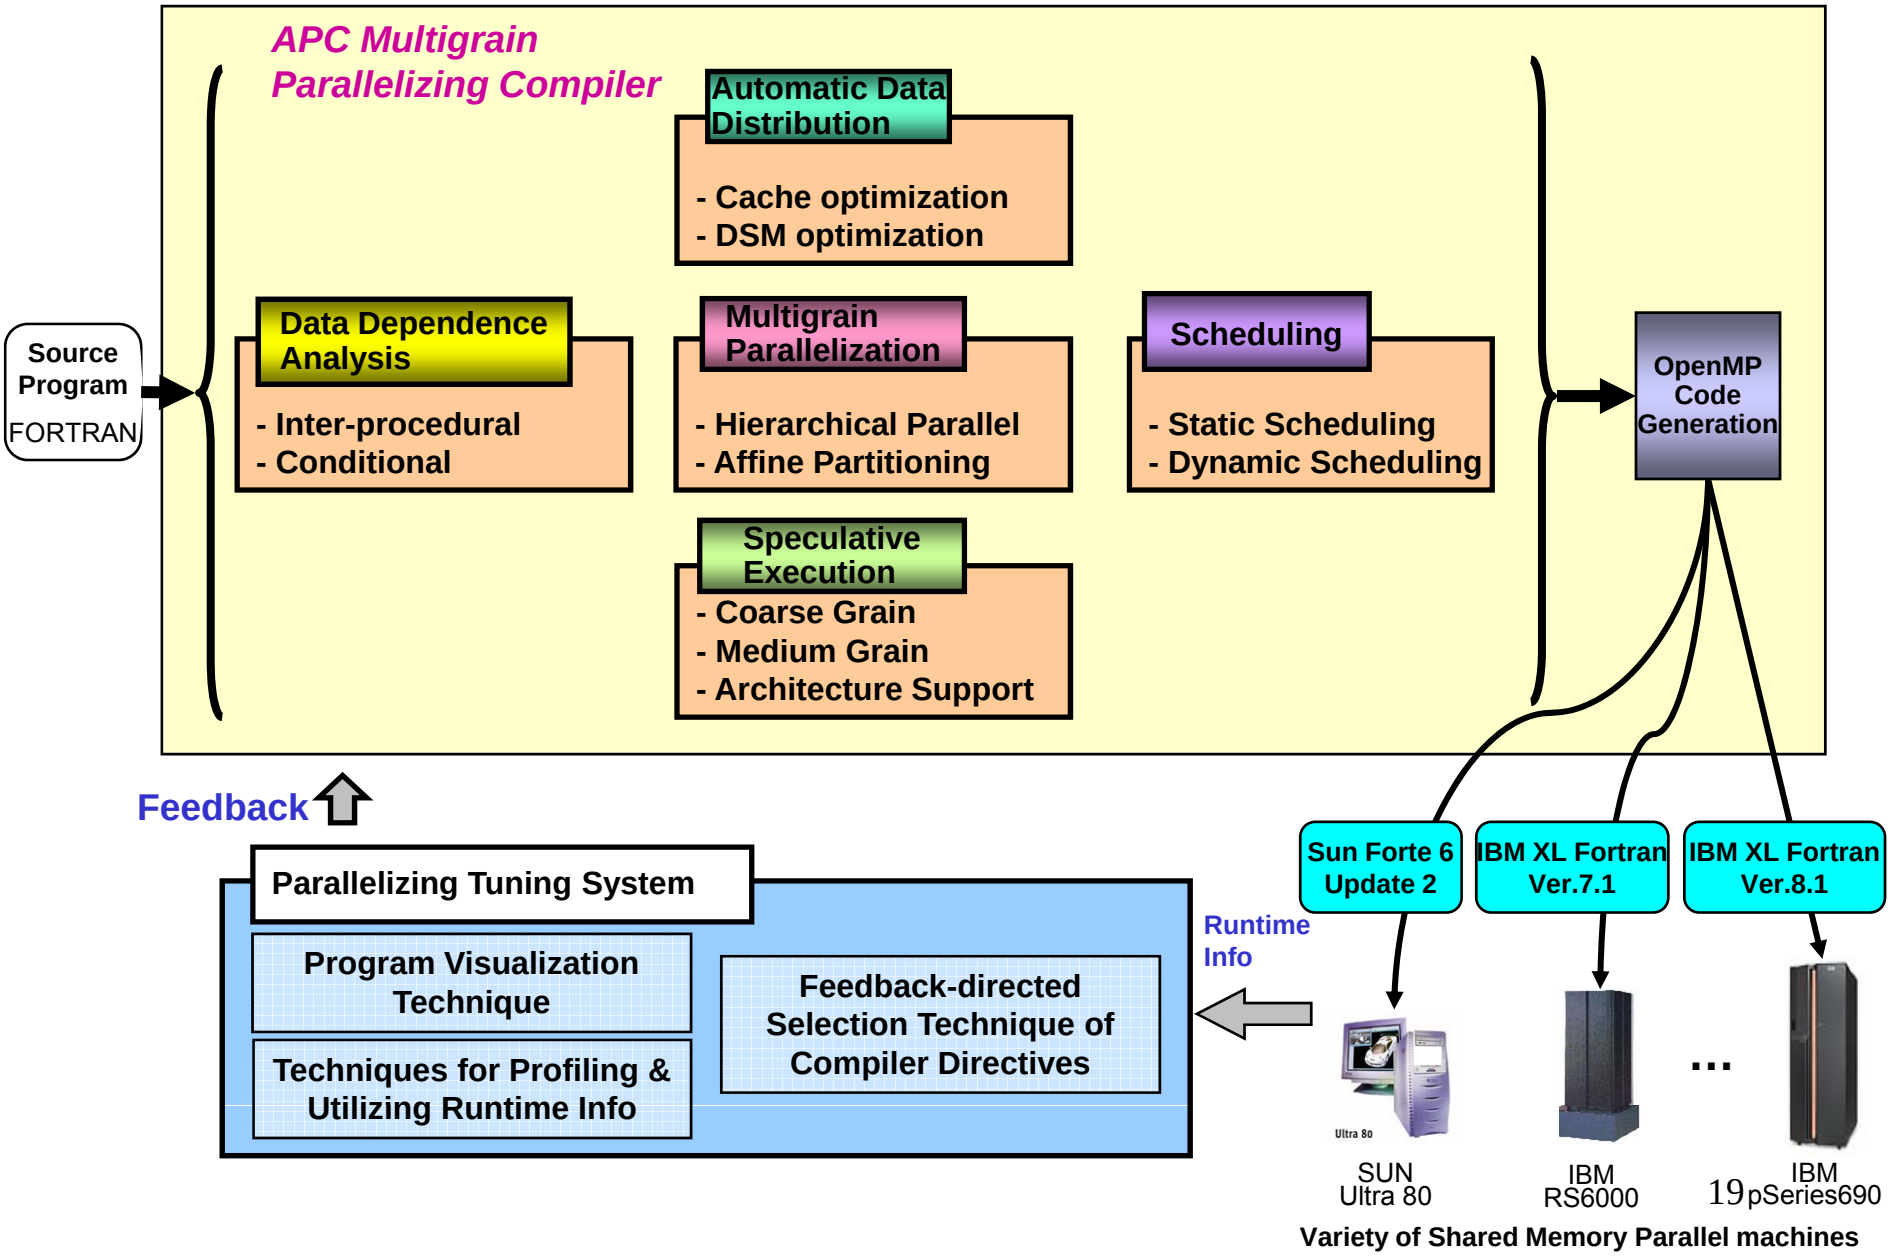

Image of Generated OpenMP Code for Hierarchical Multigrain Parallel Processing

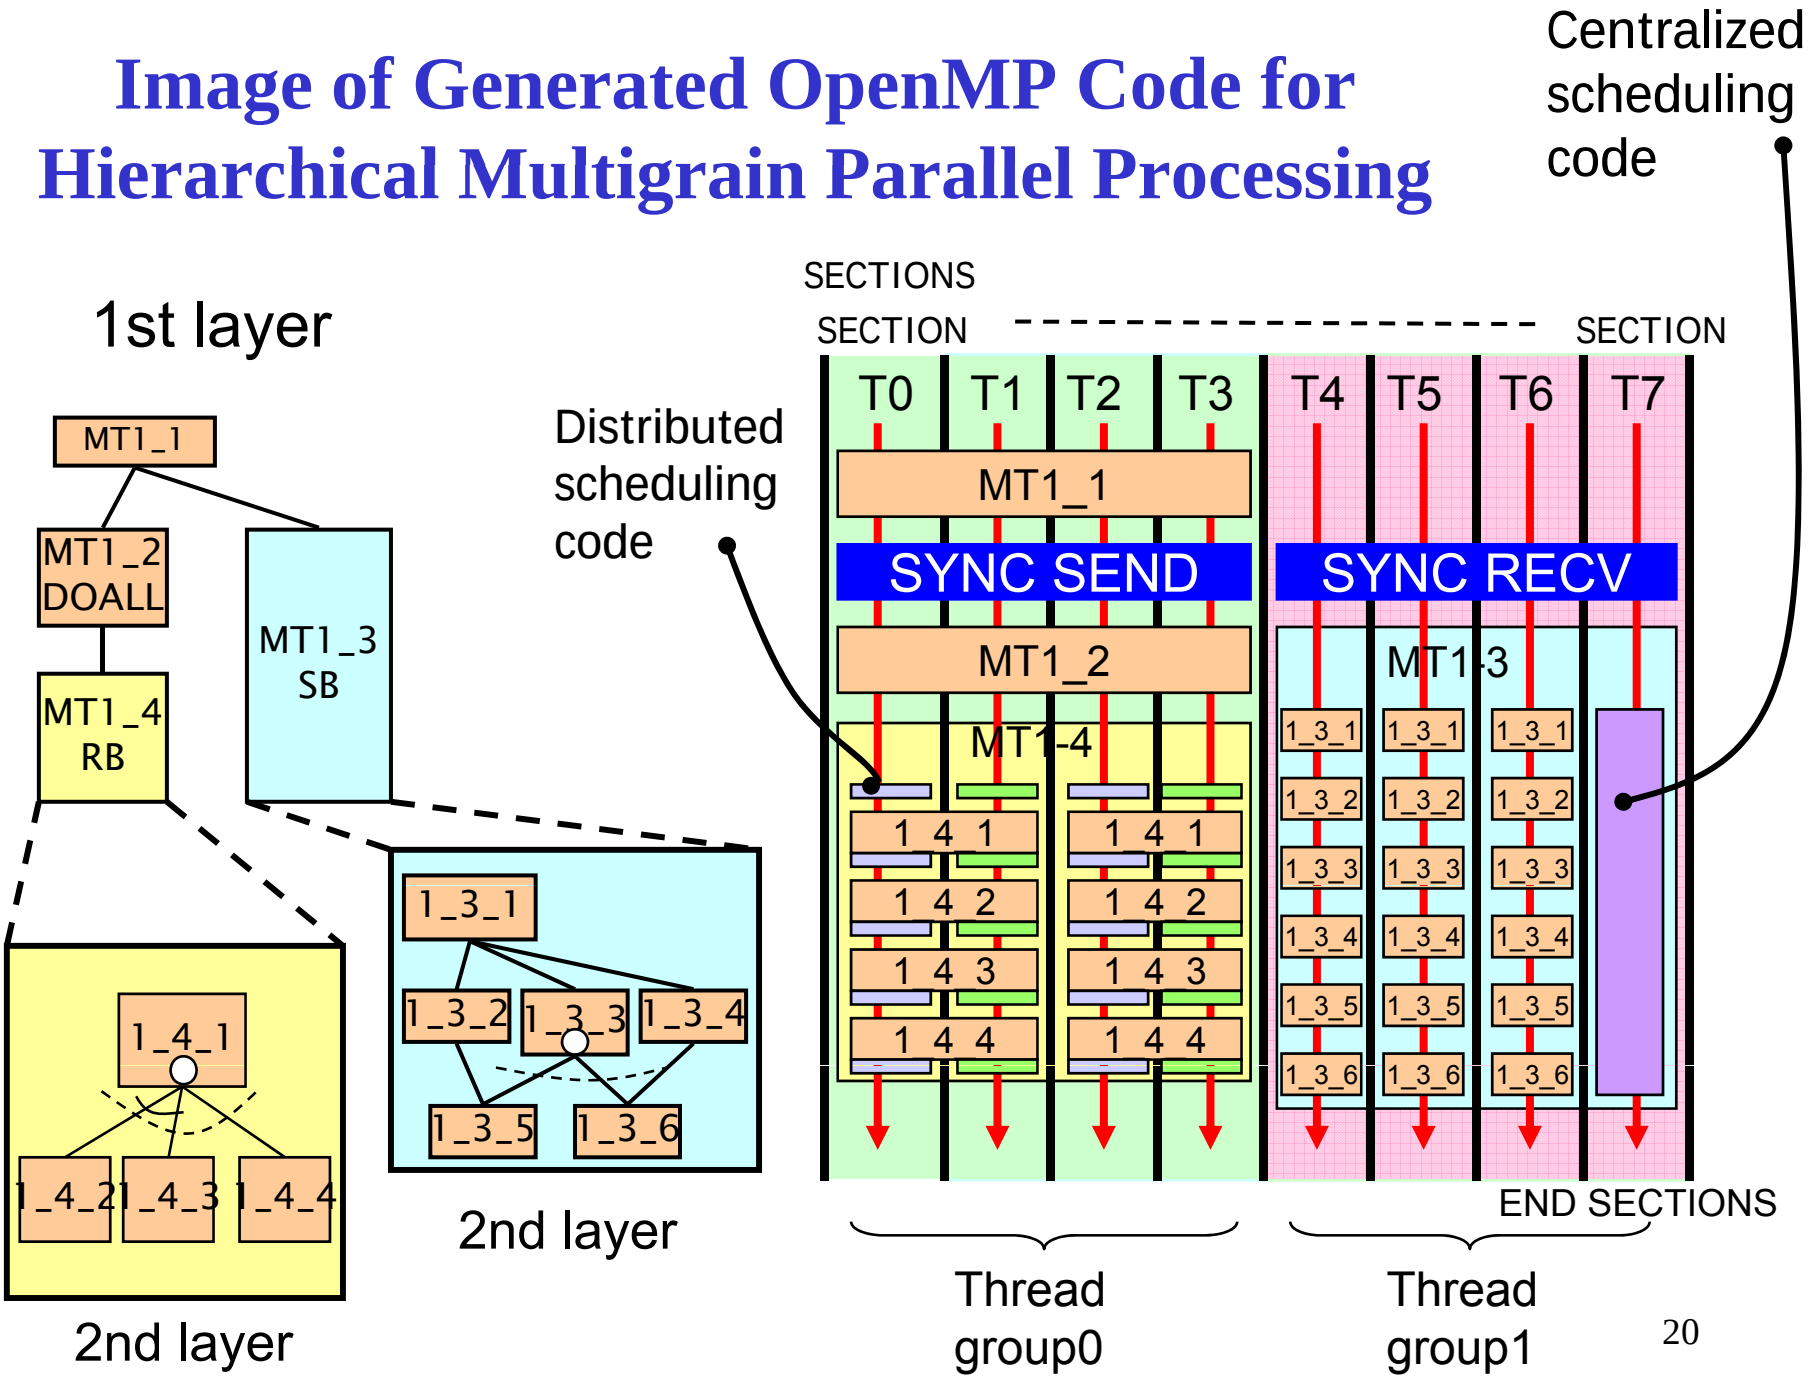

Performance OSCAR Multigrain Parallelizing Compiler on a IBM p550q 8core Deskside Server

- 2.7 times speedup against loop parallelizing compiler on 8 cores

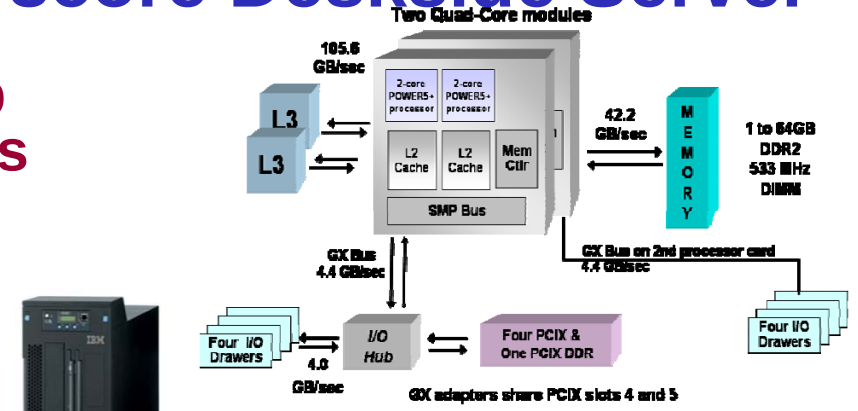

OSCAR Compiler Performance on 24 Processor IBM p690Highend SMP Server

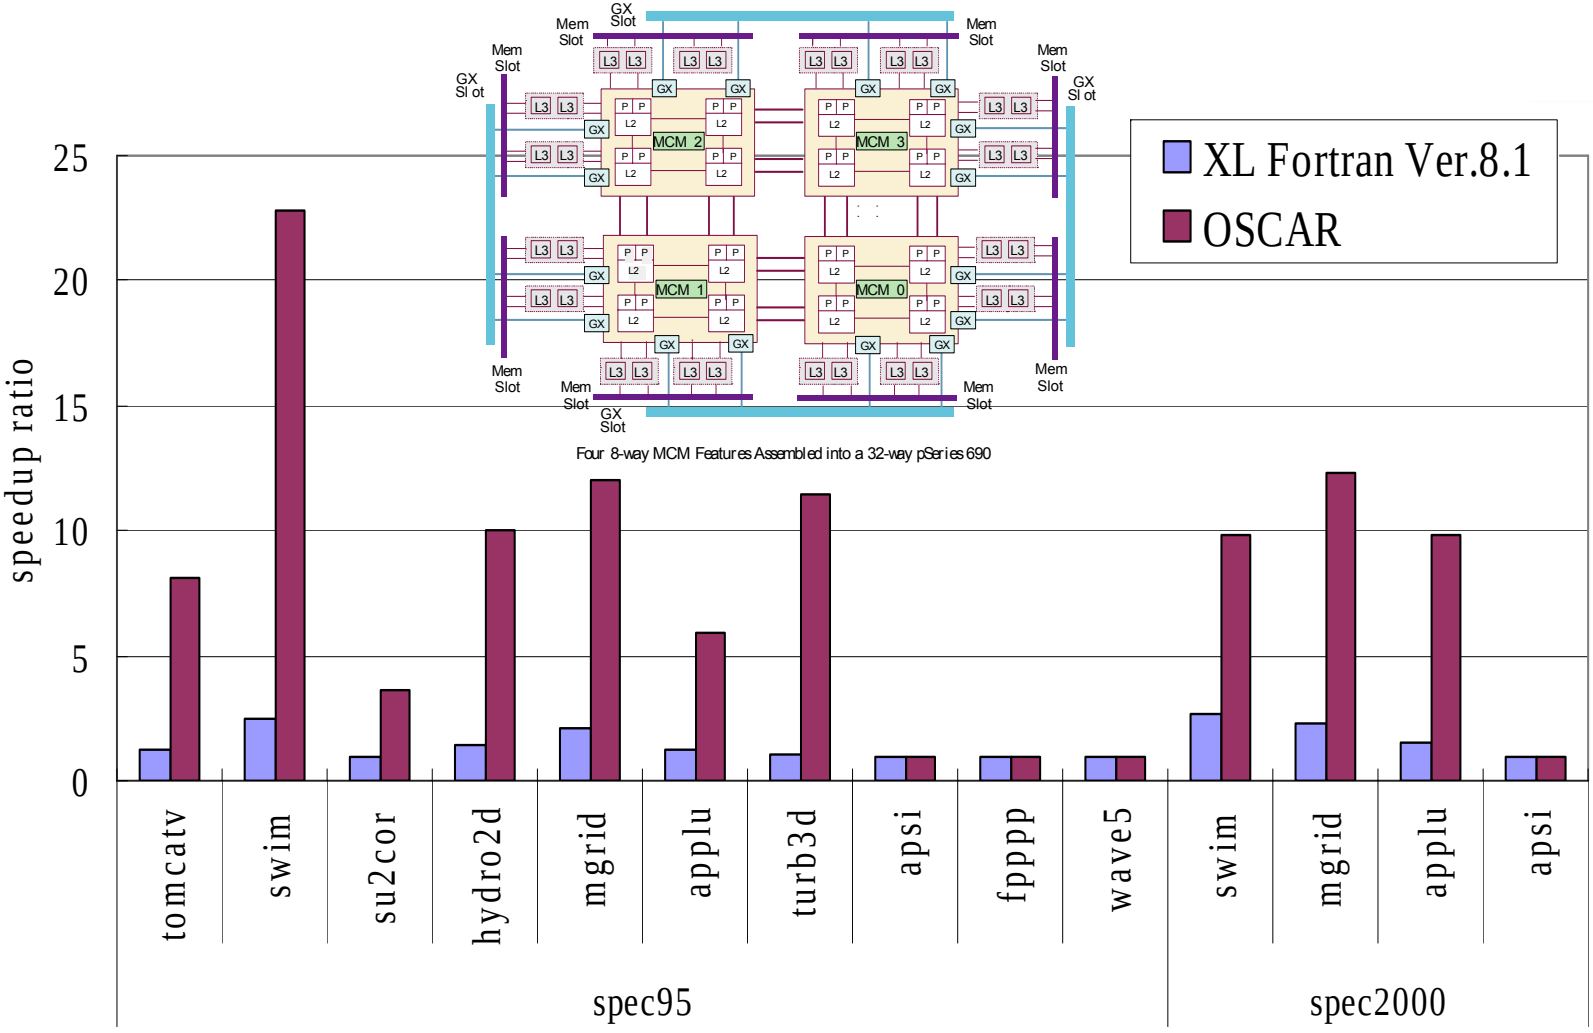

4.82 times speedup against loop parallelization

Performance on SGI Altix 450
Montecito 16 processors cc-NUMA server

Compiler options for the Intel Compiler:
for Automation parallelization, : -fast -parallel.
for OpenMP codes generated by OSCAR : -fast -openmp

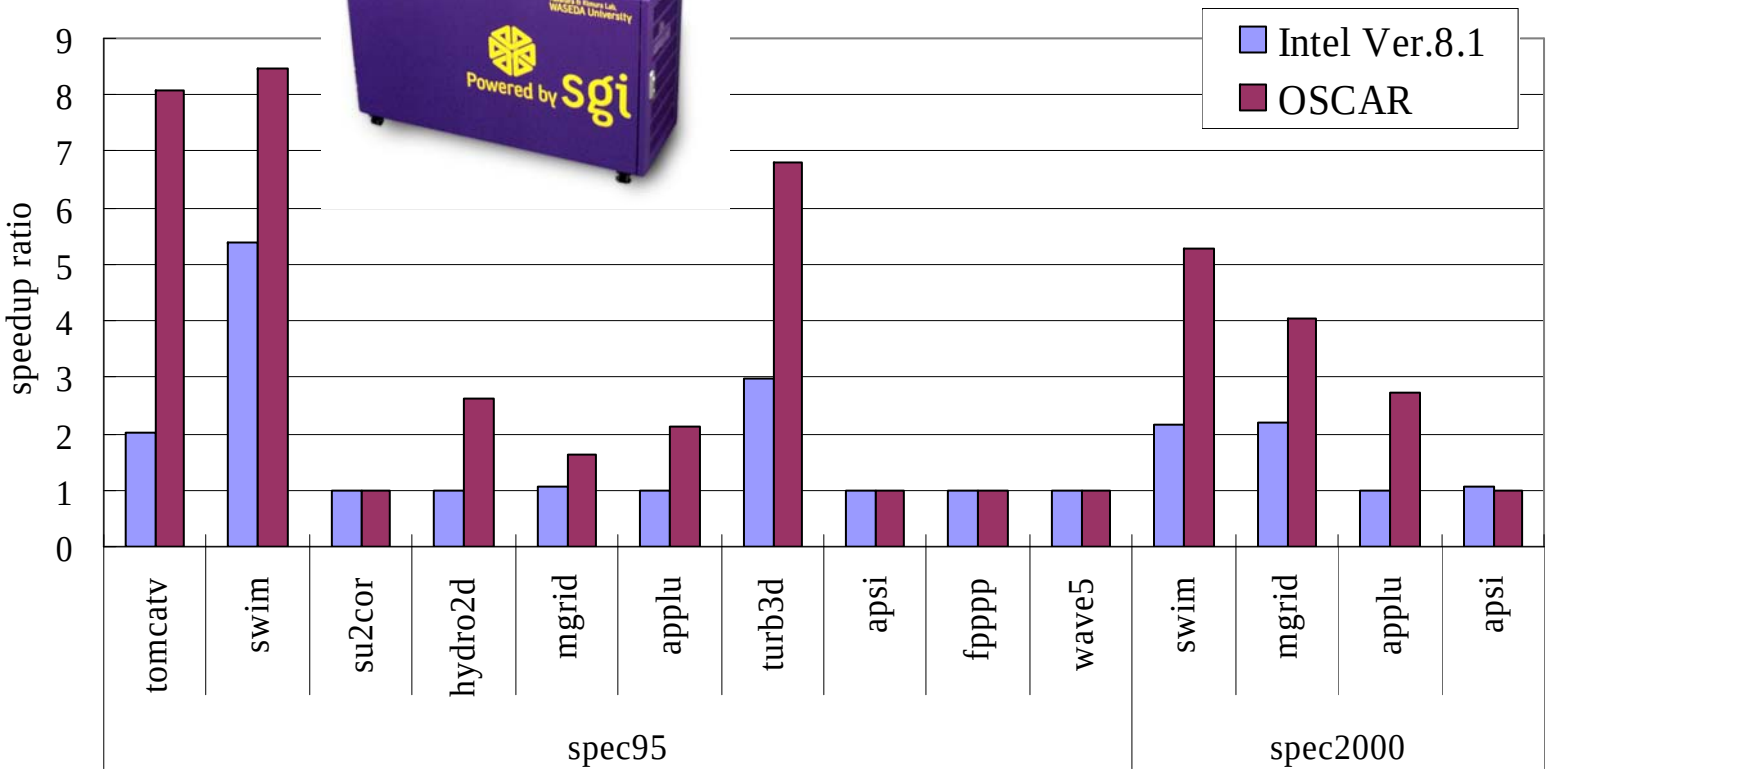

- **OSCAR compiler gave us 1.86 times speedup against Intel Fortran Itanium Compiler revision 8.1**

Performance of OSCAR compiler on 16 cores SGI Altix 450 Montvale server

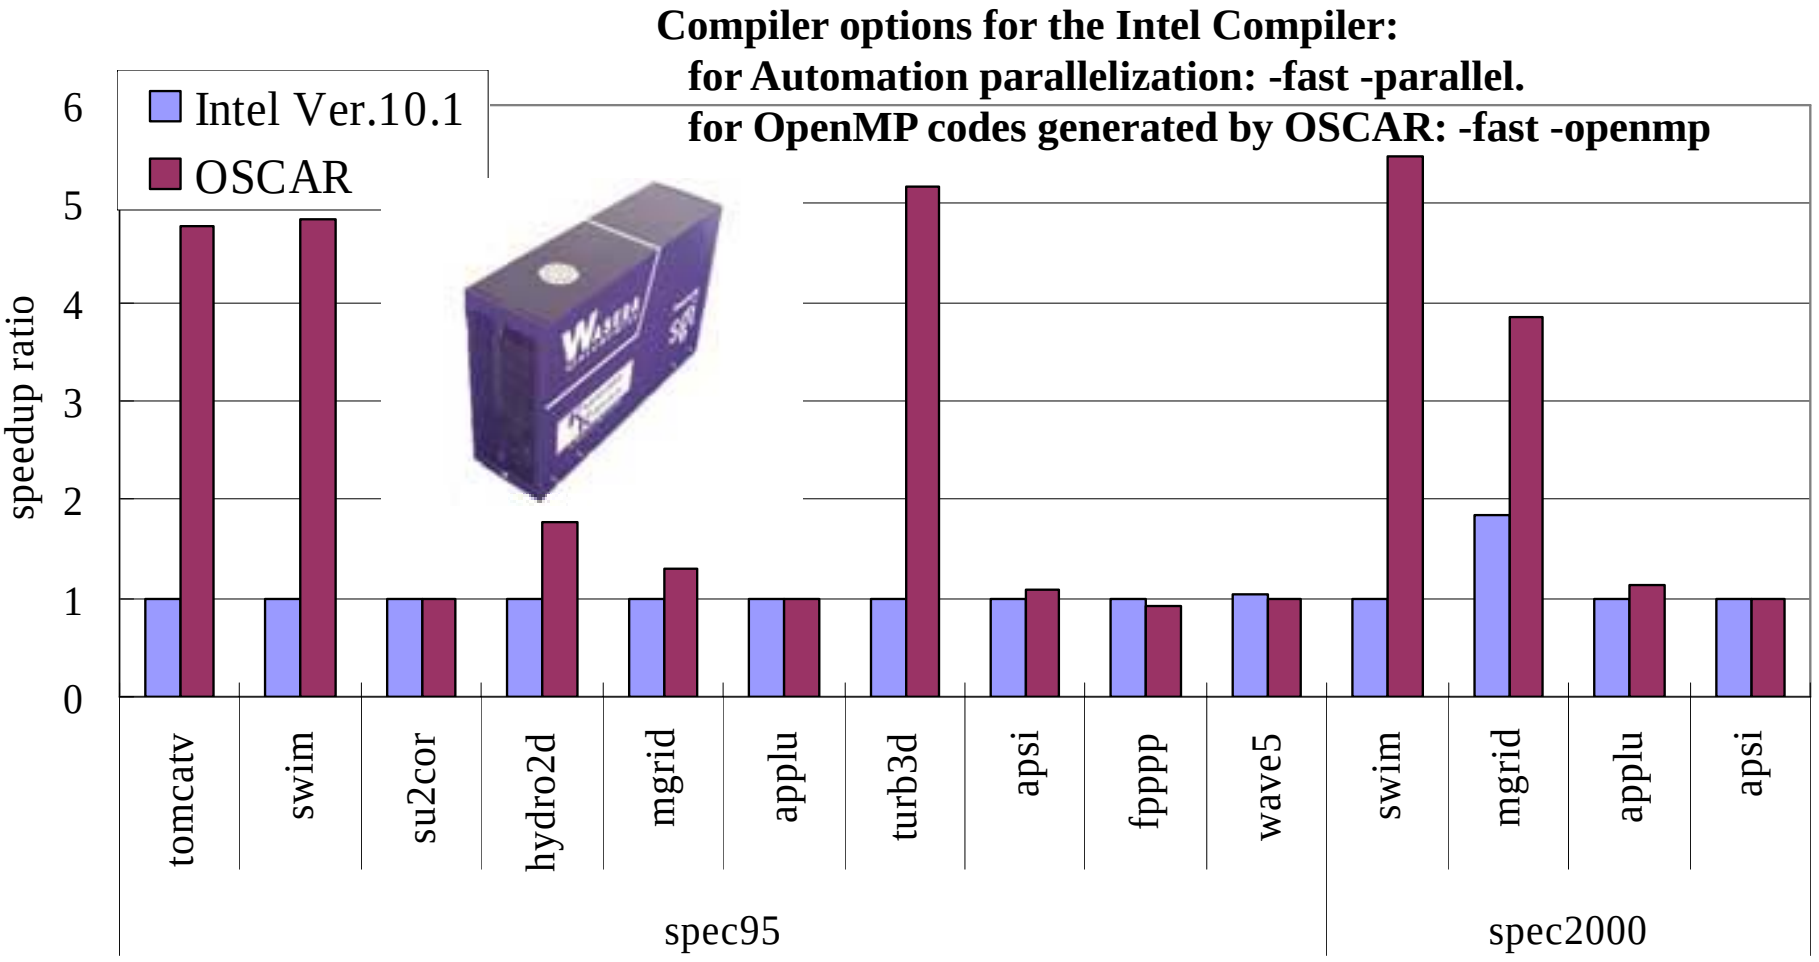

- **OSCAR compiler gave us 2.32 times speedup against Intel Fortran Itanium Compiler revision 10.1**

NEC/ARM MPCore Embedded 4 core SMP

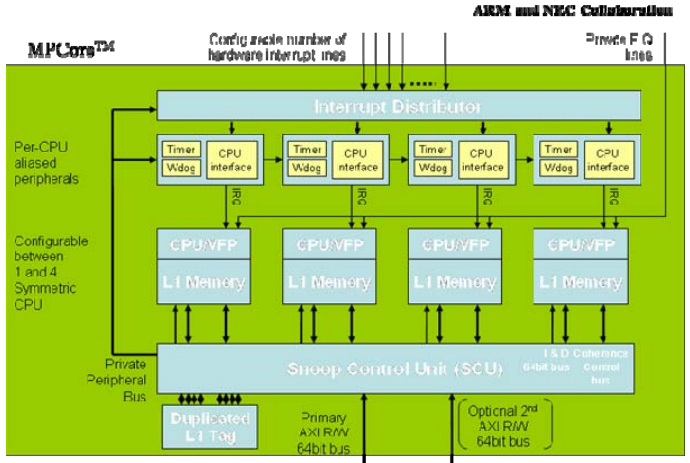

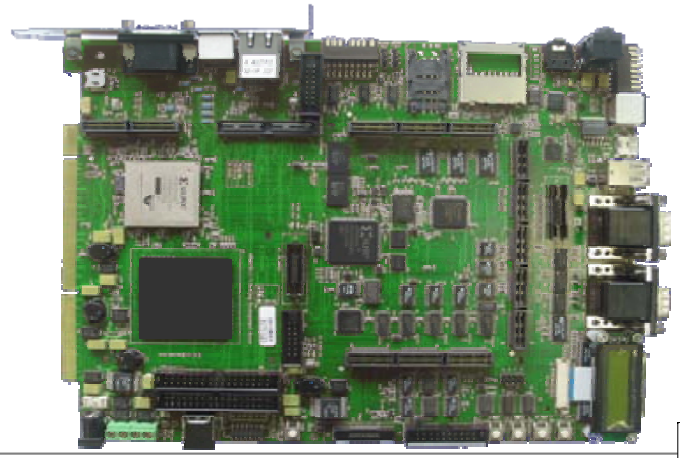

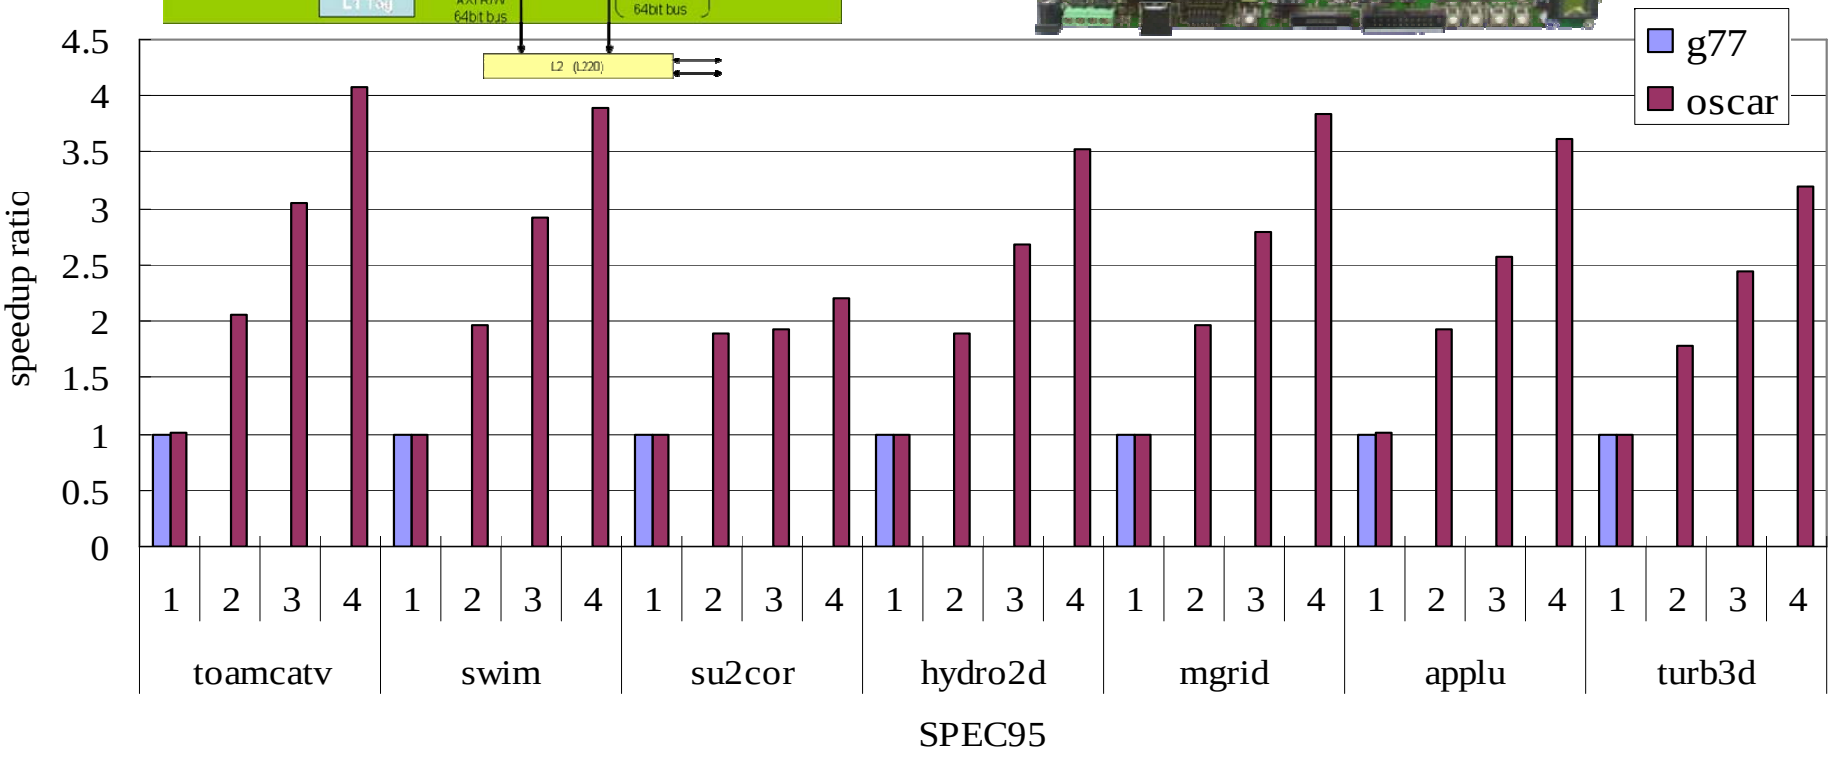

3.48 times speedup by OSCAR compiler against sequential processing

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

- Shortest execution time mode

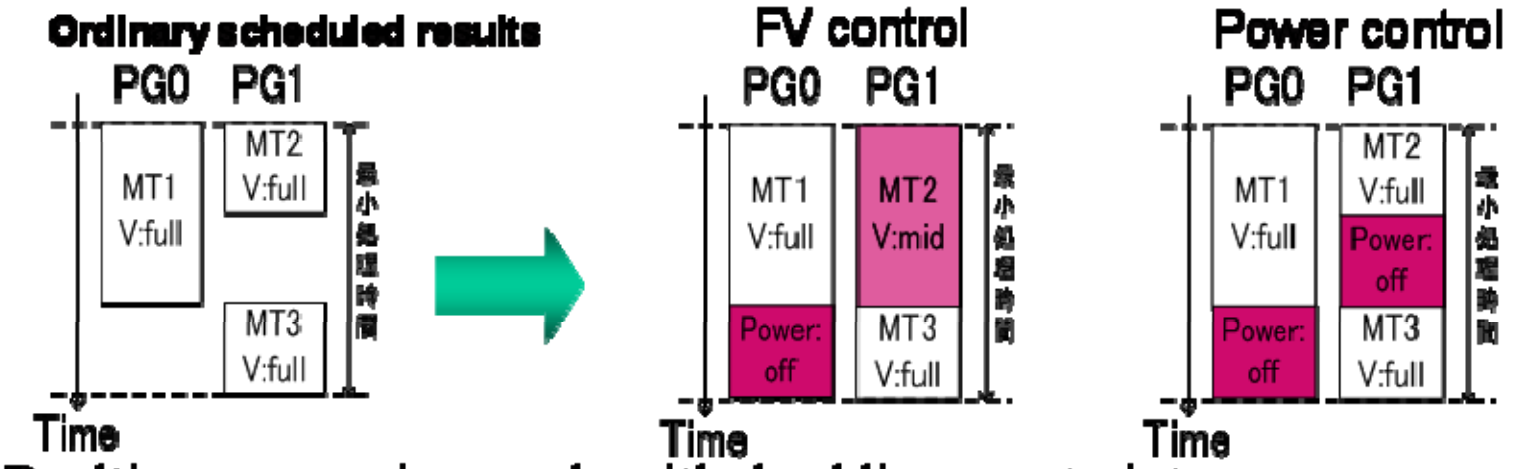

- Realtime processing mode with dead line constraints

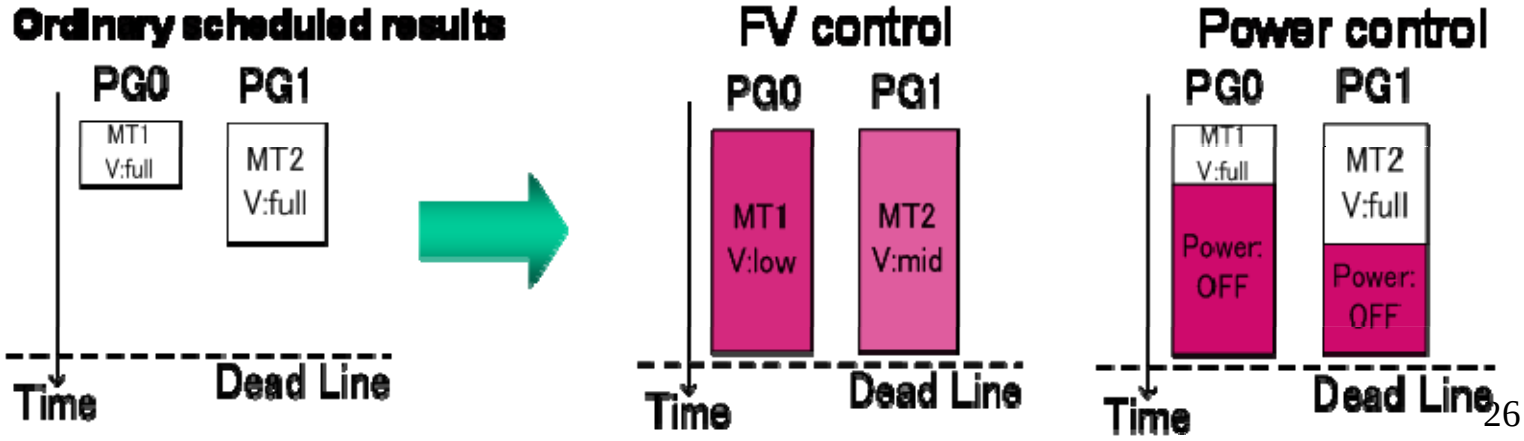

OSCAR Multi-Core Architecture

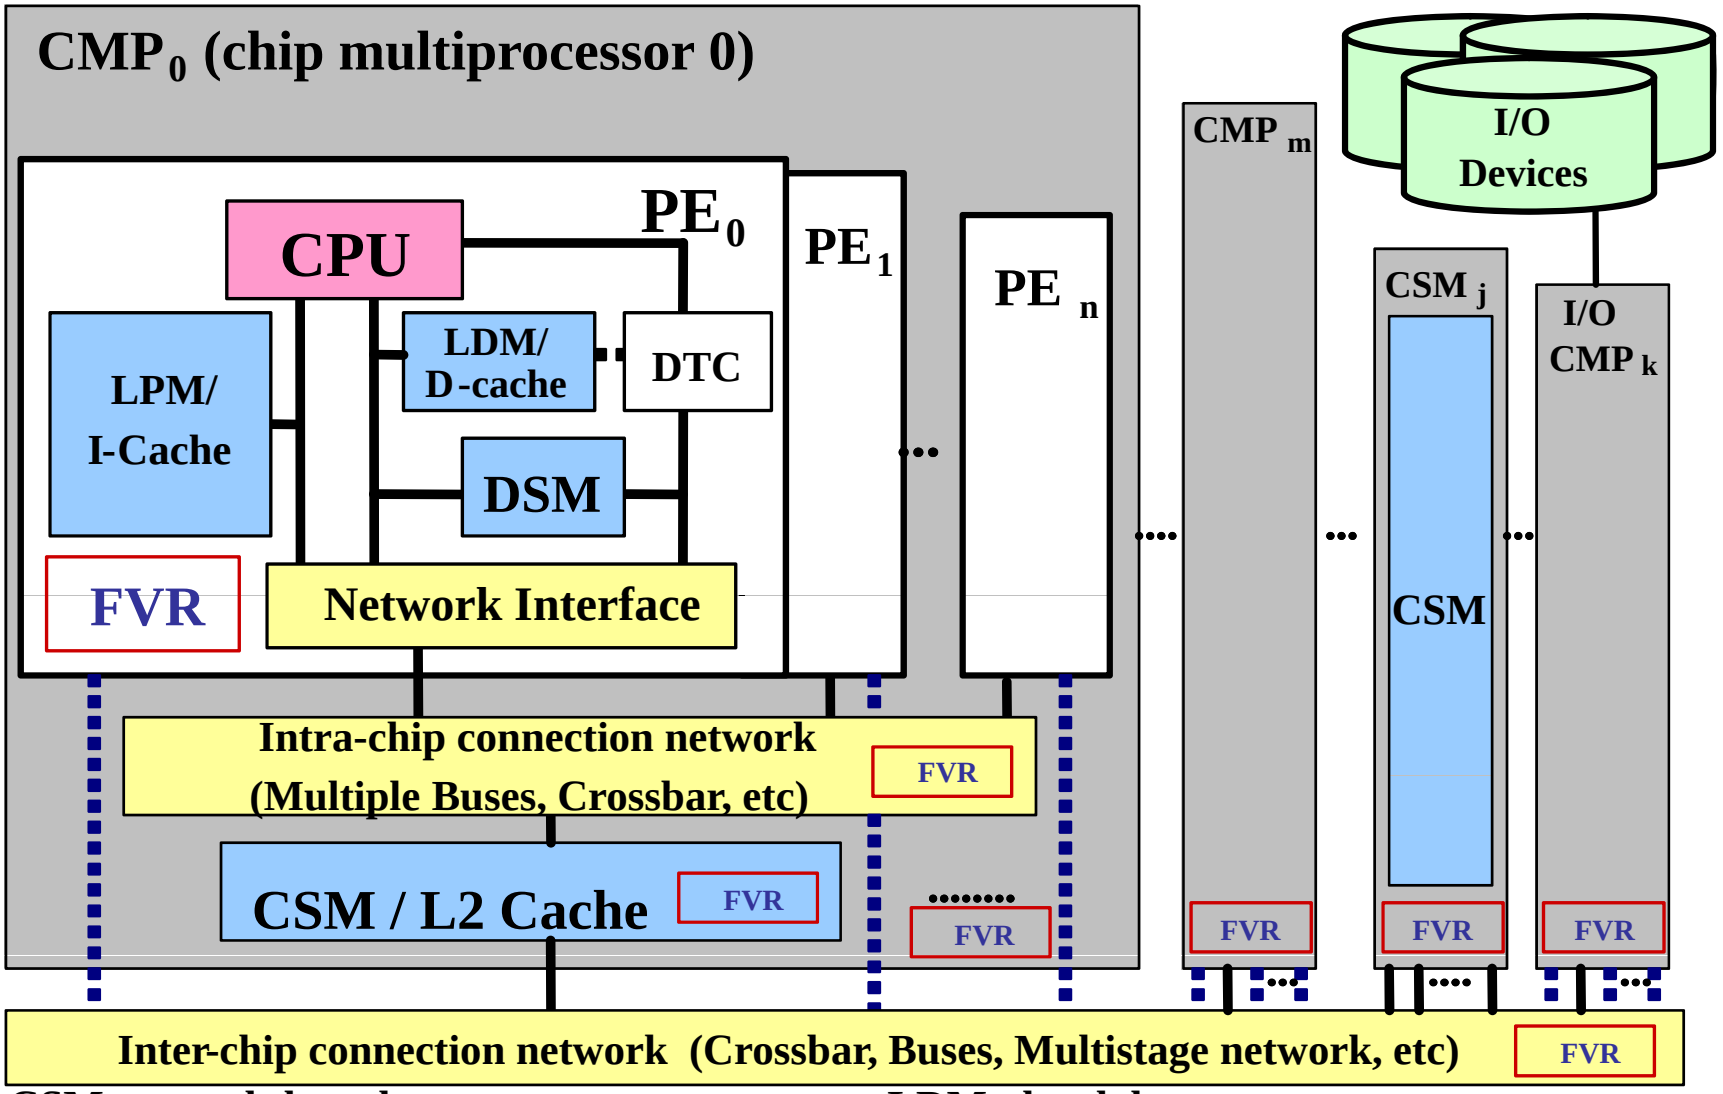

CSM: central shared mem.

DSM: distributed shared mem.

DTC: Data Transfer Controller

LDM : local data mem.

LPM : local program mem.

FVR: frequency / voltage control register

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**

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]

Speed-up in Fastest Execution Mode

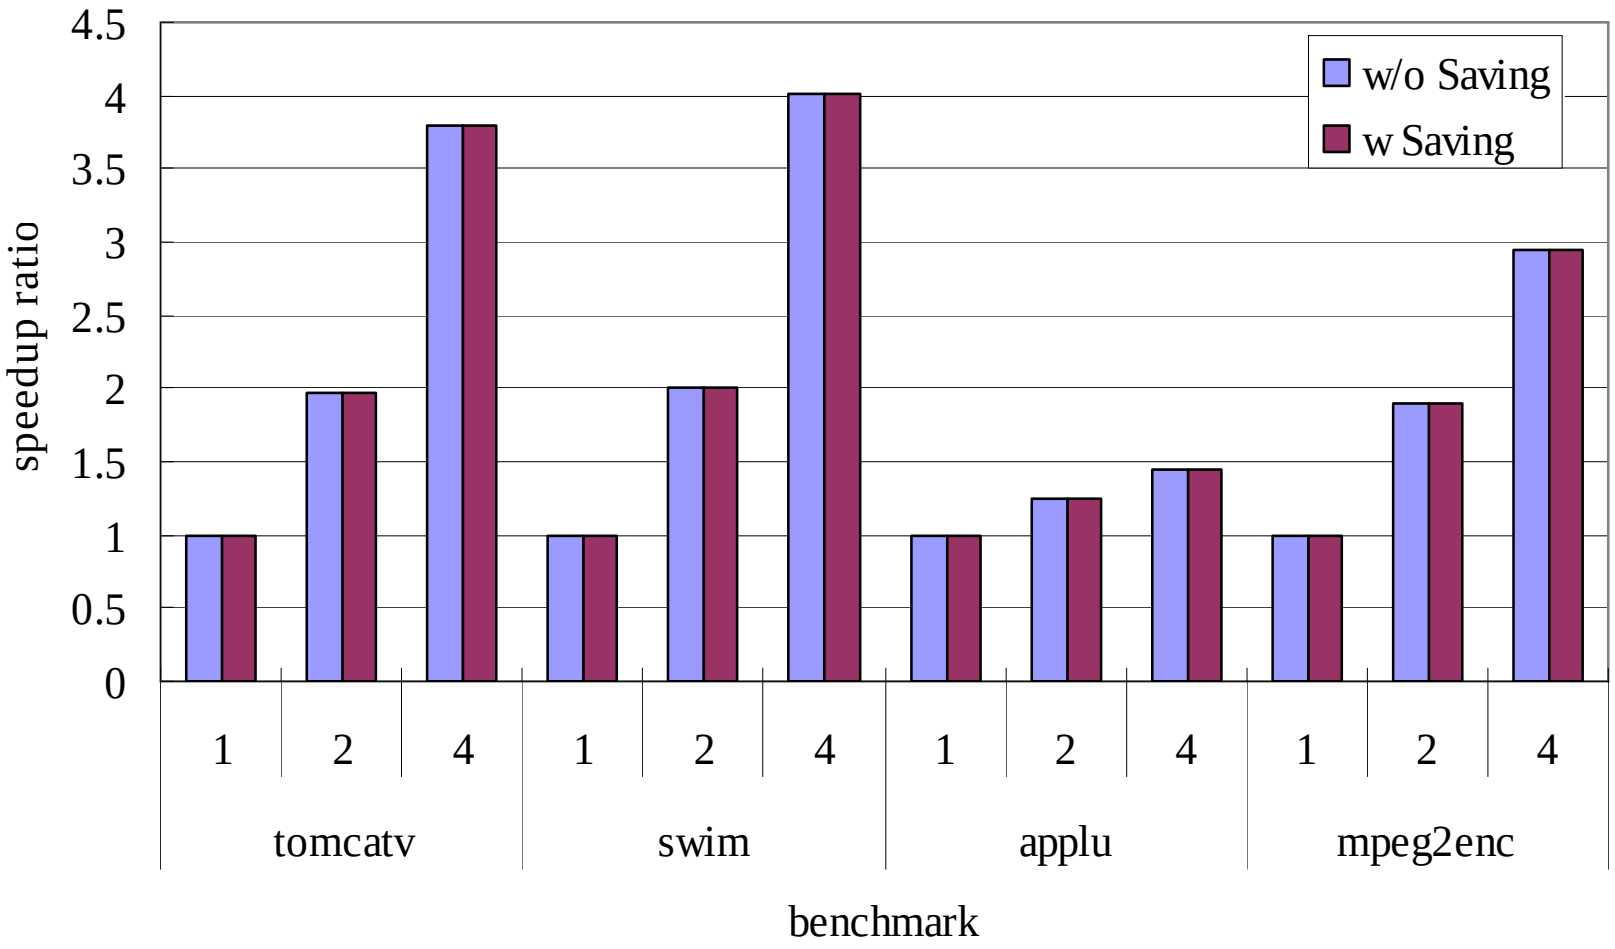

Consumed Energy in Fastest Execution Mode

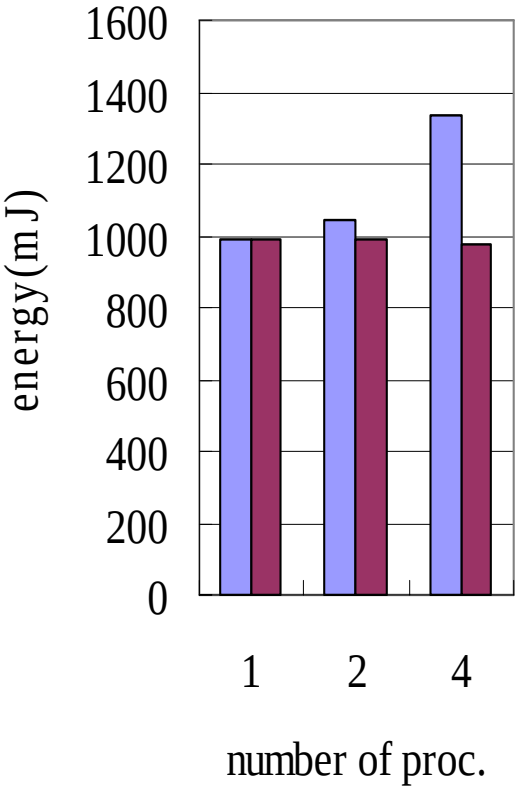

mpeg2_encode

Energy Reduction by OSCAR Compiler in Real-time Processing mode (10% Leak)

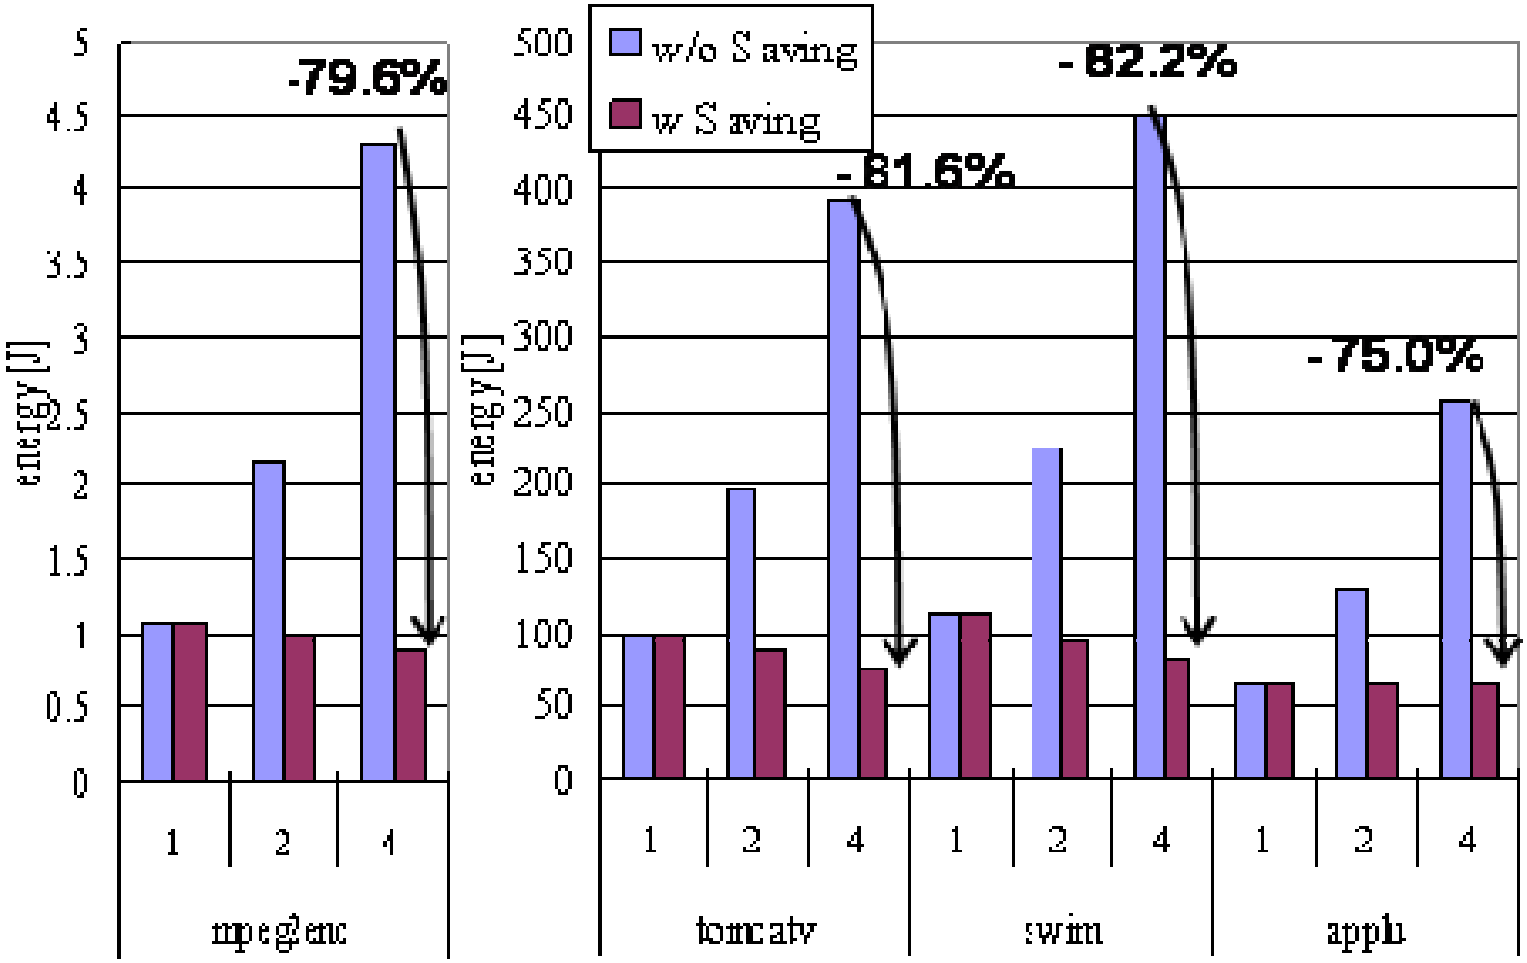

- deadline = sequential execution time, Leakage Power: 10%

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**

(2005.7～2008.3)**

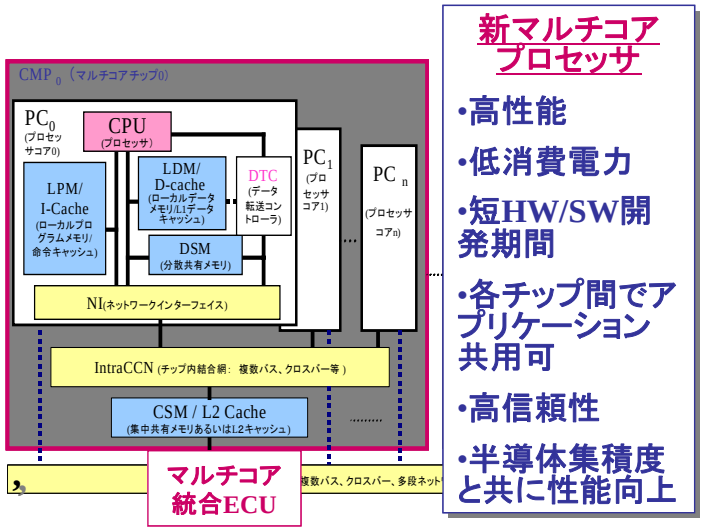

開発マルチコアチップは情報家電へ

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

NEDOMulticore Technology for Realtime Consumer Electronics R&D Organization(2005.7-2008.3)

Integrated R&D Steering Committee
Chair : Hironori Kasahara (Waseda Univ.), Project Leader
Sub-chair : Kunio Uchiyama (Hitachi), Project Sub-leader

Grant: Multicore Architecture and Compiler Research and Development
Project Leader :
Prof. Hironori Kasahara, Waseda Univ.

R&D Steering Comittee

Architecture & Compiler R&D Group.
Group Leader :
Associate Prof. Keiji Kimura, Waseda Univ.
(Outsourcing : Fujitsu)

Standard Multicore Architecture & API Committee
Chair : Prof. Hironori Kasahara, Waseda Univ.

Subsidy: Evaluation Environment for Multicore Technology
Project Sub-leader :
Dr. Kunio Uchiyama, Chief Researcher, Hitachi

Test Chip Development Group.
Group Leader :
Renesas Technology

Evaluation Sytem Development Group
Group Leader :
Hitachi

Hitachi, Fujitsu, Toshiba, NEC, Panasonic, Renesas Technology

Fujitsu FR-1000Multicore Processor

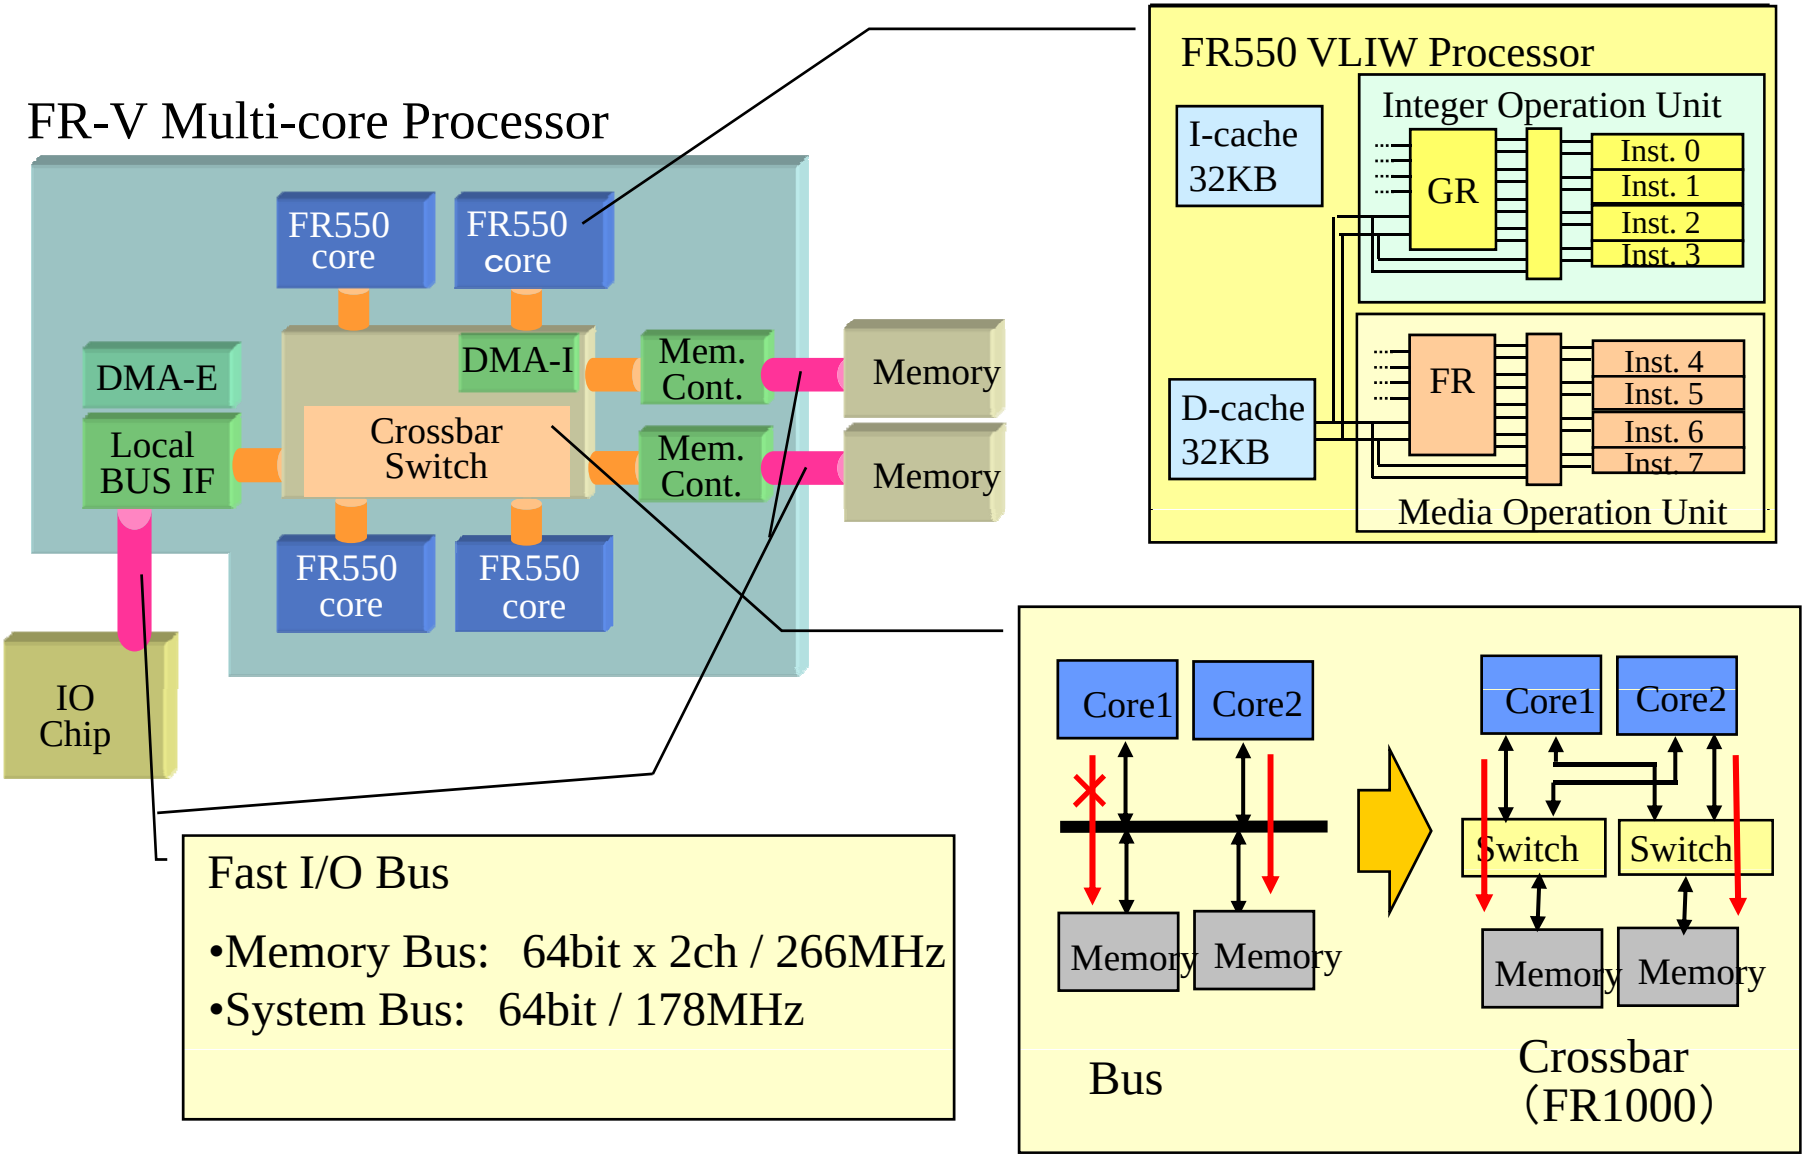

Fast I/O Bus

- Memory Bus: 64bit x 2ch / 266MHz
- System Bus: 64bit / 178MHz

Panasonic UniPhier

Scalable media processing architecture

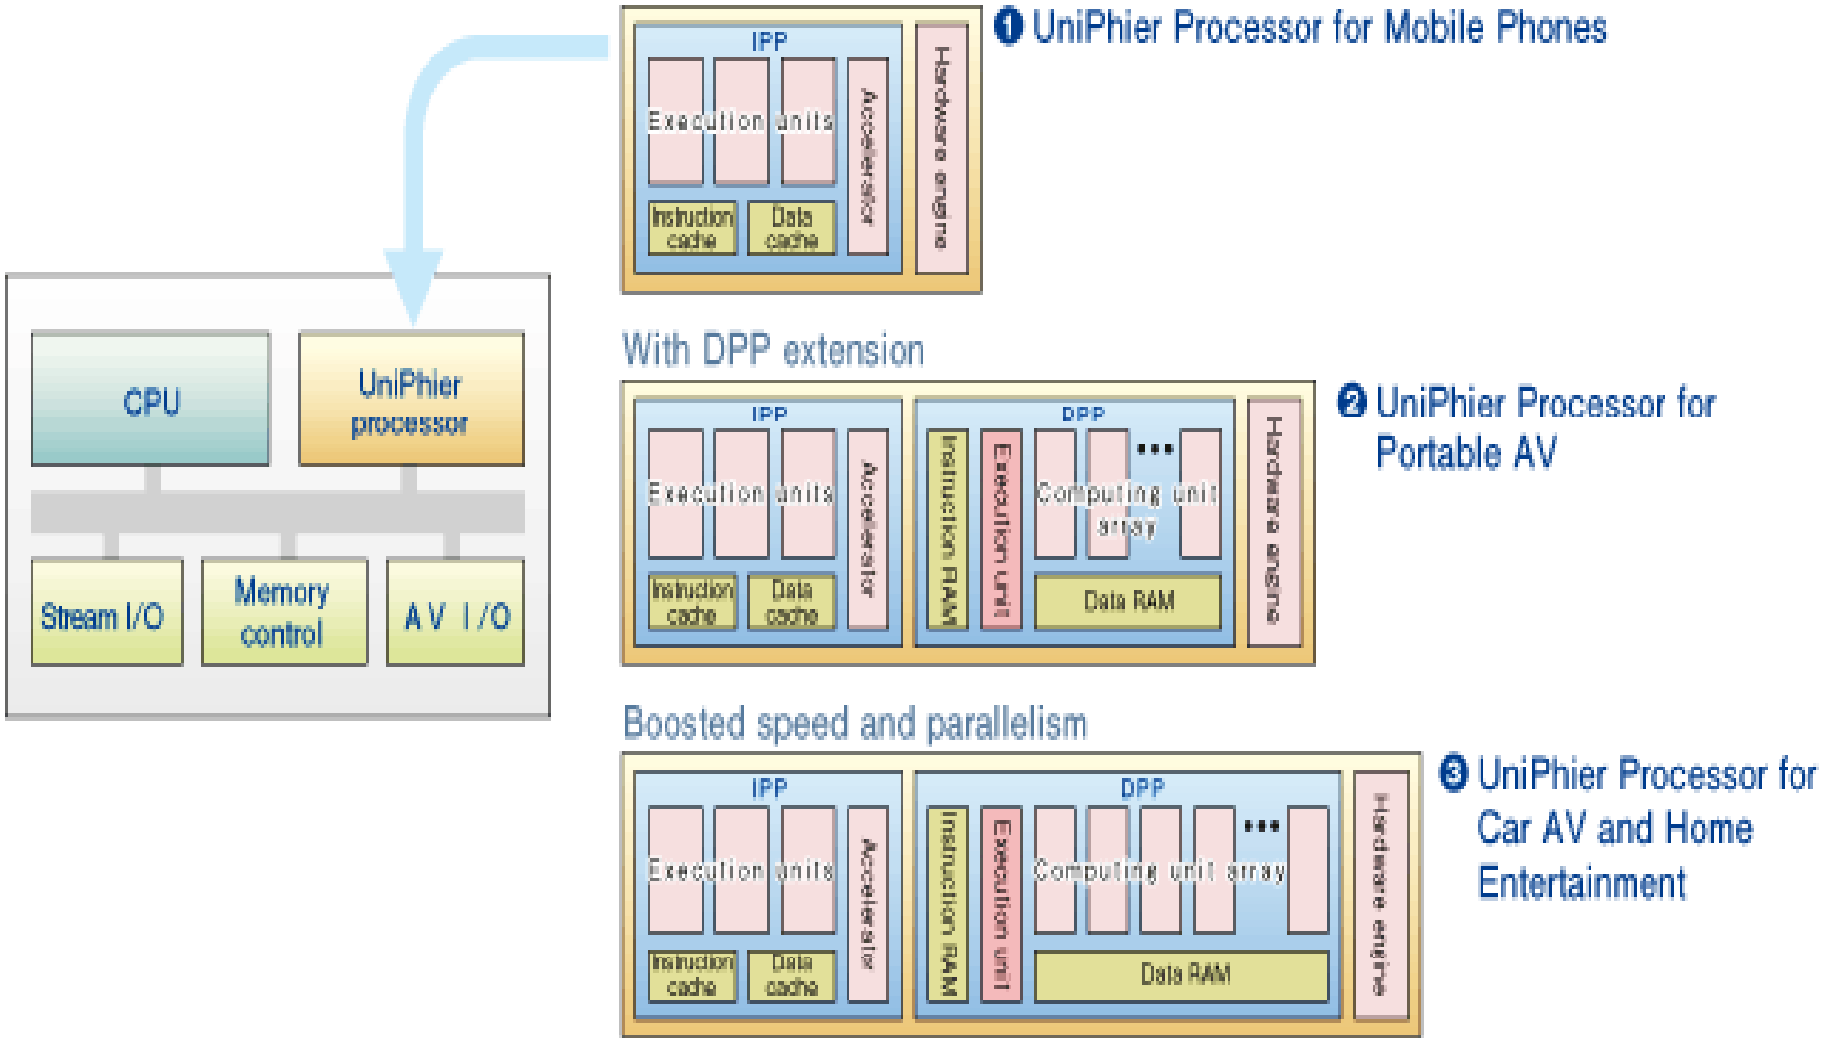

CELL Processor Overview

- **Power Processor Element (PPE)**
  - **PowerCore processes OS and Control tasks**
  - **2-way Multi-threaded**
- **Synergistic Processor Element (SPE)**
  - **8 SPE offers high performance**
  - **Dual issue RISC Architecture**
  - **128bit SIMD(16 - way)**
  - **128 x 128bit General Registers**
  - **256KB Local Store**
  - **DedicatedDMA engines**

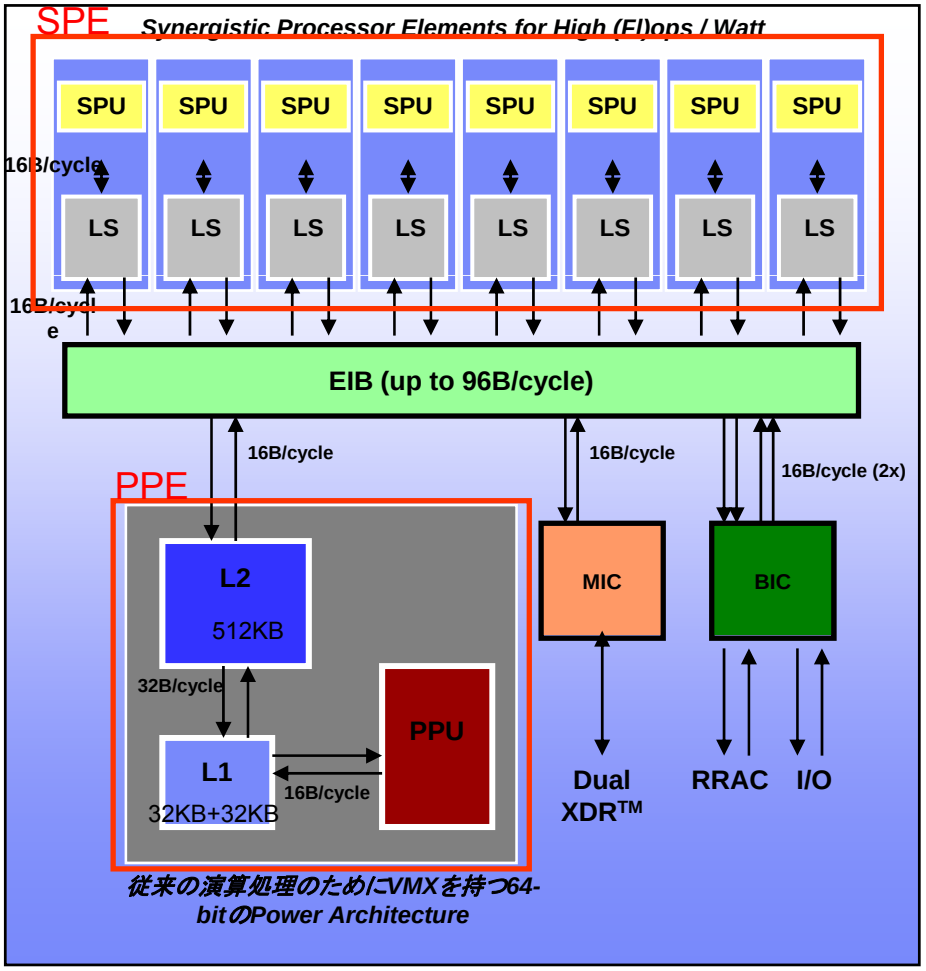

1987 OSCAR(Optimally Scheduled Advanced Multiprocessor)

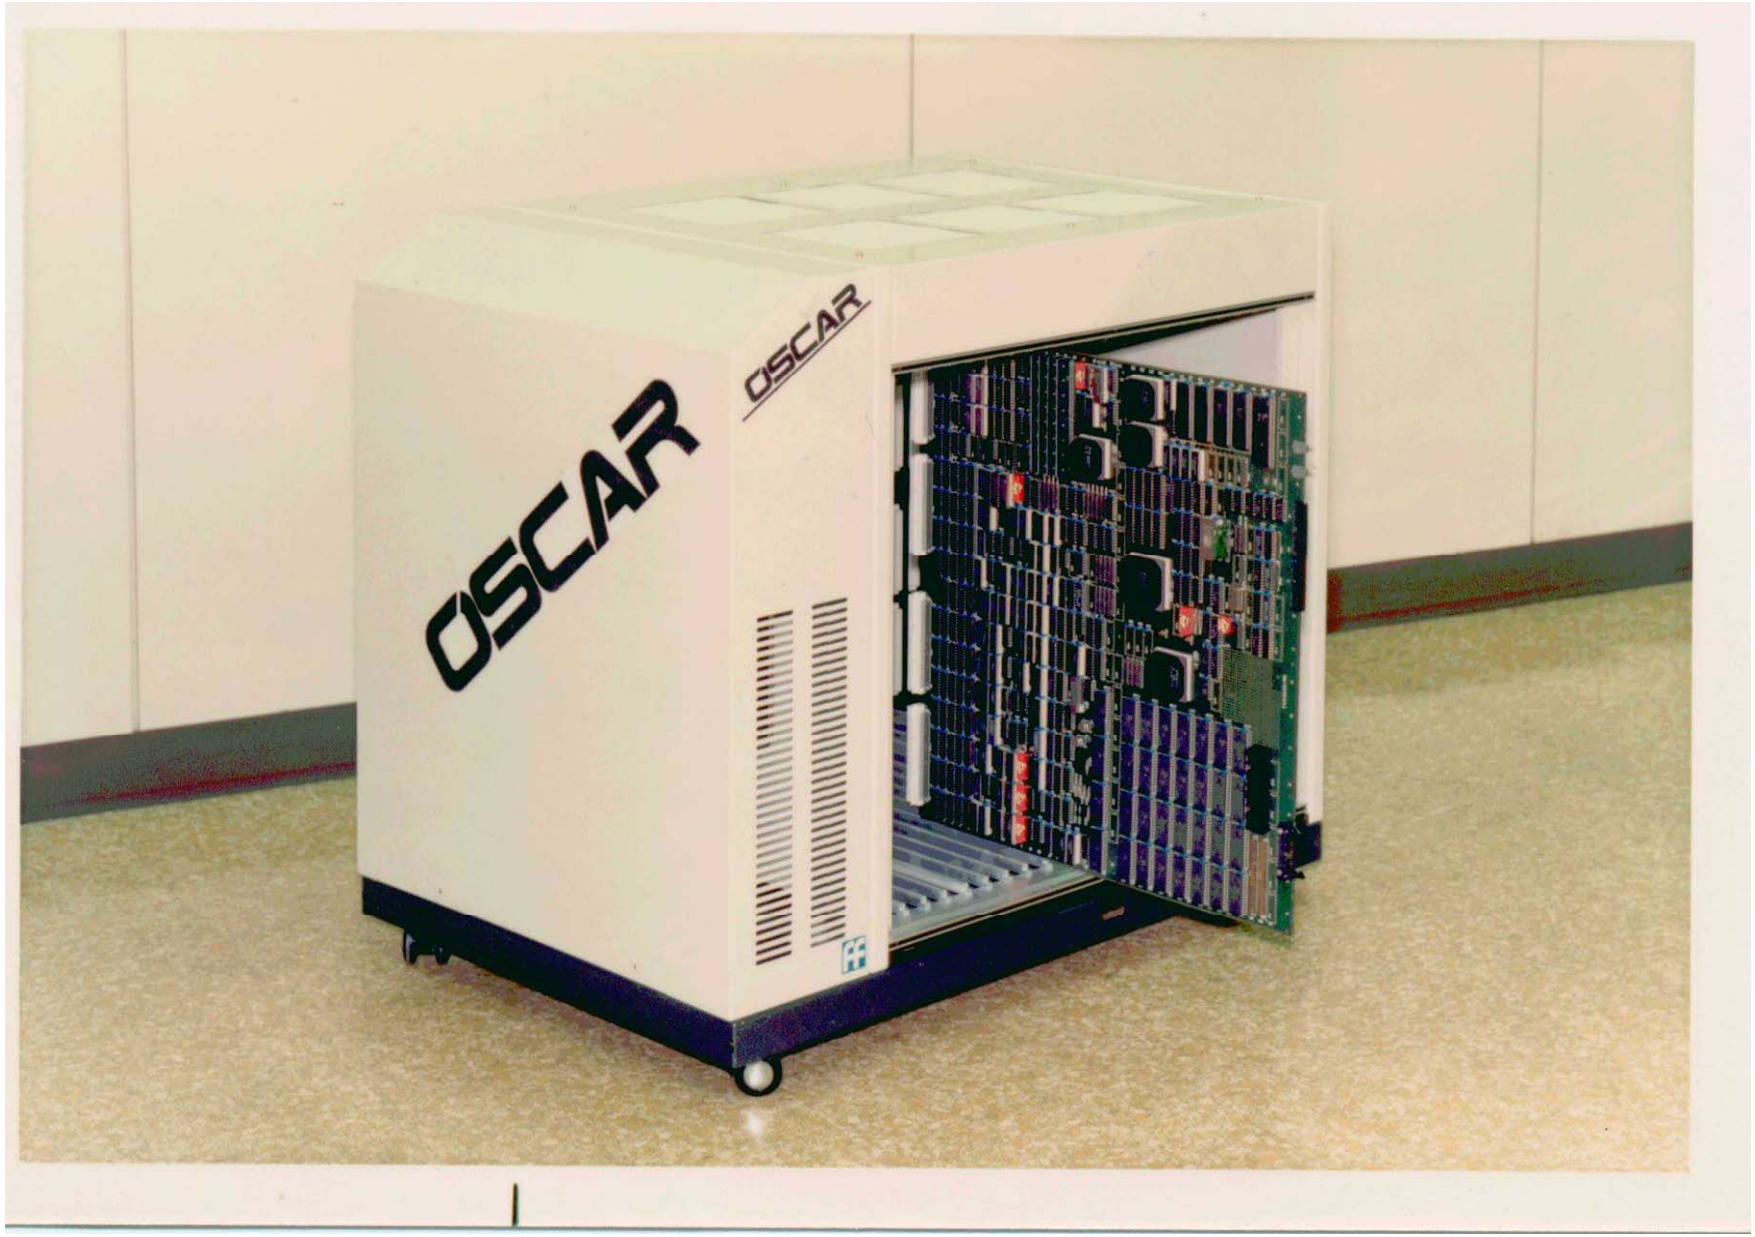

OSCAR(Optimally Scheduled Advanced Multiprocessor)

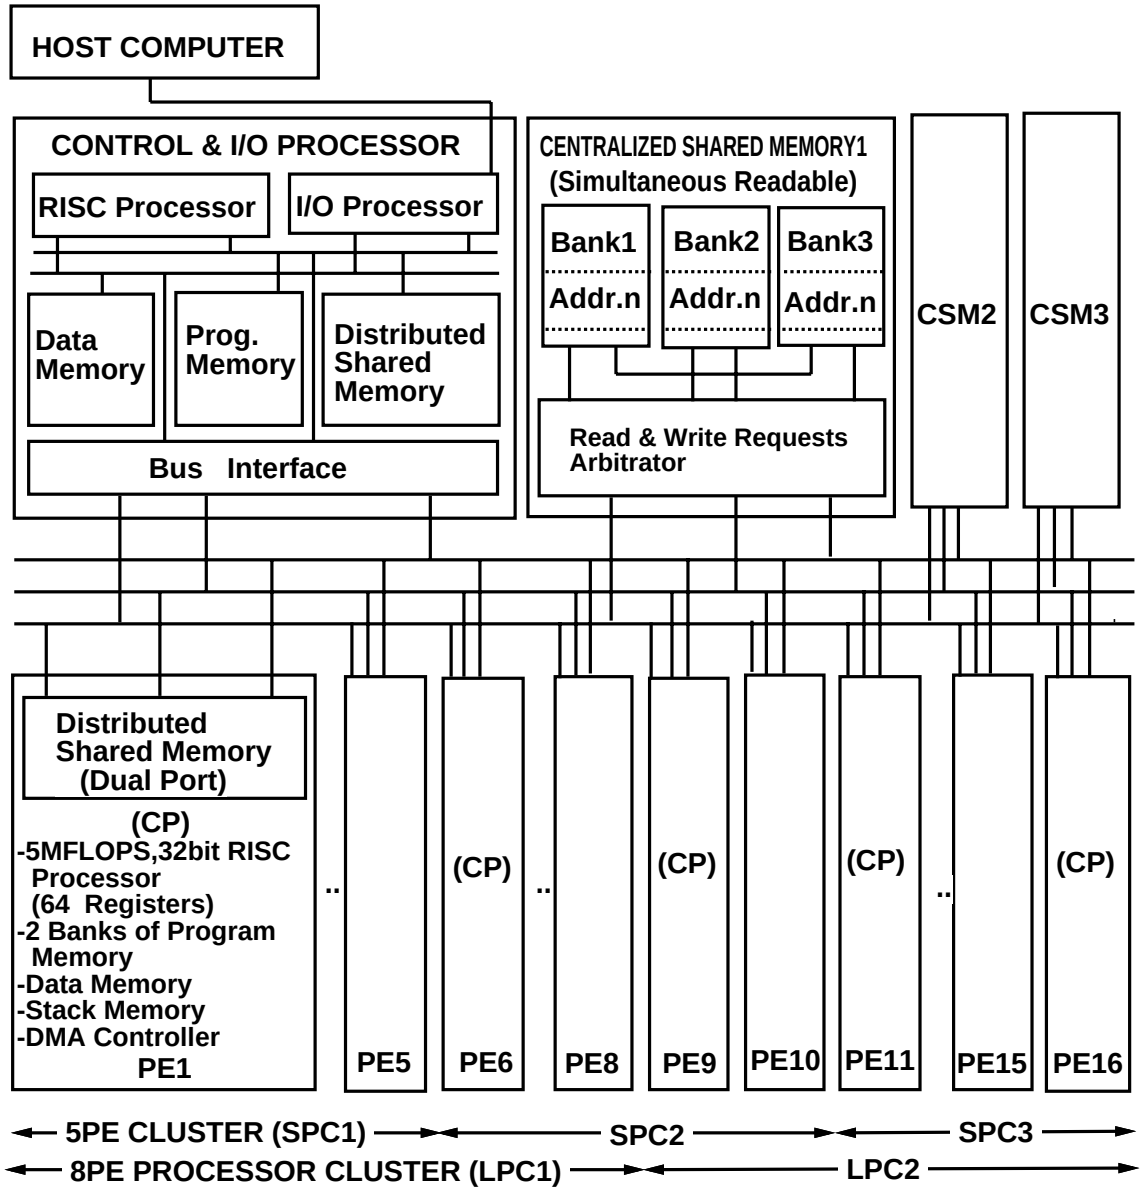

OSCAR PE (Processor Element)

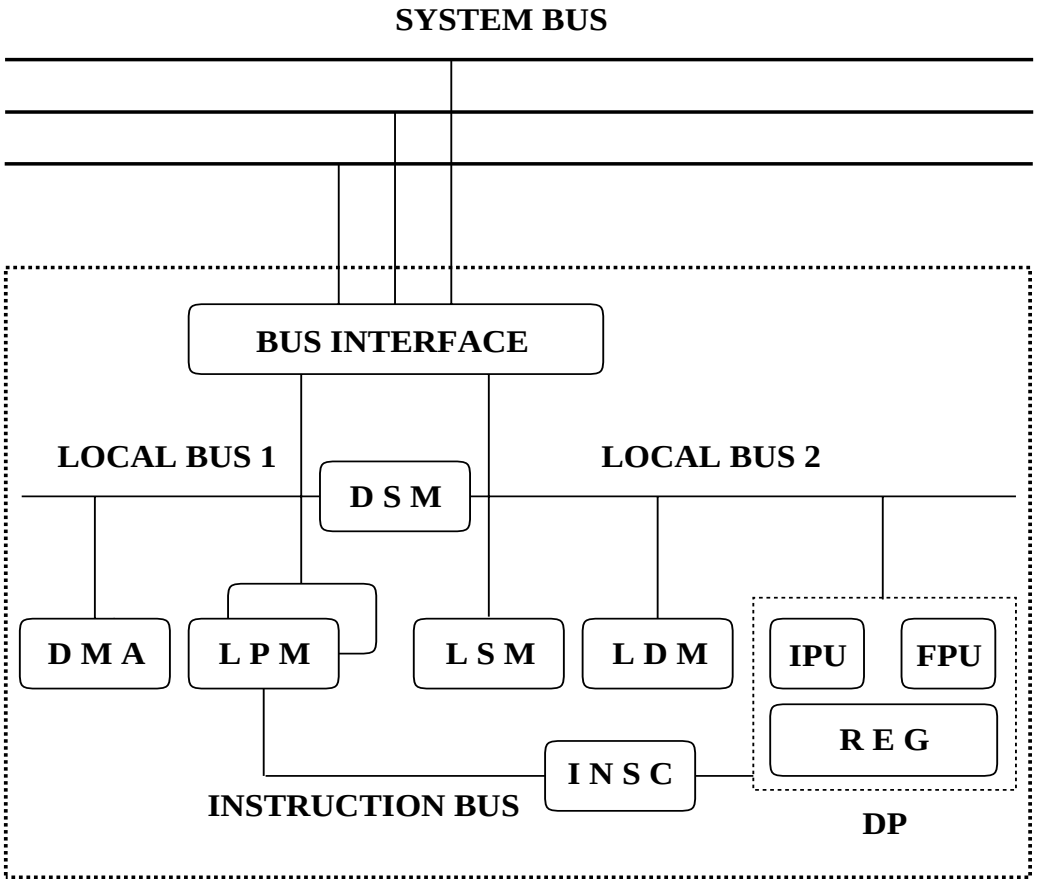

D M A : DMA CONTROLLER
L P M : LOCAL PROGRAM MEMORY (128KW * 2BANK)
I N S C : INSTRUCTION CONTROL UNIT
D S M : DISTRIBUTED SHARED MEMORY (2KW)
L S M : LOCAL STACK MEMORY (4KW)
L D M : LOCAL DATA MEMORY (256KW)
D P : DATA PATH
I P U : INTEGER PROCESSING UNIT
F P U : FLOATING PROCESSING UNIT
R E G : REGISTER FILE (64 REGISTERS)

1987 OSCAR PE Board

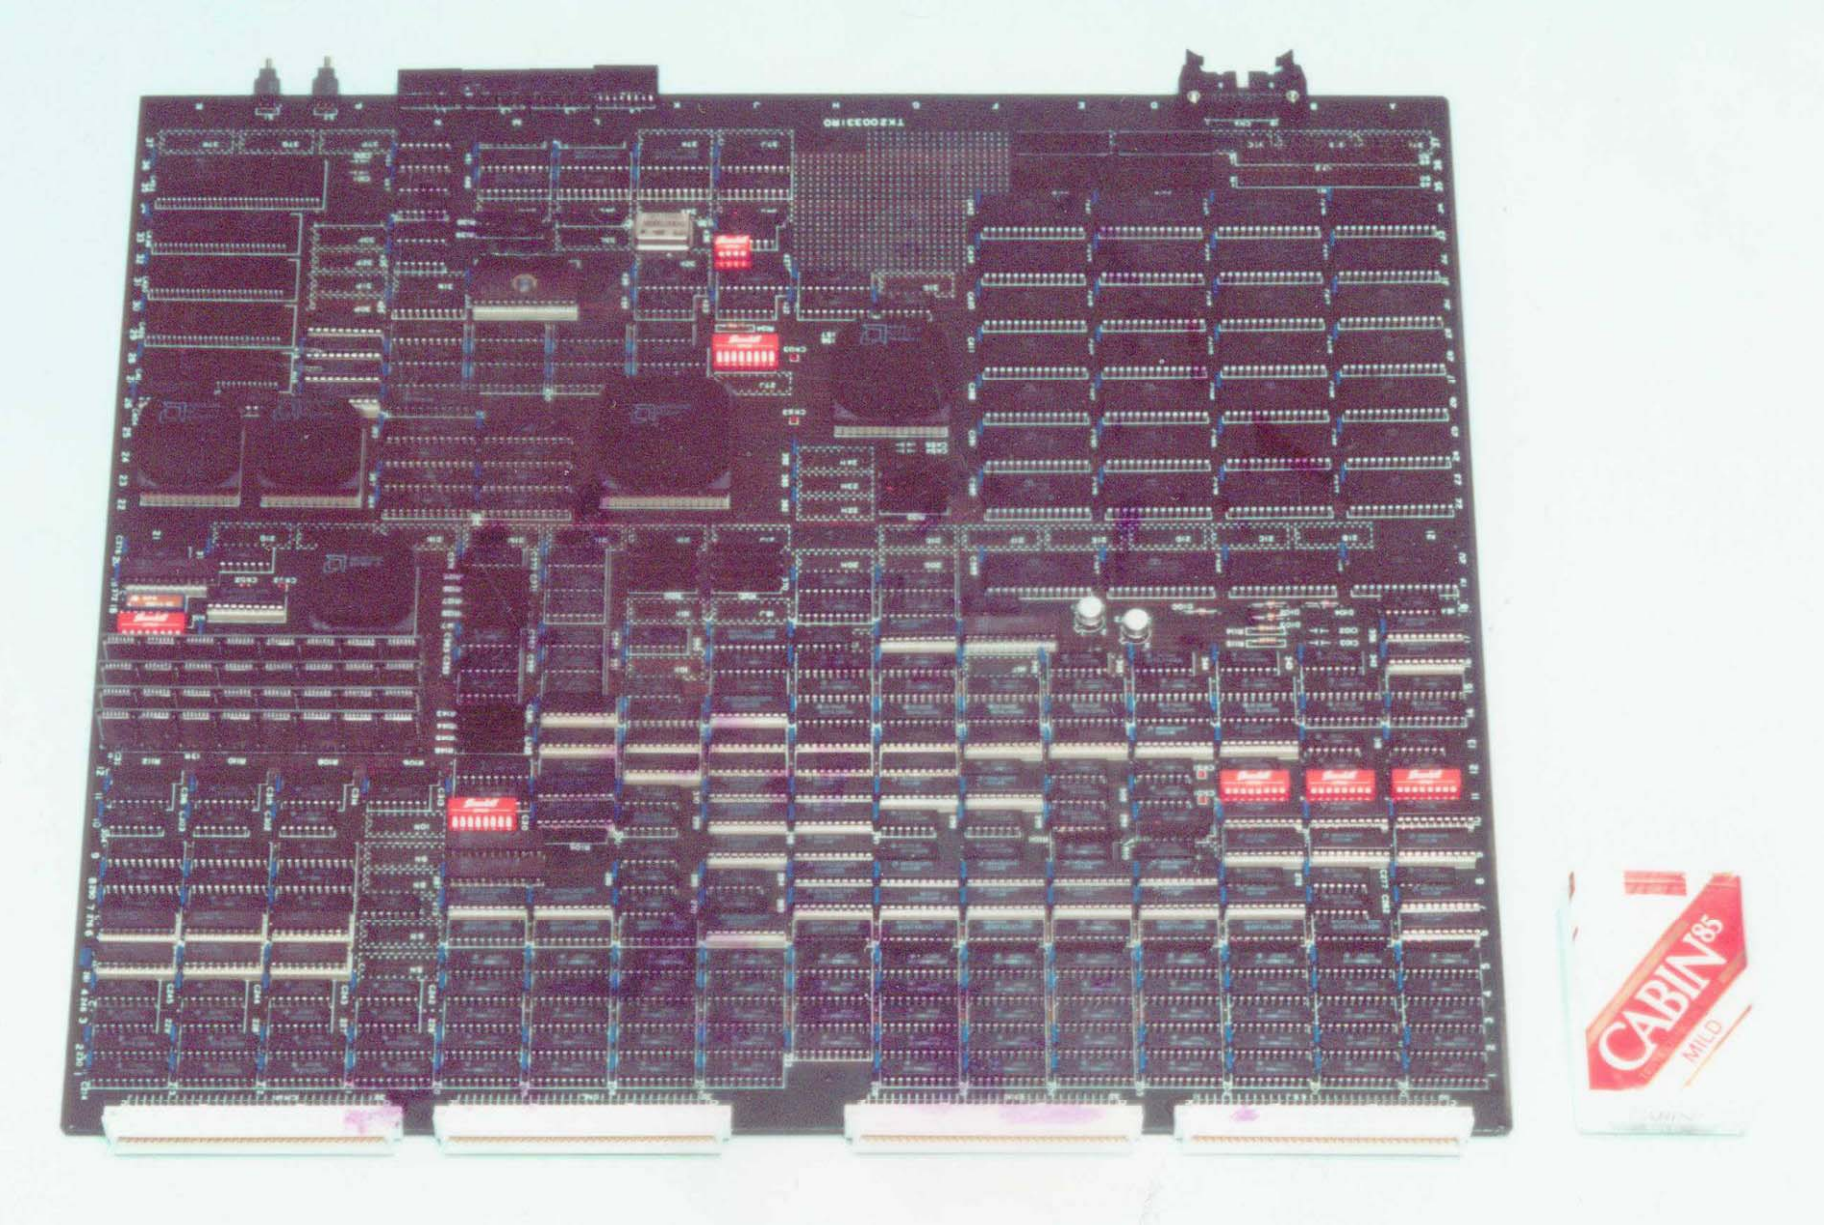

OSCAR Memory Space

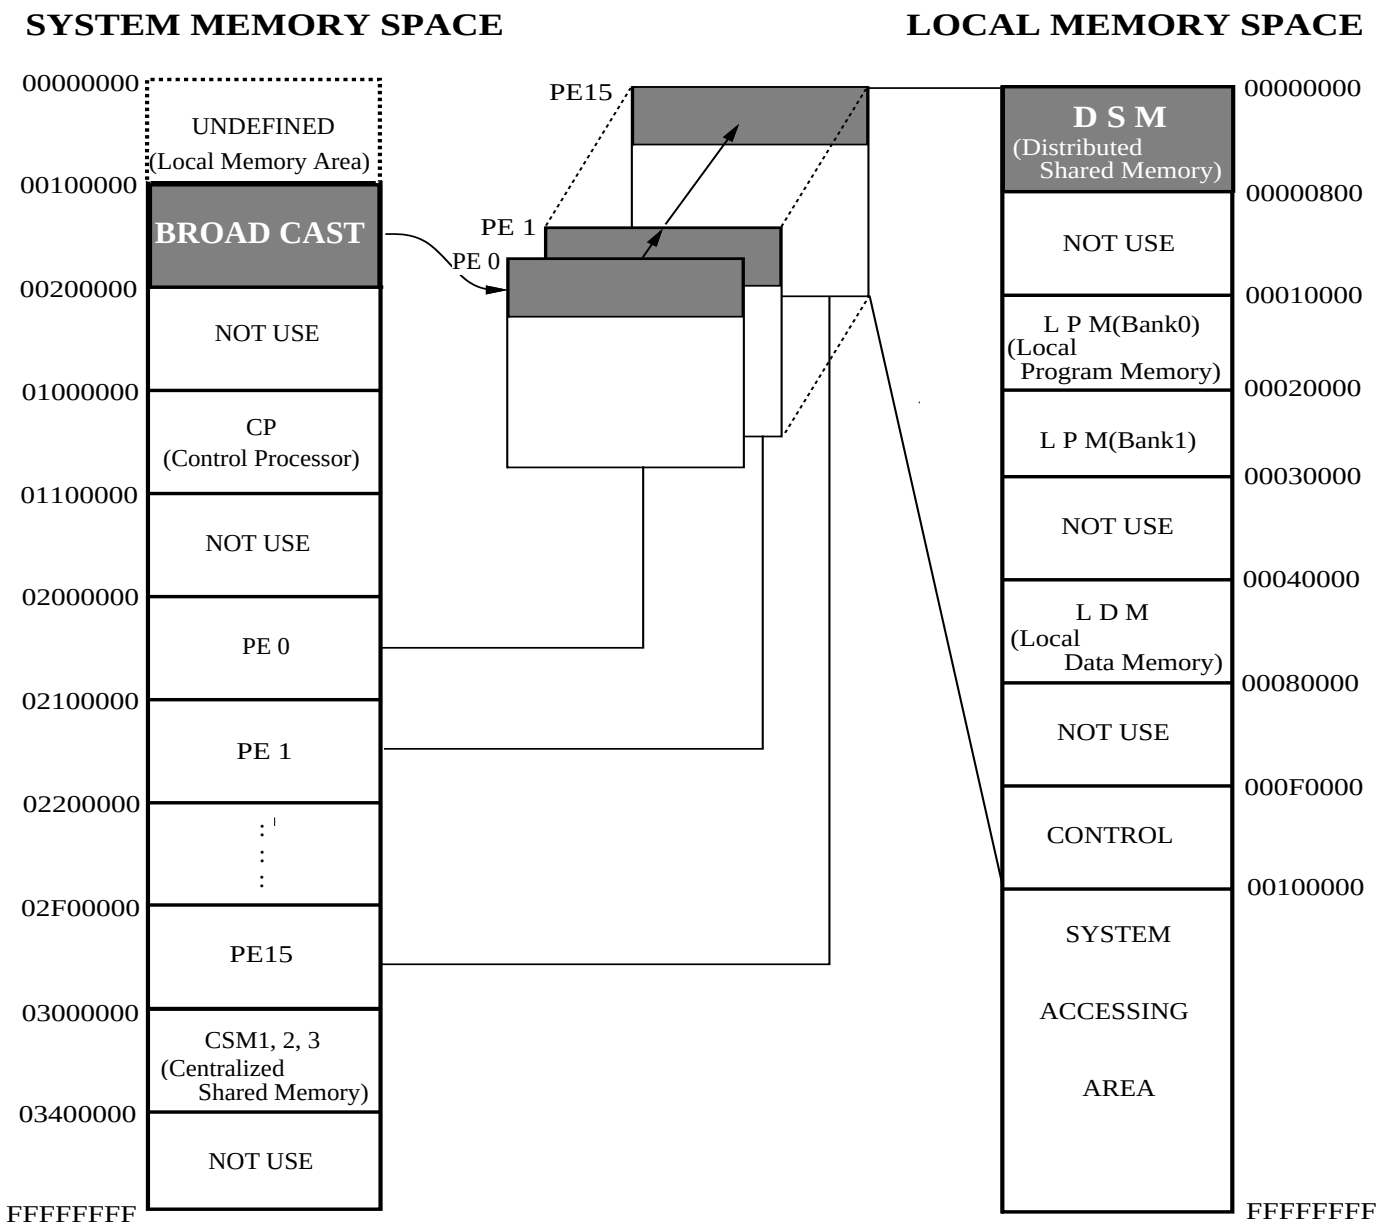

OSCAR Multi-Core Architecture

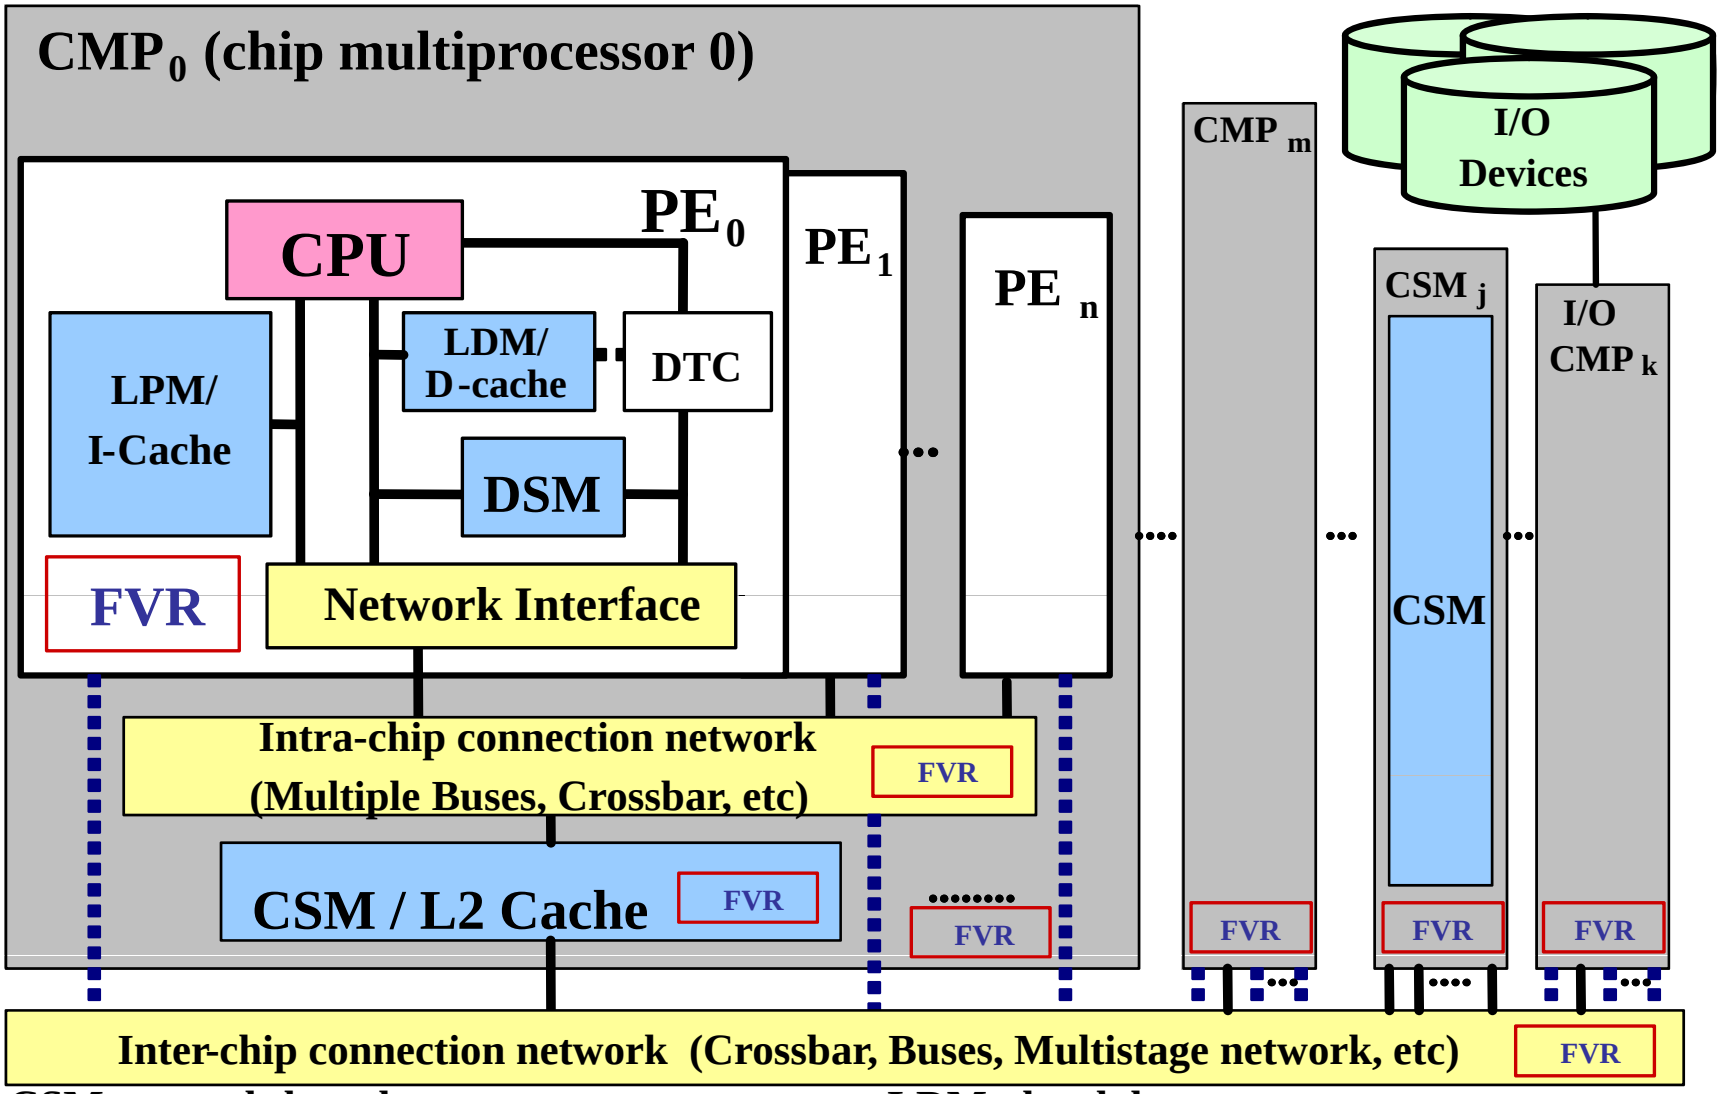

CSM: central shared mem.

DSM: distributed shared mem.

DTC: Data Transfer Controller

LDM : local data mem.

LPM : local program mem.

FVR: frequency / voltage control register

API and Parallelizing Compiler in METI/NEDO Advanced Multicore for Realtime Consumer Electronics Project

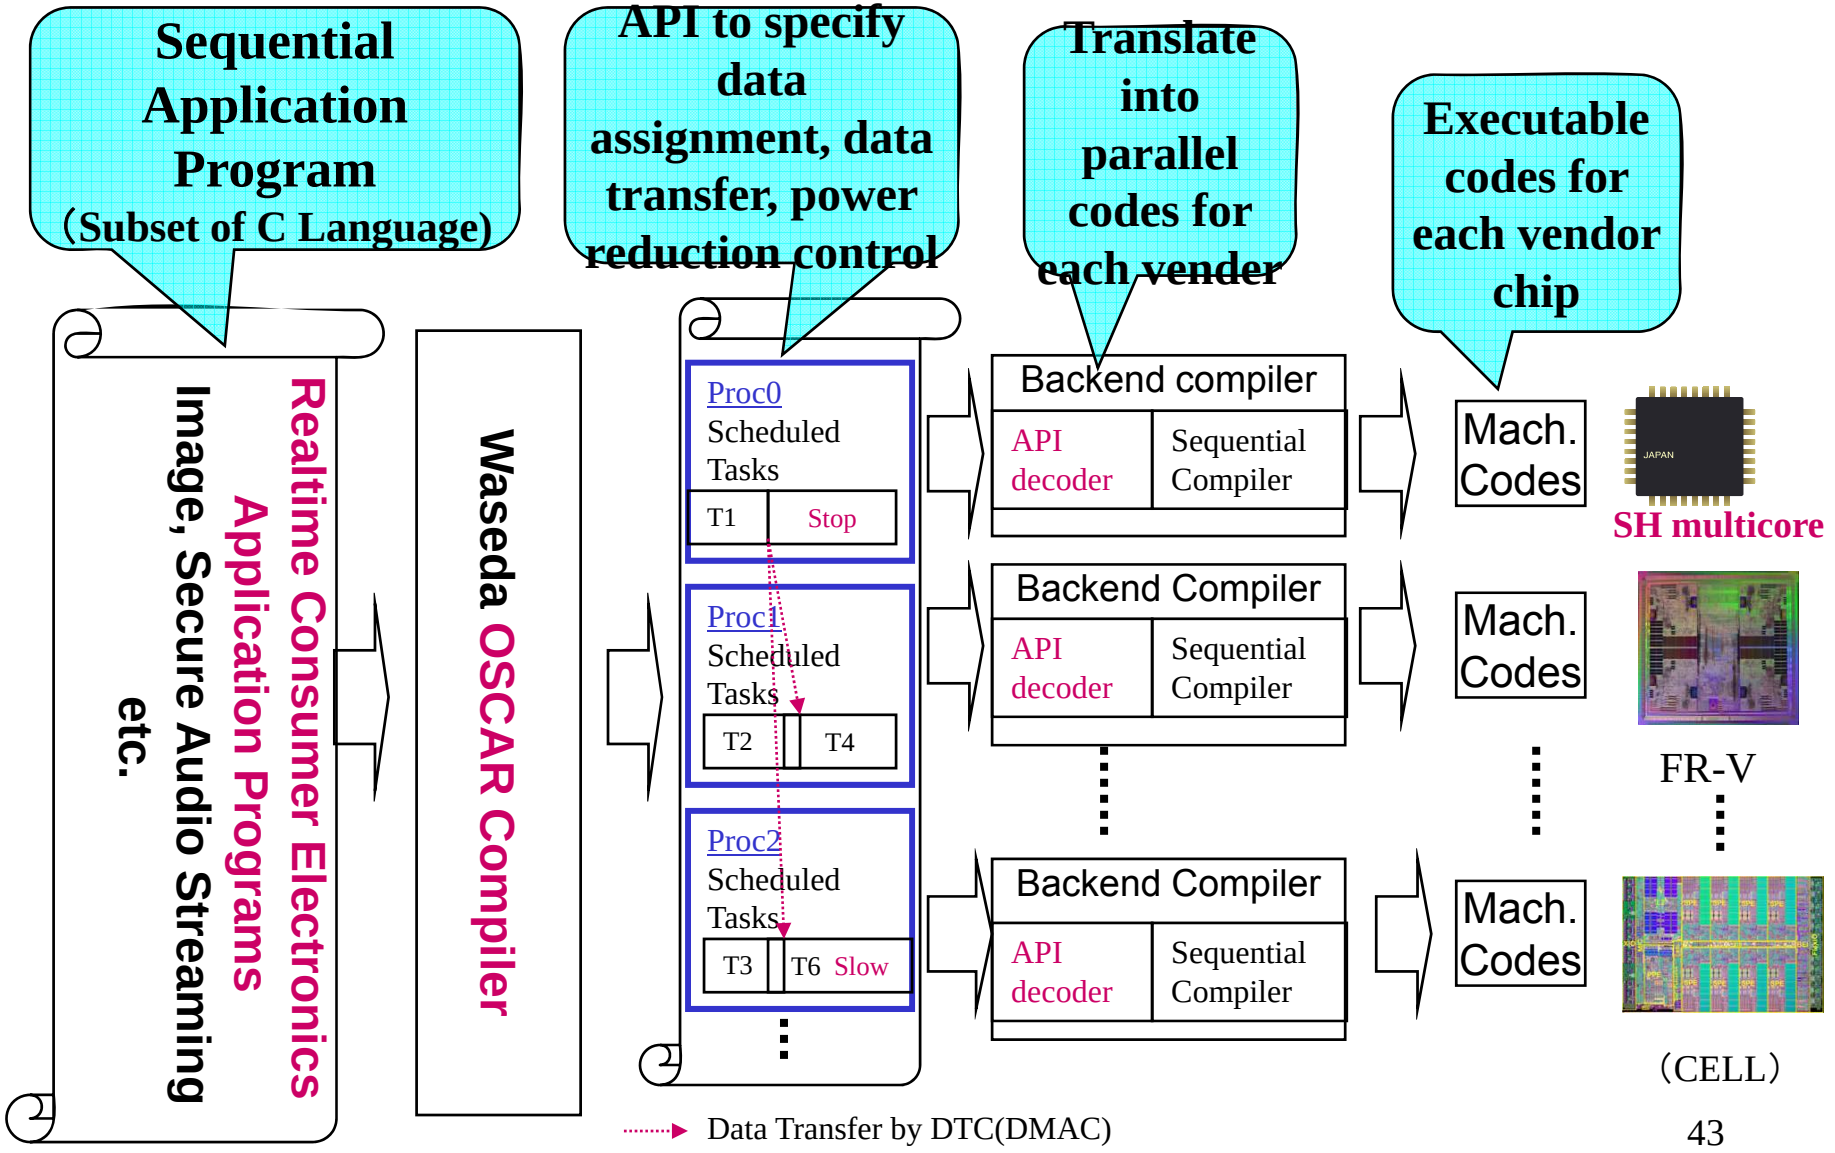

OSCAR Multigrain Parallelizing Compiler

- **Automatic Parallelization**
  - **Multigrain Parallel Processing**
  - **Data Localization**
  - **Data transfer Overlapping**
  - **Complier Controlled Power Saving Scheme**
- **Compiler cooperative Multi-core architecture**
  - OSCAR Multi-core Architecture
  - OSCAR Heterogeneous Multiprocessor Architecture
- **Commercial SMP machines**

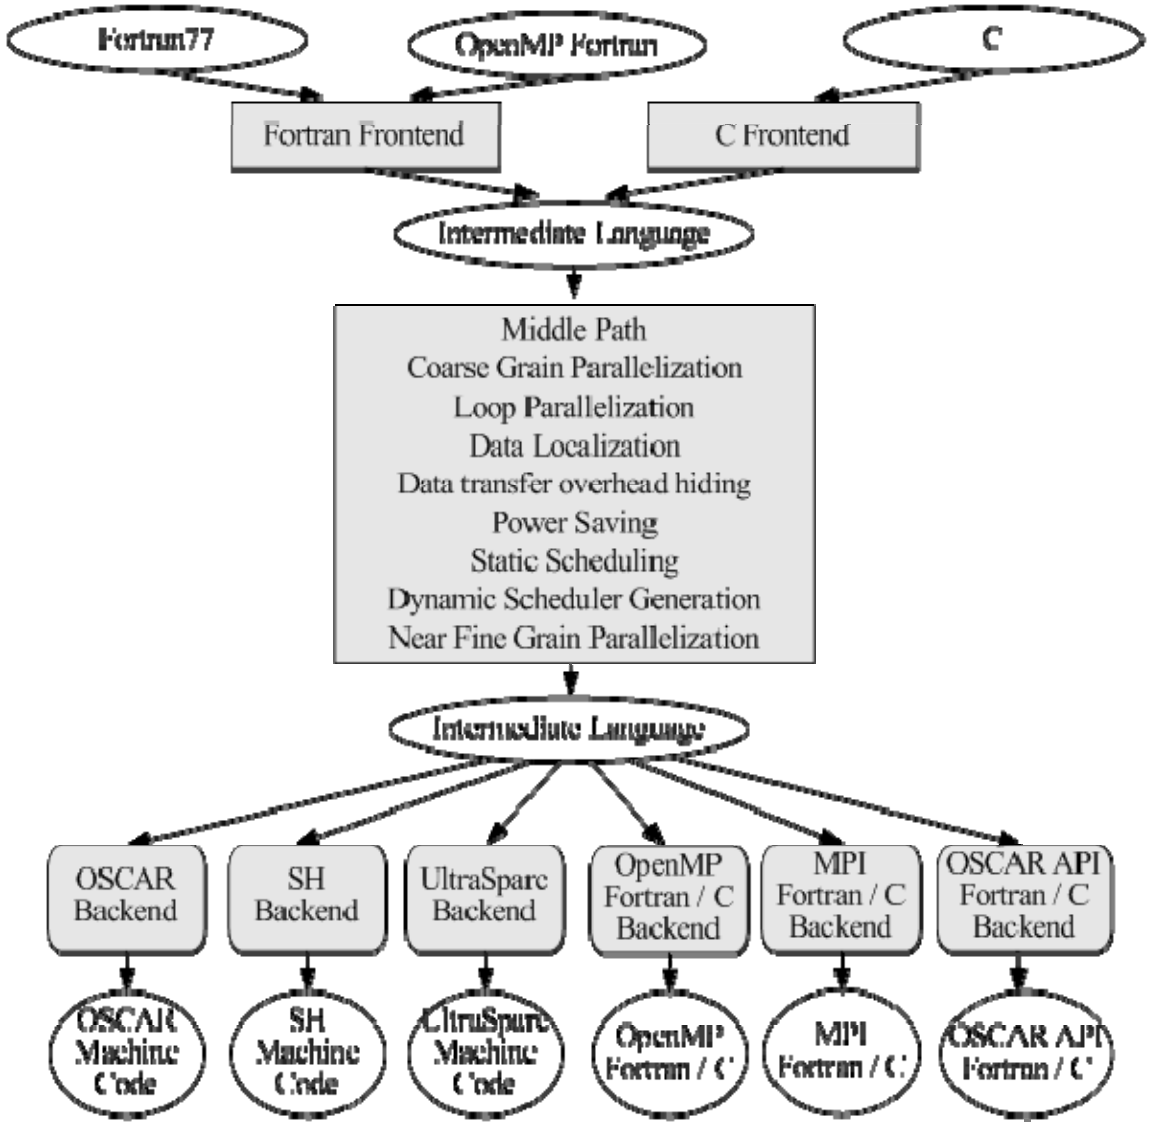

Processor Block Diagram

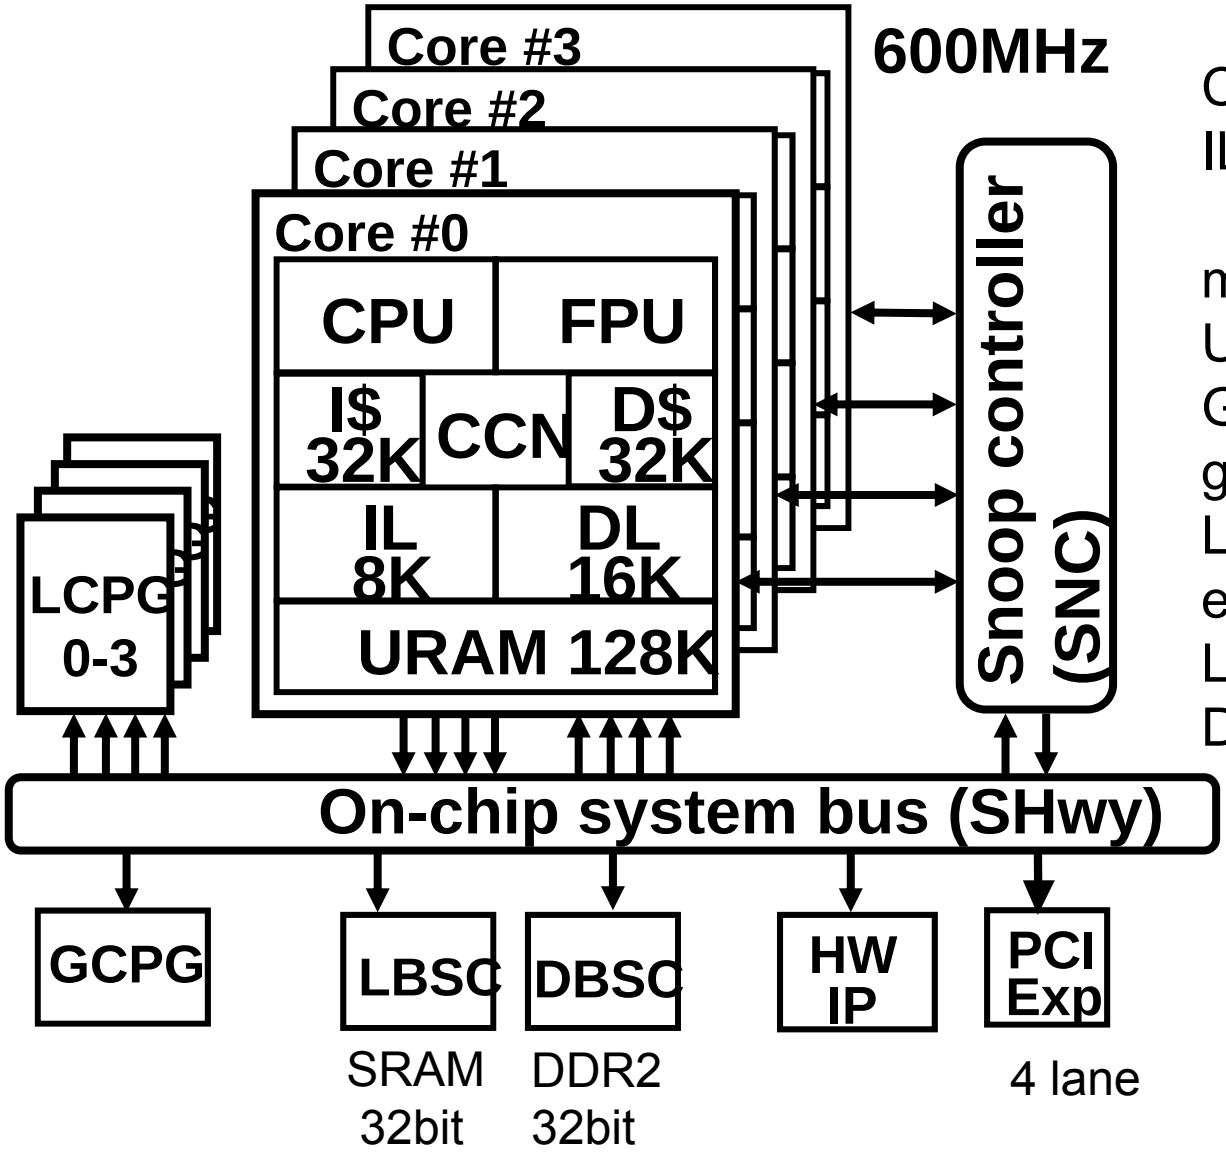

CCN:Cache controller
IL, DL: Instruction/Data local memory
URAM: User RAM
GCPG: Global clock pulse generator
LCPG: Local CPG for each core
LBSC: SRAM controller
DBSC: DDR2 controller

Chip Overview

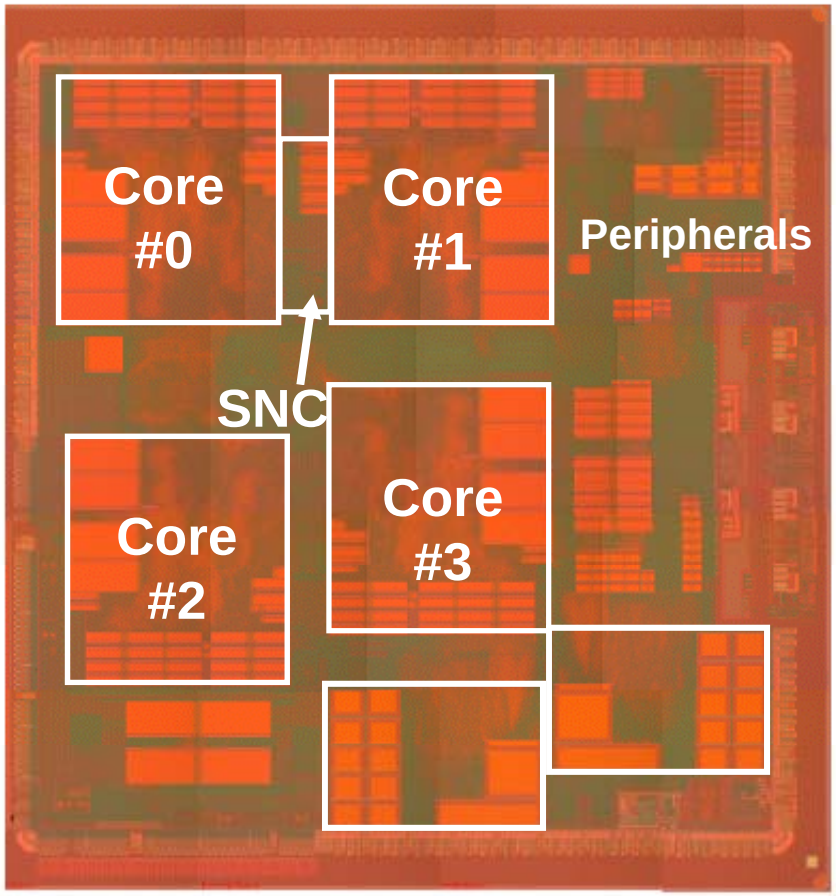

SH4A Multicore SoC Chip

Process Technology	90nm, 8-layer, triple-Vth, CMOS
Chip Size	97.6mm^2 (9.88mm x 9.88mm)
Supply Voltage	1.0V (internal), 1.8/3.3V (I/O)
Power Consumption	0.6 mW/MHz/CPU @ 600MHz (90nm G)
Clock Frequency	600MHz
CPU Performance	4320 MIPS (Dhrystone 2.1)
FPU Performance	16.8 GFLOPS
I/D Cache	32KB 4way set-associative (each)
ILRAM/OLRAM	8KB/16KB (each CPU)
URAM	128KB (each CPU)
Package	FCBGA 554pin, 29mm x 29mm

ISSCC07 Paper No.5.3, Y. Yoshida, et al., "A 4320MIPS Four-Processor Core SMP/AMP with Individually Managed Clock Frequency for Low Power Consumption"

Performance on a Developed SH Multi-core (RP1: SH-X3) Using Compiler and API

Audio AAC* Encoder

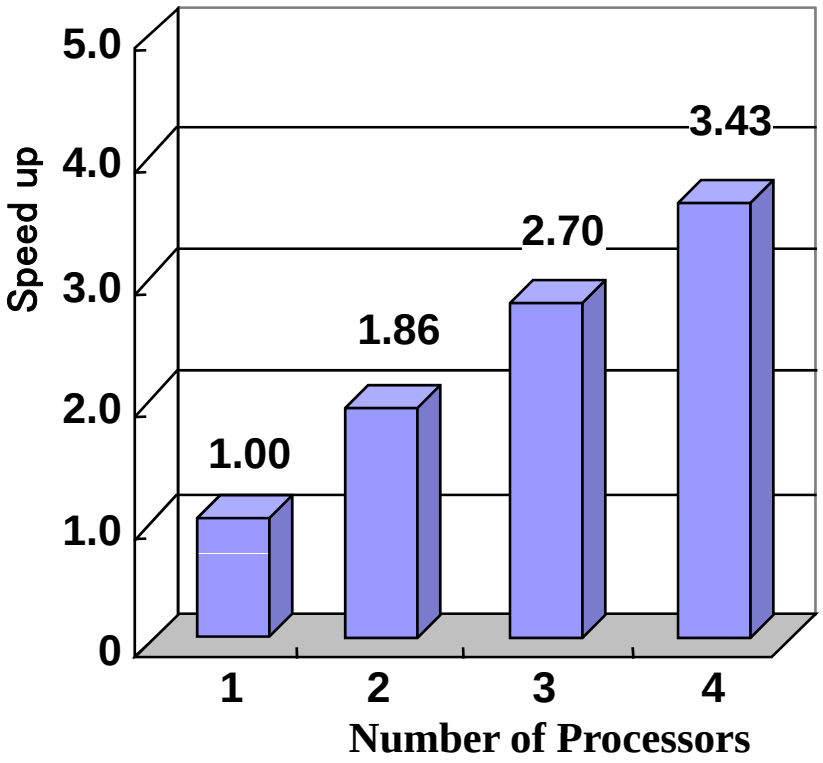

*) ISO Advanced Audio Coding :

Image Susan Smoothing

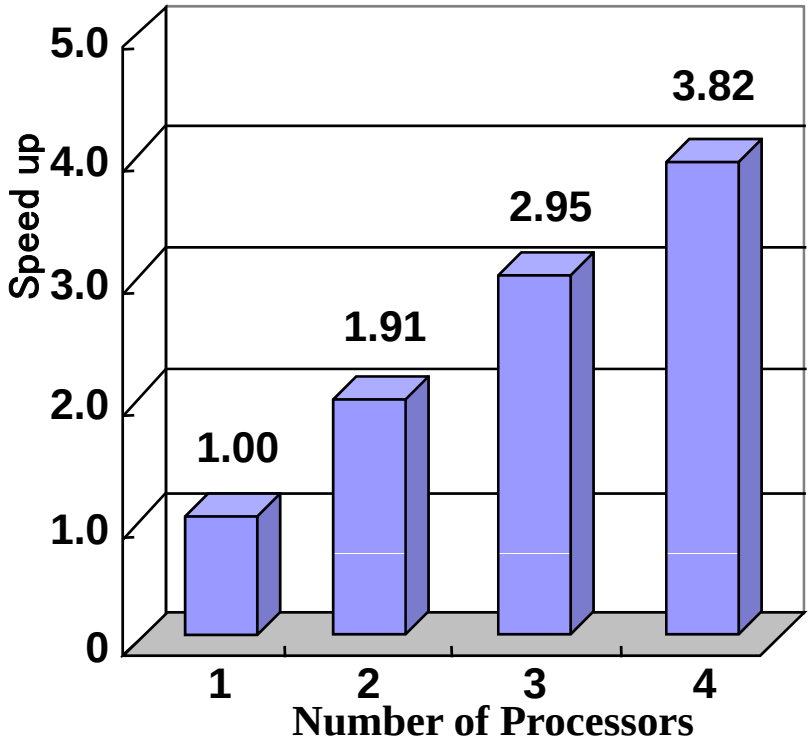

**) Mibench Embedded application benchmark by Michigan Univ.

RP2 Chip Photo and Specifications

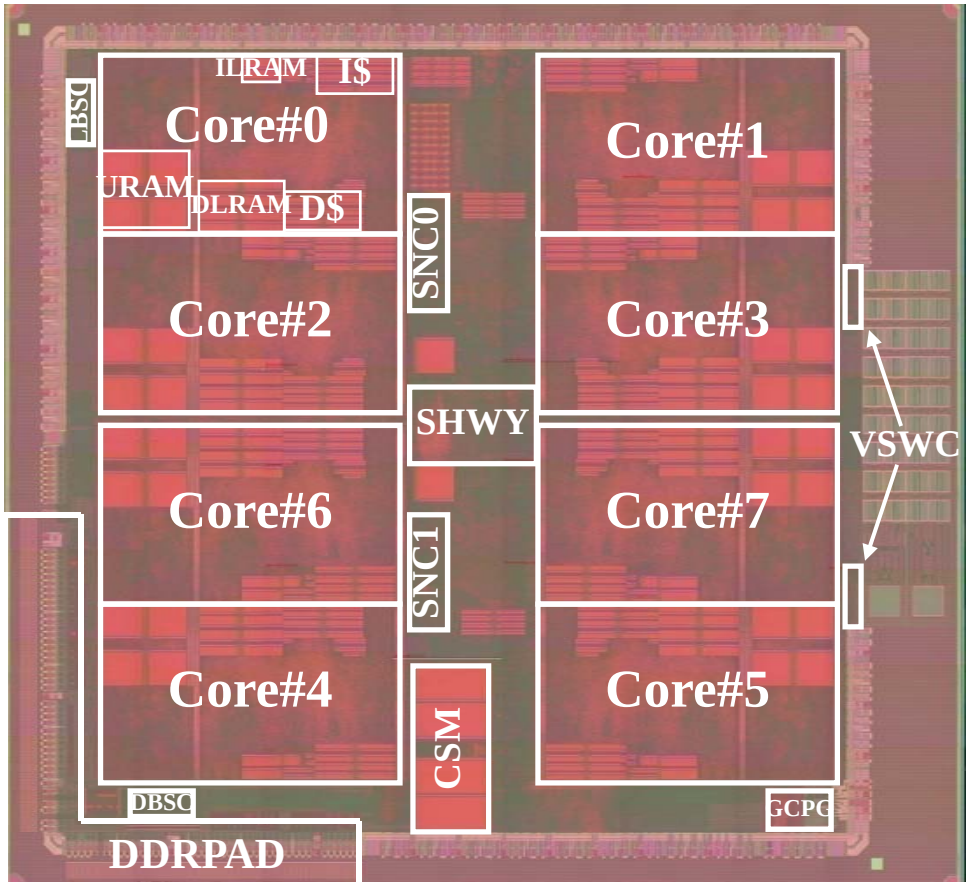

Process Technology	90nm, 8-layer, triple-Vth, CMOS
Chip Size	104.8mm^2 (10.61mm x 9.88mm)
CPU Core Size	6.6mm^2 (3.36mm x 1.96mm)
Supply Voltage	1.0V–1.4V (internal), 1.8/3.3V (I/O)
Power Domains	17 (8 CPUs, 8 URAMs, common)

8 Core RP2 Chip Block Diagram

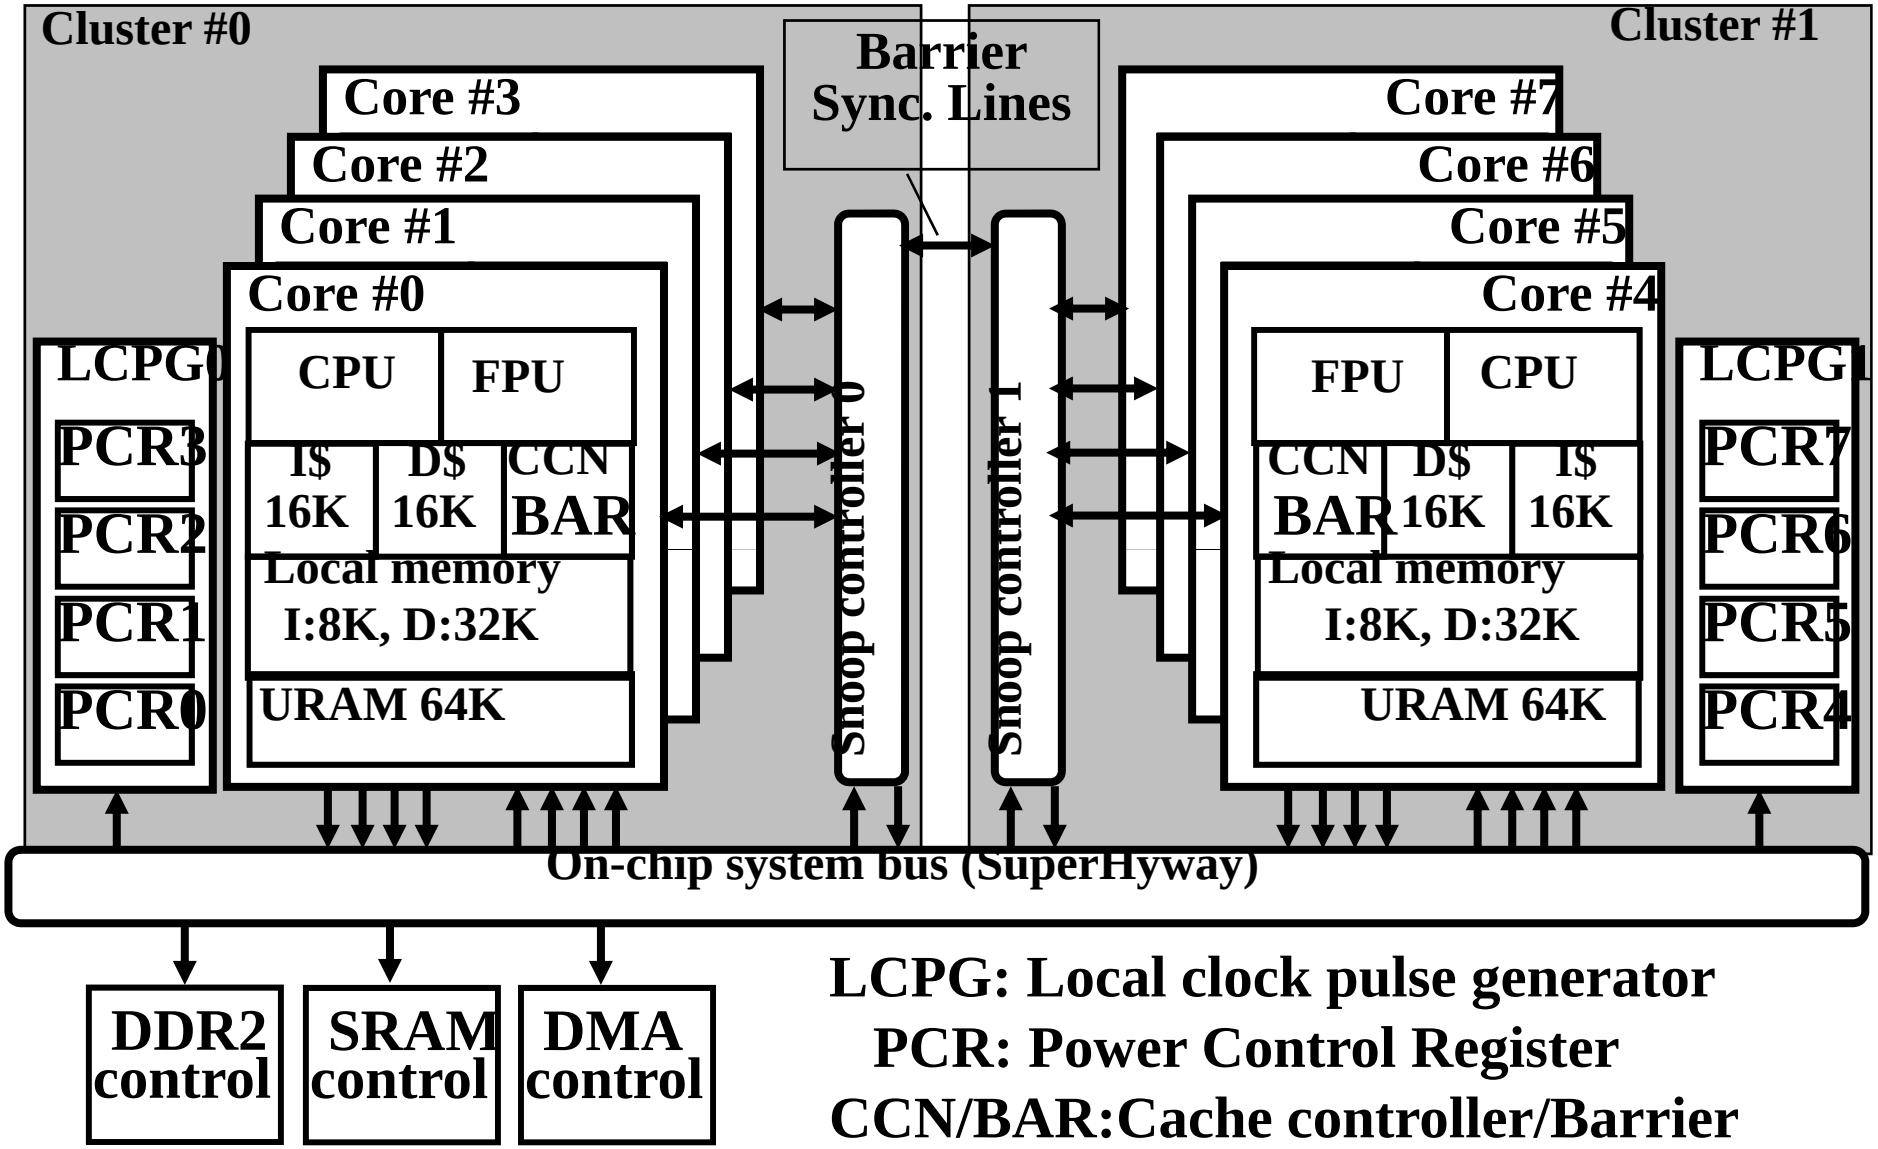

LCPG: Local clock pulse generator

PCR: Power Control Register

CCN/BAR:Cache controller/Barrier Register

URAM: User RAM

Processing Performance on the Developed Multicore Using Automatic Parallelizing Compiler

Speedup against single core execution for audio AAC encoding

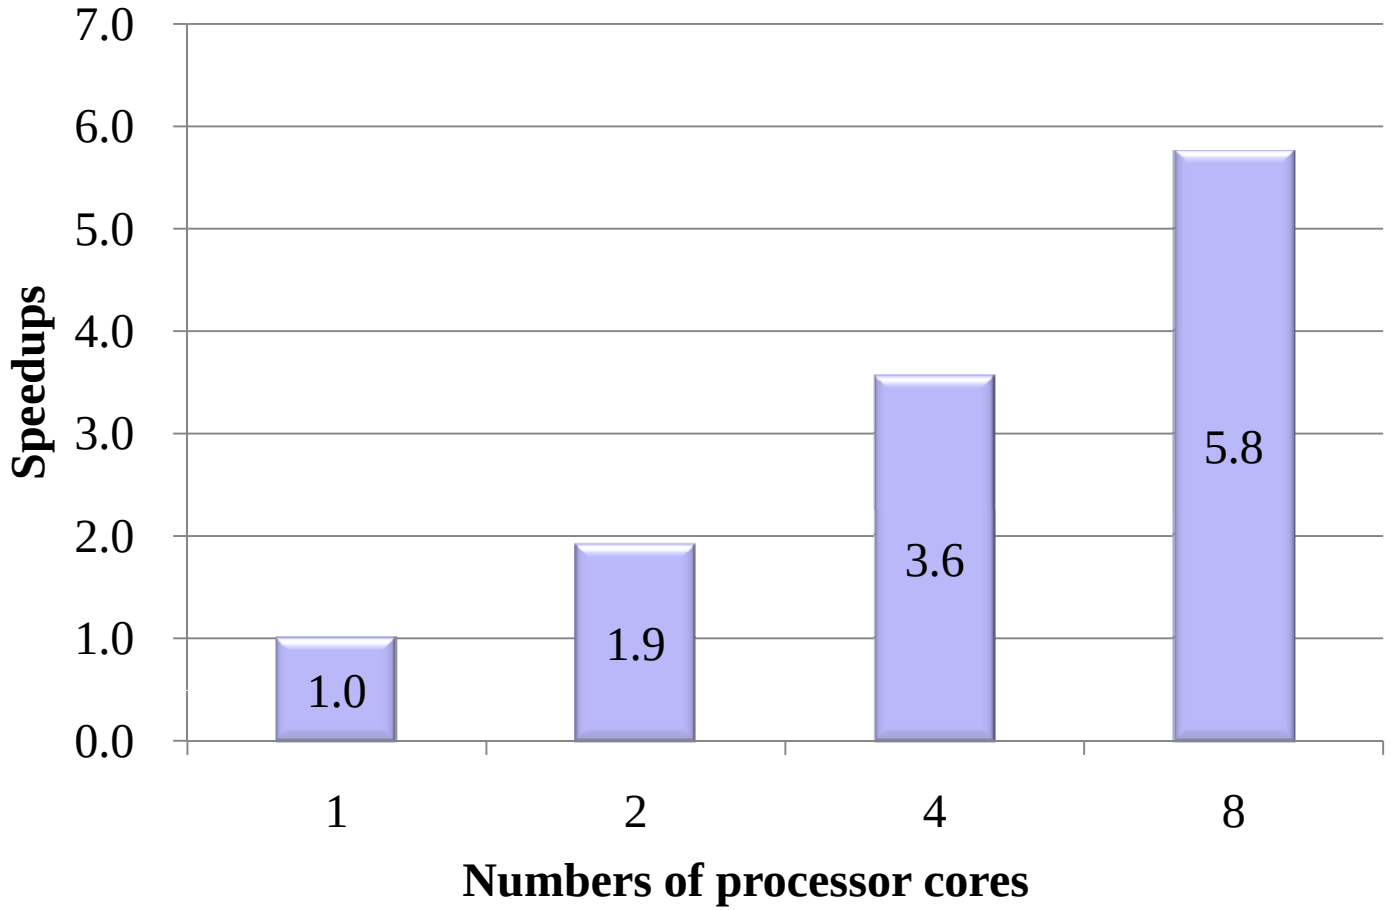

*) Advanced Audio Coding

Power Reduction Using Power Shutdown and Voltage Control by OSCAR Parallelizing Compiler

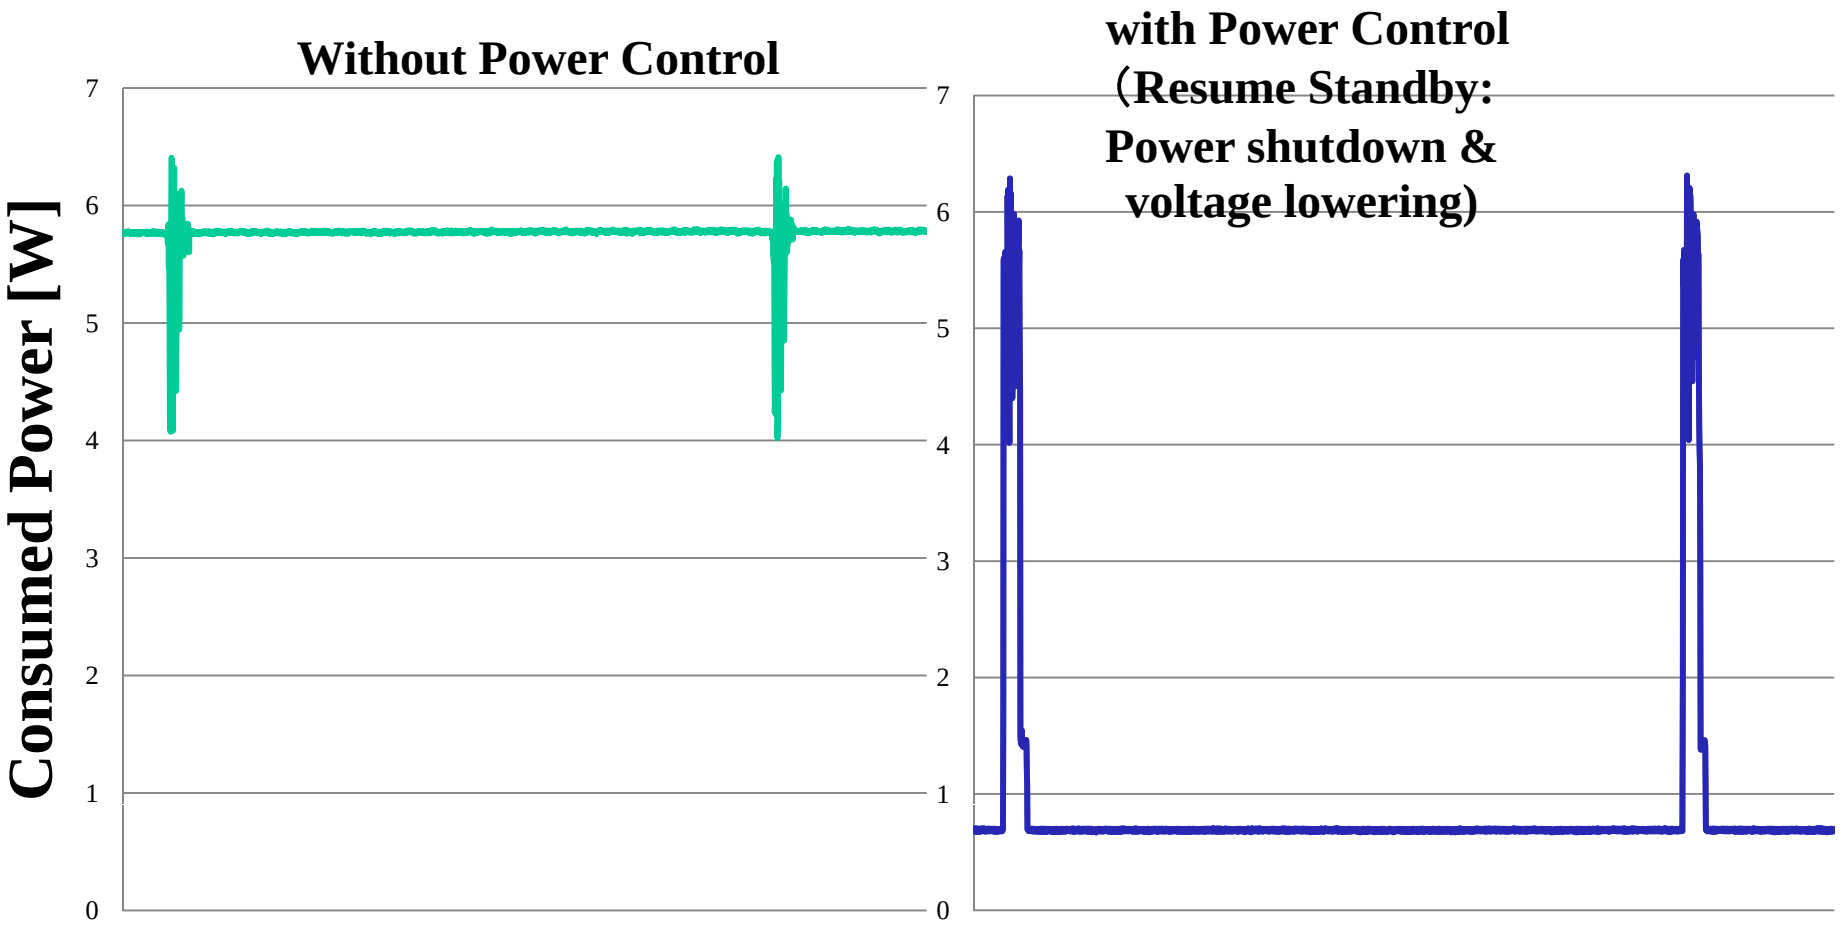

OSCAR Heterogeneous Multicore

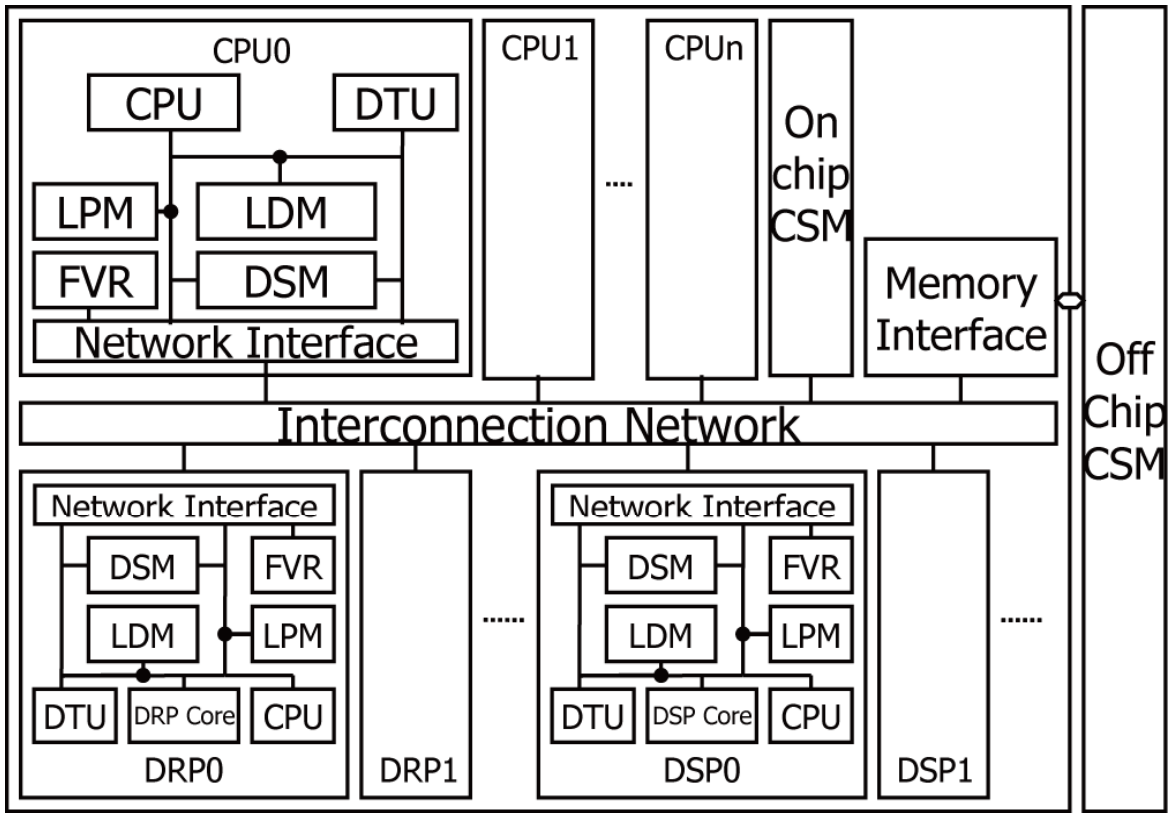

- OSCAR Type Memory Architecture
- LPM
  - Local Program Memory
- LDM
  - Local Data Memory
- DSM
  - Distributed Shared Memory
- CSM
  - Centralized Shared Memory
    - On Chip and/or Off Chip
- DTU
  - Data Transfer Unit
- Interconnection Network
  - Multiple Buses
  - Split Transaction Buses
  - CrossBar …

Static Scheduling of Coarse Grain Tasks for a Heterogeneous Multi-core

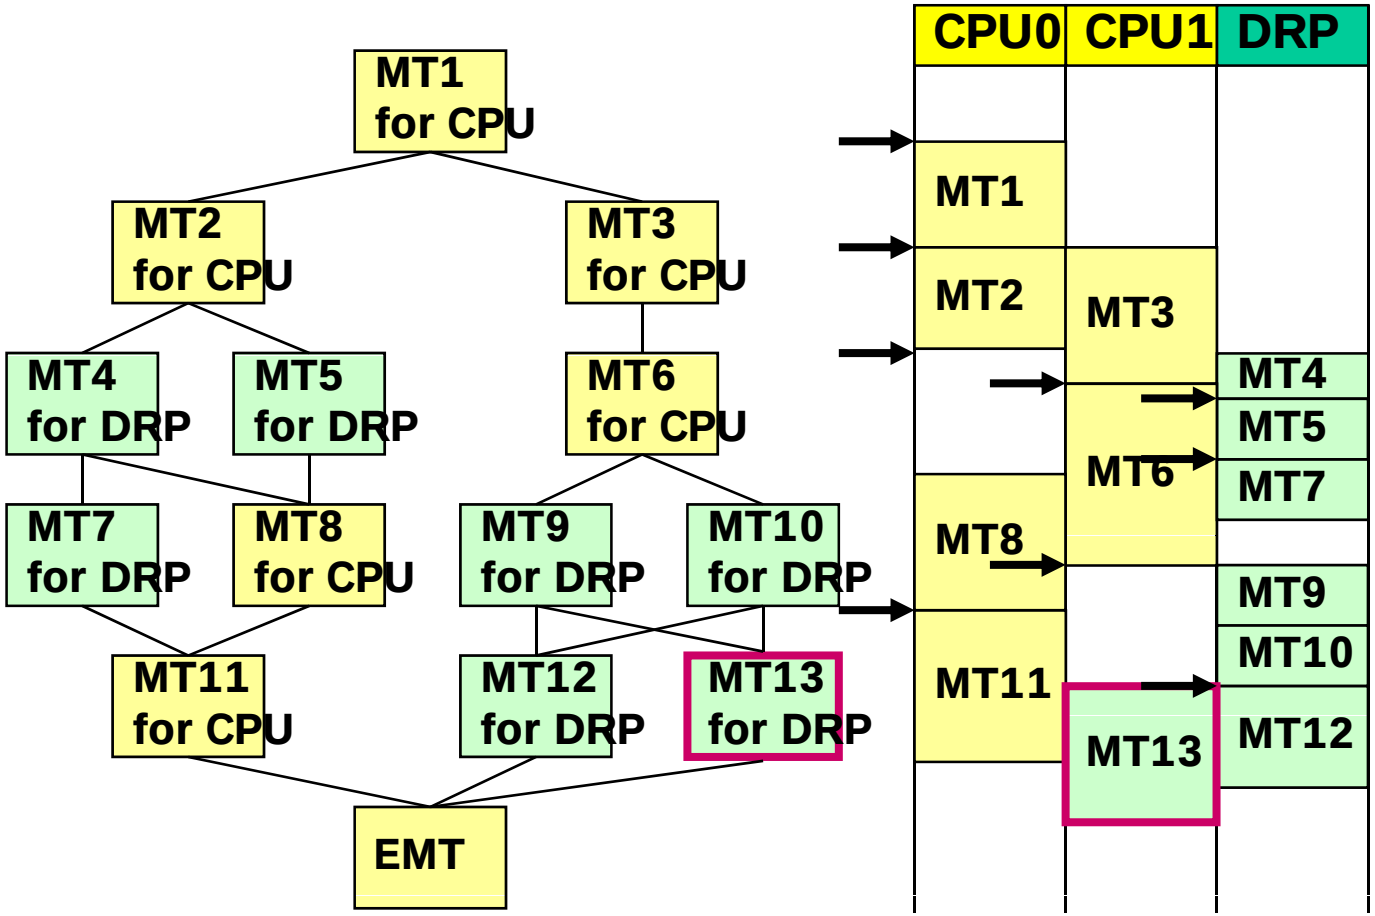

An Image of Static Schedule for Heterogeneous Multi-core with Data Transfer Overlapping and Power Control

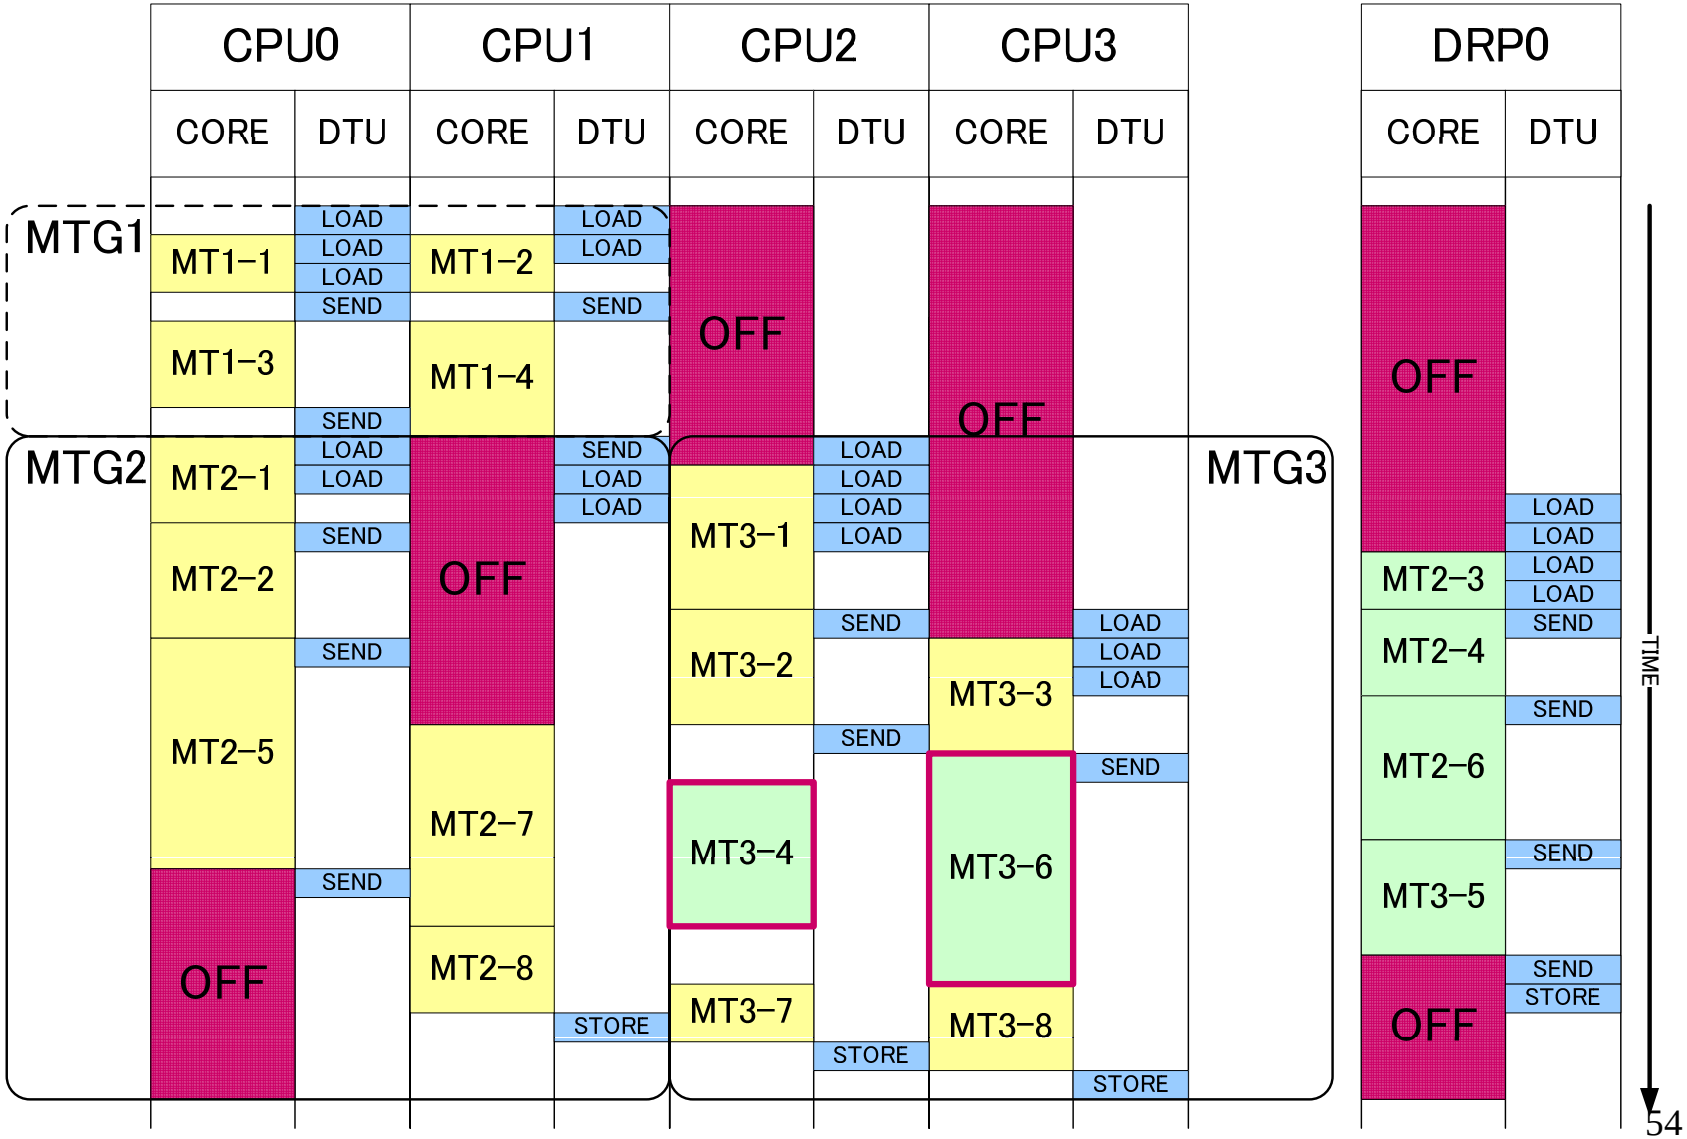

Compiler Performance on a OSCAR Hetero-multi-core

- 25.2 times speedup using 4 SH general purpose cores and 4 DRP accelerators against a single SH

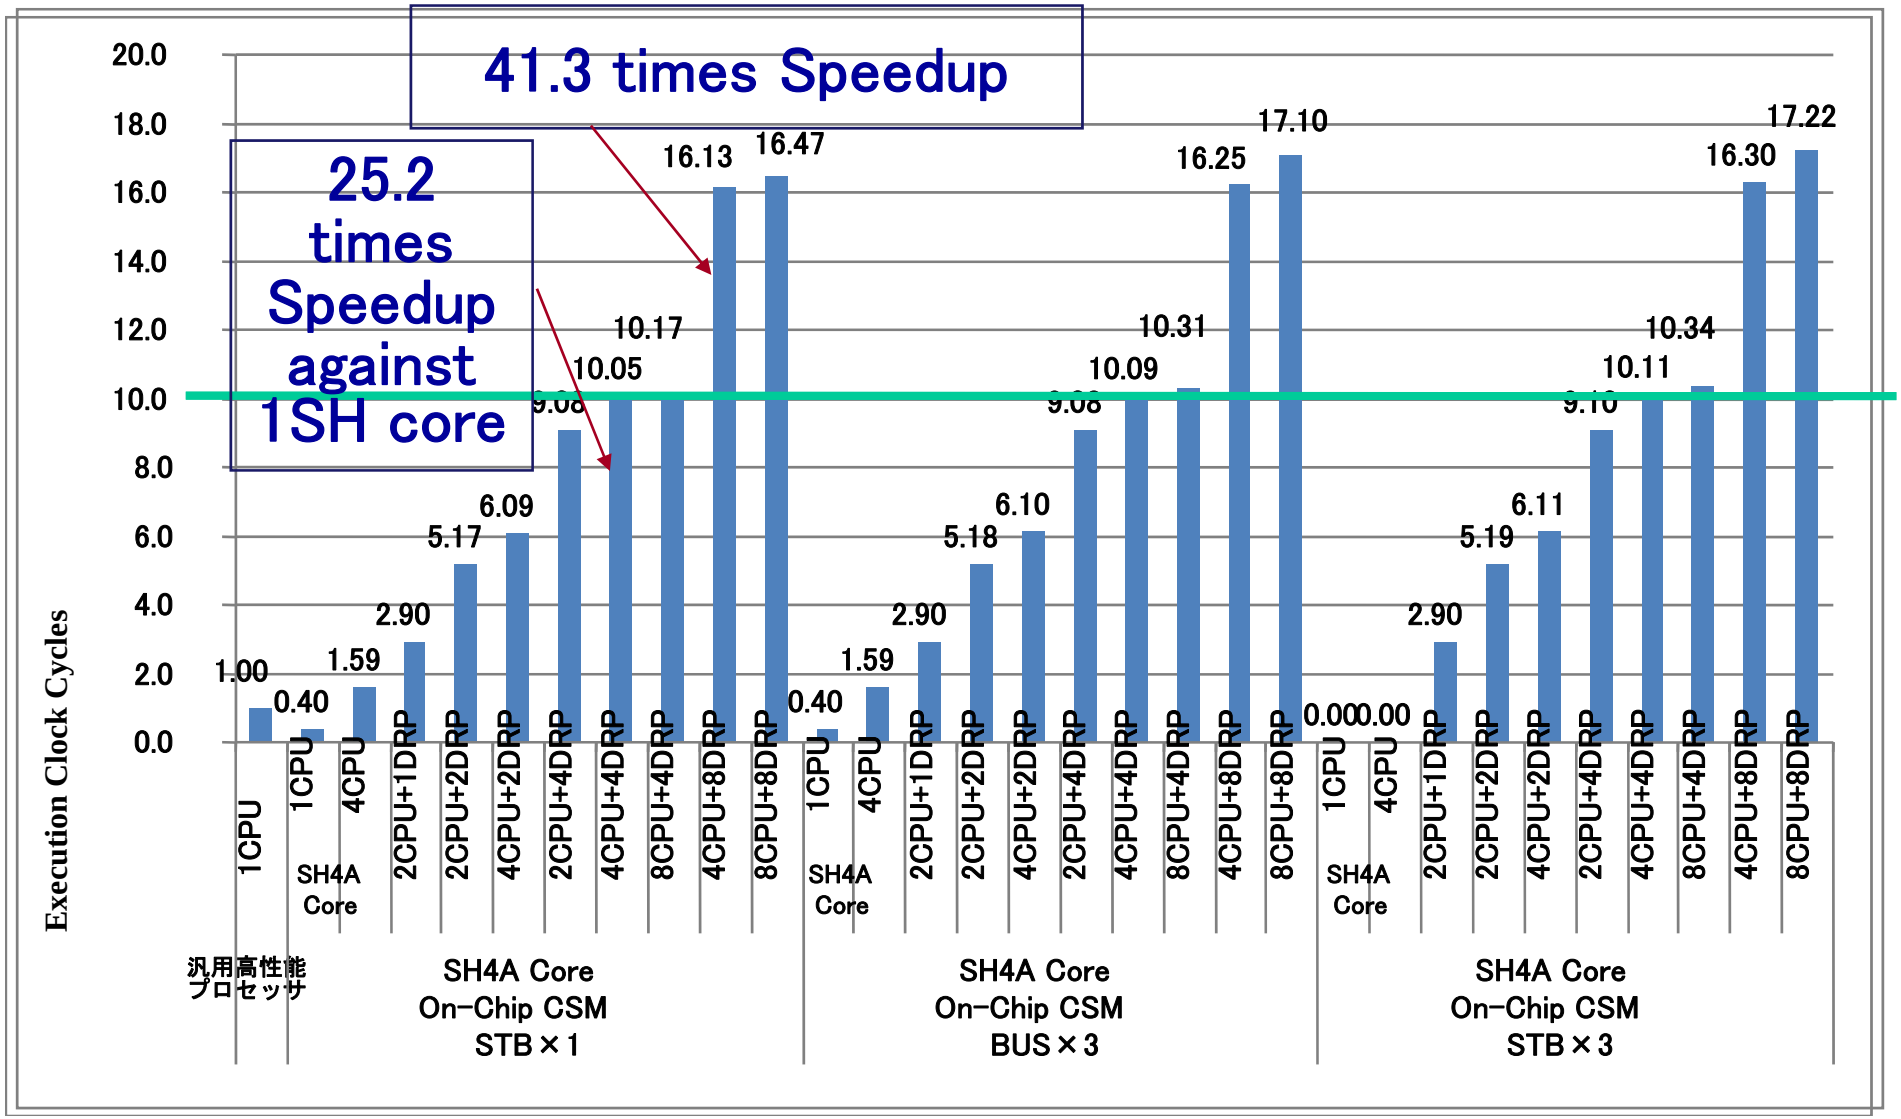

Power Reduction by OSCAR Compiler(4SHs＋4DRPs)

- 0.78 W: 22% Power reduction by Compiler Control

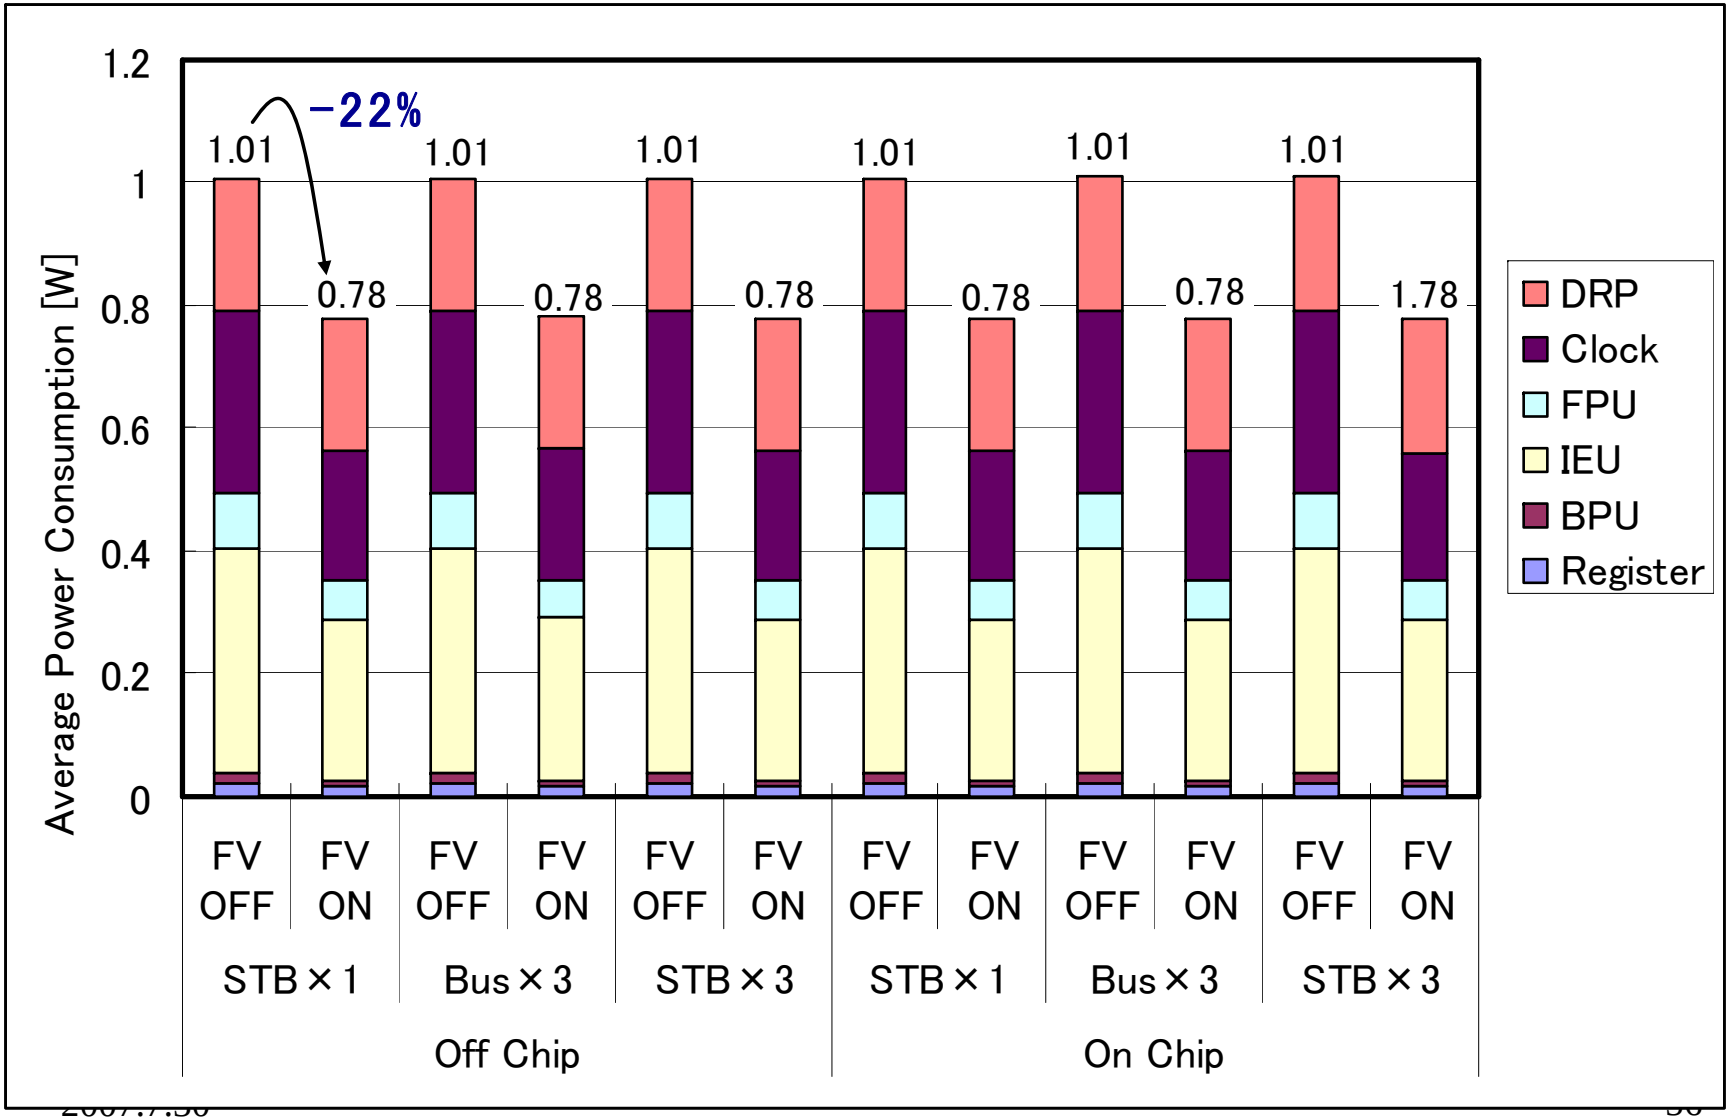

Conclusions

- Compiler cooperative low power high effective performance multi-core processors will be more important in wide range of information systems from games, mobile phones, automobiles to peta-scale supercomputers.
- Parallelizing compilers are essential for realization of
  - Good cost performance
  - Short hardware and software development periods
  - Low power consumption
  - High software productivity
  - Scalable performance improvement with advancement in semiconductor integration technology
- Key technologies in multi-core compiler
  - Multigrain parallelization, Data localization, Data transfer overlapping using DMA, Low power control technologies