アクセラレータ付きマルチコア上でのリアルタイム制御計算の自動並列化・低消費電力化コンパイラ技術

早稲田大学　理工学術院　情報理工学科　教授　笠原博徳

グリーンコンピューティング機構　アドバンストマルチコアプロセッサ研究所長

IEEE Computer Society President 2018, 早稲田大学副総長 (2018-2022)

1976 早稲田大学高等学院卒
1980 早大電気工学科卒、1982 同修士課程了
1985 早大大学院博士課程了 工学博士,学振第一回PD
　　カリフォルニア大学バークレー客員研究員
1986 早大理工専任講師、1988年 助教授
1989-1990 イリノイ大学Center for
　　Super computing R&D客員研究員
1997 教授、現在 理工学術院情報理工学科
2004 アドバンストマルチコア研究所所長
2017 日本工学アカデミー会員、 日本学術会議連携会員
2018 IEEE Computer Society会長、
　　早大副総長（-2022年9月）
2019-2023　産業競争力懇談会(COCN) 理事
2020-日本工学アカデミー理事

【受賞】
1987 IFAC World Congress Young Author Prize
1997 情報処理学会坂井記念特別賞
2005 半導体理工学研究センタ共同研究賞
2008 LSI・オブ・ザ・イヤー 2008 準グランプリ、
　Intel Asia Academic Forum Best Research Award
2010 IEEE CS Golden Core　Member Award
2014 文部科学大臣表彰科学技術賞研究部門
2015 情報処理学会フェロー
2017 IEEE Fellow,　IEEE Eta-Kappa-Nu
2019 IEEE CS Spirit of Computer Society Award
2020 情報処理学会功績賞、SCAT表彰 会長大賞
2023 IEEE Life Fellow

査読付き論文232件、招待講演234件、国際特許取得67件(米・英・中・日等)、新聞・Web記事・TV等メディア掲載 698件

【政府・学会委員等】 歴任数 295件
IEEE Computer Society President 2018、Executive Committee委員長、理事(2009-14)、戦略的計画委員長、Nomination Committee委員長、Multicore STC 委員長、IEEE CS Japan委員長、IEEE技術委員、IEEE Medal選定委員、ACM/IEEE SC'21基調講演選定委員等
【経済産業省・NEDO】情報家電用マルチコア・アドバンスト並列化コンパイラ・グリーンコンピューティング・プロジェクトリーダ、NEDOコンピュータ戦略委員長等
【内閣府】スーパーコンピュータ戦略委員、政府調達苦情検討委員、総合科学技術会議情報通信PT 研究開発基盤領域&セキュリティ・ソフト検討委員、日本国際賞選定委員
【文部科学省・海洋研】地球シミュレータ（ES）中間評価委員、情報科学技術委員、HPCI計画推進委員、次世代スパコン（京）中間評価委員・概念設計評価委員、地球シミュレータES2導入技術アドバイザイリー委員長等、JST: ムーンショットG3ロボット&AI Vice Chair, SBIRフェーズ1委員長等

TIER IV Workshop 2023 自動運転におけるAIコンピューティング, 2023年7月13日(木), 18:15-18:55

Some of papers in and just after Ph.D. Course in Waseda U.

IEEE TRANSACTIONS ON COMPUTERS, VOL. C-33, NO. 11, NOVEMBER 1984 1023

Practical Multiprocessor Scheduling Algorithms for Efficient Parallel Processing

HIRONORI KASAHARA, MEMBER, IEEE, AND SEINOSUKE NARITA, SENIOR MEMBER, IEEE

Courtesy of dexchao - Fotolia.com

104 IEEE JOURNAL OF ROBOTICS AND AUTOMATION, VOL. RA-1, NO. 2, JUNE 1985

Parallel Processing of Robot-Arm Control Computation on a Multimicroprocessor System

HIRONORI KASAHARA MEMBER, IEEE, AND SEINOSUKE NARITA, SENIOR MEMBER, IEEE

2nd International Conference on Superecomputing
Santa Clara, CA, USA May 3-8,1987

A PARALLEL PROCESSING SCHEME FOR THE SOLUTION OF SPARSE LINEAR EQUATIONS USING STATIC OPTIMAL-MULTIPROCESSOR-SCHEDULING ALGORITHMS

H.Kasahara*, T.Fujii*, H.Nakayama*, S.Narita*, and Leon O.Chua**

* Dept. of Electrical Eng., Waseda University, Tokyo, 160, Japan
** Dept. of Electrical Eng. and Computer Sciences, University of California, Berkeley, CA 94720, U.S.A.

PARALLEL PROCESSING OF ROBOT MOTION SIMULATION

H. Kasahara, H. Fujii and M. Iwata

Department of Electrical Engineering, Waseda University, 3–4–1 Ohkubo Shinjuku-ku, Tokyo 160, Japan

The First Compiler Codesigned Multiprocessor

OSCAR (Optimally Scheduled Advanced Multiprocessor) in 1987

AMD29325 32-bit Floating-point unit

H. Kasahara, "OSCAR Fortran Multigrain Compiler", Stanford University, Hosted by Professor John L. Hennessy and Professor Monica Lam, May. 15. 1995.

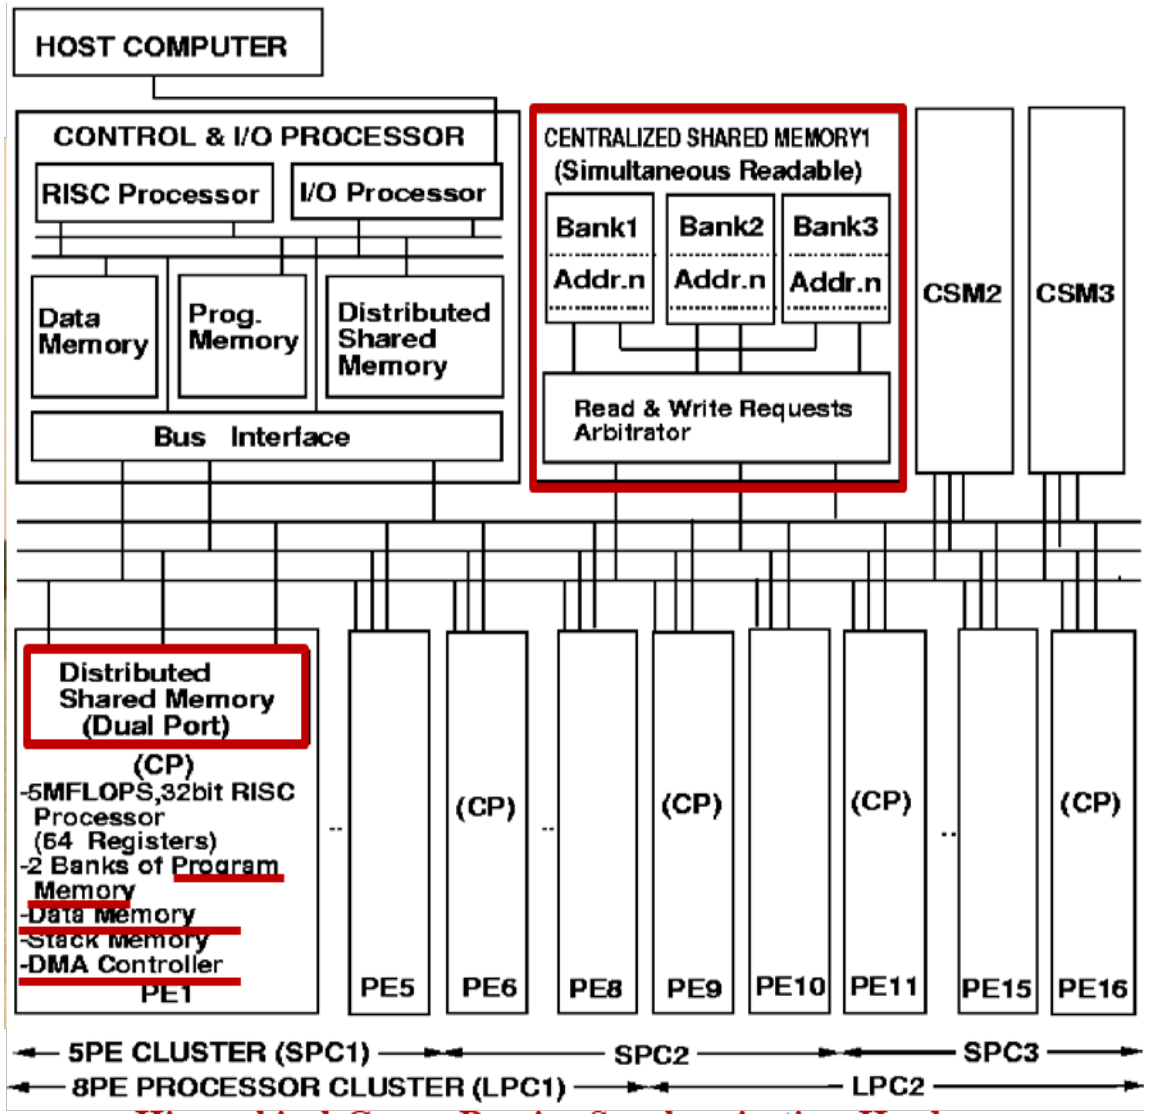

Hierarchical Group Barrier Synchronization Hardware

Cedar Supercomputer

University of Illinois at Urbana-Champaign, CSRD (Center for Supercomputing R&D)

Prof. David Kuck

Univ. Illinois at Urbana-Champaign
Prof. Emeritus
Founder of Kuck & Associates

- **ILLIAC IV (Member)**
- **Parallelizing Compiler**
  - **Data dependence analysis**
  - **Automatic vectorization**
  - **Automatic loop parallelization**
  - **Loop restructuring**
  - **Cedar supercomputer (Alliant FX8)**
  - **OpenMP**

Intel Senior Fellow (Past CTO):

- **Intel Parallel Tuning Tools & Debuggers**

- **National Academy of Engineering Member**
- **IEEE ACM Eckert-Mauchly Award**
- **IEEE Charles Babbage Award**
- **ACM/IEEE Ken Kennedy Award**
- **Okawa Prize**
- **IEEE Computer Pioneer Award**

1993年 スーパーコンピュータVPP500、数値風洞(NWT)

Mr. Hajime Miyoshi

ACM/IEEE SC '94: Washington, D.C. November, 1994にて発表

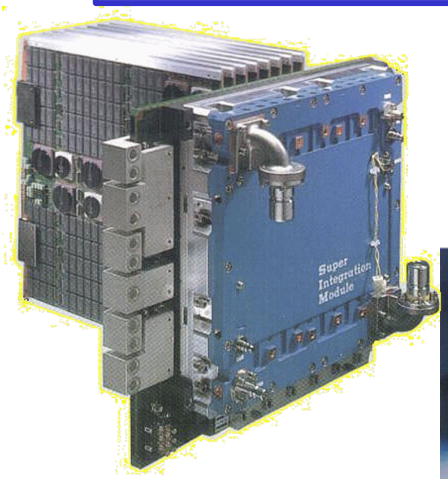

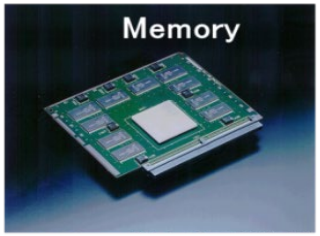

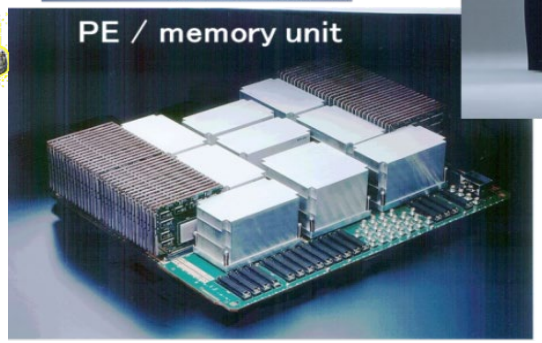

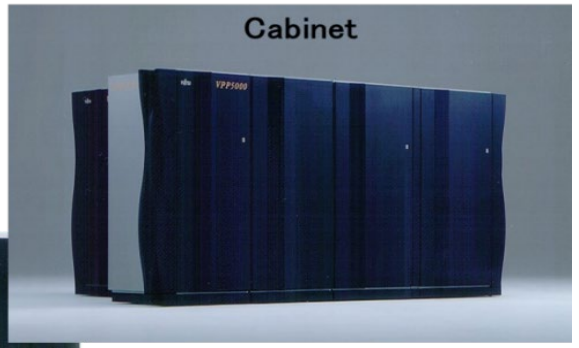

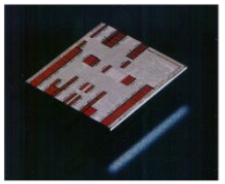

CMOS LSI

スパーコンピュータNWTの外観

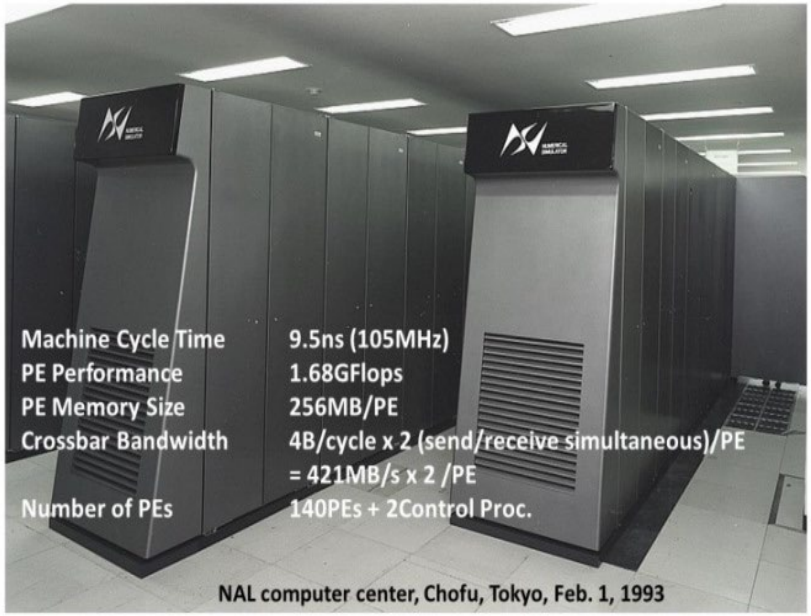

商用VPP5000 (仏気象庁他)

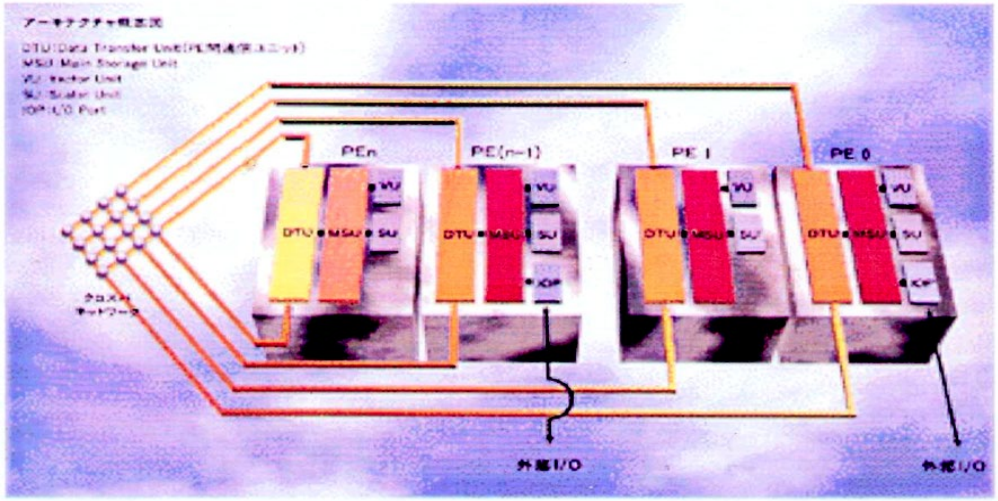

Earth Simulator

2021年ノーベル物理学賞
プリンストン大 真鍋淑郎先生
大気・海洋大循環モデル

(http://www.es.jamstec.go.jp/)

Mr. Hajime Miyoshi

- Earth Environmental simulation like Global Warming, El Nino, PlateMovement for the all lives onr this planet.
- Developed in Mar. 2002 by STA (MEXT) and NEC with 400 M$ investment under Dr. Miyoshi's direction.
  (Dr.Miyoshi: Passed away in Nov.2001. NWT, VPP500, SX6)

40 TFLOPS Peak ($40*10^{12}$)

35.6 TFLOPS Linpack

Image of Earth Simulator

4 Tennis Courts

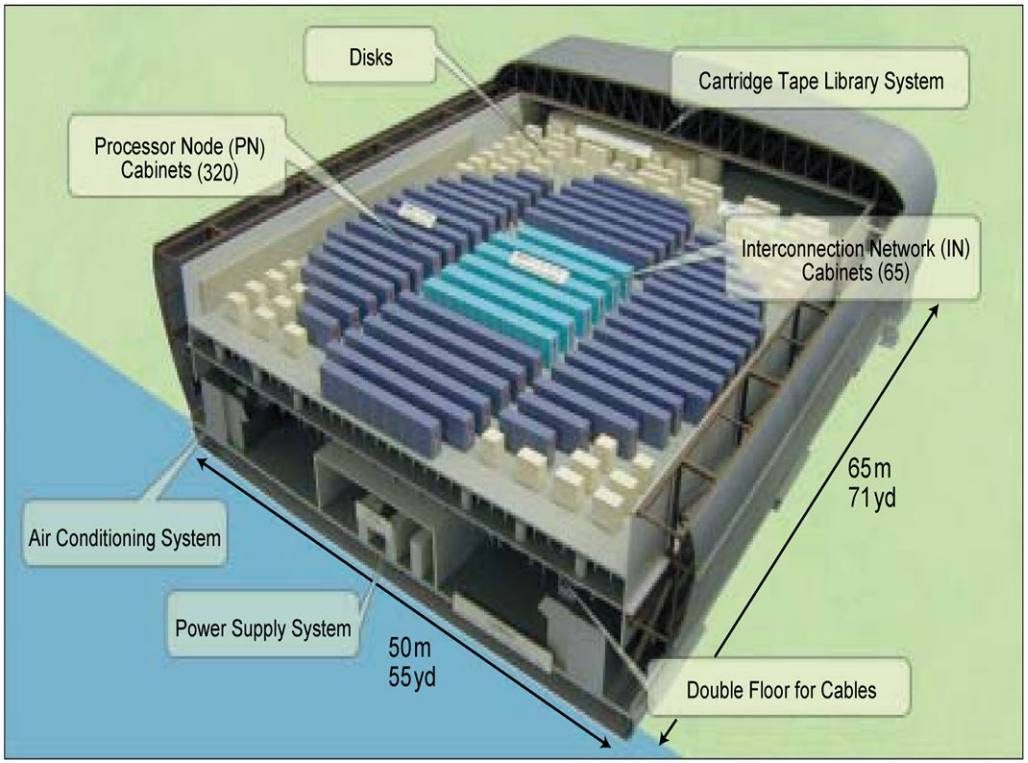

June 2002 Top1
Cores: 5,120, Rmax:35.86TFlop/s
Rpeak: 40.96TFlop/s, Power: 3.2MW

理化学研究所　神戸ポートアイランド　10PFLOPS　京のアーキテクチャ

“K” Supercomputer by Riken
No.1 in TOP500, June 20 & Nov.2,2011

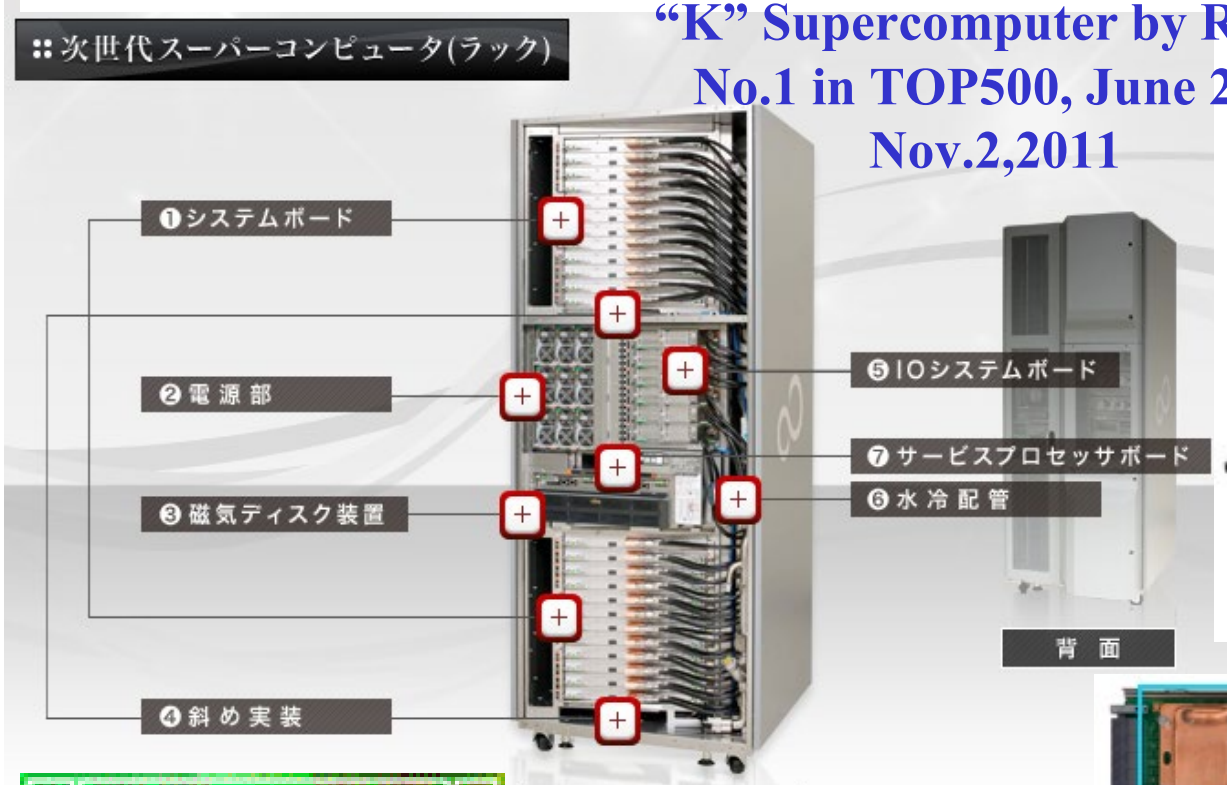

6次元メッシュ/トーラス(概念模型)
TOFUネットワーク

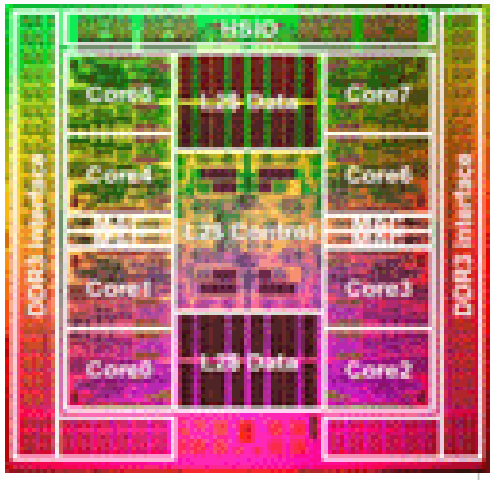

SPARC64™ VIIIfx
(提供:富士通(株))

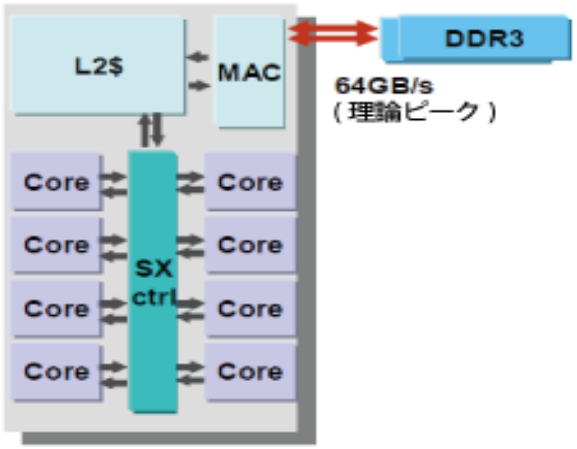

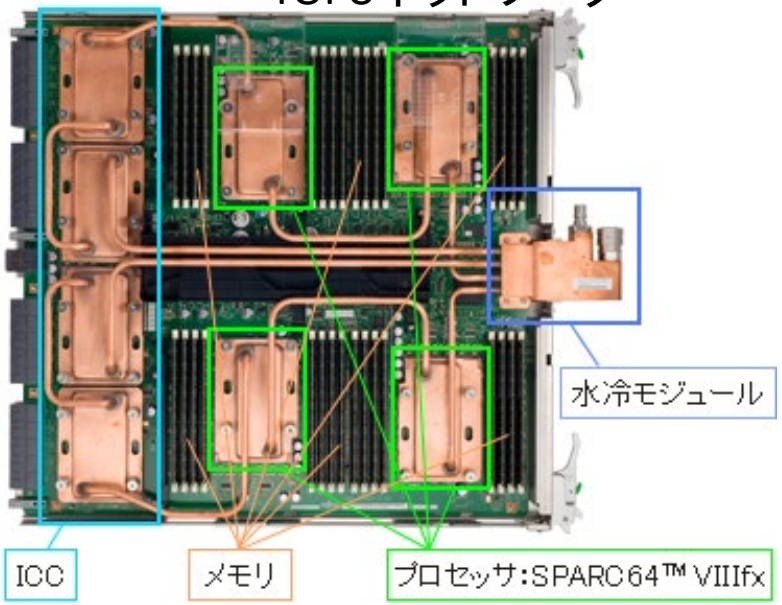

OSCAR Parallelizing Compiler

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

Multigrain Parallelization(LCPC1991,2001,04)

coarse-grain parallelism among loops and subroutines (2000 on SMP), near fine grain parallelism among statements (1992) in addition to loop parallelism

Data Localization

Automatic data management for distributed shared memory, cache and local memory
(Local Memory 1995, 2016 on RP2,Cache2001,03)
Software Coherent Control (2017)

Data Transfer Overlapping(2016 partially)

Data transfer overlapping using Data Transfer Controllers (DMAs)

Power Reduction

(2005 for Multicore, 2011 Multi-processes, 2013 on ARM)

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

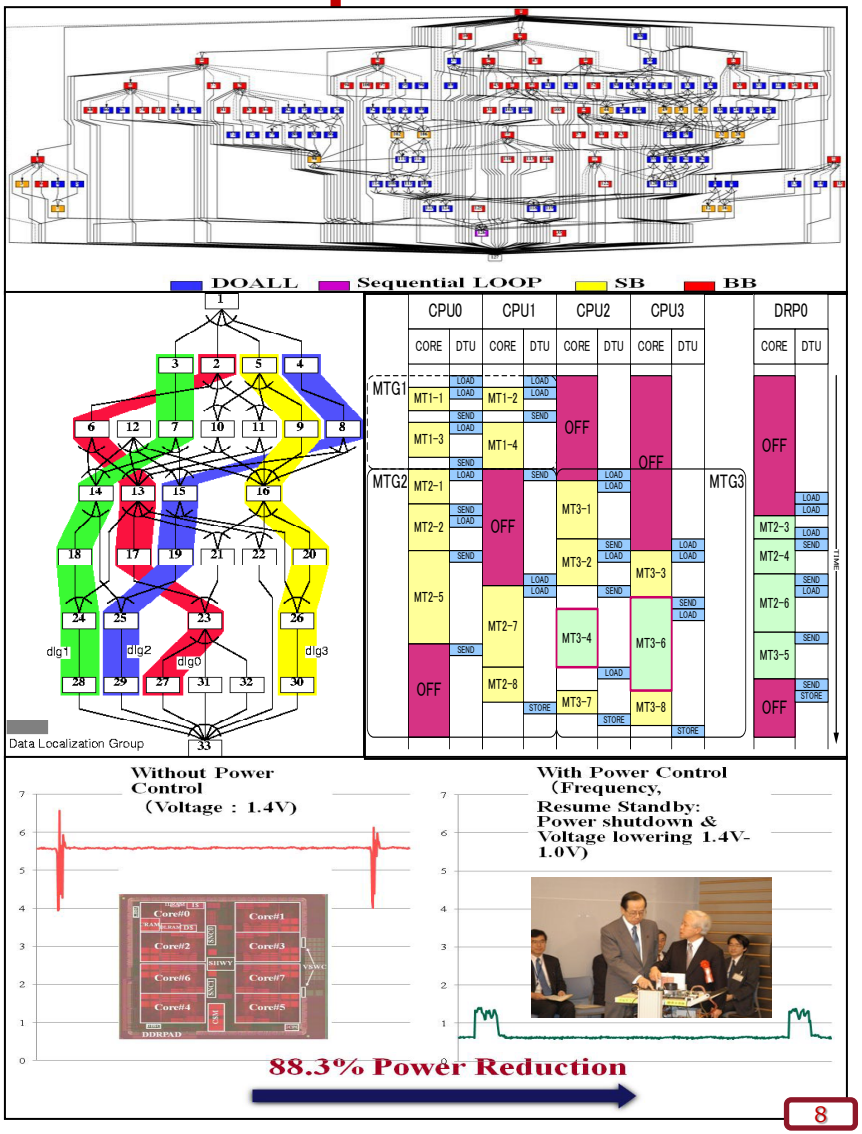

Performance of APC Compiler on IBM pSeries690 16 Processors High-end Server

- 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)

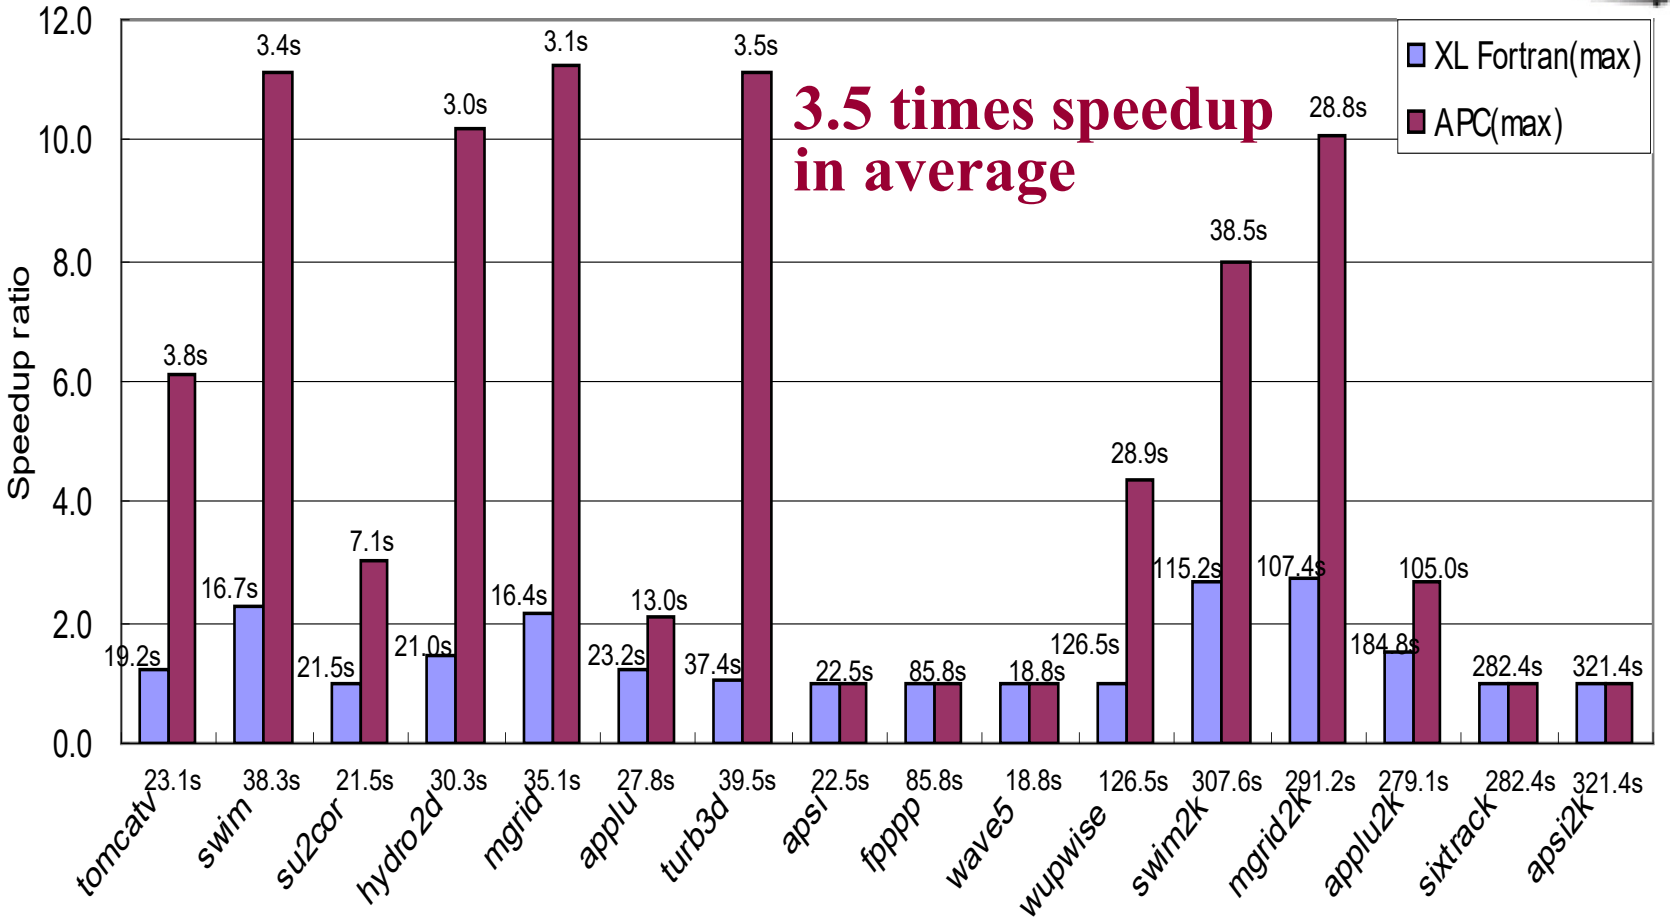

Performance of Multigrain Parallel Processing for 102.swim on IBM pSeries690

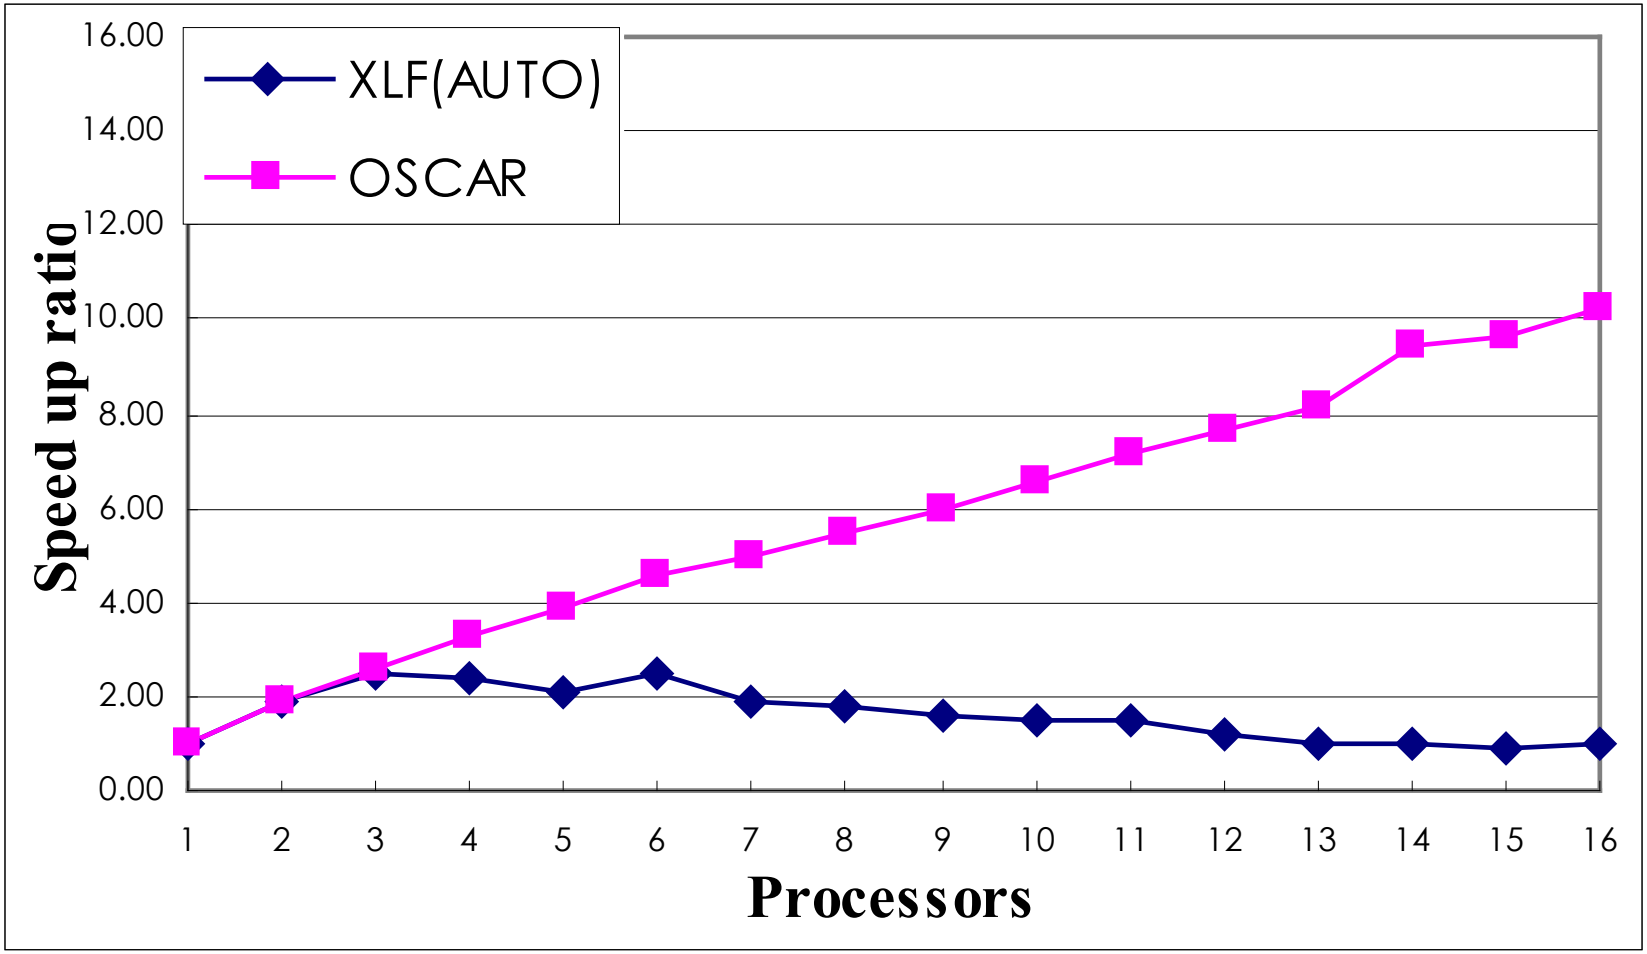

NEC/ARM MPCore Embedded 4 core SMP

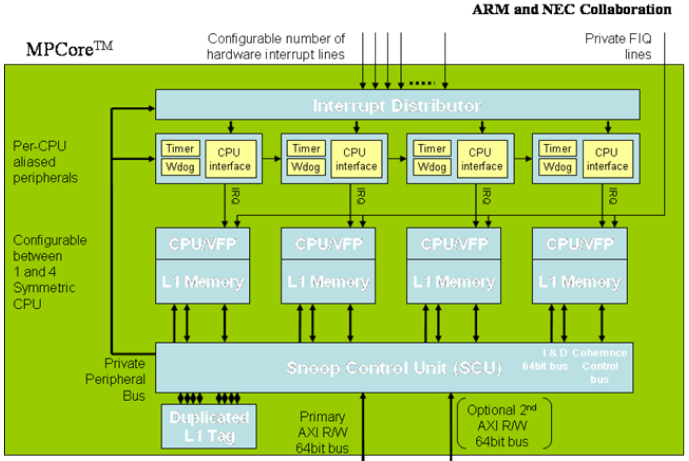

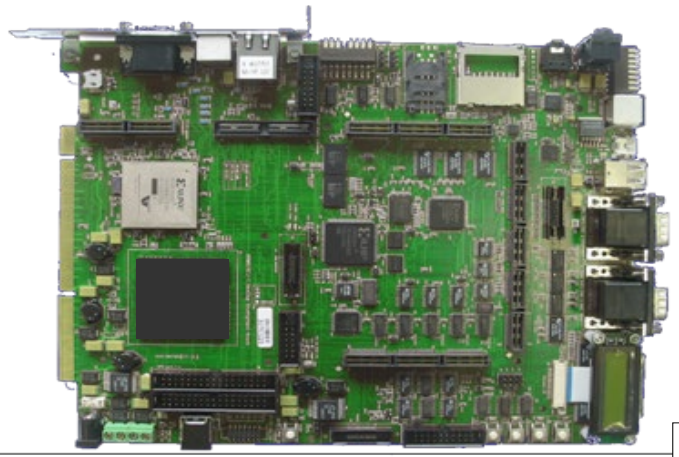

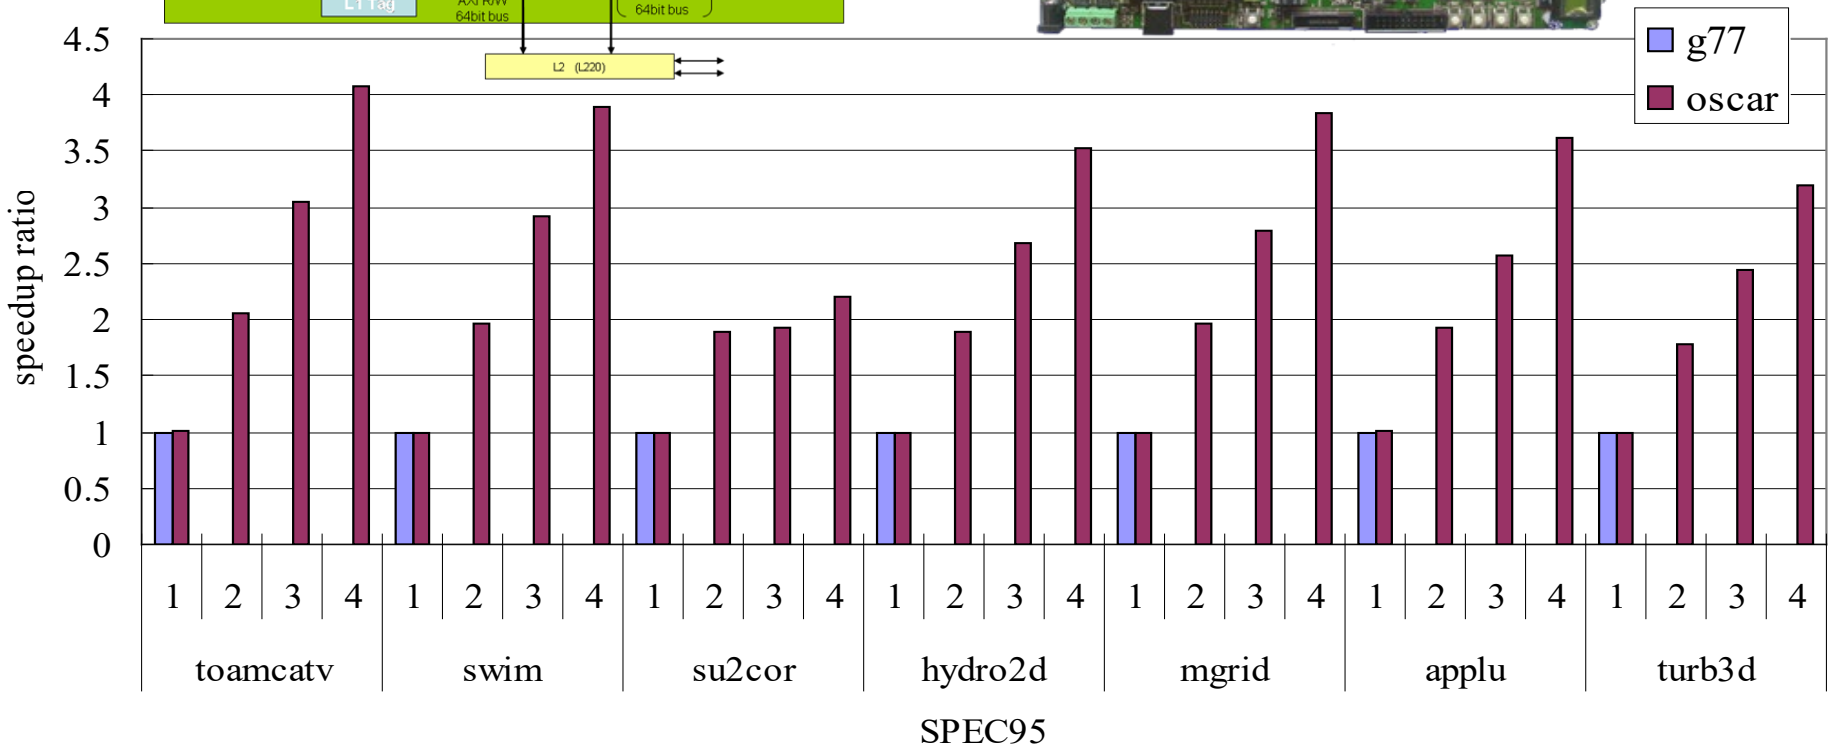

3.48 times speedup by OSCAR compiler against sequential processing

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

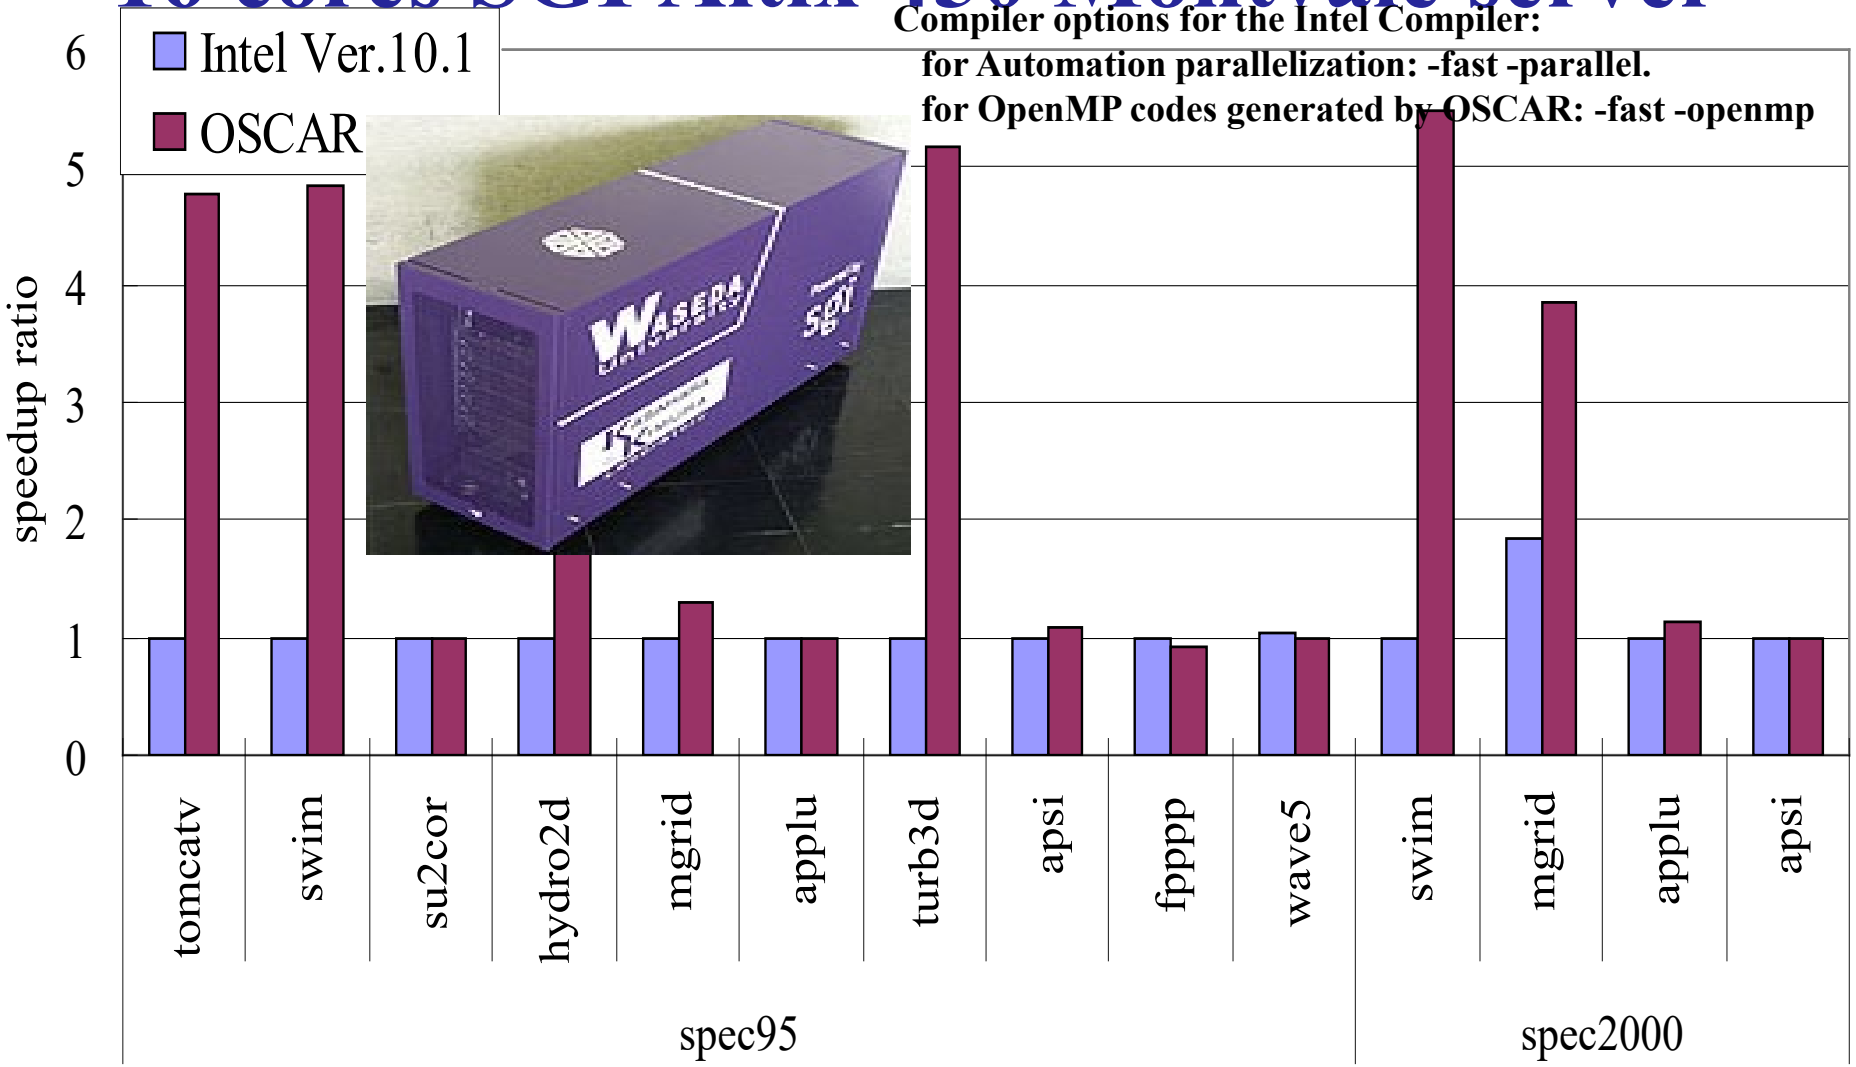

- **OSCAR compiler gave us 2.32 times speedup against Intel Fortran Itanium Compiler revision 10.1**

Distributed Shared Memory Architecture (ccNUMA)
分散共有メモリ型アーキテクチャ（ccNUMA）

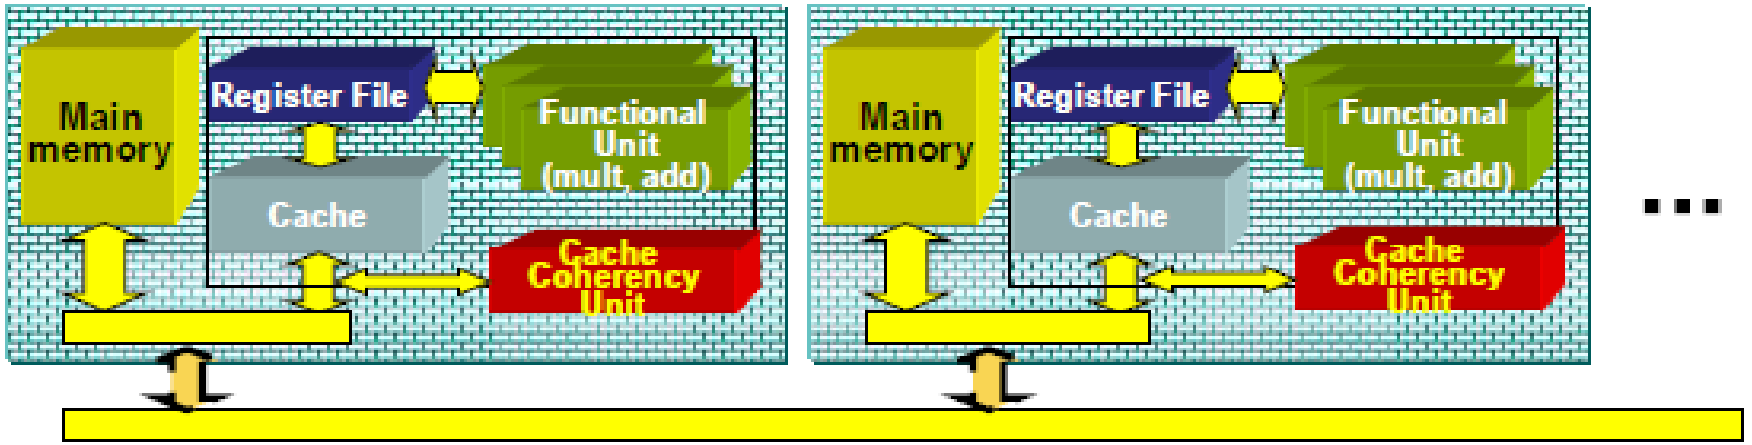

- Each processor has Local Memory.
- All memory is logically shared.
- Local memory access is faster than remote memory access.

– プロセッサはそれぞれローカルメモリを持つ
– 全てのメモリで論理的に“共有メモリ”を構成
– Non-Uniform Memory Access：cache-coherent NUMA (ccNUMA)
  • ローカルメモリへのアクセスはリモートメモリへのアクセスより速い
– 共有メモリ型アーキテクチャのように自動並列化が可能
– 分散メモリ型アーキテクチャのように拡張性がある

- Programing is easy like shared memory architectures.
- Scalable like distributed memory architectures.

2007/5/21
SGI PROPRIETARY

sgi INNOVATION FOR RESULTS

D. Lenoski, J. Laudon, K Gharachorloo, A. Gupta, J. Hennessy, “The directory-based cache coherence protocol for the DASH multiprocessor,” ACM SIGARCH Computer Architecture News 18 (2SI), 148-159, 1990

H. Kasahara, "OSCAR Fortran Multigrain Compiler", Stanford University, Hosted by Professor John L. Hennessy and Professor Monica Lam, May. 15. 1995.

What Are the Most Cited ISCA Papers?

by David Patterson on Jun 15, 2023 https://www.sigarch.org/what-are-the-most-cited-isca-papers/

Rank	Citations	Year	Title (★ means it won the ISCA Influential Paper Award)	First Author + HOF Authors	Type	Topic
1	5351	1995	The SPLASH-2 programs: Characterization and methodological considerations	Stephen Woo, *Anoop Gupta*	Tool	Benchmark
2	4214	2017	In-datacenter performance analysis of a Tensor Processing Unit	*Norm Jouppi, David Patterson*	Arch	Machine Learning
3	3834	2000	★ Wattch: A framework for architectural-level power analysis and optimizations	*David Brooks, Margaret Martonosi*	Tool	Power
4	3386	1993	★ Transactional memory: Architectural support for lock-free data structures	Maurice Herlihy	Micro	Parallelism
5	2690	2016	EIE: Efficient inference engine on compressed deep neural network	Song Han, *Bill Dally, Mark Horowitz*	Arch	Machine Learning
6	2620	2007	★ Power provisioning for a warehouse-sized computer	Xiaobo Fan, *Luiz Barroso*	Micro	Power
7	2507	1992	Active messages: a mechanism for integrated communication and computation	Thorsten von Eiken	Micro	Parallelism
8	2391	2011	Dark silicon and the end of multicore scaling	*Hadi Esmaeilzadeh, Doug Burger, Karthikeyan Sankaralingam*	Micro	Parallelism
9	2352	1995	★ Simultaneous multithreading: Maximizing on-chip parallelism	*Dean Tullsen, Susan Eggers, Hank Levy*	Micro	Parallelism
10	2243	1990	★ Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers	*Norm Jouppi*	Micro	Cache
11	1801	2009	Architecting phase change memory as a scalable DRAM Alternative	Benjamin Lee, *Doug Burger, Engin Ipek, Onur Mutlu*	Micro	NV RAM
12	1790	1990	Memory consistency and event ordering in scalable shared-memory multiprocessors	*Kourosh Gharachorloo, Anoop Gupta, John Hennessy*	Micro	Consistency/ Coherence
13	1769	2009	Scalable high performance main memory system using phase-change memory technology	*Moinuddin Qureshi*	Micro	NV RAM
14	1659	2016	ISAAC: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars	Ali Shafiee, *Rajeev Balasubramonian, Naveen Muralimanohar*	Arch	Machine Learning
15	1643	2003	★ Temperature-aware microarchitecture	Kevin Skadron	Micro	Power
16	1557	2016	Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks	Yu-Hsin Chen, *Joel Emer*	Micro	Machine Learning
17	1420	2016	Prime: A novel processing-in-memory architecture for neural network computation in ReRAM-based main memory	Ping Chi, *Yuan Xie*	Arch	Machine Learning
18	1401	2014	A reconfigurable fabric for accelerating large-scale datacenter services	Andrew Putnam, *Hadi Esmaeilzadeh*	Micro	Interconnect
19	1374	1992	The turn model for adaptive routing	Christopher Glass	Micro	Interconnect
20	1350	1995	Multiscalar processors	*Guri Sohi, T. N. Vijaykumar*	Micro	Parallelism
21	1302	2000	Memory access scheduling	Andrew Putnam, *Bill Dally*	Micro	Parallelism
22	1284	1997	★ Complexity-effective superscalar processors	Subbarao Palacharla, *Norm Jouppi, Jim Smith*	Micro	Parallelism
23	1221	2002	★ Drowsy caches: simple techniques for reducing leakage power	Krisztián Flautner, *Nam Sung Kim, Trevor Mudge*	Micro	Power
24	1210	1996	★ Exploiting choice: Instruction fetch and issue on an implementable simultaneous multithreading processor	*Hank Levy, Susan Eggers, Joel Emer, Dean Tullsen*	Micro	Parallelism
25	1201	1997	A Study of Branch Prediction Strategies	*Jim Smith*	Micro	Parallelism
26	1201	1997	The SGI Origin: A ccNUMA highly scalable server	James Laudon	Arch	Consistency/ Coherence
27	1177	2009	A durable and energy efficient main memory using phase change memory technology	Ping Zhou	Tool	Benchmark
28	1175	2014	Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors	Yoongu Kim, *Onur Mutlu, Chris Wilkerson*	Micro	Security
29	1166	2010	Debunking the 100X GPU vs. CPU myth: an evaluation of throughput computing on CPU and GPU	Victor Lee	Tool	Simulator
30	1104	2017	SCNN: An accelerator for compressed-sparse convolutional neural networks	Angshuman Parashar, *Joel Emer, Bill Dally, Steve Keckler*	Arch	Machine Learning
31	1070	2015	ShiDianNao: Shifting vision processing closer to the sensor	Zidong Du	Arch	Machine Learning
32	1051	1994	★ The Stanford FLASH multiprocessor	Jeffrey Kuskin, *Kourosh Gharachorloo, Anoop Gupta, John Hennessy, Mark Horowitz*	Arch	Parallelism
33	1027	2004	★ Transactional memory coherence and consistency	Lance Hammond, *Christos Kozyrakis, Kunle Olukotun*	Micro	Consistency/ Coherence
34	1021	1992	Lazy release consistency for software distributed shared memory	Pete Keleher	Micro	Consistency/ Coherence
35	1006	2001	Cache decay: Exploiting generational behavior to reduce cache leakage power	*Stefanos Kaxiras, Margaret Martonosi*	Micro	Power
36	993	1990	The directory-based cache coherence protocol for the DASH multiprocessor	Daniel Lenoski, *Kourosh Gharachorloo, Anoop Gupta, John Hennessy*	Micro	Consistency/ Coherence
37	991	2000	Clock rate versus IPC: The end of the road for conventional microarchitectures	Vikas Agarwal, *Doug Burger, Steve Keckler*	Micro	Parallelism
38	980	1990	Weak ordering—a new definition	*Sarita Adve, Mark Hill*	Micro	Consistency/ Coherence
39	957	2008	Technology-driven, highly-scalable dragonfly topology	*John Kim, Bill Dally*	Micro	Interconnect
40	938	2007	Adaptive insertion policies for high performance caching	*Moinuddin Qureshi, Joel Emer, Yale Patt*	Micro	Cache
41	935	1973	Banyan networks for partitioning multiprocessor systems	Rodney Goke, *Jack Lipovski*	Micro	Interconnect
42	911	2008	3D-Stacked Memory Architectures for Multi-core Processors	*Gabriel Loh*	Micro	Cache
43	895	2010	High performance cache replacement using re-reference interval prediction (RRIP)	*Aamer Jaleel, Joel Emer*	Micro	Cache
44	877	2015	A scalable processing-in-memory accelerator for parallel graph processing	Junwhan Ahn, *Onur Mutlu*	Arch	Parallelism
45	873	2000	Transient fault detection via simultaneous multithreading	*Steven Reinhardt*	Micro	Reliability
46	868	2008	Corona: System implications of emerging nanophotonic technology	Dana Vantrease, *Norm Jouppi*	Micro	Interconnect
47	863	2009	An analytical model for a GPU architecture with memory-level and thread-level parallelism awareness	Sunpyo Hong	Tool	Parallelism
48	859	2004	Single-ISA heterogeneous multi-core architectures for multithreaded workload performance	*Rakesh Kumar, Norm Jouppi, Partha Ranganathan & Dean Tullsen*	Micro	Parallelism
49	854	1974	A Preliminary Architecture for a Basic Data Flow Processor	Jack Dennis	Arch	Parallelism
50	854	1983	Very Long Instruction Word Architectures and the ELI-512	Josh Fisher	Arch	Parallelism

スーパーテクニカルサーバ「SR16000モデルVM1」

Green Computing
Research Center Project

Kasahara & Kimura Laboratory

CPU : POWER7（4.0GHz）×128Core
Memory : 1024GB
Peak Perfromance:4.1TFLOPS
Electricity Consumption : 256.8MFLOPS/W
OS: Red Hat Enterprise Linux 6

【用途】
・並列化コンパイラ開発
・低消費電力メニーコア開発
・屋上太陽光発電を用いた
グリーンコンピューティング
・災害シミュレーション

【Purpose】
・Parallelizing compiler development
・Development of low power
manycore processors
・Green Computing
by rooftop solar power
・Disaster Simulation

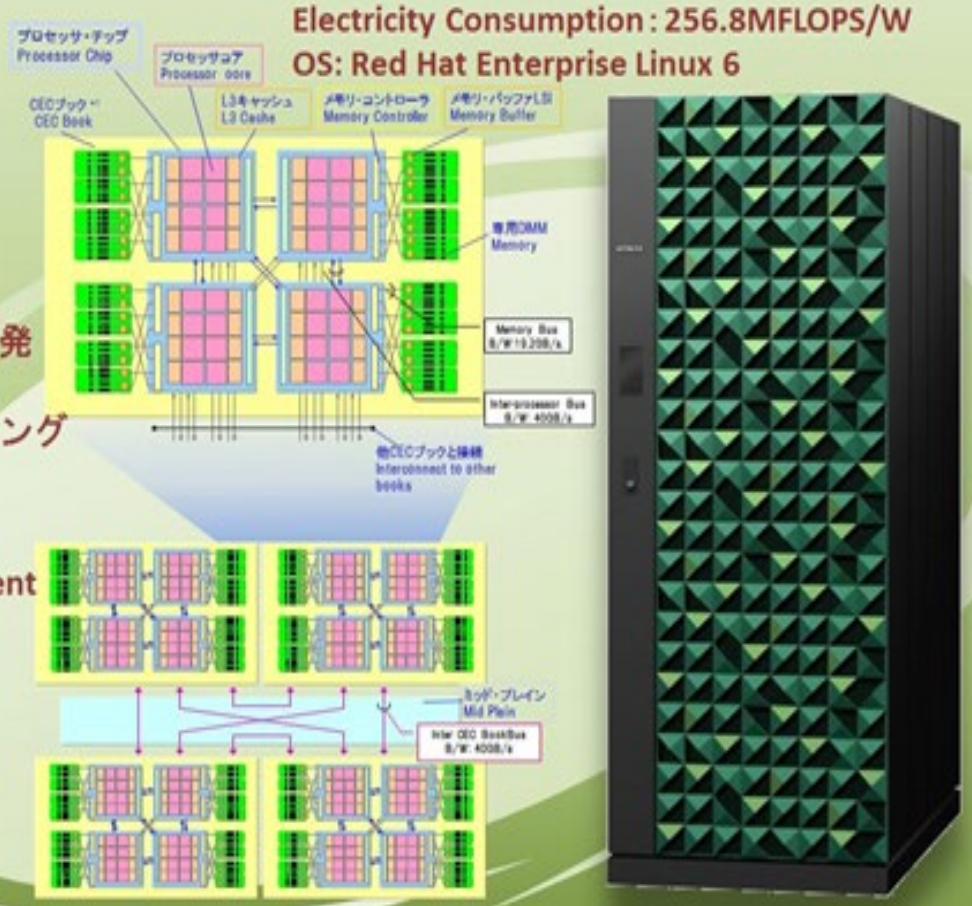

HITACHI

WASEDA Univ.

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

（Power7 Based 128 Core Linux SMP） (LCPC2015)

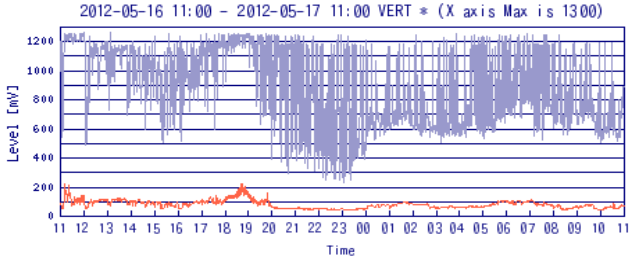

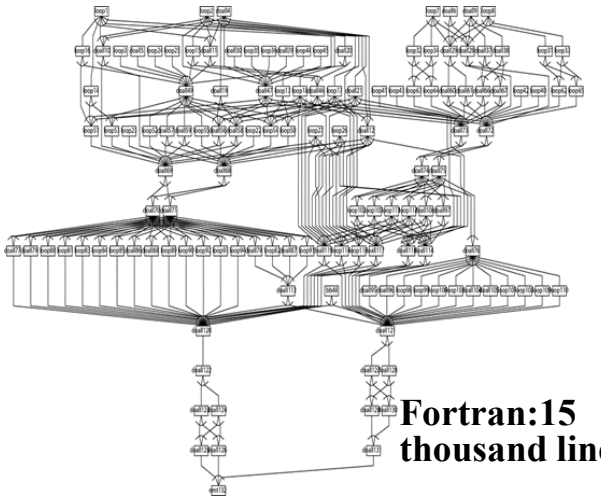

Fortran:15 thousand lines

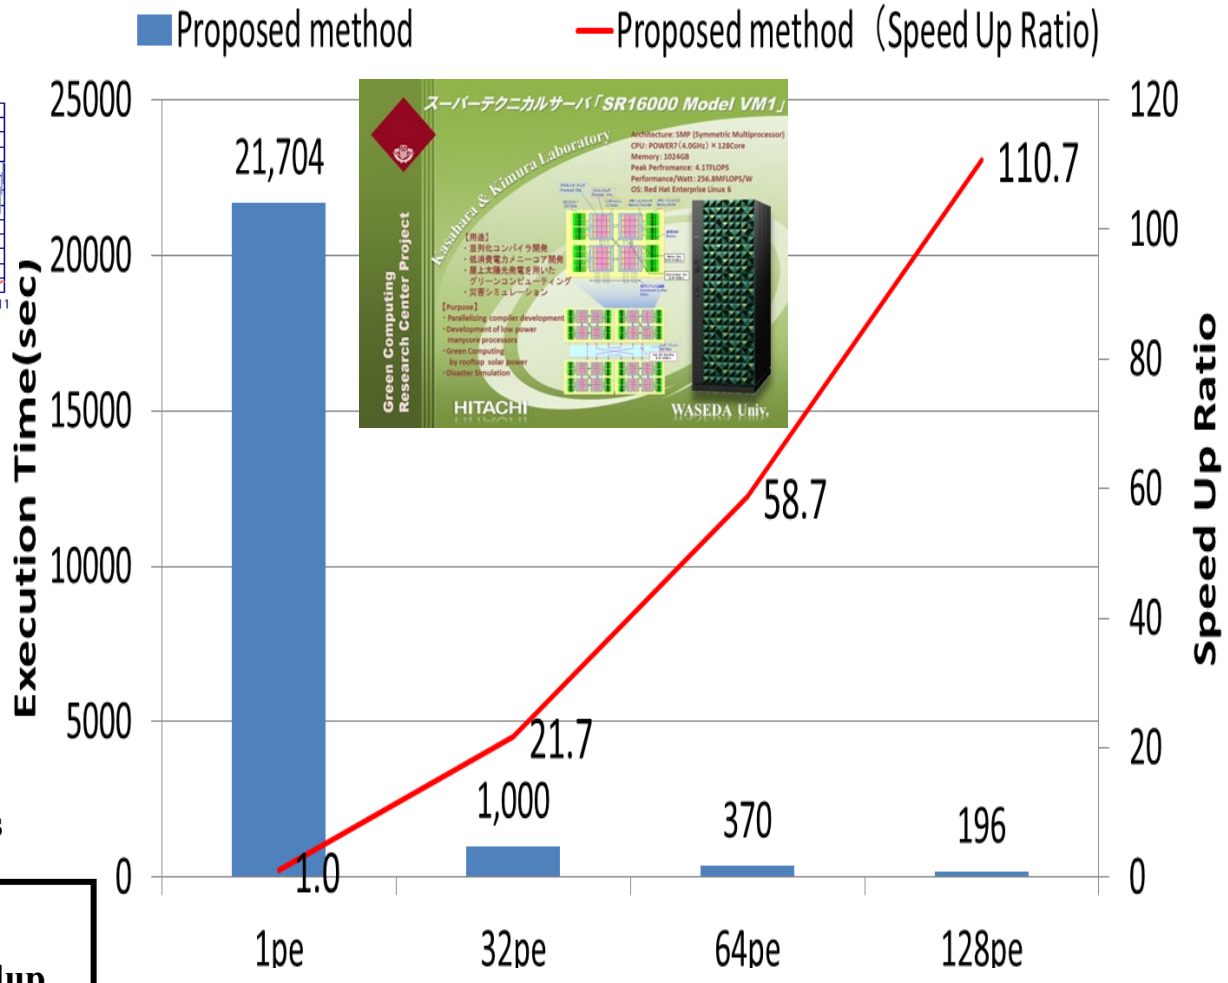

First touch for distributed shared memory and cache optimization over loops are important for scalable speedup

RIKEN/Fujitsu 富岳(Fugaku) : Tofu Interconnect

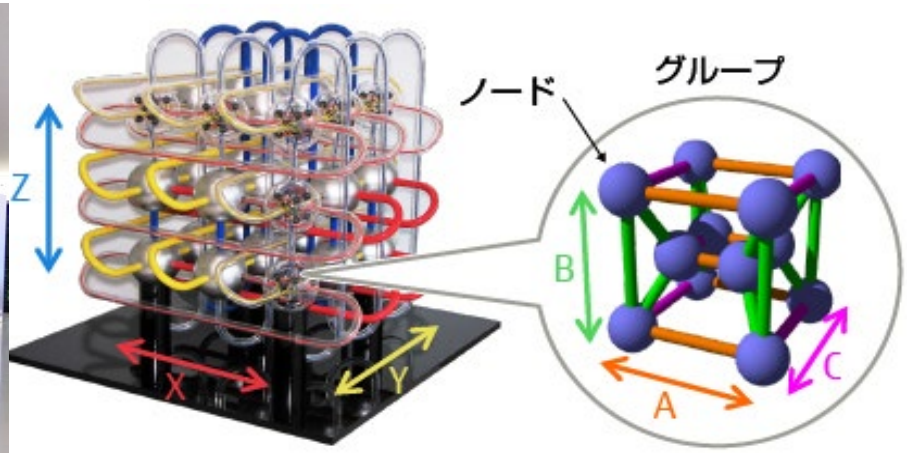

6次元メッシュ/トーラス (6D Mesh/Torus)

https://www.fujitsu.com/jp/about/businesspolicy/tech/fugaku/

理研富岳スーパーコンピュータ 2020年6月から2021年11月まで世界No.1

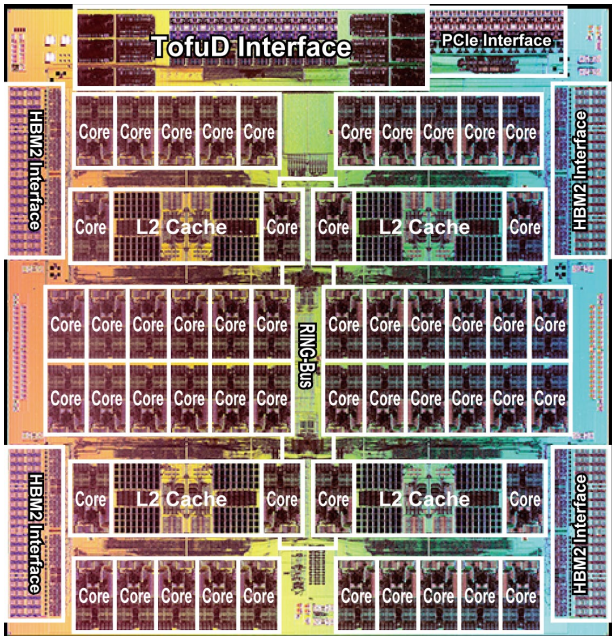

RIKEN Center for Computational Science,
Fujitsu（arm based processor）
Cores:7,299,072； Memory:4,866,048GB；
Processor:A64FX 48Cores, 2.2GHz
Interconnect: Tofu interconnect D
Linpack (Rmax)415,530 TFlop/s;
Theoretical Peak (Rpeak): 513PFLOPS
HPCG [TFlop/s]13,366.4; Power: 28.3MW

48コア/チップ, 2.2GHz, 7 nm FinFET,
約730万コア, 28MW
理論最高性能:51京回浮動小数点演算/秒, 2020年6月時点

https://www.r-ccs.riken.jp/en/fugaku/about/

https://japanese.engadget.com/arm-super-computer-fugaku-top-500-034015910.html

Fugaku A64FX

TSMCの7 nm CMOS Process

SVE（Scalable Vector Extension）: 2 Floating Units（FLA/FLB）:
512 bit SIMD, Multiply & Add operations/ Cycle,
1,536 Double Precision Floating Operations Per Cycle on a chip

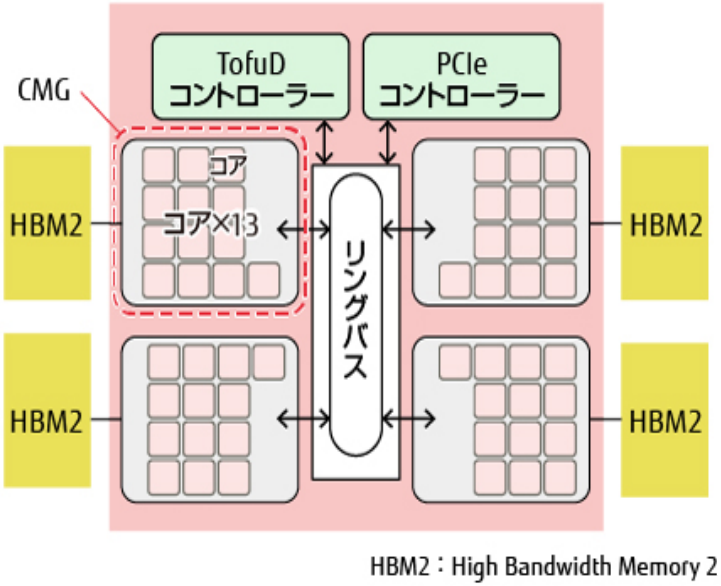

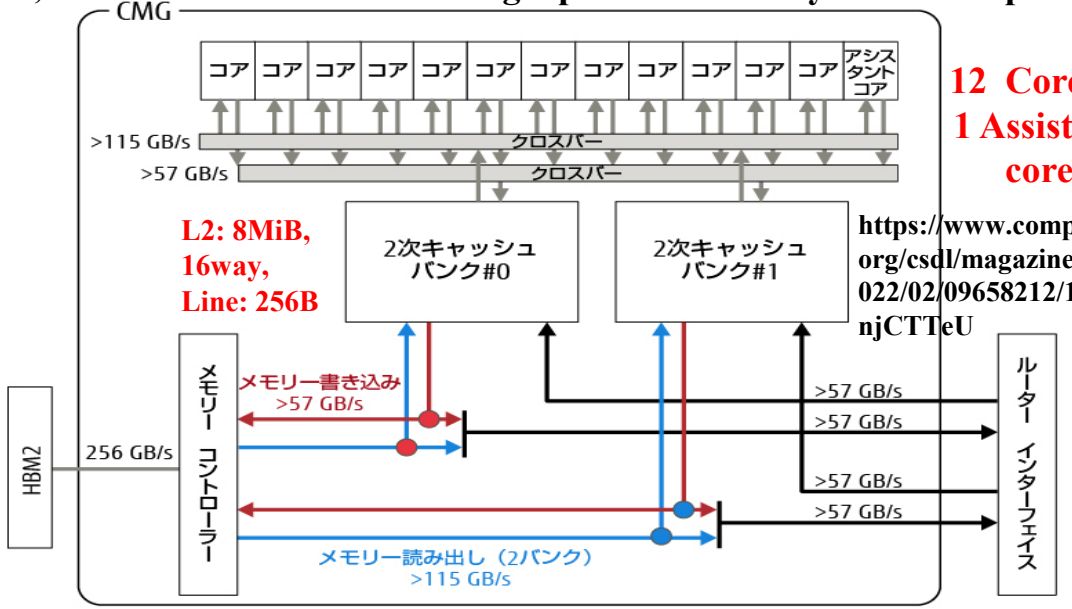

Memory Architecture of A64FX CPU: Tag Directory based CC-NUMA

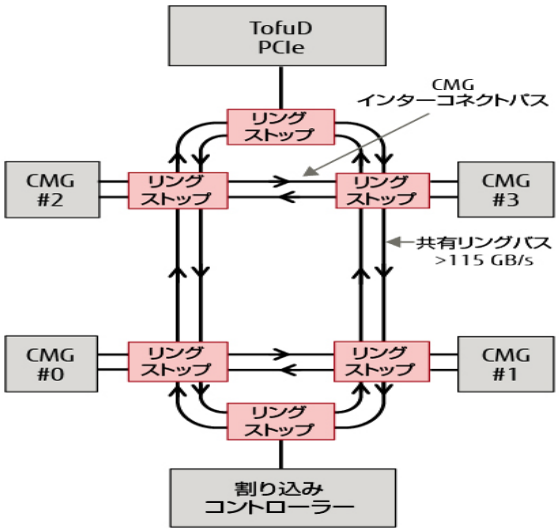

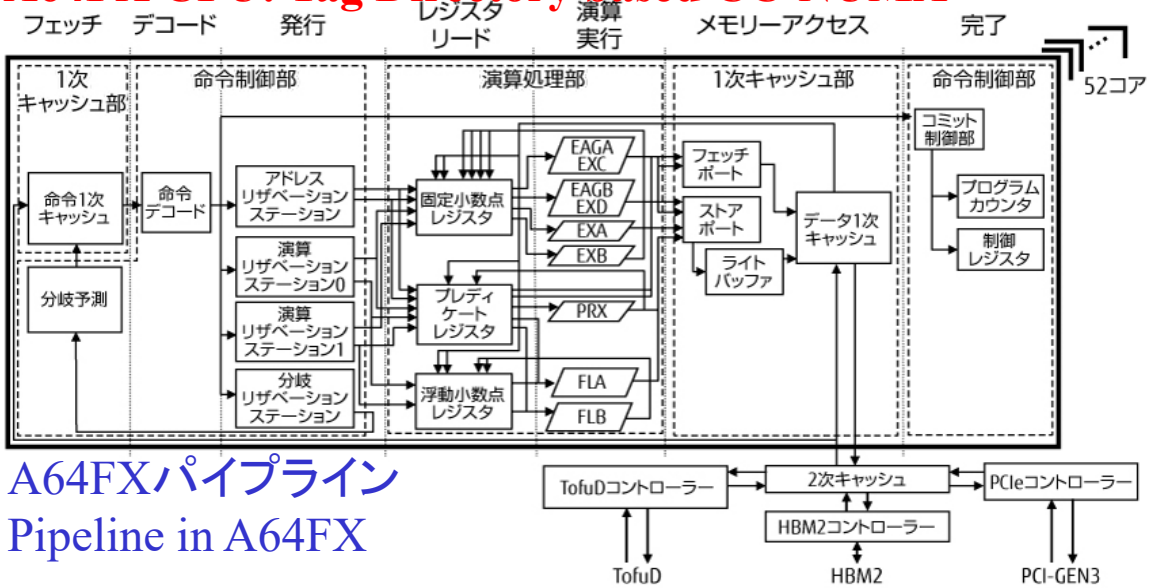

A64FXパイプライン
Pipeline in A64FX

https://www.fujitsu.com/jp/about/resources/publications/technicalreview/2020-03/article03.html

https://www.top500.org/ June 2023

PRESENTED BY

TOP 500 The List.

FIND OUT MORE AT top500.org

2021 ACM A.M. Turing Award

https://amturing.acm.org/award_winners/dongarra_3406337.cfm

For his pioneering contributions to numerical algorithms and libraries that enabled high performance computational software to keep pace with exponential hardware improvements for over four decades

JUNE 2023

			SITE	COUNTRY	CORES	RMAX PFLOP/S	POWER MW
1	Frontier	HPE Cray EX235a, AMD Opt 3rd Gen EPYC (64C 2GHz), AMD Instinct MI250X, Slingshot-11	DOE/SC/ORNL	USA	8,699,904	1,194.0	22.7
2	Fugaku	Fujitsu A64FX (48C, 2.2GHz), Tofu Interconnect D	RIKEN R-CCS	Japan	7,630,848	442.0	29.9
3	LUMI	HPE Cray EX235a, AMD Opt 3rd Gen EPYC (64C 2GHz), AMD Instinct MI250X, Slingshot-11	EuroHPC/CSC	Finland	2,220,288	309.0	6.01
4	Leonardo	Atos Bullsequana intelXeon (32C, 2.6 GHz), NVIDIA A100 quad-rail NVIDIA HDR100 Infiniband	EuroHPC/CINEC	Italy	1,824,768	238.7	7.40
5	Summit	IBM POWER9 (22C, 3.07GHz), NVIDIA Volta GV100 (80C), Dual-Rail Mellanox EDR Infiniband	DOE/SC/ORNL	USA	2,414,592	148.6	10.1

No. 1 June 2022-23 870万プロセッサ, 22.7MW

Frontier - HPE Cray EX235a, 8,699,904 total cores, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X accelerators,

HPE Slingshot-11 interconnect 168京回演算/秒

Oak Ridge National Laboratory (ORNL) , USA

Rmax: 1.19 (ExaFlop/s), Rpeak 1.68(ExaFlop/s)

HPCG:14,054[TFlop/s]

ムーアの法則の終焉 (End of Moore's Law)

ムーアの法則(Moore's law): インテル創業者の一人であるゴードン・ムーア(Gordon E. Moore: IEEE Computer Pioneer Award)が、1965年の論文で提唱した経験則:
"半導体の集積率は18か月で2倍になる" "Transistors on a chip 2× every 1.5 year"

コンピュータの高性能化と低消費電力化にはマルチコアが必須

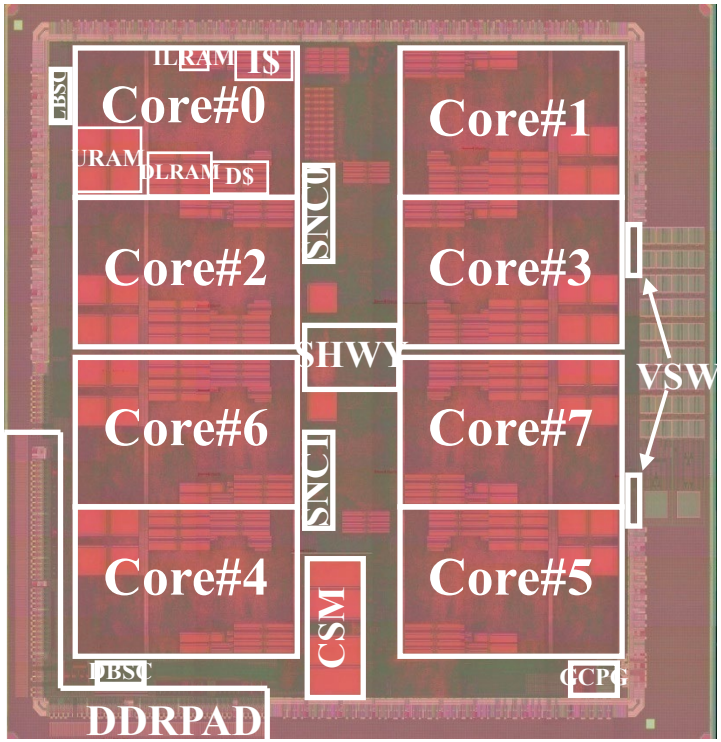

IEEE ISSCC08: Paper No. 4.5, M.ITO, … and H. Kasahara, "An 8640 MIPS SoC with Independent Power-off Control of 8 CPUs and 8 RAMs by an Automatic Parallelizing Compiler"

$$\text{Power} \propto \text{Frequency} * \text{Voltage}^2$$

⇒ (Voltage ∝ Frequency)

$$\text{Power} \propto \text{Frequency}^3$$

周波数 Frequency を1/4にすると
(Ex. 4GHz→1GHz),
性能は 1/4、消費電力は 1/64
Performance 1/4, Power 1/64

<マルチコア(Multicore)>
8cores をチップに集積すると,
Performance 2 timesでPowerは 1/8

8 Core RP2 Chip Block Diagram

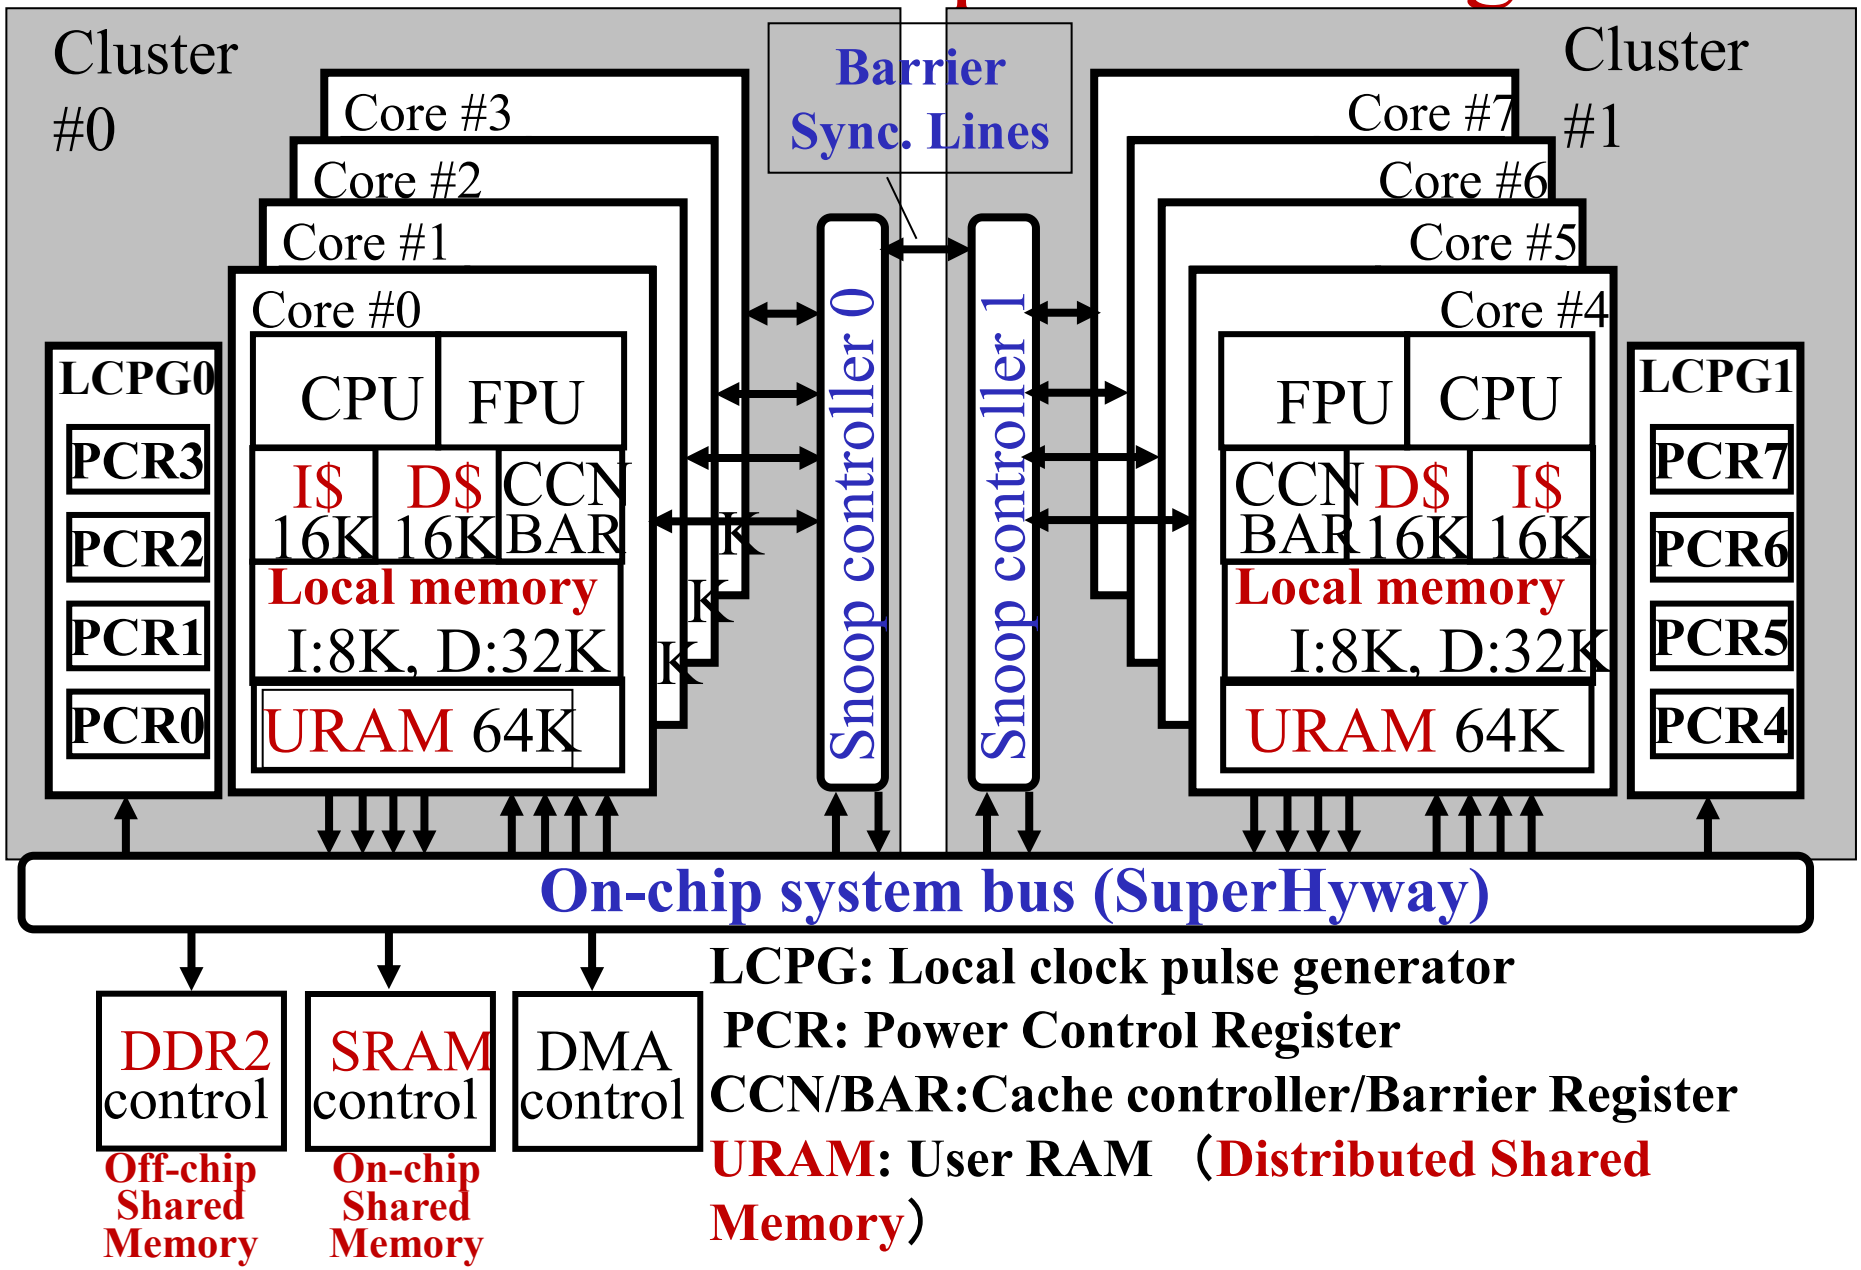

LCPG: Local clock pulse generator

PCR: Power Control Register

CCN/BAR:Cache controller/Barrier Register

URAM: User RAM （Distributed Shared Memory）

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

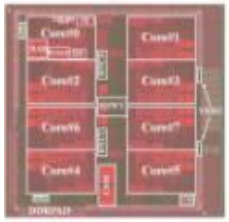

MPEG2 Decoding with 8 CPU cores

Without Power Control (Voltage : 1.4V)

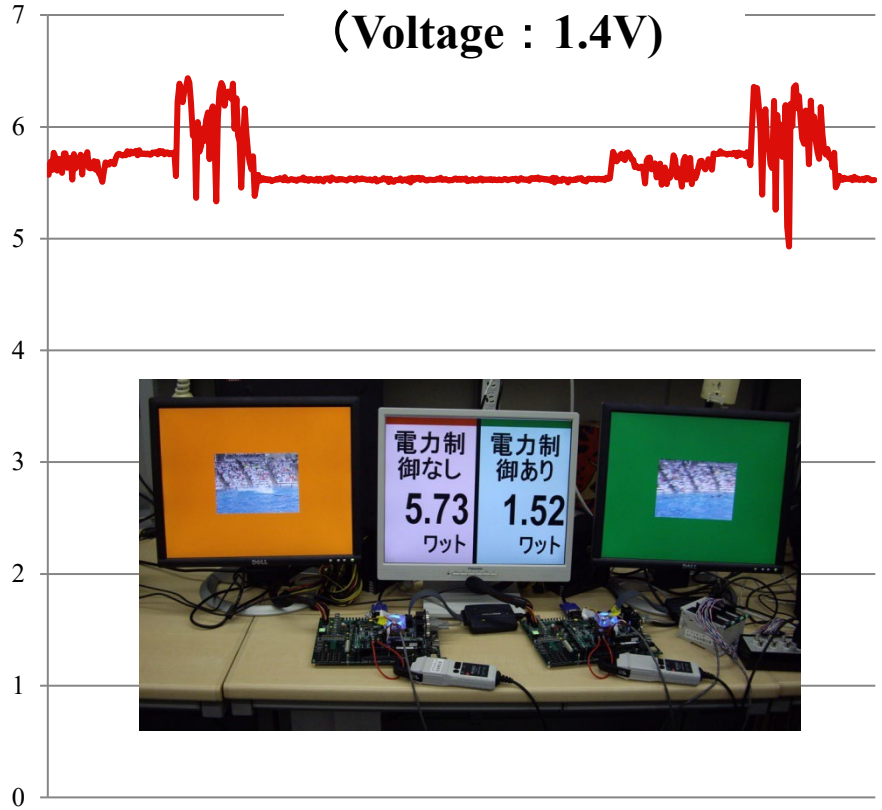

With Power Control (Frequency, Resume Standby: Power shutdown & Voltage lowering 1.4V-1.0V)

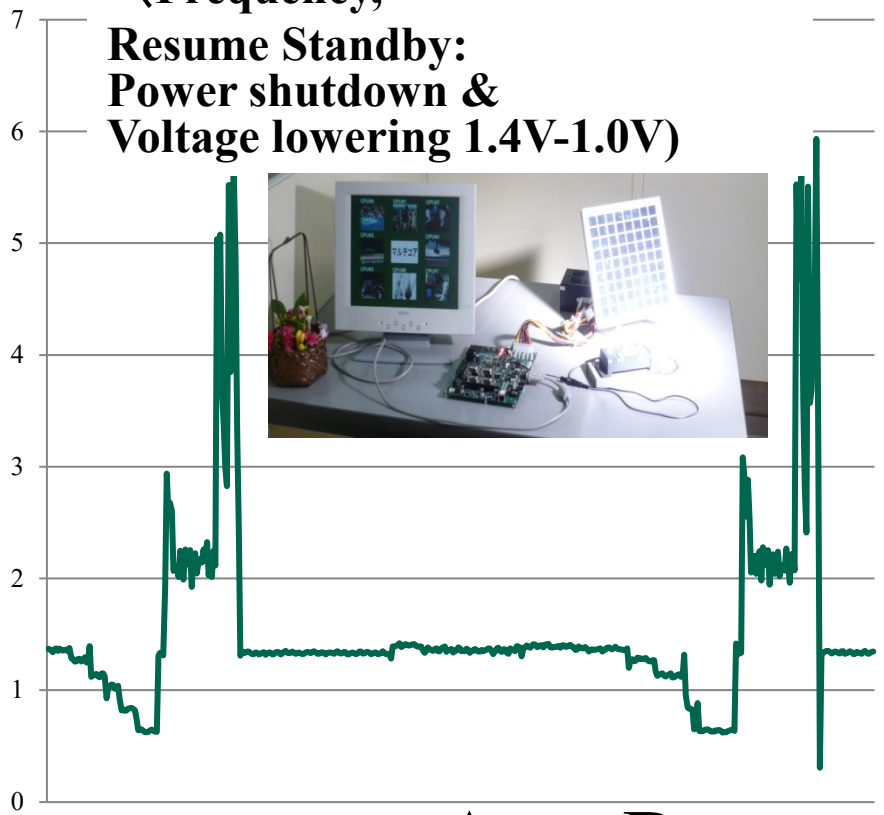

Avg. Power 5.73 [W]

73.5% Power Reduction

Avg. Power 1.52 [W]

Speedups & Power Reduction on RP-X Heterogeneous Multicore with 8 CPUs and 4 DRPs

33 Times Speedup Using OSCAR Compiler and API on Renesas RP-X with 8 CPUs & 4 DRP Accelerators

(Optical Flow with a hand-tuned library)

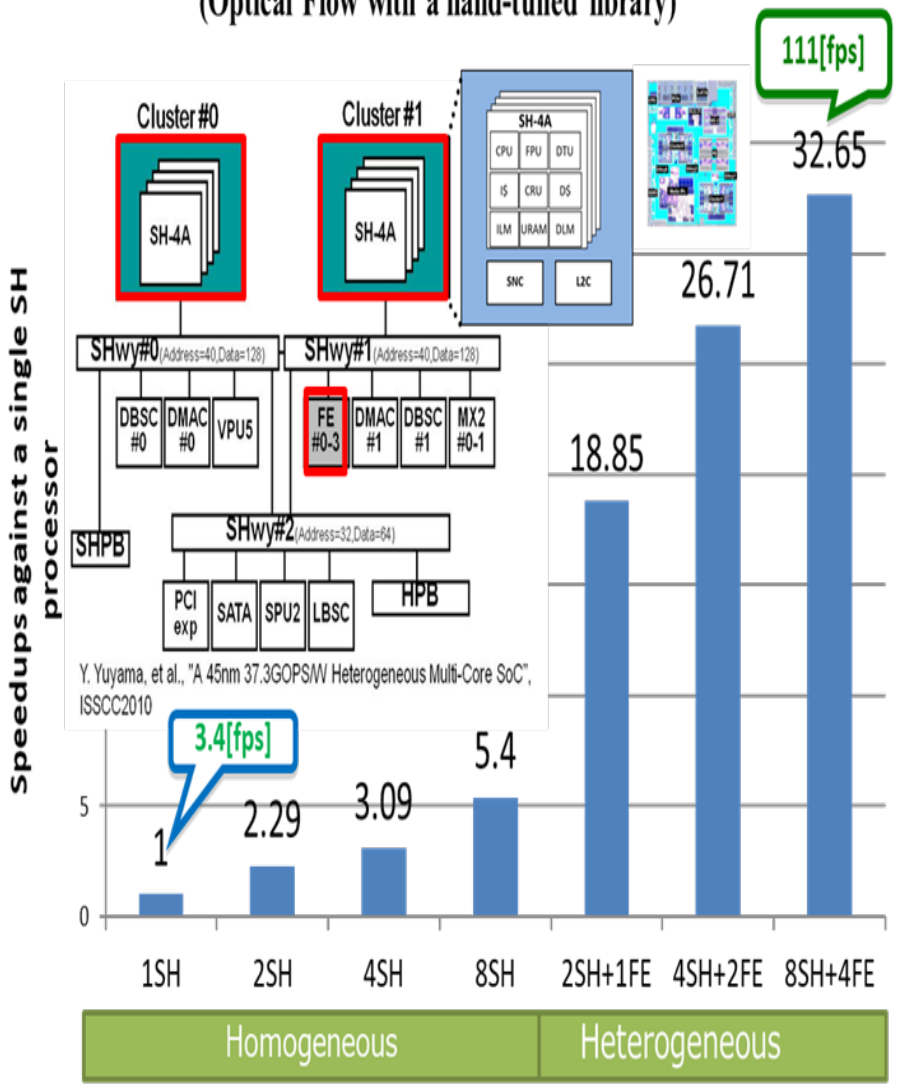

Power Reduction in a real-time execution controlled by OSCAR Compiler and OSCAR API on RP-X (Optical Flow with a hand-tuned library)

Without Power Reduction

With Power Reduction by OSCAR Compiler

70% of power reduction

Average:1.76[W]

Average:0.54[W]

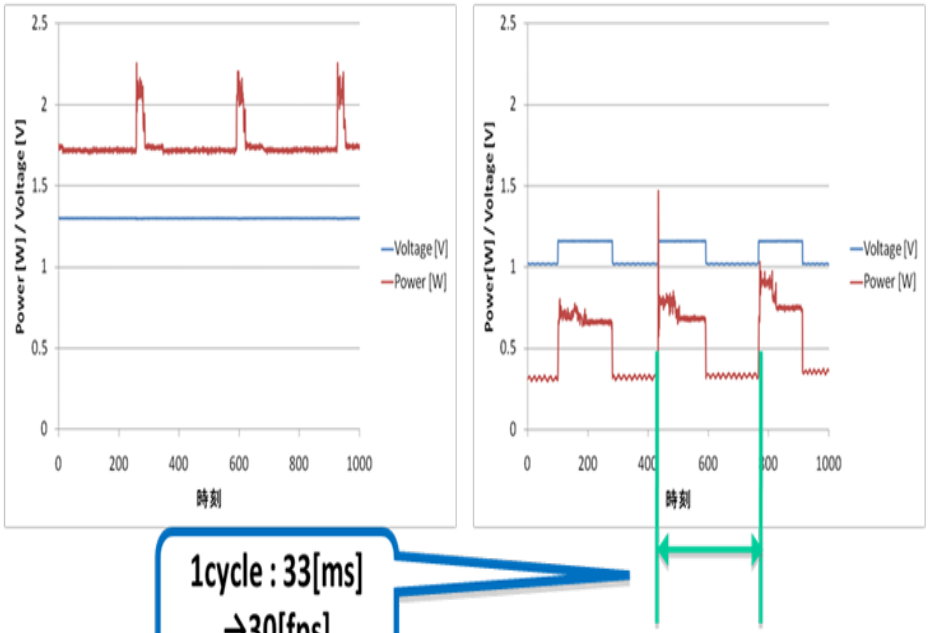

1cycle : 33[ms]
→30[fps]

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

http://www8.cao.go.jp/cstp/gaiyo/honkaigi/74index.html

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)

Codesign of Compiler and Multiprocessor Architecture since 1985

4 core multicore RP1 (2007), 8 core multicore RP2 (2008) and 15 core Heterogeneous multicore RPX (2010) developed in NEDO Projects with Hitachi and Renesas

RP-1 (ISSCC2007 #5.3)	*RP-2* (ISSCC2008 #4.5)	*RP-X* (ISSCC2010 #5.3)
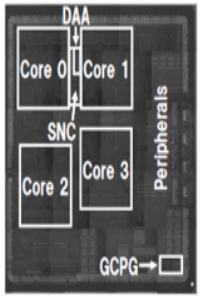	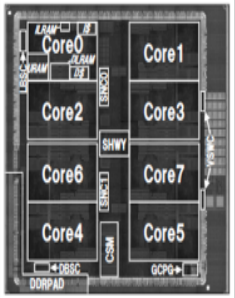	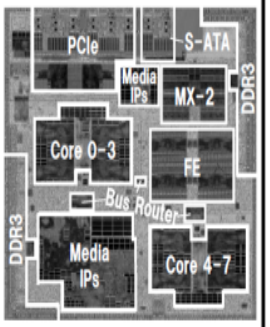
90nm, 8-layer, triple-Vth, CMOS	90nm, 8-layer, triple-Vth, CMOS	45nm, 8-layer, triple-Vth, CMOS
97.6 mm^2 (9.88 x 9.88 mm)	104.8 mm^2 (10.61 x 9.88 mm)	153.8 mm^2 (12.4 x 12.4 mm)
1.0V (internal), 1.8/3.3V (I/O)	1.0-1.4V (internal), 1.8/3.3V (I/O)	1.0-1.2V (internal), 1.2-3.3V (I/O)
600MHz ,4.32 GIPS,16.8 GFLOPS	600MHz , 8.64 GIPS, 33.6 GFLOPS	648MHz, 13.7GIPS, 115GOPS, 36.2GFLOPS
11.4 GOPS/W (32b換算)	18.3 GOPS/W (32b換算)	37.3 GOPS/W (32b換算)

Prime Minister FUKUDA is touching our multicore chip during execution.

IEEE Computer Society

The first President from outside North America in 72 years history of IEEE CS

Bjarne Stroustrup: Morgan Stanley & Columbia Univ.
2018 IEEE Computer Society Computer Pioneer Award
IEEE COMPSAC2018 Keynote & Award Ceremony

July 26, 2018, Keynote, Hitotsubashi Hall

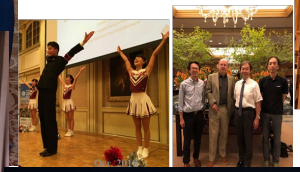

July 25, 2018 Award Ceremony Rihga Royal Hotel Tokyo

215 International Conferences

12 Magazines

35 Journals

47 Total Publications

847,000+ Articles in CSDL

6 New Standards

230 Active Standards

IEEE754, 802

12,000+ Volunteers

615 Committees/ Boards

2,352+ Meetings/ Teleconferences

373,100+ Community Members

168 Countries with CS Members

634 Chapters

Seymour Cray: Father of Supercomputers using vector pipeline

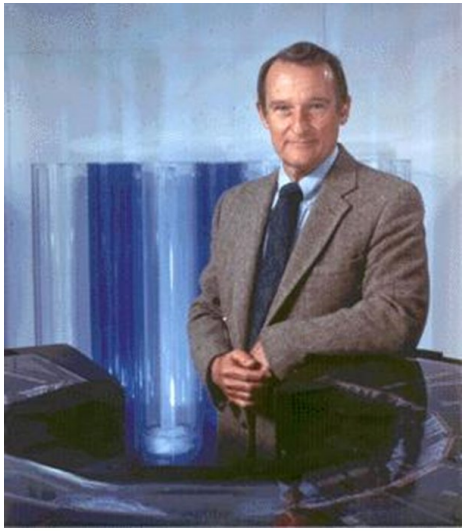

2019 Seymour Cray Award Winner:
David Kirk, NVIDIA Corporation (retired)

Cray 1

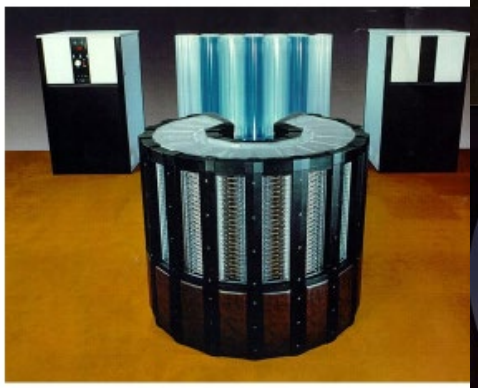

Cray 2

Cray 3

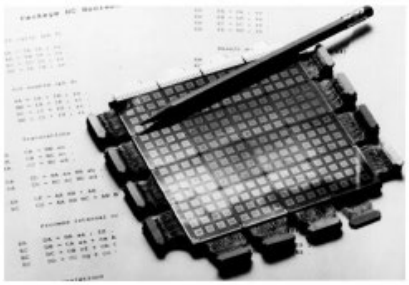

https://www.youtube.com/watch?v=Yc-VFuRWevw

IEEE COMPUTER SOCIETY

https://www.computer.org/product/education/multi-core-video-lectures-bundle

Created by Multicore STC Chair Hironori Kasahara

Multi-Core Lecture Series consists of 11 one-hour lectures by some of the world's leading researchers in the field. This series is not a course and it consists of the presentation for those who are in the research field. This is more intended for research information sharing than educational training. Topics that are covered during these lectures are listed below. This series also includes an hour discussion of the lecturers.

Video Presentations:

- Automatic Parallelization by David Padua
- Autoparallelization for GPUs by Wen-Mei Hwu
- Dependences and Dependence Analysis by Utpal Banerjee
- Dynamic Parallelization by Rudolf Eigenmann
- Instruction Level Parallelization by Alexandru Nicolau
- Multigrain Parallelization and Power Reduction by Hironori Kasahara
- The Polyhedral Model by Paul Feautrier
- Vector Computation by David Kuck
- Vectorization by P. Sadayappan
- Vectorization/Parallelization in the IBM Compiler by Yaoqing Gao
- Vectorization/Parallelization in the Intel Compiler by Peng Tu
- Roundtable Discussion by all presenters

Dependences and Dependence Analysis by Utpal Banerjee Utpal Banerjee's research interests in computer science are in the general area of parallel processing and he has published four books on loop transformations and dependence analysis, with a fifth one on instruction level parallelism on the way.

Multigrain Parallelization and Power Reduction by Hironori Kasahara. Professor Kasahara has been researching on OSCAR Automatic Parallelizing and Power Reducing Compiler and OSCAR Multicore architecture for more than 30 years, and led four Japanese national projects on parallelizing compilers, multicores, and green computing.

太陽光電力で動作する情報機器

コンピュータの消費電力をHW&SW協調で低減。電源喪失時でも動作することが可能。

リアルタイムMPEG2デコードを、8コアホモジニアス
マルチコアRP2上で、消費電力1/4に削減

世界唯一の差別化技術

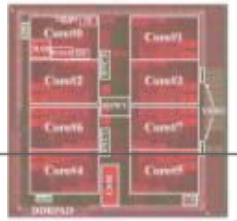

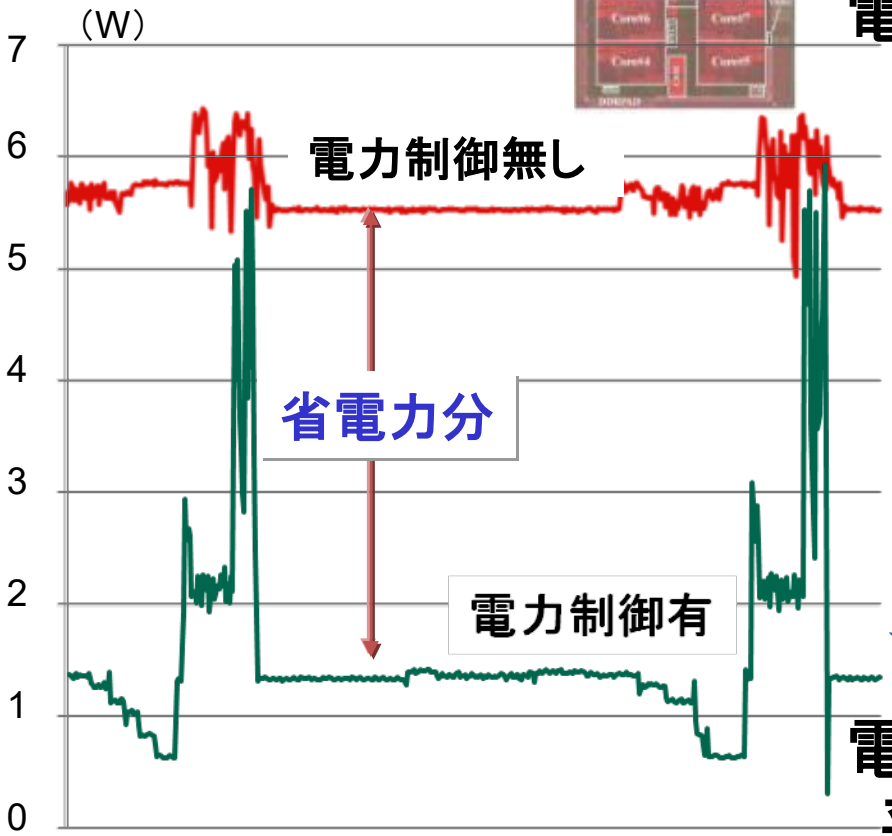

電力制御無し
平均電力
5.73 [W]

電力を
ソフトで
1/4に
削減

電力制御有
平均電力
1.52 [W]

周波数/電圧・電源遮断制御

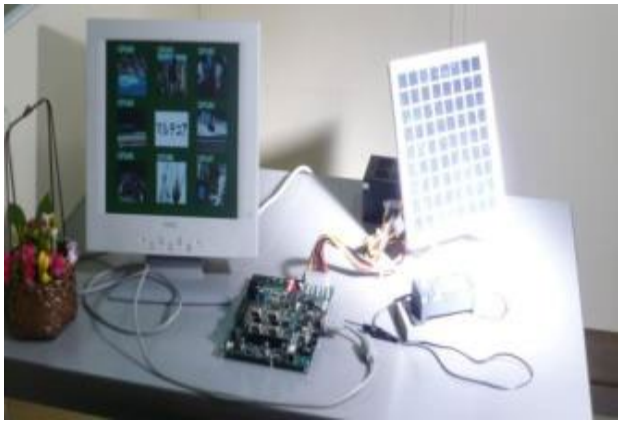

太陽電池で駆動可

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

- Shortest execution time mode

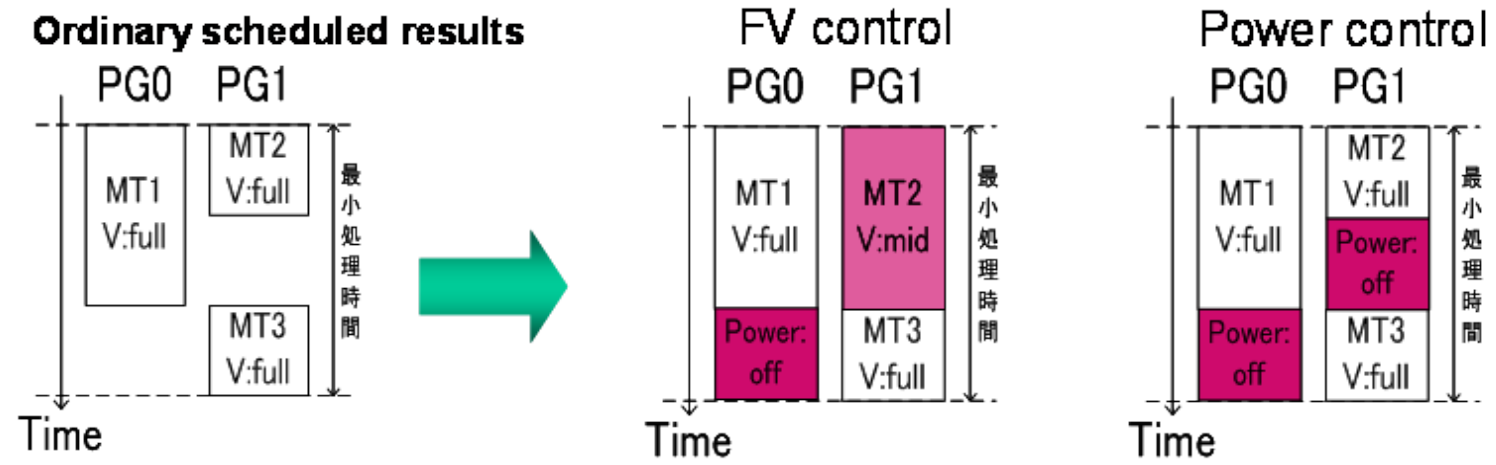

- Realtime processing mode with dead line constraints

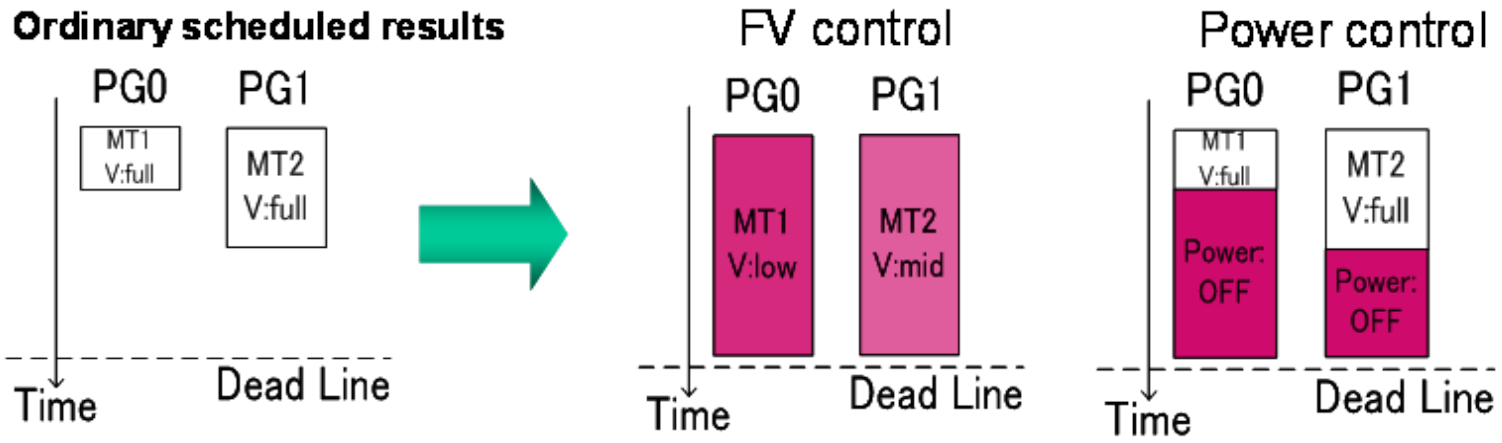

Power Reduction Scheduling

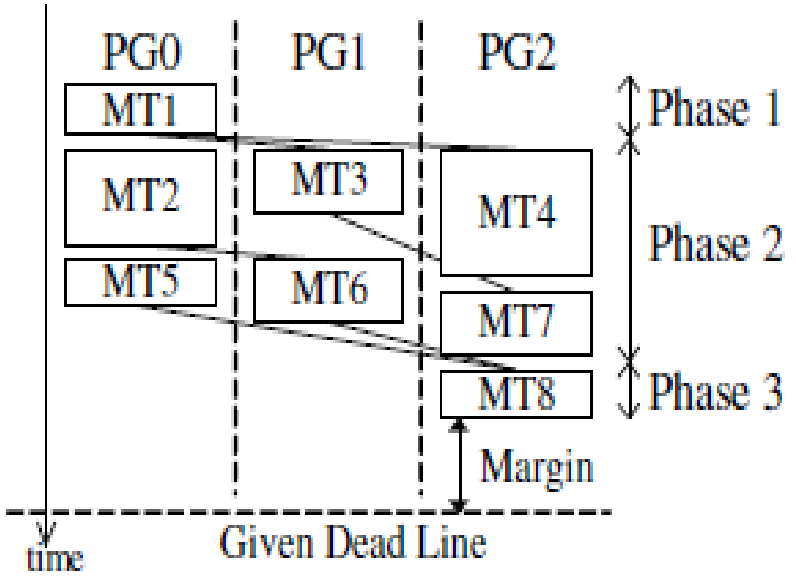

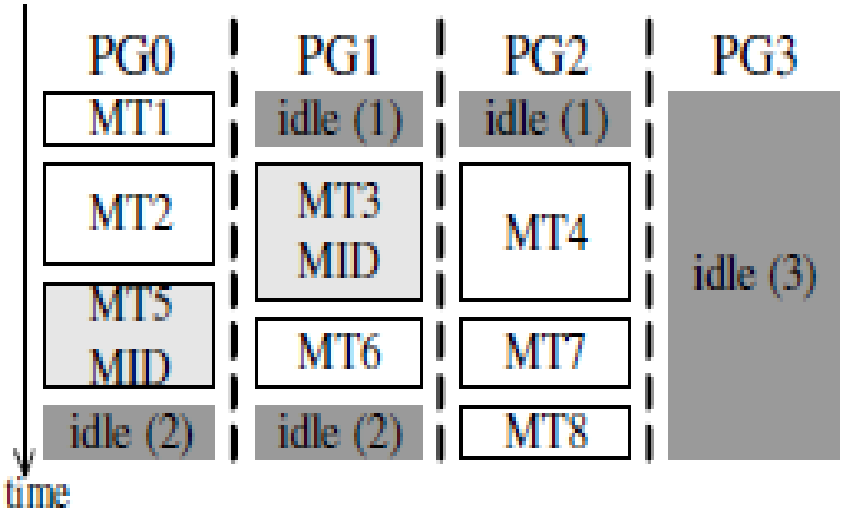

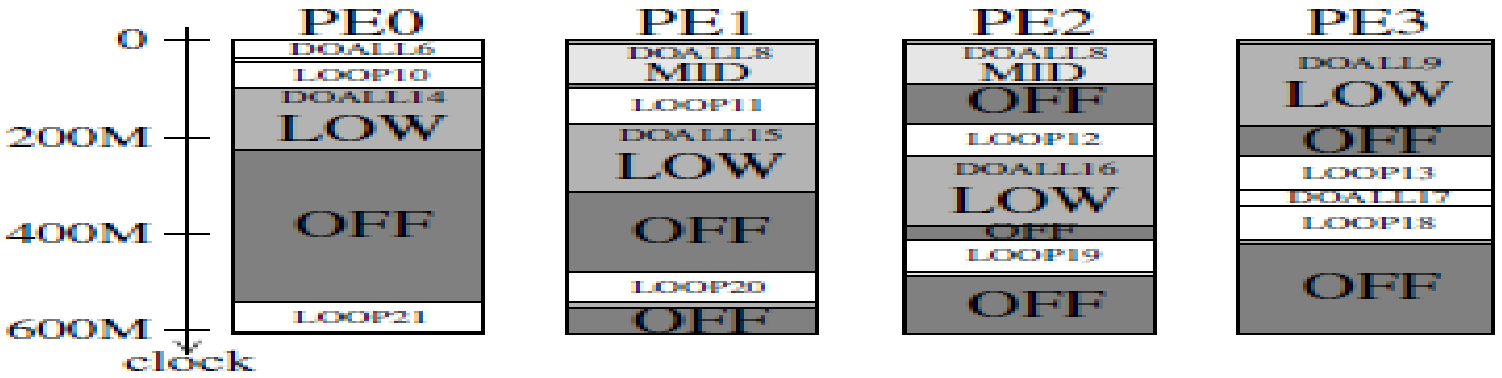

Fig. 6. V/F control of applu(4proc.)

Low-Power Optimization with OSCAR API

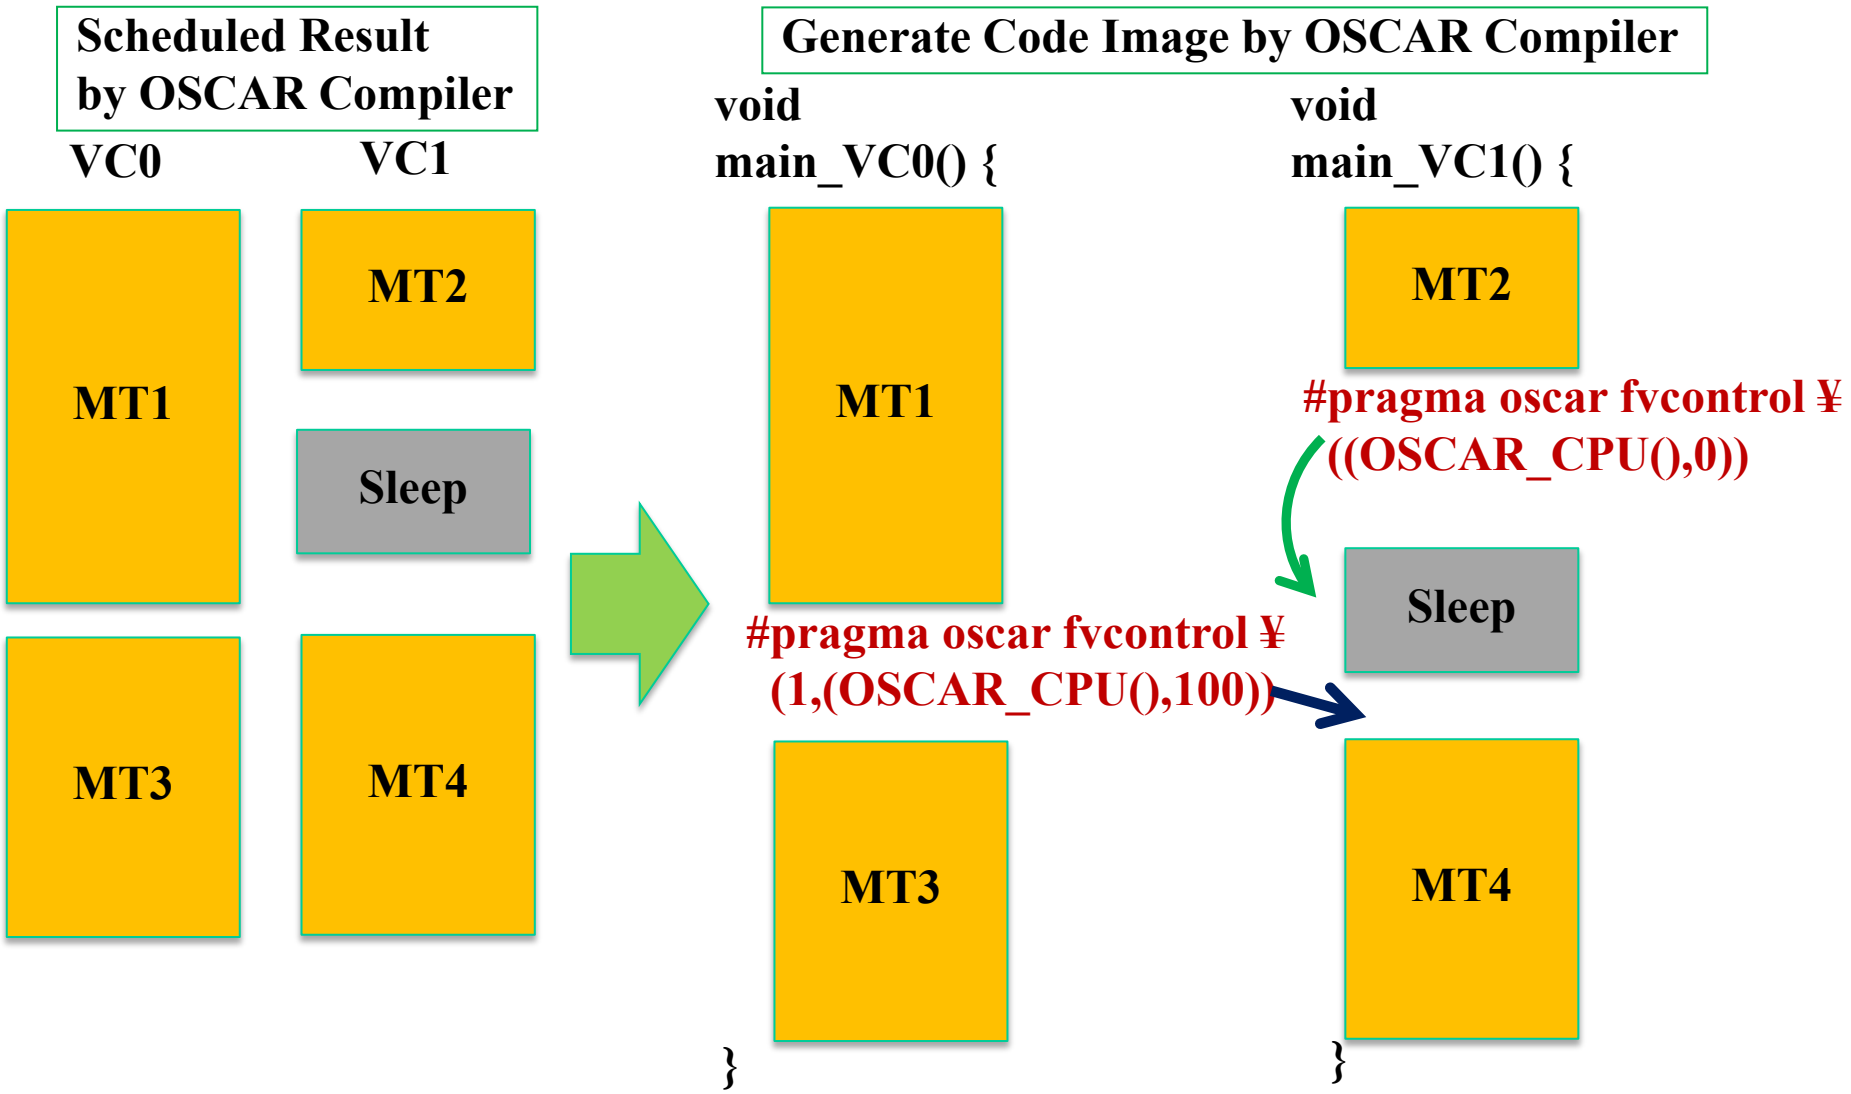

Multicore Program Development Using OSCAR API V2.0

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

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

Generation of parallel machine codes using sequential compilers

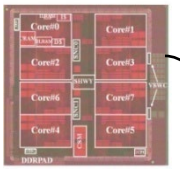

Homegeneous Multicore s from Vendor A (SMP servers)

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

Earliest Executable Condition Analysis for Coarse Grain Tasks (Macro-tasks)

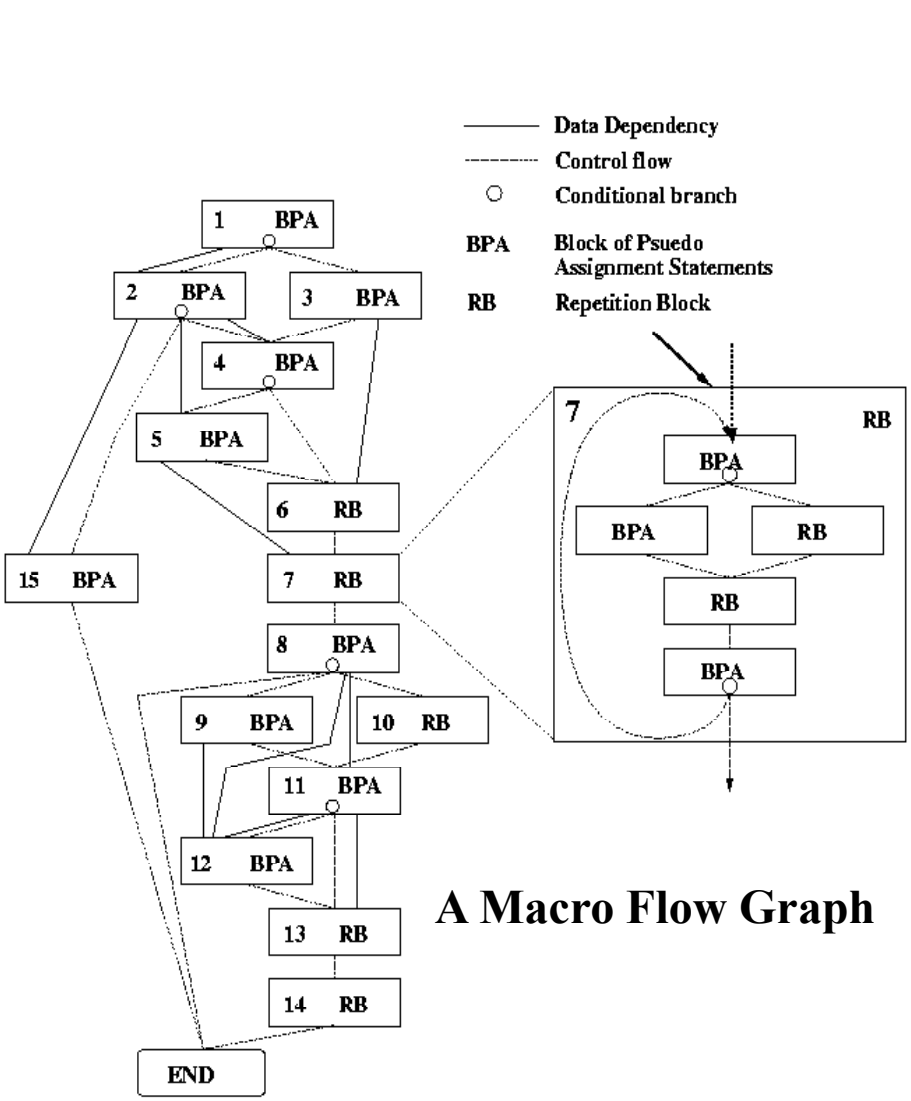

A Macro Flow Graph

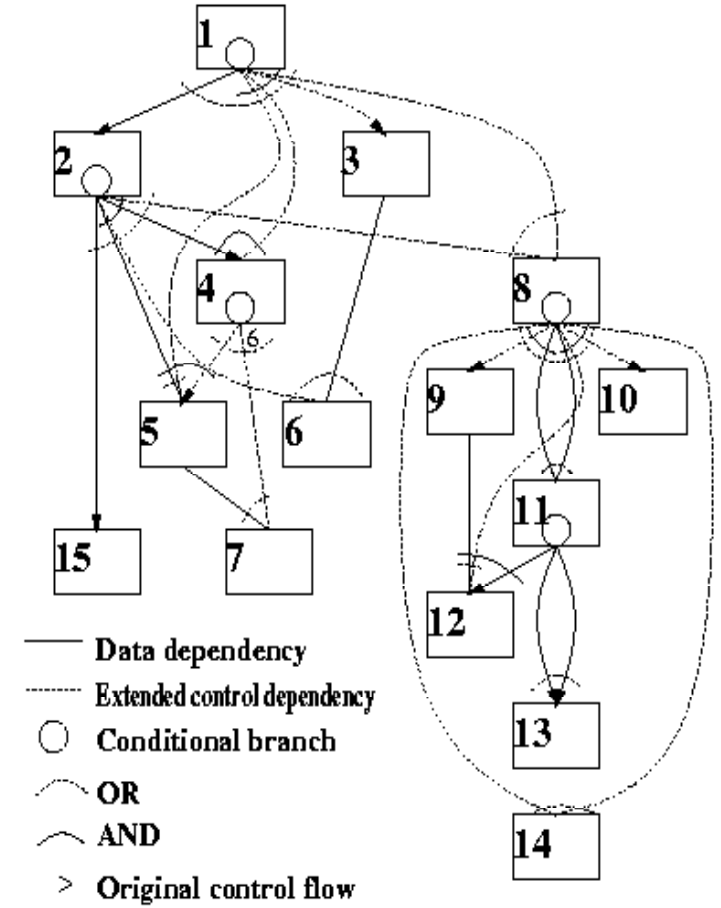

A Macro Task Graph

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

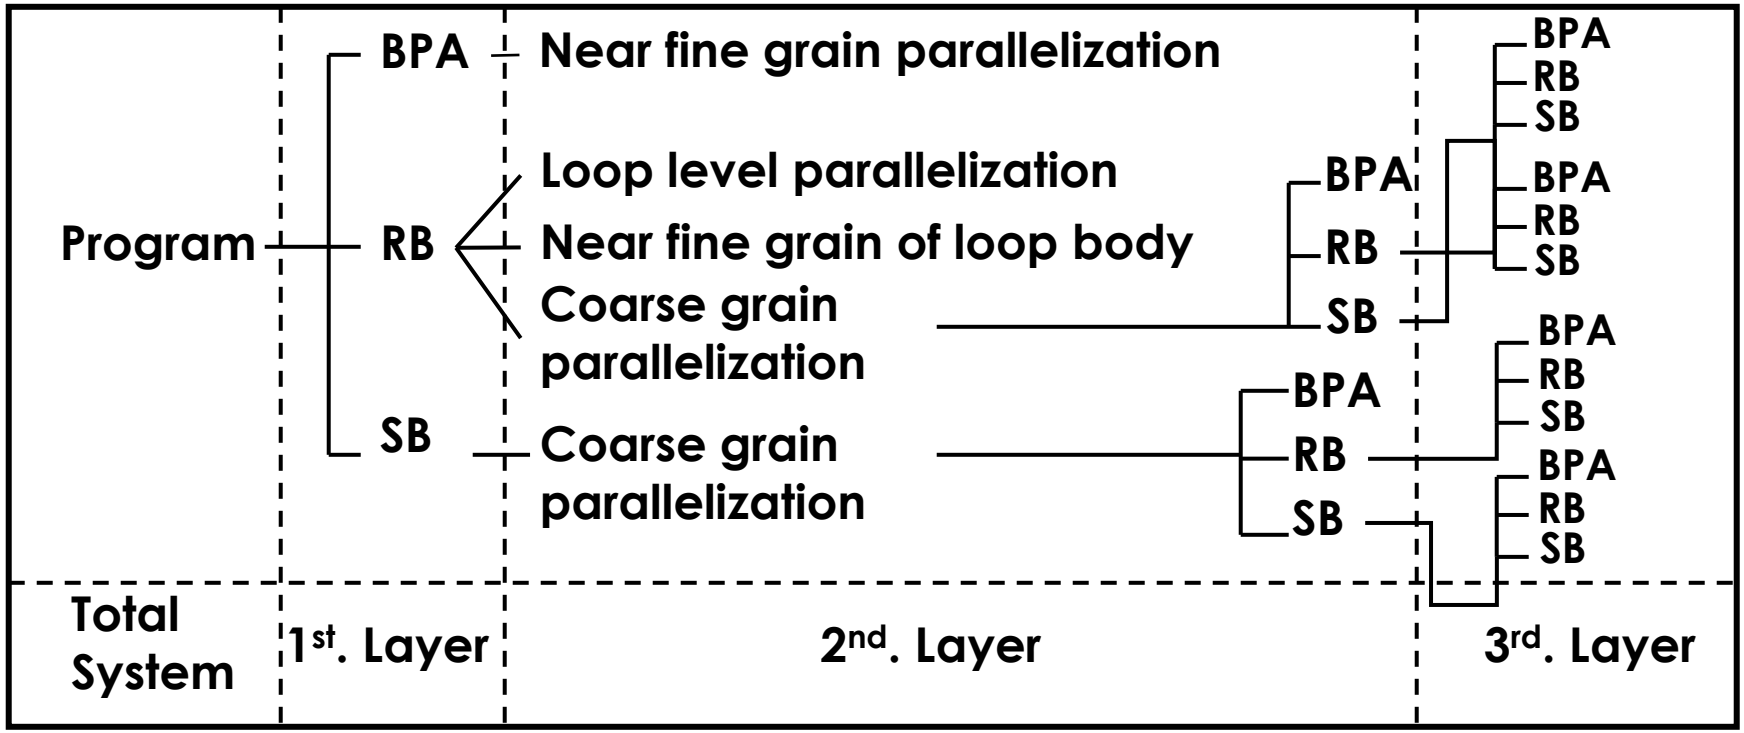

MTG of Su2cor-LOOPS-DO400

- **Coarse grain parallelism PARA_ALD = 4.3**

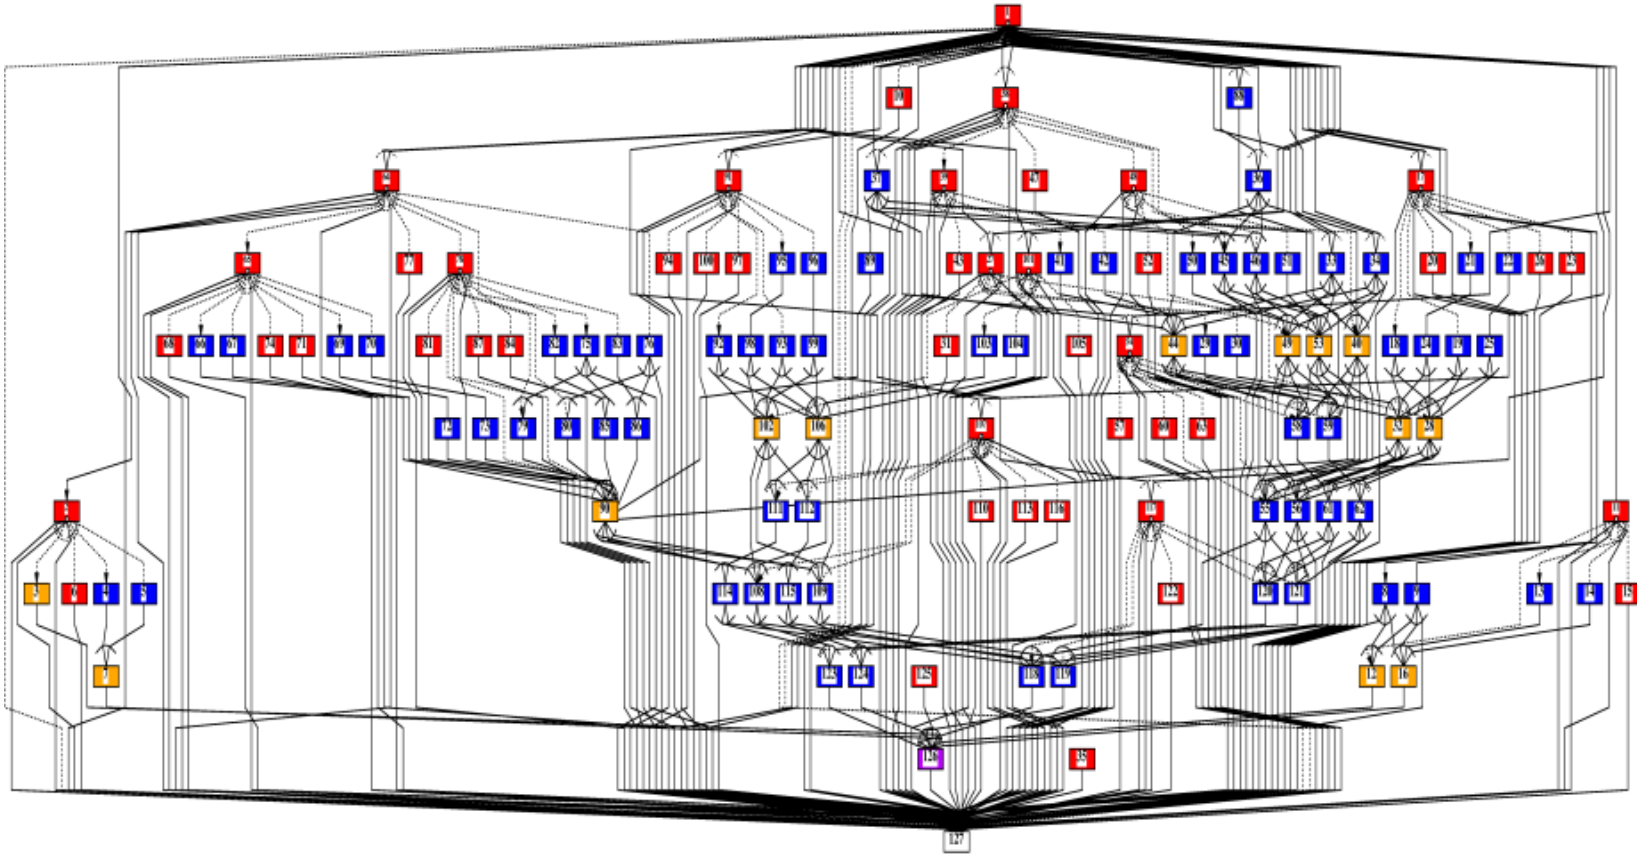

DOALL **Sequential LOOP** **SB** **BB**

Generated Multigrain Parallelized Code

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

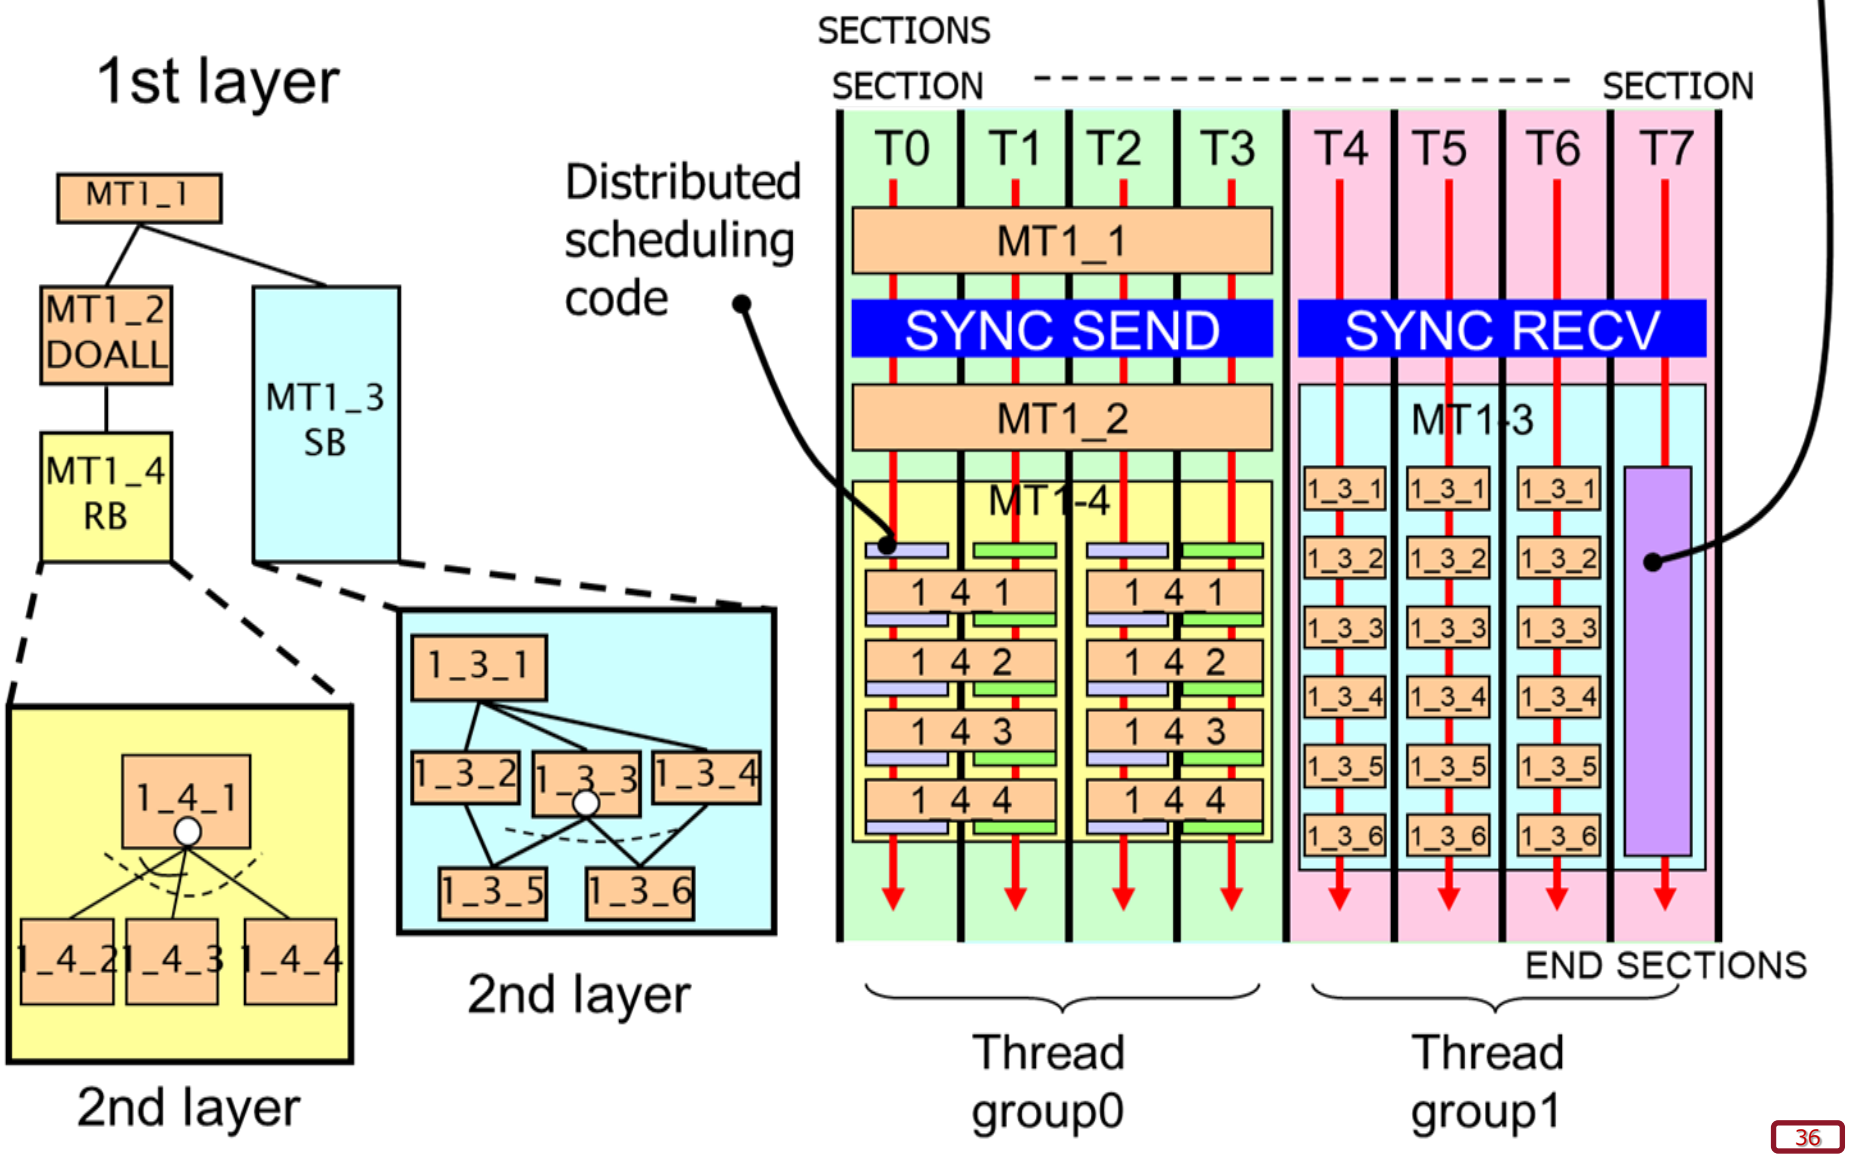

Data Localization

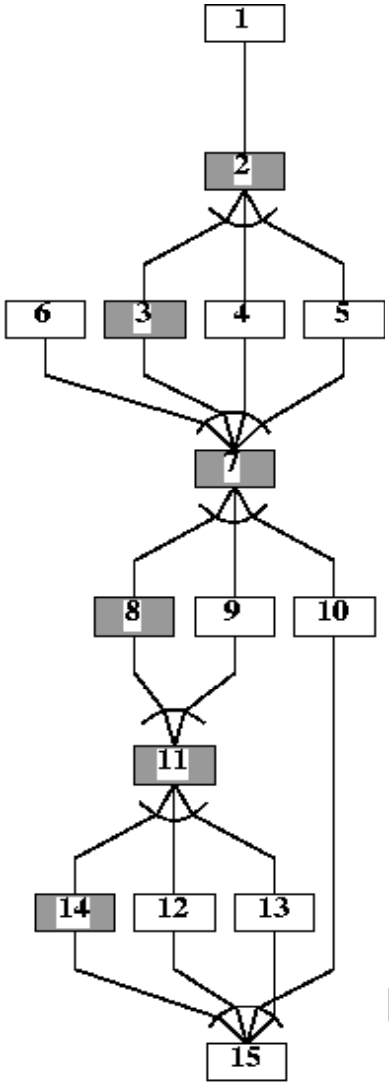

MTG

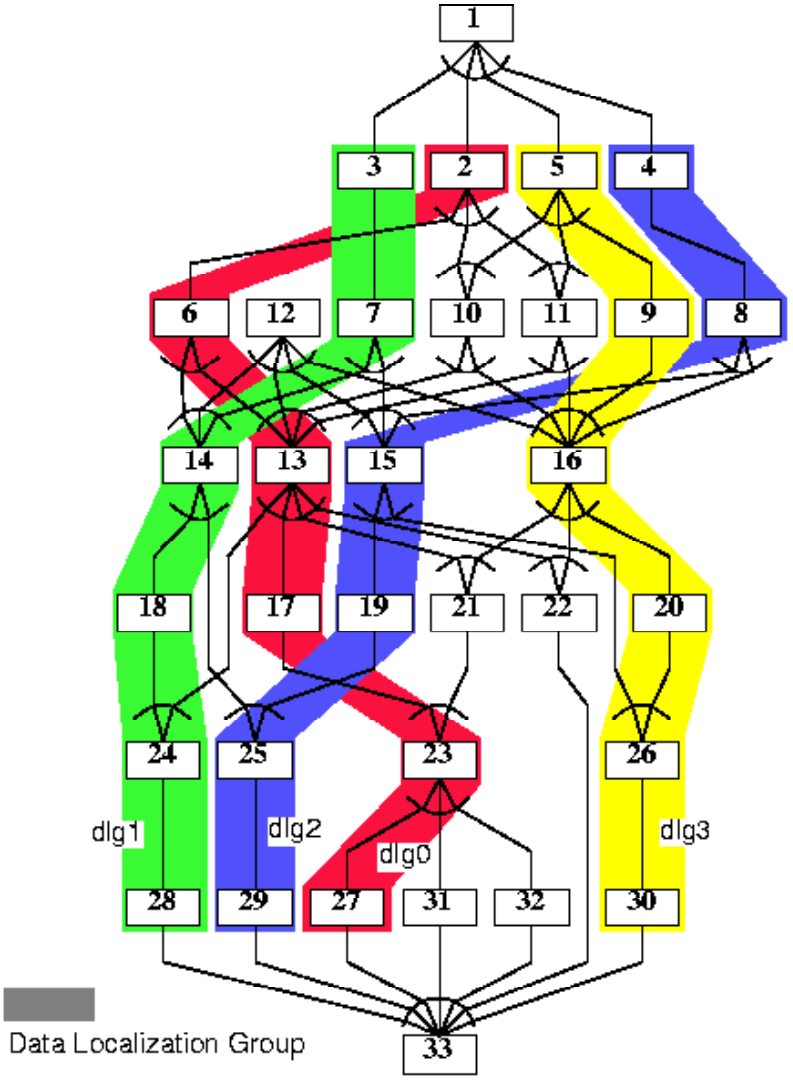

MTG after Division

A schedule for two processors

Intel Xeon E5-2650v4 – Benchmark results on upto 8 cores

- x86-64 based Architecture, 12 Cores, 2.2 GHz – 2.9 GHz
- 30 MiB shared L3 cache,
- L3 Cache: Shared by all cores
- speedup to sequential version
- gcc as backend

- **NAS parallel benchmark suite**
  - **BT: Block Tri-diagonal solver**
  - **CG: Conjugate Gradient computation**
  - **SP: Scalar Penta-diagonal solver**
- **SPEC2000**
  - **art: Image recognition / Neural networks**
  - **equake: Seismic wave propagation simulation**
  - **swim: Shallow water modelling - Fortran 77**
- **MediaBench II : MPEG2 encoding**

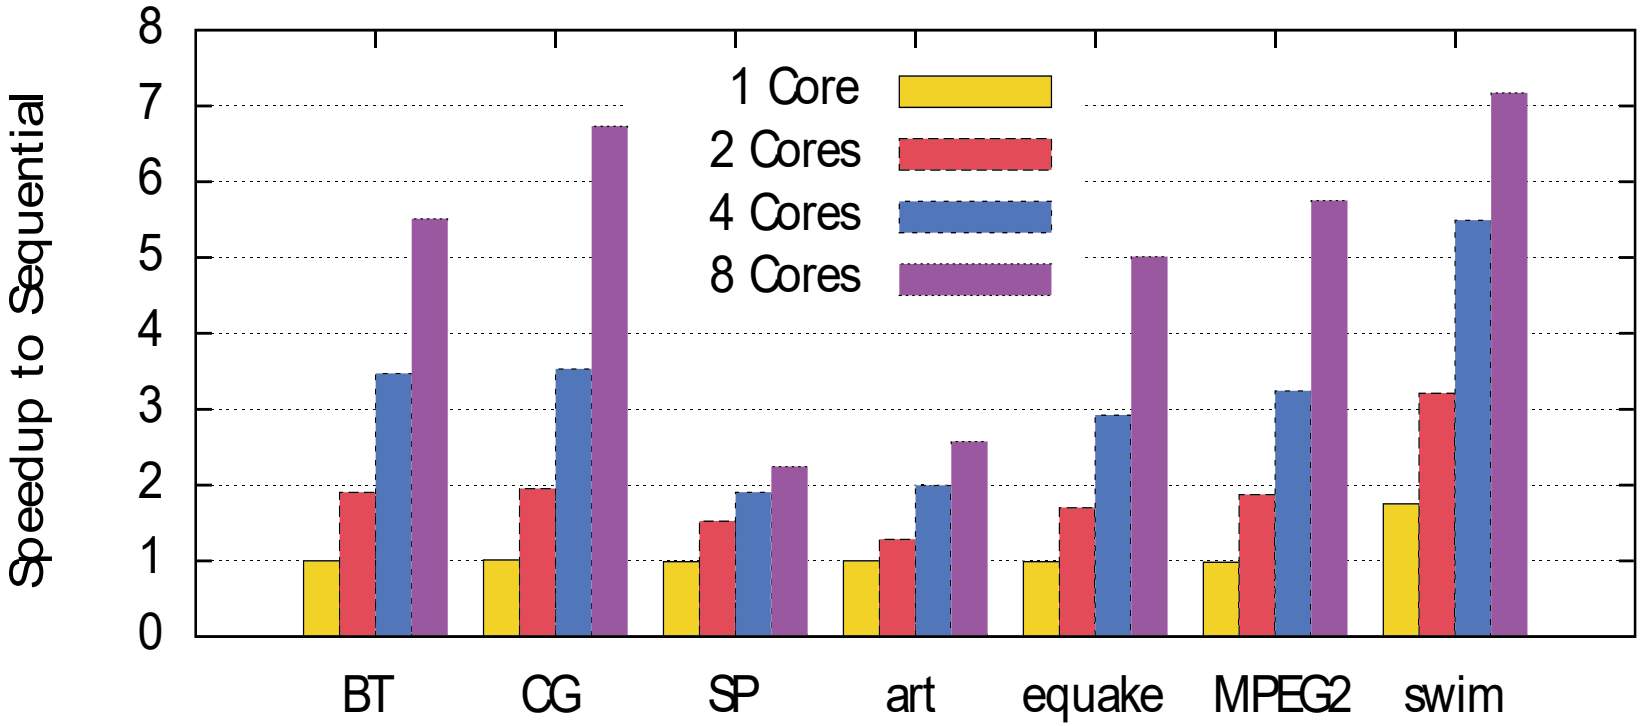

- swim shows superlinear speedup and 1 core speedup
  - seq.: 58.1 s, 1 core OSCAR: 33.2 s, 4 core OSCAR: 10.5 s

AMD EPYC 7702P – Benchmark results on upto 8 cores

- x86-64 based Architecture, 64 Cores, 2.0 GHz – 3.35 GHz
- 16 MiB L3 cache per 4 core cluster,
- shared within the cluster

- NAS parallel benchmark suite
  - BT: Block Tri-diagonal solver
  - CG: Conjugate Gradient computation
  - SP: Scalar Penta-diagonal solver
- SPEC2000
  - art: Image recognition / Neural networks
  - equake: Seismic wave propagation simulation
  - swim: Shallow water modelling - Fortran 77
- MediaBench II : MPEG2 encoding

- speedup to sequential version
- gcc as backend

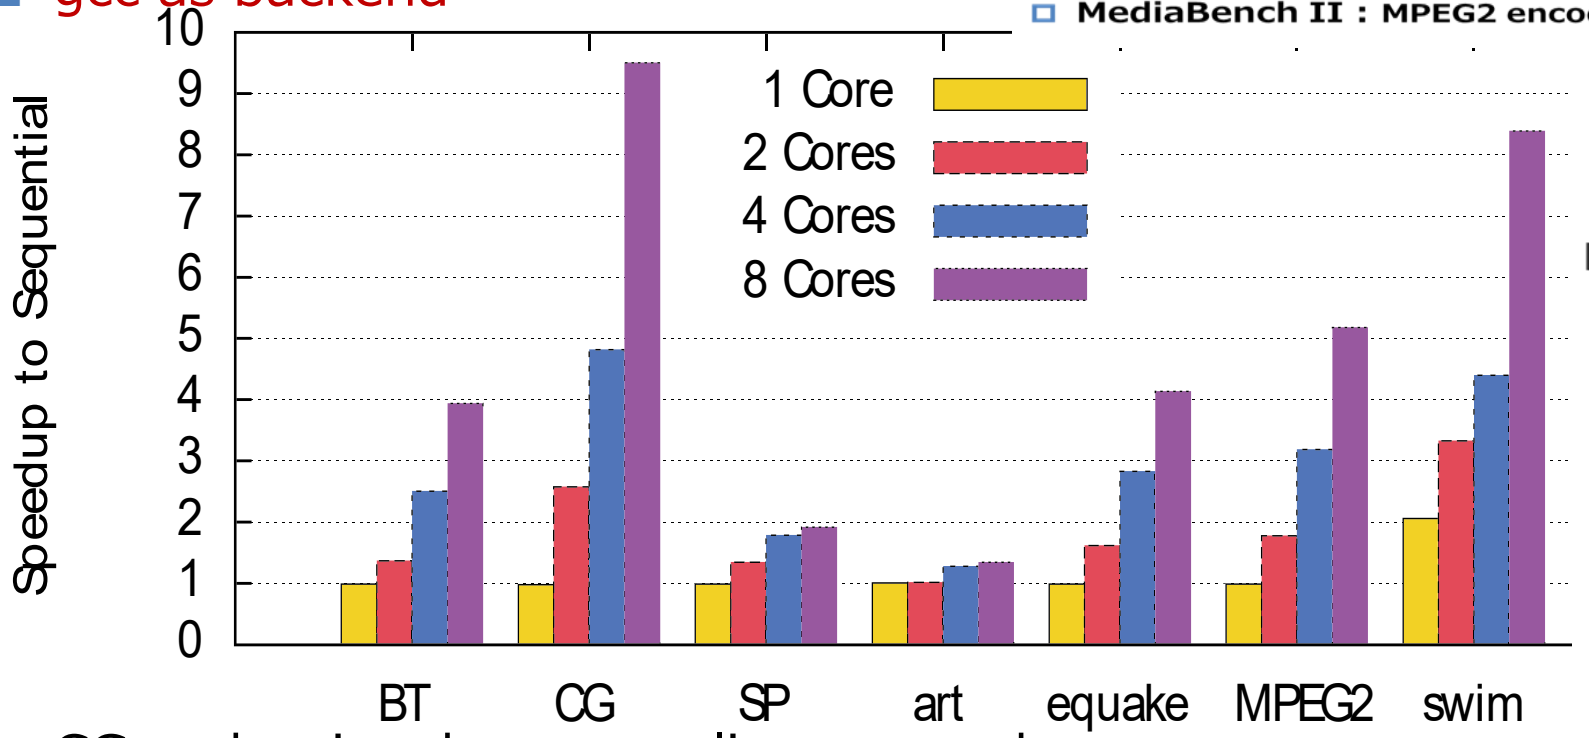

- CG and swim show superlinear speedup
  - CG: seq.: 0.86 s, 8 core OSCAR: 0.09 s

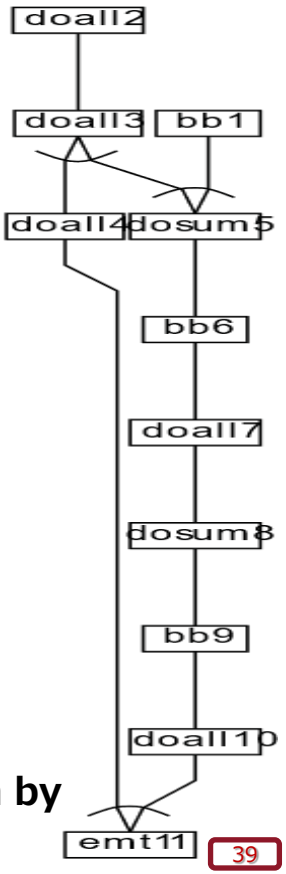

CG macrotask graph
L3 cache Optimization by Data Localization

NVIDIA Carmel ARMv8.2 – Benchmark results on upto 4 cores

- Arm v8.2 based Architecture, 6 Cores, 1.4 GHz
- 4 MiB shared L3 cache,
- L3 Cache: Shared across all cores
- speedup to sequential versior
- gcc as backend

- NAS parallel benchmark suite
  - BT: Block Tri-diagonal solver
  - CG: Conjugate Gradient computation
  - SP: Scalar Penta-diagonal solver
- SPEC2000
  - art: Image recognition / Neural networks
  - equake: Seismic wave propagation simulation
  - swim: Shallow water modelling - Fortran 77
- MediaBench II : MPEG2 encoding

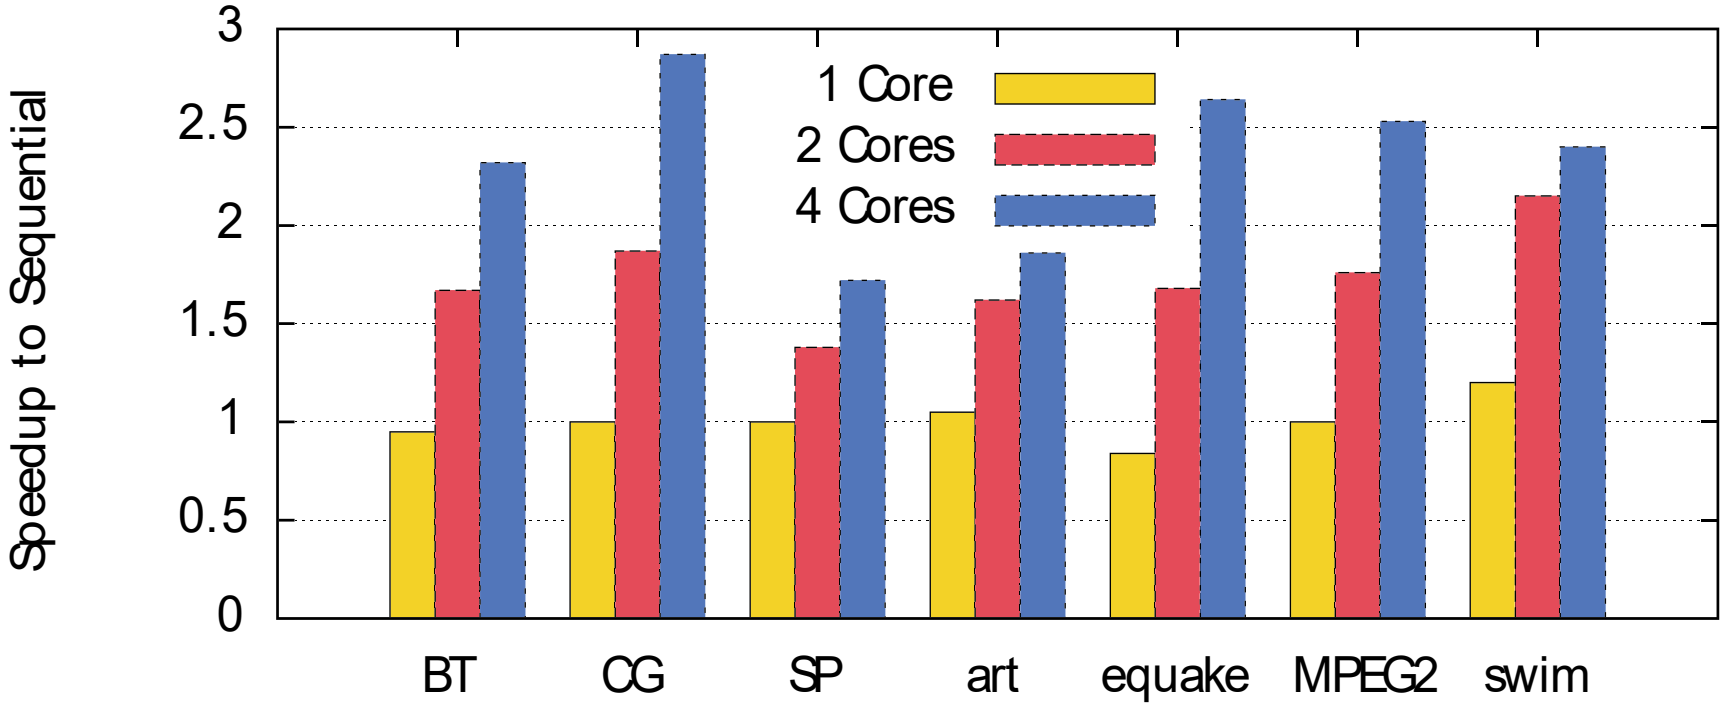

- overall good speedup is observed
  - equake: seq.: 19.0 s, 4 core OSCAR: 7.18 s

SiFive Freedom U740 – Benchmark results on upto 4 cores

- RISC-V based Architecture, 4 Cores, 1.2 GHz
- 2 MiB shared L2 cache
- L2 Cache: Shared across all cores

- NAS parallel benchmark suite
  - BT: Block Tri-diagonal solver
  - CG: Conjugate Gradient computation
  - SP: Scalar Penta-diagonal solver
- SPEC2000
  - art: Image recognition / Neural networks
  - equake: Seismic wave propagation simulation
  - swim: Shallow water modelling - Fortran 77
- MediaBench II : MPEG2 encoding

- speedup to sequential version
- gcc as backend

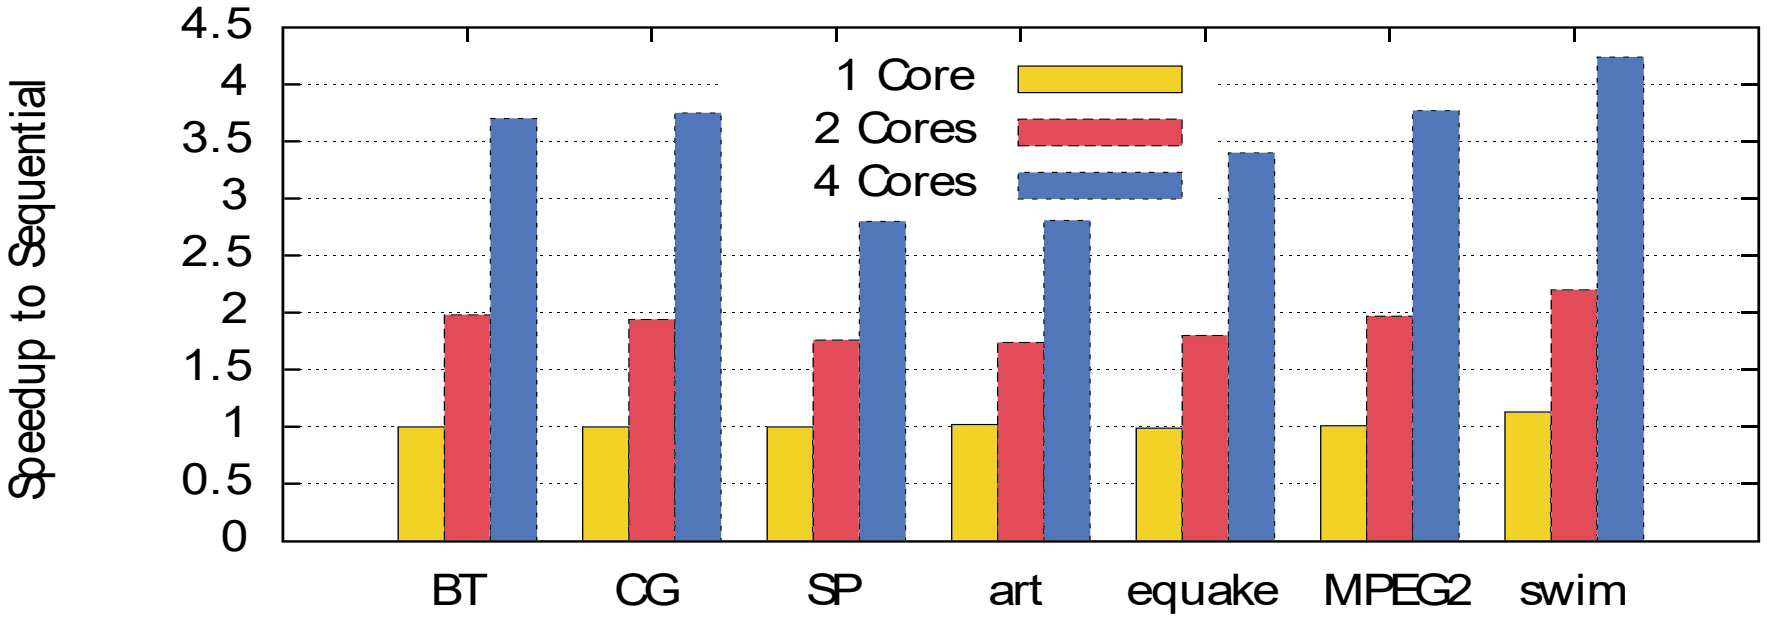

- overall good speedup is observed, swim superlinear
  - BT: seq.: 2041 s, 4 core OSCAR: 551 s

日本乗用車のエンジン制御計算をデンソー2コアECU上で、1．95倍の速度向上に成功。(見神、梅田）

欧州農耕作業車エンジン制御計算をインフィニオン2コアプロセッサ上で8.7倍の高速化に成功。

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.

- Hard real-time automobile engine control by multicore using local memories
- Millions of lines C codes consisting conditional branches and basic blocks

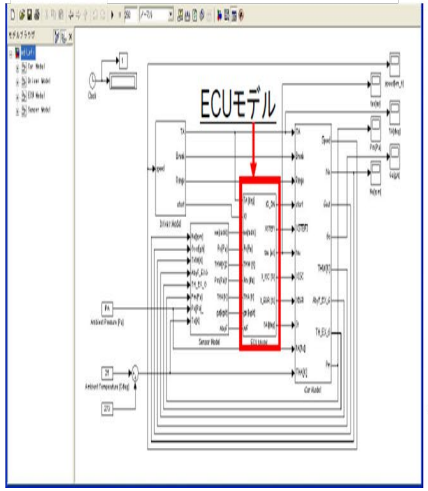

Automatic Parallelization of an Engine Control C Program with 400 thousands lines on AUTOSAR on 2 cores of Infineon AURIX TC277

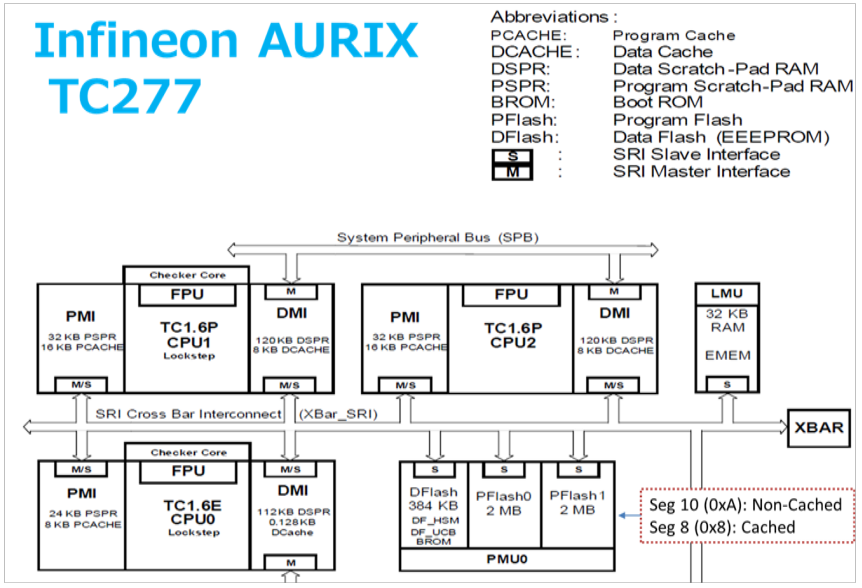

Macro Task Fusion for Static Task Scheduling

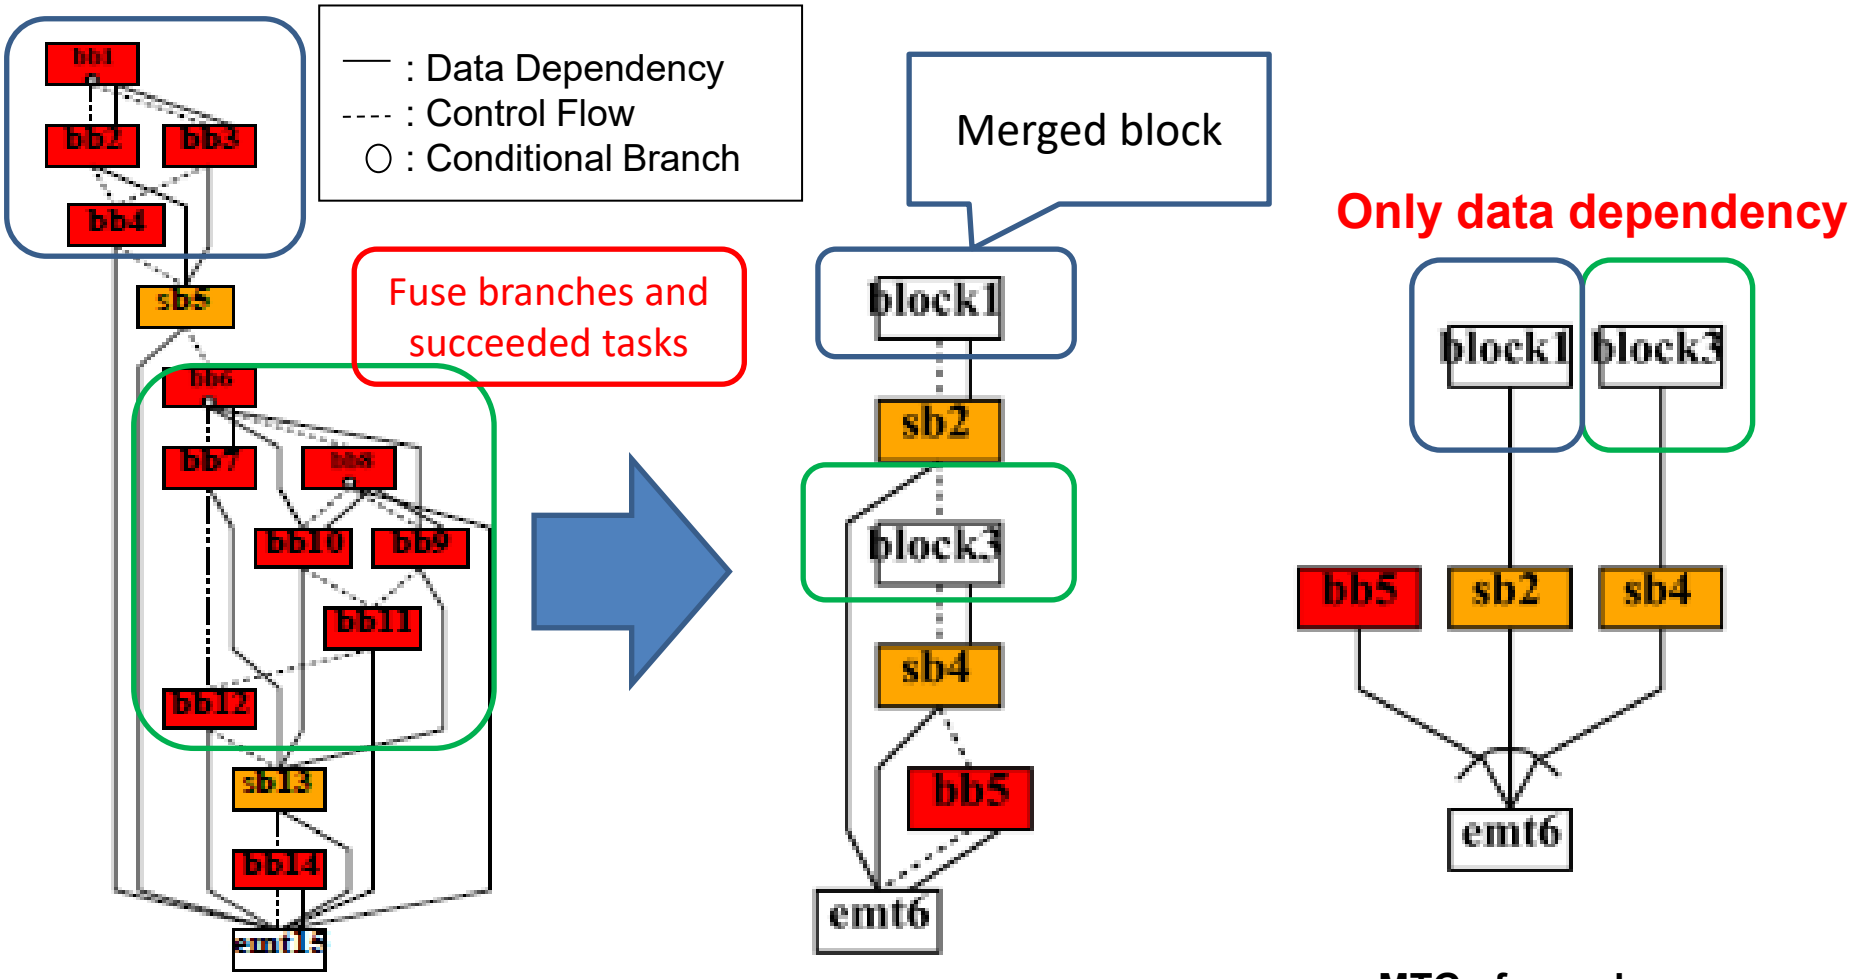

MFG of sample program before maro task fusion

MFG of sample program after macro task fusion

MTG of sample program after macro task fusion

OSCAR Compile Flow for MATLAB/Simulink

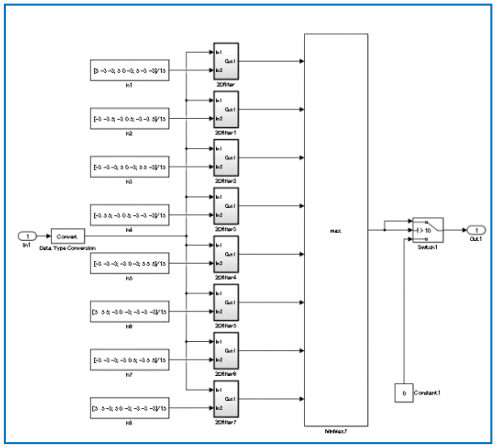

Simulink model

Generate C code
using Embedded Coder

```
/* Model step function */
void VesselExtraction_step(void)
{
  int32_T i;
  real_T u0;

  /* DataTypeConversion: '<S1>/Data Type Conversion' incorporates:
   *  Inport: '<Root>/In1'
   */
  for (i = 0; i < 16384; i++) {
    VesselExtraction_B.DataTypeConversion[i] = VesselExtraction_U.In1[i];
  }

  /* End of DataTypeConversion: '<S1>/Data Type Conversion' */

  /* Outputs for Atomic SubSystem: '<S1>/2Dfilter' */

  /* Constant: '<S1>/h1' */
  VesselExtraction_Dfilter(VesselExtraction_B.DataTypeConversion,
    VesselExtraction_P.h1_Value, &VesselExtraction_B.Dfilter,
    (P_Dfilter_VesselExtraction_T *)&VesselExtraction_P.Dfilter);

  /* End of Outputs for SubSystem: '<S1>/2Dfilter' */

  /* Outputs for Atomic SubSystem: '<S1>/2Dfilter1' */

  /* Constant: '<S1>/h2' */
  VesselExtraction_Dfilter(VesselExtraction_B.DataTypeConversion,
    VesselExtraction_P.h2_Value, &VesselExtraction_B.Dfilter1,
    (P_Dfilter_VesselExtraction_T *)&VesselExtraction_P.Dfilter1);
```

C code

OSCAR Compiler

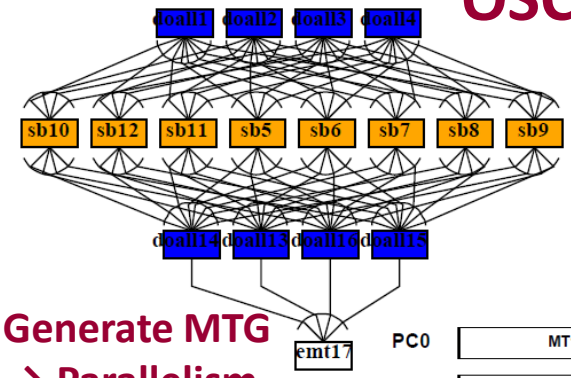

(1) Generate MTG
→ Parallelism

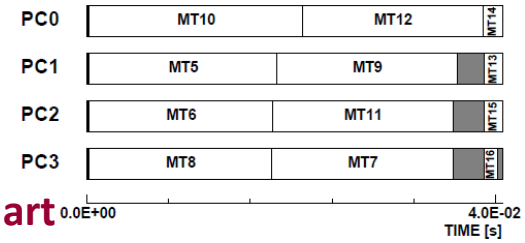

(2) Generate gantt chart
→ Scheduling in a multicore

```
void VesselExtraction_step ( )
{
    int thr1 ;
    int thr2 ;
    int thr3 ;
    {
        oscar_thread_create ( & thr1 ,
            thread_function_001 , (void*)1 ) ;
        oscar_thread_create ( & thr2 ,
            thread_function_002 , (void*)2 ) ;
        oscar_thread_create ( & thr3 ,
            thread_function_003 , (void*)3 ) ;

            VesselExtraction_step_PE0 ( ) ;

        oscar_thread_join ( thr1 ) ;
        oscar_thread_join ( thr2 ) ;
        oscar_thread_join ( thr3 ) ;
    }
}

void thread_function_001 ( void )
{
    VesselExtraction_step_PE1 ( ) ;
}
```

(3) Generate parallelized C code
using the OSCAR API
→ Multiplatform execution
(Intel, ARM and SH etc)

Speedups of MATLAB/Simulink Image Processing on Various 4core Multicores

(Intel Xeon, ARM Cortex A15 and Renesas SH4A)

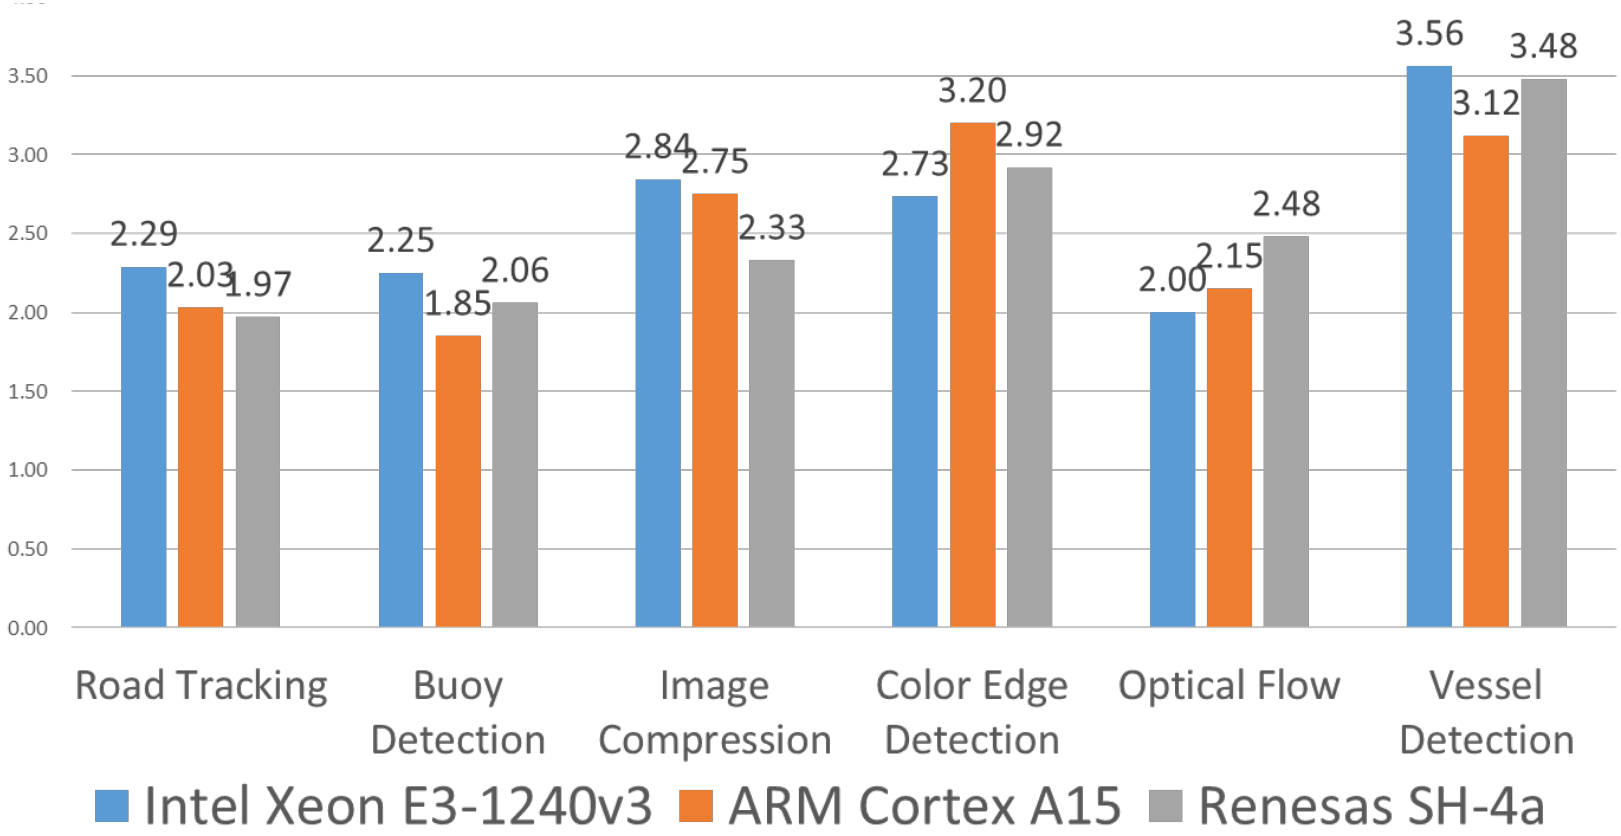

Road Tracking, Image Compression : http://www.mathworks.co.jp/jp/help/vision/examples
Buoy Detection : http://www.mathworks.co.jp/matlabcentral/fileexchange/44706-buoy-detection-using-simulink
Color Edge Detection : http://www.mathworks.co.jp/matlabcentral/fileexchange/28114-fast-edges-of-a-color-image--actual-color--not-converting-to-grayscale-/
Vessel Detection : http://www.mathworks.co.jp/matlabcentral/fileexchange/24990-retinal-blood-vessel-extraction/

Speedups of Deep Learning Winograd 2D-Convolution generated by TVM on NEC Personal Vector Supercomputer SX-Aurora TSUBASA 8 Core Type 10C

OSCAR Parallelization and NEC Vectorization gave us 363x Speedup against a Scalar Core

Parallelization of Deep Leaning C Code generated by TVM

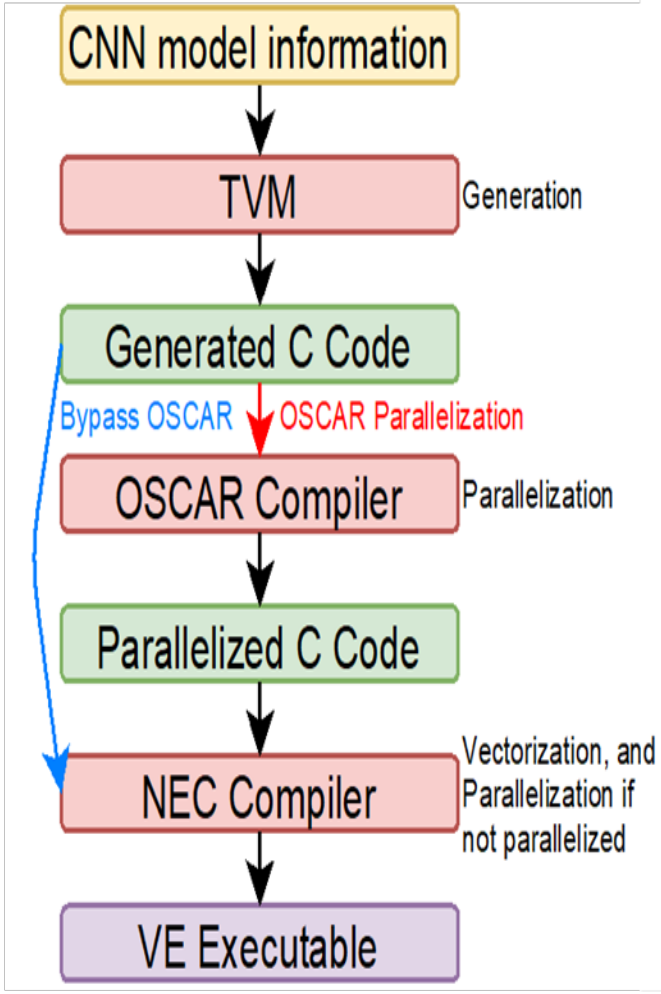

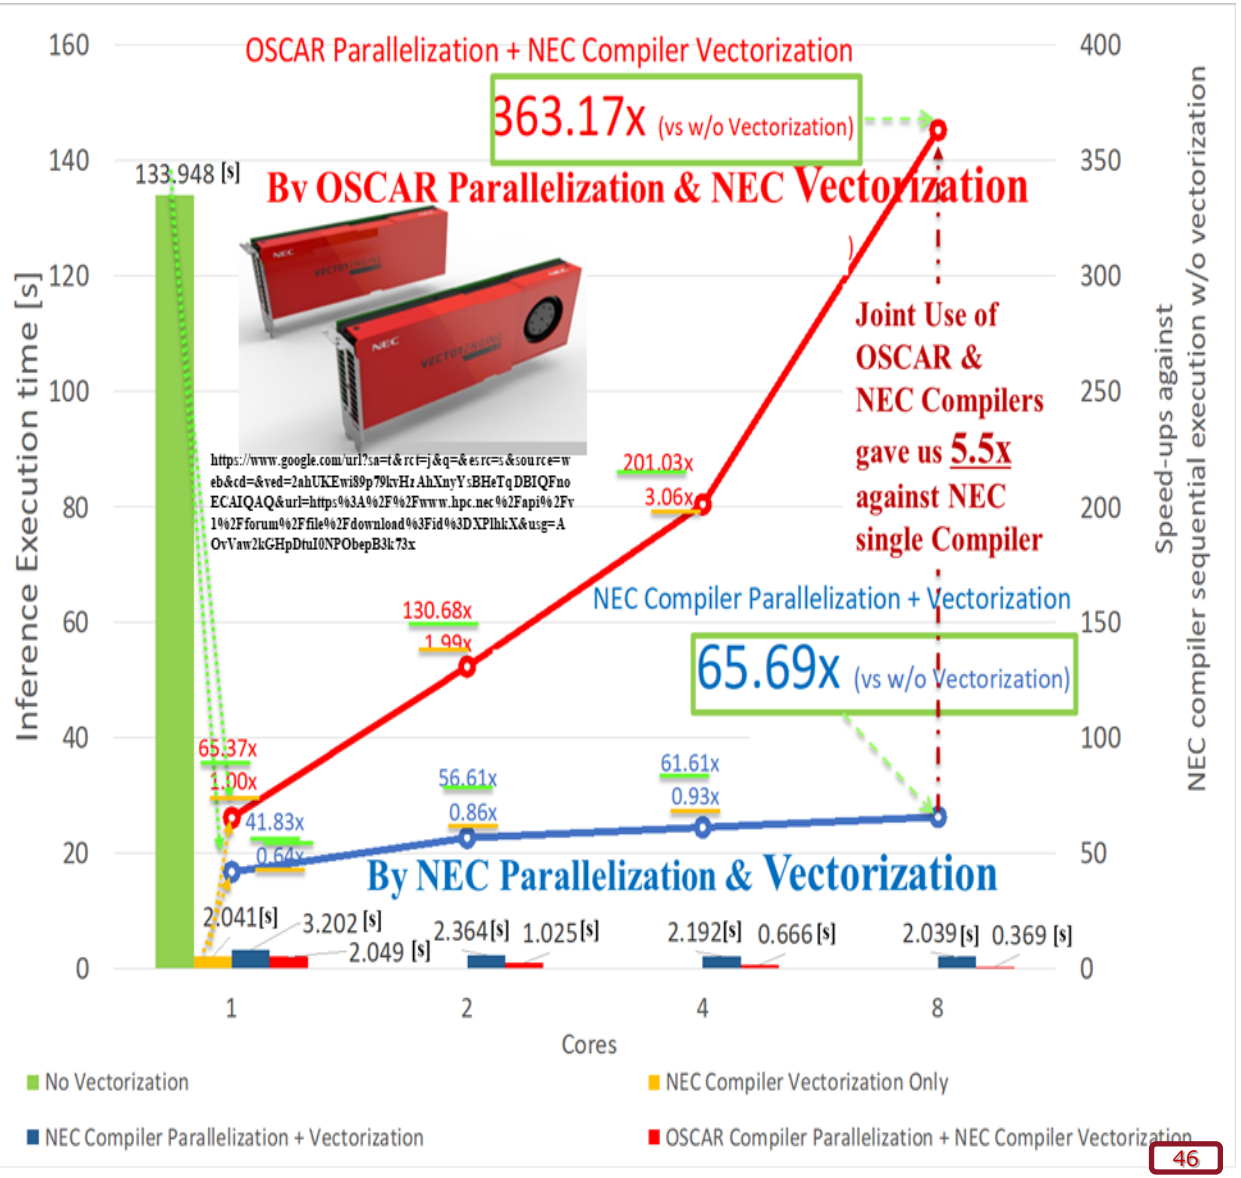

OSCAR Heterogeneous Multicore

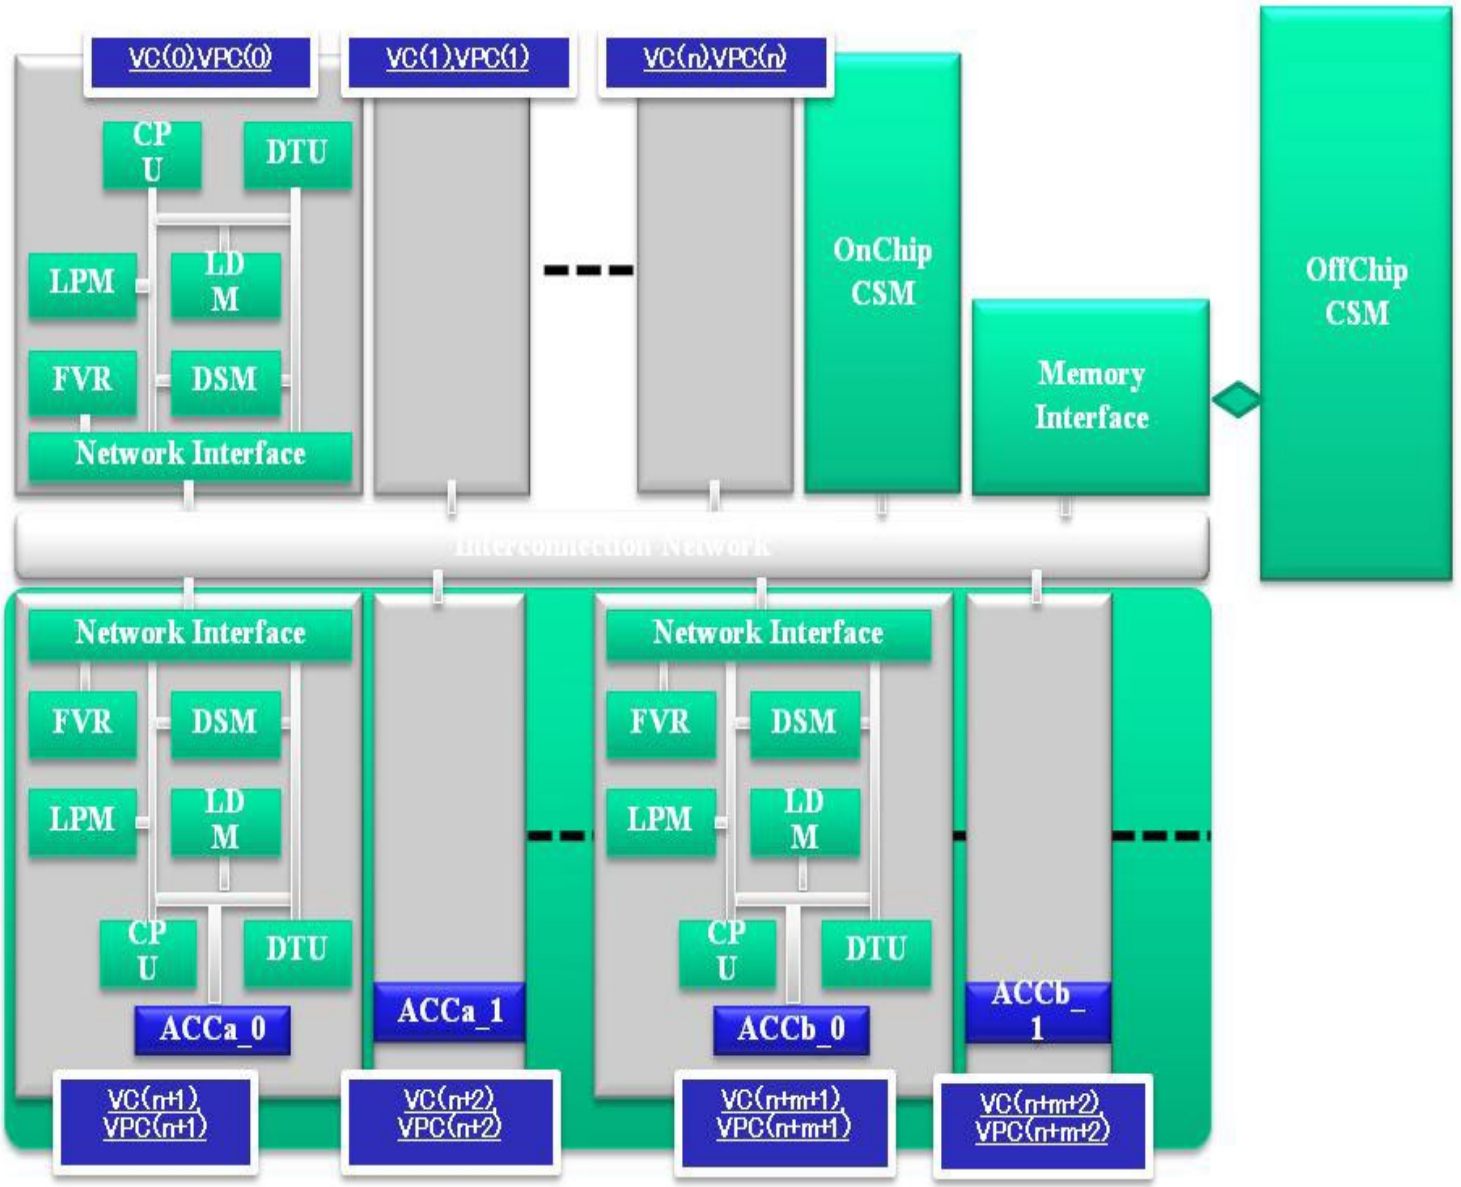

DTU
- Data Transfer Unit

LPM
- Local Program Memory

LDM
- Local Data Memory

DSM
- Distributed Shared Memory

CSM
- Centralized Shared Memory

FVR
- Frequency/Voltage Control Register

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

(LCPC2009 Homogeneous, 2010 Heterogeneous)

Specification: http://www.kasahara.cs.waseda.ac.jp/api/regist.php?lang=en&ver=2.1

List of Directives (22 directives)

- Parallel Execution API
  - parallel sections (*)
  - flush (*)
  - critical (*)
  - execution
- Memoay 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

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

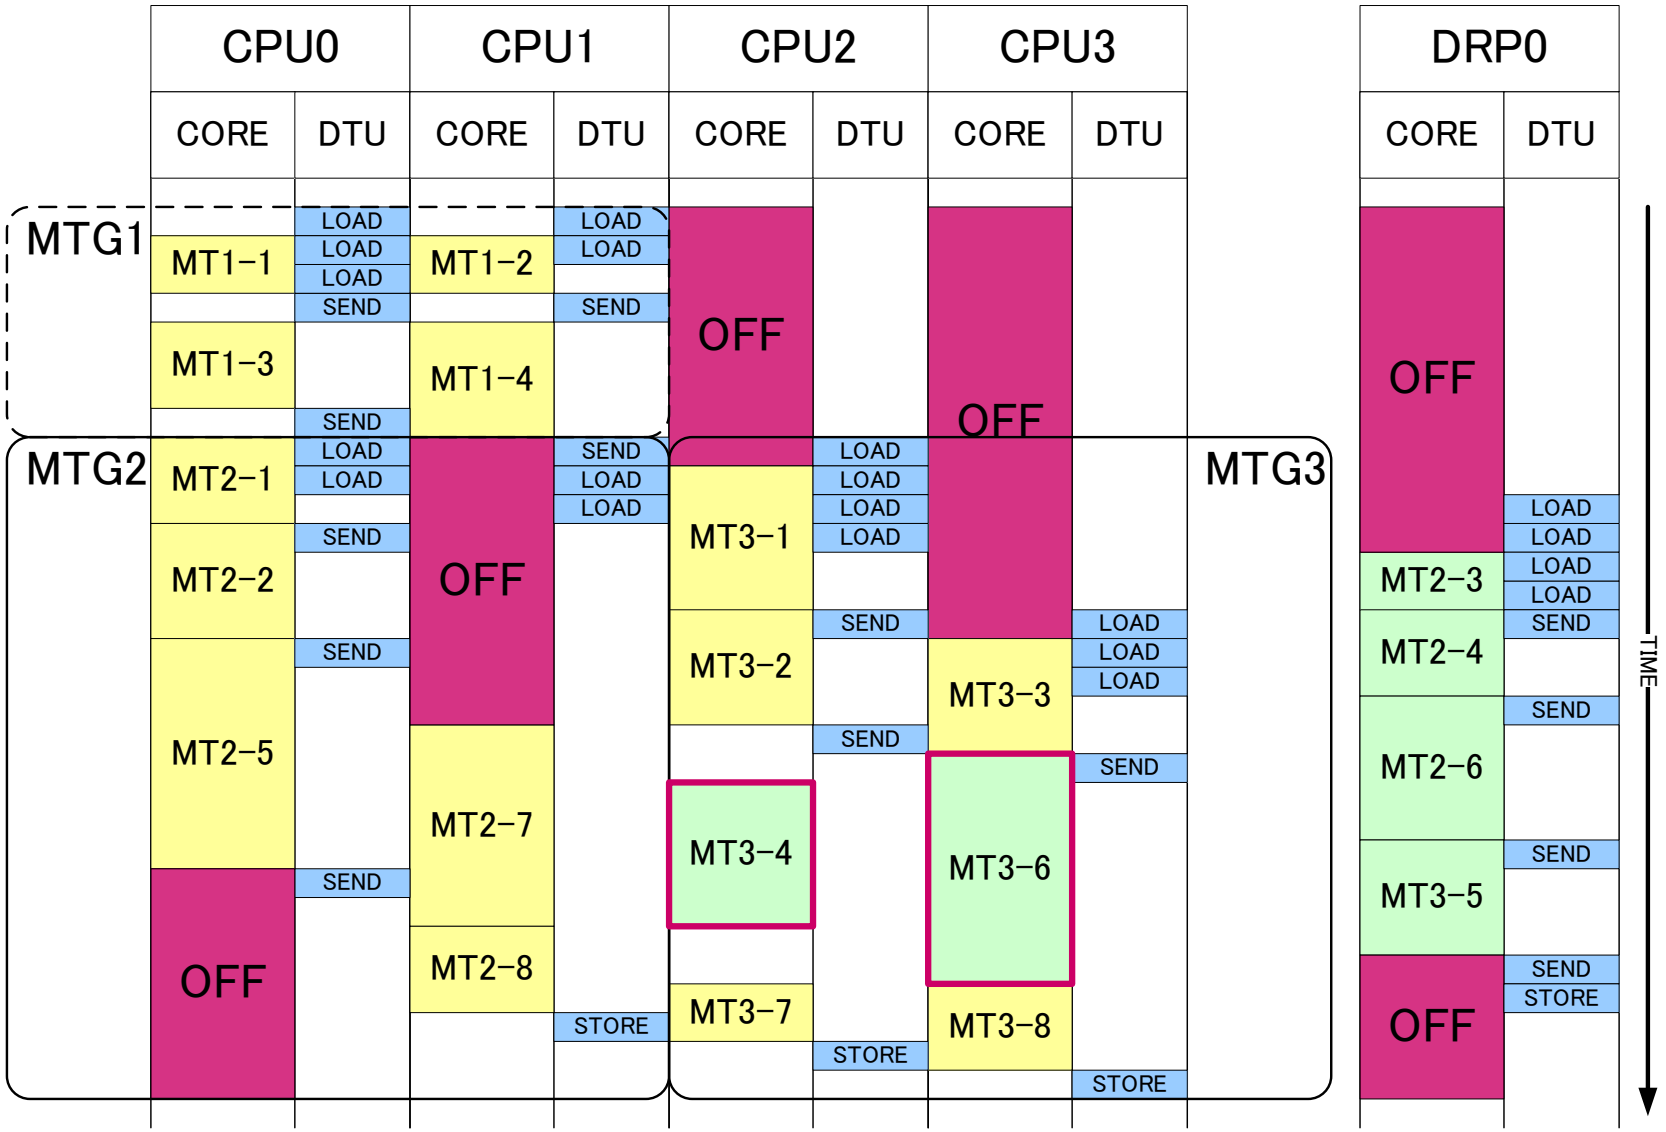

Automatic Local Memory Management

Data Localization: Loop Aligned Decomposition

- **Decomposed loop into LRs and CARs**
  - **LR (Localizable Region): Data can be passed through LDM**
  - **CAR (Commonly Accessed Region): Data transfers are required among processors**

Single dimension Decomposition

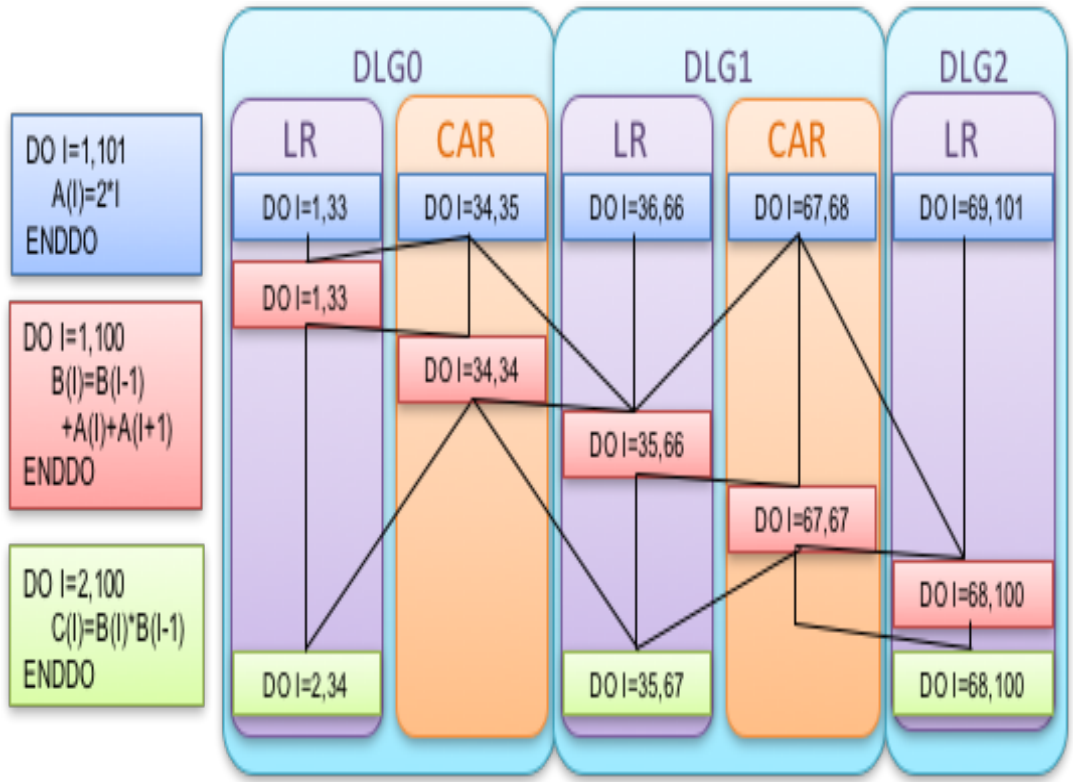

Multi-dimension Decomposition

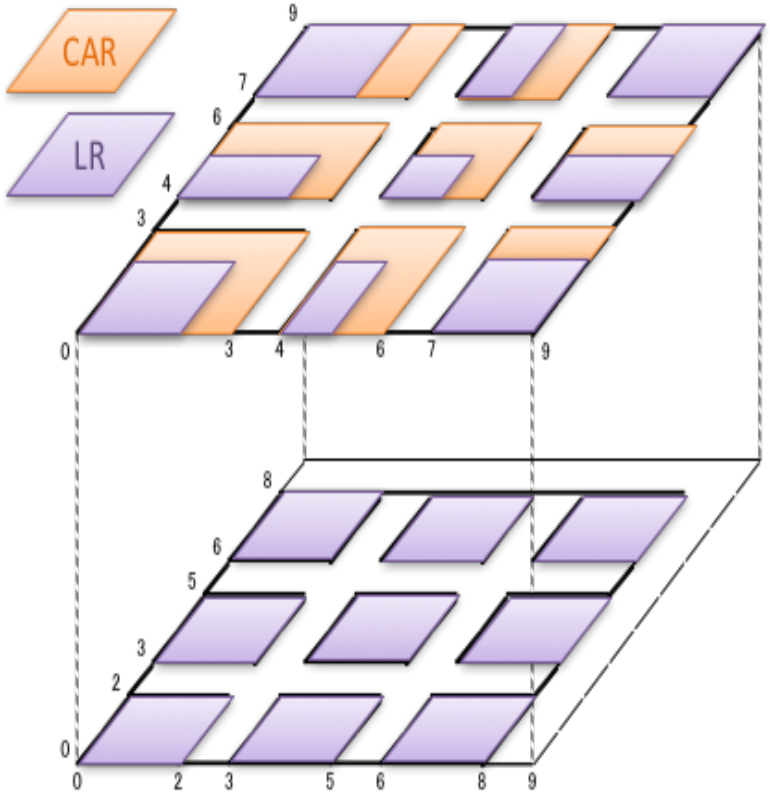

Speedups by OSCAR Automatic Local Memory Management compared to Executions Utilizing Centralized Shared Memory on Embedded and Scientific Application on RP2 8core Multicore

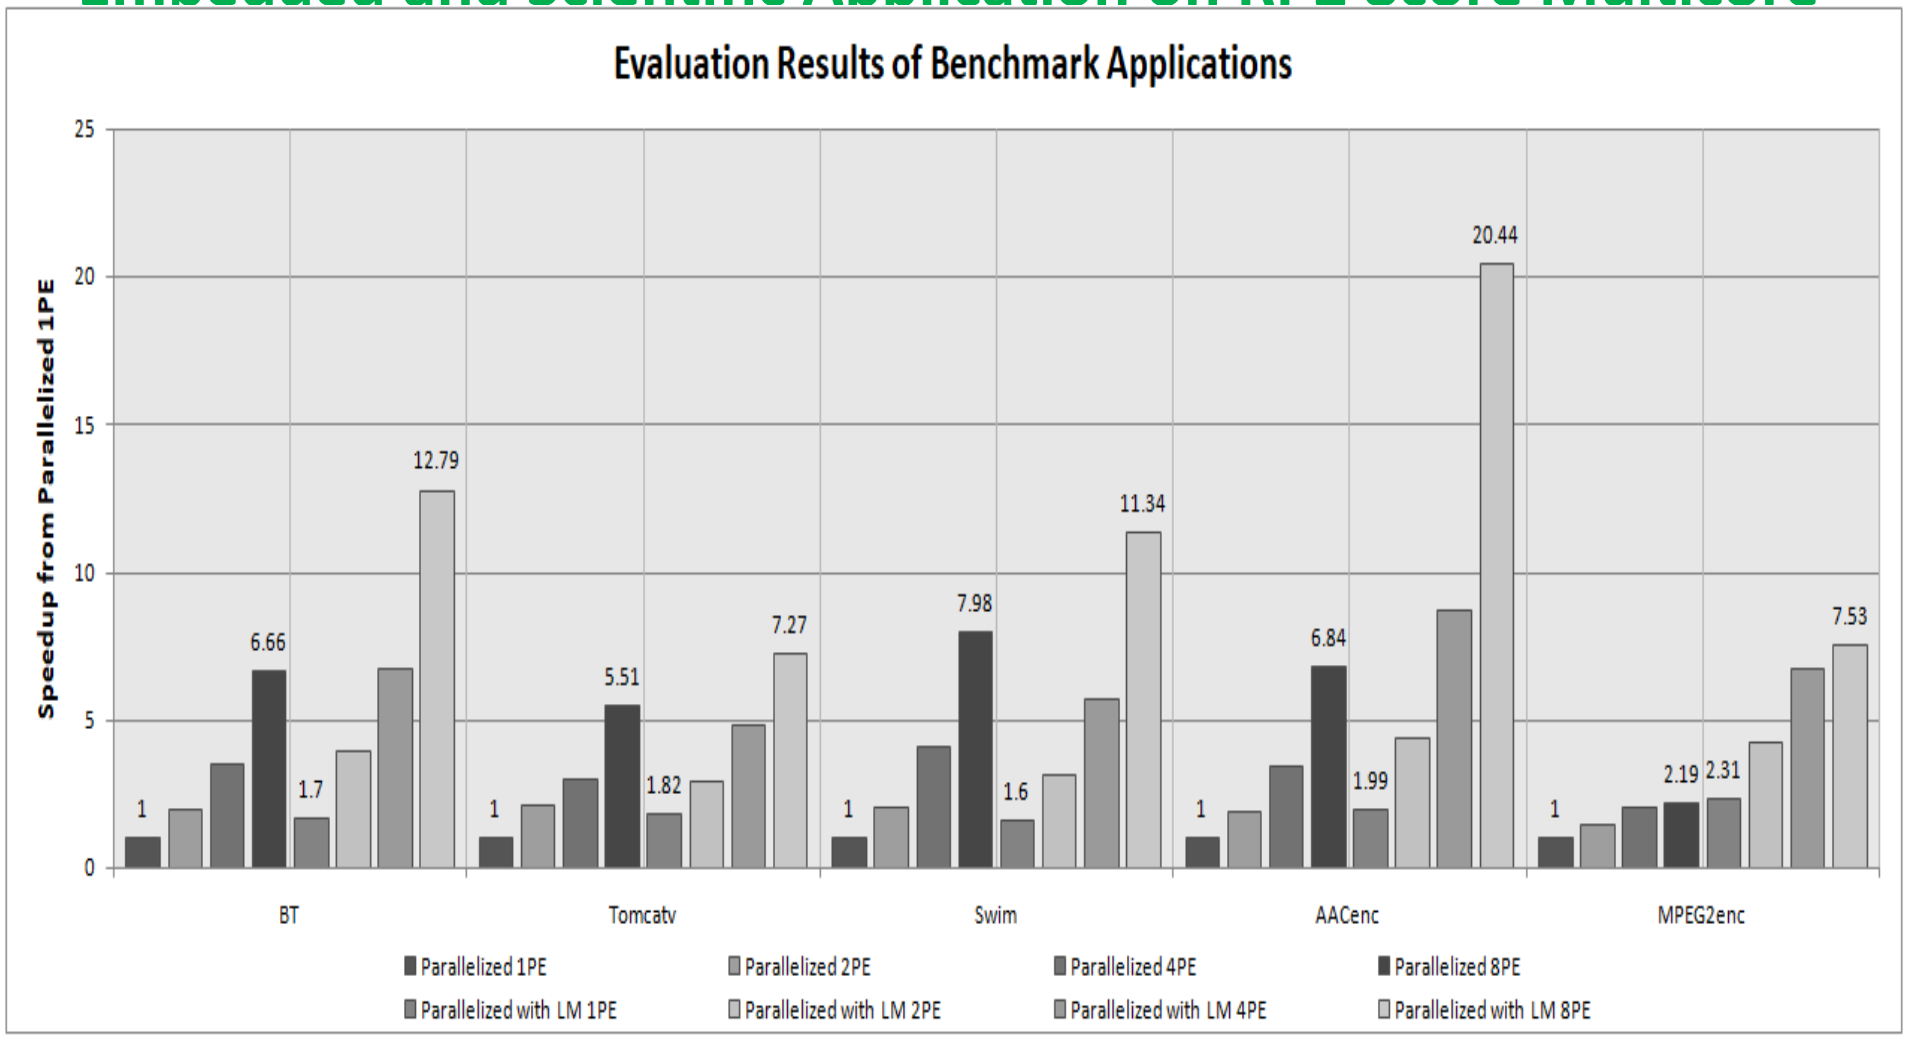

Maximum of 20.44 times speedup on 8 cores using local memory against sequential execution using off-chip shared memory

AIREC (AI-driven Robot for Embrace and Care) Led by Prof. Sugano

Supported by Japanese Government "Moonshot" Project from 2020

人生に寄り添うAIロボット：育児,家事,家業補助,看護,遠隔診断,介護

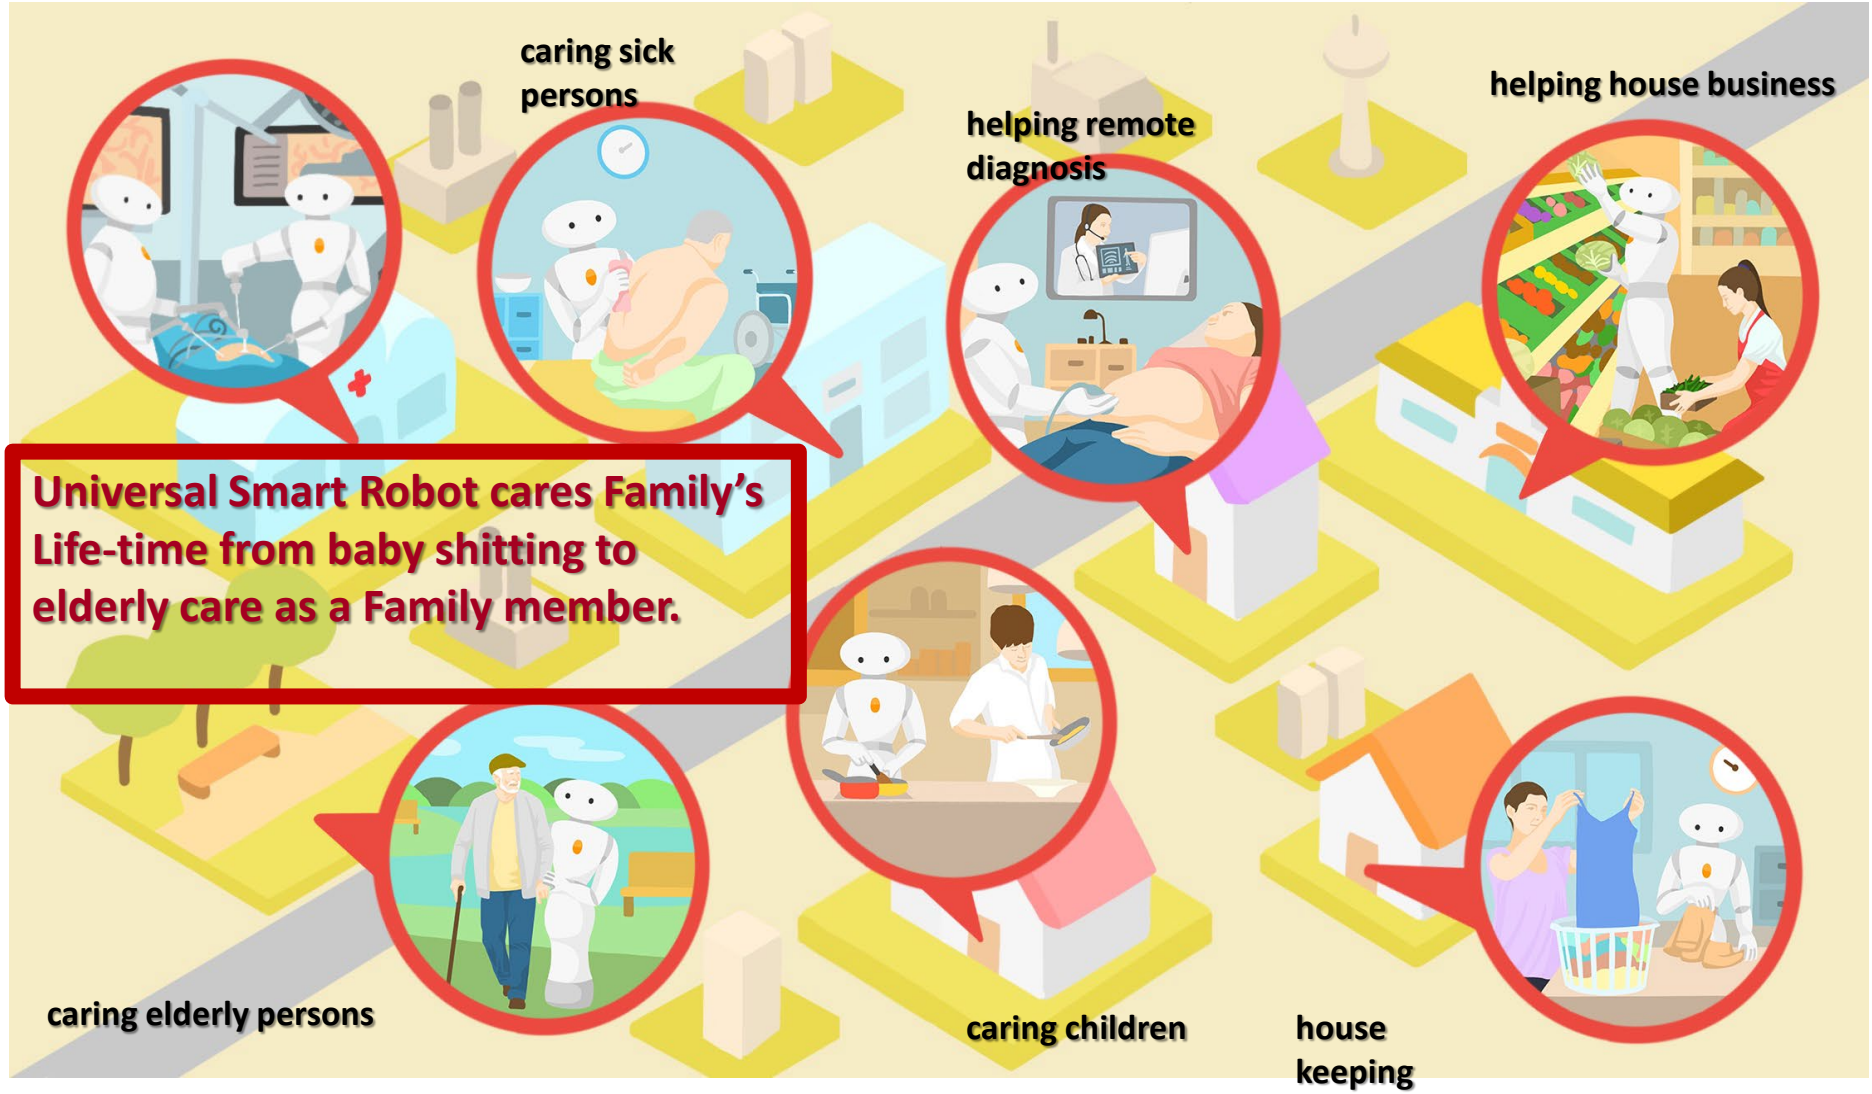

Performance for Multimedia & Scientific Applications on OSCAR Vector Accelerator

(A Vector Processor with Local Memory or Distributed Shared Memory and DMA Controller Managed by the OSCAR Compiler Redesigned Improving Japanese Supercomputer Technology in 1980-2000)

Prof. Keiji Kimura is developing a new custom chip.

OSCAR Vector Accelerator on FPGA

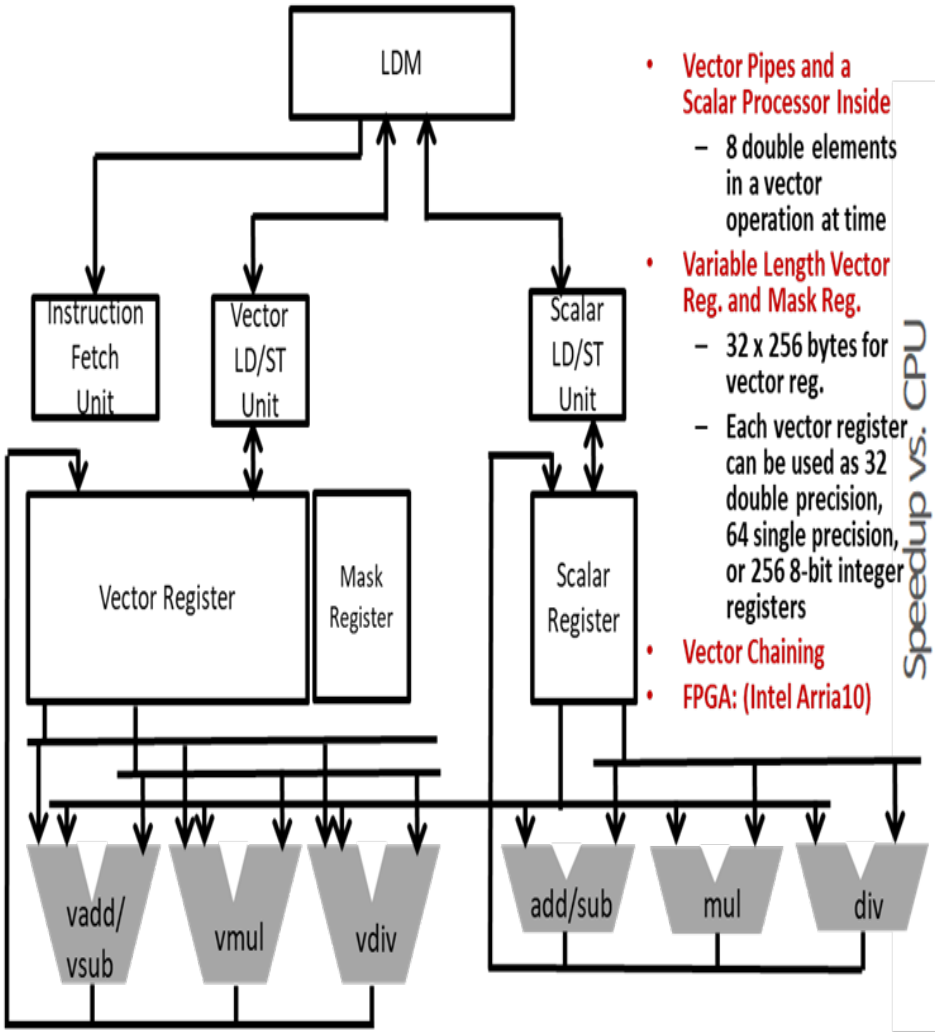

- Vector Pipes and a Scalar Processor Inside
  - 8 double elements in a vector operation at time
- Variable Length Vector Reg. and Mask Reg.
  - 32 x 256 bytes for vector reg.
  - Each vector register can be used as 32 double precision, 64 single precision, or 256 8-bit integer registers
- Vector Chaining
- FPGA: (Intel Arria10)

Speedups against General Purpose Processor by OSCAR Vector Multicore Processor

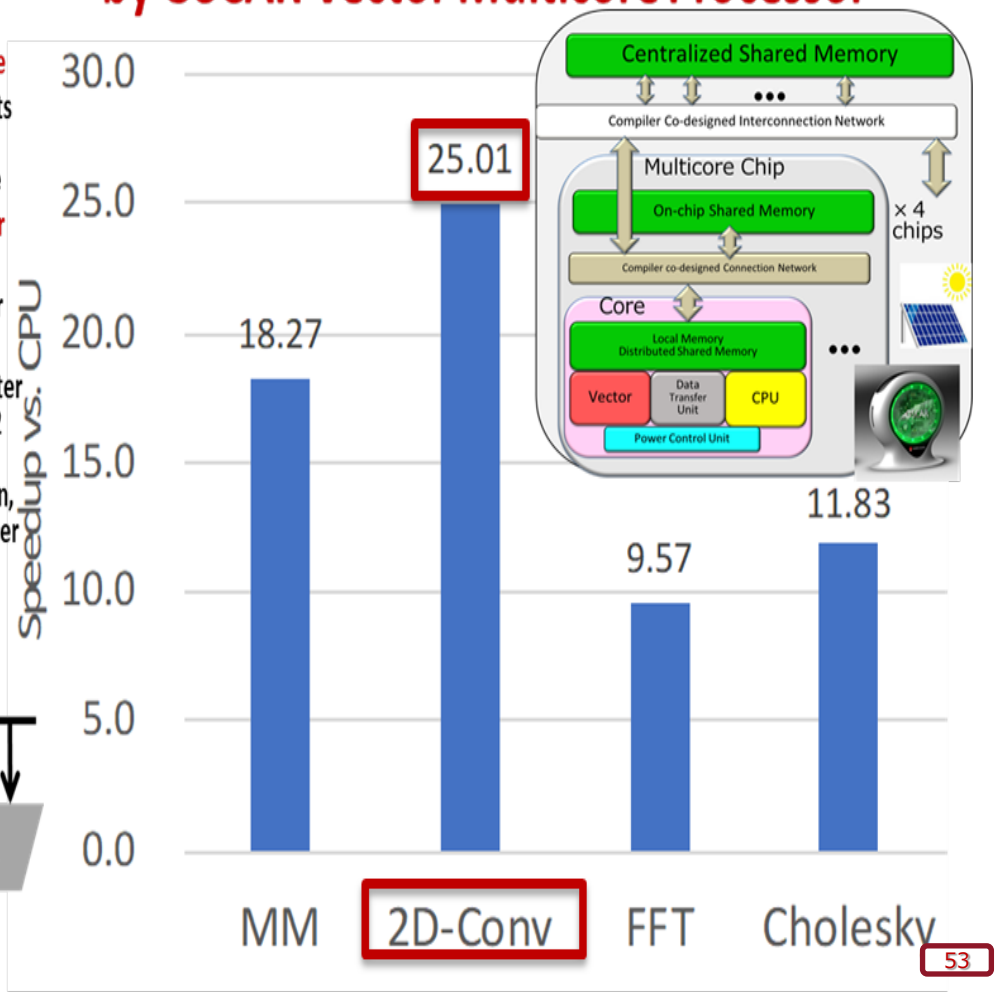

ソーラーパワー・パーソナル・スパコン：新アクセラレータ・グリーンマルチコア

（AI、ビッグデータ、自動運転車、交通制御、ガン治療、地震、ロボット）

世界最高性能・低電力化機能OSCARコンパイラとの協調

ベクトルアクセラレータ併置・
共有メモリ型マルチコアシステム

性能：8TFLOPS, 主メモリ：8TB

電力: 40W, 効率：200GFLOPS/W

- 命令拡張などのプロセッサにも付加できるベクトルアクセラレータ
  - 低消費電力で高速に立ち上がるベクトルで、低コスト設計
- コンパイラによる自動ベクトル・並列化及び自動電力削減
  - 周波数・電源電圧制御機能
  - バリア高速同期・ローカル分散メモリで無駄削減
  - ローカルメモリ利用で低メモリコスト
  - 誰でもチューニングなく使用でき、低コスト短期間ソフト開発可能

グリーン・コンピューティング： 環境に優しい低消費電力・高性能計算

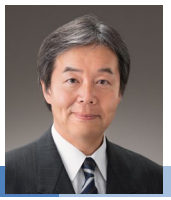

笠原博徳

制御

交通シミュレーション・信号制御 NTTデータ・日立

環境への貢献
カーボンニュートラル

生命・SDGsへの貢献

木村啓二

データセンター: 100WM(火力発電所必要)
→ 1000MW=1GW（原子力発電所必要）

グリーンスパコン

車載(エンジン制御・自動運転Deep Learning・ADAS・MATLAB/Simulink自動並列化) デンソー, ルネサス,NEC

HPC,AI,BigData高速化・低消費電力化

グリーンデータ・クラウドサーバ

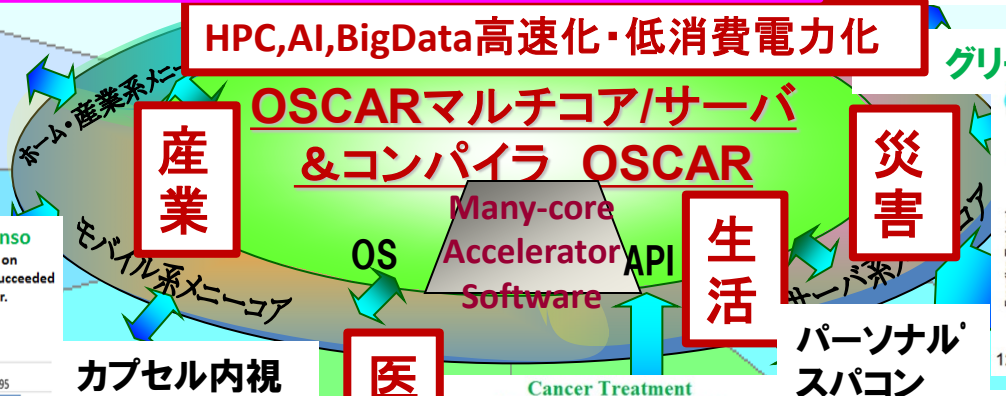

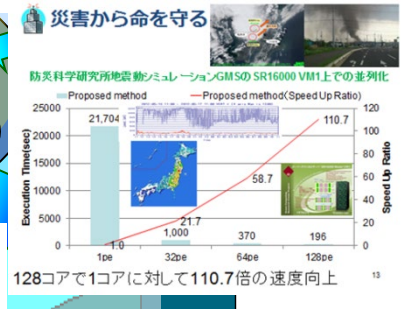

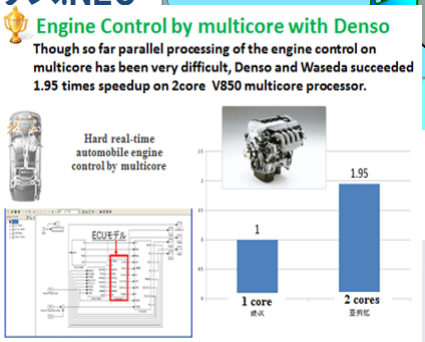

カプセル内視鏡オリンパス

パーソナルスパコン

首都圏直下型地震火災延焼、住民避難指示

スマホ

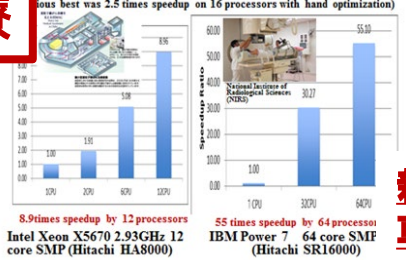

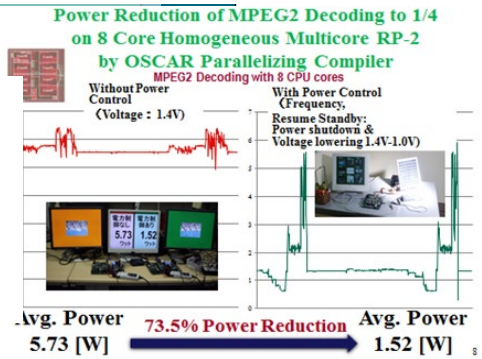

高信頼・低コスト・ソフト開発

FA 三菱

カメラ

太陽光駆動

重粒子ガン治療日立

新幹線車体設計・ディープラーニング・日立

世界の人々への貢献
安全安心便利な製品・サービス
（産官学連携・ベンチャー）

高速化

低消費電力化