

Future of High Performance Green OSCAR Multicore Computing



Hironori Kasahara, Ph.D., IEEE Fellow

IEEE Computer Society President 2018

Professor, Dept. of Computer Science & Engineering

Director, Advanced Multicore Processor Research Institute

Waseda University, Tokyo, Japan

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

1980 BS, 82 MS, 85 Ph.D. , Dept. EE, Waseda Univ.
1985 Visiting Scholar: U. of California, Berkeley
1986 Assistant Prof., 1988 Associate Prof., 1997,
Waseda Univ., Now Dept. of Computer Sci. & Eng.
1989-90 Research Scholar: U. of Illinois, Urbana-
Champaign, Center for Supercomputing R&D
2004 Director, Advanced Multicore Research
Institute, 2017 member: the Engineering Academy
of Japan and the Science Council of Japan

2005 STARC Academia-Industry Research Award
2008 LSI of the Year Second Prize
2008 Intel AsiaAcademic Forum Best Research Award
2010 IEEE CS Golden Core Member Award
2014 Minister of Edu., Sci. & Tech. Research Prize
2015 IPSJ Fellow
2017 IEEE Fellow, IEEE Eta Kappa Nu

Reviewed Papers: 214, Invited Talks: 145, Published
Unexamined Patent Application:59 (Japan, US, GB,
China Granted Patents: 30), Articles in News Papers,
Web News, Medias incl. TV etc.: 572

Committees in Societies and Government 245

IEEE Computer Society President 2018, BoG(2009-
14), Multicore STC Chair (2012-), Japan Chair (2005-
07), IPSJ Chair: HG for Mag. & J. Edit, Sig. on ARC.
【METI/NEDO】 Project Leaders: Multicore for
Consumer Electronics, Advanced Parallelizing
Compiler, Chair: Computer Strategy Committee
【Cabinet Office】 CSTP Supercomputer Strategic
ICT PT, Japan Prize Selection Committees, etc.
【MEXT】 Info. Sci. & Tech. Committee,
Supercomputers (Earth Simulator, HPCI Promo.,
Next Gen. Supercomputer K) Committees, etc.


IEEE Computer Society BoG (Board of Governors) Feb.1, 2018



<https://www.computer.org/web/cshistory/officers-2018>

Multicores for Performance and Low Power

Power consumption is one of the biggest problems for performance scaling from smartphones to cloud servers and supercomputers (“K” more than 10MW) .



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 \propto Frequency)

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

If Frequency is reduced to 1/4
(Ex. 4GHz→1GHz),
Power is reduced to 1/64 and
Performance falls down to 1/4 .

<Multicores>

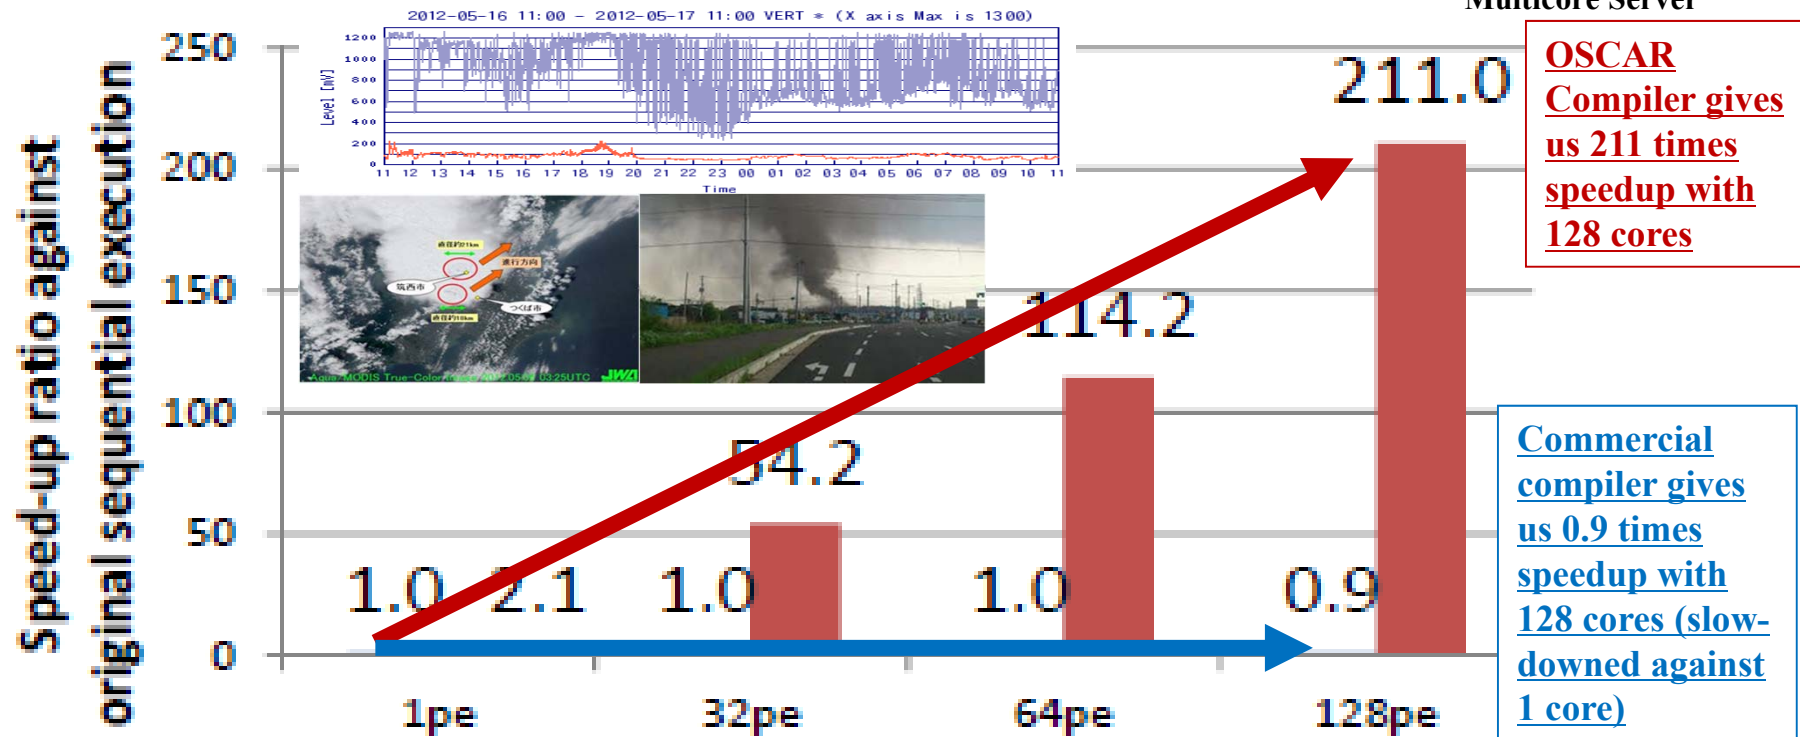
If 8cores are integrated on a chip,
Power is still 1/8 and
Performance becomes 2 times.

- Just more cores don't give us speedup
- Development cost and period of parallel software are getting a bottleneck of development of embedded systems, eg. IoT, Automobile

■ original (sun studio) ■ proposed method




Fjitsu M9000 SPARC Multicore Server



- 4


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)



Avg. Power
5.73 [W]

With Power Control
(Frequency,
Resume Standby:

Power shutdown &
Voltage lowering 1.4V-1.0V)



Avg. Power
1.52 [W]

73.5% Power Reduction



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)

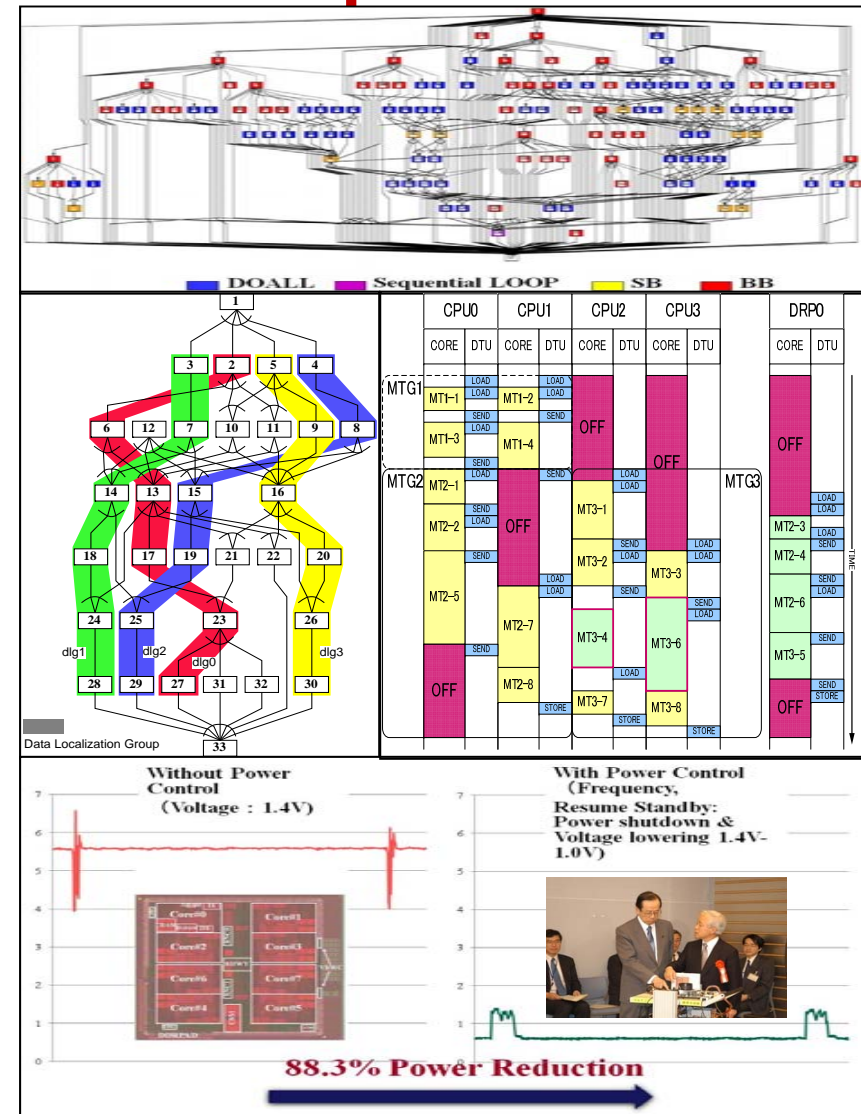
Data Transfer Overlapping^(2016 partially)

Data transfer overlapping using Data Transfer Controllers (DMAs)

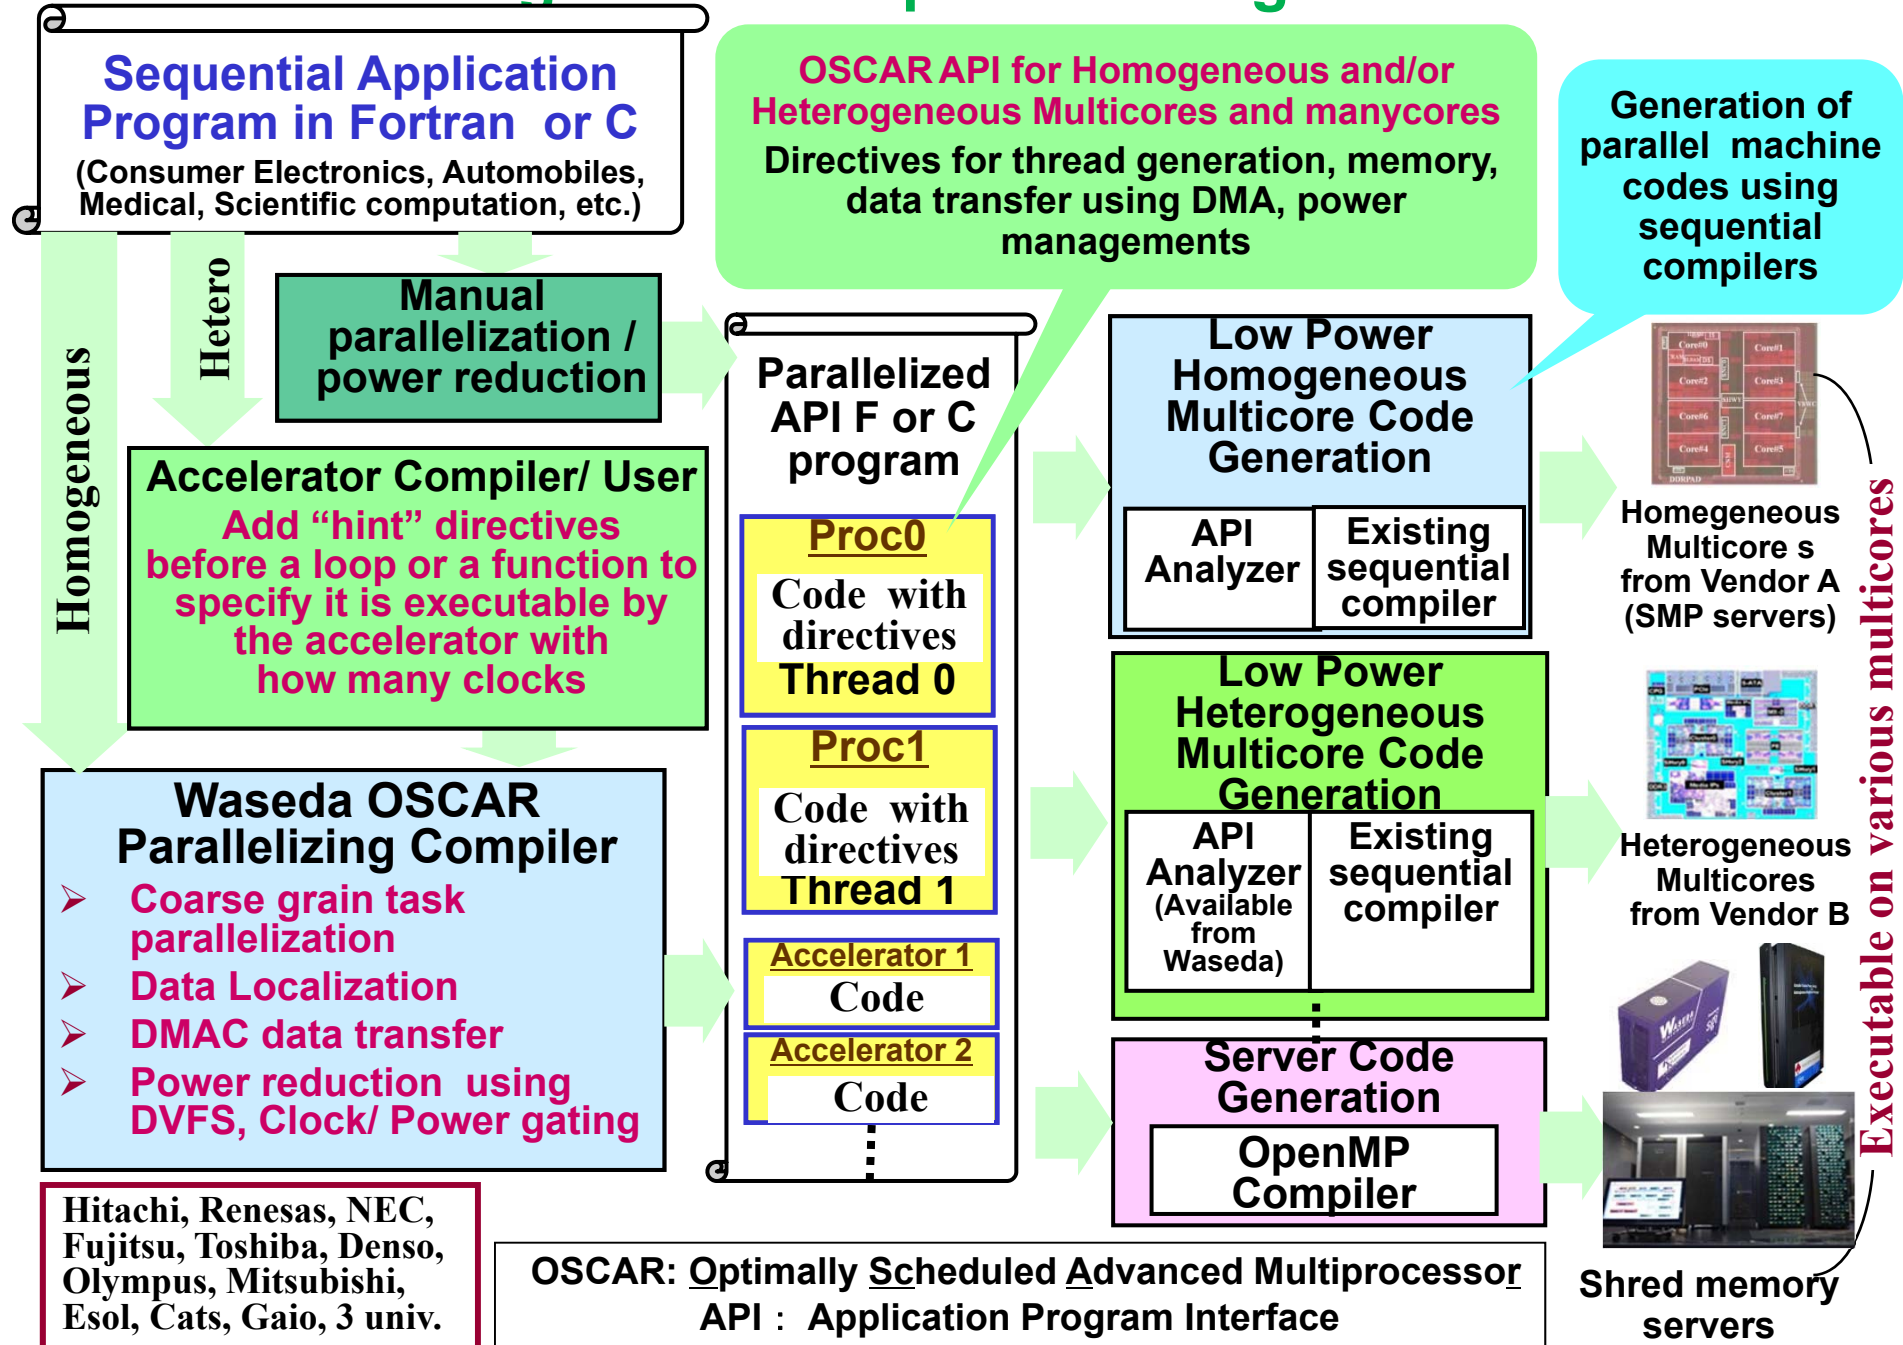
Power Reduction

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

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



Multicore Program Development Using OSCAR API V2.0







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.



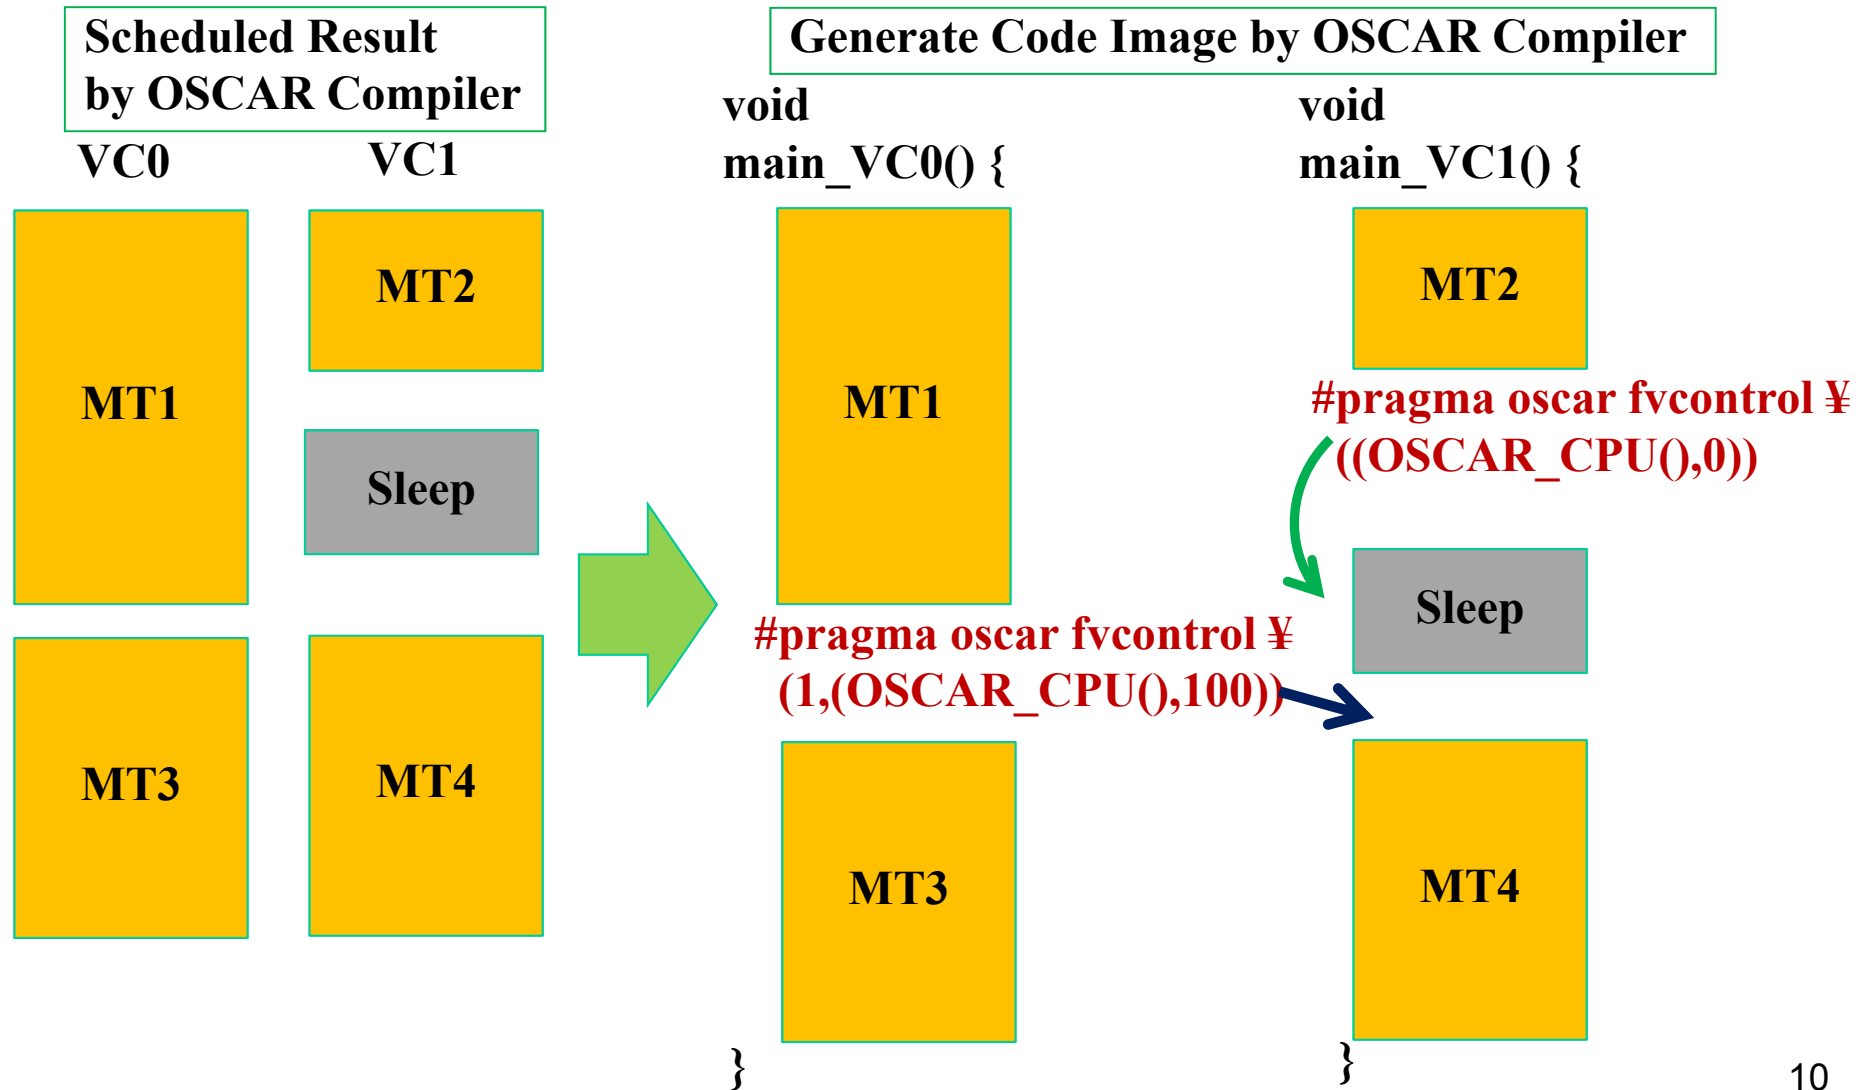
- Hard real-time automobile engine control by multicore using local memories
- Millions of lines C codes consisting conditional branches and basic blocks



Speedup ratio for H.264 and Optical Flow on ARM Cortex-A9 Android 3 cores by OSCAR Automatic Parallelization



Low-Power Optimization with OSCAR API





Automatic Power Reduction on ARM CortexA9 with Android

http://www.youtube.com/channel/UCS43INYEIkC8i_KIgFZYQBQ

ODROID X2

Samsung Exynos4412 Prime, ARM Cortex-A9 Quad core
1.7GHz~0.2GHz, used by Samsung's Galaxy S3





Power for 3cores was reduced to **1/5~1/7** against without software power control

Power for 3cores was reduced to **1/2~1/3** against ordinary 1core execution

Automatic Power Reuction on Intel Haswell


H.264 decoder & Optical Flow (3cores)

H81M-A, Intel Core i7 4770k
Quad core, 3.5GHz~0.8GHz



Power for 3cores was reduced to **1/3~1/4** against without software power control
Power for 3cores was reduced to **2/5~1/3** against ordinary 1core execution

Automatic Power Reduction of OpenCV Face Detection on big.LITTLE ARM Processor



• ODROID-XU3

• Samsung Exynos 5422 Processor



- 4x Cortex-A15 2.0GHz, 4x Cortex-A7 1.4GHz big.LITTLE Architecture

• 2GB 1 DDR2 RAM


Frequency can be changed by

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


(Power7 Based 128 Core Linux SMP) ([LCPC2015](#))



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



Parallelization of 3D-FFT for New Magnetic Material Computation on Hitachi SR16000 Power7 CC-Numa Server



OSCAR optimization

- reducing number of data transpose with interchange, code motion and loop fusion

OSCAR API Ver. 2.0 for Homogeneous/Heterogeneous Multicores and Manycores (LCPC2009Homo, 2010 Hetero)

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

▶ 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

(* from OpenMP)


Software Coherence Control Method on OSCAR Parallelizing Compiler

- Coarse grain task parallelization with **earliest condition analysis** (control and data dependency analysis to detect parallelism among coarse grain tasks).
- OSCAR compiler automatically controls coherence using following simple program restructuring methods:
 - To cope with stale data problems:
 - ◆ **Data synchronization by compilers**
 - To cope with false sharing problem:
 - ◆ **Data Alignment**
 - ◆ **Array Padding**
 - ◆ **Non-cacheable Buffer**



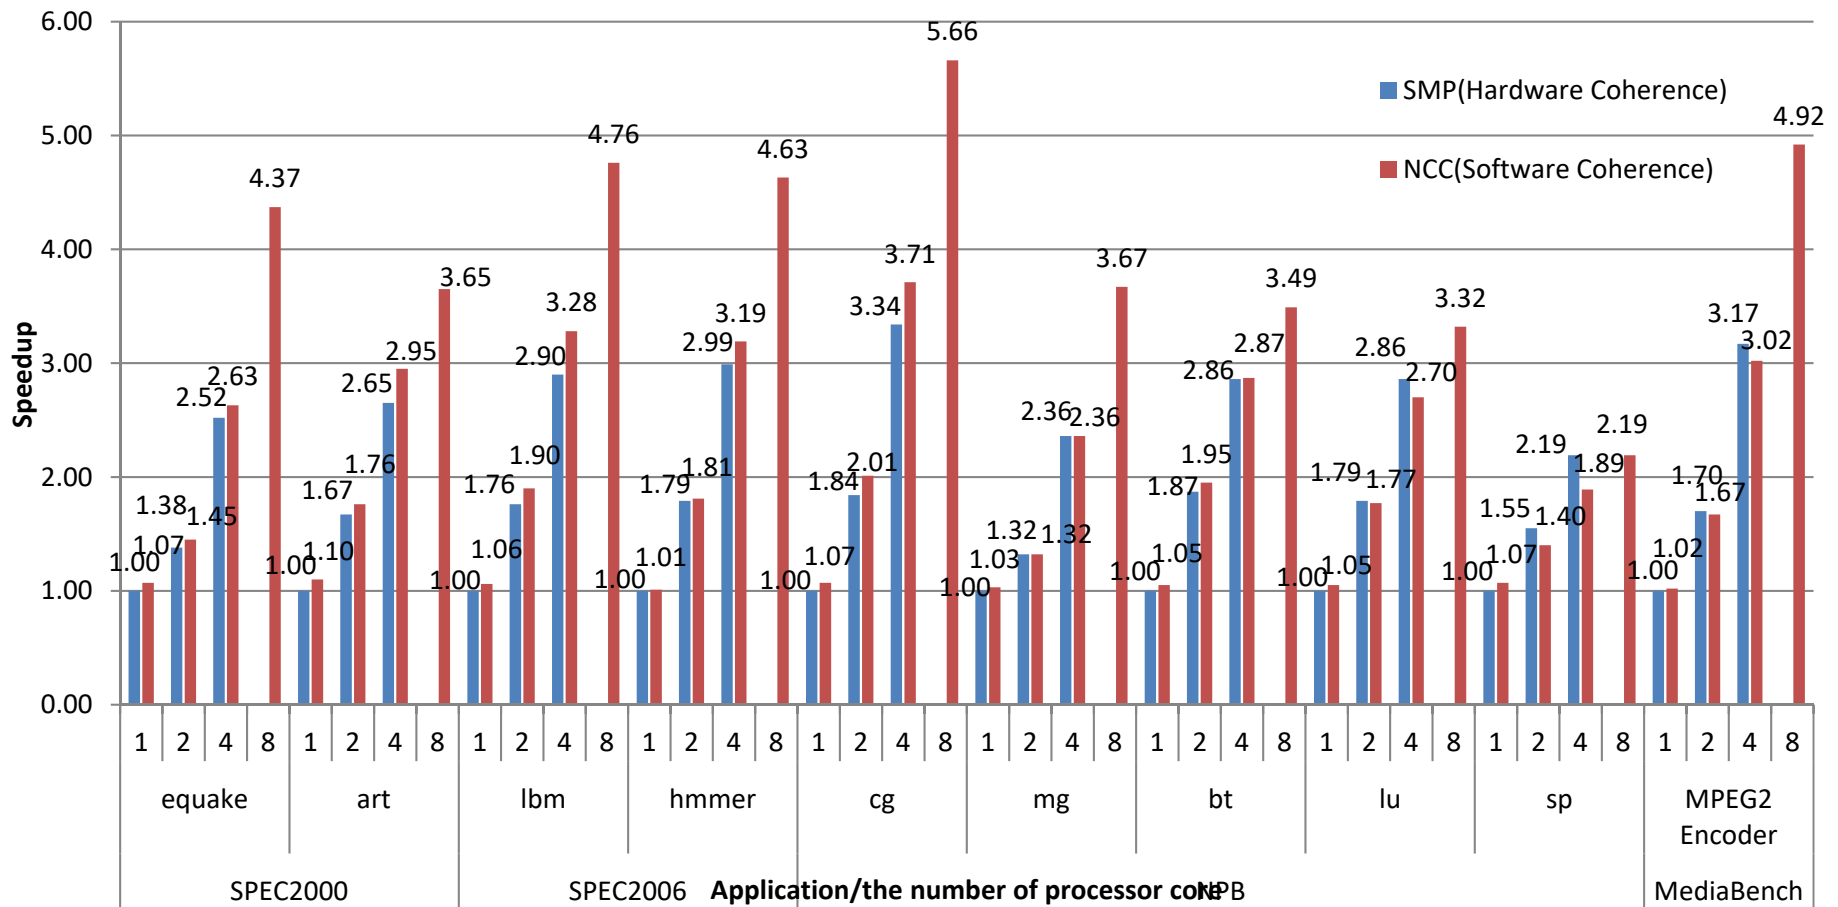
**MTG generated by
earliest executable
condition analysis**

8 Core RP2 Chip Block Diagram

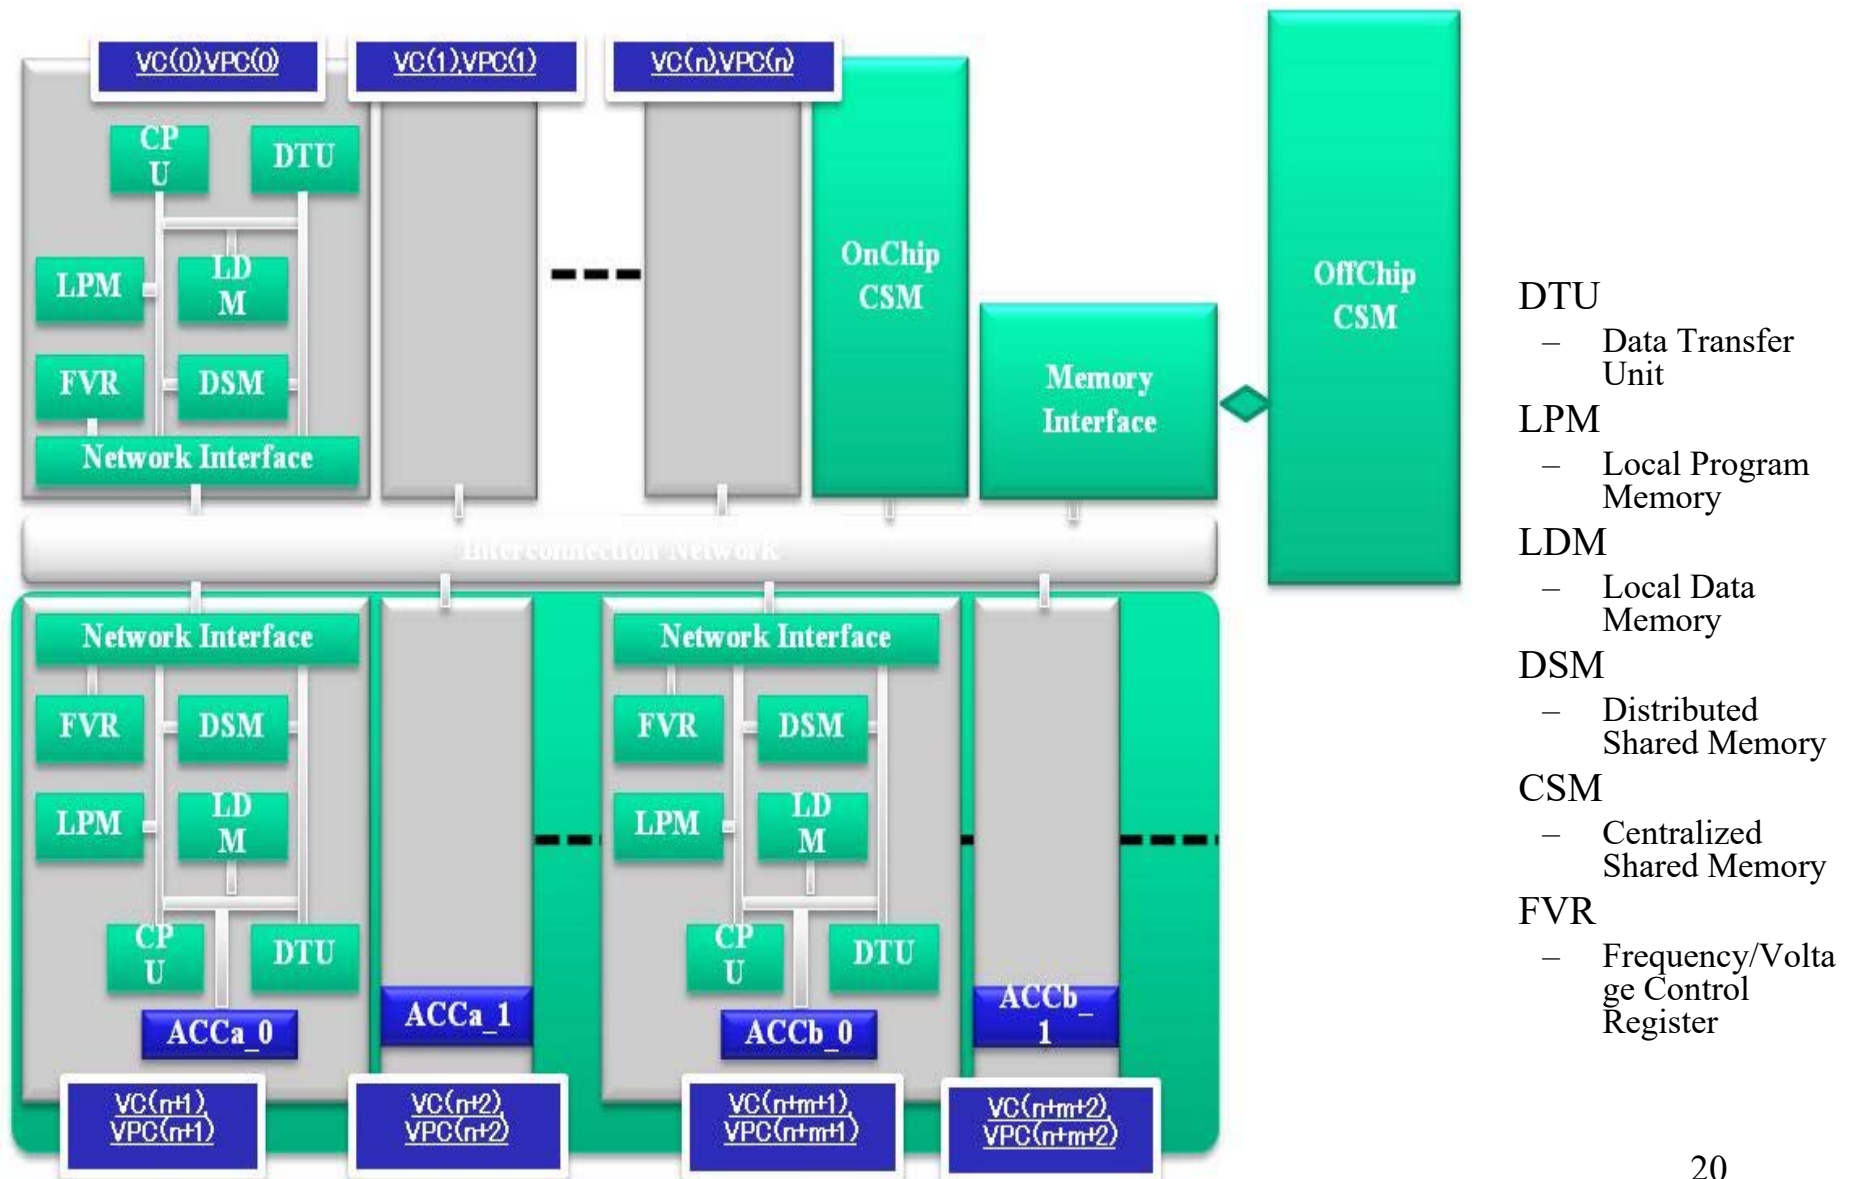


Automatic Software Coherent Control for Manycores

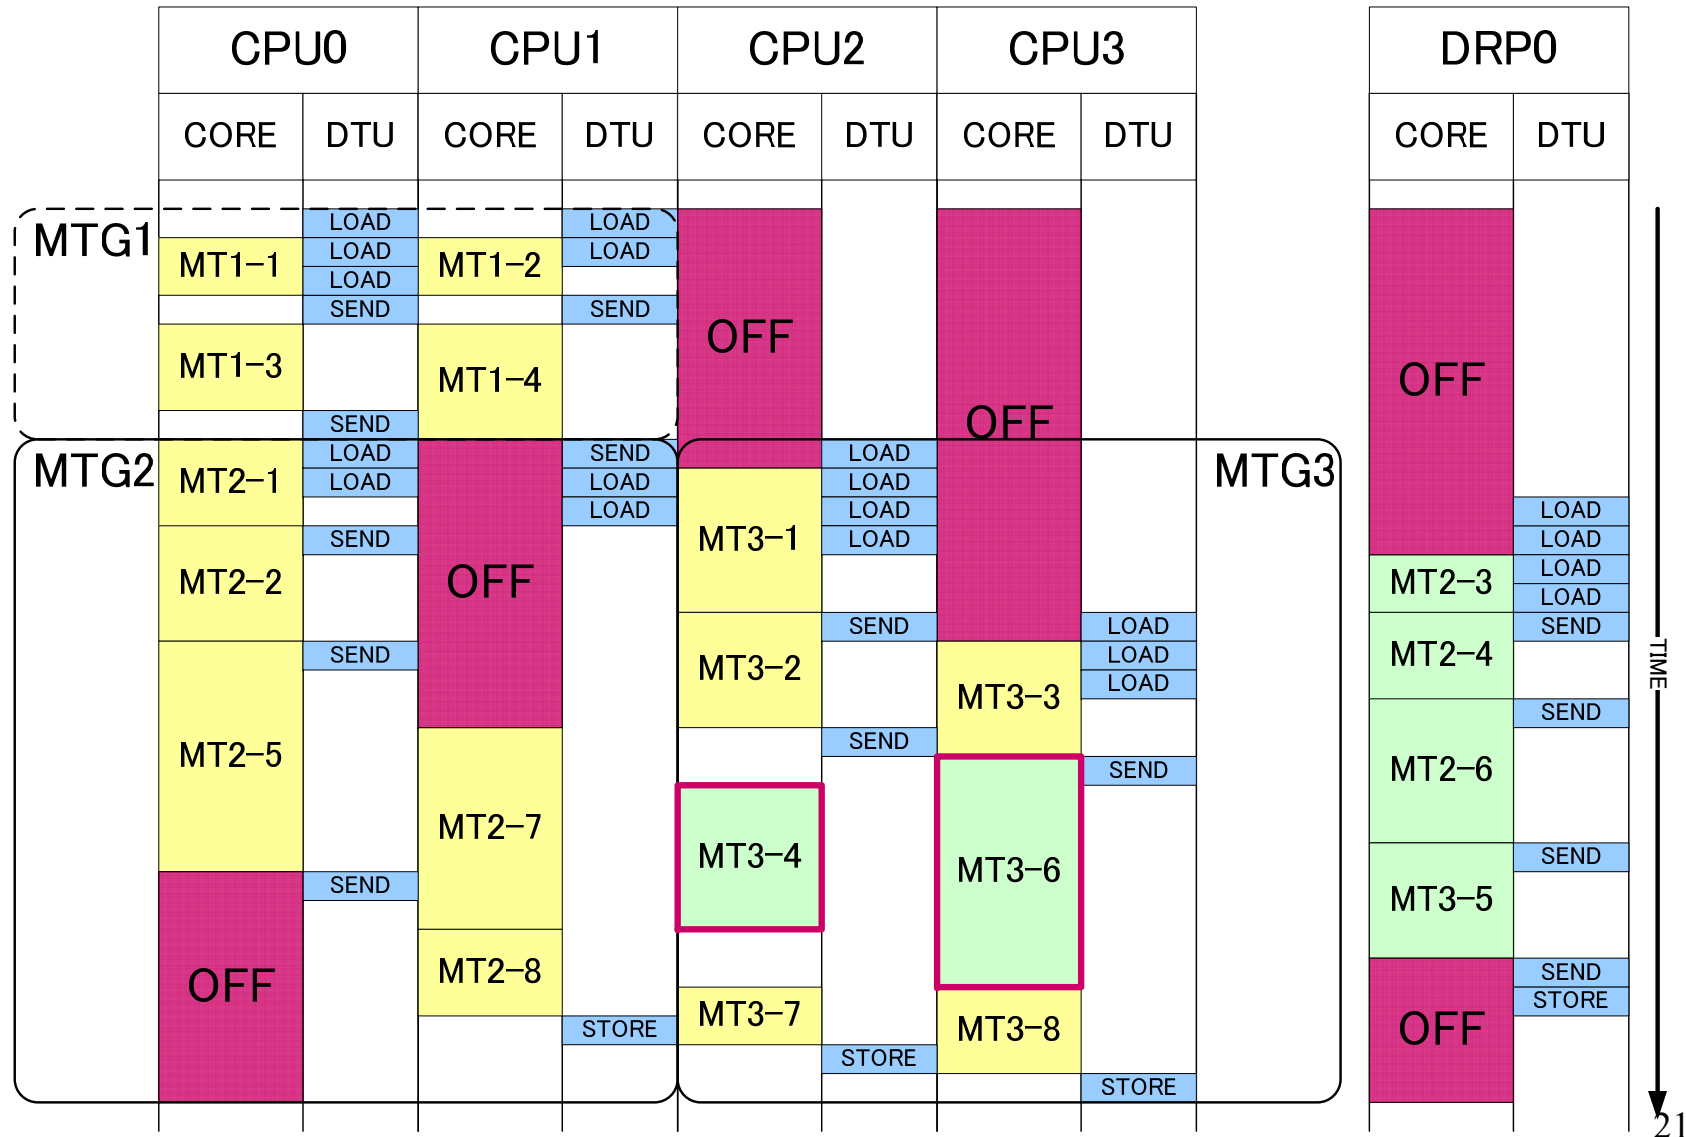
Performance of Software Coherence Control by OSCAR Compiler on 8-core RP2



OSCAR Heterogeneous Multicore




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



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

111[fps]





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]



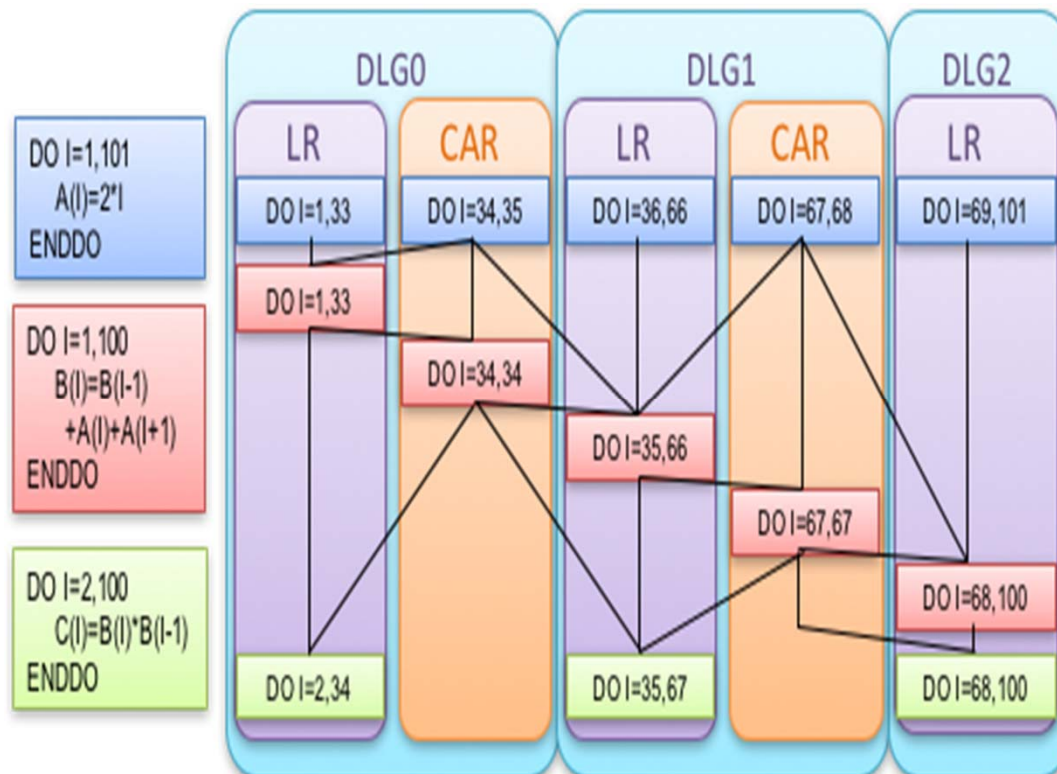
**1cycle : 33[ms]
→30[fps]**

Automatic Local Memory Management


Data Localization: Loop Aligned Decomposition

- Decomposed loop into LR 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




Adjustable Blocks

- Handling a suitable block size for each application
 - different from a fixed block size in cache
 - each block can be divided into smaller blocks with integer and scalar



Multi-dimensional Template Arrays for Improving Readability

- a mapping technique for arrays with varying dimensions
 - each block on LDM corresponds to multiple empty arrays with varying dimensions
 - these arrays have an additional dimension to store the corresponding block number
 - $TA[Block\#][\]$ for single dimension
 - $TA[Block\#][\][\]$ for double dimension
 - $TA[Block\#][\][\][\]$ for triple dimension
 - ...
- LDM are represented as a one dimensional array
 - without Template Arrays, multi-dimensional arrays have complex index calculations
 - $A[i][j][k] \rightarrow TA[offset + i' * L + j' * M + k']$
 - Template Arrays provide readability
 - $A[i][j][k] \rightarrow TA[Block\#][i'][j'][k']$



Block Replacement Policy

□ Compiler Control Memory block Replacement

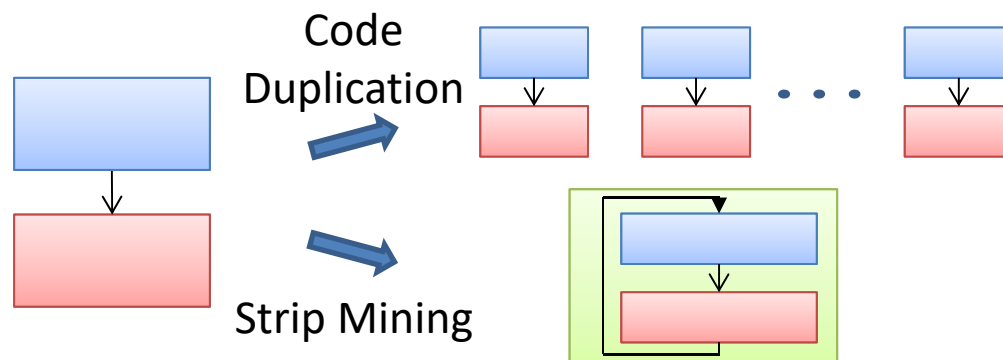
- using live, dead and reuse information of each variable from the scheduled result
- different from LRU in cache that does not use data dependence information

□ Block Eviction Priority Policy

1. (Dead) Variables that will not be accessed later in the program
2. Variables that are accessed only by other processor cores
3. Variables that will be later accessed by the current processor core
4. Variables that will immediately be accessed by the current processor core

Code Compaction by Strip Mining

- Previous approach produces duplicate code
 - generates multiple copies of the loop body which leads to code bloat
- Proposed method adopts code compaction
 - based on strip mining
 - multi-dimensional loop can be restructured




```
for (i = 0; i < 16; i++)  
  for (j = 0; j < 64; j++)  
    a[i][j] = i + j;
```

```
for (i = 0; i < 15; i++)  
  for (j = 0; j < 63; j++)  
    b[i][j] = a[i][j] + a[i+1][j+1];
```



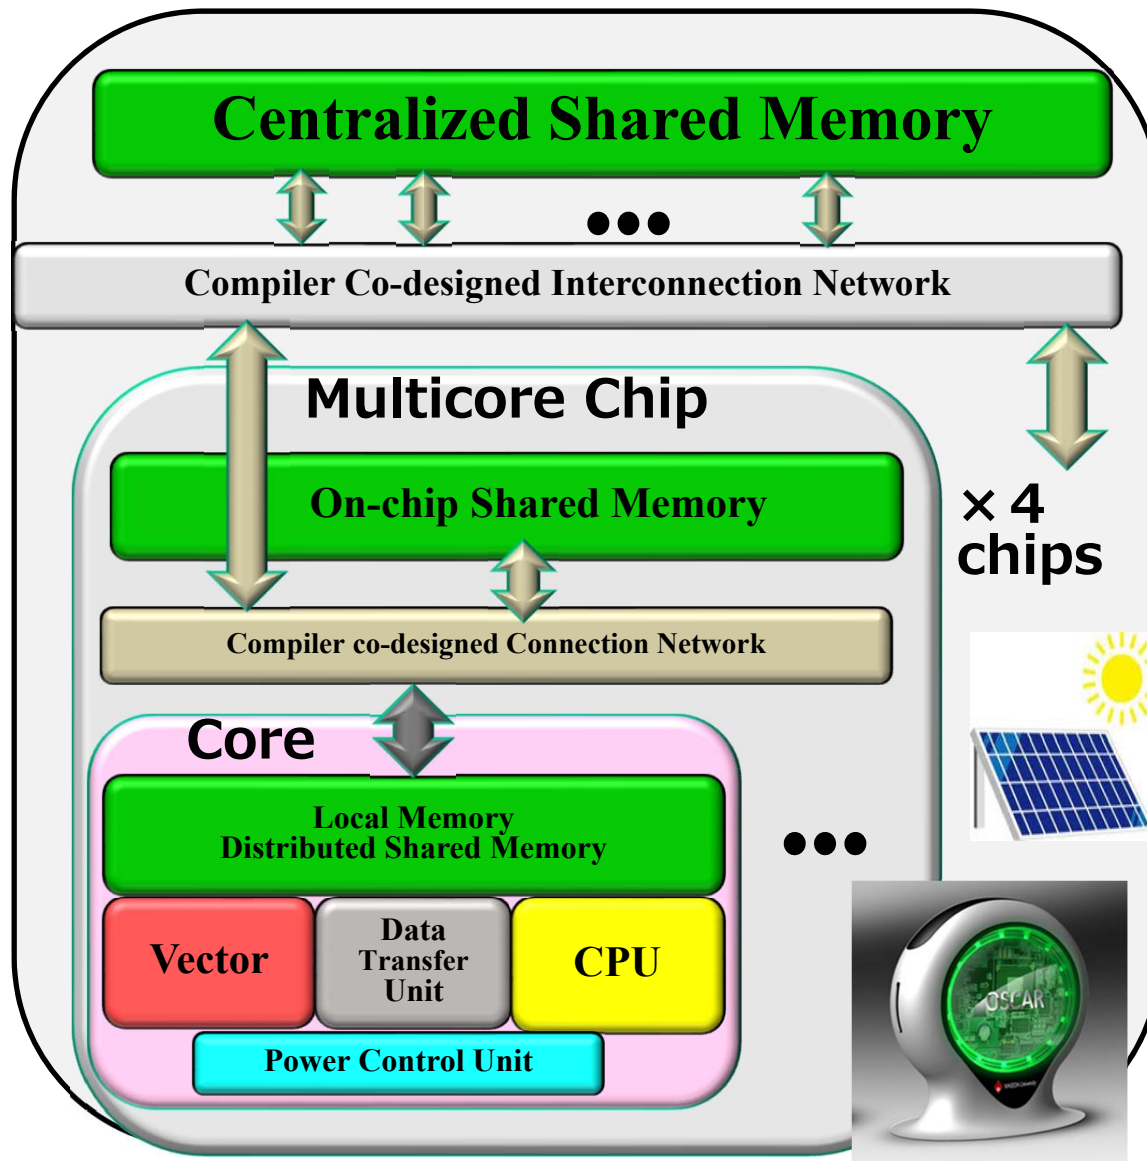
```
for (ii = 0; ii < 15; ii += 8)  
  for (jj = 0; jj < 63; jj += 32)  
    for (i = ii; i < min(15, ii + 8 + 1); i++)  
      for (j = jj; j < min(63, jj + 32 + 1); j++)  
        a[i][j] = i + j;  
    for (i = ii; i < min(15, ii + 8); i++)  
      for (j = jj; j < min(63, jj + 32); j++)  
        b[i][j] = a[i][j] + a[i+1][j+1];
```

Speedups by the Proposed Local Memory Management Compared with Utilizing Shared Memory on Benchmarks Application using RP2



20.12 times speedup for 8cores execution using local memory against sequential execution using off-chip shared memory of RP2 for the AACenc

OSCAR Vector Multicore and Compiler for Embedded to Servers with OSCAR Technology




Target:

- **Solar Powered**
- **Compiler power reduction.**
 - **Fully automatic parallelization and vectorization including local memory management and data transfer.**

Vector Accelerator

Features

- **Attachable for any CPUs (Intel, ARM, IBM)**
- **Data driven initiation by sync flags**



Function Units [tentative]

- **Vector Function Unit**
 - 8 double precision ops/clock
 - 64 characters ops/clock
 - Variable vector register length
 - Chaining LD/ST & Vector pipes
- **Scalar Function Unit**

Registers[tentative]

- **Vector Register** 256Bytes/entry, 32entry
- **Scalar Register** 8Bytes/entry
- **Floating Point Register** 8Bytes/entry
- **Mask Register** 32Bytes/entry

OSCAR Technology Corp.

Started up on Feb. 28, 2013:

Licensing the all patents and OSCAR compiler from Waseda Univ.

Founder and CEO: Dr. T. Ono (Ex- CEO of First Section-listed Company, Director of National U., Invited Prof. of Waseda U.)

Executives: **Dr. M. Ohashi : COO** (Ex- OO of Ono Sokki)

Mr. A. Nodomi : CTO (Ex- Spansion)

Mr. N. Ito (Ex- Visiting Prof. Tokyo Agricult. And Tech. U.)

Dr. K. Shirai (Ex- President of Waseda U., Ex- Chairman of Japanese Open U.)

Mr. K. Ashida (Ex- VP Sumitomo Trading, Adhida Consult. CEO)

Mr. S. Tsuchida (Co-Chief Investment Officer of Innovation Network Corp. of Japan)

Auditor: **Mr. S. Honda** (Ex- Senior VP and General Manager of MUFG)

Dr. S. Matuda (Emeritus Prof. of Waseda U., Chairman of WERU INVESTMENT)


Mr. Y. Hirowatari (President of AGS Consulting)

Advisors: **Prof. H. Kasahara** (Waseda U.)


Prof. K. Kimura (Waseda U.)



富士通ベクトルスパコンVPP5000



OSCAR TECHNOLOGY CORPORATION



Future Multicore Products



Next Generation Automobiles

- Safer, more comfortable, energy efficient, environment friendly
- Cameras, radar, car2car communication, internet information integrated brake, steering, engine, motor control

Smart phones



- From everyday recharging to less than once a week
- Solar powered operation in emergency condition
- Keep health

Advanced medical systems



Cancer treatment, Drinkable inner camera

- Emergency solar powered
- No cooling fan, No dust, clean usable inside OP room



Personal / Regional Supercomputers



Solar powered with more than 100 times power efficient : FLOPS/W

- Regional Disaster Simulators saving lives from tornadoes, localized heavy rain, fires with earthquakes

Summary

- Waseda University Green Computing Systems R&D Center supported by METI has been researching **on low-power high performance Green Multicore** hardware, software and application with industry including Hitachi, Fujitsu, NEC, Renesas, Denso, Toyota, Olympus and OSCAR Technology.
- OSCAR Automatic Parallelizing and Power Reducing Compiler has succeeded speedup and/or power reduction of scientific applications including “Earthquake Wave Propagation”, medical applications including “Cancer Treatment Using Carbon Ion”, and “Drinkable Inner Camera”, industry application including “Automobile Engine Control”, “Smartphone”, and “Wireless communication Base Band Processing” on various multicores from different vendors including Intel, ARM, IBM, AMD, Qualcomm, Freescale, Renesas and Fujitsu.
- In automatic parallelization, **110 times speedup** for **“Earthquake Wave Propagation Simulation”** on **128 cores of IBM Power 7** against 1 core, **55 times speedup** for **“Carbon Ion Radiotherapy Cancer Treatment”** on **64cores IBM Power7**, **1.95 times** for **“Automobile Engine Control”** on **Renesas 2 cores** using **SH4A or V850**, **55 times** for **“JPEG-XR Encoding for Capsule Inner Cameras”** on **Tilera 64 cores Tile64 manycore**.
 - The compiler will be available on market from OSCAR Technology.
- In automatic power reduction, consumed powers for real-time multi-media applications like **Human face detection, H.264, mpeg2 and optical flow** were **reduced to 1/2 or 1/3 using 3 cores** of ARM Cortex A9 and Intel Haswell and **1/4 using Renesas SH4A 8 cores** against ordinary single core execution.
- **Local memory management** for automobiles and **software coherent control** have been patented and already realized by OSCAR compiler.