

### Two-Year Experiences with SX-ACE and Design Space Exploration of the Next Generation Vector System

### Hiroaki Kobayashi

Deputy Directer for the HPC Strategy of Cybersicence Center Professor of Graduate School of Information Sciences Chair of Computer and Mathematics Sciences Department

> 'Tohoku University koba@tohoku.ac.jp

26th WSSP October 10-11, 2017



### Today's Agenda

★ A two-year experience with SX-ACE since 2015

- Lessons Learned from SX-ACE Operations and Applications
  Developments
- Some thoughts for design and development of the next generation vector system
  - Design Constraints
  - Domain-Specific Architecture Design in the Era of End-of-Moore's Law

Hiroaki Kobayashi, Tohoku University



#### Tohoku Univ.'s Supercomputer System (2015.2.20~) HPCI



October 10-11, 2017



### Operation Statistics of SX-ACE (Normalized by SX-9 Data)





### Applications



26th WSSP

October 10-11, 2017

![](_page_5_Picture_1.jpeg)

### Power Consumption of the Cooling System Effect of Fresh-Air Cooling

![](_page_5_Figure_3.jpeg)

![](_page_6_Picture_1.jpeg)

### Node-Core Activity:

Effect of Automatic Core-Node Activation/Deactivation Control

![](_page_6_Figure_4.jpeg)

Month/Day

![](_page_7_Picture_1.jpeg)

### Features of the SX-ACE Vector Processor

- 4 Core Configuration, each with High-Performance Vector-Processing Unit and Scalar Processing Unit
  - 272Gflop/s of VPU + 4Gflop/s of SPU per socket
    - 68Gflop/s + IGflop/s per core
  - IMB private ADB per core (4MB per socket)
    - Software-controlled on-chip memory for vector load/store
    - 4x compared with SX-9
    - 4-way set-associative
    - MSHR with 512 entries (address+data)
    - 256GB/s to/from Vec. Reg.
      - 4B/F for Multiply-Add operations
  - <sup>2</sup> 256 GB/s memory bandwidth, Shared with 4 cores
    - IB/F in 4-core Multiply-Add operations
      - $\sim$  4B/F in 1-core Multiply-Add operations
    - 128 memory banks per socket
    - 128B-Block access to the memory (16B of SX-9)

SX-ACE Processor Architecture

![](_page_7_Figure_19.jpeg)

Source: NEC

- Other improvement and new mechanisms to enhance vector processing capability, especially for efficient handling of short vectors operations and indirect memory accesses
- Dynamically and Statically reordering execution of vector gather/scatter operations
- Advanced data forwarding in vector pipes chaining
- Shorter memory latency than SX-9

![](_page_8_Picture_1.jpeg)

### Dynamically and Statically Reordering Execution of Vector Gather/Scatter Operations

#### Vectorization

- NODEP: Load & Store operations of each element of a list-vector are executed in parallel.
  - Vectorization Possible
- Giving Priority to Vector Gather(Loads) over Scatter(Vector) Stores
  - GTHREORDER
    - Compiler statically reschedules vector gather/scatter operations to give the priority to Loads over Stores

#### NOCONFLICT

• Assign a flag that there are no data hazard between loads and stores to access different arrays, and dynamic reordering becomes possible by hardware to give the priority to VLD and VGT (gather) over VST and VSC (scatter)

#### SOVERTAKE

 Assign a flag that there is no data hazard between loads and stores within the same array, and dynamic reordering become possible by hardware to give the priority to VLD and VGT over VST and VSC

![](_page_9_Picture_1.jpeg)

### Rescheduling vector gather/scatter operations by Compiler

![](_page_9_Figure_3.jpeg)

![](_page_10_Picture_1.jpeg)

Out of Order execution for vector load&gather/ store&scatter operations across iterations by hardware

![](_page_10_Figure_3.jpeg)

![](_page_11_Picture_1.jpeg)

Out of Order execution for vector load&gather/ store&scatter operations across iterations by hardware

![](_page_11_Figure_3.jpeg)

![](_page_12_Picture_1.jpeg)

### Effects of Giving the Priority to LOADs over STOREs

![](_page_12_Figure_3.jpeg)

![](_page_13_Picture_0.jpeg)

### **Future Vector Systems R&D\***

\*This work is partially conducted with NEC, but the contents do not reflect any future products of NEC

Hiroaki Kobayashi, Tohoku University

![](_page_14_Picture_1.jpeg)

### Timeline of the Cyberscience Center HPC System Development and R&D For the Future

![](_page_14_Figure_3.jpeg)

![](_page_15_Picture_1.jpeg)

### Scaling may be End, but Silicon is not End! And Use it Smart and Effective!

- ✓ We are facing the end of Moore's low due to the physical limitations, and the transistor cost is hard to reduce, however
- Silicon is still fundamental constructing material for computing platforms such as plastic, steel and concrete for automobiles, buildings and home appliances.
  - So, we have to become much more smart for design of Future HEC systems.

![](_page_15_Figure_6.jpeg)

![](_page_15_Picture_7.jpeg)

It's time to focus on Domain-Specific Architectures for computationintensive, memory-intensive, I/O intensive, low-precision computing… etc applications to improve silicon/power efficiency!

![](_page_16_Picture_1.jpeg)

#### Slide at 16th WSSP in 2012

# So What to Do ~NNGV(?) Design Pendulum~

![](_page_16_Figure_4.jpeg)

Hiroaki Kobayashi, Tohoku University

![](_page_17_Picture_1.jpeg)

#### Aurora: NEC's Next Vector Product

![](_page_17_Figure_3.jpeg)

![](_page_18_Picture_1.jpeg)

### Now "Makimoto"s Wave" also hits the HPC community.

#### Makimoto's Wave

The cyclical nature of the semiconductor industry alternates between standardization and customization

![](_page_18_Figure_5.jpeg)

IEEE Computer Dec. 2013

\* "one fits all" no longer make sense! means, shifting from General Purpose to Special Purpose

![](_page_19_Picture_1.jpeg)

![](_page_20_Picture_1.jpeg)

#### One Feasible Solution to Future Vector Computing

![](_page_20_Figure_3.jpeg)

![](_page_21_Picture_1.jpeg)

### Design Space Exploration of the Next Generation Vector Architecture

Improving Flop/s-Memory Balance for Vector Operations

- Effects of FMAs instead of traditional SX vector pipes
- ★On-chip-Memory Capacity
  - Capacity becomes a help to recover the limited memory bandwidth?

Heterogeneous Memory Subsystem Design for Bandwidth-Capacity Tradeoff

- Stacked High-Bandwidth near-memory and DDR-based Far-memory Combined.
- A stacked high-bandwidth memory subsystem becomes a backup of ADB?

Hiroaki Kobayashi, Tohoku University

![](_page_22_Picture_1.jpeg)

### Design Space Exploration of the Next Generation Vector Architecture

Effects of FMA S instead of using traditional vector pipes
 Operations of the Himeno benchmark kernel

![](_page_22_Figure_4.jpeg)

![](_page_23_Picture_1.jpeg)

### Design Space Exploration of the Next Generation Vector Architecture

![](_page_23_Figure_3.jpeg)

![](_page_24_Picture_1.jpeg)

## Increasing ADB Capacity Becomes a Help to Recover the Limited Memory Bandwidth?

![](_page_24_Figure_3.jpeg)

**HBM** only

![](_page_25_Picture_1.jpeg)

### Memory Subsystem Design Choice

HBM: Bandwidth-Oriented

Stacked high-bandwidth memory subsystem

DDR: Capacity-Oriented

DDR-based memory subsystem

HMA: Bandwidth-Capacity Heterogeneous

Stacked High-Bandwidth near-memory and DDRbased Far-memory Combined for tradeoff between BW and capacity.

How can we effectively use these under the consideration of applications characteristics!?

CPU

DDR only

![](_page_25_Picture_12.jpeg)

![](_page_25_Figure_13.jpeg)

Hiroaki Kobayashi, Tohoku University

![](_page_26_Picture_1.jpeg)

#### Memory Management is Not So Easy!

![](_page_26_Figure_3.jpeg)

![](_page_27_Picture_1.jpeg)

### Summary

#### $\star$ SX-ACE shows high sustained performance compared with SX-9 and other modern **HEC** systems

- $\checkmark$  achieved the same single core performance in practical applications even with 60% of peak performance of SX-9
- ✓ NoI. computing-efficiency and power-efficiency in the HPCG Benchmark ranking
- $\checkmark$  Pave the way to a new social infrastructure for homeland safety in Japan

 $\star$  Well balanced HEC systems regarding memory performance is the key to success for realizing high productivity and power&silicon efficiency in science and engineering simulations

 $\checkmark$  Consider domain-specific architectures with co-design with individual applications domains

• One-fits-all does not make sense any more!

![](_page_27_Picture_10.jpeg)

✓ Think different with Smart Force from Brute Force in HPC design

• Computing Quality is the first, for power-efficient HPC , not Quantity for productive HPC!

![](_page_27_Picture_13.jpeg)

Source: Toyota

![](_page_27_Picture_15.jpeg)