

Final Program

April 13, 2015 Main Hall(7th Floor), Yokohama Joho Bunka Center
(Yokohama Media & Communications Center)

- 13:30-14:30** **Special Invited Lecture 1**
Chair: Tohru Ishihara (Kyoto Univ.)
- 13:30-14:30** **Adaptive Many-Core Architectures for Speed and Power Convergence in Advanced Technology Nodes**
Edith Beigné (CEA-LETI MINATEC, France)

Abstract: With the increasing complexity of today's many-core applications, extremely high performance has become the main requirement. However, high performances do not only mean high speed but also low power. For example, in wireless internet devices, very high speed is mandatory for games or video computing while it is necessary to save dynamic and static power for low speed applications in order to improve the battery life. The convergence between high speed and low power is very difficult to reach. Most of the time, ultra low power architectures cannot reach high speed and conversely, at high speed, a lot of power is consumed. Using energy efficient architectures is the only way to achieve a good compromise between speed and power. In this talk, we will first overview fine-grain adaptive Voltage and Frequency Scaling architectures for many-core. Hardware issues will be discussed and some design solutions will be proposed for good performances results. Those architectures are, however, requiring a Wide Voltage Range of operation to reduce power and increase energy efficiency. We will then focus, during the second part of this talk, on Ultra Wide Voltage Range (UWVR) design challenges at the nanometer regime. How to improve the trade-off between leakage, variability and speed at low-voltage? Obviously the trend is to use thin film devices. Undoped thin-film planar FDSOI devices are being investigated in this presentation as an alternative to bulk devices in 28nm node and beyond, thanks to its excellent short-channel electrostatic control, low leakage currents and immunity to random dopant fluctuation. This compelling technology appears to meet the needs of nomadic devices, combining high performance and low power consumption. A major challenge for this technology is to provide various device threshold voltages (VT), trading off power consumption and speed. This presentation will finally highlight the development of an UWVR multi-VT design platform in FDSOI planar technology on Ultra Thin Body and Box (UTBB) for the 28nm node. The efficient use of an adaptive voltage and frequency scaling architecture has been proved on a 32-bit VLIW DSP exhibiting outstanding silicon results in terms of speed and energy. The use of an efficient Body Biasing (BB) shows an extremely efficient performance tuning for high energy efficiency. To conclude, this talk will give a short overview of FDSOI performances for Internet of Things future applications.

- 14:30-14:50** **Break**
- 14:50-15:50** **Special Invited Lecture 2**
Chair: Tohru Ishihara (Kyoto Univ.)
- 14:50-15:50** **System-Level Energy Management in Many-Core Systems Utilizing Distributed Speed-Power Controllers**
Anca Molnos (CEA-LETI, France)

Abstract: Energy efficiency is one of the crucial concerns today in computing systems ranging from small connected devices to large data-centers. This issue is addressed at various levels, and recently we have witnessed a lot of progress in methods to control speed and power consumption of digital circuits. Notable examples are fine-grain adaptive voltage and frequency scaling, and the adoption of new technologies such as Fully-Depleted Silicon On Insulator (FDSOI). These advances however bring new knobs to tradeoff power and speed, e.g., supply voltage, body-bias voltage, which, in turn, open interesting questions about how to fully take advantage of their potential at software level. This talk we will present methods to reduce power consumption of applications and the tradeoffs therein. As a research vehicle, we have the case of a low-power many-core architecture with several power domains and distributed speed-power controllers. We will study the impact of adaptive voltage scaling and discuss methods to determine the optimal power modes, both with benefits at system level, in the context of advanced technologies such as FDSOI.

15:50-16:10 Break

16:10-17:10 Special Invited Lecture 3
Chair: Tohru Ishihara (Kyoto Univ.)

16:10-17:10 Towards Open-Source Development of Autonomous Vehicles
Shinpei Kato (Nagoya Univ.)

Abstract: Autonomous driving is composed of perception, planning and control technologies. Perception components are supposed to understand scenes of driving in real-time with, for example, object detection and self-localization. Planning components use the results of perception to determine the path of driving including behaviors and motions. Finally, control components drive the vehicle in accordance with the plan. These components are often developed in different communities and are not designed to coordinate with each other, not being integrated as a reliable system. This talk introduces open-source software for autonomous driving, which provides all necessary components integrated as a system. Research and development of autonomous driving can build on top of this software, using provided components or adding new components to enrich the system. To be the best of my knowledge, this is the first piece of work on open solutions for autonomous driving.

17:10-17:30 Break

17:30-18:00 Special Invited Session
Co-chairs: Makoto Ikeda (Univ. of Tokyo), Fumio Arakawa (Nagoya Univ.)

17:30-17:45 Conquering the challenges of wireless small cell base stations; TI's TC16630K2L System on Chip (SoC) achieves low power footprint without sacrificing performance and flexibility
Abhijeet Chachad (TI, USA)

17:45-18:00 Functional Safety considerations for an ADAS System on chip
Zoran Nikolic, Rama Venkatasubramanian, Aish Dubey, Rahul Gulati, Neil Simpson (TI, USA)

April 14, 2015 Main Hall(7th Floor), Yokohama Joho Bunka Center
(Yokohama Media & Communications Center)

9:30-9:50 Session I

9:30-9:50 **Welcome and Opening Remarks**

Chair: Hiroyuki Igura (NEC)

Hiroaki Kobayashi, Chair of the Organizing Committee
Yoshiaki Nakano, President of IEICE Electronics Society
Yuichirou Imatomi Director of the Economic Affairs Bureau, City of Yokohama

9:50-10:40 Session II

9:50-10:40 **Keynote Presentation 1**

Co-chairs: Hiroyuki Igura (NEC), Masato Suzuki (Socionext)

Advancing Moore's Law: Opening New Horizons

Tsuyoshi Abe (Intel, Japan)

Abstract: 2015 is the 50th anniversary of Moore's Law. Since 1965, the semiconductor industry has been thriving on Moore's Law, enabling new semiconductor devices with higher functionality and complexity, while controlling power, cost, and size. The Internet-of-Things (IoT) era is right at our doorstep and industry experts predict that by the year 2020, there will be well over 50 billion devices connected to the Internet. These IoT devices together with cloud computing will continue drive demand for high performance semiconductors with low power consumption. As a result, Moore's Law will play a pivotal role in controlling the cost of semiconductors. This presentation will focus on Intel's 14nm process technology featuring the 2nd generation tri-gate transistors. Specifically, it will illustrate the advantages of the new process technology over the previous generation and how Moore's Law will bring benefit to the development of leading edge semiconductors by reducing cost per transistor and improving performance/watt.

10:40-10:50 Break

10:50-11:40 Session III

10:50-11:40 **Keynote Presentation 2**

Co-chairs: Ryusuke Egawa (Tohoku Univ.), Yukinori Sato (Tokyo Institute of Tech.)

Data Centric Systems: Architecture and Solutions for Technical Computing, Big Data, and High Performance Analytics

Michael Rosenfield (IBM Research Division, USA)

Abstract: Computing systems will need to evolve in two fundamental ways: they must target solution driven workflows and they must be designed in such a way as to explicitly accommodate the impact of big data and complex analytics. The system requirements of classic modeling and simulation (HPC or technical computing) will converge with those of big data and analytics. As an example, HPC systems will need to be optimized to perform well on modeling and simulation; but, also must focus on other important elements of the overall workflow which include data management and manipulation coupled with associated analytics. Traditional machine balance points are no longer sustainable, nor achievable, with standard approaches. At a macro level, workflows will take advantage of different elements

of the systems hierarchy, at dynamically varying times, in different ways and with different data distributions throughout the hierarchy. This leads us to a data centric design point that has the flexibility to handle these data-driven demands. This flexibility will come from balanced & composable systems built from modular components with computation distributed to all elements of the system hierarchy. Data Centric Systems (DCS) focus on the problem of data location, and the principle that moving computing to the data will lead to more cost effective and efficient systems than prior generation systems. DCS systems, characterized by heterogeneous hardware, will provide leadership capabilities for Big Data, complex analytics, modeling/simulation and cognitive computing. DCS system software will allow the hardware to be used efficiently. Fully exploiting heterogeneous high performance capabilities will require additional evolution and innovation in programming models. A central motivator for DCS is to ensure the attributes of the architecture and implementation lead to commercially viable Exascale-class systems. This means that investments in programming models, languages and software development will be preserved for the future and that new optimized code will be positioned to take advantage of Exascale features.

11:40-11:50 Break

11:50-12:20 Session IV: Poster Short Speeches

Chair: Koji Hashimoto (Fukuoka University)

- Poster 1 **Implementation and Evaluation of an Efficient Parallel Architecture for Matrix Calculations**
Yuki Murakami, Naohito Nakasato, Stanislav Sedukhin (Univ. of Aizu)
- Poster 2 **A Scalable Processor for Real-time Sound Rendering**
Tan Yiyu¹, Yasushi Inoguchi¹, Yukinori Sato¹, Yukio Iwaya², Makoto Otani³, Takao Tsuchiya⁴ (¹JAIST, ²Tohoku Gakuin Univ., ³Kyoto Univ., ⁴Doshisha Univ.)
- Poster 3 **A Capacity Sharing Mechanism for Energy Savings in MCP**
An Hsia, Jui-Hsiang Wang, Ching-Wen Chen (Feng Chia Univ., Taiwan)
- Poster 4 **Design Guide and Incremental Synthesis Technique for Low Leakage 22nm Standard Cell Libraries**
Tsung-YiWu¹, Hsin-Hui Li², Zhi-Yao Ding¹, Guan-Cheng Guo¹ (¹National Changhua Univ. of Education, ²Global Unichip, Taiwan)
- Poster 5 **Off-loading LET generator in PEACH2 : A Switching Hub for High Performance GPU Clusters**
Chiharu Tsuruta¹, Yohei Miki², Takuya Kuhara¹, Takuji Mitsuishi¹, Naru Sugimoto¹, Hideharu Amano¹ (¹Keio Univ., ²Univ. of Tsukuba)
- Poster 6 **A Configuration Memory Reduced Programmable Logic Cell**
Motoki Amagasaki, Qian Zhao, Masahiro Iida, Morihiko Kuga, Toshinori Sueyoshi (Kumamoto Univ.)
- Poster 7 **Low Leakage Router Design Using Vertical Gating Technique**
Yu-Cheng Cheng, Jin-Hao Chen, Tung-Chi Wu, Yen-Jen Chang (National Chung Hsing Univ.)
- Poster 8 **Analysis of Linpack Benchmark for Manycore Processors with Native Mode**
Naoto Fukumoto, Kohta Nakashima (Fujitsu Labs.)
- Poster 9 **Power Reduction Method for L1 Cache by Focusing on Remaining Time due to DVFS**

Masahiro Kondo, Hiroya Ochiai Ryotaro Kobayashi (Toyohashi Univ. of Tech.)

- Poster 10 **Improvement of Data Utilization Efficiency for Cache Memory by Compressing Frequent Bit Strings**
Takuro Yoshida¹, Hiroya Ochiai¹, Ryotaro Kobayashi¹, Hajime Shimada²
(¹Toyohashi Univ. of Tech., ²Nagoya Univ.)
- Poster 11 **3D Bus Architecture using Inductive Coupling ThruChip-Interface**
Akio Nomura, Yu Fujita, Hiroki Matsutani, Hideharu Amano (Keio Univ.)
- Poster 12 **Ultra Low Power Reconfigurable Accelerator CMA-SOTB-2**
Koichiro Masuyama, Yu Fujita, Hayate Okuhara, Hideharu Amano (Keio Univ.)
- Poster 13 **A Chopper Amplifier Utilizing Optimized Chopping Frequency for EEG Signal Processing Devices**
Yuki Kawamata, Daisuke Kanemoto, Makoto Ohki (Univ. of Yamanashi)
- Poster 14 **Staggered Stacking : Connecting Many Small Chips Using ThruChip Interface**
Hiroshi Nakahara¹, Tomoya Ozaki¹, Hiroki Matsutani¹, Michihiro Koibuchi², Hideharu Amano¹ (*¹Keio Univ., ²National Institute of Informatics*)
- Poster 15 **Reconfigurable Multiplier Architecture based on Memristor-CMOS Circuits**
Seungbum Baek, Byung-Suk Park, Kyoungrok Cho (Chungbuk National Univ., Korea)
- Poster 16 **Reference Picture Buffer Memory Architecture for 4K HEVC Encoders**
Yukikuni Nishida, Takayuki Onishi, Hiroe Iwasaki, Mitsuo Ikeda, Atsushi Shimizu (NTT Media Intelligence Labs.)
- Poster 17 **Parallel Processing of Graph Search by Tightly Coupled Accelerator**
Takahiro Kaneda¹, Takuji Mitsuishi¹, Yuki Katsuta¹, Takuya Kuhara¹, Toshihiro Hanawa², Hideharu Amano¹, Taisuke Boku³ (*¹Keio Univ., ²Univ. of Tokyo, ³Univ. of Tsukuba*)
- Poster 18 **A 3rd Order Delta-Sigma Modulator Utilizing Equivalent Gain of Internal Multi-bit ADC**
Daisuke Kanemoto, Naoki Fukasawa, Takahide Sato, Makoto Ohki (Univ. of Yamanashi)
- Poster 19 **Stubborn Cache: A Novel Strategy for Repeating Thrashing Access Patterns**
Hayato Nomura, Takuma Nakajima, Masato Yoshimi, Tsutomu Yoshinaga, Hidetsugu Irie (Univ. of Electro-Communications)
- Poster 20 **Measurement-based Evaluation of NVM-added Systems**
Jisun Kim, Hyokyung Bahn (Ewha Univ., Korea)
- Poster 21 **Area Efficient Memory Bandwidth Compressor for Stream Computation on FPGAs**
Tomohiro Ueno, Kentaro Sano, Satoru Yamamoto (Tohoku Univ.)
- Poster 22 **XStencil: a 7.1GFLOPS/W 16-Core Coprocessor with a Ring Structure for Stencil Applications**
Jun Yao^{1,2}, Yasuhiko Nakashima², Kazutoshi Kobayashi³, Makoto Ikeda⁴, Wei Xue⁵, Tomohiro Fujiwara², Ryo Shimizu², Masakazu Tanomoto², Yangtong Xu⁵, Xinliang Wang⁵, Weimin Zheng⁵ (*¹Huawei Technologies, China, ²Nara Institute of Science and Tech., ³Kyoto Institute of Tech., ⁴Univ. of Tokyo, ⁵Tsinghua Univ., China*)

Poster 23 **A Cost Effective FPGA Implementation of Robust Circle Estimation based on RANSAC Algorithm**
Theint Theint Thu, Jimpei Hamamura, Yuichiro Shibata, Kiyoshi Oguri (Nagasaki Univ.)

Poster 24 **Energy-Efficient Event-Driven Microcontroller with Atomic Event Quantization Unit for IoT Sensing**
Sanghyun Lee¹, Daejin Park², Joonhyun Ahn², Jeonghun Cho², Tag Gon Kim¹
(¹KAIST, ²Kyungpook National Univ., Korea)

Poster 25 **A Parameterized Many Core Simulator for Design Space Exploration**
Shohei Takeuchi, Thi Hong Tran, Shinya Takamaeda, Yasuhiko Nakashima (Nara Institute of Science and Tech.)

12:20-13:20 Lunch Time Break

13:20-13:40 Poster Open: 7th floor poster show room

13:40-14:20 Session V

13:40-14:20 **Invited Presentation 1**
Co-chairs: Yuki Kobayashi (NEC), Yasuo Unekawa (Toshiba)

Acceleration Methods of Accurate Ego-Motion Using an Image Recognition Hardware for Advanced Driver Assistance Systems
Motoki Kimura (Renesas Electronics)

Abstract: An accurate ego-motion estimation algorithm based on optical flow and stereo matching has been implemented on R-Car H2 SoC which has eight CPU cores and an image recognition hardware for ADAS (Advanced Drivers Assistance Systems) applications. The image recognition hardware is constituted from a cluster of sixteen programmable floating-point based processors and four dedicated cores for image operations, which are tightly connected through internal bus and SRAM. These two types of cores have been implemented carefully in R-Car H2 based on the analysis of the open source based computer vision libraries frequently used for prototyping and application development, so as to offer enough programmability and processing capability for embedded image recognition applications. In this talk, we will introduce the architectures of these cores in the image recognition processor, and show acceleration methods exploiting the capability of R-Car H2 SoC, in order to achieve the real-time operation of one of the most accurate ego-motion estimation algorithms in the world.

14:20-15:00 Session VI

14:20-15:00 **Invited Presentation 2**
Co-chairs: Yasuo Unekawa (Toshiba), Yuki Kobayashi (NEC)

Heterogeneous Multi-Core SoC for ADAS and Image Recognition Applications
Takashi Miyamori (Toshiba)

Abstract: In recent years, image recognition technologies have become into practical on embedded systems for automotive, digital-consumer and mobile products. For automotive applications, they are the key technologies for Advanced Driving Assistance System (ADAS) and will lead us to safer car society. We have developed heterogeneous multi-core SoCs for ADAS and other image recognition applications. Because these applications require tremendous computing power, there are big challenges to achieve such a high performance with low power consumption.

Furthermore, high accuracy recognition is also required. We proposed a heterogeneous multicore architecture that consists of energy efficient VLIW processor cores with a SIMD coprocessor and hardware accelerators. Novel image features, such as CoHOG (Co-occurrence Histograms of Oriented Gradients) and color-based features are introduced and dedicated hardware accelerators have been developed for them. In this presentation, we will introduce architectures of these image recognition SoCs. Our latest SoC is composed with two 4-core processor clusters and 14 hard-wired accelerators and achieves 1.9TOPS as its peak.

- 15:00-16:00** **Break (Poster Open: 7th floor poster show room)**
- 16:00-17:05** **Session VII: Object recognition techniques**
Co-chairs: Sugako Otani (Renesas Electronics), Hiroyuki Takizawa (Tohoku Univ.)
- 16:00-16:25 **A Keypoint-level Parallel Pipelined Object Recognition Processor with Gaze Activation Image Sensor for Mobile Smart Glasses System**
Injoon Hong, Dongjoo Shin, Youchang Kim, Kyeongryeol Bong, Seongwook Park, Kyuho Lee, Hoi-Jun Yoo (KAIST, Korea)
- 16:25-16:50 **An Energy Efficient Hybrid FPGA-GPU based Embedded Platform to Accelerate Face Recognition Application**
Santhosh Kumar Rethinagiri, Oscar Palomar, Javier Arias Moreno, Osman Unsal, Adrian Cristal (BSC-Microsoft Research Centre, Spain)
- 16:50-17:05 **Lowering the Complexity of k-means Clustering by BFS-dijkstra Method for Graph Computing**
Anna Zhang¹, Jun Yao^{1,2}, Yasuhiko Nakashima¹ (¹NAIST, ²Huawei Technologies)
- 17:05-17:15** **Break**
- 17:15-18:45** **Session VIII: Panel Discussions**
- Topics: Computing Technology for Autonomous Driving**
Organizer & Moderator: Shinpei Kato (Nagoya Univ.)
Panelists: Takashi Miyamori (Toshiba)
Tadashi Kamada (Denso)
Tsuguo Nobe (Intel)
Mandali Khalesi (HERE)

April 15, 2015 Main Hall(7th Floor), Yokohama Joho Bunka Center
(Yokohama Media & Communications Center)

9:30-10:20 Session IX

9:30-10:20 **Keynote Presentation 3**

Co-chairs: Kunio Uchiyama (Hitachi), Akihiko Hashiguchi (Sony)

How can Medical Electronics Revolutionise Health Care by 2050?

Rudy Lauwereins (IMEC, Belgium)

Abstract: Today's medical practice lacks prevention, is slow and expensive, and often treats symptoms instead of root causes. In this visionary presentation, I will predict how medicine could look like in 2050, enabled by modern electronics. I will present four dream scenarios of a healthier world: a world where illness is prevented, a world without cancer, a world where personalised spare parts are produced, and a world without neurological and psychiatric disorders. Based on advanced prototypes of today, I will motivate why these visionary scenarios might become a reality. And I will leave the audience behind with a massive amount of ethical questions.

10:20-11:10 Session X

10:20-11:10 **Keynote Presentation 4**

Co-chairs: Kunio Uchiyama (Hitachi), Akihiko Hashiguchi (Sony)

Low Power and High Speed Working Memory with Spintronics and Vertical MOSFET Technology

Tetsuo Endoh (Tohoku Univ.)

Abstract: Recently in semiconductor memories such as working memories (SRAM, DRAM) and storage memories (NAND memory), it is becoming difficult to meet the target performance only by scaling technologies. Especially for 1X nm high speed working memories and beyond, the large power consumption brings more serious issues due to rapidly increase memory capacity, operation speed and leakage current of scaled CMOS. Moreover, the speed gap between each memory levels in addition to the speed gap between the operation speed of MPUs and that of working memories have expanded year by year. In this invited talk, it is discussed about the directionality of the semiconductor memory hierarchy structure in the future from the background mentioned above. It is introduced that with using 3D stacked memories based on Vertical MOSFETs and STT-MRAMs, the current issues of cell density, speed gap and power consumption will be simultaneously overcome, and novel memory hierarchy structure will be achieved. In addition, from the viewpoint of future high-end memory system, the impact of memory technologies hybridized with Vertical MOSFETs and spintronics devices such as MTJs is discussed. Finally, nonvolatile logic as one of application of STT-MRAM is shown.

11:10-11:20 Break (Poster Open: 7th floor poster show room)

11:20-12:05 Session XI: Low Power Circuits

Chair: Hajime Shimada (Nogoya Univ.)

11:20-11:35 **Fined-Grained Body Biasing for Frequency Scaling in Advanced SOI Processes**

Johannes Maximilian Kühn¹, Hideharu Amano², Oliver Bringmann¹, Wolfgang Rosenstiel¹ (¹Univ. of Tübingen, Germany, ²Keio Univ.)

11:35-11:50 A Leakage Current Monitor Circuit Using Silicon on Thin BOX MOSFET for Dynamic Back Gate Bias Control

Hayate Okuhara¹, Kimiyoshi Usami², Hideharu Amano¹ (¹Keio Univ., ²Shibaura Institute of Tech.)

12:05-12:45 Session XII: Low Power Processors

Co-chairs: Yuetsu Kodama (Tsukuba Univ./Riken), Kotaro Shimamura (Hitachi)

12:05-12:20 Power Management on 14 nm Intel® Core™ M processor

Anant Deval, Avinash Ananthakrishnan, Craig Forbell (Intel, USA)

12:20-12:45 0.39-V, 18.26-μW/MHz SOTB CMOS Microcontroller with Embedded Atom-Switch ROM

Toshitsugu Sakamoto¹, Yukihide Tsuji¹, Munehiro Tada¹, Hideki Makiyama¹, Takumi Hasegawa¹, Yoshiki Yamamoto¹, Shinobu Okanishi¹, Keiichi Maekawa¹, Nanoki Banno¹, Makoto Miyamura¹, Koichiro Okamoto¹, Noriyuki Iguchi¹, Yasuhiro Ogasahara², Hidekazu Oda¹, Shiro Kamohara¹, Yasushi Yamagata¹, Nobuyuki Sugii¹, Hiromitsu Hada¹ (¹LEAP, ²AIST)

12:45-13:50 Lunch Time Break

13:50-14:40 Session XIII

13:50-14:30 Keynote Presentation 5

Co-chairs: Fumio Arakawa (Nagoya Univ.), Koyo Nitta (NTT Electronics)

Riding the Perfect Storm, Bringing Mobile Compute to the Data Centre

John Goodacre (ARM / Univ. of Manchester, UK)

Abstract: EUROSERVER is a European commission FP7 funded project which is combining the technology trends of nanotechnology 3D integration, low-power mobile SoC processor integration and the impossible requirements from next generation cloud and high performance compute to investigate and build a solution for scalable, cost effective and flexible ARM-based server system architecture suitable across multiple markets. This talk will introduce the vision and the goals for the project and the approach the consortium is taking to realize a ground breaking solution out of this perfect storm.

14:40-15:00 Break (Poster Open: 7th floor poster show room)

15:00-16:00 Session XIV: System and Architecture Level Low Power Techniques

Co-chair: Kyoung-Rok Cho (Chungbuk National Univ.), Masaki Gondo (eSOL)

15:00-15:15 OS-less Dynamic Binary Instrumentation for Embedded Firmware

JinSeok Oh, Sungyu Kim, Eunji Jeong, Soo-Mook Moon (Seoul National Univ., Korea)

15:15-15:30 An Energy-Efficient Dynamic Memory Address Mapping Mechanism

Masayuki Sato, Chengguang Han, Kazuhiko Komatsu, Ryusuke Egawa, Hiroyuki Takizawa, Hiroaki Kobayashi (Tohoku Univ.)

15:30-15:45 Electronic Paper Display Update Scheduler for Extremely Low Power Non-volatile Embedded Systems

Yusuke Shiota, Shiyo Yoshimura, Tatsunori Kanai (Toshiba)

- 15:45-16:00 **A Novel Energy-efficient Data Acquisition Method for Wearable Devices**
Akira Takeda, Akira Yokosawa, Shintaro Sano, Shunsuke Sasaki, Takeshi Kodaka, Takahiro Tokuyoshi, Toshiki Kizu (Toshiba)
- 16:00-16:10 **Break (Poster Open: 7th floor poster show room)**
- 16:10-16:55 **Session XV: Processor Cores and NoC**
Co-chairs: Jun Yao (Huawei Tech.), Hidetoshi Matsumura (Fujitsu Labs.)
- 16:10-16:25 **MIAOW - An Open Source RTL Implementation of a GPGPU**
Raghuraman Balasubramanian, Vinay Gangadhar, Ziliang Guo, Chen-Han Ho, Cherin Joseph, Jaikrishnan Menon Mario, Paulo Drumond, Robin Paul, Sharath Prasad, Pradip Valathol, Karthikeyan Sankaralingam (Univ. of Wisconsin-Madison, USA)
- 16:25-16:40 **An Energy-Efficient FPGA-based Soft-Core Processor with Configurable Word Size ECC Arithmetic Accelerator**
Aiko Iwasaki, Yuichiro Shibata, Kiyoshi Oguri, Ryuichi Harasawa (Nagasaki Univ.)
- 16:40-16:55 **TURO: A Lightweight TUrN-Guided ROuting Scheme for 3D NoCs**
Jun Zhou, Huawei Li, Tiancheng Wang, Ying Wang, Xiaowei Li (Chinese Academy of Sciences, China)
- 16:55-17:05 **Break**
- 17:05-17:55 **Session XVI**
- 17:05-17:55 **Keynote Presentation 6**
Co-chairs: Hideharu Amano (Keio Univ.), Makoto Ikeda (Univ. of Tokyo)

The Kalray MPPA Mission-Critical Supercomputer on a Chip

Benoît Dupont de Dinechin (Kalray)

Abstract: The Kalray MPPA-256 processors implement a supercomputer architecture on chips manufactured with 28nm CMOS technology. These processors achieve high performance, low power, and timing dependability, by distributing a total of 256 user cores and 32 system cores across 16 compute clusters of 16+1 cores, and 4 I/O subsystems with a quad-core. All cores implement the same a 5-issue VLIW architecture, which is optimized for MCU as well as DSP type of computing. Each compute cluster and I/O subsystem owns a private address space, while communication and synchronization between them is ensured by data and control Networks-On-Chip (NoC). The MPPA-256 processors are also fitted with high-performance I/O interfaces, in particular DDR3, PCI Gen3, Ethernet 10G, Interlaken and GPIO. The MPPA-256 processors programming environment includes GGC-based C, C++, FORTRAN compilers with OpenMP support, and stand-alone or Eclipse GDB-based debuggers. The local memory spaces are either visible from the applications, or transparently turned into a cache of the external DDR memory thanks to a distributed shared memory run-time system. This allows to support the global memory of the OpenCL data-parallel and task-parallel programming models. Current uses of the MPPA-256 processors include data center CPU acceleration such as real-time H.265 video encoding, Monte Carlo type of simulations, or cryptanalysis. The CPU accelerators are based on PCI-e card with 4x MPPA-256 processors for a total of 1024 cores. Advanced uses of the MPPA-256 processors are in mission-critical applications such as avionics functions, airborne image analysis, and automotive ADAS (Advanced Driver Assistance

Systems) functions.

17:55-18:35 Session XVII

17:55-18:35 Invited Presentation 3

Co-chairs: Hideharu Amano (Keio Univ.), Makoto Ikeda (Univ. of Tokyo)

ExaScaler-1: The Power-Efficient Submersion Many-Core Processor Based Supercomputer

Sunao Torii (ExaScaler / PEZY Computing)

Abstract: We have developed the proprietary many-core processor based supercomputer, “ExaScaler-1”. “Suiren”, the first installed system of ExaScaler-1 consisting of four 8U tanks, totaling 32 unit (32U) systems, achieves over 190TFlops HPL (High Performance Linpack benchmark) performance and 4.95GFlops/W power efficiency respectively. These values are ranked at 369th in the TOP500 list and at 2nd in the Green500 list of November 2014. Suiren is the only supercomputer ranked both in the TOP500 list and in the Green500 list utilizing a general-purpose many-core accelerator, which is developed by a venture company. ExaScaler-1 adopts PEZY-SC many-core processor as a calculation accelerating device. PEZY-SC integrates 1,024 of MIMD (Multiple Instruction-stream Multiple Data-stream) processing elements (PEs) on a chip. It also integrates 8-ch 64bit 2.4GHz DDR3/4 SDRAM interfaces and 34MB on-chip cache and scratch pad memory enough for memory bandwidth requirements. PEZY-SC consumes around 80W dynamic and 10W leak power consumption. Each unit of ExaScaler-1 combines two Intel Xeon E5-2660v2 processors and 8 PEZY-SC chips with 512GB SDRAM and it consumes around 1.3KW. To maintain low-temperature for the whole system and to minimize a tank machine’s installation space, we have developed a new submersion liquid cooling system from the scratch. Since it soaks whole mother-board totally in coolant, we can keep low enough temperature even for small parts such as power supply and memory modules as well as a highly power consuming device like CPU. It realizes not only reducing the chip leak current but also increasing the system reliability. Since this system adopts not popular 2-phase cooling, but 1-phase thermal conduction cooling with open roof top tank, it enables to reduce both maintenance cost and manufacturing expenditures for cooling tanks. In my talk, I will present the ExaScaler-1’s cooling system as well as PEZY-SC many-core chip architecture. Furthermore, I will explain our further development plan for the new era of Exa-scale computing system.

18:35-18:55 Poster Award and Closing Remark

Makoto Ikeda, Program Committee Co-chair (Univ. of Tokyo)