Final Program

April 20, 2022 (Japan Standard Time)

9:00-10:30 **Special Invited Lecture 1**
Chair: Takatsugu Ono (Kyushu Univ.)

9:00-10:30 **Closing the Gap between Quantum Algorithms and Machines with Hardware-Software Co-Design**
Fred Chong (Univ. of Chicago and Super.tech, USA)

Abstract: Quantum computing is at an inflection point, where 127-qubit (quantum bit) machines are deployed, and 1000-qubit machines are perhaps only a few years away. These machines have the potential to fundamentally change our concept of what is computable and demonstrate practical applications in areas such as quantum chemistry, optimization, and quantum simulation. Yet a significant resource gap remains between practical quantum algorithms and real machines. A promising approach to closing this gap is to design architectural interfaces that selectively expose to programming languages and compilers some of the key physical properties of emerging quantum technologies. I will describe some of our recent work that focuses on techniques that break traditional abstractions, including compiling programs directly to analog control pulses, adapting programs for machine variations, computing with ternary quantum bits, efficient error correction using 2.5D structures, and designing just enough tunability in qubits to avoid crosstalk.

10:30-11:00 **Break**

11:00-12:30 **Special Invited Lecture 2**
Chair: Takatsugu Ono (Kyushu Univ.)

11:00-12:30 **Computer Architecture Challenges for the Security of Persistent Memory**
Yan Solihin (Univ. of Central Florida, USA)

Abstract: Non-volatile or persistent memory (PM) is increasingly integrated into the main memory of computer systems. PM allows programmers to keep persistent data in data structures in memory instead of files in storage. Such structures can be wrapped into containers that we refer to as Persistent Memory Objects (PMO). While PMO allows fine-grain access at low latency to persistent data, it also presents several challenges. Among these challenges are: (1) new abstraction is needed to define its use, access, and sharing, (2) security vulnerabilities that arise from keeping them in memory instead of files, and (3) how it affects memory encryption and integrity verification. In this talk, I will give a broad landscape of computer architecture challenges with PM and discuss our approaches in solving these problems.

12:30-13:20 **Lunch Time Break**

13:20-14:10 **Special Session I**

13:20-14:10 **Keynote Presentation 1**
Co-chairs: Tohru Ishihara (Nagoya Univ.), Takumi Uezono (Hitachi)

Universal Chiplet Interconnect Express (UCIe): Poised to Change the Compute Landscape
Debendra Das Sharma (Intel, USA)

Abstract: High-performance workloads demand on-package integration of heterogeneous processing units, on-package memory, and communication infrastructure to meet the demands of the emerging compute landscape. Applications such as artificial intelligence, machine learning, data analytics, 5G, automotive, and high-performance computing are driving these demands to meet the needs of cloud computing, intelligent edge, and client computing infrastructure. On-package interconnects are a critical component to deliver the power-efficient performance with the right feature set in this evolving landscape. Universal Chiplet Interconnect Express (UCIe), is an open industry standard with a fully specified stack that comprehends plug-and-play interoperability of chiplets on a package; similar to the seamless interoperability on board with well-established and successful off-package interconnect standards such PCI Express®, Universal Serial Bus (USB)®, and Compute Express Link (CXL)®. In this talk, we will discuss the usages and key metrics associated with different technology choices in UCIe. We will also delve into the different layers as well as the software model associated with UCIe along with the compliance and interoperability mechanisms. We will also discuss how this open standard could potentially evolve to incorporate additional usage models in the future.

14:10-14:30 **Break**

14:30-15:20 **Special Session II**

14:30-15:20 **Keynote Presentation 2**
Co-chairs: Akihiro Hashiguchi (Sony), Teruaki Sakata (Hitachi)

RISC-V-based Parallel Processor IP with Vector Extension for Embedded Systems
Shotaro Shintani (NSITEXE)

Abstract: There are a variety of controls in automobiles, and the amount of processing is continuously on the increase due to the more complex algorithms. Moreover, the progress of autonomous driving has been remarkable in recent years, and the performance requirements have increased dramatically in addition to the diversification of processing. There are also traditional requirements in automotive systems, such as hard real-time performance, functional safety, power efficiency, software portability, and so on. On the other hand, there is the problem that software development will become more complex and less portable when several processors or accelerators are adopted to meet the automotive requirements. In order to solve it, we develop new processors with enough performance and flexibility in a single architecture that can handle the various types of processing for autonomous driving while also satisfying the traditional automotive requirements such as functional safety. In this presentation, our first-generation Data Flow Processor (DFP), DR1000C, is introduced, and it is targeted at safety-critical systems and delivers the real-time performance and high-throughput data processing required for microcontrollers for embedded systems. DR1000C is a RISC-V-based multiple instruction stream, multiple data stream (MIMD) processor with a vector processing unit. It also supports a hardware multi-threading mechanism that allows up to 16 threads which simultaneously share the vector processing unit, resulting in highly efficient resource use. Furthermore, DR1000C has the necessary functional safety modules and meets ISO 26262 ASIL B to D safety requirements without additional external special safety mechanisms. Our evaluation results show that DR1000C can

efficiently process control algorithms such as Model Predictive Control (MPC). DR1000C is also the world's first RISC-V processor with vector extensions which has achieved ISO 26262 ASIL D Ready certification. Lastly, our product roadmap and portfolio are referred to for addressing the area of autonomous driving, which is the target of the next-generation DFP.

15:20-15:30 **Break**

15:30-16:20 **Special Session III**

15:30-16:20 **Keynote Presentation 3**
Co-chairs: Yukinori Sato (Toyohashi Univ. of Tech.), Makoto Ikeda (Univ. of Tokyo)

NanoBridge-based FPGA for Space Applications
Makoto Miyamura (NanoBridge Semiconductor)

Abstract: We demonstrate FPGA based on NanoBridge, a novel resistive-change switch. NanoBridge, which is integrated in the back end of line (BEOL), features a high ON/OFF conductance ratio, weak temperature dependence of its resistance, non-volatility, and endurance against soft errors. In place of SRAM and a pass transistor, NanoBridge is utilized as a configuration switch in the FPGA. In this presentation, we will report the NanoBridge-based FPGA (NBFPGA) for applications in low-power and harsh environments, such as space application. Under JAXA's innovative satellite technology demonstration program, NBFPGA was mounted on a camera module and installed in RAPIS-1 (RAPid Innovative payload demonstration Satelite-1). Successful one-year operation in space, soft error reliability of Nanobridge device, and low-power operation will also be reported.

16:20-17:20 **Poster Short Speeches**
Chair: Koji Hashimoto (Fukuoka Univ.)

Poster 1 **Simulation Platform of Computation-in-Memory with Memory Device Non idealities for DNN**
Kazuhide Higuchi, Chihiro Matsui, Naoko Misawa, Ken Takeuchi (Univ. of Tokyo)

Poster 2 **Study of General-Purpose CGRA Architecture**
Shigeyuki Takano, Hideharu Amano (Keio Univ.)

Poster 3 **Compact and Tunable ReRAM Computation-in-Memory for Log-encoding Simulated Annealing**
Naoko Misawa, Chihiro Matsui, Ken Takeuchi (Univ. of Tokyo)

Poster 4 **Fusion of Multiple Core and Just-in-Time Compilable CGRA**
Tomoya Akabe, Hidenari Inamasu, Renyuan Zhang, Yasuhiko Nakashima (Nara Institute of Science and Tech.)

Poster 5 **LLVM Instruction Latency Measurement for Software-Hardware Interface for Multi/many-core**
Hiro Mikami, Kei Torigoe, Makoto Inokawa, Masato Edahiro (Nagoya Univ.)

Poster 6 **Towards Network-Wide Malicious Traffic Detection with Power-Effective Hardware NIDS Design**
Zhenguo Hu, Hirokazu Hasegawa, Yukiko Yamaguchi, Hajime Shimada (Nagoya Univ.)

Poster 7 **Just-In-Time Code Fusion of Unary Mathematical Calculation for Supercomputer Fugaku**

April 21, 2022 (Japan Standard Time)

9:00-9:10 **Session I**

9:00-9:10 **Welcome and Opening Remarks**
Co-chairs: Yuki Kobayashi (NEC), Hiroki Matsutani (Keio Univ.)

Makoto Ikeda	*Chair of the Organizing Committee*
Tadao Nakamura	*Chair of the Steering Committee*
Jose Renau	*Chair of IEEE/CS TCMM*
Minoru Fujishima	*Vice President and President-Elect of IEICE/ES*

9:10-10:00 **Session II**

9:10-10:00 **Keynote Presentation 4**
Co-chairs: Takuya Nakaike (IBM), Hiroki Matsutani (Keio Univ.)

The IBM Telum Enterprise-class Processor
Christian Jacobi (IBM, USA)

Abstract: The IBM Telum Processor powers scalable performance for enterprise workloads, embedded AI acceleration, and industry leading security and RAS. In this presentation, learn how Telum balances high performance and high frequency design, power efficiency, and reliability using a multi-faceted approach including micro-architecture, advanced power and thermal management and optimization of processor and system design. Telum-based systems will grow total workload capacity, provide new functionality like real-time embedded AI inferencing, and improve power efficiency compared to prior-generation IBM Z systems.

10:00-10:10 **Break**

10:10-11:40 **Session III: Panel Discussions**

Topics: "The Future of Mission-critical Mixed-criticality High-performance Embedded Systems"
Organizer and Moderator: Masaki Gondo (eSOL)
Panelists: Christian Jacobi (IBM, USA)
Timothy Mattson (Intel, USA)
Makoto Miyamura (NanoBridge Semiconductor)
Juanjo Noguera (Xilinx, USA)

11:40-11:50 **Break**

11:50-12:40 **Session IV**

11:50-12:40 **Keynote Presentation 5**
Co-chairs: Yuki Kobayashi (NEC), Yasuo Unekawa (Toshiba Electronic Devices & Storage)

Heterogeneity: A Play in Two Acts
Joseph Curley, Timothy Mattson (Intel, USA)

Abstract: Heterogeneous computing promises high performance with optimized

energy consumption. In Act 1 programmers write code with one API which our system maps onto any single processor from a range of processor-types. We'll show, for example, a single codebase moving between CPUs and GPUs. In Act 2 we look to the future where many heterogeneous processors work together to solve a single problem. We still focus on one API mapping onto many processor-types, but now we run parallel programs over distributed heterogeneous systems. Our name for this amazing system… oneAPI of course.

12:40-13:30 **Lunch Time Break**

13:30-14:20 **Session V: Power Management and Analysis**
Co-chairs: Yuetsu Kodama (Riken), Satoshi Imamura (Fujitsu)

13:30-13:55 **Body Bias Control on a CGRA based on Convex Optimization**
Takuya Kojima[1], Hayate Okuhara[2], Masaaki Kondo[3], Hideharu Amano[3] ([1]Univ. of Tokyo, [2]National Univ. of Singapore, [3]Keio Univ.)

13:55-14:20 **Power Analysis of Directly-connected FPGA Clusters**
Kensuke Iizuka, Haruna Takagi, Aika Kamei, Kazuei Hironaka, Hideharu Amano (Keio Univ.)

14:20-15:10 **Poster Break**

15:10-16:00 **Session VI**

15:10-16:00 **Keynote Presentation 6**
Co-chairs: Fumio Arakawa (Univ. of Tokyo), Masato Suzuki (Socionext)

Software-defined Architecture and Platforms - Automotive and beyond
Masaki Gondo (eSoL)

Abstract: The embedded system industries, notably the automotive but not limited to, are all talking about "Software-Defined." They typically mean "Software-Defined Architecture (SDA)" but the meaning is quite vague. In this talk, we start by reviewing the different use of the term in the industry. Then, using the next-generation automotive software platform based on AUTOSAR Adaptive as an example, we look into the background of SDA trends, and the design approach of such platforms that needs to tackle the challenge of developing highly-energy efficient, high-performance, safe, and yet affordable systems. Also, the business and market implications of such trends and approaches will be discussed briefly if the time allows.

16:00-16:10 **Break**

16:10-17:00 **Session VII: Machine Learning I - Circuit based**
Co-chairs: Shunsuke Sasaki (Toshiba), Renyuan Zhang (Nara Institute of Science and Tech.)

16:10-16:35 **A Memcapacitive Spiking Neural Network with Circuit Nonlinearity-aware Training**
Reon Oshio[1], Takuya Sugahara[1], Atsushi Sawada[1], Mutsumi Kimura[1,2], Renyuan Zhang[1], Yasuhiko Nakashima[1] ([1]Nara Institute of Science and Tech., [2]Ryukoku Univ.,)

16:35-17:00 **A 1036 TOp/s/W, 12.2 mW, 2.72 µJ/Inference All Digital TNN Accelerator in 22 nm FDX Technology for TinyML Applications**
Moritz Scherer[1], Alfio Di Mauro[1], Georg Rutishauser[1], Tim Fischer[1], Luca Benini[1,2]

([1]ETH Zurich, Switzerland, [2]Univ. of Bologna, Italy)

April 22, 2022 (Japan Standard Time)

9:00-9:40 **Session VIII**

9:00-9:40 **Invited Presentation**
Co-chairs: Ryusuke Egawa (Tokyo Denki Univ.), Daisuke Kobayashi (NTT)

AMD Ryzen 6000™ Series For Mobile
Jim Gibney (AMD, USA)

Abstract: AMD will present the AMD Ryzen™ 6000 Series processor for laptops, bringing the new "Zen 3+" core architecture together with AMD RDNA™ 2 architecture-based on-chip graphics. Fabricated using TSMC's 6nm process technology, the processor SoC delivers higher single and multi-threaded CPU performance than the last generation, new levels of built-in graphics performance, new platform features and long battery life. The updated AMD "Zen 3+" core is optimized to deliver high frequency and performance-per-watt. The SoC features up to eight high-performance cores, delivering 16 threads of processing, with clock speeds up to 5 GHz. These are the first notebook processors to feature RDNA™ 2 architecture-based built-in graphics, with performance up to 2X of the last generation. The SoC includes an all-new integrated display engine, allowing for ultra-high resolutions and refresh rates. New power management features include deep sleep states that save power and greatly extend battery life. The platform is entirely new with the SoC including support for DDR5 memory, PCIe® 4.0, USB4, and AI noise cancellation. This is the first x86 processor to fully support advanced Windows® 11 security features with the integrated Microsoft® Pluton security processor.

9:40-9:50 **Break**

9:50-10:40 **Session IX**

9:50-10:40 **Keynote Presentation 7**
Co-chairs: Yasutaka Wada (Meisei Univ.), Koyo Nitta (Univ. of Aizu)

Xilinx 7nm Edge Processors
Juanjo Noguera (Xilinx, USA)

Abstract: In this presentation, Xilinx will provide an overview of the 7nm Versal architecture, focusing on Edge applications. This talk with introduce the key target use-cases focusing on AI/ML inference applications at the edge, give a device-level Versal architecture overview and provide details on the second-generation AIE architecture (AIE-ML). We'll provide an overview of the programing abstractions for this products, and finally demonstrate some system-level performance/W comparison to different architectures.

10:40-11:30 **Session X: Machine Learning II - Applications**
Co-chairs: Hiroyuki Takizawa (Tohoku Univ.), Teppei Hirotsu (NSITEXE)

10:40-11:05 **Reinforcement Learning based Efficient Mapping of DNN Models onto Accelerators**
Shine Parekkadan Sunny, Satyajit Das (Indian Institute of Tech. Palakkad, India)

11:05-11:30 **A Low-power and Real-time 3D Object Recognition Processor with Dense RGB-D Data Acquisition in Mobile Platforms**
Dongseok Im, Gwangtae Park, Junha Ryu, Zhiyong Li, Sanghoon Kang, Donghyeon Han, Jinsu Lee, Wonhoon Park, Hankyul Kwon, Hoi-Jun Yoo (KAIST, Korea)

11:30-12:00 Poster Break

12:00-12:50 Session XI: Circuit and Systems
Co-chairs: Yasutaka Wada (Meisei Univ.), Ryuichi Sakamoto (Tokyo Institute of Tech.)

12:00-12:25 **Zero-standby-power Nonvolatile Standard Cell Memory Using FiCC for IoT Processors with Intermittent Operations**
Yuki Abe[1], Kazutoshi Kobayashi[1], Jun Shiomi[2], Hiroyuki Ochi[3] ([1]Kyoto Institute of Tech., [2]Osaka Univ., [3]Ritsumeikan Univ.)

12:25-12:50 **DXT501: An SDR-based Baseband MP-SoC for Multi-protocol Industrial Wireless Communication**
Yang Chen, Lin Liu, Xuelin Feng, Jinlin Shi (Univ. of Chinese Academy of Sciences, China)

12:50-14:20 Lunch Time Break

14:20-15:10 Session XII: Broadcasting-tech
Co-chairs: Kotaro Shimamura (Hitachi), Sugako Otani (Renesas Electronics)

14:20-14:45 **An Efficient Reference Image Sharing Method for the Parallel Video Encoding Architecture**
Ken Nakamura[1,2], Yuya Omori[1], Daisuke Kobayashi[1], Koyo Nitta[1], Kimikazu Sano[1], Masayuki Sato[2], Hiroe Iwasaki[2], Hiroaki Kobayashi[2] ([1]NTT, [2]Tohoku Univ.)

14:45-15:10 **Ultra-low Latency 8K Video-transmission System Utilizing Disaggregation Configuration**
Yasuhiro Mochida, Daisuke Shirai, Koichi Takasugi (NTT)

15:10-15:20 Break

15:20-16:10 Session XIII: Domain Specific Accelerator
Co-chairs: Yuichiro Shibata (Nagasaki Univ.), Ryohei Kobayashi (Univ. of Tsukuba)

15:20-15:45 **Hardware Acceleration of Aggregate Signature Generation and Authentication by BLS Signature over BLS12-381 Curve**
Kaoru Masada, Ryohei Nakayama, Makoto Ikeda (Univ. of Tokyo)

15:45-16:10 **Encoder-based Many-Pattern Matching on FPGAs**
Hoang-Gia Vu, Ngoc-Dai Bui (Le Quy Don Technical Univ., VietNam)

16:10-16:20 Poster Award and Closing Remarks
Program Committee Co-chairs: Yasutaka Wada (Meisei Univ.), Ryusuke Egawa (Tokyo Denki Univ.),