18th EMBEDDED SYSTEMS WEEK

7-14 October 2022 | Shanghai, China + Phoenix, US + Virtual

2022 ESWEEK PROGRAM

SPONSORING SOCIETIES:

THANK YOU TO OUR SPONSORING SOCIETIES

THANK YOU TO OUR PLATINUM SPONSORS

THANK YOU TO OUR GOLD SPONSORS

TABLE OF CONTENTS

WELCOME TO ESWEEK 2022

Aviral Shrivastava, General Chair
Arizona State University, US

Xiaobo Sharon Hu, General Co-Chair
University of Notre Dame, US

Embedded Systems Week (ESWEEK) is the premier event covering all aspects of hardware and software design for smart and connected computing systems. By bringing together three leading conferences (CASES, CODES+ISSS, EMSOFT), two symposia (NOCS and MEMOCODE), seven workshops, seven tutorials, two special sessions, an industry session, ten education classes, and ACM SIGBED SRC, ESWEEK allows attendees to benefit from a wide range of topics covering the state of the art in embedded systems research and development.

Due to the continuing travel restrictions, this 18th edition of ESWEEK is being organized in a hybrid form, with in-person events in Shanghai, China, Phoenix, USA, and online on Gathertown. We believe that this multi-site, distributed program will provide an unparalleled opportunity for the growing embedded systems community across the globe to come together, engage, interact, and celebrate the latest advances in embedded software, hardware, systems, and applications.

Highlights of the ESWEEK program this year include three distinguished keynote talks by prominent leaders in academia, industry, and funding agencies, covering relevant exciting trends for future embedded and cyber-physical systems and providing deep insights into technology drivers. Prof. Jei Li from Shanghai Jiao Tang University will address the challenges and issues in the application of big data, AI, and blockchain for Industrial Internet. Dr. Pete Wurman, Director of Sony AI, will describe how they trained AI agents for Gran Turismo that can compete with the world's best e-sports drivers. Finally, Prof. Margaret Martonosi from Princeton and the U.S. NSF will discuss key technical themes for the U.S. NSF CISE (Computer and Information Science) directorate, and how CISE is developing programmatic opportunities to advance research related to them. She will particularly note how ESWEEK topic areas relate to these technical priorities.

In order to further stimulate the ESWEEK attendees with more industry-side focused inspirational talks, ESWEEK 2022 will feature 2 SKY (short keynote) talks. Yu Huang from Huawei will present the industry perspective on AI *(Artificial Intelligence)* for EDA (Electronic Design Automation), and Tomas Evensen from AMD will talk about Open-Source Software Stacks for Heterogenous SoCs (System on Chips).

Continuing with the long-standing tradition of ESWEEK, the main program of ESWEEK 2022 will conclude with a vibrant panel discussion on "Wafer-scale Computing Systems: Are we there yet?"

Following the now established journal-integrated publication model for the three conferences (CASES, CODES+ISSS, and EMSOFT), all regular papers presented are published in a Special Issue of the IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems. To this end, ESWEEK-related journal submissions to IEEE TCAD followed a journal-style peer review process conducted in two stages with the opportunity of minor/major revisions before the final decision. In addition, the selected Work-in-Progress (WiP) track papers are published in conference-specific Proceedings.

The main program of ESWEEK on Monday, Tuesday, and Wednesday consists of two special sessions, an industry sessions, and 25 regular sessions from the three

WELCOME TO ESWEEK 2022

conferences (CASES, CODES+ISSS, and EMSOFT). There is a strong emphasis on interaction in the virtual setting. All the accepted journal-track and WiP papers along with teaser videos of them are made available to the ESWEEK participants one week before the hybrid event. During the conference week, the live technical sessions will feature 15-minute talks for the journal-track papers, and 5 minutes for the WiP papers, including the Q&A. In addition, all journal-track and WiP papers will be accompanied by an interactive virtual poster session in which attendees can further interact with presenters and authors.

The main program is preceded by tutorials on the Friday before the conference week. The tutorials provide an excellent opportunity to get in-depth knowledge on both new trends and hot topics. This year, ESWEEK will host seven, 3-hour tutorials covering a wide scope of topics, including quantum control, quantum flow, cyber-physical systems, accelerated computing, in-memory computing, hardware security, and embedded machine learning.

On Saturday and Sunday preceding the main event, there will be the ESWEEK Education Classes, which are usually highly attended. ESWEEK 2022 is organizing ten, 2-hour virtual, but live classes on the theme of "Security, Privacy and Trust". These education classes will introduce topics and concepts that have become fundamental and pervasive but are yet not in the textbooks. These education classes enable students across the globe, especially the ones that do not have access to high-quality educational content, an opportunity to learn these topics in an engaging and hands-on way from the top educators in the field.

Thursday and Friday of the conference week are the days for the symposium and workshops. This year we welcome MEMOCODE (Symposium on Formal Methods and Models for System Design) as our ESWEEK's newest symposium as they increase the critical mass for the event. MEMOCODE will be organized in parallel to our longstanding partner NOCS (Networks on Chip). In addition, we have 7 workshops covering a range of important topics in embedded systems: RSP (Rapid System Prototyping), MSC (Memory and Storage Computing), CODAI (Compilers, Deployment, and Tooling for Edge AI), HEC (Heterogeneous Edge Computing for Embedded System), SECRISC-V (Secure RISC-V Architecture Design Exploration), EIC (Edge Intelligent Computing), and EEDA (Emerging Techniques in System Design and Design Automation for Embedded Systems).

To improve the accessibility of ESWEEK to the large embedded systems community throughout the world, ESWEEK 2022 has provided registration fee waivers, allowing attendees with demonstrated needs (e.g., students and researchers with limited financial funds or resources) to attend ESWEEK for free!

We are looking forward to seeing you at the inspiring, interesting, interactive, hybrid, and accessible ESWEEK 2022!

The organization of ESWEEK 2022 was only possible with the continuous support and help from the sponsors and many volunteers: The conference program chairs with and their program committee members, the organizers of the workshops, tutorials, and symposia, all members of the organization committee, and, finally, the virtual conference chair and the web chair---without their commitment and contributions this hybrid event would not exist.

We are looking forward to seeing you at the inspiring, interesting, interactive, hybrid, and accessible ESWEEK 2022!

ESWEEK 2022 OVERVIEW

Friday, October 7	Saturday, October 8	Sunday, October 9
Tutorials	**Education Classes**	**Education Classes**
Monday, October 10	**Tuesday, October 12**	**Wednesday, October 12**
ESWEEK Opening Keynote 1 – Jie Li Industry Sessions CASES, CODES+ISSS, EMSOFT, and Special Session Posters and Industry Booth Planned Networking (Shanghai and GatherTown)	Test of Time Awards Keynote 2 – Pete Wurman Sky Talk 1 - Yu Huang CASES, CODES+ISSS, EMSOFT, and Special Session Workshops Posters and Industry Booth Planned Networking (GatherTown)	Best Paper and Other Awards Keynote 3 – Margaret Martonosi Sky Talk 2 - Tomas Evensen CASES, CODES+ISSS, EMSOFT Posters and Industry Booth Planned Networking (GatherTown) Panel on Wafer-scale Computing
Thursday, October 13	**Friday, October 14**	
NOCS and MEMOCODE Symposia Workshops and Onsite Networking (Phoenix)	NOCS and MEMOCODE Symposia Workshop SIGBED Buisness Meeting	

BEST PAPER CANDIDATES

CASES

Horae: A Hybrid I/O Request Scheduling Technique for Near-Data Processing based SSD
Jiali Li, Xianzhang Chen, Duo Liu, Lin Li, Jiapin Wang, Zhaoyang Zeng, Yujuan Tan - *Chongqing University*
Lei Qiao - *Beijing Institute of Control Engineering*

SASCHA - Sparsity-Aware Stochastic Computing Hardware Architecture for Neural Network Acceleration
Wojciech Romaszkan, Tianmu Li, Puneet Gupta - *UCLA*

SWAP: A Server-Scale Communication-Aware Chiplet-Based Manycore PIM Accelerator
Harsh Sharma - *Washington State University*, Sumit K. Mandal - *University of Wisconsin-Madison*, Janardhan Rao Doppa - *Washington State University,* Umit K. Ogras - *University of Wisconsin – Madison*, Partha Pratim Pande - *Washington State University*

CODES+ISSS

NeuroMap: Efficient Task Mapping of Deep Neural Networks for Dynamic Thermal Management in High-Bandwidth Memory
Shailja Pandey , Preeti Ranjan Panda - *Indian Institute of Technology Delhi*

FLAM-PUF: A Response Feedback-based Lightweight Anti-Machine Learning-Attack PUF
Linjun Wu - *Hunan University,* Yupeng Hu - *Hunan University*, Kehuan Zhang - *The Chinese University of Hong Kong*, Wenjia Li - *New York Institute of Technology*, Xiaolin Xu - *Northeastern University*, Wanli Chang - *University of York*

Exploring Synchronous Page Fault Handling
Yin-Chiuan Chen - *National Taiwan University*, Chun-Feng Wu - *Harvard University*, Yuan-Hao Chang - *Academia Sinica*, Tei-Wei Kuo - *Academia Sinica & National Taiwan University*

EMSOFT

Verified Train Controllers for the Federal Railroad Administration Train Kinematics Model: Balancing Competing Brake and Track Forces
Aditi Kabra, Stefan Mitsch, André Platzer - *Carnegie Mellon University*

Tinkertoy: Build your own operating systems for IoT devices
Bingyao Wang, Margo Seltzer - *University of British Columbia*

Vulnerability Detection of ICS Protocols Via Cross-State Fuzzing
Feilong Zuo - *Tsinghua University*, Zhengxiong Luo - *Tsinghua University*, Junze Yu - *Tsinghua University*, Ting Chen - *University of Electronic Science and Technology of China*, Zichen Xu - *Nanchang University*, Aiguo Cui - *Huawei Technologies Co. Ltd.*, Yu Jiang - *Tsinghua University*

TUTORIALS SCHEDULE

FRIDAY, OCTOBER 7

TUTORIALS

CST - 10/7/2022 14:00-17:00 PST - 10/6/2022 23:00-02:00 EST - 10/7/2022 02:00-05:00 CET - 10/7/2022 08:00-11:00	**T1 – Quantum Control**	**T2 – Taming Delays in Cyber-Physical Systems**	**T4 – Integrating Compute Acceleration Into Embedded System Design Using Vitis**
CST - 10/7/2022 22:00-01:00 PST - 10/7/2022 07:00-10:00 EST - 10/7/2022 10:00-13:00 CET - 10/7/2022 16:00-19:00	**T6 – Manycore processing-in-memory systems for accelerating deep learning applications**	**T7 – Hardware Security and Trust Verification**	

Tutorial T1

Tutorials on Quantum Control

Abstract: The rapid development of quantum computers has entailed extensive research in quantum controlling technologies that manipulate quantum states. Equally important to quantum controlling hardware is the development of software, which greatly improves the efficiency of quantum experiments or the applications of quantum computers in solving real world problems. In this proposal, a half-day tutorial is proposed. It contains four sessions and covers the following topics: quantum programming languages, quantum compilation, quantum controlling software, and quantum computing based on superconducting qubits. The intended audience includes but not limited to experts and students in the fields of (1) quantum computing, (2) compiler and architecture optimization, (3) hardware/software co-design. Since quantum computing is a sub-area in EDA conferences, programming language, compilation, systematic software support for controlling superconducting quantum computing systems will fit the audience from ESWEEK.

Ways to deliver: The tutorials are given via four introductory lectures. Each lecture takes 45 minutes including the Q&A discussion with attendees.

The necessary software includes Quingo and Quantify, which can be installed based on the instructions on https://gitee.com/quingo/quingo-runtime (Quingo), https://quantify-quantify-core.readthedocs-hosted.com/en/0.5.3_a/installation.html (Quantify-core) and https://quantify-quantify-scheduler.readthedocs-hosted.com/en/0.7.0/installation.html (Quantify-scheduler).

List of topics and speakers

Topic 1: Superconducting quantum computing

Abstract: In this talk, I will show our recent progress with our collaborators on superconducting multiqubits system. We designed and fabricated several versions of quantum processors, including one integrated up to 66 qubits. The fidelities of single-bit gates and two-bit gates are calibrated by randomized benchmarking or parallel cross-entropy benchmarking. For single-qubit gates, the average error is ~0.14% and that of two-qubit gates is ~0.59%. I will also show some of the multi-qubits experiment results, e.g., genuine multiparticle entanglement for 12 superconducting qubits, quantum walks on a programmable two-dimensional 62-qubit superconducting processor, and strong quantum advantage.

Instructor: Xiaobo Zhu

Biography: Prof. Xiaobo Zhu received his PhD in condensed matter physics at Institute of Physics, *Chinese Academy of Sciences* (IOP-China), Beijing, in 2003. Between 2003 and 2008, he worked in IOP-China as intern, assistant professor and associated professor. In 2008, he joined NTT Basic Research Laboratories as Research Associate and Senior Research Associate. In 2013, he became a special assigned professor in IOP-China. In 2016, he joined *University of science and technology of China* as professor. His team is dedicated to develop scalable superconducting quantum computing. He has been responsible for developing several generations of highly coherent superconducting quantum processors, setting up ultra-low noise control and readout platform at millikelvin temperatures, and building up software and hardware of the room-temperature electronics. This team developed a two-dimensional programmable superconducting quantum processor, Zuchongzhi, which is composed of 66 functional qubits in a tunable coupling architecture. Based on this state-of-the-art quantum processor, they achieved larger-scale random quantum circuit sampling, with a system scale of up to 60 qubits and 24 cycles and therefore exhibited strong quantum advantage. The achieved sampling task is about 6 orders of magnitude more difficult than that of Sycamore.

Topic 2: Quantify, an open-source control framework for quantum computing experiments

Abstract: Operating a quantum computer in the NISQ era is an often-underestimated challenge. Specifically, the tune-up and execution of quantum algorithms, which consist of physics experiments requiring access to control parameters and measured signals, as well as classical logic. Typically, these are defined at a higher level of abstraction and are not supported by current control architectures because they are limited to expressing experiments as either a series of classical pulses or variants of QASM. Here, we present Quantify, a robust and extensively documented open-source experiment platform inspired by PycQED (Rol et al. 10.5281/zenodo.160327). Quantify contains all the basic functionality to control experiments (e.g., instrument management, live plotting, data storage, etc.), as well as a novel scheduler featuring a unique hybrid control model allowing quantum gate- and pulse-level descriptions to be combined in a clearly defined and hardware-agnostic

way. The scheduler allows parameterized expressions and classical logic to make efficient use of hardware backends that support this. This opens up new avenues for efficient execution of calibration routines as well as variational quantum algorithms (VQA).

Instructor: Adriaan Rol

Biography: Adriaan received his PhD in applied physics (cum laude) from the group of Prof. Leo DiCarlo at Delft University of Technology/QuTech where he worked on control for superconducting quantum systems. He has developed new calibration and characterization protocols, as well as a new type of two-qubit gate, and made key contributions to the development of Quantum Infinity, a transmon-based full-stack quantum computer that preceded QuTech's Quantum Inspire platform. Adriaan has worked closely with experts from all layers of the stack, which was formally recognized through the Zurich Instruments Pioneer Award, and resulted in awardwinning papers. Before starting his PhD, Adriaan was on the executive board of a non-profit business consultancy. Adriaan is Director of Research & Development and one of the founders of Orange Quantum Systems, a company building diagnostic systems for QPU characterization.

Topic 3: Heterogeneous quantum-classical programming using Quingo

Abstract: Quantum computers hold the promise to accelerate solving problems intractable by classical computers. Programmers write quantum programs to solve problems, which are compiled and optimized by a quantum computer, resulting in some quantum binaries executable on quantum hardware. In this tutorial, I will start with a short introduction to the basics of quantum computing, followed by the rationale and design of a new quantum programming language, Quingo. After a short presentation on the Quingo software ecosystem, I will demonstrate how to write some heterogeneous quantum-classical programs using the Quingo framework and execute them on a Chinese quantum computing platform on the cloud based on superconducting qubits.

Instructor: Xiang Fu

Biography: Dr. Xiang Fu is an assistant researcher at the Institute of Quantum Information and State Key Laboratory of High Performance Computing, School of Computer Science and Technology, *National University of Defense Technology* (NUDT). He holds a bachelor's degree in electronics from *Tsinghua University*, a master's degree in computer architecture from NUDT, and a PhD in quantum control architecture from QuTech, Delft University of Technology in the Netherlands. His work focuses on enabling the programmability over a quantum computing system, which is highlighted by the first quantum control microarchitecture QuMA, an executable quantum computing instruction set architecture, eQASM, and the heterogeneous quantum-classical programming framework Quingo. He is honored by the MICRO-50 Best Paper Award and Top Picks 2017. His main research interests include quantum programming languages, quantum compilation, and quantum architecture.

Topic 4: Quantum compilation for near-term quantum computers

Abstract: Quantum computing can solve some problems that are intractable by classical computers. We are now entering the noisy intermediate-scale quantum (NISQ) computing era. NISQ computers have short qubit coherence time and high gate error rates, making it difficult to implement reliable quantum computation. Moreover, NISQ computers have hardware constraints such as limited qubit connectivity, making quantum circuits not directly executable. Compilation techniques are required to compile hardware-agnostic quantum circuits into hardware-compatible ones, meanwhile minimizing circuit sizes and hardware errors for highfidelity computation. In this tutorial, I will first introduce the basics of quantum compilation and then review different types of compilers for NISQ computers such as general-purpose compilers, application-specific compilers, noise-aware compilers, etc.

Instructor: Lingling Lao

Biography: Lingling Lao is a Research Fellow in Prof. Dan Browne's group at the Department of Physics and Astronomy, University College London. She holds a bachelor's degree in electronics from Harbin Institute of Technology, a master's degree in Electronics from *Northwestern Polytechnical University*, and a PhD in quantum computing from QuTech, Delft University of Technology in the Netherlands under the supervision of Prof. Koen Bertels. Her PhD thesis was entitled "Quantum computing in practice: fault-tolerant protocols and circuit-mapping techniques". Her research interests include quantum compilation, quantum error correction, fault-tolerant quantum computing, and quantum computer architecture.

Organizers
Professor Yuxin Deng has extensive experience in organizing national and international conferences. He has organized TASE 2016, FMAC 2017, FMAC 2019, and SETTA 2019. He has also served in around 40 program committees of international conferences in the area of formal methods such as ICALP 2013, ICALP 2016, ICALP 2018, CONCUR 2019, CONCUR 2021, CAV 2021, and CAV 2022. He is an invited speaker for CONCUR 2018. He is a guest editor for a special issue at Theoretical Computer Science.

Dr. Xiang Fu organized 1st workshop on Programmable Quantum Control during QIP 2020. He is an invited speaker for Quantum Algorithm and Software 2019 by Pengcheng Laboratory, Workshop on Quantum Computation and Information 2020 by Institute of Physics, *Chinese Academy of Sciences*, etc.

Tutorial T2
Taming Delays in Cyber-Physical Systems

Abstract: The advent of systems of cooperative cyber-physical systems draws attention to a central problem of networked and distributed control systems: the ubiquity of delay in feedback loops between logically or spatially distributed components, which is not adequately re ected in traditional models of hybrid-state dynamics based on ordinary differential equations and immediate transitions. The occurrence of feedback delays may significantly alter a system's dynamic response. Unmodelled delays in a control loop consequently have the potential to invalidate any stability or safety certificate obtained on a related delay-free model, which is the current practice in hybrid-system analysis. In this tutorial, we will present various approaches to the analysis and correct-by-construction design of dynamical systems subject to delayed information exchange, as pertinent to distributed hybrid systems. We will explain automatic verification procedures for invariance properties over bounded or unbounded temporal horizons. This analytical view will be complemented by a constructive one for synthesizing delay-resilient control strategies for discrete and hybrid discrete-continuous dynamics.

Motivation
Conventional embedded systems have over the past decades vividly evolved into an open, interconnected form that integrates capabilities of computing, communication, and control, thereby triggering yet another round of global revolution of the information technology. This form, now known as cyber-physical systems (CPS), has witnessed an increasing number of safety-critical systems particularly in major scientific projects vital to people's livelihood. Prominent examples include automotive electronics, health care, nuclear reactors, high-speed transportations, manned space ight, etc., in which a malfunction of any software or hardware component would potentially lead to catastrophic consequences like significant casualties and economic losses. Meanwhile with the rapid development of feedback control, sensor techniques, and computer control, time delays have become an essential feature underlying both the continuous evolution of physical plants and the discrete transition of computer programs, which may well annihilate the stability/safety certificate and control performance of embedded systems. Traditional engineering methods, e.g., testing and simulations, are nevertheless argued insufficient for the zero-tolerance of failures incurred in time-delayed systems in a safety-critical context. Therefore, how to rigorously verify and design reliable safety-critical embedded systems involving delays tends to be a grand challenge in the computer science and control theory community.

In contrast to delay-free systems, time-delayed systems yield substantially higher theoretical complexity thus rendering the underlying design and verification tasks exceedingly harder. This tutorial focuses on formal modelling, verification, and control synthesis of time-delayed, networked hybrid systems, while particularly addressing the stability/safety verification of continuous dynamics governed by delay differential equations (DDEs), the control-strategy synthesis of discrete dynamics captured by delayed safety games, and the switching-logic synthesis of hybrid dynamics modelled as delay hybrid automata, with applications in a typical set of time-delayed dynamics originated from realms of biology, control, astronautics, and transportation.

Though time delays have been extensively studied in the literature of mathematics and control theory from a qualitative perspective, automatic verification and synthesis methods addressing feedback delays in hybrid discrete-continuous systems are still in their infancy. This tutorial delivers a peek into the blueprint of hybrid systems under delays, wherein the attendees will grasp quantitative techniques, e.g., interval Taylor models, validated simulations, sensitivity analysis, linearization, spectral analysis, and invariant generation, and thereby learn how to achieve formal verification and synthesis of time-delayed systems leveraging these commonly used ingredients in the community of formal methods.

With plenty of illustrating examples provided throughout the tutorial, the attendees may realize how to make mathematical techniques contribute to real-world applications, and hopefully, be inspired by the success stories on verification and synthesis of time-delayed dynamics presented in this tutorial and start to use the proposed approaches, or even develop their own formal methods, for their future work.

List of Topics

- An overview of CPS and hybrid systems featuring networked delays, in particular, delay hybrid automata, which augment classical hybrid automata with time delays in both discrete jumps and the continuous evolution.

- Verifying stability/safety of delayed differential dynamics possibly with inputs and disturbances:
  - bounded verification based on validated simulations and the boundary propagation of reachable sets;
  - unbounded verification via interval Taylor models, spectral analysis plus linearization, and control barrier-certificate functional synthesis.

- Synthesizing safe controllers resilient to time delays:

  - for discrete dynamics:
    - safety games played under discrete delays and their decidability of controllability,
    - incremental synthesis of delay-resilient control strategies,
    - equivalence of qualitative controllability for safety games under different delay patterns;
  - for hybrid discrete-continuous dynamics:
    - local invariant generation for DDEs based on reachable-set computation,
    - global invariant generation for delayed hybrid dynamics based on fixed-point computation,
    - safe switching-logic synthesis for delayed hybrid dynamics based on invariant generation.

The above topics cover our recent results published at IEEE Trans. Automat. Contr., Acta Informatica, CAV, FM, HSCC, CDC, ATVA, FORMATS, etc.

Biographies of the Instructors

Dr. Naijun Zhan is a distinguished professor at the State Key Lab. of Computer Science, Institute of Software, *Chinese Academy of Sciences*, Beijing, China. He received the Ph.D. degree in computer science from the Institute of Software, *Chinese Academy of Sciences* in 2000, and M.Sc. in computer science and B.Sc in mathematics both from *Nanjing University* respectively in 1996 and 1993. Prior to joining the Institute of Software, he was a Research Fellow with the Faculty of Mathematics and Information, University of Mannheim, Mannheim, Germany, from 2001 to 2004. His research interests include formal techniques for the design of real-time and hybrid systems, program verification, and modal/temporal logics. He has published two books and more than one hundred papers in the leading international journals and conferences. He is a member of the editorial boards of Formal Aspects of Computing and Journal of Logical and Algebraic Methods in Programming; PC co-chair of FM 2021 and other conferences; and serves in the program committees of CAV, RTSS, HSCC, EMSOFT, HSCC, ICCPS, etc.

Dr. Mingshuai Chen is a postdoctoral researcher at the Chair for Software Modeling and Verification, Department of Computer Science, RWTH Aachen University, Aachen, Germany. He received his Ph.D. degree in computer science from the Institute of Software, *Chinese Academy of Sciences* in 2019, and B.Sc. in computer science from the School of Computer Science and Technology, Jilin University in 2013. His main research interest lies in the formal verification and synthesis of programs and hybrid systems for ensuring the reliability and effectiveness of safety-critical cyber-physical systems. He has published over twenty peer-reviewed papers at agship journals/ conferences including IEEE Trans. Automat. Contr., Acta Informatica, CAV, TACAS, IJGAR/CADE, FM, etc. He serves in the reviewer panel of AMS Mathematical Reviews and the program/repeatability committees of RTCSA, ADHS, VMCAI, SYNASC, etc. He was the awardee of the Distinguished Paper Award at ATVA 2018, Best Paper Award at FMAC 2019, and the CAS-President Special Award in 2019.

Tutorial T4

Integrating Compute Acceleration Into Embedded System Design Using Vitis

Abstract: AMD-Xilinx is the inventor of the FPGA, programmable SoCs, and now, the ACAP. Our highly flexible programmable silicon, enabled by a suite of advanced software and tools, drives rapid innovation across a wide span of industries and technologies – from consumer to cars to the cloud. The goal of this tutorial is to introduce the Vitis software development environment for designing

accelerators for embedded systems using Vitis.. Attendees will have the opportunity to learn how to use these tools, test the tutorial examples on the target boards, and explore the latest Vitis AI, Versal AIE technologies.

Topics to be covered

- Vitis development framework and design flows
- Vitis HLS project creation flow, profiling and design analysis
- Vitis embedded platform based design flow
- Introduction to Vitis AI and Versal AIE

Biographies of the speakers:
Joshua Lu is currently working for AMD AECG group (former Xilinx) as senior product application manager. He covers the Asia region, his group is part of the Xilinx University Program team which is technical group inside AECG group CTO office. He has more than 15 years' experience with Xilinx embedded solutions and actively engaged in Xilinx open source project called PYNQ. He has published Vivado, Zynq, NetFPGA related textbooks in Chinese.

Yajun Ha is currently a Professor at ShanghaiTech University, China. Before this, he was a Scientist and Co-Director, I2R-BYD Joint Lab at Institute for Infocomm Research, Singapore, and an Adjunct Associate Professor at the Department of Electrical & Computer Engineering, National University of Singapore. Between Mar 2004 and Dec 2013, he was an Assistant Professor with National University of Singapore. He received his Ph.D. degree from Katholieke Universiteit Leuven (KULeuven), Leuven, Belgium, in 2004 and worked at the same period as a researcher with the Inter-University MicroElectronics Center (IMEC), Leuven, Belgium.

Tutorial T6

Manycore processing-in-memory systems for accelerating deep learning applications

Abstract: Training machine learning (ML) models at the edge (on-chip training on end user devices) can address many pressing challenges including data privacy/security, increase the accessibility of ML applications to different parts of the world by reducing the dependence on the communication fabric and the cloud infrastructure, and meet the real-time requirements of AR/VR applications. However, existing edge platforms do not have sufficient computing capabilities to support complex ML tasks such as training large CNNs. This tutorial will consider solutions based on Resistive Random-access Memories (ReRAMs) to address these challenges and answer the following questions: (1) How to use ReRAMs as a Processing-in-memory (PIM) architecture? (2) How to map machine learning techniques to ReRAMbased manycore systems to improve performance and energy-efficiency? (3) What are the challenges associated with different types of Deep Learning applications (such as CNNs and GNNs) when mapped to PIM-based manycore architectures? (4) How can we ensure reliability in these architectures? To address these outstanding challenges, outof- the-box approaches need to be explored. By integrating machine learning algorithms, data analytics, statistical modeling, and design of advanced computing systems, this tutorial will engage a broad section of ESWEEK conference attendees. It also highlights how machine learning and embedded system researchers can join hands to design energyefficient and reliable miniaturized computing systems. This tutorial is targeted towards university researchers/professors, MS/Ph.D. students, professionals from industry, and IC designers, who want to learn how to use ReRAMs for ML applications, as well as experienced researchers looking for exciting new directions in PIM. We expect at least 30-40 attendees. We will announce the tutorial through our regular networks, as well as mailing lists of CEDA, IEEE CAS Society, IEEE Computer Society and ACM SIGDA (SIGDA E-News that reaches thousands of EDA professionals, and we will leverage it to publicize our tutorial). The target audience matches the typical ESWEEK participant very well. We also hope that this tutorial will allow ESWEEK to reach out to a wider audience and help boost attendance.

List of topics to be covered
The main objective of the tutorial is to help attendees understand the potential of emerging ReRAM-based computing systems both as memory and as a processing-in-memory architecture. We will elaborate how ReRAMs can implement complex ML applications in an energy-efficient manner. Specific topics are as follows:

- Application of ReRAMs as a processing-in-memory system
- Mapping of ML algorithms to ReRAMs for inferencing and training

- Addressing computation and communication challenges of the PIM-based manycore system
- Ensuring high performance and accuracy by addressing reliability challenges in ReRAMs

Organizers

Janardhan Rao Doppa, *Washington State University*

Jana is currently a George and Joan Berry Chair Associate Professor with *Washington State University* (WSU), Pullman, WA, USA. His current research interests are at the intersection of machine learning and computing systems design. He received a NSF CAREER Award (2019), an Outstanding Paper Award at the AAAI (2013) conference, a Google Faculty Research Award (2015), the Outstanding Innovation in Technology Award from *Oregon State University* (2015). He received the Reid-Miller Teaching Excellence Award (2018) and the Outstanding Junior Faculty in Research Award (2020) from the Voiland College of Engineering and Architecture at WSU. He is among the 15 outstanding young researchers selected to give Early Career Spotlight talk at the International Joint Conference on Artificial Intelligence (2021).

Biresh Kumar Joardar, *Duke University*

Biresh finished his PhD from *Washington State University* in 2020. He is currently a Post-doctoral Computing Innovation Fellow (CI-Fellow) at the Department of Electrical and Computer Engineering at *Duke University*. His current research interests include machine learning, manycore architectures, accelerators for deep learning, hardware reliability and security. He received the 'Outstanding Graduate Student Researcher Award' at Washington state University in 2019. His works have been nominated for Best Paper Awards at prestigious conferences such as DATE and NOCS.

Tutorial T7

Hardware Security and Trust Verification

Abstract: System-on-Chip (SoC) is the brain behind computing and communication in a wide variety of systems. Reusable hardware *Intel*lectual Property (IP) based System-on-Chip (SoC) design has emerged as a pervasive design practice in the industry to dramatically reduce design and verification cost while meeting aggressive time-to-market constraints. Growing reliance on these pre-verified hardware IPs, often gathered from untrusted third-party vendors, severely affects the security and trustworthiness of SoC computing platforms. These third-party IPs may come with deliberate malicious implants to incorporate undesired functionality (e.g., hardware Trojans), undocumented test/debug interface working as hidden backdoor, or other integrity issues. It is extremely difficult to verify integrity and trust of hardware IPs due to several reasons including (a) lack of a golden reference model or incomplete specification, (b) exponential space of diverse types of complex IPs and IP-specific vulnerabilities, (c) lack of automated and scalable CAD tools for IP trust verification, and (d) lack of security metrics to measure the security robustness of a given design or mitigation technique. While functional validation has received significant attention over the years, it is critical to perform "security and trust verification" for designing trustworthy systems.

This tutorial will provide a comprehensive overview of both fundamental concepts and recent advances in hardware security and trust validation using simulation-based approaches, formal methods as well as side-channel analysis. Specifically, the tutorial will consist of four parts. The first part will introduce security vulnerabilities (threats) and various challenges associated with trust validation of hardware IPs. It will highlight recent advances in developing trust metrics and benchmarks. The second part will cover assertion-based security validation utilizing automated generation of security assertions. It will also cover automated test generation techniques for activation of security assertions. The third part will describe how formal verification techniques (including model checking, SAT solving, theorem proving and equivalence checking) can be effectively utilized for validation of hardware security vulnerabilities. The fourth part will discuss how side-channel analysis can be effectively utilized to detect malicious implants. It will conclude with a discussion on integration of security verification in existing functional validation methodology.

Tutorial Organization

This tutorial will be based on both the prior work of the presenter and the work of other prominent researchers in this field. It will cover theoretical aspects as well as industrial practices with real-life case studies. Following is the organization of the tutorial that consists of four major components.

Part I: Introduction and Motivation

- Security and trust validation challenges (why functional validation is not enough)

- Development of threat models, trust metrics and benchmarks

Part II: Simulation-based Security and Trust Validation

- Automated vulnerability analysis of RTL models

- Assertion-based validation of security vulnerabilities

- Scalable test generation for activation of malicious implants

Part III: IP Trust Verification using Formal Methods

- Model checking to verify unwanted scenarios (security properties)

- Equivalence checking to ensure nothing more, nothing less.
- Scalable trust verification using theorem proving

Part IV: Hardware Trojan Detection using Side-Channel Analysis

- Logic testing versus side-channel analysis for detecting malicious implants

- Trojan detection using dynamic current (power), path delay, and electromagnetic emanations

Part V: Conclusions and Future Directions

- Seamless integration of security verification techniques in existing functional validation flow

- Future trends and research directions

Presenter

Prabhat Mishra is a Professor in the Department of Computer and Information Science and Engineering and a UF Research Foundation Professor at the *University of Florida*. His research interests include embedded and cyber-physical systems, hardware security and trust, energy-aware computing, formal verification, system-on-chip validation, and post-silicon debug. He received his Ph.D. in Computer Science from the University of California at Irvine in 2004. Prior to joining *University of Florida*, he spent several years in various companies including *Intel*, Motorola, Sasken, Synopsys and Texas Instruments. He has published 8 books, 35 book chapters, 21 patents/copyrights, and more than 200 research articles in premier international journals and conferences. His research has been recognized by several awards including the NSF CAREER Award from the National Science Foundation, IBM Faculty Award, three Best Paper Awards (ISQED'16, VLSID'11 and CODES+ISSS'03) as well as seven Best Paper Nominations (DATE'19, ASPDAC'17, NANOARCH'13, VLSI'13, DATE'12, DAC'09, VLSI'09), and EDAA Outstanding Dissertation Award from the European Design Automation Association. Prof. Mishra currently serves as an Associate Editor of IEEE Transactions on VLSI Systems and ACM Transactions of Embedded Computing Systems. He has served as the Deputy Editor-in-Chief of IET Computer & Digital Techniques, and as an Associate Editor of ACM Transactions on Design Automation of Electronic Systems, and IEEE Design & Test, and Springer Journal of Electronic Testing. He is an IEEE Fellow and an ACM Distinguished Scientist. ■

EDUCATION CLASSES

SATURDAY CLASSES, OCTOBER 8

CST - 10/7/2022 14:00-16:00 PST - 10/6/2022 23:00-01:00 EST - 10/8/2022 02:00-04:00 CET - 10/8/2022 08:00-10:00	**Education Class on Secure Computing Systems, by Prof. Ahmad Reza Sadeghi, TU Darmstadt (EC1)**	**Education Class on Biomedical CPS, by Prof. Partha Roop, Univ. of Auckland (EC2)**	**Education Class on Crypto and Security on FPGA, by Prof. Nele Mentens, Leiden University (EC3)**
CST - 10/8/2022 22:00-00:00 PST - 10/8/2022 07:00-09:00 EST - 10/8/2022 10:00-12:00 CET - 10/8/2022 16:00-18:00	**Education Class on Robust Federated Learning, by Prof. Farinaz Koushanfar, UCSD (EC4)**	**Education Class on Confidential Computing, by Prof. Christof Fetzer, TU Dresden (EC5)**	**Education Class on Attacks on Crypto Systems, by Prof. Debdeep Mukhopadhyay, IIT Kharagpur (EC6)**

SUNDAY CLASSES, OCTOBER 9

CST - 10/9/2022 14:00-16:00 PST - 10/8/2022 23:00-01:00 EST - 10/9/2022 02:00-04:00 CET - 10/9/2022 08:00-10:00	**Education Class on Automated Repair, by Prof. Abhik Roychoudhury, NUS (EC7)**	**Education Class on Privacy in IoT, by Prof. Thorsten Strufe, KIT (EC8)**	
CST - 10/9/2022 22:00-00:00 PST - 10/9/2022 07:00-09:00 EST - 10/9/2022 10:00-12:00 CET - 10/9/2022 16:00-18:00	**Education Class on ML for Trojan Detection, by Prof. Houman Homayoun, UC Davis (EC9)**	**Education Class on Hardware Security, by Prof. Ramesh Karri, NYU (EC10)**	

Education Class 1

A Hitchhiker's Guide to Systems Security: The Art and Science of Building and Breaking Secure Computing Systems

Instructor:
Ahmad-Reza Sadeghi, TU Darmstadt

Abstract:
The science of systems security is concerned with security aspects of computing systems at both software and hardware layers. The ever-increasing complexity of computing systems, emerging technologies such as IoT and AI, and advancing attack capabilities pose a variety of (new) challenges on the design and implementation of security concepts, methods and mechanisms.

This talk provides an overview of our journey through the systems security research universe. It points out several aspects of advancing software security and hardware-assisted security in academia and industry. A particular focus of the talk is devoted to the current crucial security threat posed by software-exploitable hardware vulnerabilities that put our critical systems, and hence our society, at risk. Finally, we discuss our future vision and new research directions in systems security.
Bio: Ahmad-Reza Sadeghi is a professor of Computer Science and the head of the System Security Lab at Technical University of Darmstadt, Germany. He has been leading several Collaborative Research Labs with *Intel* since 2012, and with Huawei since 2019.

He has studied both Mechanical and Electrical Engineering and holds a Ph.D. in Computer Science from the University of Saarland, Germany. Prior to academia, he worked in R&D of IT-enterprises, including Ericsson Telecommunications. He has been continuously contributing to security and privacy research field. He was Editor-In-Chief of IEEE Security and Privacy Magazine, and currently serves on the editorial board of ACM TODAES, ACM TIOT, and ACM DTRAP.

For his influential research on Trusted and Trustworthy Computing he received the renowned German "Karl Heinz Beckurts" award. This award honors excellent scientific achievements with high impact on industrial innovations in Germany. In 2018, he received the ACM SIGSAC Outstanding Contributions Award for dedicated research, education, and management leadership in the security community and for pioneering contributions in content protection, mobile security and hardware-assisted security. In 2021, he was honored with *Intel* Academic Leadership Award at USENIX Security conference for his influential research on cybersecurity and in particular on hardware-assisted security.

Education Class 2

A synchronous approach for the design of biomedical cyber-physical systems

Instructor:
Partha Roop, University of Auckland
Co-Instructors:
Nathan Allen, University of Auckland
Hamond Peare, NYU

Abstract:
Many biomedical systems use embedded controllers to control physical processes, and consequently form a class of Cyber-Physical Systems (CPSs). Examples range from pacemakers to automated insulin pumps. These systems must work safely at all times. In this tutorial, we will focus on a design methodology for such CPSs.

We rely on the well-known synchronous approach for modelling the biological processes i.e. the human organ in question at a suitable abstraction-level as well as the medical device. Using the synchronous approach helps in both modelling and verification as well as automated code generation. The synchronous approach provides well known benefits such as deterministic execution, which is an ideal fit for ensuring safety. This tutorial will introduce a systematic design approach starting with modelling of biomedical devices using the SCCharts synchronous language. Next, the *Intel* NIOS II platform will be introduced as a means of prototyping a given medical device. This prototype device can then be run in closed-loop with a real-time version of the associated human organ. We will use a cardiac pacemaker as a running example and will therefore briefly introduce modelling of the cardiac conduction system using compositional models. We will then demonstrate how the pacemaker can be tested, in closed-loop, with the adjoining model of the cardiac conduction system. Additionally, we will present an approach for the formal verification of this device, using the tool UPPAAL and associated properties to ensure its correct operation. We will also present how pacemakers can be secured in the face of adversarial attacks, before presenting key areas for future research.

Bio: Partha's research interests are in Digital Health, Formal Methods for Safety-Critical applications of AI and Machine Learning, and Real-Time Systems. Partha is working with colleagues from the Medical School and the Auckland Bioengineering Institute (ABI) on new techniques developed by his group known as organ on a chip. He is also interested in heart rate variability and Biofeedback.

Partha is an academic in the Department of Electrical and Computer Engineering at the University of Auckland. He is currently the Associate Dean (International) for the Faculty of Engineering.

He completed his PhD in Computer Science and Engineering at the *University of New South Wales*, Sydney, Australia, a M.Tech at Indian Institute of Technology in Kharagpur, India and a BE degree at Anna University (College of Engineering), Madras, India. Partha had visiting positions in Kiel University, Germany (collaboration with the Embedded Systems Group, French National Laboratory of Informatics and Control (SPADES team Grenoble, and Iowa State University. He collaborates with University of California, Berkeley in the PRET project.

He heads the precision timed systems research group. His group has created a tool-chin for the design of automation and embedded systems using the IEC61499 standard. The tool and associated benchmarks are available for download.

Partha co-founded APIMatic, a cloud services company for automatic SDK generation using the model-driven approach. He has co-authored two research monographs:

- Model-Driven Design Using IEC 61499: A Synchronous Approach for Embedded and Automation Systems (2015)
- Correct-by-Construction Approaches for SoC Design (2013)

Education Class 3

FPGA design for cryptography and security

Instructor:
Nele Mentens, *Leiden University* and KU Leuven

Abstract:
Field-Programmable Gate Arrays (FPGAs) are configurable hardware architectures that combine the performance of Application-Specific Integrated Circuits (ASICs) with the programmability of microprocessors. FPGAs are popular implementation platforms for cryptography and security applications. FPGAs are used as accelerators for cryptographic algorithms and network security solutions, and as patchable trusted computing platforms. This lecture will first introduce the technology and design flow of FPGAs. Next, the lecture will concentrate on the use of FPGAs as cryptographic accelerators, as network security solutions and as trusted computing platforms.

Bio: Nele Mentens is a professor at *Leiden University* in the Netherlands and KU Leuven in Beligum. Her research interests are in the field of configurable computing and hardware security. She was/is the PI in around 25 finished and ongoing research projects with national and international funding. She serves/served as a program committee member of renowned international conferences on security and hardware design. She was the general co-chair of FPL'17 and she was/is the program chair of FPL'20, CARDIS'20, RAW'21, VLSID'22 and DDECS'23. She is (co-)author in around 150 publications in international journals, conferences and books. She received best paper awards and nominations at CHES'19, AsianHOST'17 and DATE'16. Nele serves as an associate editor for IEEE TIFS, IEEE CAS Magazine, IEEE S&P, and IEEE TCAD.

Education Class 4

Robustness against Poisoning Attacks in Centralized and Federated Deep Learning Scenarios: A Survey

Instructor:
Farinaz Koushanfar, *UC San Diego*

Abstract:
Deep Learning (DL) has been increasingly deployed in various real-world applications due to its unprecedented performance and automated capability of learning hidden representations. While DL can achieve high task performance, the training process of a DL model is both time and resource-consuming. Therefore, current supply chains of the DL models assume the customers obtain pre-trained Deep Neural Networks (DNNs) from the third-party providers that have sufficient computing power. In the centralized setting, the model designer trains the DL model using the local dataset. However, the collected training data may contain erroneous or poisoned data points. The model designer might craft malicious training samples and inject

backdoors in the DL model distributed to the users. As a result, the user's model will malfunction. In the federated learning setting, the cloud server aggregates local models trained on individual local datasets and updates the global model. In this scenario, the local client could poison the local training set and/or arbitrarily manipulate the local update. If the cloud server incorporates the malicious local gradients in model aggregation, the resulting global model will have degraded performance or backdoor behaviors. In this class, we present a comprehensive overview of contemporary data poisoning and model poisoning attacks against DL models in both centralized and federated learning scenarios. In addition, we review existing detection and defense techniques against various poisoning attacks.

Bio: Farinaz Koushanfar is a professor and Henry Booker Faculty Scholar in the Electrical and Computer Engineering (ECE) department at University of California San Diego (UCSD), where she is also the co-founder and co-director of the UCSD Center for Machine-*Intelligence*, Computing & Security (MICS). Her research addresses several aspects of efficient computing and embedded systems, with a focus on hardware and system security, real-time/energy-efficient big data analytics under resource constraints, design automation and synthesis for emerging applications, as well as practical privacy-preserving computing. Dr. Koushanfar is a fellow of the Kavli Foundation Frontiers of the National Academy of Engineering and a fellow of IEEE. She has received a number of awards and honors including the Presidential Early Career Award for Scientists and Engineers (PECASE) from President Obama, the ACM SIGDA Outstanding New Faculty Award, Cisco IoT Security Grand Challenge Award, MIT Technology Review TR-35, Qualcomm Innovation Awards, as well as Young Faculty/ CAREER Awards from NSF, DARPA, ONR and ARO.

Education Class 5

Confidential Computing – protecting the confidentiality, integrity, and consistency of applications.

Instructor:
Christof Fetzer, TU Dresden

Abstract:
An application might run in a cloud, e.g., an edge cloud with limited physical security. Or, the cloud provider might be in a different jurisdiction. We introduce the fundamental concepts of confidential computing and show how one can ensure the confidentiality, integrity, and consistency of applications – even if we cannot trust the external provider. We will also show how to convert a cloud-native application into a cloud-confidential application.

Bio: Christof Fetzer received his Ph.D. from *UC San Diego* (1997). As a student he received a two-year scholarship from the DAAD and won two best student paper awards (SRDS and DSN). He was a finalist of the 1998 Council of Graduate Schools/UMI distinguished dissertation award and won an IEE mather premium in 1999. Dr. Fetzer joined AT&T Labs-Research in August 1999 and was a principal member of technical staff until March 2004. Since April 2004, he is head of the Systems Engineering Chair in the Computer Science Department at the Dresden University of Technology. He is the chair of the Distributed Systems Engineering International Masters Program at the Computer Science Department. Prof. Dr. Fetzer has published over 130 research papers in the field of dependable systems.

Education Class 6

An Exposition of Fault Based Attacks on Modern Cryptosystems

Instructor:
Debdeep Mukhopadhyay,
Indian Institute of Technology Kharagpur

Abstract:
The lecture would provide an overview on fault attacks on modern cryptosystems. We shall start with some classical fault based cryptanalysis of the Advanced Encryption Standard (AES), called Differential Fault Analysis (DFA). Later we discuss another class of fault attacks, called Differential. Fault Intensity Attacks (DFIA), and show how fault bias can be utilized to break redundancy based countermeasures based on classical fault tolerance. In the later part of the talk, we present an overview on some of the advanced fault attack techniques, namely Statistical Ineffective Fault Attacks (SIFA), and Fault Template Attacks (FTA). We also aim to discuss on suitable techniques to thwart these menacing classes of physical attacks. The tutorial would be accompanied with small demonstrations to elucidate the concepts presented.

Bio: Prof. Debdeep Mukhopadhyay is currently a Professor at the Department of CSE, *IIT Kharagpur*. At *IIT Kharagpur* he initiated Secured Embedded Architecture Laboratory (SEAL), focusing on Hardware-Security. He had worked as,

visiting scientist at K U Leuven-Belgium, NTU-Singapore, visiting Associate Professor of NYU-Shanghai, Assistant Professor at IIT Madras, and Visiting Researcher at NYU-Tandon School of Engineering, USA. He holds a Ph.D, M.S., and a B.Tech from *IIT Kharagpur*. His research interests are on the topics of Cryptographic Engineering and Hardware Security. Recently he is intrigued by adversarial attacks on machine learning and encrypted computations.

Prof. Mukhopadhyay is the recipient of the prestigious Shanti Swarup Bhatnagar Award 2021 for Science and Technology (highest science honor in India below the age of 45) and is a Fellow of the Indian National Academy of Engineers. He was awarded Khosla National Award from IIT Roorkee 2021, DST Swarnajayanti Fellowship 2015-16, INSA Young Scientist award, INAE Young Engineer award, and Associateship for the Indian Academy of Sciences.

Education Class 7

Fuzzing and automated repair of security vulnerabilities in embedded software

Instructor:
Abhik Roychoudhury, National University of Singapore

Abstract:
Fuzz testing is a popular technique for detecting security vulnerabilities in software systems. It proceeds by compile time instrumentation, along with a run-time biased random search to find crashing inputs. The biased random search may be guided by an objective function or by logical constraints leading to different testing setups. In this paper, we will share various mechanisms and viewpoints in adapting or adopting greybox fuzzing for embedded software. This is of increased importance due to increased movement of the attack surface towards the edge. Moreover, as the security vulnerabilities are found and published, the software systems suffer from increased exposure, which can be alleviated by automated program repair technologies. In a synergistic setting, the searches over program edits in program repair, and the search over program inputs can strengthen each other. The tutorial give the audience wide exposure on greybox fuzzing and whitebox fuzzing (also known as symbolic execution) technologies, as well as forward looking techniques for automated program repair, which are seeing increased adoption.

Bio: Abhik Roychoudhury is a Provost's Chair Professor of Computer Science at the National University of Singapore, where he has been since 2001 after receiving his Ph.D. in Computer Science from the State University of New York at Stony Brook in 2000. He is the Director of the National Satellite of Excellence in Trustworthy Software Systems at Singapore (2019-23). He has also helped set up the Singapore Cyber-security Consortium, which is a consortium of 25 companies in the cyber-security space engaging with academia for research and collaboration. Abhik's research focuses on software testing and analysis, software security and trust-worthy software construction. His research team is known for contributions to automated program repair, and fuzz testing. Abhik is a member of the Steering committee of the flagship conferences in Software Engineering, International Conference on Software Engineering (ICSE) and Symposium on Foundations of Software Engineering (FSE). He has served as an Associate Editor of IEEE Transactions on Software Engineering (TSE), IEEE Transactions on Dependable and Secure Computing (TDSC) and ACM Transactions on Software Engineering and Methodology (TOSEM). His former doctoral students have been placed all over the world as academics (Peking University, University College London, Max-Planck Institute, University of Melbourne, Beihang University, SUSTech, SUTD).

Education Class 8

You better act normal! Ubiquitous electronic observation: Threats and Attempted Solutions

Instructor:
Thorsten Strufe, KIT, privacy and IT security

Abstract:
The evolution from desktops over mobile devices to smart-X has brought near perfect visibility on the behavior and whereabouts of citizens, both in the digital and the real world. This comes with numerous features and simplifications of processes, and increasing utility for users and observed individuals. It does raise concerns about the potential for abuse, from unexpected identification to the disclosure of sensitive characteristics, health conditions, or personal peculiarities, despite constant claims of "anonymization" and "GDPR compliance".

In this talk we will discuss reasons of the situation, and how claimed protection has proven ineffective under scrutiny.

Bio: Thorsten Strufe is professor for IT Security at *Karlsruhe Institute of Technology* (KIT/KASTEL), and adjunct professor for Privacy and Network Security at

TU Dresden, as deputy speaker of the Excellence Centre for Tactile Internet with Human-in-the-Loop (CeTI). His research interests lie in the areas of privacy and resilience, especially in the context of social networking services and large scale distributed systems.

Recently, he has focused on studying privacy implications of user behavior and possibilities to provide privacy-preserving and secure networked services. Previous posts include faculty positions at TU Dresden, TU Darmstadt, and Uni Mannheim, as well as postdoc/researcher positions at EURECOM and TU Ilmenau.

Education Class 9

ML-Assisted Hardware Trojan Detection

Instructor:
Houman Homayoun, UC Davis

Abstract:
With the growth and globalization of IC design and development, there is an increase in the number of Designers/ Design houses. As setting up a fabrication facility may easily cost upwards of $20 billion, costs for advanced nodes may be even greater. IC design houses that cannot produce their chips in-house have no other option but to make use of external foundries that are often in other countries. Establishing trust with these external foundries can be a challenge, and these foundries are assumed to be untrusted. The use of these untrusted foundries in the global semiconductor supply chain has raised concerns about the security of the fabricated ICs that are targeted for sensitive applications. One of these security threats is the adversarial infestation of fabricated ICs with a Hardware Trojan. A Hardware Trojan (HT) can be broadly described as a malicious modification to a circuit to control, modify, disable, or monitor its logic. Conventional VLSI manufacturing tests and verification methods fall short in detecting HT due to the different and un-modeled nature of these malicious modifications. Current state-of-the-art HT detection methods utilize statistical analysis of various side-channel information collected from ICs, such as power analysis, power supply transient analysis, regional supply current analysis, temperature analysis, wireless transmission power analysis, and delay analysis. To detect HTs, most methods require a trojan-free reference golden IC. A signature from these golden ICs is extracted and used to detect ICs with HTs. However, access to a golden IC is not always feasible. Thus, a novel mechanism for HT detection is sought that does not require the golden IC. Machine Learning (ML) approaches have emerged to be extremely useful to help eliminate the need for a golden IC. Recent works on utilizing ML for HT detection have been shown to be promising in achieving this goal. Thus, in this class, we will explain utilizing ML as a solution to the challenge of HT detection. Additionally, we will describe the Electronic Design Automation (EDA) tool flow for automating ML-assisted HT detection. Moreover, to further discuss the benefits of ML-assisted HT detection solutions, we will demonstrate a Neural Network (NN)-assisted timing profiling method for HT detection. Finally, we will discuss the shortcomings and open challenges of ML-assisted HT detection methods. The following topics will be covered in this class:

1. Hardware Trojan (HT) taxonomy and overview of state-of-the-art HT detection approaches
2. Application of Machine Learning (ML) in HT detection
3. Electronic Design Automation (EDA) flow for automating utilization of ML for HT detection
4. Demo of an ML-assisted HT detection approach and EDA tool flow
5. Open Challenges of ML-assisted HT detection

Bio: Houman Homayoun is currently an Associate Professor in the Department of Electrical and Computer Engineering at University of California, Davis. He is also the director of National Science Foundation Center for Hardware and Embedded Systems Security and Trust (CHEST). He conducts research in hardware security and trust, applied machine learning and AI, data-intensive computing and heterogeneous computing, where he has published more than 200 technical papers in the prestigious conferences and journals on the subject and directed over $8M in research funding from NSF, DARPA, AFRL, NIST, US Congress, and various industrial sponsors. His work received several best paper awards and nominations in various conferences including ACM GLSVLSI 2016, IEEE ICDM and ICCAD 2019, ISVLSI 2020, IEEE DCAS 2021. His CHEST center received congressional support for research in HW security which was included in 2021 National Defense Authorization Act. Houman served as Member of Advisory Committee, Cybersecurity Research and Technology Commercialization working group in the Commonwealth of Virginia. He is also serving as core group member of hardware security body of knowledge development team supported by the Department of Defense. He was a recipient of 2010 National Science Foundation computing innovation fellow award by CCC/CRA. Since 2017 he has been serving as an Associate Editor of IEEE Transactions on VLSI. He chaired and co-chaired major conferences in ACM including Great Lake Symposium on VLSI.

Education Class 10

High-Level Approaches to Hardware Security

Instructor:
Ramesh Karri, NYU

Abstract:
Designers use third-party intellectual property (IP) cores and outsource various steps in the integrated circuit (IC) design and manufacturing flow. As a result, security vulnerabilities have been rising. This is forcing IC designers and end users to re-evaluate their trust in ICs. If attackers get hold of an unprotected IC, they can reverse engineer the IC and pirate the IP. Similarly, if attackers get hold of a design, they can insert malicious circuits and backdoors into the design. Unintended design bugs can also result in security weaknesses.

The first part of the class will outline High-Level Design for Trust techniques to prevent these and similar attacks: Locking/Obfuscation and Secure Sourcing of IPs for High-Level Integration. Locking/Obfuscation implements a built-in obfuscation mechanism in ICs to prevent reverse engineering. Secure sourcing can thwart Trojan insertion in 3rd party *Intel*lectual Properties. The second part of the class will discuss hardware security bugs, focusing on the recent common weakness enumeration (CWE) list for hardware design. We will wrap up by pointing out why hardware security is an essential objective from economics, security, and safety aspects and offer a vision of the emerging directions in hardware cybersecurity.

Bio: Ramesh Karri is a Professor of Electrical and Computer Engineering at Tandon School of Engineering, *New York University*. He has a Ph.D. in Computer Science and Engineering, from the University of California at San Diego. His research and education activities span hardware cybersecurity including trustworthy ICs, processors and cyberphysical systems; security-aware computer aided design, test, verification, validation and reliability; nano meets security; metrics; benchmarks; hardware cybersecurity competitions; additive manufacturing security.

He has over 200 journal and conference publications including tutorials on Trustworthy Hardware in IEEE Computer (2) and Proceedings of the IEEE (5). His groups work on hardware cybersecurity was nominated for best paper awards (ICCD 2015 and DFTS 2015) and received awards at conferences (ITC 2014, CCS 2013, DFTS 2013 and VLSI Design 2012) and at competitions (ACM Student Research Competition at DAC 2012, ICCAD 2013, DAC 2014, ACM Grand Finals 2013, Kaspersky Challenge and Embedded Security Challenge).

He was the recipient of the Humboldt Fellowship and the National Science Foundation CAREER Award. He is the area director for cyber security of the NY State Center for Advanced Telecommunications Technologies at NYU-Poly; Co-founded the NYU Center for CyberSecurity -CCS (cyber.nyu.edu), co-founded the Trust-Hub (trust-hub.org/) and founded and organizes the Embedded Security Challenge, the annual red team blue team event at NYU, (www.nyu.edu/csaw2016/csaw-embedded).

He co-founded the IEEE/ACM Symposium on Nanoscale Architectures (NANOARCH). He served as program/general chair of conferences including IEEE International Conference on Computer Design (ICCD), IEEE Symposium on Hardware Oriented Security and Trust (HOST), IEEE Symposium on Defect and Fault Tolerant Nano VLSI Systems (DFTS) NANOARCH, RFIDSEC 2015 and WISEC 2015. He serves on several program committees (DAC, ICCAD, HOST, ITC, VTS, ETS, ICCD, DTIS, WIFS).

He was the Associate Editor of IEEE Transactions on Information Forensics and Security (2010-2014), IEEE Transactions on CAD (2014-present), ACM Journal of Emerging Computing Technologies (2007-present), ACM Transactions on Design Automation of Electronic Systems (2014-present), IEEE Access (2015-present), IEEE Transactions on Emerging Technologies in Computing (2015-present), IEEE Design and Test (2015-present) and IEEE Embedded Systems Letters (2016-present). He served as an IEEE Computer Society Distinguished Visitor (2013-2015). He is on the Executive Committee of IEEE/ACM Design Automation Conference initiating and leading the Security@DAC initiative (2014-2017). He has delivered invited keynotes, talks, and tutorials on Hardware Security and Trust (ESRF, DAC, DATE, VTS, ITC, ICCD, NATW, LATW, CROSSING etc).

PROGRAM

MONDAY, OCTOBER 10

CST - 10/10/2022 08:30-12:00 PST - 10/09/2022 17:30-21:00 EST - 10/09/2022 20:30-00:00 CET - 10/10/2022 02:30-06:00	**T3 – Tutorial on QuantumFlow+VACSEN: A Visualization System for Quantum Neural Networks on Noisy Quantum Device**		**T5 – Embedded Machine Learning: Design, Optimizations, and Applications**	
CST - 10/10/2022 14:00-14:30 PST - 10/09/2022 23:00-23:30 EST - 10/10/2022 02:00-02:30 CET - 10/10/2022 08:00-08:30	**ESWEEK Opening**			
CST - 10/10/2022 14:30-15:30 PST - 10/09/2022 23:30-00:30 EST - 10/10/2022 02:30-03:30 CET - 10/10/2022 08:30-09:30	**KEYNOTE** Keynote 1 – Blockchain, Big Data, and AI Empower High-Quality Development of Industrial Internet by Prof. Jie Li, Shanghai Jiao Tong University			
CST - 10/10/2022 22:00-23:00 PST - 10/10/2022 07:00-08:00 EST - 10/10/2022 10:00-11:00 CET - 10/10/2022 16:00-17:00	**CASES 1 – Storage and Emerging Memory Architectures**	**CODES+ISSS 1 - Multi-dimension mapping of embedded applications**	**EMSOFT 1 - Memory and Compilers Industry Session**	
CST - 10/10/2022 23:00-00:00 PST - 10/10/2022 08:00-09:00 EST - 10/10/2022 11:00-12:00 CET - 10/10/2022 17:00-18:00	**CASES 2 – Hardware Accelerators for Neural Networks**	**CODES+ISSS 2 - Embedded systems design**	**EMSOFT 2 - Safety and Security**	**Special Session 1 - Brain-Inspired Hyperdimensional Computing**
CST - 10/11/2022 00:00 - 00:30 PST - 10/10/2022 09:00 - 09:30 EST - 10/10/2022 12:00 - 12:30 CET - 10/10/2022 18:00 - 18:30	**Posters**	**Industry Booth**	**Yoga and Escape Room (GatherTown)**	
CST - 10/11/2022 00:30 - 01:00 PST - 10/10/2022 09:30 - 10:00 EST - 10/10/2022 12:30 - 13:00 CET - 10/10/2022 18:30 - 19:00	**ESWEEK Networking**	**Yoga and Escape Room (GatherTown)**		

KEYNOTES

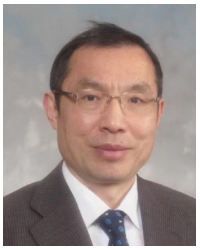

Keynote 1 – Blockchain, Big Data, and AI Empower High-Quality Development of Industrial Internet by Prof. Jie Li, Shanghai Jiao Tong University

Description

Big data, AI (Artificial *Intelligence*), and blockchain become essential for the cyber digital world. The industrial Internet is the use of trusted big data collected from smart sensors and actuators to enhance manufacturing and industrial processes with power of AI and real-time analytics. In this talk, we will overview big data, AI, blockchain, and Industrial Internet. We will address the application and challenge issues in applications of big data, AI, and blockchain for Industrial Internet.

Biography

Dr. Jie Li is a chair Professor in Department of Computer Science and Engineering, School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University (SJTU). His research interests are in big data and AI, blockchain, computer networking and security, and smart city. He is the director of SJTU Blockchain Research Centre. He is a co-chair of IEEE Big Data Community and chair of IEEE ComSoc Technical Committee on Big Data. He received the B.E. degree in computer science from Zhejiang University, Hangzhou, China, the M.E. degree in electronic engineering and communication systems from China Academy of Posts and Telecommunications, Beijing, China. He received the Dr. Eng. degree from the University of Electro-Communications, Tokyo, Japan. He was a full professor in University of Tsukuba, Japan. He is a senior member of IEEE and ACM. He has served on editorial boards of IEEE journals and transactions. He also has served in the program committees for several international conferences.

Tutorial T3

Tutorial on QuantumFlow+VACSEN: A Visualization System for Quantum Neural Networks on Noisy Quantum Device

Abstract: As one of the most popular machine learning algorithms, neural networks have been applied in a wide variety of applications, such as autonomous vehicles, simultaneous translation, and diagnostic medical imaging. With the increasing requirement on analyzing the large-scale data (e.g., 108 pixels for one 3D-CT medical image), neural networks encounter both memory-wall and compute-bound on classical computers. With the extremely high parallelism in representing and processing information, Quantum Computing is promising to address these limitations. But, how to make full use of the powerful quantum computers to accelerate neural networks is still unclear. QuantumFlow, published at Nature Communications last year, is an end-to-end framework to optimize neural networks onto a given quantum processor. Importantly, following the co-design philosophy, the developed quantum neurons in QuantumFlow demonstrate the quantum advantage. Meanwhile, VACSENis an online visualization system which provides the "easy to understand" visualization of the noise status on all available quantum computing nodes, recommends the most robust transpilation of circuit on the selected quantum computing node, and allows the real-time execution for a given quantum algorithm with noise awareness. In this tutorial, we will introduce how to conduct the co-design of neural networks and quantum circuits with QuantumFlow and VACSEN. We will have hands-on experience in implementing the neural network on the quantum circuit. Finally, targeting the near-term quantum computers, we will discuss how to leverage VACSEN to design quantum neural networks in the NISQ-Era. All attendees will leave with code examples that they can use as the backbone implementation to their own projects, and they will have access to VACSEN for the profiling of quantum devices.

Biographies of the Speakers

Weiwen Jiang: Dr. Jiang joined the ECE department at *George Mason University* as an Assistant Professor in Fall 2021. He was a Postdoctoral Associate at the University of Notre Dame. He received the Ph.D. degree from *Chongqing University* in 2019. From 2017 to 2019, he was a research scholar at the *University of Pittsburgh*. His research works have won Best Paper Awards in IEEE TCAD'21, ICCD'17, and NVMSA'15. He is the receipt of four Best Paper Nominations in ASP-DAC'16, DAC'19, CODES+ISSS'19, ASP-DAC'20, and the Top Winning Awards at IEEE Services Hackathon. He built the first co-design framework, QuantumFlow, to demonstrate the quantum advantage in designing neural network onto a quantum computer, which was published in Nature Communications. On the quantum topic, he was invited to give a contribution talk at IBM Quantum Summit 2020 and host tutorials at QuantumWeek-21, CODES+ISSS'21, ICCAD'21, and will host tutorials at DAC'22.

Qiang Guan: Dr. Guan is an assistant professor in Department of Computer Science at Kent State University, Kent, Ohio. Dr. Guan is the direct of Green Ubiquitous Autonomous Networking System lab (GUANS). He is also a member of Brain Health Research Institute (BHRI) at Kent State University. He was a computer scientist in Data Science at Scale team at Los Alamos National Laboratory before joining KSU. His current research interests include fault tolerance design for HPC applications; HPC-Cloud hybrid system; virtual reality; quantum computing systems and applications.

Yong Wang: is currently an assistant professor in School of Computing and Information Systems at Singapore Management University. He obtained his Ph.D. in Computer Science from *Hong Kong University of Science and Technology* in 2018. His research interests include information visualization, visual analytics and explainable machine learning. His research has won the Best Paper Award in IUI 2017, Best Paper Honorable Mention Awards in CHI 2022 and IEEE VIS 2021, and Best Poster Award in IEEE VIS 2019. For more details, please refer to http://yong-wang.org.

Tutorial T5

Embedded Machine Learning: Design, Optimizations, and Applications

Abstract: By integrating AI into small embedded systems, we can use the power of billions of devices that we already use in our lives without depending on extra costly equipment. We can build cheaper devices that adapt to our daily lives and have a high impact on how we deal with the environment around us. In this 3-hour tutorial, three speakers will cover the hardware/software co-design of accelerators, performance autotuning in AI chips, and some novel applications.

List of topics and speakers

Topic 1: Fundamentals of Machine Learning Algorithm and Acceleration Co-design and its Applications in Autonomous System (45 mins + 15 mins Q&A)

Abstract: High quality machine learning (ML) solutions require joint optimization of algorithms and their hardware accelerators. In the first talk, we introduce the fundamentals of both ML accelerator and algorithm design, which pave the road to efficient, practical, and system-level co-design in real applications. In particular, the talk is arranged as follows:

- Part 1 (10 – 12 min), we introduce the recipe and guidelines for high-performance ML accelerator design, such as data type and precision, parallelization, and memory system optimizations.
- Part 2 (10 – 12 min), we discuss the prevailing approaches for automated high-quality ML algorithm search, i.e., neural architecture search (NAS). Example approaches include reinforcement learning, differentiable, and evolutionary algorithms.
- Part 3 (5 min), we march into the joint area of accelerator and algorithm co-design and present the classic and representative approaches, and discuss its exciting applications in autonomous systems. More advanced techniques and applications will be discussed in the second talk.

Speaker: Dr. Cong (Callie) Hao is an Assistant Professor in the School of Electrical and Computer Engineering at Georgia Institute of Technology (GaTech), Atlanta, USA, and she currently holds the Sutterfield Family Early Career Professorship. She is the primary investigator of the software/hardware co-design lab (Sharc Lab) 1. She received her Ph.D. degree in Electrical Engineering from Waseda University, Japan in 2017, and her M.S. and B.S. degrees in Computer Science and Engineering from *Shanghai Jiao Tong University*. She was a postdoctoral researcher in ECE, GaTech, from September 2020 to August 2021, and a postdoctoral researcher in ECE at the University of Illinois at Urbana-Champaign (UIUC). Her primary research interests lie in the joint area of efficient hardware design and machine learning algorithms. She is passionate about reconfigurable and high-efficiency computing and building useful electronic design automation tools.

Topic 2: Practice on Performance Autotuning in AI Compute Chip (45 mins + 15 mins Q&A)

Abstract: Modern AI compute SOCs have abundant computation cores, interconnects & high-bandwidth memory (HBM) resources. When facing enormous different neural network architectures, how to program AI compute chips to make full use of the system resources becomes challenging. In this talk, I will share an automated code generation framework “Autotuning through design space pruning & schedule templates” to program a commercial AI compute chip to achieve improvements in both accelerator performance and coding efficiency.

Speaker: Dr. Peipei Zhou joined the *University of Pittsburgh*, ECE department as a Tenure-Track Assistant Professor starting September 2021. She obtained my Ph.D. in Computer Science from University of California, Los Angeles in 2019 supervised by Prof. Jason Cong, who leads *UCLA* VAST (VLSI Architecture, Synthesis and Technology) Group and CDSC (The Center for Domain-Specific Computing). Her major interest is in Customized Computer Architecture and Programming Abstraction for Applications including Healthcare, e.g., Precision Medicine and Artificial *Intelligence*. She has received “Outstanding Recognition in Research” from *UCLA* Samueli School of Engineering in 2019, 2019 TCAD Donald O. Pederson Best Paper Award, 2018 ICCAD Best Paper Nominee, and 2018 ISPASS Best Paper Nominee.

Topic 3: Towards Independent On-Device AI: Inference without Battery and Learning without Labels (45 mins + 15 mins Q&A)

Abstract: This talk consists of two parts: Inference without battery and learning without labels for on-device machine learning algorithms. Both directions are working towards an more independent ondevice AI. The

maturation of energy harvesting (EH) technology and the recent emergence of intermittent computing, which stores harvested energy in energy storage and supports an episode of program execution during each power cycle, creates the opportunity to build sophisticated battery-less energy-neutral sensors. By deploying lightweight DNNs onto EH-powered devices, persistent, eventdriven sensing and decision capabilities can be achieved. However, harvested energy is usually weak and unpredictable and even lightweight DNNs take multiple power cycles to finish one inference. To eliminate the indefinite long wait to accumulate energy for one inference and to optimize the accuracy, we developed a power trace-aware and exit-guided network compression algorithm to compress and deploy multi-exit neural networks to EH-powered microcontrollers (MCUs) and select exits during execution according to available energy.

Another challenge for on-device AI is how to automatically adapt to new environments without excessive interaction with human. After a model is deployed on edge devices, it is desirable to learn from unlabeled data to continuously improve accuracy. Contrastive learning has demonstrated its great potential in learning from unlabeled data. However, the online input data are usually none independent and identically distributed (non-iid) and edge devices' storages are usually too limited to store enough representative data from different data classes. We developed a framework to automatically select the most representative data from the unlabeled input stream, which only requires a small data buffer for dynamic learning.

Speaker: Dr. Jingtong Hu is currently an Associate Professor and William Kepler Whiteford Faculty Fellow in the Department of Electrical and Computer Engineering at *University of Pittsburgh*, Pittsburgh, PA, USA. Before that, he was an Assistant Professor at Oklahoma State University from 2013 to 2017. He received his Ph.D. in Computer Science from University of Texas at Dallas in 2013 and his B.E. in Computer Science and Technology from *Shandong University*, China in 2007. His current research interests include hardware/software co-design for machine learning algorithms, on-device AI, and embedded systems. His works have received Donald O. Pederson Best Paper Award from IEEE Transactions on Computer-Aided Design of Circuits and Systems in 2021 and 5 best paper nominations from DAC, ASP-DAC, and ESWEEK, etc. He is also the recipient of Oklahoma State University Outstanding New Faculty Award, Air Force Summer Faculty Fellowship, and ACM SIGDA Meritorious Service Award. He has served as TPC track chair for CASES, GLSVLSI, ASP-DAC, DAC, SAC, and on the TPC of many other international conferences such as DATE, ESWEEK, CPS-IoT Week, etc. He served as a guest editor for Sensors, IEEE Transactions on Computers, ACM Transactions on Cyber-Physical Systems, and is currently serving as executive committee member and education chair for ACM SIGDA, associate editor for IEEE Embedded Systems Letters, the Journal of Systems Architecture: Embedded Software Design, and ACM Transactions on Cyber-Physical Systems.

CASES 1 Storage and Emerging Memory Architectures
Session Chair: Per Gunnar Kjeldsberg (NTNU)

Adaptive Switch on Wear Leveling for Enhancing I/O Latency and Lifetime of High-Density SSDs
Jiaojiao Wu (Southwest University); Jun Li (Southwest University); Zhibing Sha (Southwest University); Zhigang Cai (Southwest University); Jianwei Liao (Southwest University)

[Best paper candidate] Horae: A Hybrid I/O Request Scheduling Technique for Near-Data Processing-Based SSD
Jiali Li (Chongqing University); Xianzhang Chen (Chongqing University); Duo Liu (Chongqing University); Lin Li (Chongqing University); Jiapin Wang (Chongqing University); Zhaoyang Zeng (Chongqing University); Yujuan Tan (Chongqing University); Lei Qiao (Beijing Institute of Control Engineering)

Hybrid RRAM/SRAM In-Memory Computing for Robust DNN Acceleration
Gokul Krishnan (Arizona State University); Zhenyu Wang (Arizona State University); Injune Yeo (Arizona State University); Li Yang (Arizona State University); Jian Meng (Arizona State University); Maximilian Liehr (State University of New York Polytechnic Institute); Rajiv Joshi (IBM); Nathaniel C. Cady (State University of New York Polytechnic Institute); Deliang Fan (Arizona State University); Jae-Sun Seo (Arizona State University); Yu Cao (Arizona State University)

Work-in-Progress: Efficient Low-latency Near-Memory Addition
Alexander Reaugh (University of Kentucky); Sayed Ahmad Salehi (University of Kentucky)

Work-in-Progress: A Processing-in-Pixel Accelerator based on Multi-level HfOx ReRAM
Minhaz Abedin (SUNY Polytechnic Institute); Arman Roohi (University of Nebraska–Lincoln); Nathaniel Cady (SUNY Polytechnic Institute); Shaahin Angizi (New Jersey Institute of Technology)

Work-in-Progress: DRAM-MaUT:- DRAM Address Mapping Unveiling Tool for ARM Devices
Anandpreet Kaur (IIIT Allahabad); Pravin Srivastav (IIIT Allahabad); Bibhas Ghoshal (IIIT Allahabad)

CODES+ISSS 1 Multi-Dimension Mapping of Embedded Applications
Session Chairs: Wanli Chang (Hunan University); Fei Wu (Huazhong University of Science and Technology)

An Efficient BCNN Deployment Method Using Quality-Aware Approximate Computing
Bo Liu (Southeast University); Hao Cai (Southeast University); Ziyu Wang (Southeast University); Xuetao Wang (Southeast University); Renyuan Zhang (Southeast University); Anfeng Xue (Southeast University); Qiao Shen (Southeast University); Na Xie (Southeast University); Yu Gong (Southeast University); Zhen Wang (Nanjing Prochip Electronic Technology Co. Ltd); Jun Yang (Southeast University); Hao Cai (Southeast University)

Memory-Computing Decoupling: A DNN Multitasking Accelerator With Adaptive Data Arrangement
Chuxi Li (Northwestern Polytechnical University); Xiaoya Fan (Northwestern Polytechnical University); Xiaoti Wu (Northwestern Polytechnical University); Zhao Yang (Northwestern Polytechnical University); Miao Wang (Northwestern Polytechnical University); Meng Zhang (Northwestern Polytechnical University); Shengbing Zhang (Northwestern Polytechnical University)

[Best paper candidate] NeuroMap: Efficient Task Mapping of Deep Neural Networks for Dynamic Thermal Management in High-Bandwidth Memory
Shailja Pandey (Indian Institute of Technology Delhi); Preeti Ranjan Panda (Indian Institute of Technology Delhi)

On Transferring Application Mapping Knowledge Between Differing MPSoC Architectures
Jan Spieck (Friedrich-Alexander-Universität Erlangen-Nürnberg); Stefan Wildermann (Friedrich-Alexander-Universität Erlangen-Nürnberg); Jürgen Teich (Friedrich-Alexander-Universität Erlangen-Nürnberg)

TECHNICAL SESSIONS

EMSOFT 1 Memory and Compilers
Session Chair: Heiko Falk (Hamburg University of Technology)

iNVMFS: An Efficient File System for NVRAM-Based Intermittent Computing Devices
Ying-Jan Wu (National Yang Ming Chiao Tung University); Ching-Yu Kuo (National Yang Ming Chiao Tung University); Li-Pin Chang (National Yang Ming Chiao Tung University)

Mercury: Instruction Pipeline Aware Code Generation for Simulink Models
Zehong Yu (Tsinghua University); Zhuo Su (Tsinghua University); Yixiao Yang (Capital Normal University); Jie Liang (Tsinghua University); Yu Jiang (Tsinghua University); Aiguo Cui (Huawei Technologies Co. Ltd.); Wanli Chang (Hunan University); Rui Wang (Capital Normal University);

An I/O Virtualization Framework With I/O-Related Memory Contention Control for Real-Time Systems
Niccolò Borgioli (Scuola Superiore Sant'Anna); Matteo Zini (Scuola Superiore Sant'Anna); Daniel Casini (Scuola Superiore Sant'Anna); Giorgiomaria Cicero (Scuola Superiore Sant'Anna); Alessandro Biondi (Scuola Superiore Sant'Anna); Giorgio Buttazzo (Scuola Superiore Sant'Anna);

Toward Register Spilling Security Using LLVM and ARM Pointer Authentication
Andrea Fanti (Politecnico di Torino); Carlos Chinea Perez (Huawei Technologies Oy Co. Ltd); Remi Denis-Courmont (Huawei Technologies Oy Co. Ltd); Gianluca Roascio (Politecnico di Torino); Jan-Erik Ekberg (Huawei Technologies Oy Co. Ltd);

Industry Session
Session Chair: Ziegenbein Dirk (Bosch Research)

Industry Paper: Towards Agile Design of Neural Processing Unit
Binyi Wu (Technische Universität Dresden & Infineon Technologies AG Dresden); Wolfgang Furtner (Infineon Technologies AG München); Bernd Waschneck (Infineon Technologies AG Dresden); Christian Georg Mayr (Infineon Technologies AG Dresden)

Industry Paper: Challenges in rebooting autonomy with deep learned perception
Michael Abraham (The Boeing Company); Aaron Mayne (The Boeing Company); Tristan Perez (The Boeing Company); *Chiao Hsieh (University of Illinois)*; Yangge Li (University of Illinois at Urbana-Champaign UIUC); Italo Romani De Oliveira (The Boeing Company); Huafeng Yu (The Boeing Company) Dawei Sun (University of Illinois); Sayan Mitra (University of Illinois)

Industry Paper: Surrogate Models for Testing Analog Designs under Limited Budget - A Bandgap Case Study
Roderick Bloem (Graz University of Technology); *Alberto Larrauri (Graz University of Technology)*; Roland Lengfeldner (Infineon Technologies); Cristinel Mateis (AIT Austrian Institute of Technology); Dejan Nickovic (AIT Austrian Institute of Technology); Björn Ziegler (AIT Austrian Institute of Technology)

Industry Paper: System-Level Logical Execution Time for Automotive Software Development
Kai-Björn Gemlau (TU Braunschweig); Hermann v. Hasseln (Mercedes-Benz AG); Rolf Ernst (TU Braunschweig)

CASES 2 Hardware Accelerators for Neural Networks
Session Chairs: Prabuddha Chakraborty (University of Maine), Wanli Chang (University of York)

Photonic Reconfigurable Accelerators for Efficient Inference of CNNs With Mixed-Sized Tensors
Sairam Sri Vatsavai (University of Kentucky); Ishan Thakkar (University of Kentucky)

TECHNICAL SESSIONS

CMQ: Crossbar-Aware Neural Network Mixed-Precision Quantization via Differentiable Architecture Search
Jie Peng (National University of Defense Technology); Haijun Liu (National University of Defense Technology); Zhongjin Zhao (National University of Defense Technology); Zhiwei Li (National University of Defense Technology); Sen Liu (National University of Defense Technology; Qingjiang Li (National University of Defense Technology)

[Best paper candidate] SASCHA - Sparsity-Aware Stochastic Computing Hardware Architecture for Neural Network Acceleration
Wojciech Romaszkan (University of California at Los Angeles); Tianmu Li (University of California at Los Angeles); Puneet Gupta (University of California at Los Angeles)

Work-in-Progress: Ultra-fast yet Accurate Performance Prediction for Deep Neural Network Accelerators
Konstantin Lübeck (University of Tübingen); Alexander Louis-Ferdinand Jung (University of Tübingen); Felix Wedlich (University of Tübingen); Oliver Bringmann (University of Tübingen)

Work-in-Progress: SuperNAS: Fast Multi-Objective SuperNet Architecture Search for Semantic Segmentation
Marihan Amein (McGill University); *Zhuoran Xiong (McGill University)*; Olivier Therrien (McGill University); Brett H. Meyer (McGill University); Warren J. Gross (McGill University)

CODES+ISSS 2 Embedded Systems Design
Session Chairs: Peipei Zhou (University of Pittsbugh); Weiwen Jiang (George Mason University)

Amphis: Managing Reconfigurable Processor Architectures With Generative Adversarial Learning
Weiwei Chen (Institute of Computing Technology, Chinese Academy of Sciences); Ying Wang (Institute of Computing Technology, Chinese Academy of Sciences); Ying Xu (Institute of Computing Technology, Chinese Academy of Sciences); Chengsi Gao (Institute of Computing Technology, Chinese Academy of Sciences); Yinhe Han (Institute of Computing Technology, Chinese Academy of Sciences); Lei Zhang (Institute of Computing Technology, Chinese Academy of Sciences)

Adaptive Mode Transformation for Wear Leveling in Nonvolatile FPGAs
Huichuan Zheng (Shandong University); Hao Zhang (Shandong University); Shuo Xu (Shandong University); Fanjin Xu (Shandong University); Mengying Zhao (Shandong University)

FRL: Fast and Reconfigurable Accelerator for Distributed Sound Source Localization
Xiaofeng Ding (Chong Qing University); Chengliang Wang (Chongqing University); Heping Liu (Chong Qing University); Zhihai Zhang (Chongqing University); Xianzhang Chen (Chongqing University); Yujuan Tan (Chongqing University); Duo Liu (Chongqing University); Ao Ren (Chongqing University)

HARM: A Hint-Based Assertion Miner
Samuele Germiniani (University of Verona); Graziano Pravadelli (University of Verona)

EMSOFT 2 Safety and Security
Session Chair: Hokeun Kim (Hanyang University)

Safety Analysis of Embedded Controllers Under Implementation Platform Timing Uncertainties
Clara Hobbs (University of North Carolina at Chapel Hill); Bineet Ghosh (University of North Carolina at Chapel Hill); Shengjie Xu (University of North Carolina at Chapel Hill); Parasara Sridhar Duggirala (University of North Carolina at Chapel Hill); Samarjit Chakraborty (University of North Carolina at Chapel Hill)

[Best paper candidate] Verified Train Controllers for the Federal Railroad Administration Train Kinematics Model: Balancing Competing Brake and Track Forces
Aditi Kabra (Carnegie Mellon University); Stefan Mitsch (Carnegie Mellon University); André Platzer (Carnegie Mellon University)

Verifying Controllers With Vision-Based Perception Using Safe Approximate Abstractions
Chiao Hsieh (University of Illinois at Urbana-Champaign); Yangge Li (University of Illinois at Urbana-Champaign); Dawei Sun (University of Illinois at Urbana-Champaign); Keyur Joshi (University of Illinois at Urbana-Champaign); Sasa Misailovic (University of Illinois at Urbana-Champaign); Sayan Mitra (University of Illinois at Urbana Champaign)

NASA: NVM-Assisted Secure Deletion for Flash Memory
Weidong Zhu (University of Florida); Kevin R. B. Butler (University of Florida)

CASES 3 Accelerators and Application Specific Design
Session Chairs: Partha Pande (Washington State University); Jingtong Hu (University of Pittsburgh)

CASPHAr: Cache-Managed Accelerator Staging and Pipelining in Heterogeneous System Architectures
Mochamad Asri (The University of Texas at Austin); Andreas Gerstlauer (The University of Texas at Austin)

[Best paper candidate] SWAP: A Server-Scale Communication-Aware Chiplet-Based Manycore PIM Accelerator
Harsh Sharma (Washington State University); Sumit K. Mandal (University of Wisconsin-Madison); Janardhan Rao Doppa (Washington State University); Umit Y. Ogras (University of Wisconsin - Madison); Partha Pratim Pande (Washington State University)

Energy-Efficient DNN Inference on Approximate Accelerators Through Formal Property Exploration
Ourania Spantidi (Southern Illinois University Carbondale); Georgios Zervakis (Karlsruhe Institute of Technology); Iraklis Anagnostopoulos (Southern Illinois University Carbondale); Joerg Henkel (Karlsruhe Institute of Technology)

Work-in-Progress: Toward a Robust, Reconfigurable Hardware Accelerator for Tree-Based Genetic Programming
Christopher Crary (University of Florida); Wesley Piard (University of Florida); Britton Chesley (University of Florida); Greg Stitt (University of Florida)

Work-in-Progress: Smart data reduction in SLAM methods for embedded systems
Quentin Picard (CEA, LIST); Stephane Chevobbe (CEA, LIST); Mehdi Darouich (CEA, LIST); Zoe Mandelli (CEA, LIST); Mathieu Carrier (CEA, LIST); Jean-Yves Didier (IBISC, Univ Evry, Université Paris-Saclay)

Work-in-Progress: NoRF: A Case Against Register File Operands in Tightly-Coupled Accelerators
David J. Schlais (University of Wisconsin-Madison); Heng Zhuo (University of Wisconsin-Madison); Mikko H. Lipasti (University of Wisconsin-Madison)

CODES+ISSS 3 Machine Learning in the Age of IoT
Session Chairs: Xiang Chen (George Mason University); Xun Jiao (Villanova University)

PervasiveFL: Pervasive Federated Learning for Heterogeneous IoT Systems
Jun Xia (East China Normal University); Tian Liu (East China Normal University); Zhiwei Ling (East China Normal University); Ting Wang (East China Normal University); Xin Fu (University of Houston); Mingsong Chen (East China Normal University)

Optimizing Random Forest-Based Inference on RISC-V MCUs at the Extreme Edge
Enrico Tabanelli (University of Bologna); Giuseppe Tagliavini (University of Bologna); Luca Benini (University of Bologna and ETH Zurich)

Enabling Weakly Supervised Temporal Action Localization From On-Device Learning of the Video Stream
Yue Tang (University of Pittsburgh); Yawen Wu (University of Pittsburgh); Peipei Zhou (University of Pittsburgh); Jingtong Hu (University of Pittsburgh)

Work-in-Progress: What to Expect of Early Training Statistics? An Investigation on Hardware-Aware Neural Architecture Search
Xiangzhong Luo (Nanyang Technological University); Di Liu (Yunnan University); Hao Kong (Nanyang Technological University); Shuo Huai (Nanyang Technological University); Hui Chen (Nanyang Technological University); Weichen Liu (Nanyang Technological University)

Work-in-Progress: Utilizing latency and accuracy predictors for efficient hardware-aware NAS
Negin Firouzian (McGill University); Seyyed Hasan Mozafari (McGill University); James J. Clark (McGill University); Warren Gross (McGill University); Brett H. Meyer (McGill University)

Work-in-Progress: BloCirNN: An Efficient Software/hardware Codesign Approach for Neural Network Accelerators with Block-Circulant Matrix
Yunji Qin (University of Science and Technology of China); Lei Gong (University of Science and Technology of China); Zhendong Zheng (University of Science and Technology of China); Chao Wang (University of Science and Technology of China)

EMSOFT 3 Machine Learning, Networks, and IoT
Session Chair: Chih-Hong Cheng (Fraunhofer IKS)

Adaptive Edge Offloading for Image Classification Under Rate Limit
Jiaming Qiu (Washington University in Saint Louis); Ruiqi Wang (Washington University in Saint Louis); Ayan Chakrabarti (Google Research); Roch Guérin (Washington University in Saint Louis); Chenyang Lu (Washington University in Saint Louis);

Online Rerouting and Rescheduling of Time-Triggered Flows for Fault Tolerance in Time-Sensitive Networking
Zhiwei Feng (Northeastern University); Zonghua Gu (Umeå University); Haichuan Yu (Northeastern University); Qingxu Deng (Northeastern University); Linwei Niu (Howard University)

NExG: Provable and Guided State-Space Exploration of Neural Network Control Systems Using Sensitivity Approximation
Manish Goyal (University of North Carolina at Chapel HIll); Miheer Dewaskar (Duke University); Parasara Sridhar Duggirala (Universty of North Carolina at Chapel Hill)

[Best paper candidate] Tinkertoy: Build Your Own Operating Systems for IoT Devices
Bingyao Wang (University of British Columbia); Margo Seltzer (University of British Columbia)

Special Session Brain-Inspired Hyperdimensional Computing for Ultra-Efficient Edge
Organizer(s)
Mohsen Imani (UC Irvine); Xun Jiao (Villanova University); Hussam Amrouch (University of Stuttgart); Yiannis Aloimonos (University of Marylan)

Common Abstract
Hyperdimensional Computing (HDC) is rapidly emerging as an attractive alternative to traditional deep learning algorithms. Despite the profound success of DNNs in many domains, the amount of computational power and storage that they demand during training makes deploying them in edge devices very challenging if not infeasible. This, in turn, inevitably necessitates steaming the data from the edge to the cloud which raises serious concerns when it comes to scalability, security, and privacy. Further, the nature of data that edge devices often receive from sensors is inherently noisy. However, Deep Neural Network (DNN) algorithms are very sensitive to noise, which makes accomplishing the required learning tasks with high accuracy immensely difficult. In this special session, we aim at providing a comprehensive overview of the latest advances in HDC. HDC aims at realizing real-time performance and robustness through using strategies that more closely model the human brain. HDC is, in fact, motivated by the observation that the human brain operates on high-dimensional data representations. In HDC, objects are thereby encoded with high-

dimensional vectors which have thousands of elements. In this special session, we will discuss the promising robustness of HDC algorithms against noise along with the ability to learn from little data. Further, we will present the outstanding synergy between HDC and beyond von Neumann architectures and how HDC opens doors for efficient learning at the edge due to the ultra-lightweight implementation that it needs, contrary to traditional DNNs.

Talk 1: Mohsen Imani, UC Irvine - Hyperdimensional Computing Platform for Robust and Efficient Perception and Decision Making at the Edge

Abstract
HDC has several advantages over competing learning solutions towards efficient and robust computation. We exploit HDC to develop a hardware computing platform that dynamically learn in the field and at the edge. Our general platform does not rely on task-specific accelerators or large-scale data centers. Instead, it can enable online learning from limited data, handling noise and uncertainty, and intelligent decision-making at the edge. Our HDC-based algorithm is implemented with large vectors, simplifying the "compute" in data-centric applications from linear algebra to simple, low precision, and lower power array operations that can be highly parallelized. Our edge-based hardware platform should support iso-accuracy with state-of-the-art machine ML/AI solutions while enabling (1) at least 100× faster and 1000× energy efficiency than DNNs for online perception at the edge and (2) significant robustness to lack of data or noise for intelligent edge-based decision making.

Biography
Mohsen Imani is an Assistant Professor in the Department of Computer Science at UC Irvine. He is also a director of Bio-Inspired Architecture and Systems Laboratory. He is working on a wide range of practical problems in the area of brain-inspired computing, ML, computer architecture, and embedded systems. His research goal is to design real-time, robust, and programmable computing platforms that can natively support a wide range of learning and cognitive tasks on edge devices. He received his Ph.D. from UC San Diego, and he has a stellar record of publication with over 120 papers in top conferences/journals. His contribution has led to a new direction on brain-inspired HDC that enables ultra-efficient and real-time learning and cognitive support. Dr. Imani research has been recognized with several awards, including the Bernard and Sophia Gordon Engineering Leadership Award, the Outstanding Researcher Award, and the Powell Fellowship Award. He also received the Best Doctorate Research from UCSD, the best paper award in DATE'22, and several best paper nomination awards at multiple top conferences including DAC'19, DAC'20, DATE'20, and ICCAD'20.

Talk 2: Yiannis Aloimonos, University of Maryland, College Park - Hyper-dimensional Active Perception

Abstract
Action and perception are often kept in separated spaces, which is a consequence of traditional vision being frame based and only existing in the moment and motion being a continuous entity. This bridge is crossed by the dynamic vision sensor (DVS), a neuromorphic camera that can see the motion. We propose a method of encoding actions and perceptions together into a single space that is meaningful, semantically informed, and consistent by using hyper-dimensional binary vectors (HBVs). We show that the visual component can be bound with the system velocity to enable dynamic world perception, which creates an opportunity for real-time navigation and obstacle avoidance with active perception. Actions performed by an agent are directly bound to the perceptions experienced to form its own "memory." Furthermore, because HBVs can encode entire histories of actions and perceptions— from atomic to arbitrary sequences—as constant-sized vectors, auto-associative memory was combined with deep learning paradigms for controls. We demonstrate these properties on a quadcopter drone ego-motion inference task and the MVSEC (multivehicle stereo event camera) dataset.

Biography
Yiannis Aloimonos is Professor of Computational Vision and Intelligence at the Department of Computer Science, University of Maryland, College Park, and the Director of the Computer Vision Laboratory at the Institute for Advanced Computer Studies (UMIACS). He is also affiliated with the Institute for Systems Research and the Neural and Cognitive Science Program. He was

born in Sparta, Greece and studied Mathematics in Athens and Computer Science at the University of Rochester, NY (PhD 1990). He is interested in Active Perception and the modeling of vision as an active, dynamic process for real time robotic systems. For the past five years he has been working on bridging signals and symbols, specifically on the relationship of vision to reasoning, action and language using Hyper-dimensional Computing.

Talk 3: Xun Jiao, Villanova University - Robust Hyperdimensional Computing Against Hardware Errors and Cyber Attacks

Abstract

Hyperdimensional computing (HDC), also known as vector-symbolic architectures (VSA), was recently introduced as an emerging AI method mimicking the "human brain" at the functionality level. Currently HDC research largely has two focuses: applications of HDC to different domains and hardware acceleration of HDC.

However, increasing deployment of AI methods in safety-critical systems, such as healthcare and robotics, means it is not only important to strive for high accuracy of AI models, but also to ensure its robustness under even highly uncertain and adversarial environments. In this talk, we will present our recent work in developing robust hyperdimensional computing against uncertainties in both hardware and data. Specifically, we will introduce two studies, one focuses on enhancing the robustness of HDC to voltage-induced hardware errors, while the other focuses on the adversarial testing of HDC.

Biography

Xun Jiao has been an assistant professor in ECE department of Villanova University since 2018. He obtained his Ph.D. degree from the University of California, San Diego in 2018. His research interests lie in the broad areas of embedded systems, design automation, bio-inspired computing, and machine learning. He received 6 best paper award/ nominations in international conferences such as DATE, EMSOFT, DSD, and CPSCOM. He is an associate editor in of IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (IEEE TCAD), Program Committee member of DAC, GLSVLSI, LCTES, ICESS, COINS.

Talk 4: Hussam Amrouch, University of Stuttgart, Germany - HW/SW Codesign for Efficient Brain-inspired Hyperdimensional In-Memory Computing

Abstract

In this part, we will discuss how specialized hardware accelerators beyond von-Neumann architectures, that offer processing capability in where the data resides without moving it, become more and more indispensable to implement efficient edge AI. This is, in fact, indispensable to overcome the famous von-Neumann bottleneck. We will demonstrate how the emerging Ferroelectric transistor (FeFET) technology has a great potential to realize novel beyond von-Neumann architectures that outstandingly synergize with hyperdimensional computing. We will focus on demonstrating how HW/ SW codesign is a key to build reliable HDC on unreliable beyond-CMOS transistors and how abstracted, yet accurate reliability models can be developed in a cross-layer manner bridging the gap between the underlying technology and the HDC algorithms running on top of such novel architectures.

Biography

Hussam Amrouch is a Junior Professor heading the Chair of Semiconductor Test and Reliability (STAR) at the University of Stuttgar, Germany. He received his Ph.D. degree with the highest distinction (Summa cum laude) from KIT, Germany in 2015. He serves as an Editor in Nature Scientific Reports. His main research interests are brain-inspired computing, AI processor design, impact of emerging technologies on future computing, and ML for CAD. He holds eight HiPEAC Paper Awards and three best paper nominations at top EDA conferences: DAC'16, DAC'17 and DATE'17. He has 160+ publications including 65 journals in multidisciplinary research areas across the computing stack, starting from semiconductor physics to circuit design all the way up to computer architectures. He has given so far 25+ invited talks in many international conferences and leading companies (e.g., Synopsys, Advantest, Silvaco) as well as 10+ tutorials in many top EDA conferences like DAC, ICCAD, ICCAD, DATE and others.

PROGRAM

TUESDAY, OCTOBER 11

CST - 10/11/2022 08:30-09:15 PST - 10/10/2022 17:30-18:15 EST - 10/10/2022 20:30-21:15 CET - 10/11/2022 02:30-03:15	Sky talk 1 - AI for EDA, by Dr. Yu Huang HiSilicon, Huawei			
CST - 10/11/2022 09:15-12:00 PST - 10/10/2022 18:15-21:00 EST - 10/10/2022 21:15-00:00 CET - 10/11/2022 03:15-06:00	Workshop on Memory and Storage Computing (MSC)	Workshop on Edge Inteliigent Computing (EIC)	Workshop on Emerging Techniques in System Design and Design Automation (EEDA)	
CST - 10/11/2022 14:00-15:00 PST - 10/10/2022 23:00-00:00 EST - 10/11/2022 02:00-03:00 CET - 10/11/2022 08:00-09:00	CASES 4 - Security from Transistors to Compilers	CODES+ISSS 4 - Securing Embedded Hardware	EMSOFT 4 - Theory and Security	Work-in-Progress and Industry Pitch Session
CST - 10/11/2022 15:00-16:00 PST - 10/11/2022 00:00-01:00 EST - 10/11/2022 03:00-04:00 CET - 10/11/2022 09:00-10:00	CASES 5 - Security from Transistors to Compilers	CODES+ISSS 4 - Securing Embedded Hardware	EMSOFT 4 - Theory and Security	
CST - 10/11/2022 16:00-16:30 PST - 10/11/2022 01:00-01:30 EST - 10/11/2022 04:00-04:30 CET - 10/11/2022 10:00-10:30	Posters	Industry Booth	Yoga and Escape Room (GatherTown)	
CST - 10/11/2022 16:30-17:00 PST - 10/11/2022 01:30-02:00 EST - 10/11/2022 04:30-05:00 CET - 10/11/2022 10:30-11:00	ESWEEK Networking	Industry Interaction	Yoga and Escape Room (GatherTown)	
CST - 10/11/2022 22:00-22:30 PST - 10/11/2022 07:00-07:30 EST - 10/11/2022 10:00-10:30 CET - 10/11/2022 16:00-16:30	Test of Time Awards			
CST - 10/11/2022 22:30-23:30 PST - 10/11/2022 07:30-08:30 EST - 10/11/2022 10:30-11:30 CET - 10/11/2022 16:30-17:30	Keynote 2 - Training the world's best Gran Turismo racer, by Dr. Pete Wurman, Director, Sony AI, US			

PROGRAM

TUESDAY, OCTOBER 11

CST - 10/11/2022 23:30-00:30 PST - 10/11/2022 08:30-09:30 EST - 10/11/2022 11:30-12:30 CET - 10/11/2022 17:30-18:30	**CASES 6 - System Level Design: IoT, Reliability, and Verification**	**CODES+ISSS 6 - Reliable and Secure Embedded Systems**	**EMSOFT 6 - Power, Energy and Analysis**	**Special Session 2 - Programming Autonomous Machines**
CST - 10/12/2022 00:30-01:00 PST - 10/11/2022 09:30-10:00 EST - 10/11/2022 12:30-13:00 CET - 10/11/2022 18:30-19:00	**Posters**	**Industry Booth**	**Yoga and Escape Room (GatherTown)**	

KEYNOTES

Keynote 2 – Training the world's best Gran Tursimo racer, by Dr. Pete Wurman, Director, Sony AI, US

Description

Automobile racing represents an extreme example of real-time decision making in complex physical environments. Drivers must execute complex tactical maneuvers to pass or block opponents while operating their vehicles at their traction limits. Modern racing simulations, such as the PlayStation game Gran Turismo, faithfully reproduce much of the nonlinear control challenges of real race cars while also encapsulating the complex multi-agent interactions. In this talk I will describe how our team at Sony AI trained agents for Gran Turismo that can compete with the world's best e-sports drivers. We combine state-of-the-art model-free deep reinforcement learning algorithms with mixed scenario training to learn an integrated control policy that combines exceptional speed with impressive tactics. In addition, we construct a reward function that enables the agent to be competitive while adhering to racing's important, but under-specified, sportsmanship rules.
We demonstrate the capabilities of our agent, Gran Turismo Sophy, by winning a head-to-head competition against four of the world's best Gran Turismo drivers.

Biography

Pete Wurman is the director of Sony AI America and project lead for the Gran Turismo Sophy project. He has done research in computational auctions, multi-agent systems, robotics, and reinforcement learning. He earned his Ph.D. at the University of Michigan and was a professor at North Carolina State University. He was the Co-founder and CTO of Kiva Systems, the company that pioneered the use of mobile robotics in warehouse fulfillment operations. Kiva was acquired by Amazon in 2012 and, as Amazon Robotics, has deployed more than 500,000 robots to warehouses around the world. While at Amazon, Pete founded the Amazon Picking Challenge. He is an IEEE Fellow and a AAAI Fellow and was inducted into the National Inventors Hall of Fame in 2022.

Sky talk 1 – AI for EDA, by Dr. Yu Huang, HiSilicon, Huawei

Description

Electronic Design Automation (EDA) is critical for designing and manufacturing integrated circuits. Recent advancements of AI technologies can help improving the traditional EDA technology and transferring the IC design experiences from old designs to new designs or from old technologies to the new ones. In this talk, we will take a look of a few examples that apply AI to EDA and demonstrate the advantages of such technologies.

Biography

Dr. Yu Huang is Semiconductor Scientist at Huawei Technologies Co., and EDA Chief Architect of HiSilicon. Before Joining HiSilicon, he was Sr. Key Expert of Mentor Graphics (now Siemens EDA). His research interests include VLSI SoC testing, ATPG, compression, diagnosis, yield analysis and machine learning. He got his Ph.D. in electrical and computer engineering from the University of Iowa, USA. He has more than 70 patents and published more than 140 papers on leading IEEE Journals, conferences and workshops. He is a senior member of the IEEE. He is also an adjunct professor at School of Microelectronics, Fudan University, China.

TECHNICAL SESSIONS

CASES 4 Security: from Transistors to Compilers
Session Chair: Rajshekar K (IIT Dharwad)

Cut and Forward: Safe and Secure Communication for FPGA System on Chips
Francesco Restuccia (University of California at San Diego); Ryan Kastner (University of California at San Diego)

Quantifying Information Leakage for Security Verification of Compiler Optimizations
Priyanka Panigrahi (Indian Institute of Technology Guwahati); Abhik Paul (Indian Institute of Technology Guwahati); Chandan Karfa (Indian Institute of Technology Guwahati)

A Novel Attack Mode on Advanced Technology Nodes Exploiting Transistor Self-Heating
Nikhil Rangarajan (New York University Abu Dhabi); *Johann Knechtel (New York University Abu Dhabi)*; Nimisha Limaye (New York University); Ozgur Sinanoglu (New York University Abu Dhabi); Hussam Amrouch (University of Stuttgart)

Work-in-Progress: Reliability Evaluation of Power SCADA System with Three-Layer IDS
Yenan Chen (Shanghai Jiao Tong University); Linsen Li (Shanghai Jiao Tong University); Zhaoqian Zhu (Shanghai Jiao Tong University); Yue Wu (Shanghai Jiao Tong University)

Work-in-Progress: On Evaluation of On-chip Thermal Covert Channel Attacks
Jiachen Wang (South China University of Technology); *Xiaohang Wang (South China University of Technology)*; Yingtao Jiang (University of Nevada at Las Vegas); Amit Kumar Singh (University of Essex); Letian Huang (University of Electronic Science and Technology of China); Mei Yang (University of Nevada at Las Vegas)

Work-in-Progress: Towards a Smaller than Grain Stream Cipher: Optimized FPGA Implementations of Fruit-80
Gangqiang Yang (Shandong University); Zhengyuan Shi (Shandong University); *Cheng Chen (Shandong University)*; Hailiang Xiong (Shandong University), Honggang Hu (University of Science and Technology of China); Zhiguo Wan (Zhejiang Lab); Keke Gai (Beijing Institute of Technology); Meikang Qiu (Texas A&M University Commerce)

CODES+ISSS 4 Securing Embedded Hardware
Session Chairs: Chen Liu (Intel); Wei Jiang (University of Electronic Science and Technology of China)

[Best paper candidate] FLAM-PUF: A Response Feedback-Based Lightweight Anti-Machine-Learning-Attack PUF
Linjun Wu (Hunan University); Yupeng Hu (Hunan University); Kehuan Zhang (The Chinese University of Hong Kong); Wenjia Li (New York Institute of Technology); Xiaolin Xu (Northeastern University); *Wanli Chang (Hunan University and Huawei Technologies)*

CaPUF: Cascaded PUF Structure for Machine Learning Resiliency
Hassan Nassar (Karlsruhe Institute of Technology); Lars Bauer (Karlsruhe Institute of Technology); Joerg Henkel (Karlsruhe Institute of Technology)

Enhancing the Reliability and Security: A Configurable Poisoning PUF Against Modeling Attacks
Chia-Chih Lin (National Taiwan University); Ming-Syan Chen (National Taiwan University)

BLAST: Belling the Black-Hat High-Level Synthesis Tool
Mohammed Abderehman (Indian Institute Of Technology Guwahati); Rupak Gupta (Indian Institute Of Technology Guwahati); Theegala Rakesh Reddy (Indian Institute Of Technology Guwahati); Chandan Karfa (Indian Institute Of Technology Guwahati)

EMSOFT 4 Theory and Security
Session Chair: Alessandro Biondi (Sant'Anna School of Pisa)

Online Reset for Signal Temporal Logic Monitoring
Zhenya Zhang (Kyushu University); Paolo Arcaini (National Institute of Informatics); Xuan Xie (University of Alberta)

[Best paper candidate] Vulnerability Detection of ICS Protocols Via Cross-State Fuzzing
Feilong Zuo (Tsinghua University); Zhengxiong Luo (Tsinghua University); Junze Yu (Tsinghua University); Ting Chen (University of Electronic Science and Technology of China); Zichen Xu (Nanchang University); Aiguo Cui (Huawei Technologies Co. Ltd.); *Yu Jiang (Tsinghua University)*

Characterising the Effect of Deadline Misses on Time-Triggered Task Chains
Paolo Pazzaglia (Saarland University); Martina Maggio (Saarland University)

MIDAS: Safeguarding IoT Devices Against Malware via Real-Time Behavior Auditing
Yiwen Xu (Tsinghua University); Zijing Yin (Tsinghua University); Yiwei Hou (Tsinghua University); Jianzhong Liu (Tsinghua University); Yu Jiang (Tsinghua University)

Work in Progress Session
Session Chair: Reiley Jeyapaul (ARM)

Work-in-Progress: Imprecise Computing with Hot-patching Technique for Computing Resource Saving
Haegeon Jeong (Hanyang University); Kyungtae Kang (Hanyang University)

Work-in-Progress: A Browser-Driven Sensor Service for Embedded IoT
Agnieszka Chodorek (Kielce University of Technology); *Robert Ryszard Chodorek (The AGH University of Science and Technology)*

Work-in-Progress: Toward Energy-efficient Using Near STT-MRAM Processing Architecture for Neural Networks
Yueting Li (Beihang University); Bingluo Zhao (Beihang University); Xinyi Xu (Beihang University); Yundong Zhang (Vimicro Corporation); Jun Wang (Beihang University); Weisheng Zhao (Beihang University)

Work-in-Progress: HeteroRW; A Generalized and Efficient Framework for Random Walks in Graph Analysis
Yingxue Gao (University of Science and Technology of China); Lei Gong (University of Science and Technology of China); Chao Wang (University of Science and Technology of China); Xuehai Zhou (University of Science and Technology of China)

Work-in-Progress: Scheduler for Collaborated FPGA-GPU-CPU Based on Intermediate Language
Nan Hu (University of Science and Technology of China); Chao Wang (University of Science and Technology of China); Xuehai Zhou (University of Science and Technology of China); Xi Li (University of Science and Technology of China)

Work-in-Progress: High-Performance Systolic Hardware Accelerator for RBLWE-based Post-Quantum Cryptography
Tianyou Bao (Villanova University); José L. Imaña (Complutense University); Pengzhou He (Villanova University); *Jiafeng Xie (Villanova University)*

Work-in-Progress: Lark, A Learned Secondary Index Toward LSM-tree for Resource-Constrained Embedded Storage Systems
Jianan Yuan (Shenzhen University); Huan Liu (Shenzhen University); Shangyu Wu (City University of Hong Kong); Yiquan Lin (Shenzhen University); Tiantian Wang (Shenzhen University); Chenlin Ma (Shenzhen University); Rui Mao (Shenzhen University); Yi Wang (Shenzhen University)

SuperGuard, The solution for using C standard libraries in safety-critical applications
Marcel Beemster (Solid Sands)

CASES 5 System level design: Interconnect and 3D Stacking
Session Chair: Wei Zhang (HKUST)

Secured Data Transmission Over Insecure Networks-on-Chip by Modulating Inter-Packet Delays
Jiaen Xu (South China University of Technology); *Xiaohang Wang (Zhejiang University)*; Yingtao Jiang (University of Nevada at Las Vegas); Amit Kumar Singh (University of Essex); Chongyan Gu (Queen's University Belfast); Letian Huang (School of Electronic Science and Engineering); Mei Yang (University of Nevada at Las Vegas); Shunbin Li (Zhejiang Lab)

GCIM: Toward Efficient Processing of Graph Convolutional Networks in 3D-Stacked Memory
Jiaxian Chen (Shenzhen University); Yiquan Lin (Shenzhen University); Kaoyi Sun (Shenzhen University); Jiexin Chen (Shenzhen University); Chenlin Ma (Shenzhen University); Rui Mao (Shenzhen University); Yi Wang (Shenzhen University)

Accelerating Large-Scale Graph Neural Network Training on Crossbar Diet
Chukwufumnanya Ogbogu (Washington State University); Aqeeb Iqbal Arka (Washington State University); Biresh Kumar Joardar (Duke University); Janardhan Rao Doppa (Washington State University); Hai Li (Duke University); Krishnendu Chakrabarty (Duke University); Partha Pratim Pande (Washington State University)

Work-in-Progress: RISC-V Based Low-Cost Embedded Trace Processing System
Xiao Hu (National University of Defense Technology); Yaohua Wang (National University of Defense Technology); Xuan Gao (National University of Defense Technology)

Work-in-Progress: Emulation of biological tissues on an FPGA
Jerry Jacob (University of Auckland); Sucheta Sehgal (University of Auckland); Nitish Patel (University of Auckland)

Work-in-Progress: MLGOPerf: An ML Guided Inliner to Optimize Performance
Amir H. Ashouri (Huawei Technologies Canada); Mostafa Elhoushi (Huawei Technologies Canada); Yuzhe Hua (Huawei Technologies Canada); Xiang Wang (Huawei Technologies); Muhammad Asif Manzoor (Huawei Technologies Canada); Bryan Chan (Huawei Technologies Canada); Yaoqing Gao (Huawei Technologies Canada)

CODES+ISSS 5 Emerging Embdedded Memories and Storage
Session Chairs: Liang Shi (East China Normal University); Ming-chang Yang (The Chinese University of Hong Kong)

[Best paper candidate] Exploring Synchronous Page Fault Handling
Yin-Chiuan Chen (National Taiwan University); *Chun-Feng Wu (National Yang Ming Chiao Tung University and Harvard University)*; Yuan-Hao Chang (Academia Sinica); Tei-Wei Kuo (City University of Hong Kong and National Taiwan University)

When B+-Tree Meets Skyrmion Memory: How Skyrmion Memory Affects an Indexing Scheme
Jin-Wei Chang (Yuan Ze University); *Tseng-Yi Chen (National Central University)*

Resolving the Reliability Issues of Open Blocks for 3D NAND Flash: Observations and Strategies
Qiao Li (Xiamen University); Min Ye (City University of Hong Kong); Yufei Cui (City University of Hong Kong); Tianyu Ren (City University of Hong Kong); Tei-Wei Kuo (National Taiwan University); Jason Xue (City University of Hong Kong)

WA-OPShare: Workload-Adaptive Over-Provisioning Space Allocation for Multi-Tenant SSDs
Yuhong Wen (Huazhong University of Science and Technology); You Zhou (Huazhong University of Science and Technology); Fei Wu (Huazhong University of Science and Technology); Shu Li (Alibaba Group); Zhenghong Wang (Alibaba Group); ChangSheng Xie (Huazhong University of Science and Technology)

EMSOFT 5 Energy and Efficiency
Session Chair: Alain Girault (Inria)

DynLiB: Maximizing Energy Availability of Hybrid Li-ion Battery Systems

Jiwon Kim (Yonsei University); Sungwoo Baek (LIG Nex1 Co., Ltd.); Seunghyeok Jeon (Yonsei univ.); Hojung Cha (Yonsei University)

Stochastic Guarantees for Adaptive Energy Harvesting Systems
Rehan Ahmed (Information Technology University); Stefan Draskovic (ETH Zurich); Lothar Thiele (ETH Zurich)

Throughput Maximization in Wireless Communication Systems Powered by Hybrid Energy Harvesting
Chenchen Fu (Southeast University); Xinhang Lu (The University of New South Wales); Xiaoxing Qiu (Southeast University); Sujunjie Sun (Southeast University); Xueyong Xu (North Information Control Research Academy Group Co., Ltd); Weiwei Wu (Southeast University); Chun Jason Xue (City University of Hong Kong); Song Han (University of Connecticut)

Work-in-Progress: Accuracy-Area Efficient Online Fault Detection for Robust Neural Network Software-Embedded Microcontrollers
Juneseo Chang (Seoul National University); Sejong Oh (Nvidia); Daejin Park (Kyungpook National University)

Work in Progress: Dynamic Offloading of Soft Real-time Tasks in SDN-based Fog Computing Environment
Niraj Kumar (RGIPT Jais); Arijit Mondal (IIT Patna India)

Work-in-Progress: Accelerated Matrix Factorization by Approximate Computing for Recommendation System
Yining Wu (Shanghai University); Gaole Sai (Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences); Shengyu Duan (Shanghai University)

CASES 6 System Level Design: IOT, Reliability, and Verification
Session Chair: Jana Doppa (Washington State University)

Near-Optimal Energy Management for Energy Harvesting IoT Devices Using Imitation Learning
Nuzhat Yamin (Washington State University); Ganapati Bhat (Washington State University)

ARMISTICE: Microarchitectural Leakage Modelling for Masked Software Formal Verification
Arnaud de Grandmaison (Arm); Karine Heydemann (Sorbonne University); *Quentin L. Meunier (Sorbonne University)*

GNN4REL: Graph Neural Networks for Predicting Circuit Reliability Degradation
Lilas Alrahis (New York University Abu Dhabi); Johann Knechtel (New York University Abu Dhabi); Florian Klemme (University of Stuttgart); Hussam Amrouch (University of Stuttgart); Ozgur Sinanoglu (New York University Abu Dhabi)

Work-in-Progress: ACAC: An Adaptive Congestion-aware Approximate Communication Mechanism for Network-on-Chip Systems
Shize Zhou (Nanjing University), Yongqi Xue (Nanjing University), Siyue Li (Nanjing University), Jinlun Ji (Nanjing University); Tong Cheng (Nanjing University), Li Li (Nanjing University), Yuxiang Fu (Nanjing University)

Work-in-Progress: CAMiSE: Content Addressable Memory-integrated Searchable Encryption
Arnab Bag (Indian Institute of Technology Kharagpur); Sikhar Patranabis (ETH Zurich); Debdeep Mukhopadhyay (Indian Institute of Technology Kharagpur)

Work-in-Progress: An Open-Source Platform for Design and Programming of Partially Reconfigurable Heterogeneous SoCs
Biruk B. Seyoum (Columbia University); Davide Giri (Columbia University); Kuan-Lin Chiu (Columbia University); Luca P. Carloni (Columbia University)

TECHNICAL SESSIONS

CODES+ISSS 6 Reliable and Secure Embedded Systems
Session Chairs: Jiafeng Xie (Villanova University); Ishan Thakkar (University of Kentucky);

Detecting Spoofed Speeches via Segment-Based Word CQCC and Average ZCR for Embedded Systems
Jinyu Zhan (University of Electronic Science and Technology of China); Zhibei Pu (University of Electronic Science and Technology of China); Wei Jiang (University of Electronic Science and Technology of China); Junting Wu (University of Electronic Science and Technology of China); Yongjia Yang (University of Electronic Science and Technology of China)

Combating Stealthy Thermal Covert Channel Attack With Its Thermal Signal Transmitted in Direct Sequence Spread Spectrum
Xiaohang Wang (Zhejian University and South China University of Technology); Shengjie Wang (South China University of Technology); Yingtao Jiang (University of Nevada Las Vegas); Amit Kumar Singh (University of Essex); Mei Yang (University of Nevada Las Vegas); Letian Huang (University of Electronic Science and Technology of China)

eRDAC: Efficient and Reliable Remote Direct Access and Control for Embedded Systems
Junjie Feng (Chongqing University); Xianzhang Chen (Chongqing University); Duo Liu (Chongqing University); Weigong Zhang (Capital Normal University); Jiapin Wang (Chongqing University); Rongwei Zheng (Chongqing university); Yujuan Tan (Chongqing University)

Architecting Decentralization and Customizability in DNN Accelerators for Hardware Defect Adaptation
Elbruz Ozen (University of California at San Diego); Alex Orailoglu (University of California at San Diego)

EMSOFT 6 Power, Energy, and Analysis
Session chair: Susanne Graf (CNRS)

CapOS: Capacitor Error Resilience for Energy Harvesting Systems
Jongouk Choi (University of Central Florida); Hyunwoo Joe (Electronics and Telecommunications Research Institute(ETRI)); Changhee Jung (Purdue University)

Cyber-Physical Verification of Intermittently Powered Embedded System
Rose Bohrer (Worcester Polytechnic Institute); Bashima Islam (Worcester Polytechnic Institute)

Sparsity-Aware Intelligent Spatiotemporal Data Sensing for Energy Harvesting IoT System
Wen Zhang (Texas A&M University at Corpus Christi); Mimi Xie (The University of Texas at San Antonio); Caleb Scott (The University of Texas at San Antonio); Chen Pan (Texas A&M University at Corpus Christi)

Work-in-Progress: A Resource-Aware Optimization Model for Real-Time Systems Analysis and Design
Rezwana Mamata (Ontario Tech University); Akramul Azim (Ontario Tech University)

Work-in-Progress: Boot Sequence Integrity Verification with Power Analysis
Arthur Grisel-Davy (University of Waterloo); Amrita Milan Bhogayata (University of Waterloo); Srijan Pabbi (University of Waterloo); Apurva Narayan (University of Waterloo); Sebastian Fischmeister (University of Waterloo)

Work-in-Progress: Towards a Theory of Robust Quantitative Semantics for Signal Temporal Logic
Jean-Baptiste Jeannin (University of Michigan); Jiawei Chen (University of Michigan); José Luiz Vargas de Mendonça (University of Michigan); Konstantinos Mamouras (Rice University)

Special Session Programming Autonomous Machines
Shaoshan Liu (PerceptIn); Tongsheng Geng (UC Irvine); Stéphane Zuckerman (ETIS); Xiaoming Li (University of Delaware); Jean-Luc Gaudiot (UC Irvine)

Abstract
After decades of uninterrupted progress and growth, information technology has so evolved that it can be said we are entering the age of autonomous machines. One key technical challenge in the age of autonomous machines is the programming of autonomous machines, which demands the synergy across multiple domains, including fundamental computer science, computer architecture, and robotics, and requires expertise from both academia and industry.
This special session will provide the opportunity to discuss the programming theory and practices tied to producing real-life autonomous machines. The discussion covers aspects from high-level concepts down to low-level code generation in the context of specific functional requirements, performance expectation, and implementation constraints of autonomous machines.

For instance, autonomous vehicles rely on a wealth of specialized components according to the various tasks they are required to perform, which includes (hard) real-time tasks related to the outside environment, i.e., Localization and Navigation, Object Detection and Avoidance, etc. Each task communicate its data to other tasks within strict performance envelop, and may also rely on very different hardware targets to do so, e.g., scheduling with CPUs, GPUs for neural network processing, FPGAs or DSPs for image processing, etc..

Thus, there should be expressive “languages” to describe, at a high level, what each task should consist of, the appropriate (domain-specific) semantics for it, and at the same time, the interfacing between languages so as a whole real-time and run-time contexts can freely dictate which part of the underlying hardware will be used to run them at a specific moment during the motion of the target autonomous machine. To do so, simply designing a new set of DSLs is probably necessary, but not sufficient. As these autonomous machines tend to heavily rely on Machine Learning techniques for at least some of their tasks, there should also be a way to lower the high-level description provided for each task down to an intermediate representation which would allow the compiler to produce code for a family of heterogeneous devices. These considerations, which are both high and low level, can be unified under a dataflow-oriented hierarchical view of the system.

Hence, the main topics we wish to address during this special session are related to expressing the tasks of how autonomous machines must perform at high-level, and how the expression is best translated into low-level operations:

- How to make autonomous machines more programmable, through the use of high-level languages;
- How can such high-level languages be translated, i.e., lowered, to a suitable set of intermediate representations for a compiler, in order to eventually produce executable code on heterogeneous hardware;
- How is the runtime system meant to deal with the heterogeneity of the underlying hardware and dynamic performance envelop, and how can it benefit from both high-level languages, and the knowledge of the low-level machine models it has access to.

Biographies
Shaoshan Liu is CEO and founder of PerceptIn, an intelligent robotics company. He is the founder of IEEE Special Technical Community on Autonomous Driving Technologies, an IEEE Senior Member, an IEEE Distinguished Speaker, and an ACM Distinguished Speaker, and a Fellow of the Royal Society for Public Health (UK). Dr. Liu holds a M.S. in Biomedical Engineering, a Ph.D. in Computer Engineering from the University of California, Irvine and a Master of Public Administration (M.P.A.) from Harvard University. Dr. Liu has published more than 100 research papers and holds more than 150 patents in the field of robotics and autonomous systems.

Tongsheng Geng is a senior researcher at the University of California, Irvine. Her expertise covers adaptive resource management for heterogeneous platforms, event-driven parallel computing, and machine learning techniques applied to system software. Her current research interests in autonomous machine area are domain-specific language and distributed autonomous machine simulation systems.

Stéphane Zuckerman is an Associate Professor at CY Cergy Paris University, in the ETIS Laboratory. His research focuses on event-driven execution models for parallel and distributed computing on high-performance computing

and embedded systems—in particular implemented as system software over heterogeneous, reconfigurable fabric. His recent work includes the implementation of a bio-inspired neural architecture for the localization task in autonomous vehicles on reconfigurable fabric.

Xiaoming Li is an Associate Professor at the University of Delaware. His expertise covers compiler development and code generation techniques. The focus is on the optimal resource utilization on heterogeneous computing resources. In particular, his recent research develops optimization techniques based on the unique resource utilization patterns in machine learning.

Jean-Luc Gaudiot is a Distinguished Professor at the University of California, Irvine. His research focuses on the programmability of large-scale multiprocessors, distributed structure representation, data caching and nomadic threads, high-level parallel language compilation, allocation, partitioning, and performance evaluation; processor architecture, implementation and impact of future technologies; fault-tolerant computing, benchmarking and characterization of applications, design of flexible architectures tuned for specific types of applications; and the design of Autonomous Vehicle Systems.

PROGRAM

WEDNESDAY, OCTOBER 12

CST - 10/12/2022 08:30-09:15 PST - 10/11/2022 17:30-18:15 EST - 10/11/2022 20:30-21:15 CET - 10/12/2022 02:30-03:15	Sky talk 2 - Open Source Software Stacks for Heterogeneous SoCs, by Tomas Evenson, AMD/Xilinx			
CST - 10/12/2022 09:15-12:00 PST - 10/11/2022 18:15-21:00 EST - 10/11/2022 21:15-00:00 CET - 10/12/2022 03:15-06:00	Workshop on Memory and Storage Computing (MSC)	Workshop on Heterogeneous Edge Computing (HEC)	ACM SIGBED Student Research Competition (SRC)	
CST - 10/12/2022 14:00-15:00 PST - 10/11/2022 23:00-00:00 EST - 10/12/2022 02:00-03:00 CET - 10/12/2022 08:00-09:00	CASES 7 - Non-Volatile Memory : Flash and Cache	CODES+ISSS 7 - Intermittent Computing	EMSOFT 7 -Verification and Machine Learning	ACM SIGBED Student Research Competition (SRC)
CST - 10/12/2022 15:00-16:00 PST - 10/12/2022 00:00-01:00 EST - 10/12/2022 03:00-04:00 CET - 10/12/2022 09:00-10:00	CASES 8 - Hardware Support for AI	CODES+ISSS 8 - Machine Learning Acceleration in an Embedded World	EMSOFT 8 - Theory and Control	Ph.D. Forum
CST - 10/12/2022 16:00-16:30 PST - 10/12/2022 01:00-01:30 EST - 10/12/2022 04:00-04:30 CET - 10/12/2022 10:00-10:30	Posters	Industry Booth	Yoga and Escape Room (GatherTown)	
CST - 10/12/2022 16:30-17:00 PST - 10/12/2022 01:30-02:00 EST - 10/12/2022 04:30-05:00 CET - 10/12/2022 10:30-11:00	ESWEEK Networking	Industry Interaction	Yoga and Escape Room (GatherTown)	
CST - 10/12/2022 22:00-23:00 PST - 10/12/2022 07:00-08:00 EST - 10/12/2022 10:00-11:00 CET - 10/12/2022 16:00-17:00	Keynote 3 - The Computing and Information Science and Engineering Landscape: A Look Forward, by Dr. Margaret Martonosi, National Science Foundation (NSF)			
CST - 10/12/2022 23:00-23:30 PST - 10/12/2022 08:00-08:30 EST - 10/12/2022 11:00-11:30 CET - 10/12/2022 17:00-17:30	Best Paper and Other Awards			

WEDNESDAY, OCTOBER 12

CST - 10/12/2022 23:30-00:30 PST - 10/12/2022 08:30-09:30 EST - 10/12/2022 11:30-12:30 CET - 10/12/2022 17:30-18:30	**ESWEEK Panel - Waferscale Computing Systems: Are We There Yet?**

KEYNOTES

Keynote 3 – The Computing and Information Science and Engineering Landscape: A Look Forward, by Dr. Margaret Martonosi, National Science Foundation (NSF)

Description

The United States National Science Foundation (NSF) supports a majority of US academic research in the Computer and Information Science and Engineering (CISE) topic areas. A long-time computing researcher herself, Dr. Margaret Martonosi is now serving a 4-year term leading the NSF CISE Directorate, and stewarding the CISE directorate $1B+ annual budget on behalf of research, education, workforce and infrastructure funding in CISE topic areas and for science as a whole. In this talk, she will discuss key technical themes for the field, and how CISE is developing programmatic opportunities to advance research related to them. She will particularly note how ESWEEK topic areas relate to these technical priorities. More broadly, she will discuss CISE and NSF in the context of the global research efforts, and our approach to industry and international research partnerships.

Biography

Margaret Martonosi leads the US National Science Foundation's (NSF) Directorate for Computer and information Science and Engineering (CISE). With an annual budget of more than $1B, the CISE directorate at NSF has the mission to uphold the Nation's leadership in scientific discovery and engineering innovation through its support of fundamental research and education in computer and information science and engineering as well as transformative advances in research cyberinfrastructure. While at NSF, Dr. Martonosi is on leave from Princeton University where she is the Hugh Trumbull Adams '35 Professor of Computer Science. Dr. Martonosi's research interests are in computer architecture and hardware-software interface issues in both classical and quantum computing systems. Dr. Martonosi is a member of the National Academy of Engineering and the American Academy of Arts and Sciences. She is a Fellow of the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE).

Sky talk 2 – Open Source Software Stacks for Heterogeneous SoCs, by Tomas Evensen, AMD/Xilinx

Description

AMD/Xilinx's embedded SoCs integrate a lot of heterogenous execution units, like multiple CPU clusters, AI Engines, programmable logic and other accelerators. Most traditional embedded software stacks are optimized for a single operating environment running on one or more CPUs and need to be extended to handle the new complexity. This talk will introduce the Xilinx SoCs and the various open source projects that handles both the runtimes (communication, management, separation, etc.) and tooling (compilation, configuration, debugging, etc.) for these execution units.

Biography

Tomas Evensen is Chief Technology Officer, Open Source at AMD/Xilinx. He has also been responsible for the embedded software strategy for Xilinx All Programmable SoCs. Prior to joining Xilinx, Evensen was Chief Technology Officer at Wind River for 7 years, as well as GM for the Wind River Tools Division and VP of Engineering for the VxWorks operating system. Before that he was the creator of the Diab Data C/C++ compilers. Evensen received his MSEE at the Royal Institute of Technology in Stockholm, Sweden.

TECHNICAL SESSIONS

CASES 7 Non-volatile Memory: Flash and Cache
Session Chair - Duo Liu (Chongqing University), Sandeep Chandran (IIT Palakkad)

LLSM: A Lifetime-Aware Wear-Leveling for LSM-Tree on NAND Flash Memory
Dharamjeet (Academia Sinica); *Yi-Shen Chen (Academia Sinica)*; Tseng-Yi Chen (National Central University); Yuan-Hung Kuan (Academia Sinica); Yuan-Hao Chang (Academia Sinica)

Fast and Low Overhead Metadata Operations for NVM-Based File System Using Slotted Paging
Fangzhu Lin (Chongqing University); Chunhua Xiao (Chongqing University); Weichen Liu (Nanyang Technological University); Lin Wu (Chongqing University); Chen Shi (Chongqing University); Kun Ning(Chongqing University)

MAID-Q: Minimizing Tail Latency in Embedded Flash With SMR Disk via Q-learning Model
Chenlin Ma (Shenzhen University); *Zhuokai Zhou (Shenzhen University)*; Yingping Wang (Shenzhen University); Yi Wang (Shenzhen University); Rui Mao (Shenzhen University)

Work-in-Progress: High-Precision Short-Term Lifetime Prediction in TLC 3D NAND Flash Memory as Hot-data Storage
Xiaotong Fang (Shandong University); Meng Zhang (Huazhong University of Science and Technology); Yifan Guo (Shandong University); Fei Chen (Shandong University); Binglu Chen (Shandong University); Xuepeng Zhan (Shandong University); Jixuan Wu (Shandong University); Fei Wu (Huazhong University of Science and Technology); Jiezhi Chen (Shandong University)

Work-in-Progress: Prediction-based Fine-Grained LDPC Reading to Enhance High-Density Flash Read Performance
Yajuan Du (Wuhan University of Technology); Yuan Gao (Wuhan University of Technology); *Qiao Li (Xiamen University)*

Work-in-Progress: ExpCache: Online-Learning based Cache Replacement Policy for Non-Volatile Memory
Jinfeng Yang (University of Minnesota); Bingzhe Li (Oklahoma State University); Jianjun Yuan (Expedia Group); Zhaoyan Shen (Shandong University); David Du (University of Minnesota); David Lilja (University of Minnesota)

CODES+ISSS 7 Intermittent Computing
Session Chairs: Yuan-Hao Chang (Academia Sinica); Mengying Zhao (Shandong University)

Stateful Neural Networks for Intermittent Systems
Chih-Hsuan Yen (National Taiwan University and Academia Sinica); Hashan Roshantha Mendis (Academia Sinica); Tei-Wei Kuo (National Taiwan University); Pi-Cheng Hsiu (Academia Sinica)

Intermittent-Aware Distributed Concurrency Control
Wei-Che Tsai (National Taiwan University); Wei-Ming Chen (Massachusetts Institute of Technology); Tei-Wei Kuo (National Taiwan University); Pi-Cheng Hsiu (Academia Sinica)

PVoT: Reconfigurable Photovoltaic Array for Indoor Light Energy-Powered Batteryless Devices
Jiwon Kim (Yonsei University); Eunyeong Kim (Yonsei University); Seunghyeok Jeon (Yonsei University); Junick Ahn (Yonsei University); Hyungchol Jun (Yonsei University); Hojung Cha (Yonsei University)

SENTunnel: Fast Path for Sensor Data Access on Automotive Embedded Systems
Rongwei Zheng (Chongqing University); Xianzhang Chen (Chongqing University); Duo Liu (Chongqing University); Junjie Feng (Chongqing University); Jiapin Wang (Chongqing University); Ao Ren (Chongqing University); Chengliang Wang (Chongqing University); Yujuan Tan (Chongqing University)

EMSOFT 7 Verification and Machine Learning
Session Chair: Indranil Saha (IIT Kanpur)

An MILP Encoding for Efficient Verification of Quantized Deep Neural Networks
Samvid Mistry (GitHub); Indranil Saha (Indian Institute of Technology Kanpur); Swarnendu Biswas (Indian Institute of Technology Kanpur)

Efficient Complete Verification of Neural Networks via Layerwised Splitting and Refinement
Banghu Yin (National University of Defense Technology); Liqian Chen (National University of Defense Technology); Jiangchao Liu (Ant Group); Ji Wang (National University of Defense Technology)

Tardis: Coverage-Guided Embedded Operating System Fuzzing
Yuheng Shen (Tsinghua University); Yiru Xu (Tsinghua University); Hao Sun (Tsinghua University); Jianzhong Liu (Tsinghua University); Zichen Xu (Nanchang University); Aiguo Cui (Huawei Technologies Co. Ltd.); Heyuan Shi (Central South University); Yu Jiang (Tsinghua University)

Formal Verification of Resource Synchronization Protocol Implementations: A Case Study in RTEMS
Junjie Shi (TU Dortmund); Christoph-Cordt von Egidy (TU Dortmund); Kuan-Hsun Chen (University of Twente); Jian-Jia Chen (TU Dortmund)

CASES 8 Hardware Support for AI
Session Chair - Neetu Jindal (Intel)

Hardware-Friendly Delayed-Feedback Reservoir for Multivariate Time-Series Classification
Sosei Ikeda (Kyoto University); Hiromitsu Awano (Kyoto University); Takashi Sato (Kyoto University)

Bits-Ensemble: Toward Light-Weight Robust Deep Ensemble by Bits-Sharing
Yufei Cui (McGill University); *Shangyu Wu (City University of Hong Kong)*; Qiao Li (Xiamen University); Antoni B. Chan (City University of Hong Kong); Tei-Wei Kuo (National Taiwan University); Chun Jason Xue (City University of Hong Kong)

Exploring Bitslicing Architectures for Enabling FHE-Assisted Machine Learning
Soumik Sinha (Indian Institute of Technology Kharagpur); Sayandeep Saha (Nanyang Technological University); Manaar Alam (New York University, Abu Dhabi); Varun Agarwal (Delhi Technological University); Ayantika Chatterjee (Indian Institute of Technology Kharagpur); Anoop Mishra (University of Nebraska Omaha); Deepak Khazanchi (University of Nebraska Omaha); Debdeep Mukhopadhyay (Indian Institute of Technology Kharagpur)

Work-in-Progress: Object Detection Acceleration Method by Improving Execution Efficiency of AI Device
Yoshikazu Watanabe (NEC Corporation); Yuki Kobayashi (NEC Corporation); Noboru Nakajima (NEC Corporation); Takashi Takenaka (NEC Corporation); Hiroyoshi Miyano (NEC Corporation)

Work-in-Progress: Cooperative MLP-Mixer Networks Inference On Heterogeneous Edge Devices through Partition and Fusion
Yiming Li (East China Normal University); Shouzhen Gu (East China Normal University); Mingsong Chen (East China Normal University)

CODES+ISSS 8 Machine Learning Acceleration in An Embedded World
Session Chairs: Keni Qiu (Capital Normal University); Duo Liu (Chongqing University)

A Flexible Yet Efficient DNN Pruning Approach for Crossbar-Based Processing-in-Memory Architectures
Long Zheng (Huazhong University of Science and Technology); *Haifeng Liu (Huazhong University of Science and Technology)*; Yu Huang (Huazhong University of Science and Technology); Dan Chen (Huazhong University of Science and Technology); Chaoqiang Liu (Huazhong University of Science and Technology); Haiheng He (Huazhong University of Science and Technology); Xiaofei Liao (Huazhong University of Science and Technology); Hai Jin (Huazhong

University of Science and Technology); Jingling Xue (University of New South Wales)

ReaDy: A ReRAM-Based Processing-in-Memory Accelerator for Dynamic Graph Convolutional Networks
Yu Huang (Huazhong University of Science and Technology); Long Zheng (Huazhong University of Science and Technology); Pengcheng Yao (Huazhong University of Science and Technology); Qinggang Wang (Huazhong University of Science and Technology); Haifeng Liu (Huazhong University of Science and Technology); Xiaofei Liao (Huazhong University of Science and Technology); Hai Jin (Huazhong University of Science and Technology); Jingling Xue (University of New South Wales)

Efficient Hardware Acceleration of Sparsely Active Convolutional Spiking Neural Networks
Jan Sommer (Friedrich-Alexander-Universität Erlangen-Nürnberg); M. Akif Özkan (Friedrich-Alexander-Universität Erlangen-Nürnberg); Oliver Keszocze (Friedrich-Alexander-Universität Erlangen-Nürnberg); *Jürgen Teich (Friedrich-Alexander-Universität Erlangen-Nürnberg)*

ViA: A Novel Vision-Transformer Accelerator Based on FPGA
Teng Wang (University of Science and Technology of China); Lei Gong (University of Science and Technology of China); Chao Wang (University of Science and Technology of China); Yang Yang (University of Science and Technology of China); Yingxue Gao (University of Science and Technology of China); Xuehai Zhou (University of Science and Technology of China); Huaping Chen (University of Science and Technology of China)

EMSOFT 8 Theory and Control
Session Chair: David Broman (KTH)

Differentiable Inference of Temporal Logic Formulas
Nicole Fronda (Oregon State University); Houssam Abbas (Oregon State University)

Efficient Backward Reachability Using the Minkowski Difference of Constrained Zonotopes
Liren Yang (Huazhong University); Hang Zhang (University of Wisconsin-Madison); Jean-Baptiste Jeannin (University of Michigan); Necmiye Ozay (University of Michigan)

Towards Minimum WCRT Bound for DAG Tasks Under Prioritized List Scheduling Algorithms
Shuangshuang Chang (Northeastern University); Ran Bi (Dalian University of Technology); Jinghao Sun (Dalian University of Technology); Weichen Liu (Nanyang Technological University); Qi Yu (Dalian University of Technology); Qingxu Deng (Northeastern University); Zonghua Gu (Umeå University)

Response-Time Analysis of Limited-Preemptive Sporadic DAG Tasks
Gaoyang Dai (Uppsala University); Morteza Mohaqeqi (Uppsala University); Petros Voudouris (Uppsala University); Wang Yi (Uppsala University)

Ph.D. Forum
Session Chairs: Mengying Zhao (Shandong University); Mengying Zhao (Shandong University)

Towards Accelerator Design 2.0
Shail Dave (Arizona State University)

Towards Event-driven Context: JavaScript, an Energy-Efficient Language for the Internet of Things
Fernando Oliveira (Federal University of Pelotas)
STT-RAM-based In-Memory Computing Across the Memory Hierarchy
Dhruv Gajaria (University of Arizona)

Stochastic Bit-Stream Processing Systems with New Perspectives Towards Vision and Learning Machines
Sercan Aygun (University of Louisiana at Lafayette)

Design Automation for Microfluidic Biochips
Debraj Kundu (Indian Institute of Technology, Roorkee)

A High-Level Synthesis Approach for Precisely-Timed, Energy-Efficient Embedded Systems
Yuchao Liao (University of Arizona)

ESWEEK Panel
Title: Waferscale Computing Systems: Are We There Yet?
Organizers: Puneet Gupta *(UCLA)*; Saptadeep Pal *(Auradine Inc.)*

Abstract: Fueled by the tremendous growth of new applications in the domain of big-data computing, deep learning, and scientific computing, the demand for increasing system performance is far outpacing the capability of conventional methods for system performance scaling. Waferscale computing, where an entire 300 mm wafer worth of compute and memory can be extremely tightly integrated, promises to provide orders of magnitude improvement in performance and energy efficiency compared to today's systems built using traditional packaging technologies.

In this panel, we will discuss the "Promised Land" of waferscale computing. Back in the 1980s, waferscale systems were attempted by a few companies, notable amongst them were Trilogy systems and Tandem Computers. However, yield and cost challenges of building waferscale systems led to its early demise, but the promise remained. After more than 30 years, recent academic (e.g., UCLA/UIUC) and industrial (Cerebras, Tesla) efforts have taken up the challenge again. So, are we in a waferscale technology renaissance period and nearing the days when waferscale technologies would be widely adopted? Many questions remain which we will discuss in this panel.

First and foremost, is the overall technology there yet? Does manufacturing difficulties, more so in the advanced nodes, limit choice of waferscale architectures? Would waferscale integration of heterogeneous chiplets open up more architectural choices over monolithic waferscale technologies ? Waferscale computing comes with very high power density, which means 10s of kilowatts of power need to be supplied and that heat needs to be extracted from the wafer. Is the data center infrastructure ready to accommodate such waferscale systems at scale? Are there lower power use cases for Waferscale ? Moreover, the design infrastructure, such as EDA and simulation tools are not yet truly ready for larger-than-a-reticle design. Thus, the overall design challenges of a waferscale system is humongous and so would the development be confined to a niche group? What are the applications that need waferscale computing and would benefit massively from such systems? Does the cost of waferscale systems justify adoption of these systems at volume? Are there previously untenable applications and business cases that now become feasible with waferscale computing, if so what are those? Are there edge compute use cases where the volumetric compute density would lead to adoption of waferscale systems?

Panelists:

Rakesh Kumar is a Professor in the Electrical and Computer Engineering Department at the University of Illinois at Urbana Champaign with research and teaching interests in computer architecture and system-level design automation. His research has been recognized through one ISCA Influential Paper Award, one MICRO Test-of-Time Award, one ASPDAC 10 Year Retrospective Most Influential Paper (MIP) Award, several best paper awards and best paper award nominations (IEEE MICRO Top Picks, ASPLOS, HPCA, CASES, SELSE, IEEE CAL), ARO Young Investigator Award, and UCSD CSE Best Dissertation Award. His teaching and advising have been recognized through Stanley H Pierce Faculty Award and Ronald W Pratt Faculty Outstanding Teaching Award. He often writes about issues at the intersection of technology, policy, and society; he is the author of the book Reluctant Technophiles (Sage Select: Dec 2021), one of "GQ's Best Indian Non-fiction Books of 2021". Rakesh has a BS from IIT Kharagpur and a PhD from University of California at San Diego.

Gabriel H. Loh is a Senior Fellow in AMD Research, the research and advanced development lab for Advanced Micro Devices, Inc. Gabe received his Ph.D. and M.S. in computer science from Yale University in 2002 and 1999, respectively, and his B.Eng. in electrical engineering from the Cooper Union in 1998. Gabe was also a tenured associate professor in the College of Computing at the Georgia Institute of Technology, a visiting researcher at Microsoft Research, and a senior researcher at Intel Corporation. He is a Fellow of the ACM and IEEE, recipient of ACM SIGARCH's Maurice Wilkes Award, Hall of Fame member for the MICRO, ISCA, and HPCA conferences, (co-)inventor on over one hundred US patent applications and over ninety granted patents, and a recipient of the US National Science Foundation Young Faculty CAREER Award.

Joel Hestness is a Senior Research Scientist at Cerebras Systems, an AI-focused hardware startup building the largest ever processors using wafer-scale integration. Joel helps define algorithms, performance optimizations, and scaling approaches for machine learning and NLP applications on the Cerebras Wafer-Scale Engine. Previously, Joel was a Research Scientist at Baidu's Silicon Valley AI Lab (SVAIL), where he worked on deep learning speech and language modeling. His work was the first to demonstrate predictable accuracy scaling laws for modern deep learning algorithms, sparking a trend of scaling law studies now pervasive in the field. Joel received his PhD in computer architecture from the University of Wisconsin – Madison, and his Bachelor's degrees in Mathematics and Computer Science also from UW-Madison.

Dave Nellans joined NVIDIA in 2013 where he is the Director of System Architecture Research. His research interests include building scalable computing systems that optimize node-level efficiency by improving the performance, utilization, and interaction of GPUs, CPUs, smart NICS, and storage systems. Dr. Nellans was previously an early engineering leader at Fusion-IO, one of the pioneers in PCIe-attached NAND-flash storage, where helped the company invent, develop, and ship new datacenter storage products that ultimately led to Fusion-IO's IPO in 2011. He holds a B.A. in Computer Science from Colgate University and a Ph.D. in Computer Science from the University of Utah.

PROGRAM

THURSDAY, OCTOBER 13

CST - 10/13/2022 22:00-02:00 PST - 10/13/2022 07:00-11:00 EST - 10/13/2022 10:00-14:00 CET - 10/13/2022 16:00-20:00	**20th IEEE/ ACMFormal Methods and Models for System Design (MEMOCODE)**	**33rd International Workshop on Rapid System Prototyping (RSP)**	**Workshop on Compilers, Deployment, and Tooling for Edge AI (CODAI)**
CST - 10/13/2022 22:00-03:00 PST - 10/13/2022 07:00-12:00 EST - 10/13/2022 10:00-15:00 CET - 10/13/2022 16:00-21:00	**16th IEEE/ACM International Symposium on Networks-on-Chip (NOCS)**		
CST - 10/14/2022 00:00-07:30 PST - 10/13/2022 09:00-16:30 EST - 10/13/2022 12:00-19:30 CET - 10/13/2022 18:00-01:30	**3rd International Workshop on Secure RISC-V Architecture Design Exploration (SECRISC-V)**	**On-site Networking and Social Events (Phoenix)**	

PROGRAM

FRIDAY, OCTOBER 14

CST - 10/14/2022 22:00-01:00 PST - 10/14/2022 07:00-10:00 EST - 10/14/2022 10:00-13:00 CET - 10/14/2022 16:00-19:00	**20th IEEE/ACMFormal Methods and Models for System Design (MEMOCODE)**	
CST - 10/14/2022 22:00-02:00 PST - 10/14/2022 07:00-11:00 EST - 10/14/2022 10:00-14:00 CET - 10/14/2022 16:00-20:00	**33rd International Workshop on Rapid System Prototyping (RSP)**	**SIGBED Buisness Meeting**
CST - 10/14/2022 22:00-03:00 PST - 10/14/2022 07:00-12:00 EST - 10/14/2022 10:00-15:00 CET - 10/14/2022 16:00-21:00	**16th IEEE/ACM International Symposium on Networks-on-Chip (NOCS)**	

WORKSHOPS

Workshop 1: CODAI'22: WORKSHOP ON COMPILERS, DEPLOYMENT, AND TOOLING FOR EDGE AI

Organizers: Michael J. Klaiber (Bosch Research); Sebastian (Vogel NXP Semiconductors,)

About: The goal of the CODAI workshop is to coalesce the emerging energy in the AI compiler communities and AI accelerator communities focusing on Edge AI in both academic and industrial research. These realms of research have the opportunity to deliver a pervasive and seamless end-to-end tooling that connects hardware and software development methodologies. This workshop mainly focuses on: Discussing new impulses for deployment of Neural Networks on low-end embedded systems and Cooperative development of (open source) toolchains/frameworks for neural network deployment + cooperation between industry and academia.Therefore, it complements the established conferences CASES and CODES+ISSS which have a wider and more academic research-spectrum. We also welcome work-in-progress papers of ongoing research projects and case studies from industrial applications. Topics for workshop submissions include, but are not limited to: optimization techniques and performance estimation of neural networks on embedded systems: e.g., compression, quantization techniques, virtual prototyping, compilers for Edge AI: partitioning, µC, heterogeneous systems, intermediate representations or (domain specific) languages, code-generation and hardware-backends for AI accelerators – especially for RISC-V is appreciated, applications: processing of embedded vision, time-series data, etc., novel brain-inspired algorithm for Edge-AI, and compiler and optimization techniques for beyond-von-Neumann AI accelerators.

Workshop 2: INTERNATIONAL WORKSHOP ON HETEROGENEOUS EDGE COMPUTING FOR EMBEDDED SYSTEM (HEC)

Organizers: Songwen Pei Professor (University of Shanghai for Science and Technology); Tong Liu (Shanghai University)

About: Based on various IoT devices or local edge nodes which provide strong capability of heterogeneous computation, Edge computing has become a distributed computing framework in the past few years. Especially, mobile computing devices, self-driving auto, service robot, etc. are emerging as the future applications supported by edge computing that would bring much more convenience to people and even change modern lifestyles. The workshop will bring together scientists and engineers from industry and academia who are working on heterogeneous computing and edge computing systems. The topics of the workshop will include but not limit to the following topics: microarchitecture design on heterogeneous processor/system combined with emerging memory/storage system (PCM, SSD, etc.), heterogeneous edge computing paradigms and models, energy efficient computing paradigms for embedded systems, heterogamous edge computation on mobile devices, heterogeneous computation supports for autonomous vehicle driving, or other applications using artificial intelligent algorithms (e.g. images processing, features recognition, etc.), heterogeneous edge computing in distributed datacenter, software engineering implementation on heterogeneous computing, and task scheduling algorithms on heterogeneous edge computation on distributed platforms.

Workshop 3: INTERNATIONAL WORKSHOP ON MEMORY AND STORAGE COMPUTING (MSC)

Organizers: Liang Shi (East China Normal University); Weichen Liu (Nanyang Technological University); Yuan-hao Chang (Academia Sinica

About: Memory and storage have been developed during the last decades. However, the data transfer cost between CPU and storage/memory becomes the critical challenge for the advanced systems. Storage and memory computing provides a new opportunity to solve this issue by adding computing functions beside storage or memory. In this workshop, we are not only interested in the advanced storage and memory computing technologies, but also pay special attention to the advanced storage and memory technologies. The workshop will bring together scientists and engineers from industry and academia who are working on storage and memory computing architectures and systems. The topics of the workshop will include but not limit to the following topics: storage computing architecture, memory computing architecture, systems for memory and computing, systems for storage and computing, state-of-the-art storage technologies, state-of-the-art memory technologies, FPGA enabled in-storage computing, intermittent

computing/system, design-space exploration for memory and storage systems, cross-layer design methodologies for memory hierarchy, compiler and OS optimization for emerging memory and storage, modeling, simulation and analysis for memory and storage, storage and memory design issues for emerging application scenarios, data storage and memory management interplay between embedded and IoT systems / Edge / Cloud, unified memory and storage, non-volatile computing systems, and in-memory computing applications.

Workshop 4: INTERNATIONAL WORKSHOP ON EDGE INTELLIGENT COMPUTING (EIC)

Organizers: Keni Qiu (Capital Normal University); Lei Yang (George Mason University)

About: Edge Intelligent Computing (EIC) is an emerging technology paradigm appealing for extensive Internet of Things (IoT) scenarios such as environment monitoring, disaster warning, and meta cyberspace. However, low power and constrained hardware resources impose critical design challenges in the process of EIC systems' real-world implementation. In this workshop, we are not only interested in the advanced edge computing technologies, but also bring a focus on how to implement the computation-/data-intensive machine learning algorithms on IoT edge devices. The workshop will bring together both scientists and engineers from academia and industry who are working on Edge Intelligent Computing architectures and systems. The topics of the workshop will include but not limit to the following topics: edge computing architectures, edge computing software co-design, state-of-the-art edge computing systems Intermittent computing systems, energy harvesting systems, low power ML algorithm optimization, low power Processing-In-Memory (PIM) accelerators, nonvolatile processors, FPGA-based edge intelligent computing, design-space exploration for edge intelligent computing systems, cross-layer design methodologies for edge (intelligent) computing, compiler and OS optimization for edge (intelligent) computing, modeling, simulation, and analysis for edge computing systems, and case study for edge intelligent computing systems.

Workshop 5: 33RD INTERNATIONAL WORKSHOP ON RAPID SYSTEM PROTOTYPING (RSP)

Organizers: Frédéric Rousseau (TIMA, University Grenoble-Alpes); Fabiano Hessel (PUCRS); Amer Baghdadi (IMT Atlantique/Lab-STICC); Kenneth Kent (University of New Brunswick); Sungjoo Yoo (Seoul National University)

About: The International Workshop on Rapid System Prototyping (RSP) emphasizes design experience sharing and collaborative approaches between hardware and software research communities from industry and academy. It considers prototyping as an iterative design approach for embedded hardware and software systems. The RSP series of workshops aims at bridging the gaps in embedded system design between applications, architectures, tools, and technologies to achieve rapid system prototyping of emerging software and hardware systems.

Workshop 6: INTERNATIONAL WORKSHOP ON SECURE RISC-V (SECRISC-V) ARCHITECTURE DESIGN EXPLORATION

Organizers: Michel A. Kinsy (Arizona State University)

About: Following the very successful and well-attended SECRISC-V'20 and SECRISC-V'21, the Secure RISC-V (SECRISC-V) architecture design exploration workshop is planning for an in-person for 2022. It will have both oral presentations and a poster session. The core theme for SECRISC-V'22 is embedded security, which could serve as an anchor for the in-person portion of the ESWeek 2022. SECRISC-V'22 seeks original research papers on the design, implementation, verification, and evaluation of micro-architecture security features, hardware-assisted security techniques, and secure executions around the RISC-V instruction set architecture (ISA). Submission of early work is encouraged. The RISC-V ISA based topics of specific interest for the workshop include, but are not limited to: Secure cores and multicores, ISA extensions for Security, Software and hardware obfuscation Techniques, Hardware security solutions for machine learning, Secure design for emerging applications: IoT, robotics, wearable computing, etc., Architectural designs and hardware security solutions for HPC, Data Centers and cloud computing, Hardware virtualization and isolation for security, Hardware-Software co-design solutions: graph analytics, Post-quantum cryptosystem designs, Quantum

Computing, Neuromorphic Architectures, Blockchain enabled secure computing, Classic and Modern encryption algorithms and hardware support, Hardware security support for integrity and authentication, key distribution and management, and trust platform modules, Secure execution environment, Memory subsystem organization to secure data accesses, Network-on-Chip (NoC) security feature to process and compute isolation.

Workshop 7: SIGBED x SIGDA Workshop on Emerging Techniques in System Design and Design Automation for Embedded Systems

Program Chairs: Chenren Xu (Peking University); Mingsong Chen (East China Normal University)

Advisory Committee: Xue Liu SIGBED Chair (McGill University); Jie Liu SIGBED China (Chair Harbin Institute of Technology Shenzhen); Yiran Chen SIGDA Chair (Duke University); Sharon Hu Past SIGDA Chair (University of Notre Dame); Qingfeng (Karen) Zhuge (East China Normal University)

Mission and Purpose: For a long time, embedded systems have been recognized as a key discipline including new computer and systems science foundations, new design technology and new hardware and software frameworks. Today, as 5G starts to penetrate our daily life, embedded systems become a key driving force more important than ever: 1) It is the interface connecting the physical world with cyber space and hence serves as a critical entry point cultivating new AIoT technology, application and data; 2) It is typically designed with reduced system complexity (compared with PC and server) for customized application scenario and is therefore easier towards chiplization with electronic design automation tool for high energy-efficiency; 3) Along with the paradigm shift from consumer Internet to industrial Internet, embedded system is extending its traditional application domain (e.g., aerospace, automotive, robotics) to a broader sense including transportation, logistics, manufacturing, health domains. In a nutshell, embedded systems are exhibiting ever-increasing potential in promoting interdisciplinary research, even just starting from different Special Interest Groups (SIGs) in the computing area. The ACM SIGBED Cross-SIG Forum Series #1 – SIGBED x SIGDA workshop is the kickoff attemptation to bring together the experts and expertise from both academia and industry to present the latest results in the intersection of embedded system and design automation. The focus of the workshop is to promote the research idea exchange between these two communities and generate new research direction benefiting both SIGBED and SIGDA.

SYMPOSIA

NOCS:

The International Symposium on Networks-on-Chip (NOCS) is the premier event dedicated to interdisciplinary research on on-chip, package-scale, chip-to-chip, and datacenter rack-scale communication technology, architecture, design methods, applications and systems. NOCS brings together scientists and engineers working on NoC innovations and applications from inter-related research communities, including discrete optimization and algorithms, computer architecture, networking, circuits and systems, packaging, embedded systems, and design automation.

MEMOCODE:

Over the last decade, the boundaries between computer system components, such as hardware, software, firmware, middleware, and applications, have blurred. This evolution in system design and development practices led in 2014 to a change in the title and scope of the MEMOCODE conference from its original focus on hardware/software co-design to its new focus on formal methods and models for developing computer systems and their components. MEMOCODE's objective is to emphasize the importance of models and methodologies in correct system design and development, and to bring together researchers and industry practitioners interested in all aspects of computer system development, to exchange ideas, research results and lessons learned.

ESWEEK | 2022 ORGANIZING COMMITTEE

GENERAL CHAIRS

Aviral Shrivastava
General Chair
Arizona State University, US

Xiaobo Sharon Hu
General Co-Chair
University of Notre Dame, US

PAST CHAIR

Andreas Gerstlauer
Past General Chair
University of Texas at Austin, US

Tulika Mitra
Past Past General Chair
National University of Singapore, SG

PROGRAM CHAIRS

Preeti Ranjan Panda
CASES TPC Chair
IIT Delhi, IN

Swarup Bhunia
CASES TPC Co-Chair
University of Florida, US

Chengmo Yang
CODES+ISSS TPC Chair
University of Delaware, US

Mohammad Al Faruque
CODES+ISSS TPC Co-Chair
University of California, Irvine, US

David Broman
EMSOFT TPC Chair
KTH Royal Institute of Technology, SE

Clarie Pagetti
EMSOFT TPC Co-Chair
Onera / ENSEEIHT, FR

ORGANIZATIONAL CHAIRS

Edwin Sha
ESWEEK Conference Chair
East China Normal University, CN

Sudeep Pasricha
Workshop and Tutorials Chair
Colorado State University, US

Jingtong Hu
Finance Chair
University of Pittsburgh, US

Prabhat Mishra
Special Sessions Chair
University of Florida, US

Qingfeng (Karen) Zhuge
Local Arrangement Chair
East China Normal University, CN

Puneet Gupta
Panel Chair
University of California, Los Angeles, US

Mingsong Chen
Local Arrangement Co-Chair
East China Normal University, CN

Jürgen Teich
Awards Chair
Friedrich-Alexander University Erlangen-Nürnberg, DE

Hoeseok Yang
Virtual Conference Chair
Santa Clara University, US

Akash Kumar
Education Chair
Technische Universität Dresden, DE

ESWEEK | 2022 ORGANIZING COMMITTEE (cont.)

Anup Das
Education Chair
Drexel University, US

Selma Saidi
Brainstorming Sessions Chair
Technical University of Dortmund, DE

Umit Ogras
Student Competition Chair
University of Wisconsin–Madison, US

Tosiron Adegbija
Ph.D. Forum Chair
University of Arizona, US

Amit K. Singh
Publication Chair
University of Essex, UK

Hussam Amrouch
Registration Chair
Stuttgart University, DE

Patricia Derler
Diversity and Inclusion Chair
Palo Alto Research Center, US

Kyoungwoo Lee
Web Chair
Yonsei University, KR

Lars Bauer
Publicity Chair
Karlsruhe Institute of Technology, DE

Shail Dave
Social Media Chair
Arizona State University, US

Liang Shi
Local Publicity Chair
East China Normal University, CN

Anuj Pathania
Media Chair
University of Amsterdam, NL

INDUSTRY ENGAGEMENT TEAM

Dirk Ziegenbein
Industry Engagement Chair
Robert Bosch GmbH, DE

Reiley Jeyapaul
Industry Engagement Team Member
ARM, US

Paul Bogdan
Industry Engagement Team Member
University of Southern California, US

Andreas Herkersdorf
Industry Engagement Team Member
Technical University of Munich, DE

Patricia Derler
Industry Engagement Team Member
Palo Alto Research Center, US

Hugo A. Andrade
Industry Engagement Team Member
Xilinx, US

Ujjwal Gupta
Industry Engagement Team Member
Intel Corporation, US

ESWEEK | 2022 STEERING COMMITTEE

STEERING COMMITTEE CHAIR

Nikil Dutt
University of California, Irvine

STEERING COMMITTEE

Petru Eles
Linköping University

Joerg Henkel
Karlsruhe Institute of Technology

Vinod Kathail
Xilinx

Christoph Kirsch
University of Salzburg

Wang Yi
Uppsala University

CASES | 2022 PROGRAM COMMITTEE

PROGRAM CHAIRS

Preeti Ranjan Panda
Indian Institute of Technology - Delhi

Swarup Bhunia
University of Florida

TRACK CHAIRS

Jana Doppa
Track 1: AI Systems and Applications of AI at Edge
Washington State University

Lars Bauer
Track 2: Embedded Systems and IoT/CPS Security, Safety, Reliability, and Energy-Efficiency
Karlsruhe Institute of Technology

Jingtong Hu
Track 3: Memory and Storage
University of Pittsburgh

Laura Pozzi
Track 4: Accelerators, Emerging Technologies, and Applications
USI Lugano

Christian Pilato
Track 5: Architectures, Compilers, System-level Design
Politecnico di Milano

PROGRAM COMMITTEE

Ali Akoglu
University of Arizona

Lars Bauer
Karlsruhe Institute of Technology

Ganapati Bhat
Washington State University

Oliver Bringmann
University of Tuebingen / FZI

Luca Carloni
Columbia University

Jeronimo Castrillon
TU Dresden

Sandeep Chandran
Indian Institute of Technology Palakkad

Wanli Chang
University of York

Henri-Pierre Charles
CEA

Mainak Chaudhuri
Indian Institute of Technology - Kanpur

Anup Das
Drexel University

Luciana De Micco
Universidad Nacional de Mar del Plata

Jana Doppa
Washington State University

Lide Duan
Alibaba DAMO Academy

Christophe Dubach
Univ. of Edinburgh

William Fornaciari
Politecnico di Milano - DEIB

Amlan Ganguly
Rochester Institute of Technology

Sumitha George
North Dakota State University

CASES | 2022 PROGRAM COMMITTEE (cont.)

Dibakar Gope
ARM

Puneet Gupta
University of California, Los Angeles

Yuko Hara-Azumi
Tokyo Institute of Technology

Joerg Henkel
Karlsruhe Institute of Technology

Jingtong Hu
University of Pittsburgh

Neetu Jindal
Intel

Changhee Jung
Purdue University

Ramesh Karri
New York University

Ryan Kim
Colorado State University

Per Gunnar Kjeldsberg
Norwegian University of Science and Technology

Andreas Krall
TU Wien

Akash Kumar
TU Dresden

Amit Kumar Singh
University of Essex

Jingwen Leng
Shanghai Jiaotong University

Xueqing Li
Tsinghua University

Duo Liu
Chongqing University

Scott Mahlke
University of Michigan

Avinash Malik
University of Auckland

Sai Manoj Pudukotai Dinakarrao
George Mason University

Pietro Mercati
Intel

Krishna Nandivada
IIT Madras

Partha Pratim Pande
Washington State University

Anuj Pathania
University of Amsterdam

Christian Pilato
Politecnico di Milano

Laura Pozzi
USI Lugano

Simon Rokicki
Irisa

Sourav Roy
NXP

Nader Sehatbakhsh
University of California, Los Angeles

Namita Sharma
Intel

Karthik Swaminathan
IBM

Lorenzo Hernández Tabares
Centro de Aplicaciones Tecnológicas y Desarrollo Nuclear (CEADEN)

Hiroyuki Tomiyama
Ritsumeikan University

Amit Trivedi
UIC

CASES | 2022 PROGRAM COMMITTEE (cont.)

Hao Wang
Hygon R&D

Yi Wang
Shenzhen University

Weng-Fai Wong
National University of Singapore

Wei Zhang
Hong Kong University of Science and Technology

PUBLICITY CHAIRS

Duo Liu
Chongqing University

Luciana De Micco
Universidad Nacional de Mar del Plata

Sandeep Chandran
Indian Institute of Technology Palakkad

CODES+ISSS | 2022 PROGRAM COMMITTEE

PROGRAM CHAIRS

Chengmo Yang
University of Delaware

Mohammad Al Faruque
University of California, Irvine

PROGRAM COMMITTEE

Tosiron Adegbija
University of Arizona

David Atienza
École Polytechnique Fédérale de Lausanne (EPFL)

Lars Bauer
Karlsruhe Institute of Technology (KIT)

Christophe Bobda
University of Florida

Paul Bogdan
University of Southern California

Oliver Bringmann
University of Tuebingen / FZI

Rosario Cammarota
Intel

Wanli Chang
Hunan University

Yuan-Hao Chang
Academia Sinica

Xiang Chen
George Mason University

Robert P. Dick
University of Michigan

Rainer Doemer
University of California, Irvine

Jana Doppa
Washington State University

Nikil Dutt
University of California, Irvine

Petru Eles
Linkoping University

Fabrizio Ferrandi
Politecnico di Milano

Franco Fummi
Universita' di Verona

Andreas Gerstlauer
University of Texas, Austin

Tony Givargis
University of California, Irvine

Dimitris Gizopoulos
University of Athens

Michael Glaß
Ulm University

Soonhoi Ha
Seoul National University

Jie Han
University of Alberta

Frank Hannig
Friedrich-Alexander University Erlangen-Nürnberg

Christian Haubelt
University of Rostock

Joerg Henkel
Karlsruhe Institute of Technology

Pi-Cheng Hsiu
Academic Sinica

Reiley Jeyapaul
ARM Research

Weiwen Jiang
George Mason University

Wei Jiang
University of Electronic Science and Technology of China

Ryan Kim
Colorado State University

Fadi Kurdahi
University of California, Irvine

Sébastien Le Beux
Concordia University

Kyoungwoo Lee
Yonsei University

Jenq-Kuen Lee
National Tsing Hua University

Chung-Wei Lin
National Taiwan University

Chen Liu
Intel

Weichen Liu
Nanyang Technological University

Grant Martin
Cadence Design Systems

Hiroki Matsutani
Keio University

Brett Meyer
McGill University

Prabhat Mishra
University of Florida

CODES+ISSS | 2022 PROGRAM COMMITTEE (cont.)

Wolfgang Mueller
University of Paderborn

Daniel Mueller-Gritschneder
Technical University of Munich

Gabriela Nicolescu
École Polytechnique de Montréal

Umit Ogras
University of Wisconsin - Madison

Hyunok Oh
Hanyang University

Alex Orailoglu
UC San Diego

Partha Pratim Pande
Washington State University

Sri Parameswaran
University of New South Wales

Sudeep Pasricha
Colorado State University

Andy Pimentel
University of Amsterdam

Ilia Polian
University of Stuttgart

Graziano Pravadelli
University of Verona

Keni Qiu
Capital Normal University

PUBLICITY CHAIR

Weiwen Jiang
George Mason University

Francesco Regazzoni
University of Amsterdam

Marco D. Santambrogio
Politecnico di Milano

Gunar Schirner
Northeastern University

Rishad Shafik
Newcastle University

Zili Shao
The Chinese University of Hong Kong

Liang Shi
East China Normal University

Amit Kumar Singh
University of Essex

Todor Stefanov
Leiden University

Gregory Stitt
University of Florida

Sander Stuijk
Eindhoven University of Technology

Cheng Tan
Microsoft

Jürgen Teich
University of Erlangen-Nuremberg

Ishan Thakkar
University of Kentucky

Hiroyuki Tomiyama
Ritsumeikan University

Chao Wang
University of Science and Technology of China

Fei Wu
Huazhong University of Science and Technology

Jiafeng Xie
Villanova University

Jiang Xu
Hong Kong University of Science and Technology

Jason Xue
City University of Hong Kong

Ming-Chang Yang
The Chinese University of Hong Kong

Mengying Zhao
Shandong University

Peipei Zhou
University of Pittsburgh

EMSOFT | 2022 PROGRAM COMMITTEE

PROGRAM CHAIRS

David Broman
KTH Royal Institute of Technology

Claire Pagetti
Onera / ENSEEIHT

PROGRAM COMMITTEE

James Andersson
University of North Carolina at Chapel Hill

Stanley Bak
Stony Brook University

Sanjoy Baruah
Washington University in St. Louis

Jan Beutel
University of Innsbruck

Alessandro Biondi
Scuola Superiore Sant'Anna

Timothy Bourke
Inria

Björn Brandenburg
Max Planck Institute for Software Systems (MPI-SWS)

Franck Cassez
Macquarie University

Samarjit Chakraborty
University of North Carolina at Chapel Hill

Li-Pin Chang
National Yang Ming Chiao Tung University

Sudipta Chattopadhyay
Singapore University of Technology and Design

Jian-Jia Chen
TU Dortmund

Chih-Hong Cheng
Fraunhofer IKS

Remi Delmas
Amazon Web Services

Patricia Derler
Kontrol GmbH

Petru Eles
Linköping University

Rolf Ernst
TU Braunschweig

Heiko Falk
Hamburg University of Technology

Chris Gill
Washington University in St. Louis

Alain Girault
Inria

Guy Gogniat
University of South Brittany

Susanne Graf
CNRS/Verimag

Nan Guan
City University of Hong Kong

Ichiro Hasuo
National Institute of Informatics

Oana Hotescu
ISAE-SUPAERO

Yu Jiang
Tsinghua University

Hokeun Kim
Hanyang University and UC Berkeley

Hanjun Kim
Yonsei University

EMSOFT | 2022 PROGRAM COMMITTEE (cont.)

Shan Lin
Stony Brook University

Franck Mamalet
IRT Saint Exupéry

Tulika Mitra
National University of Singapore

Frank Mueller
North Carolina State University

Geoffrey Nelissen
Eindhoven University of Technology

André Platzer
Carnegie Mellon University

Indranil Saha
IIT Kanpur

Selma Saidi
TU Dortmund

Martin Schoeberl
Technical University of Denmark

Donatella Sciuto
Polytechnic University of Milan

Stavros Tripakis
Northeastern University

Matthias Wählisch
Freie Universität Berlin

Bryan Ward
MIT Lincoln Laboratory

Haibo Zeng
Virginia Tech

WIP REVIEW COMMITTEE

Thomas Carle
IRIT, France

Mélanie Ducoffe
Airbus, France

Gerhard Fohler
TU Kaiserslautern, Germany

Divya Gopinath
KBR Inc., NASA AMES, USA

Mohamed Hassan
McMaster University, Canada

Meiyi Ma
Vanderbilt University, USA

Catherine Nemitz
Davidson College, USA

Mickael Randour
University of Mons, Belgium

Jinghao Sun
Dalian University of Technology, China

PUBLICITY CHAIR

Jan Mathieu
CEA Tech

NOCS 2022 | ORGANIZING COMMITTEE

GENERAL CHAIRS

Sergi Abadal *(Universitat Politècnica de Catalunya)*
Joshua San Miguel *(University of Wisconsin-Madison)*

TECHNICAL PROGRAM CHAIRS

Miquel Moreto *(Barcelona Supercomputing Center)*
Mahdi Nikdast *(Colorado State University)*

SPECIAL SESSIONS CHAIR

Masoumeh (Azin) Ebrahimi *(KTH Royal Institute of Technology)*

PUBLICITY CHAIRS

Jieming Yin *(Lehigh University)*
Poona Bahrebar *(Ghent University)*
Trevor Carlson *(National University of Singapore)*

VIRTUAL LOGISTICS CHAIRS

Alirez Monemi *(Barcelona Supercomputing Center)*
Abhijit Das *(IRISA, INIRA, and Université de Rennes)*

WEB CHAIR

Salvatore Monteleone (*Niccoló Cusano University)*

STEERING COMMITTEE CHAIR

Radu Marculescu (*University of Texas, Austin)*

TECHNICAL PROGRAM COMMITTEE

Shaahin Angizi *(New Jersey Institute of Technology)*
Poona Bahrebar *(Ghent University)*
Jonathan Balkind *(UCSB)*
Davide Bertozzi *(Universitá degli Studi di Ferrara)*
Sanmitra Banerjee *(NVIDIA)*
José Cano *(University of Glasgow)*
Luca Carloni *(Columbia University)*
Felipe Gohring de Magalhães *(Polytechnique Montréal)*
Paul Gratz *(Texas A&M University)*
Robert Hesse *(Cerebras)*
Biresh Kumar Joardar *(Duke)*
Nachiket Kapre *(University of Waterloo)*
Avinash Karanth *(Ohio University)*
Cédric Killian *(IRISA)*
John Kim *(KAIST)*
Ryan Kim *(Colorado State University)*
Gwangsun Kim *(POSTECH)*
EJ Kim *(TAMU)*
Michihiro Koibuchi *(National Institute of Informatics, Japan)*
Sushant Kondguli (*Meta, Facebook)*
Somayyeh Koohi *(Sharif University of Technology)*
Jinho Lee *(Yonsei)*
Xiongfei Liao *(MediaTek)*
Zhonghai Lu *(KTH Royal Institute of Technology)*
Hiroki Matsutani *(Keio University)*
Syed Minhaj Hassan *(Qualcomm)*
Mondrian Nuessle *(Extoll)*
Maurizio Palesi *(University of Catania)*
Partha Pratim Pande *(Washington State University)*
Mayank Parasar *(Samsung)*
Semeen Rehman *(TU Wien)*
Arman Roohi *(University of Nebraska – Lincoln)*
Hamid Sarbazi-Azad *(Sharif University of Technology)*
Jens Sparsø *(Technical University of Denmark)*
Tomohiro Ueno *(Riken)*
Enrique Vallejo *(University of Cantabria)*
Jiang Xu *(Hong Kong University of Science and Technology)*
Yaoyao Ye *(Shanghai Jiao Tong University)*

MEMOCODE 2022

ORGANIZATION

GENERAL CO-CHAIRS
Prof. Reinhard von Hanxleden, *Christian-Albrecht University of Kiel, Germany*
Prof. Dominique Méry, *LORIA, Université de Lorraine, France*

PROGRAMME COMMITTEE CO-CHAIRS
Prof. Jens Brandt, *Niederrhein University of Applied Sciences, Germany*
Prof. Lijun Zhang, *Institute of Software, CAS*

PUBLICITY CHAIR
Bai Xue, *Institute of Software, CAS*

PROGRAM COMMITTEE

Shuvra Bhattacharyya, *University of Maryland*
Sudipta Chattopadhyay, *Singapore University of Technology and Design*
Yu-Fang Chen, *Academia Sinica*
Thao Dang, *CNRS/VERIMAG*
Patricia Derler, *Palo Alto Research Center*
Stephen A. Edwards, *Columbia University*
Martin Fränzle, *Carl von Ossietzky Universität Oldenburg*
Goran Frehse, *ENSTA Paris*
Holger Hermanns, *Saarland University*
Eunsuk Kang, *Carnegie Mellon University*
Rick Kuhn, *Natl Institute of Standards & Technology*
Matthew Kuo, *Auckland University of Technology*
Luciano Lavagno, *Politecnico di Torino*
Elizabeth Leonard, *Naval Research Laboratory*
Xuandong Li, *Nanjing University*
Stefan Mitsch, *Carnegie Mellon University*
Katell Morin-Allory, *TIMA Laboratory*
Luan Nguyen, *University of Dayton*
Pierluigi Nuzzo, *University of Southern California*
Roberto Passerone, *University of Trento*
Doron Peled, *Bar Ilan University*
Pavithra Prabhakar, *Kansas State University*
Sanjiva Prasad, *Indian Institute of Technology Delhi*
Sven Schewe, *University of Liverpool*
Klaus Schneider, *University of Kaiserslautern*
Sandeep Shukla, *Indian Institute of Technology Kanpur (IIT Kanpur)*
Marjan Sirjani, *Malardalen University, Reykjavik University*
Jürgen Teich, *University of Erlangen-Nuremberg*
Ashutosh Trivedi, *University of Colorado Boulder*
Tiziano Villa, *Universita' di Verona*
Ji Wang, *National Laboratory for Parallel and Distributed Processing*
Tuba Yavuz, *University of Florida*
Paolo Zuliani, *Newcastle University*

VOLUNTEERS

Shailja Pandey (IIT Delhi)
Aritra Bhattacharya (University of Florida)
Rozhin Yasaei (UC Irvine)
Fanruo Meng (University of Delaware)
Tommy White (University of Delaware)
Yuhao Liu (TU Dresden)
Siddharth Gupta (TU Dresden)
Xuan Xie (University of Alberta)
Alireza Monemi (BSC)
Abhijit Das (INRIA)
Hanwei Zhang (LIS Marseille)
Adam Awale (Arizona State University)
Quoc Long Vinh Ta (Arizona State University)
Ashwin Ramaswamy (Arizona State University)
Jun Xia (East China Normal University)
Pengyu Zhang (East China Normal University)
Yingbo Zhou (East China Normal University)
Xiaoyang Zhang (East China Normal University)

EMBEDDED SYSTEMS WEEK

7-14 October 2022 | Shanghai, China + Phoenix, US + Virtual