Research Challenges in Computer Science 2022
Tuesday, January 18, 2022, 13:50 – 19:45

Webex: https://centralntua.webex.com/centralntua/j.php?MTID=m2094443eb46a64ad98af14381c2ae27c

YouTube Live Streaming: https://youtu.be/zeRLIZmdXNY

The *Division of Computer Science* (https://www.cs.ntua.gr) of the School of Electrical and Computer Engineering at the National Technical University of Athens and the *MSc Program on Data Science and Machine Learning* (https://dsml.ece.ntua.gr) organize a series of talks on recent research advances and challenges in Computer Science. The event will host seven distinguished researchers and will take place virtually, via **Webex**, on **Tuesday, January 18, 2022**, at the following link:

https://centralntua.webex.com/centralntua/j.php?MTID=m2094443eb46a64ad98af14381c2ae27c

with public broadcasting through **YouTube**: https://youtu.be/zeRLIZmdXNY

Schedule

13:50 – 14:00	**Welcome - Opening**
14:00 – 14:45	**Memory-Centric Computing** Onur Mutlu, *ETH Zürich*
14:45 – 15:30	**Decentralizing Information Technology: The Advent of Resource Based Multiparty Protocols** Aggelos Kiayias, *University of Edinburgh & IOHK*
15:30 – 16:15	**Advances in Self-Supervised Learning and Measuring Generalization** Yannis Kalantidis, *NAVER LABS Europe*
16:15 – 16:30	Short break
16:30 – 17:15	**Adventures in Property-Based Testing** Leonidas Lampropoulos, *University of Maryland*
17:15 – 18:00	**The Mayhem Cyber Reasoning System** Thanassis Avgerinos, *ForAllSecure*
18:00 – 18:15	Short break
18:15 – 19:00	**Data-Centric Debugging in Machine Learning** Theo Rekatsinas, *Apple & ETH Zürich*
19:00 – 19:45	**Automated Neural Network Design under Hardware Constraints** Dimitrios Stamoulis, *Microsoft AI*

Abstracts and Bios of the Speakers

Title: Memory-Centric Computing
Speaker: Onur Mutlu, *ETH Zürich*

Abstract: Computing is bottlenecked by data. Large amounts of application data overwhelm storage capability, communication capability, and computation capability of the modern machines we design today. As a result, many key applications' performance, efficiency and scalability are bottlenecked by data movement. In this lecture, we describe three major shortcomings of modern architectures in terms of 1) dealing with data, 2) taking advantage of the vast amounts of data, and 3) exploiting different semantic properties of application data. We argue that an intelligent architecture should be designed to handle data well. We show that handling data well requires designing architectures based on three key principles: 1) data-centric, 2) data-driven, 3) data-aware. We give several examples for how to exploit each of these principles to design a much more

efficient and high performance computing system. We especially discuss recent research that aims to fundamentally reduce memory latency and energy, and practically enable computation close to data, with at least two promising novel directions: 1) processing using memory, which exploits analog operational properties of memory chips to perform massively-parallel operations in memory, with low-cost changes, 2) processing near memory, which integrates sophisticated additional processing capability in memory controllers, the logic layer of 3D-stacked memory technologies, or memory chips to enable high memory bandwidth and low memory latency to near-memory logic. We show both types of architectures can enable orders of magnitude improvements in performance and energy consumption of many important workloads, such as graph analytics, database systems, machine learning, video processing. We discuss how to enable adoption of such fundamentally more intelligent architectures, which we believe are key to efficiency, performance, and sustainability. We conclude with some guiding principles for future computing architecture and system designs.

A short accompanying paper, "*Intelligent Architectures for Intelligent Computing Systems*", which appeared in DATE 2021, can be found here and serves as recommended reading:
https://people.inf.ethz.ch/omutlu/pub/intelligent-architectures-for-intelligent-computingsystems-invited_paper_DATE21.pdf

Bio: ***Onur Mutlu*** (http://people.inf.ethz.ch/omutlu) is a Professor of Computer Science at ETH Zürich. He is also a faculty member at Carnegie Mellon University, where he previously held the Strecker Early Career Professorship. His current broader research interests are in computer architecture, systems, hardware security, and bioinformatics. A variety of techniques he, along with his group and collaborators, has invented over the years have influenced industry and have been employed in commercial microprocessors and memory/storage systems. He obtained his PhD and MS in ECE from the University of Texas at Austin and BS degrees in Computer Engineering and Psychology from the University of Michigan, Ann Arbor. He started the Computer Architecture Group at Microsoft Research (2006-2009), and held various product and research positions at Intel Corporation, Advanced Micro Devices, VMware, and Google. He received the IEEE High Performance Computer Architecture Test of Time Award, the IEEE Computer Society Edward J. McCluskey Technical Achievement Award, ACM SIGARCH Maurice Wilkes Award, the inaugural IEEE Computer Society Young Computer Architect Award, the inaugural Intel Early Career Faculty Award, US National Science Foundation CAREER Award, Carnegie Mellon University Ladd Research Award, faculty partnership awards from various companies, and a healthy number of best paper or "Top Pick" paper recognitions at various computer systems, architecture, and security venues. He is an ACM Fellow "for contributions to computer architecture research, especially in memory systems", IEEE Fellow for "contributions to computer architecture research and practice", and an elected member of the Academy of Europe (Academia Europaea). His computer architecture and digital logic design course lectures are freely available on YouTube (https://www.youtube.com/OnurMutluLectures), and his research group makes a wide variety of software and hardware artifacts freely available online (https://safari.ethz.ch).

Title: Decentralizing Information Technology: The Advent of Resource Based Multiparty Protocols
Speaker: Aggelos Kiayias, *University of Edinburgh & IOHK*

Abstract: The bitcoin blockchain, introduced more than a decade ago, gave the first instance of a multiparty protocol that maintains its security via the incentive driven participation of a fluctuating set of holders of a particular resource: computational power. This raises the question whether it can be possible to realise any multiparty functionality out of the self-interest of computer node operators who enroll themselves to support the system's operation in exchange of rewards that are provided in the system's digital currency. In this talk we cast this as a general paradigm for designing and deploying multiparty protocols. We give two examples of this paradigm: the Ouroboros protocol as deployed in the Cardano blockchain and the Nym Mix-net. We discuss design challenges, solutions and open questions as well as we look at what lies ahead in decentralizing information technology services.

Bio: ***Aggelos Kiayias*** (https://www.kiayias.com) is chair in Cyber Security and Privacy and director of the Blockchain Technology Laboratory at the University of Edinburgh. He is also the Chief Scientist at blockchain technology company IOHK and an associate professor at the National and Kapodistrian University of Athens.

His research interests are in computer security, information security, applied cryptography and foundations of cryptography, with a particular emphasis in blockchain technologies and distributed systems, e-voting and secure multiparty protocols, as well as privacy and identity management. He has been the recipient of an ERC fellowship, a Marie Curie fellowship, an NSF Career Award, and a Fulbright Fellowship. He holds a Ph.D. from the City University of New York, he was a graduate student in ECE at NTUA and has an undergraduate degree in Mathematics from the University of Athens. He has over 150 publications in journals and conference proceedings in the area. He has served as the program chair of the Cryptographers' Track of the RSA conference in 2011 and the Financial Cryptography and Data Security conference in 2017, as well as the general chair of Eurocrypt 2013. He also served as the program chair of Real World Crypto Symposium 2020 and the Public-Key Cryptography Conference 2020. In 2021, he was elected fellow of the Royal Society of Edinburgh.

Title: Advances in Self-Supervised Learning and Measuring Generalization
Speaker: Yannis Kalantidis, *NAVER LABS Europe*

Abstract: Self-supervised learning is a highly effective way of learning representations that are useful for (or, generalise to) a wide range of downstream tasks and datasets. After introducing self-supervised learning, the first part of the talk presents MoCHi (NeurIPS 2020) a contrastive self-supervised learning method for transferable visual representations that is able to improve and speed-up the learning process by synthesising hard negatives in feature space. The second part of the talk introduces TLDR (ArXiv 2021), a method that leverages advances in self-supervised visual representation learning for the case of arbitrary vector input spaces and the task of dimensionality reduction, i.e. for compressing an already meaningful representation. The third and final part of the talk introduces ImageNet-CoG (ICCV 2021) a novel benchmark that aims at measuring concept generalization. We argue that semantic relationships between seen and unseen concepts affect generalization performance and propose a novel benchmark on the extended ImageNet-21K dataset that can evaluate models trained on the ubiquitous ImageNet-1K dataset out-of-the-box. In our extensive study, we benchmark over thirty publicly available models under the prism of concept generalization, and show how our benchmark is able to uncover a number of interesting insights.

Bio: ***Yannis Kalantidis*** (http://www.skamalas.com) is a senior Research Scientist at NAVER LABS Europe. He got his PhD in 2014 from the National Technical University of Athens. He was a postdoc and research scientist at Yahoo Research in San Francisco from 2015 until 2017, leading the visual similarity search project at Flickr and conducting research on visual representation learning and vision & language. From 2017 to 2019 he was a research scientist at Facebook AI in Menlo Park, where his research interests expanded to video understanding and deep learning architecture modeling. He joined NAVER LABS Europe in Grenoble in March 2020. His research interests currently revolve around representation learning under limited supervision and resources, as well as adaptive multi-modal systems. He is also passionate about bringing the computer vision community closer to socially impactful tasks, datasets and applications for worldwide impact and has co-organized relevant workshops like "Computer Vision for Global Challenges" (CV4GC @ CVPR 2019), "Computer Vision for Agriculture" (CV4A @ ICLR 2020) and "Wikipedia and Multi-Modal & Multi-Lingual Research" (Wiki-M3L @ ICLR 2022).

Title: Adventures in Property-Based Testing
Speaker: Leonidas Lampropoulos, *University of Maryland*

Abstract: QuickCheck-style property-based testing occupies an interesting middle ground among software correctness techniques: it is relatively lightweight and accessible, but much more expressive than conventional unit testing dealing with logical specifications rather than simple examples. I'll describe two ongoing research threads combining property-based testing with (1) coverage-guided fuzzing in the style of tools like AFL, and (2) ideas from the combinatorial testing literature. Both yield significant speedups in realistic testing scenarios.

Bio: ***Leonidas Lampropoulos*** (https://lemonidas.github.io) is an assistant professor of Computer Science at the University of Maryland, College Park. Before that, he was a Victor Basili postdoctoral fellow jointly between UMD and UPenn, under the supervision of Prof. Michael Hicks and Prof. Benjamin Pierce. He got his Ph.D. at

the University of Pennsylvania under Prof. Pierce. His research interests lie in programming languages, with an emphasis on software correctness through both random testing and verification. He is the principal author of the fourth volume in the popular Software Foundations series of online textbooks: "QuickChick: Property-Based Testing in Coq".

Title: The Mayhem Cyber Reasoning System
Speaker: Thanassis Avgerinos, *ForAllSecure*

Abstract: Mayhem is one of the first generation of autonomous computer security bots that finds and fixes vulnerabilities without any human intervention. Mayhem won the DARPA Cyber Grand Challenge (CGC) contest and $2,000,000 in August 2016 against six other finalists. The contest was the result of a two-year DARPA program, but the R&D necessary to compete stands on the shoulders of decades of basic academic and industry scientific research in program analysis, verification, and self-healing systems. The Mayhem system alone has been in development for over a decade now in academia and industry, starting in Carnegie Mellon University and then spinning off to a company named ForAllSecure. Mayhem is now being commercialized by ForAllSecure to autonomously check and protect the world's software from exploitable bugs. In this talk, we look back and give our story in creating Mayhem, discuss key techniques we discovered along the way, and also look forward to a vision where autonomous security bots like Mayhem will radically improve the security of computer systems.

Bio: ***Thanassis Avgerinos*** (https://users.ece.cmu.edu/~aavgerin) is an expert in program analysis, testing, and software security with over a decade of operational and academic experience. He is a founder and the VP of Engineering of ForAllSecure, a cybersecurity startup. Prior to co-founding ForAllSecure, he was a researcher at Carnegie Mellon University, working on developing the first Mayhem prototype, a system for autonomous cybersecurity. Thanassis holds both a Ph.D and Master's degrees from Carnegie Mellon University and a Master's and Bachelor's degrees from the National Technical University of Athens, all in Electrical and Computer Engineering.

Title: Data-Centric Debugging in Machine Learning
Speaker: Theo Rekatsinas, *Apple & ETH Zürich*

Abstract: Data is the backbone of modern Machine Learning but access to high-quality data is a key bottleneck for obtaining reliable deployments. High-effort tasks such as data validation and cleaning are essential to learning accurate and non-biased models. In this talk, I will discuss how to automate routine data validation tasks such as missing value imputation and detection of corrupted samples. First, I will discuss how one can leverage structured, statistical dependencies in the data to obtain information theoretically optimal data preparation methods, and then I will demonstrate how the widely-used Attention mechanism is key to automated data validation. This talk builds upon experience with projects such as HoloClean, FDX, and Picket and their application to different scientific and industrial use-cases.

Bio: ***Theodoros (Theo) Rekatsinas*** (https://thodrek.github.io) is currently at Apple where he is leading the Knowledge Platform – Graph ML team. He is also an Assistant Professor in the Department of Computer Science at ETH Zürich and part of the Systems Group. Previously, Theo was an Assistant Professor at the University of Wisconsin-Madison and a member of the Database Group. Theo is also a co-founder of Inductiv (acquired by Apple), a company developing AI for identifying and correcting errors in data.

Title: Automated Neural Network Design under Hardware Constraints
Speaker: Dimitrios Stamoulis, *Microsoft AI*

Abstract: Deep Neural Networks (DNNs) have traditionally been designed by human experts in a painstaking process. However, the demand for cutting-edge performance and real-world deployment has resulted in increasingly complex models, making the manual DNN design a daunting task. AutoML aims at alleviating this

engineering burden by automatically identifying the DNN hyperparameters (e.g., number of layers, type of layer-wise operations). This talk examines state-of-the-art AutoML methods and discusses their key properties: search-cost efficiency and hardware-awareness across various hardware requirements (e.g., DNN latency, memory, power, energy consumption).

Bio: ***Dimitrios Stamoulis*** (http://dimitriosstamoulis.com) is a Senior Research Manager at Microsoft AI (Mixed Reality). His research interests span various AutoML methods to automate the design of object detection models under hardware constraints. He received a PhD in Electrical and Computer Engineering from Carnegie Mellon University, USA, under the supervision of Diana Marculescu, and a Masters in ECE from McGill University, Canada. He received a Diploma in ECE from NTUA, Greece, completing his undergraduate thesis with MicroLab and under the supervision of Dimitrios Soudris.