80th Meeting of the IFIP 10.4 Working Group on Dependable Computing and Fault Tolerance

Virtual - 25 June 2021 — 27 June 2021

Theme: Topics at the Intersection of AI & Dependability

Organizer: Mootaz Elnozahy

Program at a Glance

All times are in UTC

June 25, 2021

12:00—13:00	Virtual reception
13:00—13:40	On Monitoring and resiliency for machine-learning-based autonomous systems—*Michael Paulitsch*
13:40—14:20	Security and privacy for distributed optimization and learning—*Nitin Vaidya*
14:20—15:00	Runtime Assurance of Distributed Intelligent Control Systems—*Yair Amir*
15:00—15:30	Panel Discussion

June 26, 2021

12:00—13:00	Virtual reception
13:00—13:40	Only One Percent Rubbish—*Roy Maxion*
13:40—14:20	Safety and security in AI enabled critical applications: Who is going to solve the problems?—*Henrique Madeira*
14:20—15:00	Intrusion Detection through Unsupervised Machine Learning: pros, limitations and workarounds—*Andrea Bondavalli*
15:00—15:30	Panel Discussion

June 27, 2021 Business Meeting

Part 1: Status

13:00—13:10	State of the WG 10.4 — *Mootaz Elnozahy*
13:10—13:20	Report on DSN-21 — *Sy-Yen Kuo*
13:20—13:30	Report on DSN-22 — *Yair Amir & Cristina Nita-Rotaru*

Part 2: Administrative

13:30—14:10	New Members *Kishore Trivedi for Dan Dongseong Kim* *Stephen Poledna for Wilfried Steiner* *Andrea Bondavalli for Andrea Ceccarelli* *Evgenia Smirni for Giuliano Casale (Tentative)*
14:10—14:45	Future Meetings *Group Discussion, Mootaz Elnozahy moderator*
14:45—15:00	Updated from TC-10 and IFIP, call for organization *Mootaz Elnozahy*

Part 3: Research Updates

15:00—15:15	IVDS Update— *Jay Lala*
15:15—15:30	Short Talks (tentative)

Workshop Talks and Speakers

June 25: 13:00—13:40

On monitoring and resiliency for machine-learning-based autonomous systems

Today's machine-learning-based autonomous systems have a natural resiliency against some types of faults. Despite some "natural" resiliency, establishing trust for autonomous safety-critical systems is not easy to establish. This talks looks at basic approaches and evaluations of resiliency against platform-based faults and potential mitigations. It also shows techniques of different types of system architectures using monitoring to increase trust in such systems and their application in autonomous systems.

Michael Paulitsch brings 20 years of work theoretical and applied research and technology work at university and different industries (aerospace, railway, automotive) in dependability in safety-critical and real-time systems including security aspects of all types.

Michael is Principal Engineer at Intel and files the role of a Dependability Systems Architect (Principal Engineer) at Intel, Munich, Germany, and leads the Dependability Research Lab since 2018. He pursues Dependable Machine Learning systems (resiliency), evaluates and ensures safe and dependable use of neural network models in safety-critical systems. He also looks at novel safety monitoring approaches at different system levels (chip, platform, application). From 2014 to 2018, he has been senior engineering manager and product line manager for Vital Platform (safety-critical computing and communication platform with security requirements) at Thales Ground Transportation Systems in Vienna, Austria. In these roles he has been responsible for the execution and strategy of as well as research for this vital platform. Before this, Michael has been Senior Expert of "Dependable Computing and Networks" as well as Scientific Director at Airbus corporate research in Munich, Germany. There his work focused on dependable embedded and secure embedded computing and networks. From 2003 to 2008, he worked at Honeywell Aerospace in the U.S. on software and electronic platforms in the area of business, regional, air transport, and human space avionics and engine control electronics. Michael has also been assistant professor at Technische Universitaet Wien, Vienna, Austria, 1997 to 2003. Michael published 50+ scientific papers in his area of expertise, participates in multiple international scientific conference committees and holds 35+ patents. He holds a doctoral degree in technical sciences from the Vienna University of Technology, Vienna, Austria with emphasis on dependable embedded systems and a doctoral degree in economics and social sciences with emphasis on production management aspects. He also visited University of Illinois at Urbana-Champaign.

June 25: 13:40—14:20

Security and Privacy for Distributed Optimization and Learning

Consider a network of agents wherein each agent has a private cost function. In the context of distributed machine learning, the private cost function of an agent may represent the “loss function” corresponding to the agent’s local data. The objective here is to identify parameters that minimize the total cost over all the agents. In machine learning for classification, the cost function is designed such that minimizing the cost function should result in model parameters that achieve higher accuracy of classification. Similar optimization problems arise in the context of other applications as well.

Our work addresses privacy and security of distributed optimization, with applications to machine learning. In privacy-preserving machine learning, the goal is to optimize the model parameters correctly while preserving the privacy of each agent’s local data. In security, the goal is to identify the model parameters correctly while tolerating adversarial agents that may be supplying incorrect information. When a large number of agents participate in distributed optimization, security compromise or failure of some of the agents becomes increasingly likely. The talk will provide intuition behind the design and correctness of the algorithms.

Nitin Vaidya is the McDevitt Chair of Computer Science at Georgetown University. He received his Ph.D. from the University of Massachusetts at Amherst. He previously served as a Professor and Associate Head in Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. He has co-authored papers that received awards at several conferences, including 2015 SSS, 2007 ACM MobiHoc and 1998 ACM MobiCom. He is a fellow of the IEEE. He has served as the Chair of the Steering Committee for the ACM PODC conference, as the Editor-in-Chief for the IEEE Transactions on Mobile Computing, and as the Editor-in-Chief for ACM SIGMOBILE publication MC2R.

June 25: 14:20—15:00

RADICS: Runtime Assurance of Distributed Intelligent Control Systems

We describe RADICS: Runtime Assurance of Disrtributed Intelligent Control Systems, which combines a Simplex-based, black-box monitor with a white-box monitor to ensure correct behavior and good performance of AI systems. The black-box monitor allows the system to detect when the AI controller is on a failing trajectory and use a provably safe, but less performant algorithm, to right the system. The white-box monitor predicts when the AI controller will be put on such a trajectory before it happens and helps maximize the performance of the overall system. We describe the overall approach in detail and implement a simple version of it on a case study into controlling the lights in a small traffic grid.

Joint work with Brian Wheatman, Jerry Chen and Tamim Sookoor

Yair Amir is Professor of Computer Science, and director of the Distributed Systems and Networks lab (www.dsn.jhu.edu) at Johns Hopkins University. He served as Department Chair of Computer Science at Johns Hopkins (2015-2018), as Vice Chair of the IFIP 10.4 Working Group on Dependable Computing (2016-2018), and as Program co-Chair of the 2015 IEEE/IFIP Dependable Systems and Networks (DSN) conference.

He is a creator of the Spread toolkit (www.spread.org), the first scalable group communication system with strong semantics, the Spines overlay network platform (www.spines.org), the Prime Byzantine replication engine, the first to provide performance guarantees while under attack, and the Spire intrusion-tolerant SCADA for the power grid (www.dsn.jhu.edu), the first to protect against both system-level and network-level attacks and compromises. Some of these technologies are deployed in mission critical systems, support data center applications, are included in commercial products, and are used for research and teaching in universities and research labs around the world.

Dr. Amir is a co-founder of LTN Global Communications (www.ltnglobal.com), a company that offers a global transport service for broadcast-quality live TV that is used by major broadcasters including CNN, Fox, Disney, ABC, Bloomberg, CBS, CNBC, ESPN, NBC, PBS, and Turner.

Dr. Amir holds B.Sc. (1985) and M.Sc. (1990) from the Technion, Israel Institute of Technology, and a Ph.D (1995) from the Hebrew University of Jerusalem, Israel.

June 26: 13:00—13:40

Only One Percent Rubbish

AI/ML systems "learn" to make decisions based on the data with which they are trained. Such systems are often used to make critical decisions in which mistakes can have serious consequences -- e.g., systems for approving credit, job and college applications, digital forensic procedures, and computer-user authentication. In these kinds of applications AI/ML decision algorithms are tasked with distinguishing between legitimate and fraudulent or wrong behavior.

We show that minor degradations in as little as 1-2 percent of the training data can change decision outcomes by nearly 20 percent, wrongly reversing distinctions between legitimacy and fraudulence. In one real-world application -- user authentication -- data corruption was induced by USB keyboards injecting artifacts into the data, effecting an infidelity to the true signal. We illustrate how this phenomenon was discovered and validated.

Roy Maxion is a Research Professor in computer science and machine learning at Carnegie Mellon University, where he is also the director of the Dependable Systems Laboratory. He has long been a passionate proponent of rigorous/foundational scientific methodology. He is an IEEE Fellow, and recently served as a member of the US National Academy of Sciences committee on Future Research Goals and Directions for Foundational Science in Cybersecurity. He is one of the founding members of the US Center for Statistics and Applications in Forensic Evidence.

June 26: 13:40—14:20

Safety and security in AI enabled critical applications: who is going to solve the problems?

The intensive use of Artificial Intelligence and, more specifically, Machine Learning in safety critical applications, especially in the context of cyber-physical system (CPS), involves a considerable number of grand challenges, most of them related to dependability, safety and security. The goal of the presentation is to discuss what a (classic) dependability research community, represented by IFIP 10.4 working group, should and can do to successfully address the fascinating research opportunities resulting from the use of AI in safety critical applications. The presentation will show different examples that illustrate the possibilities, and especially the difficulties and limitations, of using known recipes to assure safety and security in AI enabled critical applications. The problem is twofold: on the one hand we need to architect the safe and secure integration of AI technology in safety-critical functions in new resource demanding contexts and, on the other hand, we must demonstrate that the use of AI in safety-critical functions is compatible with strong safety and security properties. The second problem is far from being a classic verification and validation problem and, most likely, is the most demanding task to successfully use AI in safety critical scenarios.

Henrique Madeira is full professor at the University of Coimbra, where he has been involved in research on dependable computing since 1989. His research interests include software quality and software reliability, experimental evaluation and benchmarking of dependability and security, and fault injection techniques. His recent research projects involve two research directions: a) Assured AI, focusing on providing safety and security guaranties in critical applications that use AI and b) human factors in software engineering, particularly on the use of biometrics to improve software quality.
He has participated in more than 20 research projects funded by the European Commission and by the Portuguese Government and coordinated several projects, such as the AMBER (Assessing, Measuring, and Benchmarking Resilience) IST-FP7 project. Henrique was the Vice-Chair of the IFIP Working Group 10.4 Special Interest Group (SIG) on Dependability Benchmarking from the establishment of the SIG in the summer of 1999 until 2002. He was Program Co-Chair of the International Performance and Dependability Symposium track of the IEEE/IFIP International Conference on Dependable Systems and Networks, DSN-PDS2004, and has organized several Workshops and scientific events. He was Conference Coordinator of the IEEE/IFIP DSN, the major conference of the dependability area, in 2008 and Co-General Chair of the IEEE 31st International Symposium on Software Reliability Engineering (ISSRE) in 2020. Henrique Madeira is co-developer of G-SWIFIT, the technique to inject software bugs that has been adopted by most of researchers and in the field. He is also co-developer of Xception, a software implemented fault injection tool that has been used in several universities, companies and space agencies such as NASA (USA), ESA (Europe), JAXA (Japan) e CAST (China). Henrique Madeira was co-founder of the company Critical Software, S.A. (www.criticalsoftware.com).

June 26: 14:20—15:00

Intrusion Detection through Unsupervised Machine Learning: pros, limitations and workarounds

It is undeniable that new cyber-attacks are continuously crafted against essentially any kind of system and service. Systems are subject to a mix of usual practiced attacks and new ones that were not previously known, motivating the need for building Intrusion Detectors (IDs) that can effectively deal with those zero-day attacks. Different studies have been devised Unsupervised Machine Learning (ML) algorithms belonging to different families as clustering, neural networks, density-based, neighbor-based, statistical, and classification. Those algorithms have the potential to detect even unknown threats thanks to a training phase that does not rely on labels in data. The talk shows how different algorithms are better suited for the detection of specific anomalies of system indicators, which manifest when attacks are conducted against a system. Unfortunately, those algorithms show inferior detection performance of known threats with respect to supervised ML algorithms ; to fill this gap, we show improvements achieved when adopting Meta-Learning techniques. In any case, the quality of the best solution that can be devised depends strongly on the problem at hand and demands for high cost for selecting and finding the optimal set up of Unsupervised algorithms. To this end, we conclude the talk by proposing a cheap method to quantitatively understand the achievable results without exercising the full optimization activities.

Andrea Bondavalli is a Full Professor of Computer Science at the University of Firenze, previously he was a researcher of CNR in Pisa. His research activity is focused on Dependability and Resilience of critical systems and infrastructures. In particular he has been working on designing resiliency, safety, security, and on evaluating attributes such as reliability, availability and performability. His scientific activities have originated more than 250 papers appeared in international Journals and Conferences. He received a Doctor Honoris Causa award from the Budapest University of Technology and Economics – in 2019. Andrea Bondavalli supports as an expert the European Commission in the selection and evaluation of project proposals and founded a spinoff – Resiltech – which employs currently 45 people. He led various national and European projects and coordinated a few. He participates to (and has been chairing) the program committee in several International Conferences including DSN, SRDS, SAFECOMP EDCC, LADC. Finally he is a member of the IFIP W.G. 10.4 Working Group on "Dependable Computing and Fault-Tolerance" since 2000 having organized a few workshops and hosted a few meeting.