TECoSA – Trends, Drivers, and Strategic Directions for Trustworthy Edge Computing in Industrial Applications

James Gross, jamesgr@kth.se; **Martin Törngren,** martint@kth.se; **György Dán,** gyuri@kth.se; **David Broman,** dbro@kth.se; **Erik Herzog,** erik.herzog@saabgroup.com; **Iolanda Leite,** iolanda@kth.se; **Raksha Ramakrishna,** rakshar@kth.se; **Rebecca Stower,** stower@kth.se; and **Haydn Thompson,** haydn.thompson@thhink.com
Copyright ©2022 by James Gross, Martin Törngren, György Dán, David Broman, Erik Herzog, Iolanda Leite, Raksha Ramakrishna, Rebecca Stower, and Haydn Thompson. Published by INCOSE with permission.

ABSTRACT
TECoSA – a university-based research center in collaboration with industry – was established early in 2020, focusing on Trustworthy Edge Computing Systems and Applications. This article summarizes and assesses the current trends and drivers regarding edge computing. In our analysis, edge computing provided by mobile network operators will be the initial dominating form of this new computing paradigm for the coming decade. These insights form the basis for the research agenda of the TECoSA center, highlighting more advanced use cases, including AR/VR/Cognitive Assistance, cyber-physical systems, and distributed machine learning. The article further elaborates on the identified strategic directions given these trends, emphasizing testbeds and collaborative multidisciplinary research.
KEYWORDS: edge computing, cyber-physical systems, trustworthiness, systems engineering, innovation eco-systems

INTRODUCTION

Several trends and drivers interact in the digitalization shift, including edge computing, connectivity, artificial intelligence, and big data loops, where field data are gathered to update software systems continuously. This transformation offers unprecedented innovation and product development opportunities and enables industrial companies to meet their targets for sustainable development goals. The need to address all dimensions of sustainability is highlighted by the recent European Commission initiative on Industry 5.0, emphasizing that previous efforts, such as Industry 4.0, have predominantly focused on productivity (EC Industry 5.0, 2022). A concrete example of what CPS can do for sustainability is the "tools" available to facilitate circularity, as an example, with traceability and predictive capabilities to support decisions regarding maintenance and recycling. However, digital transformation also increases system complexity. It introduces challenges of a socio-technical nature, such as risks related to technical systems acting in open environments, including ethical considerations related to fairness and personal integrity, INCOSE (2021), Törngren (2021). Specifically, our future societies will depend on increasingly sophisticated infrastructures where **edge computing** will act as a new tier, complementing the cloud and embedded systems. TECoSA, a research center on trustworthy edge computing systems and applications, was formed in 2020 to address the corresponding key challenges (TECoSA 2022; Törngren et al. 2021). The center brings together multiple research teams at KTH Royal Institute of Technology and (currently) 15 industrial partners spanning several industrial domains. The discussions among the center partners form the basis for the results presented in this paper.

TECoSA is active in industrial digitalization with a focus on edge computing systems. The aim is to provide methods, tools, and theories for building trustworthy systems relying on edge computing. The emphasis during the initial phase of the center—in the context of trustworthiness—has been on safety, cyber-security, and predictability (see Figure 1). Trustworthiness has traditionally been associated with human-machine interactions and security, referring to how we (humans) perceive trust in services and machines.

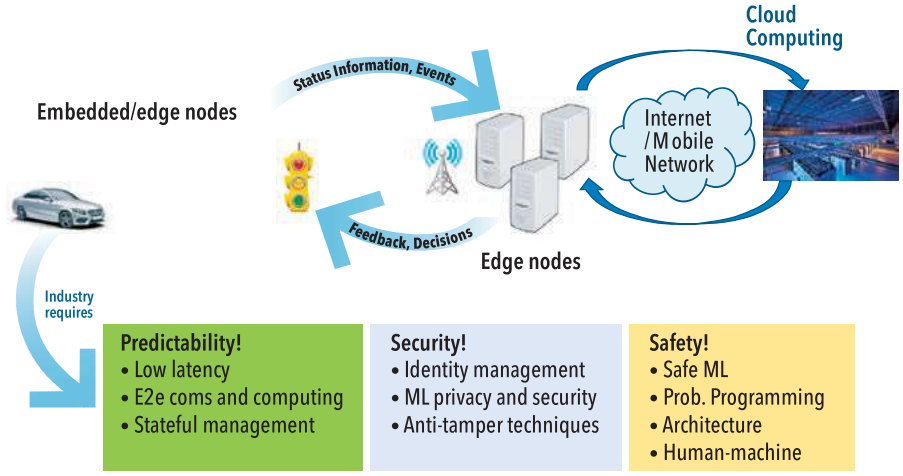

Figure 1. Edge computing as a new tier complementing embedded systems (device edge) and the cloud, illustrating initial trustworthiness properties and challenges addressed by the TECoSA center

Trustworthiness has evolved as an umbrella term encompassing the concept of dependability, associated with properties like reliability, availability, maintainability, safety, and security, and properties associated with artificial intelligence, such as transparency, explainability, and fairness (AI HLEG 2021).

The primary purpose of this article is to initially summarize and assess the current trends and drivers regarding edge computing. These insights form the basis for the research agenda of the TECoSA center. As a second purpose, the article elaborates and discusses the identified strategic directions given these trends.

EDGE COMPUTING STATE-OF-THE-ART

Edge computing is best understood in contrast to cloud computing. In the 2000s, the client-server approach dominated the first-generation internet architecture. Most clients were desktop PCs on private or corporate premises, connected via the Internet to web servers. A private or corporate entity intending to offer information or services on the Internet had to acquire server hardware and software, install and maintain it on corresponding premises, and set up a matching Internet connection. By 2010, this division had changed dramatically.

On the one hand, an increasing fraction of the clients were mobile devices, connecting through mobile networks like 3.5G and the upcoming 4G (LTE) to the Internet. On the other hand, web service offerings moved more and more to cloud providers, where very large pools of server hardware were brought together, allowing a scalable and efficient operation of web services from an installation, maintenance, and connectivity point of view. Web service operators moved from hosting and maintaining servers (with the content) locally on-premise to only curating content while renting the hardware and software for the web service from cloud providers. As a result, cloud computing centers of corresponding providers often ended up in locations where physical space, energy supply, and backbone connectivity were cheap, resulting in relatively remote locations. By and large, this is the dominating service model of the Internet as of today.

In this context, edge computing is primarily defined as computing services in "closer physical proximity" to clients compared to cloud computing, that is, offering computing services towards the "edge" of the Internet / wide-area networks. Given the dominant presence of 4G and 5G mobile networks as primary access networks of most clients in today's Internet, edge computing is realized by placing corresponding compute resources within the mobile network core or even within a radio access network, depending on the preferred proximity. In this line of thinking, proximity is traded with scale and cost: The higher the desired proximity of edge compute resources to the mobile clients, the more physical locations for placements of edge compute resources will be required, typically leading to fewer computational resources available per edge compute location.

Visions associated with edge computing have in various academic/industrial communities been given different names, including, for instance, multi-access edge computing (MEC) (related to telecommunications and 5G, earlier referred to as mobile edge computing), (Abbas et al. 2018), fog computing (with localized computations through communication devices such as routers and gateways in collaboration with the cloud), (Bonomi et al. 2012), and cloudlets (small scale localized data centers), (Satyanarayanan 2017). In the current discourse, edge computing has been associated with either locality, computing technologies, or both (Varghese et al. 2021). While the many projections for edge computing may appear confusing, this situation is not surprising since we are in the early stages of edge computing with an ongoing market positioning.

Our analysis from a commercial point-of-view is that edge computing provided by mobile network operators will be the initial dominating form of this new computing paradigm for the upcoming decade. Beyond that, new concepts might arise that exploit a continuum of available compute points from mobile clients to cloud centers (Duranton et al. 2021). *In the following, we refer to edge computing as the provisioning of additional computing resources through mobile networks.* Edge computing could be introduced to decrease hardware costs in mobile devices, such as industry robots and civilian or military surveillance systems while meeting latency and predictability demands. In this sense, edge computing adds computational resources that complement the existing capabilities of devices (embedded systems) and the cloud, belonging to a tier of a digitalized infrastructure. For a presentation of more detailed use cases, see the following discussion below.

Edge computing was arguably introduced roughly twenty years ago under the synonym "cyber foraging" (Balan et al. 2002). Since then, a set of various arguments have been brought up highlighting the potential benefits of edge computing:

- The original cyber-foraging research was motivated to **improve energy efficiency** if compute-intensive jobs could be offloaded from battery-powered mobile clients to stationary but close-by cloudlets, decreasing network-wide energy consumption. Either code or input data is offloaded through a mobile network to cloudlets, sending the computation back to the client. Cloud computing is seen in this context as having too long latency and unreliable, necessitating edge computing.
- A second argument, related to the above, can be made about the relative distance of cloud computing centers and, therefore, a much **lower access delay** in the case of edge computing. For compute tasks that are either too complex for mobile clients or require input from multiple mobile clients while being latency-sensitive, edge computing provides a clear advantage in providing lower round-trip delays. This case is, for instance, often made in the context of augmented or extended reality applications.

- Edge computing can also drastically **reduce the bandwidth** required for certain analysis services that run in the cloud. In this case, cloudlets are used as primary processing units, for instance, with respect to video analytics in detecting certain events or states in the video stream. Instead of conveying the entire stream to a cloud center, leading to a large bandwidth requirement as more and more endpoints are included in the service, only indices of the video frames and the detected objects are provided upstream to the cloud center. The corresponding video frames are nevertheless stored at the edge and can be retrieved by the cloud center. Similar cases can be made for predictive maintenance, IoT systems, and distributed machine learning applications.
- Finally, edge computing systems come with **different security and privacy features**. While typical concerns of security and privacy regarding cloud computing centers do not carry over to edge computing, new aspects such as physical access and manipulation become more relevant in the case of edge computing. Related to this shift towards more "local" aspects of security and privacy are also advantageous of edge computing with respect to **regulatory frameworks**. Due to the geographical proximity of deployed cloudlets and corresponding clients, edge computing offerings might guarantee the manipulation and storage of data within a specific regulatory framework, which a general-purpose cloud provider might not be able or willing to guarantee (in contrast to sovereign cloud offerings).

From these diverse drivers and advantages discussed in the academic/industrial community over the last ten years, for the first wave of commercial edge computing offerings foreseeable today, the regulatory and bandwidth-saving aspects are likely the main drivers. Concerning B2B customers, edge computing offerings of mobile network providers, referred to as Telco edge, as well as cloud providers, referred to as the regional cloud, will offer guarantees for the computing and storage location and, therefore, the regulatory conditions under which data is manipulated and stored. In addition, hybrid edge-cloud solutions are emerging that push the bulk of the processing to edge cloudlets while integrating the results of local cloudlet-based computing with cloud services. In both cases, "best effort" service level agreements (SLAs) between the service provider and customer are sufficient for successful commercialization. Beyond these B2B offerings, in the B2C space, a prominent commercialization case for edge computing appears to be online multi-player gaming, where depending on the location of the players and the placement of the game backend process, significantly higher quality of experience can be achieved. Still, corresponding offerings in the gaming domain will be run under best-effort SLAs.

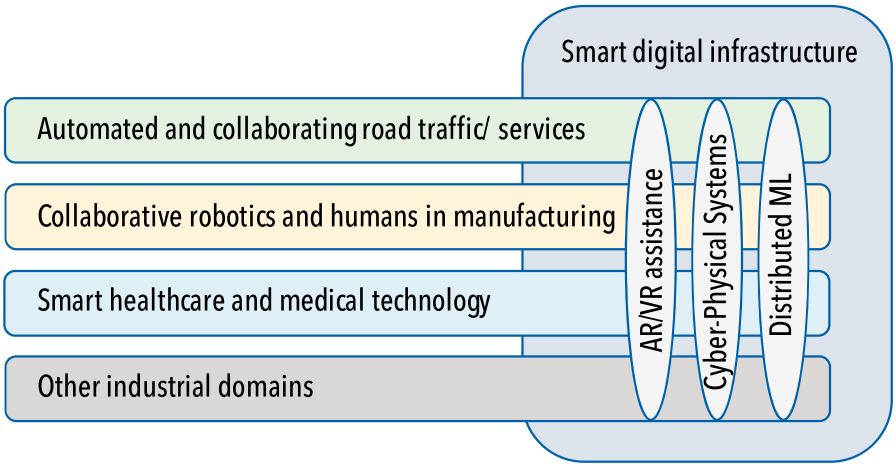

Figure 2. Various application domains, use-cases of cross-domain relevance, and interactions with a digital infrastructure (providing edge computing, communication, and other capabilities such as positioning)

BEYOND "BEST EFFORT"

More advanced use cases exist that could benefit from edge computing but where different challenges exist today, including both technical/scientific as well as related to business models. The commercial viability of these opportunities thus remains uncertain, and the TECoSA center has identified three types of use cases as particularly interesting where more research is needed. These use cases all demand more localized computing power, providing incentives for edge computing. The use cases are also relevant in several application domains, driving setups in which a digitalized edge computing-based infrastructure promises added value (see Figure 2). In manufacturing, for example, many ongoing field tests involve using private 5G networks and edge computing, representing such digitalized infrastructures. We first elaborate on these use cases and then discuss approaches to address them.

- **Use case 1: Mobile AR/VR/Cognitive Assistance**: The first use case concerns the advantages of future edge computing deployments in human-in-the-loop applications like virtual reality (VR) and augmented reality (AR). These are closed-loop systems where different "status" information is conveyed upstream to the point of computation (that is, the cloudlet). The provided status information is used for generating feedback at the backend, which is then transmitted back to the application client. AR and VR applications are generally characterized by (1) high data rate requirements upstream and/or downstream, (2) complex backend processing taking place at the cloudlet, and (3) quality-of-experience (QoE) of the application is directly related to the responsiveness of the entire loop (upstream communication, compute, and downstream communication). Subtle differences exist concerning the workloads and QoE requirements for AR systems versus VR systems, where VR systems require higher bandwidths in the downlink. Generally speaking, the latency requirements are also higher due to the level of immersion. The specific challenges for both application types relate to the following:
  1. **Efficient application support:** Due to the interplay between communication and compute elements over the offloading loop, many trade-offs exist to manage end-to-end delays at runtime dynamically. These trade-offs are largely unexplored, particularly about quality-of-experience implications in the short- and long-term. Managing end-to-end delays with respect to QoE over a heterogeneous set of active AR/VR applications is a further challenge, as is the question of optimal placement of the compute backend or efficient and reliable mobility support for such applications. To a large extent, the efficient support of such applications also hinges on the degree of control the application will be able to execute over the mobile network. In the past, mobile network systems have offered only very limited APIs (application

programming interfaces) as QoE requirements for voice, video, or web applications have been similar and hence easy to manage. However, for AR or VR applications, more complex trade-offs are likely to be only known to the application at runtime. Hence, a more powerful API for resource control enables a significantly more efficient operation.

2. **Scalable life-cycle support of applications and end system acceptance:** While several SDKs exist for AR and VR systems, devising a new application over programming, deployment, and updates is highly complex and requires deep software engineering and platform knowledge. This contrasts with the corresponding life-cycle support of smartphone apps of various ecosystems currently in the market. From the perspective of the supply side of future AR/VR applications, a significant simplification of the life-cycle support is likely to be established over the following years. Due to the above limitations, AR technology commercialization has been limited. Advanced designs, combined with a changing sentiment in the group of early adopters, might lead over the following years to a breakthrough in these applications. A scalable provisioning of backend compute capabilities via edge computing paired with near-ubiquitous mobile network access will undoubtedly lift the technological bottlenecks for widespread adoption.

▪ **Use case 2: Cyber-physical systems (CPSs):** CPSs represent the "integration of computation, networking, and physical processes." While CPSs have been around since the 1970s with the integration of microprocessors with physical systems, these systems now see unprecedented potential in their capabilities (Thompson and Reimann 2018). Representative examples include automated vehicles and future manufacturing systems. In such CPSs, additional sensors, communications, and collaboration can enhance context awareness and planning. The role of edge computing comes into play to provide the needed computational and analytics support, providing the potential for handling large amounts of data for real-time applications and supporting CPS collaboration. TECoSA has identified many applications in domains such as those depicted in Fig. 2, supporting enhanced quality and new functionalities, for example, by ensuring that the right assembly tools are used for the right parts in a manufacturing process. For CPSs, we identify the following challenges:

1. **Holistic management of computing and communication resources:** Industrial applications have demanding requirements on real-time (predictable and short enough) latencies, availability, and error detection and handling. This requires novel end-to-end resource management capabilities, including exploiting an interplay between applications and infrastructure and considering energy consumption as a key metric. With such considerations, edge computing promises to minimize/reduce the overall energy consumption of applications.
2. **Trustworthy applications based on edge computing:** As already introduced, trustworthiness has evolved to become an umbrella term. Given the evolution of CPS, most of the trustworthiness properties will be relevant for future CPS. Incorporating edge computing into future CPS poses new challenges, given new failure modes and cyber-security risks (vulnerabilities) of edge computing-based infrastructures and applications. The dependencies and trade-offs between trustworthiness properties require specific attention, especially for open and collaborative CPS with potential conflicts between cyber-security, safety, availability, and data sharing. The uncertainty involved in such open further CPS requires run-time risk assessment and handling/adaptation to balance safety and availability/performance appropriately. Certification and re-certification of (evolving and adapting) edge-based CPS also represent an open challenge.
3. **Collaborating systems and scalability**: Collaborating systems, often referred to as systems of systems (SoS), lack a central authority responsible for systems integration and where the constituent units evolve independently (for example, in the domain of roads, actors such as vehicles and the physical and digital infrastructures of the roads) (Maier 1998). This leads to challenges regarding the overall design and responsibilities of such SoS and strongly relates to the business model(s) and liability if something goes wrong. The "intelligent transport systems" example has shown the difficulty of establishing such SoS. We believe that the introduction of 5G and beyond as a digital infrastructure, with its provision for low latency and quality of service, may help to create the momentum needed to establish the required models for collaboration.

▪ **Use Case 3: Distributed ML:** Machine learning (ML) is widely considered an efficient tool for optimization, prediction, and classification tasks found in various industrial and consumer applications, among others in AR/VR systems (UC1) and CPSs (UC2). The use of ML in these systems could be limited to applying the pre-trained model for performing a certain task on data received from end devices, referred to as inference. More generally, it can entail periodic training of the model to adapt it to changing environmental conditions. The use of ML for inference usually involves upstream traffic and may involve downstream traffic also if/when the inference leads to decisions that, in turn, affect devices. Training of a model may also involve downstream traffic if the updated model is to be distributed to end-user equipment. ML algorithms are often represented as execution graphs and can be deployed on various devices spanning the edge-to-cloud continuum. Such distribution of ML primitives enables capabilities previously unattainable in energy and computationally-constrained environments. For example, by placing parts of the execution graph having challenging real-time requirements and low computation complexity on end-user equipment and computationally intensive parts in the edge cloud, one can obtain low-latency ML algorithms with limited computational resources. At the same time, distributed ML comes with a variety of challenges, in particular:

1. **Interoperability:** Interfaces for interconnection are needed to enable interoperability between components from different vendors and to make system integration more cost-efficient. Since ML algorithm development is in its early stages, it is challenging to establish interfaces that will last years or decades.
2. **Systems architecting:** Systems architecting aspects and algorithmic issues will become key in ever more complex installations. It needs to be clarified how to formulate architectural and design principles for complex, ML-enabled systems to ensure functional and non-functional requirements and simultaneously allow for efficient life-cycle management. Sustainability in terms of energy consumption and the environmental footprint of the computing and communications infrastructure needed for ML integration is a closely related issue.

3. **Robustness and cybersecurity.** Robustness to adversarial environments and the lack of privacy guarantees could also hinder the wide-scale adoption of ML-enabled systems. ML algorithms are vulnerable to adversarial inputs, for example minor perturbation of the data, unnoticeable manipulations of algorithm parameters, and trained ML models may also reveal confidential information about the data set used for creating them (Ramakrishna 2022). These issues related to trustworthiness remain to be solved.

The TECoSA center approach to address these challenges: Successful research centers have been reported to exhibit characteristics including collaborative multidisciplinary research involving multiple domains, use of testbeds/demonstrators, and having a strong connection to education (Patterson 2014). We agree that these characteristics are important. TECoSA has emphasized creating a knowledge ecosystem with the involved stakeholders and aims to develop testbeds as experimental and open infrastructures in automated and connected road traffic and collaborative robotics in the coming period. These testbeds will be used to support collaborative research and education. A critical aspect of the testbeds is stimulating the interplay between applications – potentially involving all the mentioned use cases – and digital infrastructures (Figure 2). This interplay corresponds to interactions between different research teams and organizations/companies, places the focus on platforms and services (the interfaces between applications and the infrastructures), and has the potential to be used in education and stimulate open debate on the socio-technical implications.

CONCLUSIONS AND FUTURE WORK (CONVERGENCE OF THE USE CASES AND MORE)

We have discussed trends and drivers related to edge computing as a new computing tier overcoming limitations of and complementing the cloud and resource-constrained embedded systems. The multitude of concepts such as MEC, cloudlets, fog computing, "near/far/nano/enterprise edge," and "distributed cloud" (Heinen 2021), while partly confusing, is natural considering that we are in the early stages of edge computing with an ongoing market positioning. In our analysis, edge computing provided by mobile network operators will be the initial dominating form of this new computing paradigm for the coming decade. In this sense, edge computing adds computational resources that complement the existing capabilities of devices (embedded systems) and the cloud, belonging to a new tier of a digitalized infrastructure. Regulatory aspects, bandwidth saving, and soft real-time interactions such as gaming will likely drive the first wave of commercial edge computing offerings. We highlight that in these cases, "best effort" SLAs between the service provider and customer are sufficient for successful commercialization.

We have also discussed more advanced use cases, including AR/VR/Cognitive Assistance, CPSs, and distributed ML, and corresponding challenges that require further research. The three presented use cases will, in many ways, be part of the same system, for example, with humans in the loop (such as "cobots"—humans and robots collaborating) in the context of CPS and with data gathering and distributed machine learning taking place in parallel with the other use cases.

In addressing these use cases and challenges, the identified key role of a university-led research center is to maintain and grow a knowledge ecosystem to support innovation, research, and education in trustworthy edge-based CPS and to develop corresponding technological foundations and methodologies. ■

REFERENCES

- Abbas N., Y. Zhang, A. Taherkordi, and T. Skeie. 2018. *Mobile Edge Computing: A Survey. IEEE Internet of Things Journal* 5 (1): 450–465.
- AI HLEG 2021. *High-Level Expert Group on AI of European Commission. Overview of deliverables from the AI HLEG.* Web page reference: https://digital-strategy.ec.europa.eu/en/policies/expert-group-ai (accessed 2022-08-17).
- Balan R., J. Flinn, M. Satyanarayanan, S. Sinnamohideen, and H. Yang. 2002. *The case for cyber foraging. Proc. 10th workshop on ACM SIGOPS European workshop (EW 10).* Association for Computing Machinery, New York, US-NY: 87–92. https://doi.org/10.1145/1133373.1133390
- Bonomi F., R. Milito, J. Zhu, and S. Addepalli. 2012. *Fog Computing and Its Role in the Internet of Things. Proceedings 1st Edition MCC Workshop on Mobile Cloud Computing* (Helsinki, Finland) (MCC '12). ACM, New York, US-NY: 13–16. https://doi.org/10.1145/2342509.2342513
- Duranton M., M. Malms, and M. Ostasz. 2021. *The continuum of computing.* Hipeac Vision 2021. https://doi.org/10.5281/zenodo.4719341
- EC Industry 5.0. 2022. Web page reference: https://research-and-innovation.ec.europa.eu/research-area/industry/industry-50_en (accessed 2022-08-17).
- INCOSE. 2021. *Systems Engineering Vision 2035.* https://www.incose.org/about-systems-engineering/se-vision-2035
- Heijnen A. et al. 2021. *IoT and Edge Computing: opportunities for Europe.* Report by the NGIoT project (Next Generation Internet of Things). Retrieved from https://www.ngiot.eu/
- Maier M. 1998. Architecting principles for systems-of-systems. *Systems Engineering Journal.* 1 (4): 267–284, 1998.
- Patterson D. 2014. *How to build a bad research center.* Commun. ACM 57(3): 33–36. https://doi.org/10.1145/2566969
- Satyanarayanan M. 2017. The Emergence of Edge Computing. *IEEE Computer* 50 (1).
- Törngren M. 2021. *Cyber-physical systems have far-reaching implications.* Hipeac Vision 2021. https://doi.org/10.5281/zenodo.4710500
- TECoSA, 2022. Web page reference: https://www.tecosa.center.kth.se/ (accessed 2022-08-17).
- Ramakrishna R. and G. Dán. 2022. *Inferring Class-Label Distribution in Federated Learning.* ACM Workshop on Artificial Intelligence and Security (AISec).
- Ramli R. and M. Törngren. 2022. *Towards an Architectural Framework and Method for Realizing Trustworthy Complex Cyber-Physical Systems.* Joint Proceedings of RCIS 2022 Workshops and Research Projects Track, Barcelona, ES: May 17-20, 2022.
- Thompson H. and M. Reimann. 2018. *Platforms4CPS Key Outcomes and Recommendations.* https://www. platforms4cps.eu
- Törngren M., H. Thompson, E. Herzog, R. Inam, J. Gross, and G. Dán. 2021. *Industrial Edge-based Cyber-Physical Systems – application needs and concerns for realization.* Proc. of ACM Symp. on Edge Computing Workshop on Trustworthy Edge Computing.
- Varghese B. et al. 2021. *Revisiting the Arguments for Edge Computing Research.* IEEE Internet Computing, doi: 10.1109/MIC.2021.3093924.

ABOUT THE AUTHORS

James Gross received his PhD degree from TU Berlin in 2006. Since November 2012, he has been with the Electrical Engineering and Computer Science School, KTH Royal Institute of Technology, Stockholm, where he is professor for machine-to-machine communications. At KTH, James is currently associate director of the Digital Futures Research Center, as well as co-director of the VINNOVA competence center on Trustworthy Edge Computing Systems and Applications (TECoSA). His research interests are in mobile systems and networks. He has authored over 150 (peer-reviewed) papers in international journals and conferences.

Martin Törngren is a professor in embedded control systems at the Mechatronics division at KTH since 2002, and with PhD in machine design/mechatronics also from KTH. Prior to becoming a professor at KTH he co-founded the company Fengco Real-time Control AB, specializing in advanced tools for developers of embedded control systems and related consultancy, and also did a postdoc period at the EU-JRC, Institute for Systems, Informatics and Safety, Ispra, Italy. He has particular interest in trustworthiness and dependability of cyber-physical systems and their design methodologies. Networking and multidisciplinary research have been characteristic throughout his career. He is the principal initiator of the Innovative Centre for Embedded Systems (www.ices.kth.se), launched in 2008 (and served as its director until 2020) – with close interactions with the Swedish chapter of INCOSE. He is the director of the TECoSA Swedish national competence center on Trustworthy Edge Computing Systems and Applications (initiated in 2020). In 2011/2012 he was visiting scholar at UC Berkeley (2011/12) and in 2018 at Stevens Institute of Technology (Hoboken, New Jersey, 2 months).

György Dán (M'07, SM'17) is professor of teletraffic systems at KTH Royal Institute of Technology, Stockholm, Sweden. He received the MSc in computer engineering from the Budapest University of Technology and Economics, Hungary in 1999, the MSc in business administration from the Corvinus University of Budapest, Hungary in 2003, and the PhD in telecommunications from KTH in 2006. He worked as a consultant in the field of access networks, streaming media, and videoconferencing 1999-2001. He was a visiting researcher at the Swedish Institute of Computer Science in 2008, a Fulbright research scholar at University of Illinois at Urbana-Champaign in 2012-2013, and an invited professor at EPFL in 2014-2015. He served as area editor of Computer Communications 2014-2021, and has been editor of IEEE Transactions on Mobile Computing since 2019. His research interests include the design and analysis of content management and computing systems, game theoretical models of networked systems, and cyber-physical system security and resilience.

David Broman is a professor at the Department of Computer Science, KTH Royal Institute of Technology and an associate director faculty for digital futures. He received his PhD in computer science in 2010 from Linköping University, Sweden. Between 2012 and 2014, he was a visiting scholar at the University of California, Berkeley, where he also was employed as a part-time researcher until 2016. His research focuses on the intersection of (i) programming languages and compilers, (ii) real-time and cyber-physical systems, and (iii) probabilistic machine learning. He has worked several years within the software industry, co-founded the EOOLT workshop series, and is a member of IFIP WG 2.4, Modelica Association, a senior member of IEEE, and a board member of Forskning och Framsteg.

Iolanda Leite is an associate professor at the School of Electrical Engineering and Computer Science at KTH Royal Institute of Technology. She holds a PhD in information systems and computer engineering from IST, University of Lisbon. Prior to joining KTH, she had postdoctoral appointments at Yale University and Disney Research. Her goal is to develop social robots that can perceive, learn from, and respond appropriately to people in real-world situations, allowing for truly efficient and engaging long-term interactions with people.

Raksha Ramakrishna received the BE degree in electronics and communications engineering from the Rashtreeya Vidyalaya College of Engineering, Bangalore, India, in 2014, and the MS and PhD degrees in electrical engineering from Arizona State University, in 2017 and 2020, respectively. She is currently a postdoctoral researcher with the Division of Network and System Technology, KTH Royal Institute of Technology, Stockholm, Sweden. Her research interests include the domains of statistical signal processing, data analytics for power systems, and security and privacy in federated machine learning systems.

Rebecca Stower is a postdoctoral researcher at KTH Royal Institute of Technology. She holds a PhD in psychology from Jacobs University, Bremen, Germany and a BSc in psychology from the University of Queensland in Australia. She is passionate about the intersection of psychology and technology and how psychological research methods can be applied to digital industries. Dr. Stower is working on swarm robotics and is interested in the conceptualization of social intelligence in robots and the design of social robot behavior.

Professor Haydn Thompson, BSc, PhD CEng has over 35 years experience working in a mixture of senior industrial research and development roles in flight control systems, space programmes and radar signal processing applications for leading companies. From 1993- Feb. 2013 he was the program manager of the Rolls-Royce Control and Systems University Technology Centre. He is managing director and founder of the THHINK Group of companies. He is recognised and used by the European Commission as an expert in many fields, CPS, IIoT, AI, aerospace, automotive, autonomous vehicles, smart agriculture, advanced electronics, and more. Dr. Thompson is a consultant to a range of companies and government bodies. He defines strategic technology roadmaps across Europe and for companies such as Rolls-Royce. He has over 100 publications, has written two books and contributed to several others.