



LUND UNIVERSITY

The 6G Computing Continuum (6GCC): Meeting the 6G computing challenges

Tärneberg, William; Fitzgerald, Emma; Bhuyan, Monowar; Townend, Paul ; Årzén, Karl-Erik; Östberg, Per-Olov; Elmroth, Erik; Eker, Johan; Tufvesson, Fredrik; Kihl, Maria

Published in:

1st International Conference on 6G Networking (6GNet)

DOI:

[10.1109/6GNet54646.2022.9830459](https://doi.org/10.1109/6GNet54646.2022.9830459)

2022

Document Version:

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Tärneberg, W., Fitzgerald, E., Bhuyan, M., Townend, P., Årzén, K.-E., Östberg, P.-O., Elmroth, E., Eker, J., Tufvesson, F., & Kihl, M. (2022). The 6G Computing Continuum (6GCC): Meeting the 6G computing challenges. In *1st International Conference on 6G Networking (6GNet)* IEEE - Institute of Electrical and Electronics Engineers Inc.. <https://doi.org/10.1109/6GNet54646.2022.9830459>

Total number of authors:

10

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

The 6G Computing Continuum (6GCC): Meeting the 6G computing challenges

William Tärneberg¹, Emma Fitzgerald¹, Monowar Bhuyan³, Paul Townend³, Karl-Erik Årzén²,
Per-Olov Östberg³, Erik Elmroth³, Johan Eker², Fredrik Tufvesson¹, and Maria Kihl¹

¹Department of Electrical and Information Technology, Lund University, Sweden

²Department of Automatic Control, Lund University, Sweden

³Department of Computing Science, Umeå University, Sweden

Abstract—6G systems, such as Large Intelligent Surfaces, will require distributed, complex, and coordinated decisions throughout a very heterogeneous and cell free infrastructure. This will require a fundamentally redesigned software infrastructure accompanied by massively distributed and heterogeneous computing resources, vastly different from current wireless networks. To address these challenges, in this paper, we propose and motivate the concept of a 6G Computing Continuum (6GCC) and two research testbeds, to advance the rate and quality of research. 6G Computing Continuum is an end-to-end compute and software platform for realizing large intelligent surfaces and its tenant users and applications. One for addressing the challenges or orchestrating shared computational resources in the wireless domain, implemented on a Large Intelligent Surfaces testbed. Another simulation-based testbed is intended to address scalability and global-scale orchestration challenges.

Index Terms—6G, Large Intelligent Surfaces, Computing Continuum, Distributed Orchestration, Computing at Scale

I. INTRODUCTION

The evolution of 6G requires a fundamentally redesigned software infrastructure [1]. 6G will require massive amounts of computing resources to make distributed, complex, and coordinated decisions throughout the whole infrastructure, from hyper-local computing in proposed systems such as Large Intelligent Surfaces (LIS) [2], to global computing, as illustrated in Figure 2. The infrastructure is required to provide the foundation for hosting both the massive computing capacity required for the wireless network infrastructure itself, as well as an elastic computing platform for hosting demanding user-centric 3rd-party applications. The resulting infrastructure will be highly heterogeneous and distributed, and the paradigm shift to cell-free Radio Access Network (RAN) technologies as LIS will pose novel orchestration and infrastructure challenges.

In response, we propose 6G Computing Continuum (6GCC), a system that addresses several crucial end-to-end compute and software challenges for realizing 6G. We specifically target computing infrastructures and network functions for LIS [2]. The LIS concept was first introduced by researchers from Lund University in 2017 [2]. Presently, LIS is a candidate 6G RAN technology, on the merits that LIS offers lower latency and vastly increased throughput, improved reliability, and spectral efficiency over existing 5G RAN technologies. The core concept of LIS is increased spatial

radio diversity, realized by a very large number of tightly synchronized antennas, forming a cell-free RAN. By which set of antennas a User Equipment (UE) is served is orchestrated by the LIS infrastructure, in a distributed manner.

6GCC calls for both compute and networking research for LIS. It requires research on global-scale orchestration challenges for resilient massive scale Software Defined Networks (SDNs) data-driven operations, anomaly detection, and energy-efficient monitoring, as well as research on application software design for 6G infrastructures. To conduct meaningful research on 6GCC, we also propose and detail a LIS-based test-bed and a SDN Extended Network Simulator. The LIS test-bed will consist of a real LIS deployment with applications serving connected users. Further, the proposed SDN Extended network Simulator will be used to demonstrate results on the global-scale challenges of energy-efficient monitoring and orchestration of resources and applications.

II. LARGE INTELLIGENT SURFACES

The work on 6G has just begun, but we already know that it will include extreme requirements on data rates, latency, dependability, and power consumption [3]. LIS is a candidate technology for 6G systems that has the potential to meet these requirements. In LIS, many antennas are densely distributed around an area, and managed in a distributed, cell-free manner. Ongoing LIS research [4] is aimed at geographically constrained spaces, e.g., sports arenas or factories. However, the concept itself allows for much larger networks where users can seamlessly move between antenna panels within, e.g., a city, see Figure 1.

LIS deployments will support ultra-reliable radio access with imperceptible user-perceived delays [5] to many users. This is due to the dense and distributed antenna deployment. With a very high degree of spatial diversity, favorable propagation can be achieved with high probability to users in any location in the deployment area, yielding a consistently good channel quality and low bit error rates. The substantial number of antennas, potentially an order of magnitude more than current massive MIMO deployments, gives a corresponding increase in capacity if users can be effectively spatially multiplexed by controlling the antennas coherently. For many emerging use cases for 6G, the number and density of users




Fig. 1: An illustrative overview of Large Intelligent Surfaces (LIS). A fictive production space with two robotic arms, connected with LIS, realized with antenna elements on the perimeter of the space. Here interconnected, base-band units server multiple antennas.

will potentially be very high, allowing the system to take advantage of its full capacity. Examples of such use cases are tracking of individual goods in warehouses, or augmented reality for crowds at sporting events. The proximity and density of antennas, along with the abundance of available radio resources, also leads to reduced latency for users.

Further, LIS deployments can be used not only for communications but will also enable precise positioning services [5] and wireless energy transfer [6], enabling energy neutral Internet of Things (IoT) devices. Many antennas, distributed over a wide area but controlled together, give a large aperture for positioning that can achieve far greater accuracy than existing antenna array-based positioning systems. For wireless power transfer, the closer the antennas are to the users, the better, since even in free space the transferred power drops off quickly with distance, and this is exacerbated by multipath effects. With LIS, a user is likely to be close to several antenna panels at any given time, even for mobile users, leading to highly efficient wireless power transfer.

In envisioned LIS systems, such as RadioWeaves [7], groups of distributed antennas are aggregated by *base-band* units. The base-band units carry out signal processing, such as beam-forming, and are often implemented in ASIC or Field-Programmable Gate Arrays (FPGAs). These units are interconnected to a point where the system can be managed and traffic routed, and from which the system can be connected to the wider Internet.

III. COMPUTING CHALLENGES IN 6G

LIS deployments will impose extreme requirements on the underlying compute infrastructure. The expected requirements on processing capacity for next generation RAN technologies, such as LIS, will greatly surpass what can be provided today. This calls for a much more efficient and tightly integrated compute infrastructure than what is serving 5G today. Adding

energy constraints and sustainability goals, makes it an even harder problem to solve. Therefore, 6GCC takes a holistic view of the compute infrastructure needed to support the realization of 6G, and also enables a new generation of services and applications with extreme timing and availability requirements.

A. Heterogeneity and orchestration

To realize LIS and its applications, a highly distributed compute infrastructure is required, which should handle all compute needs for both radio-side processing, resource orchestration, and for hosting user-facing applications. It will span all available compute resources in a LIS deployment, from hardware accelerators as FPGAs connected to each antenna panel, to several layers of CPUs, to resources beyond the deployment ranging from edge nodes to large public clouds, see Figure 2. The infrastructure must adapt its resources to its dynamic users, as well as providing early data aggregation that eases the communication burden on the rest of the core network [4]. Therefore, there must be a dynamic and efficient autonomous resource orchestration, that makes optimal decisions on where to allocate compute resources, on all levels in the infrastructure, and when to migrate application components somewhere else, due to changed user demands or infrastructure conditions. While there exist numerous research on edge computing and resource management for cloud infrastructures [8]–[10], there are no works on these topics that target LIS or other next-generation RAN technologies.

B. Scaling to billions

A 6G compute infrastructure will incorporate millions (*potentially billions*) of connected devices, and traditional monitoring and management approaches in such systems will incur enormous energy and performance overhead. Typically, monitoring systems in massive-scale distributed systems generate enormous amounts of data [11] requiring significant energy to process and transfer; as 5G and 6G systems evolve, this will become unsustainable. Therefore, it is essential that new techniques be developed to allow energy efficient collection and analysis of key metrics at a massive scale. Further, the complex architecture and the requirements of the applications require anomaly detection that can offer high precision and on-demand network management, service isolation, and multi-tenancy. Learning models, as [12], [13] can enable extremely efficient, rapid, and proactive decision-making for both anomaly detection and fault/attack diagnosis and mitigation, adapting to trade-offs between probing intervals and network performance. This could offer a significant step towards making large-scale networks with zero-outage. However, so far, there are no solutions in the literature specifically aimed at 6G compute infrastructures.

C. Availability and reliability

6G applications require extreme availability and reliability, despite running on nodes that are potentially unprotected from




Fig. 2: Basic overview of the 6G Computing Continuum (6GCC). An intent to control a robotic arm is submitted to the system. The continuum synthesizes an application (a controller), for appropriate hardware platforms and then dynamically placed in the system. The robotic arm is controller from the controller deployed in the system and the control signal is passed over LIS.

external impact, such as from weather or sabotage [14]. Consequently, to manage such scenarios, the computing infrastructure for 6G must provide extreme availability and reliability demands, e.g., real-time continuous backups within nodes, real-time process and data replication to neighboring nodes, etc. [15], [16]. However, the infrastructure will also provide an abundance of resources on the global scale, which enables opportunities for performance improvements, by allowing the system to trade performance with availability and reliability.

D. Quality-elastic applications

To meet stringent Quality of Service (QoS) requirements in a highly heterogeneous 6GCC, application software also needs to be re-imagined. Application cannot be tailor made by developers for each hardware platform and resource scale in the system. While taking into account the complexity of the system's orchestration objectives. Instead applications will arguably have to be synthesized from an intent, such as an objective to control a robotic are in certain manner given a set of constraints, see Figure 2. The system would synthesize the intent into a quality elastic implementation that can be deployed to and function where it deemed needed by the system orchestrator. The applications need to be both extremely resilient and quality-elastic [17], i.e., dynamically adapt to prevailing resource availability wherever they happen to reside in the infrastructure without failures. For applications with time-critical feedback control, missing computation deadlines can be catastrophic to the system under control [18]. Here, recent work on the weakly hard model has indicated that control tasks can also be inherently robust to deadline misses if designed correctly [19]. Also, quality-elasticity can be applied for both short and long time horizons. Quality elasticity for controllers for cloud environments has been studied in [20]. Because of a high degree of user mobility, the infrastructure's conditions (for example channel conditions in a LIS), applications may need to be, in a sense, spatio-elastic, and continuously evaluate

and migrate to a compute node where their outputs can be improved.

IV. PROPOSED 6GCC TEST BEDS

Because of the scale of the infrastructure and the nature of challenges we propose, a *6GCC test-bed* and a *simulator* for research on the topics presented in Section III. Below, we detail the two test-beds, which challenges they intend to investigate and how.

A. LIS Test-bed

The *LIS test-bed* will be based on and augment a LIS deployment, implementing lower stack functionality, currently under development. The LIS deployment will provide the hardware and radio architectures. To realize a 6GCC, we propose the addition of compute, networking resources as well as application architectures. The resulting 6GCC test-bed will demonstrate a full-stack end-to-end LIS deployment, with heterogeneous hardware meeting the diverse need of the system, unified with a computing continuum platform. Further, the test-bed will host a resilient and quality-elastic application in the form of a control system for a set of Unmanned Aerial Vehicles (UAVs), see Figure 3. The altitude control, localization, and object tracking are performed remotely in the infrastructure, where the physical compute location depends on the demand and resource availability. The test-bed will be built around and for the sake of conducting research on the below topics. The scale of the LIS test-bed is intended to cover the physical (cubic-meter) space in which the UAVs will operate.

1) *A LIS Compute Fabric:* The test-bed will be an end-to-end full-stack LIS deployment, building on the architecture and radio solutions from the LIS deployment. A MAC layer tailored for the targeted applications, will be implemented. Also, in cell-free RANs, user mobility will cause major challenges for the Networking layer. Therefore, we propose a Networking layer that enables the delivery of IP-traffic and




Fig. 3: The proposed LIS test-bed

IP routing between the distributed nodes in a LIS deployment. There needs to be a tight cross-layer coupling between the Physical, MAC, and Networking communication layers with domain specific knowledge of properties of the targeted applications. Further, virtualization technologies suitable for LIS will expose the infrastructure to user-facing applications in a cloud-like manner. More specifically, previously obscured FPGAs and Central Processing Units (CPUs) close to the antenna elements will be made available to applications through a simple virtualized platform. This will enable us to study how best to vitalize network functionality in such a system; although some work has investigated the application of virtualisation to FPGAs [21], optimal “cloudification” of LIS systems will require an even lighter, more responsive approach. We propose new, lightweight technologies, which can be applied to LIS and other cell-free RANs. This will involve deeper investigation of the efficacy and practicality of software container-based solutions, Function as a Service (FaaS), etc.

2) *6GCC Integration & Orchestration*: We propose new technology to enable dynamic and efficient autonomous orchestration, together with the resource management mechanisms necessary to integrate LIS with the wider system. This is essential, as user-facing applications, network functions, and the radio channels’ algorithms are all highly location-dependent, and it is essential that the required QoS is met for every UE. To realize this, we propose to investigate and design data-driven holistic compute orchestrators based on cross-layer data collection and machine learning. Additionally, we also propose to investigate online fair allocation methods [22], where not only resources but also applications and their preferences are online. The work package will further investigate the interface between LIS and the wider 6G continuum – developing methods for off-loading software tasks in an efficient manner, reporting QoS metrics off-site, etc.

3) *Application Software Design*: We propose to design and implement highly resilient software architectures and design principles for time-critical applications hosted in the system. Our main use case will be time-critical applications with feedback control over the cloud [23], which will be demonstrated

with a UAV scenario in the final LIS testbed. We plan to extend the methodology that we have proposed for control systems [24] to more generic time-critical applications. In particular, we plan to identify application-level and application-dependent requirements, which for a generic application are specific to the domain, but can be translated into a requirement for the infrastructure and its real-time characteristics. Also, we propose to address applications’ quality-elasticity [17]. Rather than failing when resources are scarce, a quality-elastic application adapts the quality of its output to the best of its ability given the resources it is allocated. We propose to build on [20] and investigate more dimensions of feedback control over the cloud. Further, we propose to investigate how applications in the clouds can be designed to be spatio-elastic, i.e. be able to evaluate its current situation and migrate to another resource to achieve better performance.

B. Simulator

Due to the complexity and scale of the global 6GCC, it is not feasible to build a physical test-bed for demonstration of results on a global scale. For this reason, we propose a *simulator* for SDN-based Extended Network systems. To capture the dynamics of massive-scale heterogeneous systems in shared resource environments, the simulator will include detailed models of the full-stack LIS deployment, resource orchestration strategies, resilient SDN systems, as well as provide interfaces for experimentation with the resilient and quality-elastic applications, and use the energy-efficient monitoring strategies. With the simulator, we can address scalability challenges of the end-to-end 6GCC and go beyond local scenarios to study the global scenario.

1) *Resilient SDNs at Massive Scale*: We aim to develop the technology necessary to create and manage highly efficient and fault/attack resilient SDNs that span the system. Specifically, using the data supplied by monitoring mechanisms, we propose a ML-based knowledge plane that maintains a local and global view of networks for inducing intelligence across the continuum by identifying and extracting relevant features and building appropriate learning models. This will include investigating contrastive reinforcement learning [12] for reinforcing monitoring of each component across layers (e.g., core networks, functions, slices), and self-supervised learning [13] for anomaly detection. Anomaly detection will be demonstrated on both the LIS testbed and the SDN simulator by demonstrating our system monitoring approach (and its performance overhead) contrasted with existing schemes.

2) *Energy Efficiency at Scale*: We aim to develop new technology to enable dynamic and highly energy efficient monitoring and management in all parts of the system. Typically, monitoring systems in massive-scale distributed systems generate enormous amounts of data [11] requiring significant energy to process and transfer; as 5G and 6G systems evolve, this will become unsustainable. This will help researchers identify the key monitoring metrics that are required at each layer of the system, and develop extremely lightweight and

distributed monitoring agents that will dynamically change recording and reporting frequency and granularity based on balancing local and global energy requirements with user service-level agreements. This extends the ideas of [11], [25], moving from generic Fog-based systems to dynamic cell-free RAN environments, and focusing on reduction in energy consumption. This will help us explore optimum deployment methods, including both containerised and function-as-a-service approaches. An autonomous management system will be developed that will compute and assess energy requirements at both local and global level with the aim of allowing for priority-based optimisation of monitoring at all layers of the conceptual stack.

V. CONCLUSIONS

Due to the non-fulfilment of future wireless network requirements with 5G, 6G has emerged to redesign the software infrastructure by inducing appealing features for offering services in massively heterogeneous wireless environments. In this paper, we present the foundational concepts on the 6G computing continuum with an architecture. Subsequently, how LIS can meet the 6G requirements is explained with real-time deployment benefits. The LIS test-bed, a real-time test-bed for illustrating and validating 6G realization and a simulation test-bed to mitigate the scalability issues in real-time environments. Finally, we explored and identified the opportunities and challenges in 6GCC for providing reliable, robust, scalable, and end-to-end compute and services to tenant's users and applications in massively heterogeneous wireless environments.

ACKNOWLEDGMENTS

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, the SEC4FACTORY project, funded by the Swedish Foundation for Strategic Research (SSF), and the IMMINENCE Celtic Next project funded by Sweden's Innovation Agency (VINNOVA). The work of Emma Fitzgerald was also supported by the National Science Centre, Poland, under the grant no. 2019/35/D/ST/7/00350 "Quality of Service in IoT Applications Using Large Scale Antenna Systems". The authors are part of the Excellence Center at Linköping-Lund on Information Technology (ELLIIT), and the Nordic University Hub on Industrial IoT (HI2OT) funded by NordForsk.

REFERENCES

- [1] P. Beckman, J. Dongarra, N. Ferrier, G. Fox, T. Moore, D. Reed, and M. Beck, "Harnessing the computing continuum for programming our world," *Fog Computing: Theory and Practice*, 2020.
- [2] S. Hu, F. Rusek, and O. Edfors, "The potential of using large antenna arrays on intelligent surfaces," in *Vehicular Technology Conf. (VTC Spring)*. IEEE, 2017.
- [3] Ericsson AB, "Joint communication and sensing in 6G networks," <https://www.ericsson.com/en/6g>, 2021.
- [4] L. Van der Perre, E. G. Larsson, F. Tufvesson, L. De Strycker, E. Björnson, and O. Edfors, "RadioWeaves for efficient connectivity: analysis and impact of constraints in actual deployments," in *Asilomar Conf. on Signals, Systems, and Computers*. IEEE, 2019.
- [5] S. Hu, F. Rusek, and O. Edfors, "Beyond massive mimo: The potential of positioning with large intelligent surfaces," *IEEE Transactions on Signal Processing*, 2018.
- [6] O. L. López, H. Alves, R. D. Souza, S. Montejó-Sánchez, E. M. G. Fernández, and M. Latva-Aho, "Massive Wireless Energy Transfer: Enabling Sustainable IoT Toward 6G Era," *IEEE Internet of Things Journal*, 2021.
- [7] "REINDEER (REsilient Interactive applications through hyper Diversity in Energy Efficient RadioWeaves technology)," <https://reindeer-project.eu/>, online, accessed 2021-07-30.
- [8] D. Espling, L. Larsson, W. Li, J. Tordsson, and E. Elmroth, "Modeling and Placement of Structured Cloud Services," *IEEE Transactions on Cloud Computing*, 2016.
- [9] W. Tärneberg, A. Mehta, E. Wadbro, J. Tordsson, J. Eker, M. Kihl, and E. Elmroth, "Dynamic Application Placement in the Telco-cloud," *Future Generation Computer Systems*, 2017.
- [10] C. Nguyen, C. Klein, and E. Elmroth, "Elasticity Control for Latency-Intolerant Mobile Edge Applications," in *The 5th ACM/IEEE Symposium on Edge Computing (SEC-20)*, 2020.
- [11] P. Townend, S. Clement, D. Burdett, R. Yang, J. Shaw, B. Slater, and J. Xu, "Improving Data Center Efficiency Through Holistic Scheduling In Kubernetes," in *Int. Conf. on Service-Oriented System Engineering*, 2019.
- [12] G. Liu, C. Zhang, L. Zhao, T. Qin, J. Zhu, J. Li, N. Yu, and T. Liu, "Return-Based Contrastive Representation Learning for Reinforcement Learning," *CoRR*, 2021. [Online]. Available: <https://arxiv.org/abs/2102.10960>
- [13] C.-L. Li, K. Sohn, J. Yoon, and T. Pfister, "CutPaste: Self-Supervised Learning for Anomaly Detection and Localization," *CoRR*, 2021.
- [14] Ericsson, "Building a stronger company long term," *Ericsson Annual Financial Report*, 2019. [Online]. Available: <https://www.ericsson.com/495c1f/assets/local/investors/documents/2019/ericsson-annual-report-2019-en.pdf>
- [15] V. Petrov, M. A. Lema, M. Gapeyenko, K. Antonakoglou, D. Moltchanov, F. Sardis, A. Samuylov, S. Andreev, Y. Koucheryavy, and M. Dohler, "Achieving end-to-end reliability of mission-critical traffic in softwarized 5G networks," *IEEE Journal on Selected Areas in Communications*, 2018.
- [16] S. Kollberg, E. Lakew, P. Svärd, E. Elmroth, and J. Tordsson, "Spreading the heat: Multi-cloud controller for failover and cross-site offloading," in *The 2nd Int. Workshop on Recent Advances for Multi-Clouds and Mobile Edge Computing (M2EC)*. Springer Int. Publishing, 2020.
- [17] L. Larsson, W. Tärneberg, C. Klein, and E. Elmroth, "Quality-elasticity: Improved resource utilization, throughput, and response times via adjusting output quality to current operating conditions," in *Int. Conf. on Autonomic Computing (ICAC)*. IEEE, 2019.
- [18] M. Maggio, A. Hamann, E. Mayer-John, and D. Ziegenbein, "Control-system stability under consecutive deadline misses constraints," in *32nd Euromicro Conf. on Real-Time Systems (ECRTS)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [19] M. Hertneck, S. Linsenmayer, and F. Allgöwer, "Stability analysis for nonlinear weakly hard real-time control systems," *CoRR*, 2020.
- [20] W. Tärneberg, P. Skarin, K.-E. Årzén, and M. Kihl, "Factory automation meets the cloud: Realizing resilient cloud-based optimal control for industry 4.0," *Submitted to IEEE Transactions on Industrial Informatics*, 2021.
- [21] X. Long, B. Liu, F. Jiang, Q. Zhang, and X. Zhi, "FPGA virtualization deployment based on Docker container technology," in *5th Int. Conf. on Mechanical, Control and Computer Engineering*, 2020.
- [22] M. Aleksandrov and T. Walsh, "Online fair division: A survey," in *AAAI Conf. on Artificial Intelligence*, 2020.
- [23] P. Skarin, W. Tärneberg, K.-E. Årzén, and M. Kihl, "Control-over-the-cloud: A performance study for cloud-native, critical control systems," in *Int. Conf. on Utility and Cloud Computing (UCC)*. IEEE, 2020.
- [24] N. Vreman, A. Cervin, and M. Maggio, "Stability and performance analysis of control systems subject to bursts of deadline misses," in *Euromicro Conf. on Real-Time Systems (ECRTS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [25] A. B. S. Forti, M. Gaglianese, "Lightweight self-organising distributed monitoring of Fog infrastructures," *Future Generation Computer Systems*, 2021.