Book Section:

Preface

We are glad to introduce the proceedings of the 17th International Conference on the Economics of Grids, Clouds, Systems, and Services (GECON 2020). GECON 2020 was held September 15–17, 2020, virtually hosted by the University of Ljubljana, Slovenia, due to the 2020 COVID-19 pandemic. The conference, which is held annually, is now firmly established as a place of convergence among economics and computer science researchers, with the ultimate aim of building a strong multidisciplinary community in this increasingly important area.

Nowadays, economics plays a pervasive role in the ICT world and is essential in strategic decisions concerning the development of new technologies. It influences its deployment, rollout plans, and is concerned in everyday operations and resource allocation optimization. However, the relationship between ICT and economics is really a two-way one, since the integration of technologies such as the Internet of Things, Artificial Intelligence, Edge, Fog and Cloud computing, and Blockchain is going to change the way economic transactions are carried out. A conference such as GECON, therefore, plays a leading role due to its blending of skills and knowledge from both worlds.

We received 40 submissions in response to our call for papers. Each paper was peer-reviewed by at least three members of the International Program Committee (PC). Based on significance, novelty, and scientific quality, we selected 11 full papers (a 27% acceptance rate), which are included in this volume. Additionally, three shorter work-in-progress papers and six extended abstracts describing the work shown on posters during the conference are also integrated in the volume.

This volume has been structured following the six sessions that comprised the conference program:
- Smartness in Distributed Systems
- Decentralising Clouds to Deliver Intelligence at the Edge
- Digital Infrastructures for Pandemic Response and Countermeasures
- Dependability and Sustainability
- Economic Computing and Storage
- Poster session

The first session on "Smartness in Distributed Systems" had four full papers. The session started with the best paper award entitled "A Consensus for Message Queue Based on Byzantine Fault Tolerance" by Jiahui Zhang et al. In this paper, the authors propose a consensus mechanism based on Byzantine fault tolerance in the context of a consortium blockchain. The objective of the work is to increase the resiliency of message queue-based consensus algorithms, since these algorithms do not tolerate malicious nodes. In the next work "Automatic QA-pair generation for incident tickets handling- an application of NLP" by Mick Lammers et al., the authors describe their experience of building a question answer system for IT support based on a large ticket dataset. The paper "ProtectDDoS: A Platform for Trustworthy Offering and Recommendation of Protections" by Muriel Franco et al. proposes a Blockchain-based solution for offering and selecting counter-DDoS services based on trustworthy service information and recommendations. Finally, the paper "Delivering Privacy-Friendly Location-Based Advertising over Smartwatches: Effect of Virtual User Interface" by Emiel Emanuel and Somayeh Koohborfardhaghighi, presents an empirical study of users' perception of a new design for enhancing privacy of location-based services.

The second session focused on "Decentralising Clouds to Deliver Intelligence at the Edge" and included two work-in-progress papers and two extended abstracts. The first work-in-progress paper

"GEM-Analytics: Cloud-to-Edge AI-Powered Energy Management" by Daniele Tovazzi et al. presents a platform that exploits fog computing, to enable AI-based methods for energy analysis at the edge of the network. The next work-in-progress paper, entitled "Using LSTM Neural Networks as Resource Utilization Predictors: The case of training Deep Learning models on the Edge" by John Violos et al., tackles the cost reduction issue by proposing an edge resource prediction model, which is used to optimize the resources by minimizing resource waste and preventing QoS infringements. The next work presented is an extended abstract entitled "Towards a semantic edge processing of sensor data in a smart factory" by Paula-Georgiana Zălhan et al. This work presents an experimental work, to support the decision making on the point in the pipeline where the semantic annotation should take place. Finally, the next extended abstract "Distributed Cloud Intelligence: Implementing an ETSI MANO-Compliant Predictive Cloud Bursting Solution using Openstack and Kubernetes" by Francescomaria Faticanti et al. presents a cloud platform for cloud bursting based on deep learning.

The third session on "Digital Infrastructures for Pandemic Response and Countermeasures" had two full papers. In "A MDE Approach for Modelling and Distributed Simulation of Health Systems", Unai Arronategui et al. present a model-driven engineering approach to support the modelling of healthcare systems in epidemic episodes combining different perspectives, and the translation to efficient code for scalable distributed simulations. This was followed by the paper "South Korea as the role model for Covid-19 policies? An analysis of the effect of Government policies on infection chain structures" by Alexander Haberling et al. which provides an analysis of the impact of policies to the COVID-19 pandemic based on a social network analysis, applying a combined method of network analysis and multiple regression.

The fourth session covered the "Dependability and Sustainability" topic and included four full papers. In "A Network Reliability Game", Patrick Maillé et al. present a game theory model on the interactions among participants (nodes) in an ad-hoc network, considering reliability investments and demand of all nodes for reliable access to a given point. In "NuPow: Managing Power on NUMA Multiprocessors with Domain-Level Voltage and Frequency Control", Changmin Ahn et al. introduce NuPow, a hierarchical scheduling and power management framework for architectures with multiple cores per voltage and frequency domain and non-uniform memory access (NUMA) properties. The paper demonstrates the feasibility of the hierarchical design with domain-specific control for the cores, the frequency and voltage domains by providing and evaluating a working implementation on the Intel Single-Chip Cloud Computer (SCC). In "Multi-Tier Power-Saving Method in Cloud Storage Systems for Content Sharing Services", Horleang Choeng et al. explore the notion of power-saving for storage systems through a novel methodology, targeting cloud content sharing services, with the aim to reduce the number of accesses in disks that are in power-saving mode. In "Instant Live Migration", Changyeon Jo et al. present the design of a distributed shared memory system that is tailored to virtual environments, e.g. cloud data centres, where live migration of virtual machines is key. The proposed system uses remote direct memory access, to reduce remote memory access latency and runs virtual machines without a significant performance degradation.

The "Economic Computing and Storage" topic was covered in the fifth session which included two full papers. In "Towards Economical Live Migration in Data Centers", Youngsu Cho et al. argue that live migration-aware service-level objectives and smart live migration open new opportunities for economical live migration in warehouse-scale computing, and present a framework based on machine learning techniques to predict key parameters of virtual machines live migration. In "Index-Selection for Minimizing Costs of a NoSQL Cloud Database", Sudarshan Chawathe proposes a cost-

model for secondary indexing in provider-hosted Cloud NoSQL stores and, in particular, for Amazon Web Services (AWS) DynamoDB with the aim of helping cloud users' cost optimization.

The conference program also included a poster session for authors to present their on-going work. In "The Influence of Online Reviews on API Adoption: A Developer-centric View", Konrad Kirner et al. discuss how online customer reviews affect developers when selecting an online API and propose a framework to study such influence. In "Power of the Chunks: Bridging Education Services and Consumer Expectations through Reusable Knowledge Modules", Djamshid Sultanov et al. discuss a new way of optimizing curriculum composition (as an educational service) by considering the consumer expectations. System dynamics is used for investigating stakeholders' value creation and exchange in terms of educational service offerings and fulfilment of consumer expectations. In "Exascale Computing Deployment Challenges", Karim Djemame et al. argue that Exascale computing deployment drives a once in a generation shift of computation, and, therefore, requires fundamentals of computer science to be re-examined. Moreover, it requires a significant degree of co-design and close attention to the economics underlying the deployment challenges ahead.

Additionally to these topic sessions, this year's GECON featured two keynotes, evenly distributed, and a Wild-and-Crazy-Ideas session.

The keynote speaker on the first day was Prof. Schahram Dustdar, from TU Wien, Austria. Prof. Dustdar keynote "The New Convergence - Research challenges in the IoT/Edge/Fog/Cloud Continuum" presented Edge Computing, IoT, AI and Human Augmentation as major technology trends, driven by recent advancements in Edge Computing, IoT, and AI accelerators. As humans, things and AI continue to grow closer together, systems engineers and researchers are faced with new and unique challenges. In this talk, Dustdar analysed the role of edge computing and AI in the cyber-human evolution and identified challenges that edge computing systems will consequently be faced with. A closer look is needed at how a cyber-physical fabric will be complemented by AI operationalisation, to enable seamless end-to-end edge intelligence systems.

The keynote speaker on the second day was Iain James Marshall, CEO of Amenesik. His keynote "Building the European Cloud Service Federation for Business Needs" is concerned with the current state of cloud application development and deployment. As cloud application deployment is still very much a manual process today, it requires a substantial software development team to be engaged and maintained throughout the entire cloud application life cycle. In his talk, Marshall envisages a technological rupture from this theme in the form of fully automated cloud service federation, which does not only support both specialized and generic applications but also fully automated cost and revenue sharing between the collaborating federation members. He illustrated how this technology, which emerged during the European Union's Horizon 2020 BASMATI project, can be applied in several use case specific service federations, which are defined in the context of the SERENE H2020 Innovation Action project, in conjunction with an underlying generic cloud service federation offer. Moreover, he detailed the environment and processes, which are needed to make this technology successful and sustainable.

The Wild and Crazy Ideas session was organised by the conference host Assoc. Prof. Vlado Stankovski on the topic of the new Horizon 2020 Next Generation Internet project entitled "OntoChain: Trusted, Transparent, and Traceable Ontological Knowledge on Blockchain".

Acknowledgments