IEEE COMMUNICATIONS SOCIETY
Cognitive Networks Technical Committee
http://committees.comsoc.org/tccn

TCCN Newsletter

Vol. 5, No. 1, May 2019

CONTENTS

Chair's Message

Dear Fellow TCCN Colleagues,

I would like to express my sincere thanks to the 2017-2018 TCCN Chair, Prof. Jianwei Huang, and other TCCN officers for their enthusiastic support and services to the community. Together with our fellow TCCN members, TCCN has achieved a number of milestones in the past two years such as recertification, awards, SIGs, etc.

I am very happy writing to you for the first time as the TCCN Chair in the Newsletter. I would like to take this opportunity to introduce to you the newly elected officers for 2019 – 2020:

- Chair: Yue Gao, Queen Mary University of London, UK
- Vice-Chair Europe/Africa: Oliver Holland, Kings College London, UK
- Vice-Chair Asia Pacific: Lingyang Song, Peking University, China
- Vice-Chair Americas: Daniel Benevides da Costa, Federal University of Ceará, Brazil
- Secretary: Lin Gao, Harbin Institute of Technology, Shenzhen, China

All elected officers have been very active in the TCCN community, and some of them have served in various TCCN leadership roles during the past few years.

We have also appointed several officers during the past few months, including:

- Standards Liaison:
  Oliver Holland, Kings College London UK.
- Publicity Board:
  Vijay Rao, Delft University of Technology, Netherlands
  Yuan Man, Shenzhen University, China.
- Student Competition Program:
  Lucio Marcenaro, University of Genova, Italy
  Sai Huang, Beijing University of Posts and Telecommunications, China

The elected and appointed officers will work together with our members to try our best to serve the TCCN community.

For the Newsletter, I am glad that it has become an important electronic platform for TCCN members to exchange research ideas and brainstorm about the future of the community. Starting from this issue, TCCN Vice-Chair, Daniel Benevides da Costa, will serve as the Newsletter Director, who has a great passion for the platform. I am sure that he will be able to take the Newsletter to the next level.

We are looking for more volunteers to actively engage in various aspects of the TC, including but not limited to:

- Organize Special Interest Groups (SIGs) (contact: Yue Gao, Vijao Rao)
- Organize Special Issues for the TCCN Newsletter (contact: Daniel Benevides da Costa)
- Contribute to the publicity efforts of TCCN (contact: Lin Gao, Yuan Ma)
- Contribute to student competition program (contact: Lucio Marcenaro, Sai Huang)
- Involve TCCN in Comsoc conference organization (contact: Lingyang Song)
- Involve TCCN in Comsoc journal special issues (contact: Yue Gao)

As always, I welcome any suggestions from TCCN members regarding how to make TCCN a better community. Please feel free to contact me at yue.gao@qmul.ac.uk if you have any suggestions.

Thanks and best regards,

Yue Gao
Chair, IEEE ComSoc TCCN
EPSRC Fellowship Award Holder (2018-2023)
Queen Mary University of London
https://wmc.eecs.qmul.ac.uk

Director's Message

Since December 2015, this Newsletter has presented and discussed some emerging topics related to the TCCN areas of interest. More specifically, it has covered a broad range of applications and techniques, for instance, non-orthogonal multiple access, ultra-reliable low-latency communications (URLLC), millimeter wave communications, unmanned aerial vehicle (UAV) communications, and massive machine-type communications (mMTC). We have interviewed over a dozen experts in these fields, included several interesting position papers, and provided state-of-the-art reviews. My sincere thanks to all the previous directors for their contributions and help which have made this Newsletter a great success.

From my side, I have contributed in two previous TCCN Newsletter editions as Feature Editor, and this TCCN Newsletter issue is the first one that I am acting as Director. It has been a great pleasure and honor for me, and I am excited to cover two areas that will likely have impact in 5G and beyond: a) Blockchain and b) Internet of Things (IoT). In the Blockchain area, we have interviewed Prof. Dusit Niyato, from Nanyang Technological University, Singapore, and Dr. Bhaskar Krichnamachari, from USC, who are leading experts in this area. We have also had the pleasure to get a position paper from Dr. Bhaskar Krichnamachari. Within the context of IoT, we have interviewed Prof. Luiz A. da Silva, from Trinity College Dublin, Dr. Samir Perlaza, from INRIA, France, and Prof. Sergey Andreev, from Tampere University, Finland, who provided us with their outlook on the opportunities and challenges of IoT. Finally, Prof. Giancarlo Fortino and Dr. Claudio Savaglio, from Università della Calabria, Italy, provided a position paper that presents ACOSO-Meth (Agent-based Cooperating Smart Objects Methodology), the first agent-based methodology that specifically and seamlessly supports the main phases of engineering of IoT ecosystems and related services.

Finally, I would like to thank our two feature topic editors: Prof. Walid Saad, from Virginia Tech - USA, and Prof. Pedro H. J. Nardelli, from Lappeenranta University of Technology - Finland, for their efforts in arranging the content of this Newsletter. Moreover, we want to thank all authors and interviewees for sharing with us their experience and time. I would finally like to acknowledge the gracious support from the TCCN chair, Dr. Yue Gao and all TCCN officers. If you have any suggestion, feel free to contact me at: danielbcosta@ieee.org. We hope that you enjoy the material of this Newsletter!

Daniel Benevides da Costa
Director, IEEE ComSoc TCCN Newsletter
Federal University of Ceará, Brazil

Feature Topic: Blockchain

Editor: Walid Saad

Department of Electrical and Computer Engineering, Virginia Tech., USA
Email: walid@vt.edu

Blockchain and distributed ledger technologies (DLT) are seen as one of the most important technological breakthroughs of the past decade. Originally conceived as a technology to enable crypto-currency, blockchain concepts have now gone way beyond their initial use in Bitcoin and are rapidly becoming a pillar of many industries ranging from the Internet of Things to healthcare and critical infrastructure.

The adoption of blockchains across these industries requires overcoming a plethora of technical challenges across computation, communication, security, and optimization. In particular, blockchain concepts admit a plethora of applications in the wireless networking and communication domains, ranging from improving spectrum sharing to the introduction of new, decentralized incentive mechanisms for managing wireless resources. In addition, blockchains can provide trustworthy authentication and communication in large-scale, decentralized wireless networks. Along with their technically rich application domain, blockchains can also enhance the economics of wireless networking through micropayments and related ideas such as smart contracts.

To reap the benefits of blockchains for wireless networking, it is imperative to identify the main challenges and opportunities related to wireless-oriented blockchain research. In consequence, this feature topic of this TCCN newsletter brings together input from two leading experts in the various areas of blockchains and DLT, so as to put forward a wireless-oriented research agenda for blockchain technologies.

First, we interview Dr. Bhaskar Krichnamachari from USC to get his view on the various research and technical challenges related to blockchains. Then, Dr. Krichnamachari provides a holistic position paper that outlines the importance of blockchain concepts for networked systems. The position paper provides a forward-looking view on how blockchains can be integrated into tomorrow's wireless and communication networks while also identifying four key research opportunities in this area. Then, we provide a second interview with Dr. Dusit Niyato from Nanyang Technological University in Singapore who shares with us his insights on the use of blockchains in wireless networks. He particularly shares his expertise on the design of new analytics to understand the performance improvements and gains that blockchain can provide across a broad range of wireless applications.

In a nutshell, this feature topic gathers together key insights on the emerging area of blockchains, with a focus on wireless and networking applications. We believe that the insights and ideas discussed in this feature topic will pave the way towards a plethora of novel research directions and will shed more light on the synergies between communications and blockchains.

Walid Saad (S'07, M'10, SM'15, F'19) received his Ph.D degree from the University of Oslo in 2010. Currently, he is a Professor at the Bradley Department of Electrical and Computer Engineering at Virginia Tech, where he leads the Network sciEnce, Wireless, and Security

(NEWS) laboratory, within the Wireless@VT research group. His research interests include wireless networks, machine learning, game theory, unmanned aerial vehicles, cybersecurity, and cyber-physical systems. Dr. Saad is the recipient of the NSF CAREER award in 2013, the AFOSR summer faculty fellowship in 2014, and the Young Investigator Award from the Office of Naval Research (ONR) in 2015. He was the author/co-author of seven best paper awards at major conferences. He is the recipient of the 2015 Fred W. Ellersick Prize from the IEEE Communications Society, of the 2017 IEEE ComSoc Best Young Professional in Academia award, and of the 2018 IEEE ComSoc Radio Communications Committee Early Achievement Award. From 2015-2017, Dr. Saad was named the Stephen O. Lane Junior Faculty Fellow at Virginia Tech and, in 2017, he was named College of Engineering Faculty Fellow. He currently serves as an editor for the IEEE Transactions on Wireless Communications, IEEE Transactions on Mobile Computing, IEEE Transactions on Cognitive Communications and Networking, and IEEE Transactions on Information Forensics and Security. He is an Editor-at-Large for the IEEE Transactions on Communications.

Interview with Prof. Dusit Niyato

Nanyang Technological University (NTU), Singapore
Email: dniyato@ntu.edu.sg

Q1: What is, in your opinion, the most appropriate way to define the concept of Blockchains in order to allow our readers to navigate the various seemingly similar terminologies such as Blockchain, Bitcoin, Distributed Ledgers?

A1: **Distributed ledger (DLT)** is the basic concept that eliminates any central decision-making authority. When decision-making is decentralized, EVERY single participant is necessarily involved and decisions are made (transactions are confirmed) based on consensus. There are various types of consensus algorithms available and every network chooses its own, the network participants, follow that consensus model. **Blockchain** is a technology which implements this above concept and also respects other critical tenets of DLT like security and immutability. **Bitcoin** is the first and the most top application of blockchain.

Q2: What do you think are the three most important technical challenges in the area of Blockchain as it stands today?

A2: There are three important technical challenges which limit blockchain technology unusable for mainstream applications.
(1) **Limited scalability**: The total number of transactions that can be processed is so limited that cannot be widely used in the most of Internet of Thing scenarios.
(2) **Privacy**: Transactions on public blockchains are recorded on public ledger, which is possible to link your identity to the address by observations.
(3) **Access to external data**: Blockchain services cannot inherently make arbitrary network requests to access data outside the network. Suppose if blockchain service retrieves some information from an external source, this retrieval is then to be performed repeatedly and separately by each node. But because this source is outside of the blockchain, there is no guarantee that every node will receive the same answer.

Q3: How do you view the role of Blockchains and related concepts in the context of wireless networks? How can Blockchain benefit communication networks and vice versa, how can communication networks sustain Blockchain technologies?

A3: (1) Blockchain is a distributed ledger in nature which can ensure the security of resource trading, e.g., computing resources, in a decentralized manner in the context of wireless networks. (2) The development of communication networks can promote the development of consensus algorithms to a certain extent, e.g., faster consensus process and less energy consumption during consensus. While the blockchains enable more secure and reliable resource sharing and trading thus improve the performance of resource cooperation and utilization. (3) The communication networks support the block propagation and verification, and information communication among blockchain users, miners and verifiers, etc.

Q4: What is the most disruptive Blockchain concept that has emerged in the past two years?

A4: In my opinion, the most disruptive blockchain concept is Internet of Value. Internet of value means an exchange anything of value like foreign currency payment, stocks, securities, intellectual property rights, scientific discoveries, etc. should be processed instantly much like what information has been doing for decades. So in the internet of value, anything valuable such as mentioned above can be transferred in a jiffy. Usually there is a middle man when you want to transfer money, such as banks, and the government. But in the internet of value, you can transfer money directly, and quicker. You can use Bitcoin for example where there is no third party involved.

Q5: When do you think we will start to see actual Blockchain implementations outside of crypto currency and which domains will benefit the most from the technology?

A5: I think that we had actual blockchain implementations, for example, the China's tech conglomerate Tencent has officially incorporated a blockchain electronic invoice system in the WeChat, the company's popular messaging, social media and payment app with over 1 billion users, in a bid to simplify the process of reimbursing company employee expenses.
In the near future, we can see many actual blockchain implementations except the crypto

currency. I think that the domain will benefit the most from blockchain is healthcare. Blockchain technology can allow hospitals to safely store data like medical records and share it with authorized professionals or patients. This can improve data security and can even help with accuracy and speed of diagnosis. Gem and Tierion are two companies that are working on disrupting the current healthcare data space.

Q6: Do you think AI will have a role to play in Blockchain systems? What is that role and why would AI be well-poised to play it?

A6: AI is the accelerator for blockchain systems. AI can be used to overcome the following limitations of the existing blockchain systems: (1) optimizing energy consumption using AI-based algorithms; (2) improving scalability through distributed or decentralized AI algorithms; (3) ensuring block data and blockchain system security using AI-based intrusion detection schemes.

Q7: Could you please briefly introduce the most recent research project(s) that you have done in this area? (Please explain the key idea(s) and interesting findings)?

A7: Recently, we had designed an incentive mechanism for secure block verification in DPoS-based blockchain through joint reputation and contract theory optimization. This work shows that the security of block verification can be significantly improved through reputation-based verifier selection scheme using contract theory. More details can be found in the following paper.

J. Kang, Z. Xiong, D. Niyato, D. Ye, D. I. Kim and J. Zhao, "Toward Secure Blockchain-Enabled Internet of Vehicles: Optimizing Consensus Management Using Reputation and Contract Theory," in *IEEE Transactions on Vehicular Technology*, vol. 68, no. 3, pp. 2906-2920, March 2019.

Q8: Beyond your own work, are there any resources that you would like to recommend, especially to those who are new in this field and want to learn more about Blockchains? Are there any specific resources that you recommend related to Blockchains in the context of wireless and communication networks?

A8: For the researchers, they can read some survey papers about blockchain. Such as

[1] R. Yang, F. R. Yu, P. Si, Z. Yang and Y. Zhang, "Integrated Blockchain and Edge Computing Systems: A Survey, Some Research Issues and Challenges," in *IEEE Communications Surveys & Tutorials*. doi: 10.1109/COMST.2019.2894727.

[2] Liu, Ziyao, et al. "A Survey on Applications of Game Theory in Blockchain." *arXiv preprint arXiv:1902.10865 (2019).*

[3] Wang, Wenbo, et al. "A survey on consensus mechanisms and mining management in blockchain networks." *arXiv preprint arXiv:1805.02707 (2018).*

For the developers, they can try to learn and develop some blockchain applications through Solidity (Link: https://solidity.readthedocs.io/en/v0.5.5/).

Q9: What are the most important open problems and future research directions in this area?

A9: The most important open problem is the **scalability** challenge in IoT scenario. The future research directions in this area include: interaction issues in cross chain, security issues of side chain, AI for scalable blockchain.

Q10: Do you think Blockchains are just a hype or will they sustain their seemingly revolutionary role in the next decade?

A10: I think the blockchain will sustain their seemingly revolutionary role in the next decade for the goal of realizing Internet of Value.

Dusit Niyato is currently a professor in the School of Computer Science and Engineering and, by courtesy, School of Physical & Mathematical Sciences, at the Nanyang Technological University, Singapore. He received B.E. from King Mongkuk's Institute of

Technology Ladkrabang (KMITL), Thailand in 1999 and Ph.D. in Electrical and Computer Engineering from the University of Manitoba, Canada in 2008. He has published more than 380 technical papers in the area of wireless and mobile networking, and is an inventor of four US and German patents. He has authored four books including "Game Theory in Wireless and Communication Networks: Theory, Models, and Applications" with Cambridge University Press. He won the Best Young Researcher Award of IEEE Communications Society (ComSoc) Asia Pacific (AP) and The 2011 IEEE Communications Society Fred W. Ellersick Prize Paper Award. Currently, he is serving as a senior editor of IEEE Wireless Communications Letter, an area editor of IEEE Transactions on Wireless Communications (Radio Management and Multiple Access), an area editor of IEEE Communications Surveys and Tutorials (Network and Service Management and Green Communication), an editor of IEEE Transactions on Communications, an associate editor of IEEE Transactions on Mobile Computing, IEEE Transactions on Vehicular Technology, and IEEE Transactions on Cognitive Communications and Networking. He was a guest editor of IEEE Journal on Selected Areas on Communications. He was a Distinguished Lecturer of the IEEE Communications Society for 2016-2017. He was named the 2017, 2018 highly cited researcher in computer science. He is a Fellow of IEEE.

Interview with Prof. Bhaskar Krishnamachari

University of Southern California, USA
Email: bkrishna@usc.edu

Q1: What is, in your opinion, the most appropriate way to define the concept of Blockchains in order to allow our readers to navigate the various seemingly similar terminologies such as Blockchain, Bitcoin, Distributed Ledgers?

A1: While Satoshi Nakamoto's Bitcoin Cryptocurrency was the original Blockchain protocol, in light of further developments, it is helpful to define Blockchain more broadly. At their core, Blockchain technologies are about maintaining an immutable distributed ledger using consensus. In light of newer protocols adopting more sophisticated data structures such as directed acyclic graphs (DAG) beyond linear chains, Distributed Ledger Technology might indeed be a more suitable name, but "Blockchain" is how the entire field has come to be known popularly. It is important to be aware that there are now both open and permissioned Blockchains, where the latter, typically proposed for industry enterprise use cases beyond cryptocurrency such as supply chain tracking, assume that each node involved in maintaining the ledger and submitting transactions has a known identity.

Q2: What do you think are the three most important technical challenges in the area of Blockchain as it stands today?

A2: I would identify three most important technical challenges as 1) Performance 2) Decentralization and 3) Interoperability. With respect to the first challenge of performance, today's blockchain protocols suffer from low transaction throughput, high latency, high energy consumption. While performance could be significantly improved by reducing the level of decentralization, the second challenge is to do so while making sure the systems are decentralized with respect to the number of independent nodes or entities that are needed to maintain and secure operation. The third challenge is to improve interoperability and connect not only many different blockchain systems to each other, but also to existing network protocols and systems so that rich new applications can be designed and deployed.

Q3: How do you view the role of Blockchains and related concepts in the context of wireless networks? How can Blockchain benefit communication networks and vice versa, how can communication networks sustain Blockchain technologies?

A3: Blockchain technologies can be beneficial to wireless and general communication networks in many ways - they can be used, for example, to increase trustworthiness of application-layer services, to decentralize spectrum resource allocation, to enable implementation of economic mechanisms within networks, and to enable micropayments for data and compute services over networks. In the other direction, as well, there are many opportunities: improvements in the latency of communication networks can be used to enhance the throughput and latency performance of Blockchain consensus protocols, and wireless localization services can be used to develop energy-efficient proof of location consensus protocols.

Q4: What is the most disruptive Blockchain concept that has emerged in the past two years?

A4: It's been more than two years in the making, but I think the most powerful new idea in Blockchain technologies is that computation can be made more transparent and trustworthy through the use of smart contracts, which are autonomous pieces of code that run in a decentralized and secure manner. They make possible many rich new ideas to be tried and tested on blockchains, from escrow services to enable trusted digital payments to prediction and curation markets for data and other services.

Q5: When do you think we will start to see actual Blockchain implementations outside of crypto currency and which domains will benefit the most from the technology?

A5: These are already starting to happen, though still largely at a proof of concept level, production-level implementation may take 5-10 years more because of technical challenges associated with scaling as well as non-technical challenges associated with getting buy-in from many partners in industry consortiums. Some of the domains that are likely to benefit from the technology are data monetization and marketplaces, supply chain, decentralized curation and prediction.

Q6: Do you think AI will have a role to play in Blockchain systems? What is that role and why would AI be well-poised to play it?

A6: I think this is somewhat speculative at this time perhaps in that I haven't seen many good examples of AI playing a role in Blockchain systems today, but potentially AI tools could be used to enhance the ease and security of writing code for or interacting with Blockchain-based decentralized applications, or the decentralized applications leveraging the trustworthiness and transparency provided by a Blockchain protocol themselves may be AI-based autonomous applications. There are also some proposals to enable decentralized training of machine learning models using blockchain to provide privacy in cases where the training data is inherently spread across multiple parties.

Q7: Could you please briefly introduce the most recent research project(s) that you have done in this area? (Please explain the key idea(s) and interesting findings)?

A7: I have been particularly excited about using Blockchain technology to a) improve trust in and b) provide economic value for data streams that naturally cross organizational, economic and trust boundaries. This is valuable for large-scale multi-party Internet of Things (IoT) applications such as for smart cities and supply chains. Under this project, supported by the USC Viterbi Center for Cyber-Physical Systems and the Internet of Things (https://cci.usc.edu), we have been developing, analyzing and evaluating several new protocols and systems, including 1) SDPP - a streaming data payment protocol, 2) a dual-deposit escrow smart contract for buying and selling digital goods, 3) Trinity - a system that allows decentralized operation of publish-subscribe brokers guaranteeing that all subscribers see the same stream, 4) PayFlow - a mechanism to allow flows to pay an SDN controller for QoS reservations, and 5) DDM - a framework for decentralized data marketplaces.

Q8: Beyond your own work, are there any resources that you would like to recommend, especially to those who are new in this field and want to learn more about Blockchains? Are there any specific resources that you recommend related to Blockchains in the context of wireless and communication networks?

A8: I recently compiled an annotated bibliography of papers on blockchain and distributed ledger technologies that can be useful for beginning researchers in this area, it can be found online at http://tiny.cc/bcbib. Slides from a tutorial I gave at MobiHoc 2018 with further pointers to the literature can be found at http://tiny.cc/bctut. The literature connecting Blockchain to networks is still in its infancy, but I highly recommend the Blockstack paper by Muneeb Ali et al., USENIX 2016, which shows how application layer services on the Internet such as DNS and PKI could be made more secure and trustworthy by decentralizing them using Blockchain.

Q9: What are the most important open problems and future research directions in this area?

A9: The good news for researchers is that today's state of the art in Blockchain technology is far from mature. The analogy I often give is that developing Blockchain-based applications today is like trying to fly an airplane that is still being built. There are many research opportunities, ranging from developing more scalable consensus mechanisms (higher transaction throughput, lower delay, lower energy resource consumption) to enhancing decentralization and security, to developing novel algorithms including game-theoretic mechanisms that leverage the capabilities provided by autonomous smart contracts to enable new capabilities for network protocols and networked applications. I would encourage networking researchers to also study the various components that are being developed and deployed as part of blockchain technologies and think about creative ways in which they could be applied to problems in networking.

Q10: Do you think Blockchains are just a hype or will they sustain their seemingly revolutionary role in the next decade?

A10: It is important to acknowledge that because of their deep connection with cryptocurrency speculation, and also because of relentless marketing in some quarters, there is indeed a lot of hype mixed in with real potential today. Though there are many technical and social challenges to adopting blockchain technologies at scale, I am optimistic that they will contribute and become a part of our engineered systems because they do bring fundamentally new models and capabilities with respect to trust and decentralization. I would counsel patience regarding the timeline over which sustainable impact will be seen. Like with other engineering fields like Artificial Intelligence, Communication Theory, or Networking itself, the technology and its widespread application may continue to mature over the course of many years, possibly even decades, as we develop a deeper understanding and develop real-world applications of importance to society.

Position Paper: Blockchain Technologies for Networked Applications

Bhaskar Krishnamachari
Ming Hsieh Department of Electrical and Computer Engineering
Center for Cyber-Physical Systems and the Internet of Things
Viterbi School of Engineering, University of Southern California

1. Abstract

I present an overview of Blockchain technology, summarizing recent technology and application developments. Using case studies from our recent and ongoing research in this area, I illustrate some topics at the intersection of Blockchain and Networking research, and identify some future research opportunities.

2. Overview of Blockchain

In 2009, Bitcoin, the original Blockchain protocol was released anonymously by some person(s) under the pseudonym Satoshi Nakamoto [1]. While it has been shown that Bitcoin itself drew on decades of academic work in many areas from cryptography to distributed systems [2], it was an enormously impressive achievement combining several elements in a very sophisticated manner. In a nutshell, Bitcoin provides for a public, ordered, essentially immutable ledger represented by a hashed chain of transaction-containing blocks that is maintained in a distributed manner using consensus by thousands of P2P nodes around the world, allowing anyone to post transactions anonymously. It provides economic mechanisms to incentivize the operation of nodes, uses public key cryptography to achieve anonymity and employs distributed solution of proof of work puzzles (referred to as mining) to provide security against Sybil attacks.

Since then, the underlying blockchain technology itself has been enhanced and developed in several directions. It has been extended to incorporate more general quasi-Turing complete on-chain computation, most notably by Ethereum [3], where scripts for such computation are referred to as smart contracts. The original construct of a linear hashed chain of blocks has been replaced by more general data structures such as Directed Acyclic Graphs, in protocols such as IOTA [4] and more recently, Avalanche [5]. In light of the significant energy requirement of Proof of Work, alternative protocols have been proposed that employ Proof of Stake, such as Ouroboros [6] and Algorand [7]. Some blockchain protocols aim to provide greater levels of anonymity and privacy, such as ZCash [8].

For enterprise use cases that go beyond cryptocurrency, industry has also pioneered the design and deployment of "permissioned" blockchain protocols such as Hyperledger Fabric [9] in which anonymity is abandoned (obviating the need for Sybil control), allowing more traditional forms of distributed Byzantine fault-tolerant consensus protocols to be employed. Unlike open blockchains, such protocols are intended to be deployed by consortiums belonging to particular industry verticals.

Even as the technology is developing at a rapid pace, applications are being explored in many directions. Open blockchains from Bitcon onwards have focused significantly on cryptocurrency transactions and this has been a primary use case. Beyond this, the availability of smart contracts allows the development of other distributed applications ("dapps") that utilize tokens for various uses, from prediction markets (e.g., Augur), to incentivizing content creation and list curation (e.g., Steemit and Adchain) to more frivolous entertainment-oriented use cases such as Cryptokitties. Another class of applications that has been explored are decentralized marketplaces (e.g., OpenBazaar) that obviate the need for centralized third-party platforms mediating between buyers and sellers.

Permissioned Blockchain protocols such as Hyperledger Fabric have largely focused on the maintenance and use of distributed ledgers for supply chain applications ranging from tracking the provenance and quality of farm-to-table products to maintaining real estate records to the industrial manufacturing of electronics and automobiles to the shipping container industry. Financial applications of blockchain beyond cryptocurrency have also been explored by consortiums of banks to speed up reconciling and settlement of accounts. Applications of blockchain are also being explored in the smartgrid community for distributed energy generation and trading. The automotive industry has been exploring use cases and standardization through efforts such as the MOBI alliance. There is also some work happening at the intersection of Blockchain and networks, such as Blockstack [10], which presents a decentralized naming service for the Internet.

3. Blockchain and Networked Applications: Case Studies

As examples of research into and with Blockchain technologies in the context of networks, I highlight below some recent work we have been doing at the University of Southern California, at the Autonomous Networks Research Group and the Viterbi Center for Cyber-Physical Systems and the Internet of Things.

a. Data Monetization: The original Internet was designed to provide value-neutral plumbing for data. While it has been successful for many applicaitons there are also many use-cases where the ability to monetize the flow of data at a protocol-level through micropayments from the recipient to the provider can be very useful. As a motivating example, consider the creation of a community IoT data marketplace [11], which allows the owners of IoT devices in a smart city to make available data streams from their device to potential customers that can use those streams for their own novel application. Today, such payment from buyers to sellers must go over a traditional (e.g., credit card-based) payment rail, which poses two challenges: a) the transaction fees can be prohibitive, effectively preventing microtransactions over data, and b) the time and manual interaction overhead associated with the transaction means that it is difficult to pay for data from devices encountered ephemerally (say, due to mobility).

To enable micropayments for data, we have developed SDPP, the Streaming Data Payment Protocol [12]. SDPP is an application layer protocol that allows a data-buying client to connect to server and provides for a full value-based transaction - including getting a menu of data streams and their unit prices, ordering from the menu of data, getting the data, invoices for the data, making payment and getting and storing receipts for the payment. By providing these capabilities at the application layer, it obviates the need for reinventing the wheel each time for a different application. SDPP combines a traditional TCP socket connection with a blockchain-agnostic micro-payment channel and a record medium (which could be implemented using any distributed ledger technology, or if acceptable, even a centralized database).

In [13], we further show how blockchain technologies can be used to build a decentralized data marketplace, such as for smart cities. By decentralizing data product postings and ratings of buyers and sellers, such a marketplace can enable participants to trust that they are not being manipulated by a third-party market operator.

b. Decentralized Publish-Subscribe: While traditional network applications such as HTTP, FTP, SMTP are designed to be based on one to one client-server communications, in the context of sensor networks and IoT, many to one, one to many and many to many real-time data flows are quite common. To support such general communication patterns, typically publish-subscribe protocols are utilized. A well-known example of such a publish-subscrbe application layer protocol is MQTT. MQTT is typically deployed as a central broker to which publish clients send messages associated with particular topics. Any clients that are interested send the broker a subscribe message for the corresponding topics, and from that point on all messages on the topic are delivered by the broker to the subscribers.

When we consider distributed IoT applications that cross organizational, and therefore trust boundaries, this simple single-broker architecture becomes problematic. Whoever owns and operates the central broker in principle has the ability to tamper with the messages (including supression and reordering, and if the messages do not include cryptographic integrity mechanisms, even modification), so that different subscribers may see potentially different versions of messages for the same topic. We recently proposed Trinity [14], a framework for decentralizing publish-subscribe brokers so that different organizations participating in a consortium (say for a supply chain appliction) can each host their own copy of a broker. The brokers in all these organization then work together using a common disttributed consensus protocol, to ensure that all streams are consistent, greatly increasing trust in the system without requiring a trusted central party.

c. Micropayments for Network QoS: In [15], we present PayFlow, a micropayment framework for software defined networks. In this system, a node desiring to reserve a certain amount of bandwidth (or more generally, any defined QoS level) for its end to end flows through a network can do so with micropayments by talking to an SDN controller. PayFlow is implemented over OpenFlow, and has been demonstrated to allow, in principle, bandwidth reservations on short time scales, on the order of seconds.

d. Proof of Location: Permissionless or open Blockchain protocols must provide a defense against Sybil attacks, so that some malicious entity doesn't game the consensus mechanism by pretending to be multiple nodes. The original Bitcoin protocol, Ethereum, and others utilize proof of work computational puzzles to provide Sybil control. In SENATE [16], we show that the ability to detect wireless signals within proximity of a device (or more generally, to localize such nodes) provides another avenue for Sybil control. By overlaying on top of location esimation a geographic election process (similar to a US Senate election where two senators are elected for each state, this mechanism elects a fixed number of devices within each geographic region), and allowing only the elected nodes to participate in the consensus process, we are able to provide such a defense.

e. Consensus in frequently partitioning networks: In another ongoing project at the intersection of Blockchain and networks, we are exporing how to implement a distributed consensus-based ledger for networks that are constantly subjected to partitions (splits) and mergers of collections of nodes. Such networks may be encountered for instance in the context of distributed robotics or UAV swarms. The crux of our approach, referred to as SwarmDAG [17] is to stitch together collections of blockchains for each partition into a directed acyclic graph and allow only a subset of transactions, namely those that have a sufficient quorum within a given partition, to be allowed into the ledger at any given time.

4. Research Opportunities

Despite 10 years of extensive development since the appearance of Bitcoin, Blockchain and distributed ledger technologies are still at a relatively immature stage. There are significant opportunities for research and development on many fronts:

- Enhancing underlay and overlay networks to improve blockchain protocol performance
- Proposing novel core blockchain protocols and higher-layer mechanisms to improve transaction throughput, confirmation latency, reduce storage requirements
- Improving energy utilization by developing alternatives to Proof of Work, particularly schemes that are well-suited to different types of networks and classes of devices
- Incorporating the capabilities provided Blockchain technologies such as immutable logging, monetization, decentralized consensus, and autonomous smart contract code into network protocols and applications

5. Conclusion

This article has presented a brief survey of blockchain technology and applications, including several case studies from research at USC at the intersection of Blockchain technologies, network protocols and applications. I hope the discussion also sheds some light on potential research directions that may be of interest to the networking community. For further reading, please find an annotated bibliography of blockchain protocols online at [18].

References

[1] Nakamoto, Satoshi. "Bitcoin: A peer-to-peer electronic cash system." (2008) https://bitcoin.org/bitcoin.pdf

[2] Narayanan, Arvind, and Jeremy Clark. "Bitcoin's academic pedigree." *Communications of the ACM* 60.12 (2017): 36-45.

[3] Wood, Gavin. "Ethereum: A secure decentralised generalised transaction ledger." *Ethereum project yellow paper* 151 (2014): 1-32. https://github.com/ethereum/yellowpaper

[4] Popov, Serguei. "The tangle." (2016). https://www.iota.org/research/academic-papers

[5] Team Rocket. "Snowflake to avalanche: A novel metastable consensus protocol family for cryptocurrencies." (2018).

[6] Kiayias, Aggelos, et al. "Ouroboros: A provably secure proof-of-stake blockchain protocol." *Annual International Cryptology Conference*. Springer, Cham, 2017.
[7] Gilad, Yossi, *et al.* "Algorand: Scaling byzantine agreements for cryptocurrencies." *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM, 2017.
[8] Hopwood, Daira, *et al.* "Zcash protocol specification." *Tech. rep. 2016–1.10. Zerocoin Electric Coin Company, Tech. Rep.*(2016).
[9] Androulaki, Elli, et al. "Hyperledger fabric: a distributed operating system for permissioned blockchains." *Proceedings of the Thirteenth EuroSys Conference*. ACM, 2018.
[10] Ali, Muneeb, et al. "Blockstack: A global naming and storage system secured by blockchains." *2016 {USENIX} Annual Technical Conference ({USENIX}{ATC} 16)*. 2016.
[11] Krishnamachari, Bhaskar, et al. "I3: an IoT marketplace for smart communities." *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 2018.
[12] Radhakrishnan, Rahul, and Bhaskar Krishnamachari. "Streaming Data Payment Protocol (SDPP) for the Internet of Things." Workshop on Blockchain for IoT in conjunction with IEEE Blockchain (2018).
[13] Ramachandran, Gowri Sankar, Rahul Radhakrishnan, and Bhaskar Krishnamachari. "Towards a Decentralized Data Marketplace for Smart Cities." *2018 IEEE International Smart Cities Conference (ISC2)*. IEEE, 2018.
[14] Ramachandran, Gowri S. et al., "Trinity: A Byzantine Fault-Tolerant Distributed Publish-Subscribe System with Immutable Blockchain-based Persistence," IEEE International Conference on Blockchain and Cryptocurrency (ICBC), Seoul, South Korea, May 2019.
[15] Chen, David *et al.* "PayFlow: Micropayments for Bandwidth Reservations in Software Defined Networks," In 1st International IEEE Workshop on the Economics of Fog, Edge and Cloud Computing (ECOFEC) in conjunction with IEEE INFOCOM, Paris, France, Apr. 2019.
[16] Jiang, Zhiyuan, et al. "SENATE: A Permissionless Byzantine Consensus Protocol in Wireless Networks." *arXiv preprint arXiv:1803.08694* (2018).
[17] Tran, Jason A. *et al.*, "SwarmDAG: A Partition Tolerant Distributed Ledger Protocol for Swarm Robotics" Symposium on Blockchain for Robotic Systems, MIT, Boston, USA, Dec. 2018.
[18] Krishnamachari, Blockchain and Distributed Ledger Protocols—a Reading List, Apr. 2019. https://medium.com/@bhaskark2/blockchain-and-distributed-ledger-protocols-a-reading-list-3c00dfc66404

Bhaskar Krishnamachari is Professor of Electrical and Computer Engineering at the Viterbi School of Engineering at the University of Southern California. He is Director of the USC Viterbi Center for Cyber-Physical Systems and the Internet of Things (CCI). He has expertise in wireless networks, IoT protocols and applications, distributed computing, machine learning, and blockchain technologies, spanning both theory and software systems. He has co-authored over 300 papers, collectively cited more than 24,000 times. He has received the NSF CAREER Award and the ASEE Terman Award for outstanding electrical and computer engineering educators, and several best paper awards including at ACM/IEEE IPSN and ACM MobiCom. In 2011, he was listed in MIT technology review magazine's TR-35 list of top 35 innovators under the age of 35, and in 2015 was named one of Popular Science magazine's "Brilliant 10".

Feature Topic: Internet of Things

Editor: Pedro H. J. Nardelli

School of Energy Systems
LUT, Finland
Email: pedro.nardelli@lut.fi

Internet of Things (IoT) has already become an established research topic within IEEE Communications Society. However, IoT is a diverse research field so that it means different things for different communities. This might range from physical layer design (e.g. how to support massive communication of machine-type) to application layer (e.g. how to improve the user interface based on graphical visualization of data collected by wearable sensors). The IEEE IoT Journal, co-sponsored by IEEE Sensors Council, Communications Society and Computer Society, indicates this broad scope of IoT and its relevance (impact factor of 5.863).

In terms of research, IoT is usually related to sensor networks and machine-type communications (machine-to-machine, or human-machine, communications). IoT is then related to short messages that can be classified in different regimes related to the application in hand. They can be related to connectivity as in massive machine-type communications (mMTC) or in ultra-reliable low-latency communications (URLLC). These regimes related to applications are at the core of the upcoming 5G and is also dominating the initial discussions of what is going to be 6G. These developments indicate that the usually human-centric mobile networks shall be designed considering machines, therefore supporting the network of such things. There is also a move from traditional "agnostic" analyses to industry-specific solutions via vertical domains (e.g. Industry 4.0, Energy, Health, Transportation, Smart City among others.).

In the next sections, we provide a better understanding of the research ideas guided the term IoT by presenting contributions of four active experts in the field – but with reasonable differences in their background. Three experts, namely Luiz Da Silva, Samir Perlaza and Sergey Andreev answered seven questions about IoT, their effects in society and future of research in the topic. Another contribution is a position paper by Giancarlo Fortino and Claudio Savaglio covering fundamental challenges for the deployment of IoT.

Pedro H. J. Nardelli received the B.S. and M.Sc. degrees in electrical engineering from the University of Campinas, Brazil, in 2006 and 2008. In 2013 he received his doctoral degree from University of Oulu, Finland, and University of Campinas following a dual-degree agreement. Nowadays he is assistant professor in IoT in energy systems (tenure track) at the Laboratory of Control Engineering and Digital Systems, School of Energy Systems, LUT University, Finland, as well as adjunct professor (docent) in information processing and communication strategies for energy systems at Centre for Wireless Communications, University of Oulu.

Interview with Prof. Luiz A. da Silva

Trinity College Dublin
Email: dasilva@tcd.ie

Q1: Internet of Things (IoT) is a term that covers many technologies across the different network layers, from physical layer to application layer. How do you define Internet of Things in relation to the "past" technologies (e.g. sensor networks, embedded systems) and the state-of-the-art?

A1: To me, what is different about IoT is the focus on services. Of course, IoT builds on a lot of progress made in sensor networks (in areas like connectivity and energy efficient communications and protocols) and in embedded systems. The Internet of Things brings all of those technologies together with a clear focus on the service that is being provided to the end user, to a vertical, or to a community.

Q2: In the Communications Society community, IoT is usually associated to machine-to-machine or, more broadly, machine-type communications. These introduce research challenges involving a more holistic design by considering the particularities of the data sources and final applications in lower layers. For example, short-message communication, ultra-reliability, low latency, massive connectivity seem mainstream now. In these new scenarios, traditional assumptions that have simplified mathematical analysis (e.g. ergodicity, long messages, and application-agnostic traffic models) are becoming obsolete. In your opinion, what are the most promising mathematical theories and computational methods to carry out research in Communications Theory in the upcoming 10 years?

A2: Different types of IoT introduce very different technical challenges, all of which disrupt the traditional ways in which we have developed communications systems, so exciting times ahead for researchers. The kind of IoT where a very large number of devices produces infrequent traffic challenges some basic assumptions of orthogonal multiple access, and we are starting to see the communications community come up with clever alternatives for that. The kind of IoT where ultra reliability and low latency are key requires flexible migration of network functions, appropriate MAC mechanisms, and in my opinion even new metrics for network reliability. So there is room for innovation across the protocol stack, from communications theory to network resource management. There is also emerging interest in nano networks, and what some are calling the Internet of Nano Bio Things. One of my colleagues, Sasi Balasubramanian, at Waterford Institute of Technology, is making great advances in molecular communications, with applications from DNA storage to inhibiting the formation of bacterial film. In this area, communications are often not through electromagnetic waves and the potential for innovation in communication theory there is huge.

Q3: 5G is becoming a reality and research in 6G is starting. In both, IoT is viewed as an enabler of different classes of applications (the so-called "verticals"). Do you think this way of conceptualizing the functional role of IoT is suitable? What is the impact in the research community in systematizing the analysis in "verticals"?

A3: IoT services can vary widely, from low data rate as a result of infrequent monitoring, to high data rate for video surveillance, for example; and from entertainment and largely best-effort to mission critical and highly demanding of reliability and latency guarantees. Conceptualizing this in terms of verticals can help understand and define the challenges that arise under each of these different cases. The challenges of dealing with massive numbers of IoT devices that are only intermittently active are very different from those of dealing with a few IoT devices with stringent dependability of latency requirements, and the verticals are one way to tease out these diverse requirements.

Even thinking more broadly than IoT, one of the big transitions now is towards networks that need to deliver dependability, in addition to the traditional objectives of coverage and capacity. To really define what dependability means in terms of communications and networking services requires that we better understand the verticals that will be the ultimate customers of those services.

Q4: A big thing in IoT is data ownership, and the ethical and legal issues this implies. Probably the core question is: Is data a commodity to be traded in markets? If yes, who should own the data? If not, how should data possession and usage be governed? Although there is not a single clear-cut answer, it would be interesting to learn your views about this topic, which is becoming more and more relevant due to the current growth of Artificial Intelligence, Machine Learning, Deep Learning, blockchain and the new business models they are enabling. In this case, what your opinion about ownership of data that is acquired, processed and disseminated via IoT-enabled devices and networks?

A4: Public policy often moves much more slowly than the technology, and I think we are still in the very early stages of figuring out all the implications of the data we can now collect with IoT. The privacy issues are enormously important, so informed consent is critical: citizens need to be aware of what data about them is collected, for what purposes it is being used, whether it is being shared, what measures are being taken to protect it, etc. Community groups should also be involved in the design of these services, especially in the smart communities space, from the start. My group was recently started to engage with researchers in Maynooth University with expertise in urban ethics, to investigate some of these issues.

One of the PhD students in CONNECT, Fiona McDermott, is particularly interested in questions surrounding data governance. She brought to my attention the city government practices and open source policies being pioneered in Barcelona, viewing the data collected by the city, including IoT data, as a public utility. Another example is New York, where the mayor's office for technology has also developed guidelines for privacy standards around the deployment of IoT devices that use city assets of are deployed in public spaces.

Q5: Back to research, could you suggest one work in IoT that you find essential to anyone interested in the topic?

A5: Too many to mention, depending on what aspect of IoT you are interested in, but one good place to start is John Stankovic's 'Research Directions for the Internet of Things', published in the IEEE IoT Journal.

Q6: What is your most important contribution in the topic?

A6: The research centre that I direct in Ireland, CONNECT, has deployed what to the best of my knowledge is the only IoT testbed with countrywide coverage designed, deployed, and operated by an academic research group. Our testbed is called Pervasive Nation, and uses LoRAWAN to provide coverage to the entire Republic of Ireland. It has been extremely useful in testing out new ideas and building collaborations with Industry to actually deploy IoT services. I would also highlight some recent work of one of my PhD students, Jernej Hribar, who is using deep learning to develop mechanisms that increase the lifetime of the network by intelligently taking advantage of correlation between information collected by multiple IoT devices, and inspired by Age of Information ideas. This is part of a broader collaboration we currently have with Tsinghua University.

Q7: What are your own short-term and long-term plans in relation to research in IoT?

A7: One of the research themes in our research centre, CONNECT, is Sustainable IoT, viewed from the point of view of economic, environmental, and social sustainability. We envision the co-design of network and device technologies, so that we can address issues that span from new energy harvesting for devices to new communication protocols to privacy and security concerns. This is part of our long-term plan, and will be done in collaboration with my colleagues at Tyndall National Institute and University College Cork.

In the shorter term, we are a partner in the Horizon 2020 project ORCA (Orchestration and Reconfiguration Control Architecture). There, we are looking at virtualization and end-to-end network slicing to support a number of services, and one of these services is mission-critical IoT with strict latency requirements.

Luiz A. DaSilva holds the personal chair of Telecommunications at Trinity College, where he is the Director of CONNECT, a telecommunications and networks research centre funded by the Science Foundation Ireland. Prior to joining TCD, Prof DaSilva was a tenured professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech. His research focuses on distributed and adaptive resource management in wireless networks, and in particular radio resource sharing and the application of game theory to wireless networks. Prof DaSilva is a principal investigator on research projects funded by the National Science Foundation, the Science Foundation Ireland, and the European Commission under Horizon 2020. Prof DaSilva is an IEEE Communications Society Distinguished Lecturer, and a Fellow of Trinity College Dublin. He is also a Fellow of the IEEE, for contributions to cognitive networks and to resource management in wireless networks.

Interview with Dr. Samir Perlaza

INRIA, France
Email: samir.perlaza@inria.fr

Q1: Internet of Things (IoT) is a term that covers many technologies across the different network layers, from physical layer to application layer. How do you define Internet of Things in relation to the "past" technologies (e.g. sensor networks, embedded systems) and the state-of-the-art?

A1: The use of the term "Internet of Things" is relatively new. Note that Internet was born in silico during the sixties and the first "thing" connected to the Internet was a Coca-Cola vending machine at Carnegie Melon University, back in 1982. From the very first "thing" connected to internet up to today, we have seen absolutely everything we might have imagined. I believe that soon every single electronic device would be a "connected thing" and thus, compatible with at least one radio access network to be part of the Internet. But, from my perspective, connecting to the Internet is just a first step. The next step is how these devices connect to us. In the near future, I see the IoT leading us to work towards "things" that would jump the barriers between "humans and machines" to interact directly with our central nervous systems. People would prefer to feel the temperature at home rather than reading a number on a screen, in part because it reduces the energy put into thinking and taking decisions. We are just in the first part: connecting the devices around us.

From this perspective, the IoT is not more than the natural evolution of the Internet. The term has been coined essentially to highlight the fact that devices other than desktops and laptops can be connected to the internet and it has been proved useful for humankind. Hence, sensor networks, embedded systems, cellular networks, artificial intelligence are just enablers of this technological evolution. In the future, probably another name will pop up to point out that other things different to electronic devices, e.g., human brains, can also be connected to the Internet. I really hope to live long enough to see this happen.

Q2: In the Communications Society community, IoT is usually associated to machine-to-machine or, more broadly, machine-type communications. These introduce research challenges involving a more holistic design by considering the particularities of the data sources and final applications in lower layers. For example, short-message communication, ultra-reliability, low latency, massive connectivity seem mainstream now. In these new scenarios, traditional assumptions that have simplified mathematical analysis (e.g. ergodicity, long messages, and application-agnostic traffic models) are becoming obsolete. In your opinion, what are the most promising mathematical theories and computational methods to carry out research in Communications Theory in the upcoming 10 years?

A2: As a theoretician, I can highlight a number of (applied) mathematical problems that exhibit an astonishing simplicity but for which we ignore their solutions. Consider for instance the simplest abstraction of an IoT network, i.e., an interference channel (IC). This canonical model, introduced by Claude Shannon in 1961, consists in two point-to-point links that are subject to mutual interference. For this model, determining the set of all possible information rates that can be simultaneously achieved by both transmitter-receiver pairs (a.k.a. capacity region) is an open problem. Despite its simplicity, the capacity region of this channel is not known today, and only approximations are available in the block length asymptotic regime. A more realistic IoT model of two point-to-point links subject to mutual interference would be to consider the IC under the assumption that the communication must last a finite number of channel uses (latency constraint) and the decoding error probabilities at each receiver should not exceed certain thresholds (reliability constraints). Nonetheless, such a model is still not well understood and very little is known about it. Therefore, if the fundamental limits on the information transmission rates of a two transmitter-receiver pairs is an open problem, what can we say about the fundamental limits of the millions and millions of devices that are part of the IoT? From a theoretical point of view, we do not have the answer, despite the fact that the problem is by now well formulated. The crucial point in this regard is essentially that the mathematical tools we have been manipulating for studying these models do not allow us to progress any further. We probably need to improve the existing tools, use different tools or invent new ones.

From the real-system implementation side, we observe the need of highly reliable and low latency systems, but from the theoretical side, we ignore the fundamental limits and trade-offs among these constraints to formally guide the analysis, even in the simplest case of the two-user IC. From this standpoint, the development of communications systems is being let to blindly evolve. That is, technology is evolving and we are capable of building more and more performing communication systems, but we are unable to determine whether or not this improves are close or far from optimal operating points. This makes no difference with letting a person walk into a dark room! That person might get its way out, but by bumping into the obstacles.

I do believe that the most promising mathematical theories and computational methods to carry out research in Communications Theory in the upcoming 10 years are those leading to clarify our understanding of communications and data processing systems from a theoretical perspective. Of course, I am not suggesting at all stoping the design and construction of future systems, e.g., 6G or beyond, until we fully develop the theory needed to understand them. My suggestion is that in order to make progress faster, more effort should be put to develop the mathematical theories that lead us to a principled design of such systems.

Note that the mathematical theory introduced by Shannon, which shaped the digital world we know today, was introduced taking into account assumptions that do not longer hold. For instance, Shannon tacitly assumes that networks are centralized and each network component can be told exactly what to do by a central controller. During a long time, this model actually hold as cellular systems until 4G can be considered centralized systems. Nonetheless, today the assumption of centralized networks is not longer valid. Networks are made of devices whose autonomous behavior is led by their individual parameter configurations.

Another assumption was that transmitters and receivers were granted with vast amounts of energy, whereas today, energy availability is one of the most pressing challenges in IoT due to the need of batteries in most of electronic devices. From this perspective, the evolution of the IoT is left over without a mathematical background to that guarantees principled design. I do believe that the most impactful research that can be carried out today aims at developing the mathematical foundations of communications systems in the age of big data, limited energy, low-latency reliable communications, decentralized networks and humans that are ready to start delegating most of their daily-life decisions to the artificial intelligence.

Q3: 5G is becoming a reality and research in 6G is starting. In both, IoT is viewed as an enabler of different classes of applications (the so-called "verticals"). Do you think this way of conceptualizing the functional role of IoT is suitable? What is the impact in the research community in systematizing the analysis in "verticals"?

A3: I am interested in the theoretical foundations of communications. The classification of the applications into verticals appears to me adequate but unrelated to the theoretical foundations of communications. I am sure that this would be subject to changes in a few years to integrate other possible applications that today are not yet into the main stream, e.g., cyber-biological systems.

Q4: A big thing in IoT is data ownership, and the ethical and legal issues this implies. Probably the core question is: Is data a commodity to be traded in markets? If yes, who should own the data? If not, how should data possession and usage be governed? Although there is not a single clear-cut answer, it would be interesting to learn your views about this topic, which is becoming more and more relevant due to the current growth of Artificial Intelligence, Machine Learning, Deep Learning, blockchain and the new business models they are enabling. In this case, what your opinion about ownership of data that is acquired, processed and disseminated via IoT-enabled devices and networks?

A4: If we understand the IoT as "connecting things to the Internet", we are probably missing half of the landscape. Beyond a technological trend, the IoT is also a social phenomenon dictating the behaviors of humans and the way they interact with each other. The IoT, at the same time that it fosters comfort and economical progress, also unlocks several threads for the privacy of individuals; safety and security of large infrastructures; the health of the economy; and the stability of democracy. From this perspective, all the societal references, the Law and the education systems must keep the pace of the transformations induced by the IoT. The European Union and in particular, France, has taken very seriously this matter. At INRIA, an interdisciplinary group (PRIVATICS) has been formed in order to provide guidance to the government and the European Parliament in all dimensions of privacy, including legal, ethical and social dimensions. Within this context, data ownership is one of the essential topics to be discussed in the light of some fundamental principles: equality, privacy, dignity, autonomy and free will. Of course, there are some trade-offs to deal with in the sense that individual privacy cannot prevail over the public security, cyber-physical security, and stability of democracy.

I personally believe that the progress of most technologies related to algorithmic decision systems depends upon the exploitation of personal data of millions of individuals and social phenomena. Interestingly, this massive recollection of data does not necessarily imply a dangerous threat to the privacy of individuals, independently of who possess the ownership of data. If at the same time that data is collected, an effort is made to anonymize it, I do believe that the data can be still useful for the purposes of algorithmic decision making at the same time that the privacy of the individual is protected. In this regard, information theory has a lot to say and in the last years a lot of progress has been made in the development of the mathematical foundations of privacy. Unfortunately, in this particular case, technology seems to be ahead of the theoretical progress.

Q5: Back to research, could you suggest one work in IoT that you find essential to anyone interested in the topic?

A5: In information theory, the term IoT is rarely used because the formulation of the mathematical problems arising from it can be described, up to some additional considerations, in terms of canonical models that date back to the early sixties, e.g, the interference channel, the multiple access channel, the broadcast channel, among others. Some of these additional considerations have been mentioned above, essentially, communications are decentralized and take place during a short time with a small energy budget. This contrasts with the block-length asymptotic theory introduced by Shannon in 1948. From this perspective, I do believe that the non-asymptotic analysis of communications systems is essential for everyone interested in IoT. By these days, the literature on this topic is rather abundant and actively developing. Another, piece of work that I believe is of paramount importance is the consideration of limited energy budgets. This is particularly relevant due to the fact that the "things" connected to the Internet are often equipped with batteries, and thus, quite limited in terms of energetic autonomy. This topic is less studied and only a few groups around the world are studying this, probably due to its mathematical difficulty. In general, all these works together are the first steps towards the characterization of fundamental limits of IoT. Using these results, technological developments could be compared to a benchmark for determining its optimality. Only by knowing the fundamental limits, technological advances can be judged as optimal or suboptimal.

On another note, IoT is far from being just a technological trend. It is by now sufficiently developed to acknowledge that it dramatically influences our lives and those of our descendants. From this perspective, the study of IoT from the standpoint of social sciences, psychology, law and education is also fundamental.

Q6: What is your most important contribution in the topic?

A6: My contributions to the IoT are essentially on the analysis of its fundamental limits. On one hand, together with my students, we have studied this canonical multi-user channel we mentioned above, the two-user Gaussian interference channel (G-IC), which is indeed the building block of the IoT. Our work consists in an approximation of the information capacity region, that is, the set of all possible information rate pairs that can be simultaneously achieved. The importance of this work is that noisy channel-output feedback was considered from the receivers to the corresponding transmitters. From this standpoint, this result generalizes all the previous approximations of the capacity region of the G-IC without feedback and perfect output feedback. Our work revealed that previous studies on feedback have been too optimistic and when noise is present in the feedback links, the enlargement of the capacity region is far from what was initially claimed in both centralized and decentralized networks.

On the other hand, we have been among the first research groups to obtain the fundamental limits of simultaneous information and energy transmission (SIET), also known as simultaneously wireless information and power transfer (SWIPT), but I do not believe that this second appellation is correct, despite its vast utilisation. The simplest instances of the problem of SIET is a point-to-point communication between a transmitter and a receiver, in which aside to the information transmission task, the transmitter is engaged to transmitting energy to an energy harvester (EH) at a given energy rate (power). In multi-user channels, the problem of SIET the problem is similar: aside to the information transmission, the transmitters seek to guaranteeing a minimum energy transmission rate to an EH. The fundamental limits of SIET are known as the information-energy capacity region, which consists in all information transmission rates and energy transmission rates that can be simultaneously achieved. Our work has revealed the fundamental tradeoffs between information rates and energy rates in the block-length asymptotic regime. More specifically, for some canonical multiuser channels, mainly the G-IC and the Gaussian multiple access channel (G-MAC), we have fully characterized the scenarios in which both information transmission and energy transmission are competing tasks and thus, increasing the energy rate necessarily implies reducing the information transmission rate, and vice versa. More recently, we have studied SIET in point-to-point channels in the non-asymptotic regimes and some preliminary results have been obtained for simple memoryless binary channels. Nonetheless, this is still ongoing work.

Q7: What are your own short-term and long-term plans in relation to research in IoT?

A7: In the short term, I would like to complete the non-asymptotic analysis of SIET I mentioned earlier in order to obtain relevant fundamental limits for the IoT. I am very interested in studying multi-user channels and more importantly to build prototypes of SIET systems in order to observe the performance of existing transmission schemes and compare them with the fundamental limits. This has never done before precisely because the fundamental limits are yet unknown.

In the long term, I do believe that one of the most difficult problems in the IoT is that of algorithmic decision systems. That is, problems in which autonomous systems obtain information about their environment and must take decisions that determine their behavior and the interactions with other autonomous systems, for instance humans. Decision making processes are too ramified to be amenable to an exact mathematical treatment. Indeed, recent advancements on machine learning highlight that while extraordinary performance has been achieved on several application domains, there is still a lack of explanatory and fundamental principles for hallmark machine learning techniques such as deep learning. In the case of humans the distinction between the available data and the available information in a decision making process is illuminating. Two humans provided with the same data might implement different information processing mechanisms and as a result make different choices due to the differences in their information sets. Similarly, two machines, built for performing the same task, might decide completely different depending on the data they have been trained upon. At the core of this problem lies the difficulty of characterizing the amount of information and the value of that information in the decision making process. My goal is to address that shortcoming by developing an information-theoretic framework for this paradigm.

Samir Perlaza is an INRIA chargé de recherche and a visiting research scholar at the Department of Electrical Engineering of Princeton University, Princeton, (NJ, USA). He received the M.Sc. and Ph.D. degrees from École Nationale Supérieure des Télécommunications (Telecom ParisTech), Paris, France, in 2008 and 2011, respectively. Previously, from 2008 to 2011, he was a Research Engineer at France Télécom - Orange Labs (Paris, France). He has held long-term academic appointments at the Alcatel-Lucent Chair in Flexible Radio at Supélec (Gif-sur-Yvette, France); at Princeton

University (Princeton, NJ) and at the University of Houston (Houston, TX). He currently serves as an Editor of the IEEE TRANSACTIONS ON COMMUNICATIONS. He has been awarded with an Alban Fellowship in 2006 and a Marie SkłodowskaCurie Fellowship in 2015, both by the European Commission. His research interests lie in the overlap of signal processing, information theory, game theory, cyber-physical systems and wireless communication systems,

Interview with Prof. Sergey Andreev

Tampere University, Finland
Email: sergey.andreev@tuni.fi

Q1: Internet of Things (IoT) is a term that covers many technologies across the different network layers, from physical layer to application layer. How do you define Internet of Things in relation to the "past" technologies (e.g. sensor networks, embedded systems) and the state-of-the-art?

A1: The IoT is a long-standing paradigm and thus has multiple interpretations coming from rather different angles. In order to define it comprehensively, we need to consider the rich history of the IoT domain. Initially, legacy radio frequency identification (RFID) technology provided the devices with unique identifiers and wireless tracking capabilities. In the years that followed, wireless sensor network (WSN) solutions equipped dissimilar objects with the means to communicate under little-to-no human intervention. As the phenomenon of IoT embraced and further expanded the RFID and WSN realms, it presently constitutes a complex and integrated ecosystem where various 'things' are named, tracked, connected, and involved into meaningful autonomous interactions.

Q2: In the Communications Society community, IoT is usually associated to machine-to-machine or, more broadly, machine-type communications. These introduce research challenges involving a more holistic design by considering the particularities of the data sources and final applications in lower layers. For example, short-message communication, ultra-reliability, low latency, massive connectivity seem mainstream now. In these new scenarios, traditional assumptions that have simplified mathematical analysis (e.g. ergodicity, long messages, and application-agnostic traffic models) are becoming obsolete. In your opinion, what are the most promising mathematical theories and computational methods to carry out research in Communications Theory in the upcoming 10 years?

A2: Machine-type communication, which broadly features machine-to-machine and machine-to-human *modus operandis*, is the fabric of today's IoT applications. Being very different from the conventional human-type interactions, it requires a whole new set of dedicated mathematical and computational tools to capture the unique properties of machines. These need to cover the entire range of the prospective machine-specific features, from small and infrequent data to ultra-reliable and low latency streaming. While support for massive connectivity requires novel random access mechanisms, the stringent latency and reliability guarantees call for revisiting classical communication-theoretic principles. A promising premise here is grant-free access that allows for randomized and possibly non-orthogonal radio resource management without prior reservation, among many other methods that enable adequate reliability via diversity.

Q3: 5G is becoming a reality and research in 6G is starting. In both, IoT is viewed as an enabler of different classes of applications (the so-called "verticals"). Do you think this way of conceptualizing the functional role of IoT is suitable? What is the impact in the research community in systematizing the analysis in "verticals"?

A3: From the perspective of its vertical applications, today's IoT use cases are categorized as either massive or reliability-/latency-critical. Together with further evolution of mobile broadband, these three classes of services constitute the popular '5G triangle', where individual applications are positioned around its corners. However, going beyond 5G we may need to accommodate scenarios that are high-bandwidth and at the same time require highly reliable operation of a large number of devices. This is quite different from today's 5G vision where these features are facilitated by individual enabling technologies and will require new solutions to handle the three 5G angles simultaneously. Example use cases include mobile augmented and virtual reality, large fleets of autonomous vehicles, cooperating drone swarms, and collaborative moving robots.

Q4: A big thing in IoT is data ownership, and the ethical and legal issues this implies. Probably the core question is: Is data a commodity to be traded in markets? If yes, who should own the data? If not, how should data possession and usage be governed? Although there is not a single clear-cut answer, it would be interesting to learn your views about this topic, which is becoming more and more relevant due to the current growth of Artificial Intelligence, Machine Learning, Deep Learning, blockchain and the new business models they are enabling. In this case, what your opinion about ownership of data that is acquired, processed and disseminated via IoT-enabled devices and networks?

A4: Indeed, this question does not have a straightforward answer. For the sake of space, let us only touch upon the recent advances in artificial intelligence for wireless. Due to massive training data available in the increasingly capable IoT devices, wireless edge and fog infrastructures can be efficiently leveraged as an integrated communication—computation substrate for collaborative learning. However, numerous new challenges emerge in this context along the lines of making mobile communication more reliable, mitigating the imbalance between the involved IoT objects, incentivizing broader device participation, and protecting the privacy of personal datasets, among many others. This demands further research on enabling distributed artificial intelligence over wireless.

Q5: Back to research, could you suggest one work in IoT that you find essential to anyone interested in the topic?

A5: Continuing on the above question regarding new theories and methods for future IoT, the important rationale behind grant-free access has been offered by Petar Popovski *et al.* in “Wireless Access for Ultra-Reliable Low-Latency Communication: Principles and Building Blocks”. This contribution appeared in *IEEE Network* last year to offer the guiding system design principles that are instrumental to construct mission-critical applications. It reviews the problem at hand within the framework of information theory to coin efficient enablers for random access protocols and facilitate optimized signaling for ultra-reliable and low latency communication. An important conclusion of this work is that in latency-constrained access the traditional communication systems engineering approaches may need to be rethought from the perspective of redundancy by integrating various sources of diversity.

Q6: What is your most important contribution in the topic?

A6: One of our recent IoT-centric lines of research conceptualizes the vision of dense moving fog facilitated by increasingly denser geographical distribution of fog functionality, beyond the conventional cloud and edge computing paradigms. The key challenge here is that the more intelligent IoT objects, such as autonomous cars and drones, may move unpredictably and at high speeds. Fortunately, we confirm non-incremental benefits of the moving fog infrastructure for collaborative data processing in vehicular and airborne fog computing. These initial findings are documented in “Dense Moving Fog for Intelligent IoT: Key Challenges and Opportunities” that was published in *IEEE Communications Magazine* earlier this year.

Q7: What are your own short-term and long-term plans in relation to research in IoT?

A7: Going further and beyond collaborative communication, we explore the co-design of wireless connectivity with dynamic control of moving IoT formations, such as autonomous fleets of aerial and terrestrial vehicles. This includes joint communication, positioning, computation, storage, navigation, and security, which should entail the development of novel evaluation tools that incorporate integrated network and robot emulation under realistic mobility. In the long run, we envision the possibility to employ smarter IoT devices like networked cars and drones for strategic densification of beyond-5G wireless layouts. Such moving networks may become a powerful operator asset to deliver on-demand capacity, content, and coverage during unpredictable and temporary events, while static deployment can be made more affordable by scaling it down for the median loading.

Sergey Andreev [SM'17] (sergey.andreev@tuni.fi) is an assistant professor of electrical engineering at Tampere University, Finland. Since 2018, he has also been a Visiting Senior Research Fellow with the Centre for Telecommunications Research, King's College London, UK. He received his Ph.D. (2012) from TUT as well as his Specialist (2006) and Cand.Sc. (2009) degrees from SUAI. He serves as editor for IEEE Wireless Communications Letters (2016-) and as series editor of the IoT Series (2018-) for IEEE Communications Magazine. He (co-)authored more than 150 published research works on intelligent IoT, mobile communications, and heterogeneous networking.

Position Paper: ACOSO-Meth: a full-fledged methodology for the agent-based Internet of Things

Prof. Giancarlo Fortino, Dr. Claudio Savaglio
Department of Informatics, Modeling, Electronics and Systems (DIMES), Università della Calabria, Italy

I. Abstract

The development of Internet of Things (IoT) systems is a complex task featured by manifold issues (large scale deployment, heterogeneity, cyber-physicality, interoperability, etc.). Therefore, an adequate and multi-disciplinary methodological approach is required to satisfy these requirements while reducing the probability of failure and time-to-market. Along this line, Artificial Intelligence (AI) is likely to be the best source of resources – such as algorithms, technologies, methodologies, and paradigms – enabling the development of next-generation Smart Objects (SOs) and IoT systems. In particular, the Agent-based computing (ABC) paradigm has been effectively exploited for modeling, programming and simulating IoT systems. This paper presents ACOSO-Meth (Agent-based Cooperating Smart Objects Methodology), the first agent-based methodology that specifically and seamlessly supports the main phases of engineering of IoT ecosystems and related services. The effectiveness and efficiency of the proposed approach have been assessed through use cases related to different application scenarios.

II. Introduction

The IoT is a dynamic, decentralized and unstructured ecosystem, where billions of SOs (i.e. everyday objects reinforced with sensing, computation, communication and actuation capabilities) are connected on a global scale and provide pervasive cyber-physical services to human users or other machines [1]. The development of interoperable and intelligent IoT systems, however, represents a complex task with many requirements and issues. In this context, a systematic and multidisciplinary development approach is necessary to face the cyber-physical nature of the IoT and to guarantee an adequate level of smartness. AI considers the theory and implementation of computing systems that display intelligence by analyzing the environment and acting - with some degree of autonomy - to achieve specific goals.

A plethora of scientific and industrial fields (e.g. security, space, transport, health, Industry 4.0) can benefit from AI, which currently integrates several areas, like machine learning, computer vision, cognitive and autonomous systems, just to mention a few. Among these AI paradigms, ABC is widely recognized as a comprehensive, effective support for the development of decentralized, dynamic, and cooperating IoT systems, particularly in conjunction with other complementary paradigms, e.g. cloud, edge, cognitive and autonomic computing, business process management [2]. The main features of the agent (autonomy, social capacity, responsiveness, proactivity and mobility), in fact, perfectly match with the generic and specific requirements of a SO and, therefore, a multi-agent system is perhaps the most natural way to approach the development of complex, dynamic, context-aware and autonomous IoT systems.

As matter of facts, the ABC has been exploited for modeling, programming and simulating IoT applications and systems, and thus systematically driving and speeding-up their development. Indeed, better than other computing paradigms (object-oriented, service-oriented, component-oriented) and both at things and at system levels, ABC allows modeling IoT systems at different degrees of details, facilitating autonomicity, distributed intelligence and

- technical interoperability, through shared resource/communication interfaces;
- syntactical interoperability, through a shared message format, because ACL is adopted across FIPA standard obeying platforms for message envelope, while XML and JSON are used for message content;
- semantic interoperability, through shared ontology and knowledge representation.

In addition, the joint exploitation of agent-oriented modeling and network-based simulation allows understanding overall dynamics, estimating performance, and validating models, protocols and algorithms featuring under-development IoT systems [1].

III. ACOSO-Meth

The development of IoT ecosystems is a complex and complex process. Several methodologies have been proposed over the years, but, as shown in Table 1, none of them systematically supports the main engineering phases (analysis, design, implementation) or provide an associated tool for the development of IoT systems and services.

Table 1 Comparison of Agent-Based Methodology (Y = totally supported, P = partially supported, Blank = not supported)

	Supported development phase (**A**nalysis, **D**esign, **S**imulation, **I**mplementation) and associated **T**ool				
Surveyed Work	**A**	**D**	**I**	**S**	**T**
Zambonelli, 2016 [3]	P	Y			
Manate, 2014 [4]	P	Y			
Spanoudakis, 2015 [5]	P	P			P
Cini., 2017 [6]	Y	Y			
ACOSO-Meth	Y	Y	Y	Y	Y

With the aim of improving the state of the art, ACOSOMeth has been defined (Agent-based COoperating Smart Objects Methodology) [7], the first methodology that fully supports the development of IoT systems of different complexity and degrees of intelligence. ACOSO-Meth adopts the agent paradigm and a set of metamodels placed at different levels of abstraction, which are specialized and detailed from the analysis phase to the implementation phase. The proposed approach is based on software agents since these represent the ideal solution to model the SOs and provide them with an adequate level of intelligence.

The agents, in fact, are able to implement the principles of Autonomic and Cognitive Computing within the SOs, autonomously manage their resources (sensors, actuators, knowledge base), and favor technical, syntactic and semantic interoperability between SOs with databases, communication interfaces and heterogeneous ontologies. As shown in Fig.1, ACOSO-Meth foresees:

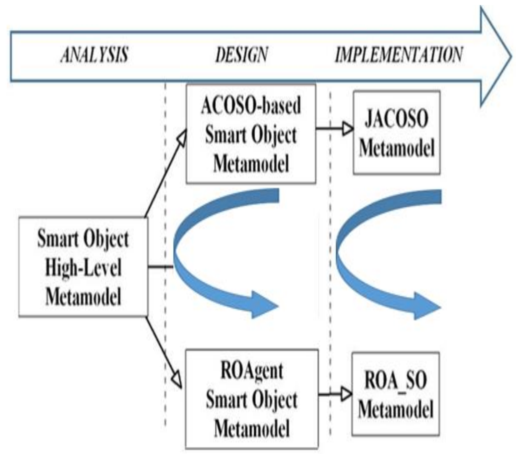

Figure 1 ACOSO-Meth development phases

- in the analysis phase, the use of a high level SO metamodel, called SO High-Level Metamodel, compliant with the main architectural standards/IoT domain models (IEEE P2430, AIOTI, IoT-A) and describes in general terms (non-)functional aspects of the SO;
- in the design phase, the use of an SO metamodel named ACOSO Metamodel which, to illustrate the functional components of the system, their relationships and interactions, specializes the metamodel of the analysis phase using the agent paradigm. In order to verify the design choices, the ACOSO Metamodel can be mapped in an OMNeT ++ model and then simulated, so as to obtain a preliminary evaluation of the designed system [1] also in light of network level problems (e.g. wireless channel modeling, interference) otherwise difficult to model; finally,
- in the implementation phase, the use of an OS metamodel named JACOSO (JADE-based ACOSO) Metamodel which specializes the ACOSO Metamodel with respect to a particular implementation based on the JADE agent platform (used in different IoT frameworks).

Each phase introduces new features and a greater degree of detail, while maintaining strong relationships with the other metamodels: this allows the translation of the analysis models into design models to platform independent agents which, in turn, can be refined into agent implementation models but platform-dependent. To provide practical and not just theoretical support, the methodology is supported by ACOSO, a middleware that simplifies the development, management and implementation of cooperative SOs. ACOSO provides an agent-oriented programming model to implement, in any context, IoT applications that require distributed computing, proactivity, knowledge management and interaction between SOs, sensors and actuators.

IV. ACOSO-Meth use cases

ACOSO-Meth was used (from the high-level analysis phase of the system to the JACOSO-based implementation) to develop the case study of a complex SO, called SmartUniCal [7]. The SmartUniCal, which in turn includes heterogeneous SOs of different scales, was developed in a real scenario (the University of Calabria) and provides cyber-physical services related to the structural, environmental and well-being monitoring of people. The application of the agent paradigm has allowed the development of intelligent SOs based on different operating systems (Android, TinyOs, Windows) and cooperating through different communication protocols (Wi-Fi, Bluetooth, IEEE 802.15.4). The systematic application of ACOSO-Meth has greatly facilitated and speeded up all the development phases of the SmartUniCal:

(i) the analysis metamodel supported the high-level analysis of the main features and functionalities of the SmartUniCal;

(ii) the agent-oriented metamodel at the design level has provided adequate effectiveness to meet the fundamental requirements both at the system level and at the individual device level;

(iii) the JADE-based implementation metamodel allowed rapid and efficient prototyping of SmartUniCal ecosystem.

In parallel, ACOSO-Meth allows the re-engineering of existing IoT systems, improving maintainability, reusability and extensibility (features that cannot be underestimated in the constantly evolving IoT scenario with ever new devices and services). In this direction, [8] presents (i) the integration in ACOSO-Meth of the ROA (resource-oriented agent) framework, which complies with the IETF Constrained RESTful Environment (CoRE) specifications and allows the development of agent applications on devices with limited hw / sw resources; and (ii) the re-engineering of a smart mobility application through ACOSO-Meth.

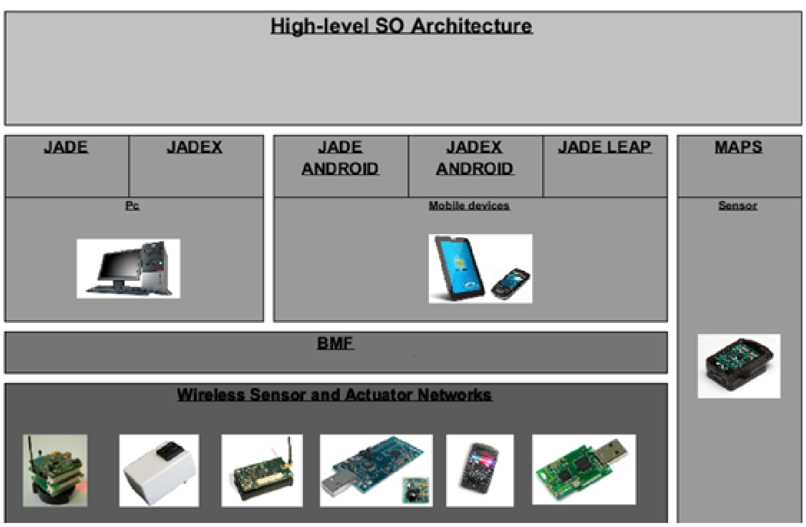

Figure 2 ACOSO middleware Architecture

V. Towards Opportunistic Services with

Collective Intelligence

Services fundamentally contributed to the evolution of the Internet and, likewise, promise to play a crucial role within the IoT ecosystems. In fact, SOs, conventional computer systems and people (hereinafter IoT Entities), supported by pervasive and global connectivity, will take part in innovative and advanced cyber-physical services (hereinafter, IoT services), which will revolutionize every application scenario.

Following the analysis of the state of the art on IoT services, ACOSO-Meth has been extended [9] to propose an innovative approach that supports the development of collective IoT services. In this approach, an IoT service is configured as an interface to access the cyber-physical functionalities of the various IoT Entities located in a specific physical space (IoT Environment) and bound to a particular context (IoT Context). In particular, the proposed IoT service model is the first that explicitly considers "opportunistic" properties (crucial to capture the real potential of the IoT service but largely overlooked so far) of IoT services and enables collective intelligence. In fact, an IoT service can be described as the functional composition of simpler services that require self-adaptive and / or self-organized behavior, spatio-temporal coordination, and awareness of available resources. The development of an integrated framework to support formal verification, simulation and implementation of opportunistic and collective IoT services before their distribution represents the line of research currently pursued.

References

[1] Giancarlo Fortino, Raffaele Gravina, Wilma Russo, and Claudio Savaglio. Modeling and simulating internet-of-things systems: a hybrid agent-oriented approach. Computing in Science & Engineering, 19(5):68–76, 2017.

[2] Claudio Savaglio, Giancarlo, Fortino, Maria Ganzha, Marcin Paprzycki, Costin Bădică, Miriana Ivanović. Agent-based computing in the Internet of Things: a survey. In International Symposium on Intelligent and Distributed Computing, pages 307-320, 2017. Springer, Cham.

[3] Zambonelli, F. (2016). Towards a general software engineering methodology for the Internet of Things. arXiv preprint arXiv:1601.05569.

[4] B. Manate, F. Fortis and P. Moore, "Applying the Prometheus Methodology for an Internet of Things Architecture," 2014 IEEE/ACM 7th International Conference on Utility and Cloud Computing, London, 2014, pp. 435-442.

[5] Spanoudakis, N., & Moraitis, P. (2015). Engineering ambient intelligence systems using agent technology. IEEE Intelligent Systems, 30(3), 60-67.

[6] Cila, N., Smit, I., Giaccardi, E., & Kröse, B. (2017, May). Products as agents: metaphors for designing the products of the IoT age. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (pp. 448-459). ACM.

[7] Giancarlo Fortino, Wilma Russo, Claudio Savaglio, Weiming Shen, and Mengchu Zhou. Agent-oriented cooperative smart objects: From iot system design to implementation. IEEE Trans. on Systems, Man, and Cybernetics: Systems, (99):1–18, 2017.

[8] Claudio Savaglio, Teemu Leppänen, Wilma Russo, Jukka Riekki, and Giancarlo Fortino. Reengineering iot systems through acoso-meth: the IETF core based agent framework case study. In WOA, volume 2215 of CEUR Workshop Proc., pages 81–89, 2018.

[9] Roberto Casadei, Giancarlo Fortino, Danilo Pianini, Wilma Russo, Claudio Savaglio, and Mirko Viroli. Modelling and simulation of opportunistic iot services with aggregate computing. Future Generation Computer Systems, 91:252–262, 2019.

Giancarlo Fortino (SM12) is Full Professor of Computer Engineering at the Dept of Informatics, Modeling, Electronics, and Systems of the University of Calabria (Unical), Italy. He received a PhD in Computer Engineering from Unical in 2000.He is also guest professor at Wuhan University of Technology (China), high-end expert at HUST (China), and senior research fellow at the ICAR-CNR Institute. He is the director of the SPEME lab at Unical as well as co-chair of Joint labs on IoT established between Unical and WUT and SMU Chinese universities, respectively. His research interests include agent-based computing, wireless (body) sensor networks, and IoT. He is author of 400+ papers in int'l journals, conferences and books. He is (founding) series editor of IEEE Press Book Series on Human-Machine Systems and EiC of Springer Internet of Things series and AE of many int'l journals such as IEEE TAC,IEEE THMS,IEEE IoTJ, IEEE SJ, IEEE SMCM, Information Fusion, JNCA,EAAI, etc. He is cofounder and CEO of SenSysCal S.r.l., a Unical spinoff focused on innovative IoT systems. Fortino is currently member of the IEEE SMCS BoG and of the IEEE Press BoG, and chair of the IEEE SMCS Italian Chapter.

Claudio Savaglio received his B.S., M.S. and Ph.D. degrees in Computer Engineering respectively in 2010, 2013 and 2018 from the University of Calabria, where he is currently serving as PostDoctoral Researcher. In 2013 he was Visiting Researcher at University of Texas at Dallas (TX, U.S.A.), in 2016 at New Jersey Institute of Technology, (NJ, U.S.A.), and in 2017 at Universitat Politecnica de Valencia (Valencia, Spain). His research interests include the Internet of Things, Edge Computing, network simulation, and agent-oriented middleware and development methodologies.

TCCN Newsletter Editorial Board

TCCN NEWSLETTER DIRECTOR

Daniel Benevides da Costa
Federal University of Ceará, Sobral-CE, Brazil.

FEATURE TOPIC EDITORS

Walid Saad, Virginia Tech, USA
Pedro H. J. Nardelli, Lappeenranta University of Technology, Finland

TCCN Officers

CHAIR

Yue Gao
Queen Mary University of London
UK

VICE CHAIRS

Daniel Benevides da Costa
Federal University of Ceará
Brazil
(TCCN Vice-Chair Americas)

Lingyang Song
Peking University
China
(TCCN Vice-Chair Asia Pacific)

Oliver Holland
King's College London
UK
(TCCN Vice-Chair Europe/Africa)

SECRETARY

Lin Gao
Harbin Institute of Technology
China