



Semantic-Effectiveness Filtering and Control for Post-5G Wireless Connectivity

Petar Popovski^{1*}, Osvaldo Simeone², Federico Boccardi³, Deniz Gündüz⁴ and Onur Sahin⁵

Abstract | The traditional role of a communication engineer is to address the technical problem of transporting bits reliably over a noisy channel. With the emergence of 5G, and the availability of a variety of competing and coexisting wireless systems, wireless connectivity is becoming a commodity. This article argues that communication engineers in the post-5G era should extend the scope of their activity in terms of design objectives and constraints beyond connectivity to encompass the semantics of the transferred bits within the given applications and use cases. To provide a platform for semantic-aware connectivity solutions, this paper introduces the concept of a semantic-effectiveness (SE) plane as a core part of future communication architectures. The SE plane augments the protocol stack by providing standardized interfaces that enable information filtering and direct control of functionalities at all layers of the protocol stack. The advantages of the SE plane are described in the perspective of recent developments in 5G, and illustrated through a number of example applications. The introduction of a SE plane may help replacing the current “next-G paradigm” in wireless evolution with a framework based on continuous improvements and extensions of the systems and standards.

1 Introduction

1.1 Motivation

Since the early days of communication systems engineering, the complexity of the communication process has motivated a compartmentalization of the subject into separate disciplines. Shannon and Weaver famously identified three levels of problems within the broad subject of communication¹:

- (A) *Technical problem* The symbols conveying information should be reliably transmitted to the recipient.
- (B) *Semantic problem* The meaning conveyed by the transmitted symbols should accurately reflect the intentions of the sender.

- (C) *Effectiveness problem* The conduct or action of the system in response to communications should be effective in accomplishing a desired task.

We illustrate these three levels through two contemporary examples. In online shopping, the technical problem amounts to the establishment of a reliable end-to-end connection; the semantic problem to the definition of the most informative and appealing display of products on the web page; and the effectiveness problem relates to the rate of success in terms of sales figures or targeted advertising. In a remote monitoring application via an Internet-of-things (IoT) system, the technical problem is the transmission of collected data from the IoT sensors to a cloud processor;

¹ Department of Electronic Systems, Aalborg University, Aalborg, Denmark.

² Centre for Telecommunications Research, King's College London, London, UK.

³ University at Bristol, Bristol, UK.

⁴ Imperial College London, London, UK.

⁵ InterDigital, London, UK.

*petarp@es.aau.dk

the semantic problem is the extraction of knowledge from data through data analytics; and the effectiveness problem concerns the impact of the action taken by a human or by relevant actuators in response to the data collected from the sensors.

Traditionally, communication engineers have been solely concerned with the technical problem. This focus has been arguably one of the key reasons for the unprecedented success of the communication technology in the last decades. However, as communication technologies have matured into the status of a commodity, the bar-center of technology and business has migrated in the direction of the semantic and the effectiveness problems. The best indicator for this trend is the economic, as well as political clout, of the most valuable digital tech companies—Facebook, Apple, Amazon, Netflix, Google (FAANG)—that are built to address these problems rather than the technical problem. Confining their activities to the technical problem may cause today’s communication engineers to be stuck in what we may refer to as “Shannon’s trap”. The ongoing migration of communication engineers towards methods, or even job positions, related to machine learning (ML) and artificial intelligence (AI) provides a tangible evidence of this trend.

1.2 An Augmented Protocol Architecture

The key thesis of this article is that, while designing solutions for post-5G wireless connectivity, wireless communication engineers should expand their efforts beyond the technical problem to address the semantic and effectiveness problems. In light of the above trends, as illustrated in Fig. 1, we propose to move away from the current protocol architecture whereby the interfaces between users, sensors, and actuators carrying information relevant to semantic and effectiveness problems are limited to the application layer. In contrast, in the proposed architecture, the protocol stack is augmented with a plane that exposes Application Programming Interfaces (APIs) between, on the one end, users, sensors, and actuators, and, on the other end, all layers of the Radio Access Network (RAN) protocol stack. The APIs enable the extraction of information to be processed via data analytics tools and the direct control of functionalities at any layer. This in turns allows the standardized integration in the communication protocol of data-driven, i.e., ML/AI, methods at a level that is within the purview of the wireless communication engineer.

As we will discuss, a related effort has been recently made by the 3rd Generation Partnership

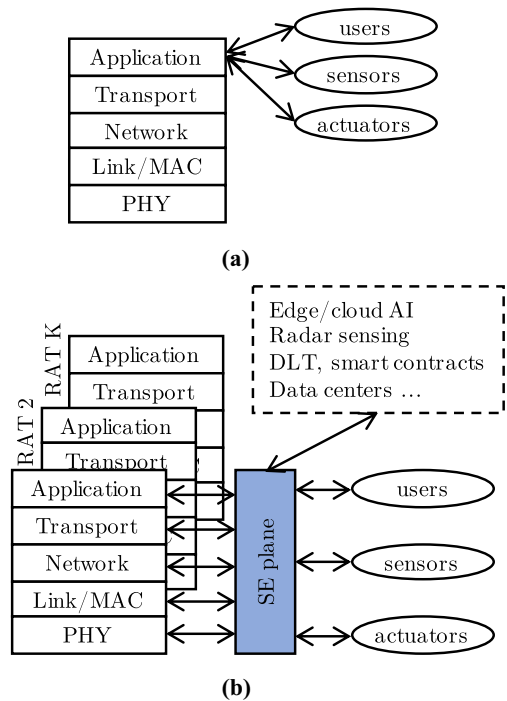


Figure 1: Interaction of the users, sensors, and actuators with the protocol stack. a 2G–5G cellular networks; b proposed architectural evolution in the post-5G wireless systems that includes the SE plane. The figure emphasizes the fact that the SE plane can generally interact and coordinate multiple Radio Access Technologies (RATs), rather than being dedicated to a single cellular interface.

Project (3GPP) at the level of the 5G core network. With Release 16, 3GPP has, in fact, introduced standardized interfaces between network functions within the control plane of the core network. These interfaces carry data on protocol operations and performance metrics that can be used for data analytics and resource management. On the one hand, the proposed architecture extends the idea of defining standardized interfaces to the RAN protocol stack, and, on the other, it allows not only for information extraction but also for direct control. We refer to this new plane—whose main desired properties will be described here along with some use cases—as semantic-effectiveness (SE) plane. The SE plane covers functionalities concerning both semantic and effectiveness problems since the two aspects are often intertwined in terms of both data collection and control.

Concepts that are closely related to the SE plane have appeared previously in different contexts, most notably the knowledge plane² and cognitive resource manager (CRM) from³.

Similar to our proposal, both works introduce the planes in a way that minimizes the positive advantages of layering. The knowledge plane² targets the broader internet architecture and harbors functionalities such as, for example, fault diagnosis and mitigation, automatic reconfiguration, or intrusion detection. Different from this, our work is focused on the wireless cellular systems and the specific wireless service requirements that have emerged in relation to 5G and beyond-5G communication. The CRM³ is closer to our work, as it is also chiefly focused on the wireless interfaces. The objective of CRM is to provide intelligent resource management by using information collected at the different layers of the protocol stack. While³ is mostly focused on radio resource management for a multitude of wireless interfaces, our work is to a large extent motivated by wireless connectivity in cyber-physical systems with sensors and actuators. In addition, we are deriving the need for semantic plane based on the technological development and knowledge created with the emergence of 5G wireless systems.

The rest of the paper is organized as follows. In Sect. 2, we provide context and discuss basic functionalities of the SE plane. Sect. 3 describes a number of exemplifying use cases. In Sect. 4, we place the proposal in the context of current standardization efforts within 3GPP. Finally, Sect. 5 concludes the paper by covering key implementation challenges and by offering an outlook.

2 SE Plane: Context and Overall Functionality

In this section, we first provide some context to further motivate the proposed augmented protocol stack illustrated in Fig. 1, and then we describe the main properties and functionalities of the proposed SE plane.

2.1 Connectivity Trends

To frame the proposal in the proper technological context, we briefly summarize here some relevant current trends in wireless connectivity. We will return to some of the main points in later sections of the paper.

- In the post-5G era, systems for wireless connectivity will be abundant and mature, shifting the performance bottleneck—and, with it, the focus of research—to layers and services that leverage physical-layer connectivity primitives.

- Recent advances in ML, also popularly referred to as AI, will enable the extraction of information by means of pattern recognition within complex data streams at all layers of the protocol stack (see, e.g., ⁴).
- There is an increased demand for wireless services with ultra-low latency and real-time interactive communications, for use cases such as Virtual Reality (VR), Augmented Reality (AR), and remotely controlled robots and drones. The need for ultra-low latency will increase the focus on predictive communications and networking based on ML techniques.
- As opposed to classical latency, in an increasing number of applications, a more relevant metric is the Age of Information (AoI), which captures the freshness of information by taking into account the dynamics, and hence semantics, of the information source ⁵.
- The architecture of 5G systems, and very likely of post-5G networks, is evolving away from mere softwarization and virtualization of well-defined Network Functions (NFs) towards a Service-Based Architecture (SBA). SBA will enable self-contained and reusable NFs to exchange data via efficient Service-Based Interfaces (SBIs), hence, enabling data-driven interactions among NFs ⁶.
- With the convergence of heterogeneous data- and task-oriented services over the cellular infrastructure, data streams will contain a large fraction of semantic or effectiveness redundancy, that is, of data that is delivered to the application layer, but ends up not being relevant or useful. A notable example is the transmission of data for the training of ML models whose source cannot be trusted—a semantic problem—or whose inclusion does not lead to significant updates in the current model—an effectiveness problem.
- 5G networks will be characterized by an increased protocol overhead to ensure security, privacy, provenance, as well as trust through a Distributed Ledger Technology (DLT)⁷. This protocol overhead will contain potentially significant semantic and/or effectiveness redundancy. For example, the provenance of data that is no longer relevant for an application of interest need not be validated using, e.g., access to DLT.

2.2 SE Plane: Basic Functionalities

To facilitate integrated solutions to the technical, semantic, and effectiveness problems, we propose

the introduction of a SE plane in the communication protocol architecture as illustrated in Fig. 1. As we will detail, the SE plane provides standardized APIs that enable information filtering and direct control at all layers of the protocol stack based on semantic and effectiveness metrics.

Through its standardized interfaces to all layers of the protocol stack, the SE plane makes it possible to

- filter information packets that are irrelevant, which is a semantic problem, or not useful, which is an effectiveness problem, by cross-checking against the state of the ML/AI models or trusted DLT tools such as blockchains and smart-contract ledgers. The effect of SE filtering is illustrated on Fig. 2: by removing SE redundancy, SE filtering aims at reducing overhead at all layers of the protocol stack;
- control the operation of the communication protocol by providing a direct interface between communication functionalities, at one end, and actuators and sensors, at the other. SE control at all layers of the protocol stack can enable a more efficient use of communication and computing resources, and it opens the way for a principled and controlled introduction of ML/AI tools and techniques into the design of the wireless communication protocol. As discussed in the examples in the next section, SE control can enhance both semantic and effectiveness performance metrics.

As we review in Sect. 4, the current standardization efforts have a limited scope in terms of the type of information that can be accessed by ML/AI tools. In particular, the standardized interface only concerns the core network, and it is one-directional, from network functions (NFs) in the core network to cognitive functionalities based on ML/AI. As a result, innovative “AI” solutions relying on network analytics are constrained to take place solely within the core networks and within the network elements of a given operator. The latter constraint stems from the absence of standardized control interfaces that are directed from the SE plane to the NFs. The SE plane outlined here addresses both limitations by extending the scope of the standardized APIs to the RAN and by allowing also for direct control of communication functionalities.

At a high level, the functionalities introduced for the SE plane pave the way for the evolution of




Figure 2: The role of SE filtering in decreasing both application and protocol overheads.

communication systems according to the following principles and objectives:

1. to provide new use cases for users and businesses;
2. to embrace ML with the aim of transforming communication networks into adaptive and predictive distributed systems;
3. to open the value chain in the direction of “open radio” with re-defined standardization models and releases.
4. to ensure security, privacy and trust by allowing a better control of the data and its use.

In the next section, we cover a number of specific examples and use cases to illustrate applications of the SE plane.

3 Implementing the SE Plane

While the specification of the detailed functional blocks and operation of the SE plane is an elaborate task that is outside of the scope for this paper, this section provides several instances that illustrate the functionality and the benefits from the SE plane beyond the current state of the art.

3.1 SE Filtering for Immersive and Tactile Applications

Consider an augmented reality (AR) application, in which collocated users have both physical and digital interactions, e.g., within social networks or video games. Broadband and low-latency traffic packets are required to be carried over wireless connections, since videos and audio signals are complemented by haptic and other sensory signals for an immersive experience. Low-latency transmissions require almost instant access to radio resources. This can be ensured by reserving dedicated resources for low-latency

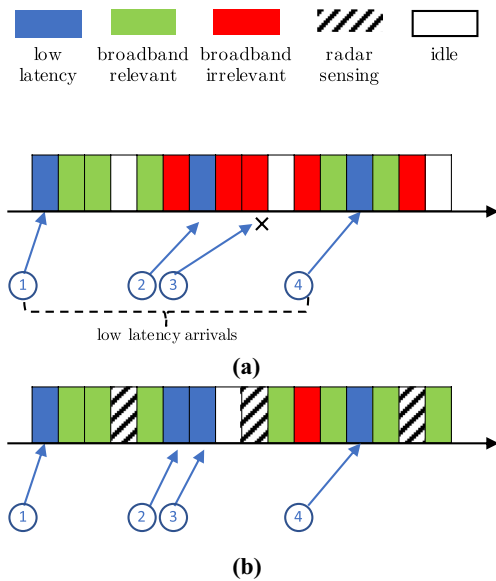


Figure 3: Illustration of multiplexing of low-latency wireless, broadband, and radio sensing in a MAC frame. **a** Low-latency resources are reserved periodically every three slots, and, if a low-latency packet needs to wait more than three slots, it is dropped; **b** SE filtering allows broadband packets carrying irrelevant information to be dropped, freeing resources for low-latency packets or other applications such as radar sensing.

transmissions. This is illustrated through the simple Medium Access Control (MAC) frame in Fig. 3a, where every third slot is reserved for low latency transmission. However, as also seen from the figure, if the low-latency traffic requests are random and intermittent, a portion of these reserved resources would be wasted (see, e.g., ⁸). Conversely, upon too many arrivals, some of them may not be served as the latency deadline is exceeded before a dedicated resource becomes available, such as for the arrival #3 on Fig. 3a. To mitigate this problem and enhance the spectral efficiency, the transmission of both types of traffic can be jointly designed and adapted to current traffic conditions by leveraging SE plane functionalities. For example, semantic redundancy may arise from the transmission of multimedia signals that are never used since they carry details of the scene that the user does not interact with. SE plane functions can detect and remove this redundancy, hence freeing resources for low-latency traffic. This is illustrated in Fig. 3, where the packets denoted as “broadband irrelevant” in the top part are filtered by the SE plane in the bottom figure, enabling low-latency access for all packets.

3.2 Semantic Control for Integrated Communication and Radar Sensing

With the growing presence of mmWave and the advent of THz communications, there is an increasing focus on approaches to repurpose communication resources for radar sensing, or more generally for sub-millimeter-wave sensing and processing ⁹. A key problem is how to enable the control of radio resources for sensing without affecting the operation of the communication protocol. A standardized interface between the SE plane and the lower layers of the protocol stack would address this issue, making it possible to seamlessly integrate solutions by different vendors of both radar sensing and communication systems.

As an example, the MAC frame could include some blank slots, with suitable time and bandwidth guard bands, that may be directly controlled by the sensing applications run by the SE plane. Alternatively, the physical layer may include new modulation and coding schemes that offer degrees of freedom on the transmitted waveform that can be used for sensing applications. This idea generalizes the principle of protocol coding introduced in ¹⁰, whereby degrees of freedom in the transmission protocol, e.g., the ordering of packets, can be leveraged to encode additional information on the communication signals.

SE filtering and control can be integrated in a feedback cycle to reinforce one another. To illustrate this point, consider the examples of SE filtering and control covered so far in this section. As illustrated in Fig. 3, SE filtering of an AR application can identify idle resources and free them by dropping irrelevant data, so that they can be re-purposed for radar sensing. In turn, sensing signals can be used to predict users’ physical interactions in the AR application; and thus, enable enhanced SE filtering of irrelevant broadband packets and reduction in the protocol overhead.

3.3 Semantic Control for Physical-Layer Computing

In distributed computing over a wireless network, a number of devices hold local data, and the end goal is for a receiver to compute a function of the data distributed across these devices. Examples of functions include aggregated statistics via sums or extremal values. Function computation can be considered a semantic task that is handled by current communication protocols solely at the application layer. Accordingly, each device would communicate its separate statistic to the receiver,

which would then compute the desired function at the application layer. In contrast, an SE control interface to the lower layers of the communication protocol can implement over-the-air computing techniques¹¹, hence, significantly enhancing the spectral efficiency of the system.

To provide a concrete example, consider implementing a supervised learning algorithm across wireless devices, each with its own local dataset. Current solutions for such tasks require devices to offload their local datasets to a cloud server, where a powerful learning algorithm, e.g., a deep neural network, is executed on the accumulated dataset, and the learned model is shared with the devices. However, privacy concerns, or the sheer size of the dataset, may prevent this centralized solution. Alternatively, in federated learning¹², the devices can be orchestrated by a parameter server, e.g., a base station or another wireless device, to learn a model in a distributed manner. At each iteration of the algorithm, each device updates and sends its local model or gradient estimate, and the goal of the parameter server is to compute the average of these to update the global model, which is multicast back to the devices.

To provide an illustration, in Fig. 4, the accuracy of a single-layer neural network trained for the classical MNIST digit classification problem by means of gradient averaging is shown versus the number of training iterations¹¹. We compare the performance of the standard method that computes the average at the application layer, and of over-the-air computation enabled through an SE plane. In the latter case, the average gradient is directly estimated from the received signal thanks to the superposition of wireless signals at the physical layer. We observe that over-the-air computation results in a significantly higher training accuracy, with a gain that increases with the number of devices. On the contrary, the accuracy of application layer computation decreases with the number of devices, due to the decreasing amount of spectral resources per device to send their local estimates separately.

In a practical protocol, physical-layer computing needs to include a process for estimating the variable wireless channels, either through dedicated protocol steps, or as part of the overall distributed learning algorithm.




Figure 4: Comparison of the training accuracy achieved by application layer computation vs. over-the-air computation for federated MNIST classification across wireless devices.

3.4 Other Examples of SE Filtering and Control

Having detailed a few applications of SE filtering and control, here we briefly list a few additional use cases.

- *Physical-layer provenance filtering* Validate data provenance using radio fingerprinting or location-based authentication through base-band signal processing, while removing the need for higher-layer operations (see, e.g.,¹³).
- *Physical-layer remote radio control* Allow access to physical resources, such as smart locks, via radio signatures controlled directly by the SE plane through trust mechanisms such as smart contracts. As an example, once rent is paid, a smart contract controls the physical layer transmission of a radio signature to open a smart lock¹⁴.
- *Layer-2 retransmission control* Retransmit only data that is expected to be still relevant on the basis of the internal state of ML/AI models by means of SE filtering (see, e.g.,¹⁵).
- *Network-layer traffic-based routing* Classify traffic on the basis of its network-level traces and configure routers accordingly to filter out protocol overhead.
- *Transport-layer semantic-based access and congestion control* Inject data in the network only when relevant and trusted, hence, reducing protocol overhead. A notable example is given by sensors and actuators that operate on the basis of input from local or cloud-based pre-




Figure 5: State of the art for ML/AI interfaces in the core network of 5G networks: interaction of NWDAF with other network functions.

dictive models. Accordingly, based on past data collected by sensors and actuators, an ML/AI module can instruct a sensor on how to carry out sampling and transmission and can inform an actuator of any change in its control module only when the local model is outdated.

- *Application-level aggregation for DLT transactions* Aggregate transactions from multiple applications running on a device or a network node in a single verifiable unit (e.g., block) prior to communications within a distributed ledger for, e.g., financial transactions or smart contracts, hence, reducing the protocol overhead.
- *Intent-based networking* Upon receiving users' instructions via natural language or visual interfaces, the network automatically reconfigures itself at all layers of the protocol stack to carry out the described task in the most efficient manner.

4 SE Functionalities in 5G and 3GPP

The proposed SE plane has precursors in the architectures that are currently under discussion within 3GPP. This section provides a brief overview of the corresponding state of the art.

The 5G SBA is designed to support the interaction of control-plane NFs. The Network Data Analytics Function (NWDAF) ^{6, 16} is one of the NFs in the 5G core system architecture. NWDAF is used for data collection and analytics, and utilizes the service-based model to communicate with other NFs and the Operation, Administration and Maintenance (OAM) system. The purpose of standardizing NWDAF, rather than relying on proprietary implementations, is to allow mobile network operators to handle analytics for NFs and third-party services. For example, NWDAF can be used to extract data from a NF, to derive statistical information about the past, or predictive information

about future events, and to provide this information to other NFs or to the OAM. These NFs or the OAM system can then take analytics-driven decisions to implement, e.g., proprietary self-organizing network tools. We emphasize that 3GPP does not standardize the data analytics algorithms, but only the interfaces and related procedures that support inter-NF communications, such as subscribe/unsubscribe procedures, as well as the type of information NWDAF can examine such as user equipment (UE) mobility and user plane congestion. To be more specific, in the 5G New Radio (NR) standard, as defined by the current Release 16 ⁶ and illustrated in Fig. 5, NWDAF interacts with and receives analytics-related information from different entities in the network, including the Access and Mobility Function (AMF), the Session Management Function (SMF), the Network Exposure Function (NEF), and the Policy Control Function (PCF). AMF allows NWDAF to access UE and base station location and mobility analytics; SMF allows NWDAF to receive session-related analytics; NEF allows NWDAF to access third party-related (e.g., YouTube, Facebook, etc.) service analytics; whereas PCF allows NWDAF to obtain policy-related information.

More recently, as part of Release 16 discussions, the 3GPP RAN group has started a study item on "RAN-centric data collection and utilization for NR", with the aim of broadening the set of standard-based solutions for enhanced data collection and utilization capabilities ¹⁷. The study item, which is currently in its early stages, has the objective to investigate features that could enable interaction of the RAN and the enhanced NDWAF blocks. It primarily aims to identify relevant data, such as Layer-1 and Layer-2 measurement quantities, and to investigate procedures that collect and utilize that data for automated network functionalities. Examples of applications under investigation include mobility optimization, random access channel optimization, load sharing/balancing-related optimization, coverage and capacity optimization, and minimization of drive testing.

5 Discussion and Conclusions

We conclude this paper with a brief discussion about challenges and implications of the proposed introduction of a SE plane.

We highlight the following two main key challenges towards the implementation of the SE plane concept.

- *Cross-Layer Recovery Mechanisms* The direct control of lower layers by the SE plane risks breaking mechanisms that have been carefully designed based on a hierarchy between different layers. As an example, the semantic layer may apply a filtering decision at Layer 2, that can have an implication on TCP retransmissions and reordering at Layer 4. Design of proper cross-layer mechanisms and interfaces should address this important technical problem.
- *Security* Data plane security today is mainly based on application, transport, and network layer mechanisms. In that context, it is not clear what are the security implications of exposing Application Programming Interfaces (APIs) below the network layer. For example, is there a risk in opening some new backdoors? Possible solutions to this problem could handle security directly at the SE plane and could be embedded in the API exposed to users, sensors, and actuators. For example, the SE plane may carry out physical layer-based authentication prior to granting access to some functionalities.

The standardization of the SE plane architecture introduced in this paper would by and large revolve around the definition of effective interfaces between communication layers and application programming interfaces towards users, actuators, and devices. On the one hand, as past experience has shown, any architectural change that affects the paradigm of low interaction between the different layers is significantly challenged in terms of widespread acceptance and standardization. On the other hand, considering the context of the post-5G wireless systems, the availability of well-defined software interfaces for SE filtering and control may invalidate the current “next-G paradigm” for mobile wireless evolution and standardization, ushering in an era of continuous “open-source” improvements and extensions. In this profoundly different ecosystem, defined by the separation between hardware and software and by standardised data analytics and control interfaces, the availability of open-source software running on general purpose hardware and the maturity of software defined radio solutions will lower the barriers for new entrants and enable faster innovation.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Acknowledgements

This work has received funding from the European Research Council (ERC) under the European Union Horizon 2020 research and innovation program (Grant agreements 725731, 677854 and 648382).

Received: 25 December 2019 Accepted: 19 March 2020
Published online: 5 May 2020

References

1. Shannon CE, Weaver W (1964) The mathematical theory of communication. University of Illinois Press, Champaign
2. Clark DD, Partridge C, Ramming JC, Wroclawski JT (2003) A knowledge plane for the internet. In: Proceedings of the 2003 conference on applications, technologies, architectures, and protocols for computer communications, ser. SIGCOMM '03. ACM, New York, pp 3–10. <https://doi.org/10.1145/863955.863957>
3. Sooriyabandara M, Farnham T, Mahonen P, Petrova M, Riihijarvi J, Wang Z (2011) Generic interface architecture supporting cognitive resource management in future wireless networks. *IEEE Commun Mag* 49(9):103–113
4. Simeone O (2018) A very brief introduction to machine learning with applications to communication systems. *IEEE Trans Cogn Commun Netw* 4(4):648–664
5. Altman E, Azouzi RE, Menasché DS, Xu Y (2010) Forever young: Aging control in DTNs. *CoRR*. [arXiv:1009.4733](https://arxiv.org/abs/1009.4733). [arXiv:1009.4733](https://arxiv.org/abs/1009.4733)
6. 3GPP (2019) Architecture enhancements for 5G System (5GS) to support network data analytics services; Stage 2 (Release 16). In: 3rd generation partnership project (3GPP), technical specification (TS) 23.288, 2, version 0.1.0
7. Christidis K, Devetsikiotis M (2016) Blockchains and smart contracts for the internet of things. *IEEE Access* 4:2292–2303
8. Popovski P, Trillingsgaard KF, Simeone O, Durisi G (2018) 5g wireless network slicing for embb, urllc, and mmcc: a communication-theoretic view. *IEEE Access* 6:55 765–55 779
9. Petrov V, Fodor G, Kokkonen J, Moltchanov D, Lehtomaki J, Andreev S, Koucheryavy Y, Juntti M, Valkama M (2019) On unified vehicular communications and radar sensing in millimeter-wave and low terahertz bands. *arXiv e-prints*, [arXiv:1901.06980](https://arxiv.org/abs/1901.06980)
10. Popovski P, Simeone O (2010) Protocol coding for two-way communications with half-duplex constraints. In: 2010 IEEE global telecommunications conference GLOBECOM 2010. IEEE, pp 1–5
11. Amiri MM, Gunduz D (2019) Machine learning at the wireless edge: distributed stochastic gradient descent over-the-air. *arXiv e-prints*. [arXiv:1901.00844](https://arxiv.org/abs/1901.00844)

12. Konecny J, McMahan HB, Ramage D, Richtárik P (2016) Federated optimization: distributed machine learning for on-device intelligence. CoRR. [arxiv:1610.02527](https://arxiv.org/abs/1610.02527). [arXiv:1610.02527](https://arxiv.org/abs/1610.02527)
13. Xu Q, Zheng R, Saad W, Han Z (2016) Device fingerprinting in wireless networks: challenges and opportunities. *IEEE Commun Surv Tutor* 18(1):94–104
14. Roush W, Pontin M (2018) *Twelve tomorrows*. MIT Press, Cambridge
15. Zhu G, Liu D, Du Y, You C, Zhang J, Huang K (2018) Towards an intelligent edge: wireless communication meets machine learning. *arXiv preprint*. [arXiv:1809.00343](https://arxiv.org/abs/1809.00343)
16. 3GPP (2018) System architecture for the 5G system; Stage 2 (Release 15). 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501, 12 2018, version 15.4
17. 3GPP (2019) Study on RAN-centric data collection and utilization for LTE and NR (Release 16). In: 3rd Generation Partnership Project (3GPP), Technical Report (TR) 37.816, 3 2019, version 0.2.0



Petar Popovski (S'97–A'98–M'04–SM'10–F'16) is a Professor at Aalborg University in Denmark. He received his Dipl. Ing and Magister Ing. degrees in communication engineering from the University of Sts. Cyril and Methodius in Skopje, and the

Ph.D. degree from Aalborg University. He has over 350 publications, over 30 patents. He received an ERC Consolidator Grant (2015), the Danish Elite Researcher award (2016), IEEE Fred W. Ellersick prize (2016), Highly Cited Researcher title (2018) and IEEE Stephen O. Rice prize (2018). He is a Steering Committee Member of IEEE SmartGridComm and IEEE Transactions on Green Communications and Networking, an Area Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and a Member-at-Large of the Board of Governors of the IEEE Communications Society. Prof. Popovski was the General Chair for IEEE SmartGridComm 2018 and IEEE Communication Theory Workshop 2019. His research interests are in wireless communication and communication theory.



Osvaldo Simeone (F'16) is a Professor of Information Engineering with the Centre for Telecommunications Research at the Department of Informatics of King's College London. He received an M.Sc. degree (with honors) and a Ph.D. degree from

Politecnico di Milano, Milan, Italy, in 2001 and 2005, respectively. From 2006 to 2017, he was a faculty member at New Jersey Institute of Technology (NJIT). He is a co-recipient of the 2019 IEEE Communication Society Best Tutorial Paper Award, the 2018 IEEE Signal Processing Best Paper Award, the 2017 JCN Best Paper Award, the 2015 IEEE Communication Society Best Tutorial Paper Award and of the Best Paper Awards of IEEE SPAWC 2007 and IEEE WRECOM 2007. He was awarded a Consolidator grant by the European Research Council (ERC) in 2016. He is a Fellow of the IET and of the IEEE.



Federico Boccardi is a Principal Advisor in the Ofcom Technology Group, where he leads the Mobile, Wi-Fi and IoT program and the Emerging Technologies program. He received a M.Sc. and a Ph.D. degree in Telecommunication Engineering from the

University of Padova, Italy, in 2002 and 2007, respectively, and a master in Strategy and Innovation from the Oxford Sad Business School in 2014. Before joining Ofcom, he was

with Bell Labs (Alcatel-Lucent) from 2006 to 2013 and with Vodafone R&D in 2014. Federico has published about 80 international research papers (with more than 6000 citations) and holds about 25 patents. He is a member of the Industrial Advisory Board for Department of Informatics at the Kings College London, a Honorary Visiting Professor at the University of Bristol and a Fellow of the IET.



Deniz Gündüz (SM13) is a Reader in information theory and communications at Imperial College London, UK. He received his Ph.D. degree from NYU Tandon School of Engineering in 2008. He is an Area Editor of the IEEE Transactions on Communica-

tions, and an Editor of the IEEE Transactions on Wireless Communications and the IEEE Transactions on Green Communications and Networking. He is the recipient of the IEEE Communications Society—Communication Theory Technical Committee Early Achievement Award in 2017, a Starting Grant of the European Research Council (ERC) in 2016, IEEE Communications Society Best Young Researcher Award for the Europe, Middle East, and Africa Region in 2014, Best Paper Award at the 2019 IEEE GlobalSIP, 2016 IEEE WCNC, and Best Student Paper Awards at 2018 IEEE WCNC and 2007 IEEE ISIT.



Onur Sahin received his B.S. degree in electrical and electronics engineering from Middle East Technical University, Ankara, Turkey, in 2003 and Ph.D. degree in electrical engineering from the Polytechnic Institute of New York University, USA in 2009.

He is currently a Senior Staff Engineer at Innovation Labs, InterDigital Europe. His primary research and development interests are on the next generation telecommunication and wireless systems (including 5G and beyond). Dr. Sahin has held technical lead positions at multiple projects on next generation cellular and Wi-Fi systems including 5G NR, LTE-A and IEEE 802.11 standards. He currently leads Beyond-5G ultra-high throughput (Tbps/THz) wireless technology design and development at InterDigital. Dr. Sahin is the co-author of over 50 peer-reviewed scientific articles and co-inventor of 25 patents and patent applications. He is co-recipient of the 2018 IEEE Signal Processing Society Best Paper Award, 2016 Journal of Communication Networks Best Paper Award, and InterDigital Innovation Awards in 2012 and 2015.