

IntellIoT: Intelligent IoT Environments

Arne Bröring¹, Vivek Kulkarni¹, Andreas Zirkler¹, Philippe Buschmann¹,
Konstantinos Fysarakis², Simon Mayer³, Beatriz Soret⁴, Lam Duc Nguyen⁴,
Petar Popovski⁴, Sumudu Samarakoon⁵, Mehdi Bennis⁵, Jérôme Härri⁶,
Martijn Rooker⁷, Gerald Fritz⁷, Anca Bucur⁸, Georgios Spanoudakis², and
Sotiris Ioannidis⁹

¹ Siemens AG, Technology, Munich, Germany

² Sphynx Analytics Ltd., Nicosia, Cyprus

³ University of St. Gallen, St. Gallen, Switzerland

⁴ Aalborg University, Aalborg, Denmark

⁵ University of Oulu, Oulu, Finland

⁶ EURECOM, Sophia Antipolis, France

⁷ TTTech Computertechnik AG, Vienna, Austria

⁸ Philips Research, Eindhoven, Netherlands

⁹ Technical University of Crete, Chania, Greece

Abstract. Traditional IoT setups are cloud-centric and typically focused around a centralized IoT platform to which data is uploaded for further processing. Next generation IoT applications are incorporating technologies such as artificial intelligence, augmented reality, and distributed ledgers to realize semi-autonomous behaviour of vehicles, guidance for human users, and machine-to-machine interactions in a trustworthy manner. Such applications require more dynamic IoT environments, which can operate locally without the necessity to communicate with the Cloud. In this paper, we describe three use cases of next generation IoT applications and highlight associated challenges for future research. We further present the IntellIoT framework that comprises the required components to address the identified challenges.

Keywords: internet of things; artificial intelligence; autonomous systems; human-computer interaction; trust;

1 Introduction

In today's Internet of Things (IoT) deployments, cloud-based platforms are typically central points of data collection and processing. However, this cloud-centric IoT model has limitations [13, 16, 27]: (i) unreliable cloud connectivity impedes dependable end-to-end applications, (ii) limited bandwidth restricts the amount of data that can be processed, (iii) high round-trip times prevent real-time operation, (iv) high cost of data transport and intake, as well as (v) privacy and trust concerns. Moreover, typical hierarchical setups of IoT cloud platforms (vi) hinder use cases with dynamically changing context due to lacking self-awareness of the individual subsystems and the overall system.

To enable next generation IoT applications, these issues can be overcome through localized IoT environments comprised of heterogeneous devices (e.g., edge computers as well as resource-constrained devices) that can collaboratively execute semi-autonomous IoT applications, which include functions for sensing, acting, reasoning, and control. However, since IoT applications cannot be completely autonomous in what they decide and act, they need to keep the human-in-the-loop for control and optimization of their Artificial Intelligence (AI).

In this paper, we derive research challenges from three key classes of Next Generation (NG) IoT use cases: 1) a fleet of agricultural vehicles (e.g., tractors) is semi-autonomously operated in conjunction with supporting devices (e.g., drones), 2) patients are semi-autonomously guided by artificial advisors based on IoT device input; and 3) semi-autonomous machine-to-machine collaboration in industrial plants (e.g., robot arms and machinery). In all three use case areas, a human expert plays a key role in controlling, monitoring and teaching AI-enabled autonomous systems.

The remainder of this paper is organised as follows: Section 2 presents the three use case classes from agriculture, healthcare and manufacturing. Section 3 describes the IntellIoT framework to enable NG IoT application development. Section 4 presents current research around key enablers to fulfill the vision and highlights associated research challenges. Section 5 concludes this paper and points to future work.

2 Next Generation IoT Use Cases

Due to the dimensions of variability, we selected three distinct use cases that stand exemplary for a broad range of NG IoT applications.

2.1 Autonomous Operation of *Agriculture Vehicle Fleets*

Fig. 1 describes the use case of a semi-autonomous agricultural vehicle fleet. This use case entails the provision of new functionalities (e.g., AI algorithm implementations) for IoT applications by technology providers (e.g. tractor manufacturers) via step (1). A human operator (Farmer or Agriculture fleet management) specifies a goal for autonomous activities (e.g., ploughing or spraying a certain farm field) of the tractor on an Edge infrastructure as depicted in step (2). From the defined goal, a plan for IoT application instantiation is derived and a deployment of the required functions to the involved devices, e.g., tractor or drones shown in step (3), is triggered. Next, the deployed AI operates the involved vehicles, which includes dealing with blockages or other adversarial events using sensors of the vehicle (e.g., cameras or LIDAR) via step (4). This can be facilitated by sensing the environment from multiple neighbouring vehicles to collectively train their models and identify objects in a faster and more robust manner. The sensed data can also serve as a training dataset for continually improving the underlying AI models.

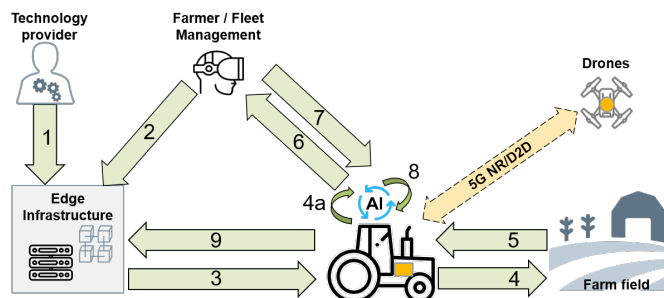


Fig. 1. IntelliIoT-enabled NG IoT agriculture use case.

If an obstacle is detected and the tractor cannot determine how to traverse it (step 5), then control of the tractor is handed over to a human operator (step 6). Data from the cameras and sensors is transmitted to the human operator who can use AR/VR technology to have a surround view of the situation. Direct and indirect strategies for taking over remote control are both necessary for step (7): Direct interaction with the tractor (i.e. remotely driving the vehicle) requires a reliable and high-speed connection enabling real-time interaction between operator and tractor. Indirect control with VR tooling aims to identify a feasible trajectory around the tractor, and control is given back to the tractor, while the human operator supervises the tractor remotely while it traverses the newly defined trajectory correctly. Based on the input from the human operator, the vehicle can refine its AI models by continually learning (step 8) how to overcome such obstacles in the future in addition to potentially sharing the learned model with other vehicles.

In the future, service providers will also offer such semi-autonomous vehicles (e.g., to provide farming services). Then, contractual agreements need to be set up using distributed ledger technology (DLT) (e.g., ownership of the farmer's land needs to be confirmed). This information will constitute a digital evidence that the field owner authorized the requested services and the area in which the smart equipment operates. Storing performed agricultural activities back in the DLT as historic evidence (step 9) can then be utilized in business models.

2.2 Collaborative intelligence for remote patient monitoring

Advances in AI and in IoT-enabled systems may lead to significant benefits in healthcare, enabling physicians to efficiently improve patient outcomes, safety and comfort, for instance by leveraging the new technology to remotely guide their patients through recovery and rehabilitation at home. The solution empowers patients to focus on their recovery, giving them the confidence that they are safe, that the tools can support and inform them all the way, and that their physicians are always in the loop when needed.

Fig. 2 describes a system that leverages IoT device inputs to give clinical experts accurate information on the health status of their patients and provides

AI-assisted recommendations and interventions to patients, with clinical expert oversight. Patient users are equipped with wearable sensing devices measuring relevant data that is transferred to a personal IoT device (e.g., smart watch) (step 1). The AI on board of this IoT device analyses the data in step (2) to identify the need for interventions or recommendations, according to the initial AI model and the intervention workflows previously defined as goals by the clinicians in charge of the patient. The model is applied on the collected data, and when the need for an intervention is detected, either a recommendation is sent to the patient (step 3a), or the case is escalated to involve a clinician, leading to the human-in-the-loop (step 3b). Privacy and security-compliant exchange is thereby crucial. The system may further implement a model for monitoring and diagnosing technical issues with the constrained devices (step 4).




Fig. 2. IntellIoT-enabled NG IoT healthcare use case.

In the IntellIoT solution, when an escalation takes place and a clinical expert is notified, the clinician may decide to contact the patient as shown (step 5a), respond to the personal device (step 5b), or raise an alarm (step 5c). The clinician provides feedback which is used to validate and re-train the AI model locally on the personal IoT device (step 6). Model updates are then contributed to the aggregated model at the edge infrastructure (e.g., of a hospital) (step 7). This distributed AI can be implemented using federated learning and the model update is communicated to all IoT devices, either in a device to device fashion (step 8a) or through distribution of the aggregated model (step 8b). Further, federated learning can be done between hospitals (step 9). All the involved communications and interactions need to be covered by state-of-the-art security and privacy provisions, catering for the intricacies of the private-sensitive user data. Digital consent management to drive the interactions of the system (patients, clinicians, devices) can be managed e.g. via smart contracts.

2.3 Autonomous Collaboration of *Production* Machines

This use case (Fig. 3) describes an example of machine to machine collaboration. A customer of a shared manufacturing plant orders a product by specifying a

manufacturing goal (step 1). In step (2), a machine orchestration and associated process plan is determined to manufacture the desired product from a workpiece. The event-based process planner monitors the manufacturing process and reacts when the health state of a machine changes. If the process planner cannot find a solution for the manufacturing goal on its own (step 3) it can request support from a human plant operator or eventually customer. In step (4), a robot or AGV is tasked to transport the workpiece to the next production step and manufacturing process data is sent to involved machines. As these machines may be operated by the plant owner or a third-party operator, contractual arrangements need to be set up, for which a distributed ledger is used. Further, comprehensive security mechanisms are applied to ensure privacy and security of customer data.




Fig. 3. IntelliIoT-enabled NG IoT manufacturing use case.

In step (5), a local AI on board of the robot decides how the robot picks a workpiece and places it in the next machine. If the confidence-level of the local AI is low and it cannot pick and place the workpiece, it can request support from a human plant operator or machine owner again (step 6). Utilizing AR/VR technology, the human can virtually grab the workpiece to support the robot. A tactile communication needs to be established for this interaction, under consideration of security and privacy. Additionally, 3D cameras can be used to generate an accurate enough reconstruction of the surroundings and the robot itself which allows the full control and visual information about the parameters of the robot. For grabbing and haptic feedback to the user, special user input devices (e.g., a stylus or glove) are needed. If support from a remote operator is needed, a tactile communication may not be possible through long-distance internet connection. Hence, the operator would be able to control a virtual robot, rendered in the local edge, with delayed movement of the real robot. From the human handling of the work piece, local AI on the robot re-trains itself (step 7), and federates the learned process parameters to other robots through model update on edge (step 8).

3 A Next Generation IoT Framework: the IntellIoT approach

Analyzing the above use cases, a pattern of NG IoT applications can be extrapolated, as shown in Fig. 4: NG IoT applications generally consist of multiple heterogeneous devices that collaborate in a semi-autonomous way through AI. Users interact with the system to provide knowledge and thereby may (re-)train the AI. Interaction among the devices and with the user may have to happen in a tactile way, with low latencies and high bandwidth.




Fig. 4. Pattern of an NG IoT Application.

Tackling the above use cases in a holistic manner is the driver of the IntellIoT project¹⁰. It aims to develop a framework for the management of intelligent IoT environments and their IoT applications, which is realized through an architecture comprising three building blocks (see Fig. 5): (a) distributed, self-aware, & semi-autonomous **IoT applications**; (b) a **human-in-the-loop** to define and support the autonomy in (a), and (c) an efficient, reliable and trustworthy **computation and communication infrastructure** that enables (a) & (b).




Fig. 5. IntellIoT concept for enabling intelligent IoT environments

3.1 Distributed, self-aware, semi-autonomous IoT applications

Autonomous software agents of a novel hypermedia-based multi-agent system (HyperMAS) [8] execute IoT applications. Interoperable access to agents and

¹⁰ <http://intelliott.eu>

functions is given through standardized interfaces that are hosted by IoT or edge devices, e.g., based on W3C Web of Things specifications [20]. Using these functions as building blocks, software agents can autonomously create distributed IoT applications and execute these applications while flexibly reacting to environment dynamics. To further facilitate IoT application development, interoperability is supported through components that are able to translate between communication technologies, protocols and vocabularies. Software agents are self-aware and observe each other, e.g., to detect and autonomously mitigate failures. They participate in a distributed ledger to enable contractual relations and monetization. Leveraging reinforcement and federated learning [22], distributed AI is enabled by on-device training and inference that are subject to the device’s resource constraints.

3.2 Autonomy defined by a Human-in-the-Loop

The human-in-the-loop provides expertise to the IoT environment and is therefore crucial to the system: At design time, the human defines goals and requirements. Then, a mechanism automatically deducts and translates an IoT application workflow into IoT/edge device interactions with associated network constraints. At runtime, the human observes the AI-enabled autonomous behavior and provides input [14] to improve it. For that, the human needs to leverage tactile interactions through AR/VR to refine the model and avoid blockages by e.g., teaching an industrial robot how to handle a product.

3.3 Efficient, reliable and trustworthy computation & communication infrastructure

Intelligent IoT environments must operate upon a communication & computation infrastructure capable of flexibly supporting the capabilities described in 3.1 & 3.2 above, whereby resource-constrained IoT devices and more powerful edge assets must be efficiently managed, optionally integrating cloud-based services, and also supporting complex, cost-intensive computations (e.g., AI inference/training, as well as AR rendering). Edge resources will be diverse [28], e.g., Multi-access Edge Computing (MEC) offered through 5G functionalities, or an industrial edge offered by networked computing devices in a manufacturing plant. Computation & communication form a closed-loop system through which the infrastructure will be optimized in an integrated way (i.e., deployment of application functions on IoT/edge resources must be optimized under consideration of network constraints, and the network must be dynamically managed and reconfigured to optimally serve the purpose of the application and the IoT/edge devices). The infrastructure will enable ultra-reliable and low-latency communication through dynamic network management, through heterogeneous network technologies (e.g., 5G NR [17], NB-IoT [15], or D2D [1]). The wireless front end will be specifically designed to support communication requirements of advanced techniques, such as DLTs and federated learning. Finally, security

& privacy assurance concepts will be included by design to ensure reliability and overall trustworthiness of the developed solution.

3.4 Bringing all together - the IntellIoT high-level architecture

Integrating the above concepts, a high-level view of IntellIoT’s logical architecture has been derived, shown in Fig. 6. The three key concepts highlighted (i.e., Collaborative IoT, Human-in-the-loop and Trustworthiness) are prominently featured in the architecture.

In total, five core component groups have been identified, with individual components falling into one of the following groups:

- **Collaborative IoT enablers:** Components that realize IntellIoT’s Collaborative IoT pillar, focusing on the cooperation of various semi-autonomous entities to execute IoT applications.
- **Human-in-the-Loop enablers:** Components involved in IntellIoT’s Human-in-the-Loop pillar, which focuses on involving the human in the process; e.g., to solve complex situations.
- **Trust enablers:** Components that are part of IntellIoT’s Trust pillar. This pillar focuses on privacy, security, and ultimately building trust into the IntellIoT framework.
- **Infrastructure management:** The computation & communication infrastructure and its management capabilities, enabling the deployment and management of edge applications.
- **Use-Case deployment:** Components which are use case-specific, (i.e., pertaining to the use case environment deployment), such as edge devices, edge apps, and edge AI models.

For more details on the individual components comprising the architecture, we defer the reader to the publicly-available architecture specification of IntellIoT [9].

4 State of the Art and Research Challenges

To achieve its vision, IntellIoT improves the state of the art in the related research areas; the key enablers and resulting research challenges are highlighted in the subsections that follow.

4.1 Autonomy and distributed intelligence

Next generation IoT applications require a paradigm shift from classic ML to distributed, low-latency and reliable ML at the wireless network edge [22]. Federated learning (FL) is a decentralized learning technique where private-sensitive training data is distributed across learning agents[19]. Agents share their local models, instead of the training data, reducing communication latencies during ML training. Nevertheless, except few works, such as [26, 7], most of the existing literature assumes ideal client-server communication conditions, overlooking




Fig. 6. IntelliIoT high-level architecture

channel dynamics and uncertainties. FL uses stochastic gradient decent (SGD) techniques (e.g., elastic SGD, entropy SGD) for local training, each with its own intrinsic characteristics (computation requirements, accuracy etc.). However, the impact of different ML algorithms in real-world applications is an open research area. The continual discovery and interaction of agents within their environment requires the use of multi-agent Reinforcement Learning (RL). Furthermore, the branch of deep RL (DRL) [5] addresses issues arising from the larger state dimensions. Model-free, value/policy-based, and actor-critic RL within DRL exhibit efficient and accurate decision-making capabilities over classical RL. Yet the aspects of computation-communication limitations and privacy in distributed multi-agent RL are still not well-understood and require further investigation. The involvement of a human expert in data collection, training, testing, and validation is the fundamental philosophy behind the human-in-the-loop for ML. With distributed AI techniques, the interaction between the agents and the human needs to consider time sensitivity, learning procedure and ability to control and train/teach remotely in the scalable systems. Hence, it is mandatory to investigate various transfer learning methods [24] as well as suitable optimizers and the human-in-the-loop of the training [14].

For the IntellIoT framework and realization of the envisioned use cases (Section 2), distributed ML needs to consider application-specific target accuracies and worst-case training latencies under tolerable number of failures (reliability and robustness guarantees), wireless resources availability, on-device energy, storage, or computing restrictions. In addition, studying the control stability (plant, string, swarm- stabilities) of both single and multi-agent systems will be mandatory. Investigating the co-design of ML, communication-computation and control are crucial for developing novel distributed AI solutions. For fully enabling the human-in-the-loop, the fusion between transfer learning, optimization and FL/RL are being incorporated.

4.2 Next generation IoT computation and communication infrastructure

IoT applications are moving from the cloud to the edge, so that computing happens in closer proximity to the data producers and consumers [28]. Relevant solutions include concepts such as fog computing or multi-access edge computing (MEC), where computing resources are part of a 5G network. This has the potential to address the concerns of response time requirements, battery life constraints, bandwidth cost saving, or safety and privacy (e.g., [27, 13]). In NG IoT applications, a key challenge for the computation infrastructure is to decide on which computing resource to handle a specific workload (e.g., execute an AI algorithm). There are multiple existing allocation strategies, e.g., [6, 21], which optimize different performance metrics, e.g., response time, bandwidth, availability, or energy consumption. Therefore, components are needed for advanced, dynamic resource management which can be flexibly applied in various private edge environments. IntellIoT develops a mechanism for optimized allocation of workloads to computing resources (i.e., mapping of IoT applications to devices). It consists of a flexible algorithmic framework that builds on prior work [23] and is adjustable to different optimality criteria at runtime. Further, it needs to be dynamically adapting to network changes based on high-level application requirements [4]; i.e., establishing closed-loop infrastructure management.

While IoT/edge devices can provide the computation side of the infrastructure, the communication side needs to be driven by advanced networking technologies, such as 5G New Radio (NR) and its extensions towards private networks and Industrial IoT. Further, 5G eV2X, as a complete redesign of the LTE V2X, can play a role for cooperative automated mobility and 'sidelink' device-to-device (D2D) [1] communications, e.g., between robots and machines. For "tactile" communication links, a major challenge is to design a steer-/control-based communication framework for real-time transmission of haptic information (touch, motion, etc.) in addition to conventional audio-visual and data traffic. The provided solutions need to enable efficient spectrum usage [3] in downlink, with massive sensory feedback on the uplink, targeting 1ms downlink, 10ms uplink with 99.99% availability. In order to support ultra-reliable and low latency (URLL) communication towards TSN for 5G Industrial IoT, 5G mmWave

radio (with fiber data speed, real-time reactivity and massive sensorics capacity) [12] as well as support for IEEE TSN, are being investigated. Regarding distributed networking support, functions for ad-hoc scheduling capabilities for enhanced D2D communication do not include a specific scheduler, hence, IntellIoT develops a wireless TSN-grade D2D scheduler providing deterministic QoS for decentralized computing in the IoT context.

Building on the computation and communication infrastructure, IoT artifacts need to be able to discover and interact with one another. A first major step towards this goal has been the Web of Things (WoT) [20], where interactions between devices are based on the Web architecture. Crucially, however, interoperability on the semantic level is a central requirement in the future evolution of the Web. Based on efforts of the W3C WoT, new means to use hypermedia for designing evolvable Web APIs and general-purpose clients are being explored. IntellIoT will build up on these developments towards integrating them with research on multi-agent systems (MAS) towards enabling a hypermedia-based MAS (HyperMAS) [8] that are vertically and horizontally scalable with respect to the number of agents, devices, and interactions among these components. It will support self-aware agents within IoT environments and semi-autonomous IoT systems.

4.3 Humans and trust in intelligent IoT

The wide adoption of IoT technologies in a plethora of domains, necessitates considering security, privacy, and trust requirements early in the design phase [10]. Even securely initialised devices can be compromised, allowing attackers to affect connected devices, the network, or collaborative applications. Trust-based mechanisms can be used to defend against such attacks by monitoring the behaviour of each participant. An IoT deployment must also have the intelligence to protect itself proactively, e.g., through Moving Target Defence (MTD) techniques [25], where AI-driven agents periodically alter the network topology and/or configuration to counter attacks. Thereby, security assurance evaluations for IoT systems are still in their infancy (e.g., [2]). Therefore, IntellIoT provides security and trustworthiness by design, via a combination of: (i) an evidence-based continuous security assurance, integrating hybrid assessments which considers different attack surfaces and vulnerabilities; ii) trust-based computing mechanisms that will act as distributed intrusion detection system, and (iii) MTD strategies with security-context aware processes.

Supporting these security and trust mechanisms, IntellIoT uses distributed ledger technologies (DLT) to encode transaction logic and policies, which include the requirements and obligations of the party requesting access to an IoT resource as well as its provider [18]. This can lead to a wave of novel applications, enabling trusted access to IoT resources. Therefore, the state-of-the-art is being progressed in three aspects (building on previous work [11]): (1) circumventing devices' resource constraints; (2) advancing uplink-dominated IoT network designs; (3) providing interoperability with third-party devices. In IoT, an edge

gateway with DLT solutions is equipped with the necessary computational intelligence. Yet, devices in a blockchain should keep a copy of the DLT record, which can be large and increasing over time and limits the scalability of the system. Moreover, the transactions associated with smart contracts require two-way communication traffic, which violates the common assumption that the IoT systems are dominated by an uplink traffic. An IoT/edge device in a blockchain network should be capable of verifying information in the blockchain, which is associated with the downlink traffic. Therefore, the IntellIoT framework provides an architecture that aims at trading off complexity of the device, achieved trust and network capabilities, and maintain the trust when a device belongs to a third party.

5 Conclusions & Future Work

The key contributions of this paper is the analysis of three classes of Next Generation IoT use cases, the extrapolation of a common pattern, the presentation of the IntellIoT framework, and the postulation of key research challenges associated with it. All three use cases are based on semi-autonomous behaviour of the IoT system. Multiple heterogeneous devices are interacting and autonomous control of their collaboration is provided through AI, which can be (re-)trained through human intervention. This pattern can be assumed for many next generation IoT applications.

The described *pattern* spreads over three key areas: (1) providing the distributed artificial intelligence for autonomous behaviour, (2) providing efficient and reliable communication and computation resources, and (3) incorporating the human (by providing trust in the system) and learning from his input. The described framework of IntellIoT addresses all three fields. The presented high-level architecture combines software components to realize functionalities required by these fields. The full implementation of this architecture is currently in process. Thereby, the key research challenges that are being faced are outlined and describes the further path of research for this project and beyond.

Acknowledgment

This work has received funding from the European Union’s Horizon 2020 research and innovation programme H2020-ICT-56-2020, under grant agreement No. 957218 (Project *IntellIoT*).

References

1. Rafay Iqbal Ansari, Chrysostomos Chrysostomou, Syed Ali Hassan, Mohsen Guizani, Shahid Mumtaz, Jonathan Rodriguez, and Joel JPC Rodrigues. 5G D2D networks: Techniques, challenges, and future prospects. *IEEE Systems Journal*, 12(4):3970–3984, 2017.

2. Claudio A Ardagna, Ernesto Damiani, Julian Schütte, and Philipp Stephanow. A case for IoT security assurance. In *Internet of Everything*, pages 175–192. Springer, 2018.
3. Mehdi Bennis, Mérouane Debbah, and H Vincent Poor. Ultrareliable and low-latency wireless communication: Tail, risk, and scale. *Proceedings of the IEEE*, 106(10):1834–1853, 2018.
4. Arne Bröring, Jan Seeger, Manos Papoutsakis, Konstantinos Fysarakis, and Ahmad Caracalli. Networking-aware IoT application development. *Sensors*, 20(3):897, 2020.
5. Fanyu Bu and Xin Wang. A smart agriculture IoT system based on deep reinforcement learning. *Future Generation Computer Systems*, 99:500–507, 2019.
6. Valeria Cardellini, Vincenzo Grassi, Francesco Lo Presti, and Matteo Nardelli. Optimal operator placement for distributed stream processing applications. In *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems*, pages 69–80. ACM Press, 2016.
7. Mingzhe Chen, Zhaohui Yang, Walid Saad, Changchuan Yin, H Vincent Poor, and Shuguang Cui. A joint learning and communications framework for federated learning over wireless networks. *IEEE Transactions on Wireless Communications*, 2020.
8. Andrei Ciortea, Simon Mayer, Fabien Gandon, Olivier Boissier, Alessandro Ricci, and Antoine Zimmermann. A Decade in Hindsight: The Missing Bridge Between Multi-Agent Systems and the World Wide Web. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1659–1663. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
9. IntelliIoT consortium. Deliverable D2.3 - High level architecture (first version). <https://intelliott.eu/wp-content/uploads/2021/10/D2.3-High-level-architecture-first-version.pdf>.
10. Mauro Conti, Ali Dehghantanha, Katrin Franke, and Steve Watson. Internet of things security and forensics: Challenges and opportunities, 2018.
11. Pietro Danzi, Anders E Kalor, Rene B Sorensen, Alexander K Hagelskjær, Lam D Nguyen, Cedomir Stefanovic, and Petar Popovski. Communication aspects of the integration of wireless iot devices with distributed ledger technology. *IEEE Network*, 34(1):47–53, 2020.
12. Marco Giordani, Michele Polese, Arnab Roy, Douglas Castor, and Michele Zorzi. Initial access frameworks for 3GPP NR at mmWave frequencies. In *2018 17th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, pages 1–8. IEEE, 2018.
13. Kiryong Ha, Zhuo Chen, Wenlu Hu, Wolfgang Richter, Padmanabhan Pillai, and Mahadev Satyanarayanan. Towards wearable cognitive assistance. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*, pages 68–81. ACM, 2014.
14. Andreas Holzinger, Markus Plass, Michael Kickmeier-Rust, Katharina Holzinger, Gloria Cerasela Crişan, Camelia-M Pinteau, and Vasile Palade. Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Applied Intelligence*, 49(7):2401–2414, 2019.
15. Bing-Zhi Hsieh, Yu-Hsiang Chao, Ray-Guang Cheng, and Navid Nikaein. Design of a UE-specific uplink scheduler for narrowband Internet-of-Things (NB-IoT) systems. In *2018 3rd International Conference on Intelligent Green Building and Smart Grid (IGBSG)*, pages 1–5. IEEE, 2018.

16. Mohammad Manzurul Islam, Sarwar Morshed, and Parijat Goswami. Cloud computing: A survey on its limitations and potential solutions. *International Journal of Computer Science Issues (IJCSI)*, 10(4):159, 2013.
17. Florian Kaltenberger, Guy de Souza, Raymond Knopp, and Hongzhi Wang. The OpenAirInterface 5G New Radio Implementation: Current Status and Roadmap. In *WSA 2019; 23rd International ITG Workshop on Smart Antennas*, pages 1–5. VDE, 2019.
18. Minhaj Ahmad Khan and Khaled Salah. IoT security: Review, blockchain solutions, and open challenges. *Future Generation Computer Systems*, 82:395–411, 2018.
19. Jakub Konečný, H Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.
20. M. Kovatsch, R. Matsukura, M. Lagally, T. Kawagucchi, K. Toumura, and K. Kajimoto. Web of Things (WoT) Architecture. <https://w3c.github.io/wot-architecture>, 2019.
21. Nitinder Mohan and Jussi Kangasharju. Edge-Fog cloud: A distributed cloud for Internet of Things computations. In *2016 Cloudification of the Internet of Things (CIoT)*, pages 1–6. IEEE, 2016.
22. Jihong Park, Sumudu Samarakoon, Mehdi Bennis, and Mérouane Debbah. Wireless network intelligence at the edge. *Proceedings of the IEEE*, 107(11):2204–2239, 2019.
23. Jan Seeger, Arne Bröring, and Georg Carle. Optimally self-healing iot choreographies. *ACM Transactions on Internet Technology (TOIT)*, 20(3):1–20, 2020.
24. Liting Sun, Cheng Peng, Wei Zhan, and Masayoshi Tomizuka. A fast integrated planning and control framework for autonomous driving via imitation learning. In *Dynamic Systems and Control Conference*, volume 51913, page V003T37A012. American Society of Mechanical Engineers, 2018.
25. Xin-Li Xiong, Lin Yang, and Guang-Sheng Zhao. Effectiveness Evaluation Model of Moving Target Defense Based on System Attack Surface. *IEEE Access*, 7:9998–10014, 2019.
26. Howard H Yang, Zuozhu Liu, Tony QS Quek, and H Vincent Poor. Scheduling policies for federated learning in wireless networks. *IEEE transactions on communications*, 68(1):317–333, 2019.
27. Shanhe Yi, Zijiang Hao, Zhengrui Qin, and Qun Li. Fog computing: Platform and applications. In *2015 Third IEEE Workshop on Hot Topics in Web Systems and Technologies (HotWeb)*, pages 73–78. IEEE, 2015.
28. Ashkan Yousefpour, Caleb Fung, Tam Nguyen, Krishna Kadiyala, Fatemeh Jalali, Amirreza Niakanlahiji, Jian Kong, and Jason P. Jue. All one needs to know about fog computing and related edge computing paradigms: A complete survey. *Journal of Systems Architecture*, 98:289 – 330, 2019.