Yuan He, Junchen Guo, and Xiaolong Zheng

From Surveillance to Digital Twin

Challenges and recent advances of signal processing for the industrial Internet of Things

INTERNET OF THINGS—ISTOCKPHOTO.COM/IAREMENKO
CIRCUITS—IMAGE LICENSED BY INGRAM PUBLISHING

With the recent advances in the Internet of Things (IoT), the significance of information technologies to modern industry is upgraded from purely providing surveillance-centric functions to building a comprehensive information framework of the industrial processes. Innovative techniques and concepts emerge under such circumstances, e.g., digital twin, which essentially involves data acquisition, human–machine-product interconnection, knowledge discovery and generation, and intelligent control, etc. Signal processing techniques are crucial to the aforementioned procedures but face unprecedented challenges when they are applied in the complex industrial environments. In this article, we survey the promising industrial applications of IoT technologies and discuss the challenges and recent advances in this area. We also share our early experience with Pavatar, a real-world industrial IoT system that enables comprehensive surveillance and remote diagnosis for ultrahigh-voltage converter station (UHVCS). Potential research challenges in building such a system are also categorized and discussed to illuminate the future directions.

Introduction

With the prosperity of various embedded sensors, low-power wireless communications, and efficient signal processing techniques, the IoT has achieved explosive development and proliferation in recent years. The IoT offers opportunities to bridge the physical world and cyberspace, enabling fine-grained sensing of objects and environments, continuous data gathering, comprehensive information fusion, deep analysis, and real-time feedback or control over the connected targets. According to Gartner's report, there are approximately 8 billion connected things providing smart services in our daily life, for example, in assisted living, building surveillance, traffic control, and environment monitoring, etc [1].

The ever-developing IoT attracts the interest of both industry and academia. Technology giants have already taken their steps. Huawei, one of the world's largest telecom equipment makers, has devoted itself to standardizing the narrow-band IoT (NB-IoT) as the next generation of low-power wide-area networks to

Digital Object Identifier 10.1109/MSP.2018.2842228
Date of publication: 28 August 2018

fulfill its long-term strategy for building a better-connected world [2]. Practical IoT platforms have also been vigorously promoted recently, e.g., Android Things (Google), Predix [General Electric (GE)], Azure IoT Suite (Microsoft), etc. In the meantime, academia focuses on exploring cutting-edge techniques to boost the application and development of the IoT, such as wireless sensing, indoor localization, low-power networking, backscatter communication, visible light communication, mobile computing, edge computing, privacy and security, etc.

Among all of the promising scenarios, applying IoT technologies in modern industry has great potential and practical significance. In 2011, Industry 4.0 is proposed to equip traditional manufacturing with cyberphysical systems to start a new industrial revolution. GE formally put forward the concept of the Industrial Internet in 2012 [3]. GE then established the Industrial Internet Consortium with AT&T, Cisco, Intel, and IBM, bringing together the world's leaders in the manufacturing, telecom, networking, semiconductor, and computer industries, respectively, to promote the industrial IoT systems.

Due to the prosperity of IoT techniques in the past few years, digital twin has recently gained extensive attention. Digital twin represents a dynamic digital replica of physical assets, processes, and systems, which comprehensively monitors their whole life cycle. The backbone technology of digital twin is the IoT for real-time and multisource data collection. In addition, it integrates artificial intelligence and software analytics to create digital simulation models that dynamically update and change along with their physical counterparts. Moreover, digital twin adopts modern data visualization schemes such as virtual reality (VR) and augmented reality that can provide more illustrative and user-friendly views.

Therefore, compared to traditional surveillance systems, digital twin provides more sensing modalities with stricter timeliness guarantees, and integrates more intelligent data analysis and friendlier display and interaction. With digital twin, we can not only better understand and predict the performance of machines and systems, but also optimize business operations for equipment suppliers and consumers. However, it is a nontrivial task to achieve such comprehensive monitoring along with requirements such as timeliness, accuracy, scalability, and interoperability in industrial IoT. We summarize potential research challenges as follows.

- First, digital twin pushes the boundary of sensing capabilities toward the physical world. Sensing methods that monitor diverse physical metrics but rely on less resources are deemed to be more practical in industrial IoT. Wireless and battery-free sensing integrating efficient techniques of data cleaning and signal processing can support lightweight and robust monitoring. How to extend the sensing capabilities of wireless signals [4]–[7] and how to refine the sensing precision from vulnerable readings [8]–[10] have triggered numerous research motivations over the past few years.
- Second, visual sensing is extremely informative for the surveillance of physical assets and their surroundings. In digital twin, intensive networked cameras are deployed at a high density to provide seamless monitoring. On one hand, processing intensive networked videos need the upgrade of computing architecture for timeliness requirements, e.g., collaborative edge computing [11]–[13]. On the other hand, enabling a resource-constrained IoT device with modern analysis techniques, e.g., deep learning, can also release the pressure of cloud infrastructure and save the network bandwidth [14]–[17].
- Third, new forms of communication and networking is anticipated in digital twin for efficient data transmission. Recent advances in low-power wireless networking such as low-power wide-area networks [2], parallel backscatter transmissions [18], and software-defined low-power wireless [19] has drawn much attention. In this section, we emphasize the research challenges and opportunities on direct communication among heterogeneous wireless technologies that share the same frequency band [20]–[23], and their upper-layer protocols as well as applications [24]–[26].
- Last but not least, comprehensive data analysis and system diagnosis need innovative and dedicated signal processing methods. For example, anomaly detection and repairing of time-series data [27], feature selection from heterogeneous stream data [28], and fault analysis based on incomplete data [29] should also be well designed.

Practical industrial IoT and signal processing

Signal processing algorithms are indispensable in almost every layer of industrial IoT. In this section, we survey the most recent research works and corresponding signal processing techniques, to provide an overview of the current progress from sensing, networking to data analysis in industrial IoT.

Wireless and battery-free sensing

In practical industrial scenarios, many physical metrics need to be closely monitored, such as temperature and humidity, vibration and noise, rotation speed, liquid leakage, etc. Although the advances of modern sensor technologies enable the sensibility of more metrics, a part of these metrics cannot be provided due to the complicated operational environments in real-world deployments that have the special characteristics that are given next.

- *Requirement of nonintrusive sensing*: Adding dedicated sensors into the existing equipments costs too much because these intrusive sensors may trigger hardware updates or even redesigns. Hence, nonintrusive sensing methods are more preferred.
- *Large-scale sensing targets*: The large number of targets to be monitored makes it unaffordable to deploy dedicated sensors at all the monitoring points. Novel low-cost sensing solutions are desired.
- *Limited sensing capability*: Physical metrics can be very fast changing, but most nonintrusive sensors can usually provide undersampled data. How to fill this gap remains a challenge.

Because traditional sensors are mostly intrusive, those approaches cannot be deployed with an operational machine that hasn't been initially equipped with such a capability. Other high-resolution approaches, e.g., cameras and lasers, suffer from the line-of-sight problem and are restricted in the application context.

Moreover, audio-based sensing is sensitive to environmental noises, which is therefore impractical for real-world industrial applications. Wireless and battery-free sensing, e.g., radio-frequency identification (RFID), which leverages backscattered radio-frequency signals to carry information, has received plenty of attention in the past few years, due to its low-cost, nonintrusive, and easy-deployment properties. A typical RFID system, as shown in Figure 1, consists of RFID tags that store information in nonvolatile memories, and two-way radio transmitter–receivers called *RFID readers* that send signals to tags and receive their responses.

Recent advances in RFID offer a promising technique for cross-modal sensing where many physical metrics are sampled with only battery-free devices and wireless signals [4]–[7]. In the meantime, the resolution of RFID sensing—especially battery-free localization and tracking—has been well improved over the past few years [8]–[10].

Cross-modal sensing with RFID

Apart from parsing the information encoded in backscatter signals from tags, widely employed commercial RFID readers, e.g., ImpinJ Speedway R420, Alien ALR-9900, and Zebra FX9500, can interrogate the readings of received signal strength indicator (RSSI) and phase values at the frequency of approximately 40 Hz. The changes in RSSI and phase offer space for the cross-modal sensing of other physical metrics, e.g., vibration [4] and eccentricity [6] of rotating machines, liquid category [5], and human–object interaction [7]. However, the relatively low interrogating frequency offered by commercial readers brings in additional research challenges in industrial scenarios.

A recent battery-free work, TagBeat, offers inexpensive and pervasive cross-modal sensing of mechanical vibration frequency with commercial RFID devices [4]. The phase shifts caused by micro vibration are too tiny to distinguish, and the high-frequency vibration is hard to capture with the limited-frequency readings. Thus, TagBeat first magnifies weak vibration signals without losing their features and then leverages compressive sensing (CS) to recover the high-frequency signals with the low-frequency samplings. To guarantee safety, another work, TagScan [5], utilizes the differences of RF signals when traversing different kinds of liquid to classify them. In this work, a feature that only relates to the liquid material is extracted from RSSI and phase values with a signal propagation and attenuation model.

High-precision RFID localization and tracking

In industrial automation, object localization and tracking is one of the most critical demands. Wireless and battery-free backscattering offers a lightweight and low-cost solution for localization and tracking of the materials in warehouses and products on production lines. Early works achieving a median accuracy of tens of centimeters either rely on RSSI for distance estimation and fingerprint map construction, or calculate the angle of arrival (AoA) for continuous localization. Recent proposals integrate reference tags or antenna arrays to calculate phase changes for centimeter-scale precision. Here we survey the most recent works on RFID localization and tracking, which improve not only the task precision but also the robustness and the practicability of sensing systems [8]–[10].

A recent work, OmniTrack [8], solves the problem of the precision degradation caused by the phenomenon of the antenna polarization when the orientation of a RFID tag changes. To achieve centimeter-level localization and orientation of a mobile tag, OmniTrack models the linear relationship between the tag orientation and the phase change of the backscattered signal. To deal with high-noise and complicated multipath environments and to soften the deployment restricts of antennas, Xiao et al. propose a double-tag system for accurate and robust object localization and tracking [9]. The work demonstrates that the phase difference of closely deployed double tags can effectively exclude the impact of undesired signals such as device noises and multipath interferences. RFind [10] manages to use time-of-flight (ToF) for RFID tag localization. To

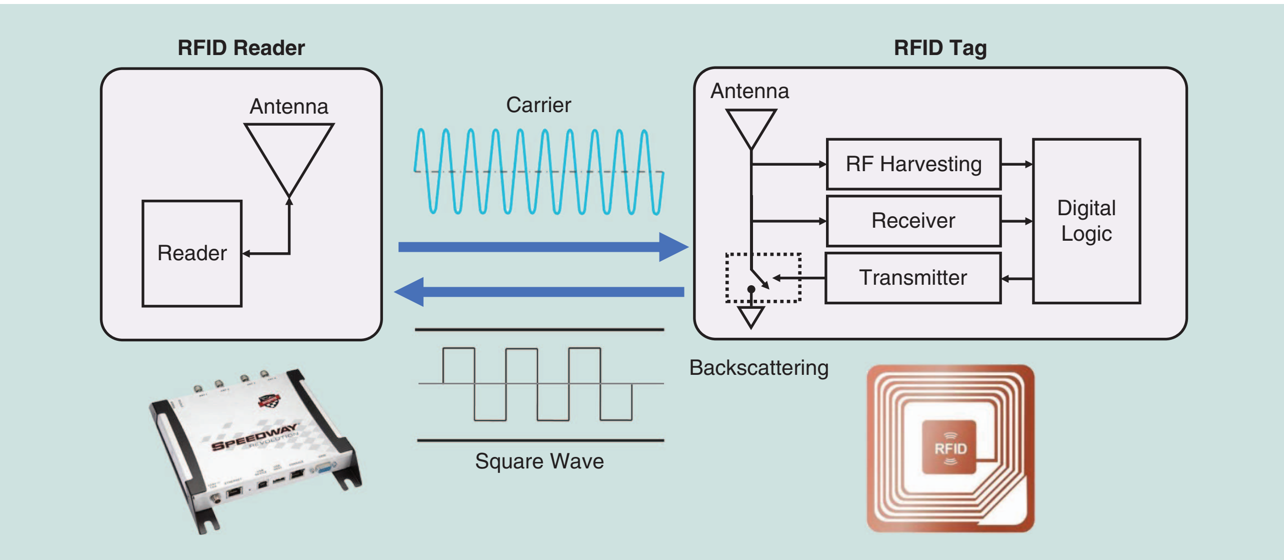

FIGURE 1. Backscattering with RFID systems.

achieve subcentimeter accuracy, a very large bandwidth of multiple gigahertz is often needed in ToF-based methods, which, however, is not compliant with Federal Communications Commission (FCC) regulation and RFID protocols. Thus, RFind generates a virtual ultrawide bandwidth by importing extremely low-power but efficient hopping localization frequencies outside the industrial scientific medical (ISM) band while keeping the normal communication band for powering up commercial tags.

To summarize, battery-free RFID sensing offers a new paradigm that not only can measures specific physical metrics with just wireless signals, but also can provide high-precision results. Except extending physical modalities, improving resolution, timeliness, and reliability of battery-free RFID sensing offers prime candidates for further studies. Besides, new nonintrusive wireless tag systems are increasingly gaining more attention recently, e.g., LiveTag [7] designs multiple metallic structures of a Wi-Fi tag to disturb ambient Wi-Fi channels for information expression. Further, it leverages customized multiantenna beamforming algorithms to sense the human–object interaction. Moreover, we will show our preliminary explorations of designing RFID systems for real-world industrial IoT in the section "Case Study: Pavatar."

Visual sensing from intensive networked videos

Surveillance cameras are one of the most commonly used IoT devices in industrial IoT because the visual sensing provides numerous informative clues. In modern industries, cameras are deployed with a high density to seamlessly monitor the status of machines and the activities of workers. The characteristics of visual sensing in industrial IoT is as follows.

- *Timeliness requirement*: Video analysis usually has a strict requirement of timeliness in modern industries. How to fulfill the real-time processing on resource-limited devices while reducing the transmission latency remains a challenge.
- *Information sparsity*: Camera surveillance systems generate intensive video data, but the spatiotemporal sparsity of significant information needs efficient processing.
- *Seamless cooperation*: Visual clues provided by one single camera is partial and limited, thus seamless cooperation among the networked cameras is desired to perform complicated sensing tasks.

Visual sensing applications on a large-scale camera network need not only the optimized allocation of the computation resources but also the efficiency and the accuracy of the vision tasks. In this section, we first introduce a rising computation paradigm, edge computing for multimedia IoT data processing [11]–[13] and then discuss efficient and accurate video analysis algorithms of resource-constrained embedded devices [14]–[17].

Edge computing for large-scale networked video processing

The networked cameras are expected to cooperate for a comprehensive understanding of the monitoring targets. However, uploading all of the multimedia data stream to the cloud is infeasible due to its limited processing capacity of the cloud, the unpredictable latency induced by the network transmission, and the unaffordable cost of the network bandwidth. Edge computing, a new computation paradigm between embedded computing and cloud computing, performs data processing and analyzing at the edge of networks. Large-scale networked video analytics is considered the killer app of edge computing [11].

Recent practical video analytics systems start adopting edge computing to deal with large-scale networked video, although there has not yet been a universally standard architecture. In [11], a practical system rocket for traffic monitoring in Bellevue, Washington, is proposed to discuss potential prospects of edge computing for the surveillance video processing. Model predictive control is used to allocate limited computation and network resources between the edge servers and the cloud server. A recent edge-computing architecture [12] introduces another offloading mode, where multiple edge servers cooperatively serve one camera and build a performance model with the compression ratio as the input. Then it separates the NP-hard problems of the edge server selection and the compression ratio selection, and solves them with heuristic algorithms. Besides offloading and scheduling, information sparsity can be leveraged to reduce computation costs among resource-constrained edge servers. The recent work ViTrack [13] proposes a spatiotemporal CS algorithm to recover the camera-level trajectories for the monitored vehicles by processing just 1/50 of the raw frames.

Practical video analytics with embedded deep learning

Recent advances in deep learning, especially convolutional neural networks (CNNs), have pushed the boundaries of computer vision. Basically, existing CNN applications purely rely on cloud infrastructures. However, problems such as network transmission delay, expensive but limited bandwidth, user privacy and costs

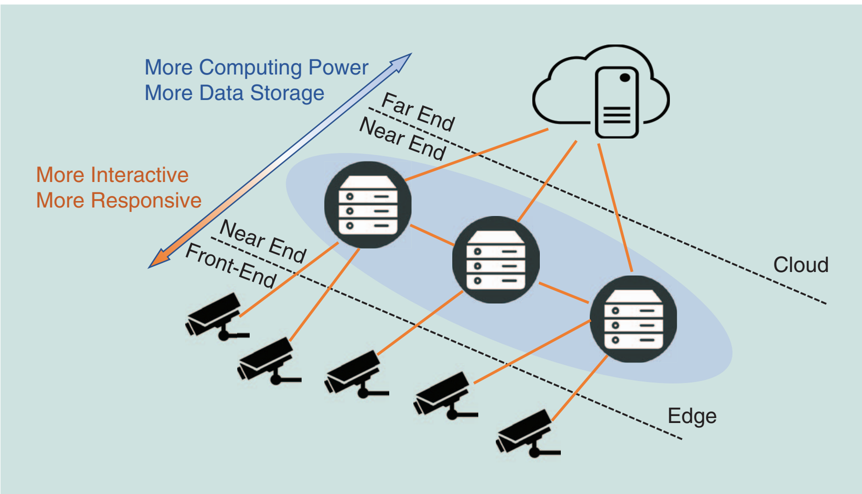

FIGURE 2. Edge computing for large-scale networked video processing.

of high-performance cloud servers make cloud-based solutions infeasible for large-scale video analysis in industrial scenarios. One potential trend to solve this dilemma is to enable real-time deep learning directly on end devices.

In a typical CNN model, convolutional layers that extract features consume much of the executing time because of the window-by-window convolutional operations, while fully connected layers that conduct the classification tasks take up much of the model weights because of the dense connections among neurons. Thus, to satisfy the requirement of the low-latency performance, we can adopt different strategies to optimize different modules, e.g., the structure pruning for the deep models [14], [15] and the runtime optimization of the inference frameworks [16], [17]. Model structure pruning methods such as DyNS [14] and DeepIoT [15] try to eliminate the redundancy in the model parameters through a three-step procedure: importance estimation, parameter pruning, and model retraining. Unimportant parameters are pruned to speed up computation and save storage space. Runtime

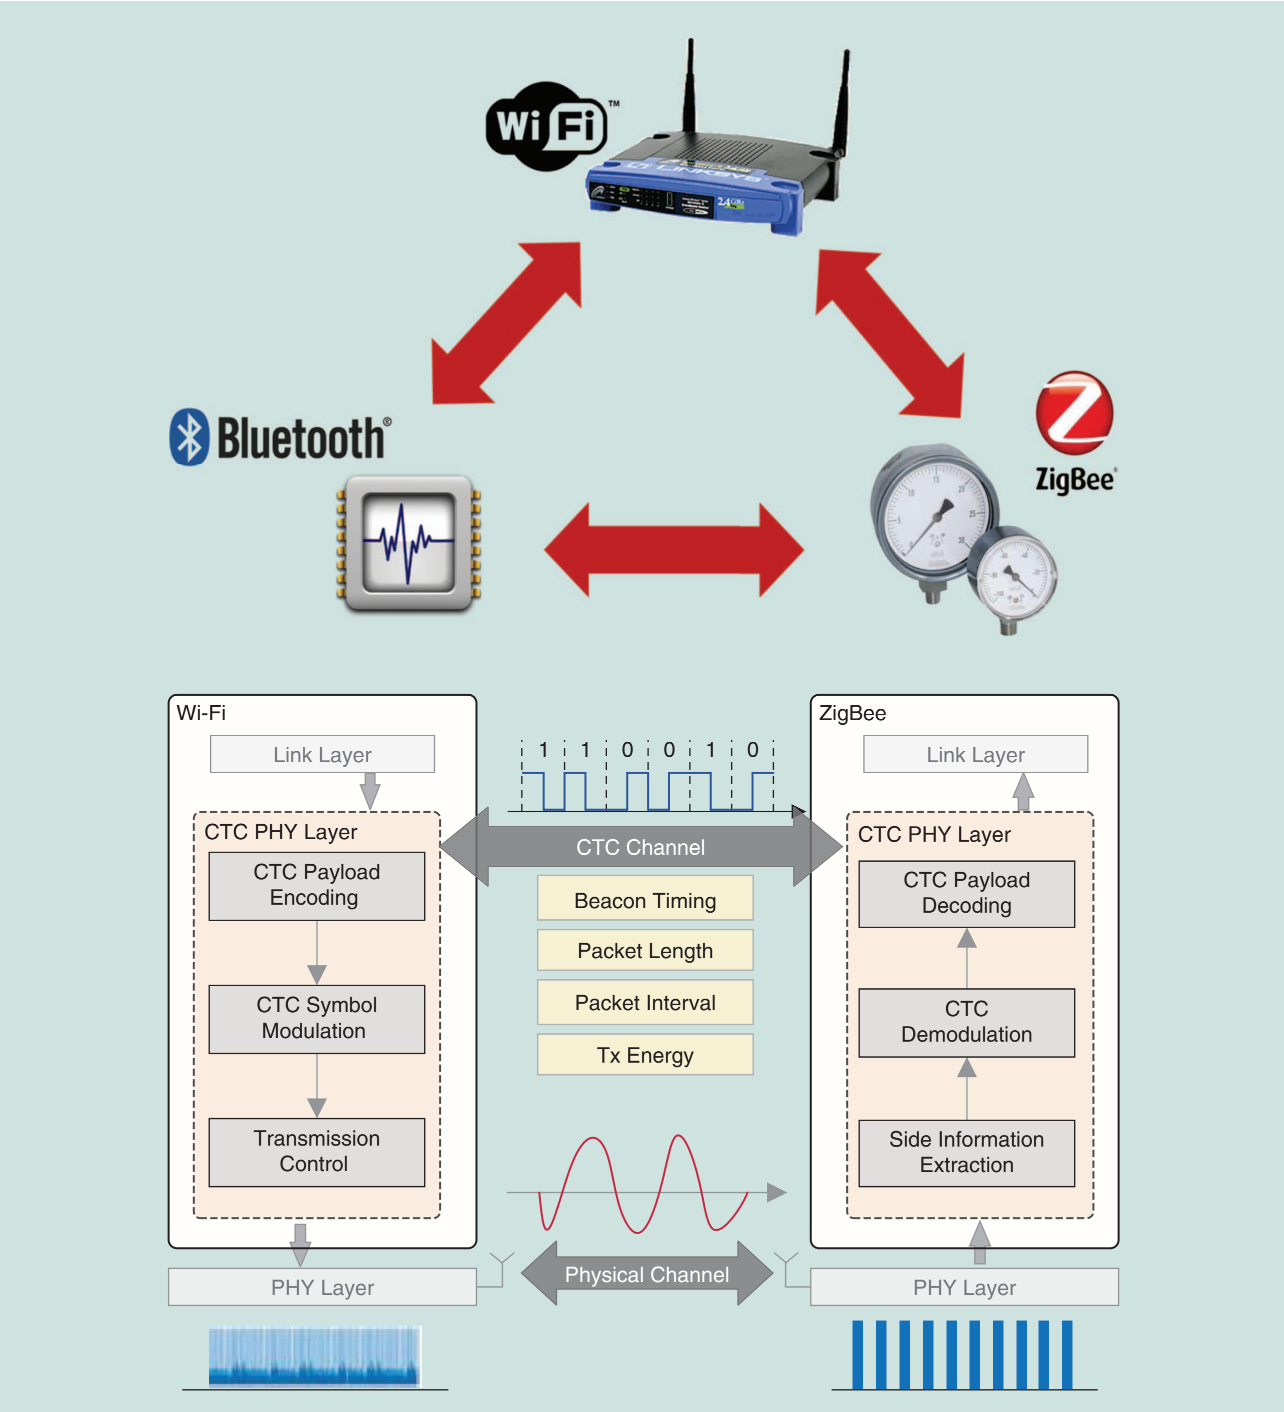

FIGURE 3. The architecture of CTC.

optimization utilizes computation parallelization and pipeline scheduling [16], intermediate-result caching [17], etc. to optimize the computation procedures of deep-learning inference frameworks on IoT devices.

In summary, edge computing is deemed as a promising architecture for practical visual sensing on ubiquitous surveillance cameras, and deep-learning algorithms with amazing analysis capabilities can be tailored to edge devices. In industrial IoT, effectiveness and timeliness are two dispensable, but mutually exclusive, performance indicators. Therefore, to maximize both of them, we believe enabling deep learning at the edge of networks is a promising direction.

Cross-technology heterogeneous wireless communication

In digital twin for smart factories, embedded sensors with various sensing capabilities are networked together to monitor the same area. These sensors might adopt heterogeneous wireless communication technologies, such as Wi-Fi, Bluetooth, ZigBee, and long-term evolution (LTE). The characteristics of heterogeneous networking environment are as follows.

- *Heterogeneous interference*: The majority of popular wireless technologies share the same frequency band, e.g., 2.4-GHz ISM band. Therefore, heterogeneous interferences and collisions are very likely to occur.
- *High-density deployment*: In many cases, networked sensors are densely deployed, which induces nontrivial challenges in collecting data in real time.
- *Interconnecting heterogeneous devices*: Due to the complicated operating states of industrial machinery, multiple devices need to exchange information in suit for a real-time understanding of current states.

Today, how to organize, manage, and cooperate heterogeneous IoT devices is increasingly drawing attention. A simple solution is to deploy a gateway with various radio interfaces for access control and information exchange among heterogeneous devices. Possible communication bottleneck and extra hardware cost drive researchers to explore the direct communication ability among different technologies, thus cross-technology communication (CTC) is proposed. With CTC, heterogeneous devices can directly exchange information for fast and effective control and cooperation, which perfectly satisfies the timeliness and interconnection requirements in industrial IoT.

The basic idea of CTC is that, although heterogeneous wireless technologies can't directly decode the packets from another technology, side-channel information of wireless transmissions, e.g., transmission time, beacon shifting [20], RSSI amplitude [21] etc., can be leveraged to encode bits. These corresponding methods are called *packet-level modulation* because one or more original packets should be transmitted to modulate one bit. Although recent CTC works manage to improve the throughput by deeply exploiting the coexistence environment to encode more bits simultaneously, packet-level modulation still offers a relatively low throughput, compared to the original wireless technologies.

Therefore, a new trend of CTC called *physical-level emulation* tries to emulate the heterogeneous signals directly in the physical layer to achieve a throughput comparable to the original wireless technology, e.g., up to 250 kilobits/s for ZigBee. WEBee [22] is the most representative work that meticulously fills the payload of a transmitted Wi-Fi frame to directly emulate ZigBee frames. The feasibility of the bit emulation is ensured by a redundancy coding technology of ZigBee called *direct sequence spread spectrum* (*DSSS*). The inevitable differences of the emulated chip sequences and the predefined chip sequences can be tolerated by the DSSS symbol matching.

The aforementioned CTC works have successfully established direct communication among heterogeneous wireless technologies in the physical layer. Protocols and applications can be further built upon these infrastructures. A basic scenario is to use CTC packets as a medium access control protocol for the channel coordination in coexisting environments, which is one of the primary intuitions of this emerging technique. ECC [24] introduces a cross-technology clear-to-send signal to negotiate an aggregated white space for better ZigBee communication. Moreover, Crocs [25] leverages CTC to directly synchronize Wi-Fi and ZigBee devices. To achieve a more robust and accurate time synchronization, Crocs first incorporates a short CTC beacon based on Barker code for a more accurate time alignment and then sends tim stamps via CTC transmissions. StripComm [26] applies CTC to a more densely coexisting environment and faces severe challenges with dynamic wireless interferences. To make the energy-encoding CTC more robust against unavoidable wireless interferences, StripComm encodes bits with Manchester coding, and decodes bits after the interference cancellation based on specific signal similarities.

In a nutshell, recent advances in CTC have experienced two stages, from packet-level modulation to physical-layer emulation. The validity and practicability of these approaches have been verified by the throughput comparable to the original wireless in [22]. Enhancing more modern wireless technologies, e.g., LTE and NB-IoT with the ability of CTC, as one of the future directions of CTC, faces the new challenges of the bandwidth asymmetry and the mismatch of the transmission rates. Moreover, facilitating the upper-layer standards and protocols to build cross-technology networks is also a very fascinating topic.

Data analytics

The data analytics layer plays a vital role in industrial IoT to provide smart services. The sensing layer samples raw data of physical metrics, the networking layer conveys data and, finally, the data analytics layer identifies patterns or mines the principles behind. The data analytics in industrial IoT have the following characteristics.

- *Low quality of raw data*: Due to the hardware imperfection or the unreliable wireless transmissions, the raw data generated by IoT devices are usually of low quality, which brings challenges for the accurate analytics.
- *Multisource data*: The data from multiple sensors may be redundant and even contradictory. Obtaining the truth from multisource data desires more advanced signal processing methods.

- *Partially labeled data*: In industrial scenarios, the high-frequency and continuous stream data is very difficult and impractical for manual labeling. Dealing with partially labeled data is also very challenging.

Analyzing IoT stream data with these characteristics is deeply associated with advanced signal processing algorithms, including data cleaning [27], feature selection [28], and event classification and system diagnosis [29].

Anomaly correction of time series data

Anomaly detection (or further anomaly correction of time series data) is an indispensable preprocessing step for upper-layer applications, such as event detection and fault diagnosis. In [27], Zhang, et al. suggest that simply filtering out anomalies will damage the continuity of time-series data, and the intermittent and incomplete time series would possibly affect subsequent classifiers. Different from existing rule-based repairing, e.g., the speed-constraint model and the autoregression model, an iterative minimum repairing (IMR) algorithm based on sparse-labeled ground truth is proposed. The sparse-labeled truth points, which can be obtained by a reliable sensor with a relatively long sampling period or manual labeling, can better fix continuous errors. Rather than sequentially repairing one error point for just one time, an IMR algorithm iteratively adjusts error points until the global convergence.

Data-driven feature selection

The multisource data can be redundant for upper-layer applications. Apart from the guidance of the physical models, data-driven feature selection can improve the final performance. In [28], Li et al. point out that traditional feature selection methods either consider only the informativeness of features regardless of sample labels, or are optimized for some particular classifiers. Hence, they leverage the sample labels and propose a novel information greedy feature filter (IGFF) method that is independent from the classifiers. With rigorous mathematical proofs, IGFF selects the optimal subset of features by maximizing mutual information between the candidate variables and the fault labels. The experiments on the real-world data set about air-handling units of a smart building shows that, regardless of back-end classifiers, IGFF can achieve a much higher improvement in the classification accuracy than the traditional methods and the empirical selection.

Event classification with partial labeled data

Fault detection is an event classification problem that classifies a short time series data from multiple sources into normality or particular faults. Current methods are mainly based on supervised learning. In industrial scenarios, however, the high-frequency and continuous stream data are almost unlabeled. Manual labeling by domain experts means considerable labor costs, which is impractical for real-world systems. In [29], a hidden structure semisupervised machine (HS^3M) is proposed to deal with sparsely labeled industrial IoT data. HS^3M incorporates fully labeled data, partially labeled data, and unlabeled data with a unified-format loss function, thus it can fully utilize all available data sets to learn a more generic model. Tested on an industrial IoT data set of a practical power distribution system, HS^3M can achieve at least 9% gain of accuracy and 10% gain of false positive in comparison to the runner-up method.

In summary, advanced signal processing technologies are indispensable to deal with fallible, multisource, and partially labeled industrial IoT data. Moreover, we believe practical data analytics is deeply associated with the characteristics of the target systems, which will be addressed in the next section.

Case study: Pavatar

In this section, we introduce our early experience with a real-world industrial IoT system, Pavatar [30]. Pavatar is an IoT system for UHVCS management. The UHVCS, built at the hub points of the ultrahigh-voltage power grid, efficiently performs dc/ac transformation of clean energy, e.g., wind, solar, water, and nuclear power. Globally connected UHVCSs are expected to construct the backbone of the Global Energy Internet (GEI), which is deemed to alleviate energy problems such as the exhaustion of fossil fuel, environmental pollution, and supply–demand imbalance. A large rotating machine called a *synchronous compensator* is the core component of an UHVCS. Its critical function is to stabilize the outgoing current by generating or absorbing reactive power, in response to unpredictable voltage fluctuations, and thus ensuring GEI's stability, safety, and reliability. Clearly, proper operation of synchronous compensators is of vital importance to GEI. There have been various conventional solutions for power plant monitoring, e.g., manual checking and video surveillance. However, those solutions are generally inefficient, inaccurate, and costly.

Our team collaborates with the State Grid Corporation of China to launch the Pavatar project in one UHVCS located in Hunan, China. Aiming to build a digital twin of this UHVCS, Pavatar monitors the entire operation process in real time and provides decisions and support for UHVCS administrators. The functionality of Pavatar generally includes the following key aspects:

- Comprehensive sensing of synchronous compensators and their cooling systems, operation environments, and surrounding human activities.
- Heterogeneous data visualization in the form of VR.
- System error prediction, anomaly detection, and root-cause diagnosis.

Figure 4 shows the architecture of Pavatar. Pavatar collects data from both built-in and ambient sensors in UHVCSs. Typical internal sensor readings include temperature, pressure, vibration, rotation, etc., which provide the key metrics for decision making. In the surrounding environment, low-power and battery-free sensors are deployed to sense temperature, humidity, noise, air quality, and liquid leakage, etc., as supplementary information. In addition, networked cameras are deployed to cover walkable areas. The maximum density of sensor deployment is about $50/m^2$, the highest sampling frequency of internal sensors is around 10 KHz, and the total data volume size per day is over 1 TB. The high-frequency and big-volume stream data are collected and transmitted through heterogeneous networks to

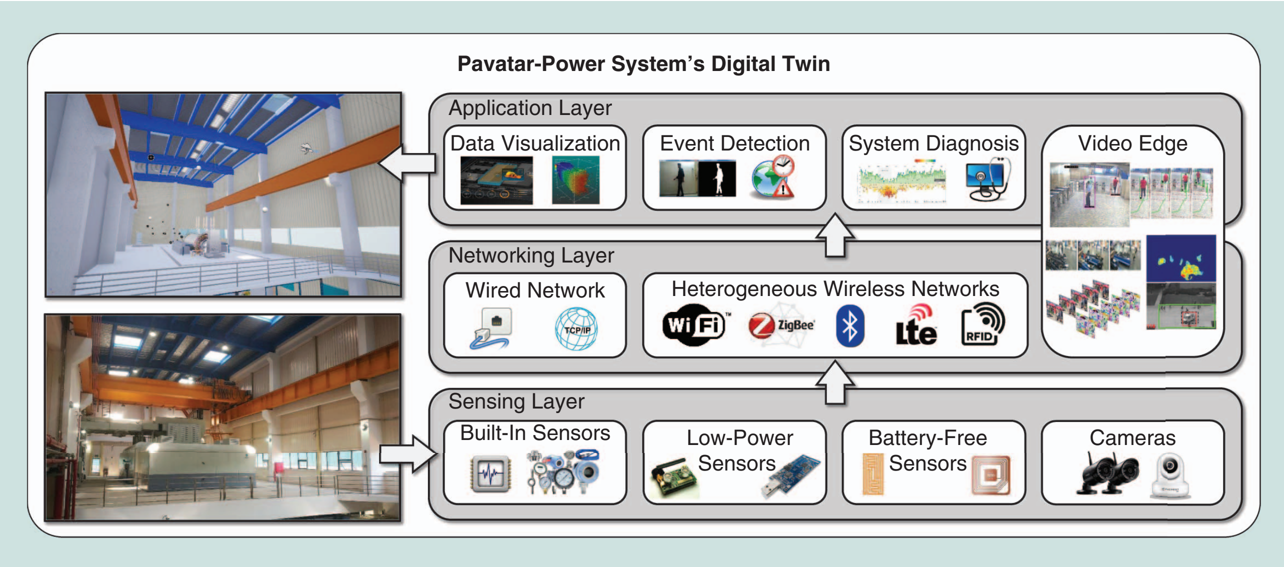

FIGURE 4. The architecture of Pavatar.

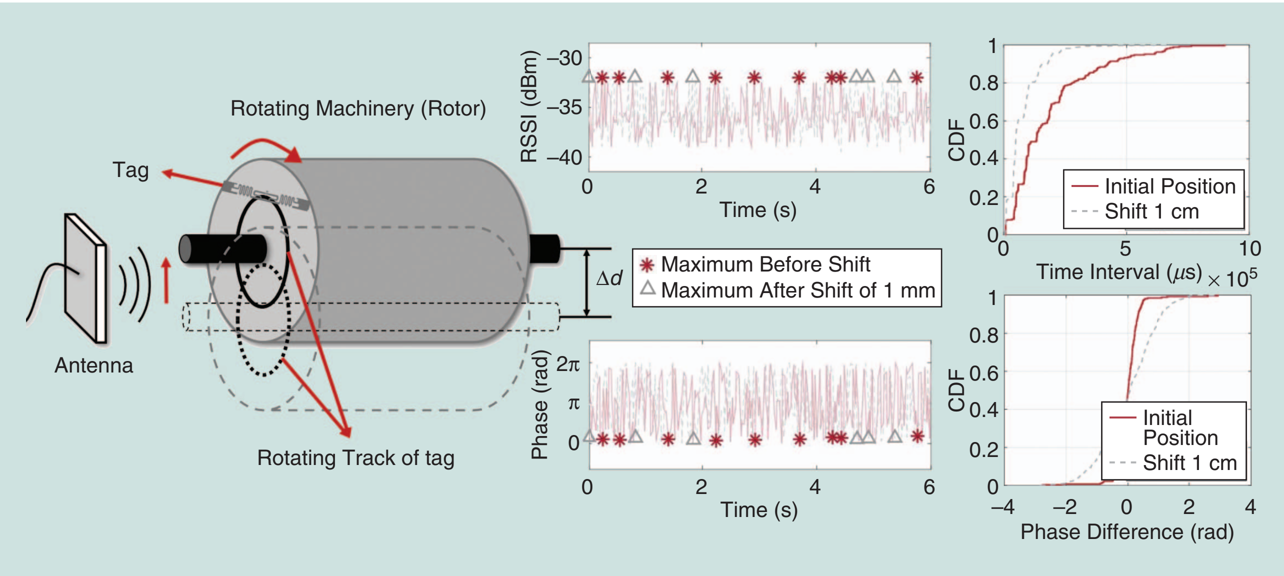

FIGURE 5. Due to the mismatch between the interrogation frequency and the rotation frequency, RED extracts statistic features for eccentricity (centroid shift) detection.

fulfill upper-level applications such as data visualization, event detection, and system diagnosis. Moreover, a three-layer edge-computing architecture is proposed to process massive video data. In the following, we present some of our recent works regarding Pavatar, which leverage advanced signal processing methods to deal with the problems of industrial IoT.

Battery-free sensing for eccentricity detection

Eccentricity, which stands for the displacement of rotating center, is essential for rotating machines, e.g., synchronous compensators in Pavatar. Traditional techniques based on special embedded sensors are either hard to deploy or not practical. Our recent work, RFID-based eccentricity detection (RED), proposes a battery-free RFID sensing system tailored to the classification of the eccentricity status [6]. As shown in Figure 5, RED first extracts features of statistic characteristics e.g., the cumulative distribution functions of the phase difference and the time interval between measured signal peaks, then constructs a Markov model to process stream data without training for a specific environment.

Parallel backscatter transmissions

RFID tags are deployed in Pavatar with the density up to $40/m^2$ for liquid leakage detection. The dense deployments in industrial IoT require new networking techniques for efficient data collection. Thus, we recently proposed a practical system called *FlipTracer* that decodes collided signals to achieve reliable parallel backscatter transmissions [18]. We found that the transition of tag states is usually caused by the discrete signal flip of a single tag. Thus, instead of the direct classification, the states can be inferred by modeling the transition probabilities. As shown in Figure 6, FlipTracer constructs a one-flip graph (OFG) in the in-phase and quadrature (IQ) domain to model the transition patterns and then tracks the OFG to resolve the collided signals. FlipTracer is able to achieve an aggregated throughput of 2 megabits/s, which is six times higher than the existing methods.

Harnessing channel state information for CTC

Compared to Wi-Fi, ZigBee has an orders-of-magnitude smaller maximum transmission power, and a much thinner channel bandwidth. These asymmetries of different communication standards make direct transmissions from ZigBee to Wi-Fi challenging. Our recent work ZigFi leverages channel state information (CSI), an indicator of Wi-Fi channel quality, to enable Wi-Fi to hear low-power ZigBee transmissions [23]. Figure 7 shows, when ZigBee transmissions interfere with Wi-Fi preambles, the changes of CSI amplitude offer a promising encoding space. In ZigFi, a Wi-Fi device decodes bytes by detecting the appearance and the absence of ZigBee signals at specific channels. By dedicatedly training time-series data classifiers, ZigFi can achieve a throughput of 215.9 bits/s, which is 18 times faster than the state of the art.

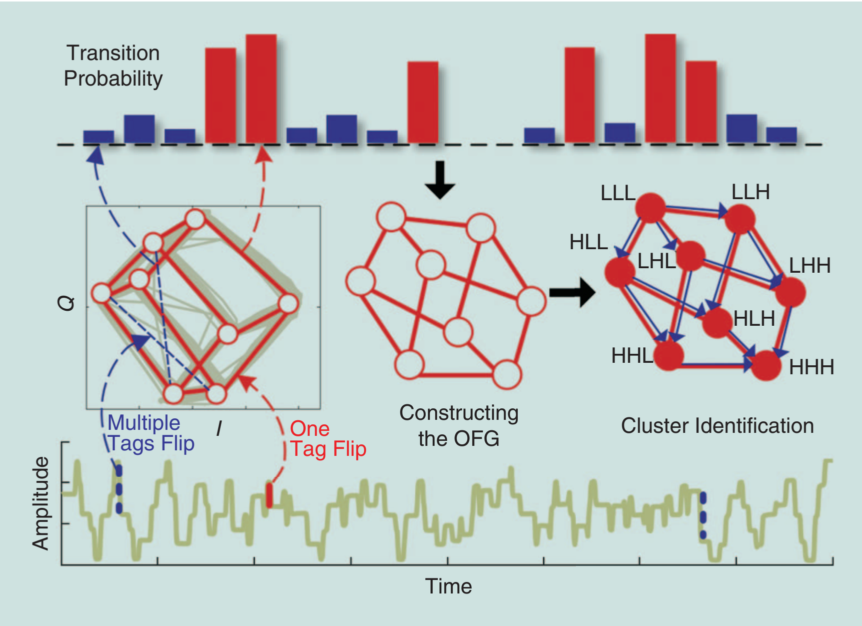

FIGURE 6. The workflow of FlipTracer. FlipTracer first constructs OFG in the IQ domain by selecting edges with large transition probabilities, then assigns each cluster to a parallel bits representation (three tags in this figure), and finally decodes the parallel bits with the OFG.

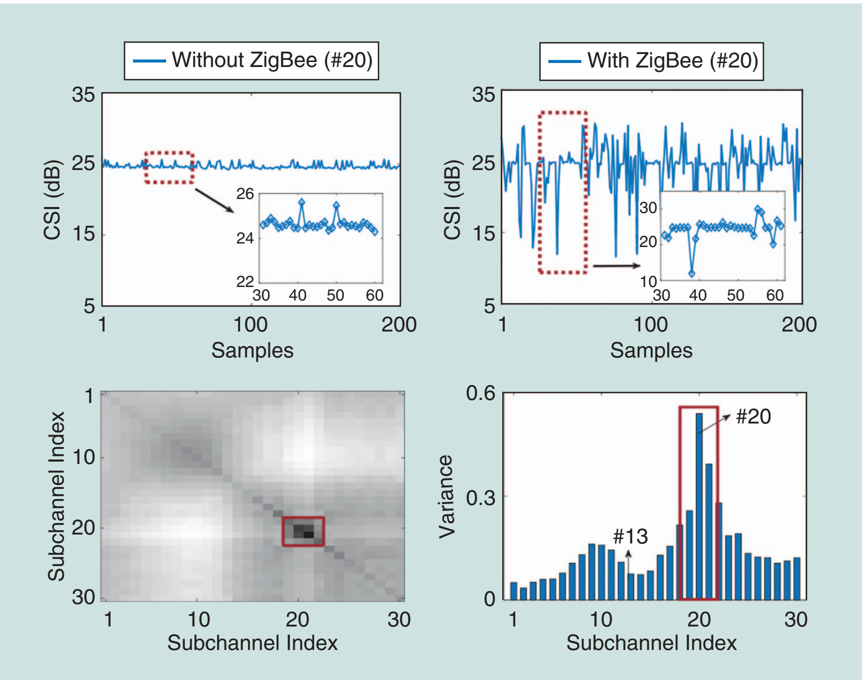

FIGURE 7. ZigFi: CTC from ZigBee to Wi-Fi. ZigFi discloses that ZigBee signals can interfere with Wi-Fi preambles and change the CSI pattern of specific subchannels, e.g., channel 20. Selected subchannels are used for CTC encoding.

Ongoing works

As mentioned previously, deep learning can provide effectiveness while edge computing can offer efficiency. A universal edge-computing architecture for real-time large-scale video analytics is desperately needed in Pavatar. Moreover, data sampling in Pavatar faces a severe problem of the category imbalance, since the anomaly states of synchronous compensators are very scarce. Therefore, modern learning techniques such as online imbalanced and hard sample mining for multisource time-series data can be further tailored to this problem.

Summary and conclusions

In this article, we surveyed and discussed the challenges and recent works toward digital twin, from sensing, networking, to analytics layer. We also presented Pavatar, a real-world IoT system for UHVCSs. We introduced our experience with Pavatar, and discussed the research issues as well as the future directions of industrial IoT. Industrial IoT is of great significance to the innovation of traditional industry. It envisions that we could automatically monitors and comprehensively simulates the factory throughout the entire life cycle, from production and manufacturing, operation to maintenance, to liberate the workforce and provide credible decision supports for industrial operations.

Acknowledgments

This work was jointly supported by National Key R&D Program of China 2017YFB1003000, the research fund from the State Grid of China on "Automatic Surveillance and Remote Diagnosis of Synchronous Compensator," the National Nature Science Foundation (NSF) of China numbers 61772306 and 61672320, NSF of China Key Projects 61432015 and 61632008, and China Postdoctoral Science Foundation number 2016M601034.

Authors

Yuan He (heyuan@tsinghua.edu.cn) received his B.E. degree in computer science and technology from the University of Science and Technology of China, Hefei, in 2003, his M.E. degree in computer software and theory from the Institute of Software, Chinese Academy of Sciences, Beijing, in 2006, and his Ph.D. degree in computer science and engineering from the Hong Kong University of Science and Technology, China, in 2010. He is currently an associate professor with the School of Software and Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing, China. His research interests include the Internet of Things, wireless networks, and mobile and ubiquitous computing. He is a Senior Member of the IEEE and a member of the Association for Computing Machinery.

Junchen Guo (juncguo@gmail.com) received his B.E. degree in computer science and technology from the University of Science and Technology of China, Hefei, in 2016. He is currently pursuing his Ph.D. degree in software engineering with the School of Software, Tsinghua University, Beijing, China. His research interests include the Internet of Things and mobile and ubiquitous computing.

Xiaolong Zheng (xiaolong@greenorbs.com) received his B.E. degree in software engineering from the Dalian University of Technology, China, in 2011 and his Ph.D. degree in computer science and engineering from the Hong Kong University of Science and Technology, China, in 2015. He is currently a postdoctoral research associate with the School of Software, Tsinghua University, Beijing, China. His research interests include the Internet of Things, wireless networks, and ubiquitous computing.

References

[1] P. Middleton, T. Tsai, M. Yamaji, A. Gupta, and D. Ruebe. (2017). Forecast: Internet of things – Endpoints and associated services. [Online]. Available: https://www.gartner.com/doc/3840665/forecast-internet-things–endpoints.

[2] Huawei Technologies Co. Ltd. NB-IoT. (2018). [Online]. Available: http://developer.huawei.com/ict/en/site-iot/product/nb-iot

[3] P. C. Evans and M. Annunziata, "Industrial internet: Pushing the boundaries general electric reports," General Electric Reports, 2012, pp. 1–37.

[4] L. Yang, Y. Li, Q. Lin, X.-Y. Li, and Y. Liu, "Making sense of mechanical vibration period with sub-millisecond accuracy using backscatter signals," in *Proc. ACM Mobile Computing and Networking*, New York, Oct. 3–7, 2016, pp. 16–28.

[5] J. Wang, J. Xiong, X. Chen, and D. Fang, "TagScan: Simultaneous target imaging and material identification with commodity RFID devices," in *Proc. ACM Mobile Computing and Networking*, Snowbird, UT, Oct. 16–20, 2017, pp. 288–300.

[6] Y. Zheng, Y. He, M. Jin, X. heng, and Y. Liu, "RED: RFID-based eccentricity detection for high-speed rotating machinery," in *Proc. IEEE Int. Conf. Computer Communications*, Honolulu, HI, Apr. 15–19, 2018.

[7] C. Gao, Y. Li, and X. Zhang, "LiveTag: Sensing human-object interaction through passive chipless WiFi tags," in *Proc. USENIX Networked Systems Design and Implementation*, Renton, WA, Apr. 9–11, 2018, pp. 533–546.

[8] C. Jiang, Y. He, X. Zheng, and Y. Liu, "Orientation-aware RFID tracking with centimeter-level accuracy," in *Proc. IEEE/ACM Information Processing in Sensor Networks*, Porto, Portugal, Apr. 11–13, 2018, pp. 290–301.

[9] F. Xiao, Z. Wang, N. Ye, R. Wang, and X.-Y. Li, "One more tag enables fine-grained RFID localization and tracking," *IEEE/ACM Trans. Netw.*, vol. 26, no. 1, pp. 161–174, 2018.

[10] Y. Ma, N. Selby, and F. Adib, "Minding the billions: Ultra-wideband localization for deployed RFID tags," in *Proc. ACM Mobile Computing and Networking*, Snowbird, UT, Oct. 16–20, 2017, pp. 248–260.

[11] G. Ananthanarayanan, P. Bahl, P. Bodík, K. Chintalapudi, M. Philipose, L. Ravindranath, and S. Sinha "Real-time video analytics: The killer app for edge computing," *IEEE Comput.*, vol. 50, no. 10, pp. 58–67, 2017.

[12] C. Long, Y. Cao, T. Jiang, and Q. Zhang, "Edge computing framework for cooperative video processing in multimedia IoT systems," *IEEE Trans. Multimedia*, vol. 20, no. 5, pp. 1126–1139, 2017.

[13] L. Cheng, J. Wang, Z. Cao, and Y. Liu, "ViTrack: Efficient tracking on the edge for commodity video surveillance systems," in *Proc. IEEE Int. Conf. Computer Communications*, Honolulu, HI, Apr. 15–19, 2018, pp. 1–9.

[14] Y. Guo, A. Yao, and Y. Chen, "Dynamic network surgery for efficient DNNs," in *Proc. Neural Information Processing Systems*, Barcelona, Spain, Dec. 5–10, 2016, pp. 1379–1387.

[15] S. Yao, Y. Zhao, A. Zhang, L. Su, and T. Abdelzaher, "DeepIoT: Compressing deep neural network structures for sensing systems with a compressor-critic framework," in *Proc. ACM Sensor Systems*, Delft, The Netherlands, Nov. 5–8, 2017, pp. 43–56.

[16] N. D. Lane, S. Bhattacharya, P. Georgiev, C. Forlivesi, L. Jiao, L. Qendro, and F. Kawsar, "DeepX: A software accelerator for low-power deep learning inference on mobile devices," in *Proc. IEEE/ACM Information Processing in Sensor Networks*, Vienna, Austria, Apr. 11–14, 2016, pp. 1–12.

[17] A. Mathur, N. D. Lane, S. Bhattacharya, A. Boran, C. Forlivesi, and F. Kawsar, "DeepEye: Resource efficient local execution of multiple deep vision models using wearable commodity hardware," in *Proc. ACM Mobile Systems, Applications, and Services*, Niagara Falls, NY, June 19–23, 2017, pp. 68–81.

[18] M. Jin, Y. He, X. Meng, Y. Zheng, D. Fang, and X. Chen, "FlipTracer: Practical parallel decoding for backscatter communication," in *Proc. ACM Mobile Computing and Networking*, Snowbird, UT, Oct. 16–20, 2017, pp. 275–287.

[19] J. Guo, Y. He, and X. Zheng. "Pangu: Towards a software-defined architecture for multi-function wireless sensor networks," in *Proc. IEEE Int. Conf. Parallel and Distributed Systems*, Shenzhen, China, Dec. 15–17, 2017, pp. 730–737.

[20] S. M. Kim and T. He, "FreeBee: Cross-technology communication via free side-channel," in *Proc. ACM Mobile Computing and Networking*, Paris, France, Sept. 7–11, 2015, pp. 317–330.

[21] X. Guo, X .Zheng, and Y. He. "WiZig: Cross-technology energy communication over a noisy channel," in *Proc. IEEE Int. Conf. Computer Communications*, Paris, France, May 1–4, 2017, pp. 1–9.

[22] Z. Li and T. He, "WEBee: Physical-layer cross-technology communication via emulation," in *Proc. ACM Mobile Computing and Networking*, Snowbird, UT, Oct. 16–20, 2017, pp. 2–14.

[23] X. Guo, Y. He, X. Zheng, L. Yu, and O. Gnawali, "ZigFi: Harnessing channel state information for cross-technology communication," in *Proc. IEEE Int. Conf. Computer Communications*, Honolulu, HI, Apr. 15–19, 2018, pp. 1–9.

[24] Z. Yin, Z. Li, S. M. Kim, and T. He, "Explicit channel coordination via cross-technology communication," in *Proceedings of ACM Mobile Systems, Applications, and Services*, Munich, Germany, June 10–15, 2018.

[25] Z. Yu, C. Jiang, Y. He, X. Zheng, and X. Guo, "Crocs: Cross-technology clock synchronization for Wifi and Zigbee," in *Proc. Embedded Wireless Systems and Networks*, Madrid, Spain, Feb. 14–16, 2018.

[26] X. Zheng, Y. He, and X. Guo, "StripComm: Interference-resilient cross-technology communication in coexisting environments," in *Proc. IEEE Int. Conf. Computer Communications*, Honolulu, HI, Apr. 15–19, 2018, pp. 1–9.

[27] A. Zhang, S. Song, J. Wang, and P. S. Yu, "Time series data cleaning: From anomaly detection to anomaly repairing," in *Proc. Very Large Data Bases Endowment*, Munich, Germany, Aug. 28–Sept. 1, 2017, pp. 1046–1057.

[28] D. Li, Y. Zhou, G. Hu, and C. J. Spanos, "Optimal sensor configuration and feature selection for AHU fault detection and diagnosis," *IEEE Trans. Ind. Informat.*, vol. 13, no. 3, pp. 1369–1380, 2017.

[29] Y. Zhou, R. Arghandeh, and C. J. Spanos, "Partial knowledge data-driven event detection for power distribution networks," *IEEE Trans. Smart Grid*, 2017. doi: 10.1109/TSG.2017.2681962.

[30] Tsinghua University (2018). Pavatar project. [Online]. Available: http://tns.thss.tsinghua.edu.cn/sun/pavatar.html

SP