

Rethinking Wireless Network Management Through Sensor-driven Contextual Analysis

Shazal Irshad, Eric Rozner

University of Colorado Boulder

{shazal.irshad,eric.rozner}@colorado.edu

ABSTRACT

Wireless network management is important to ensure the performance, utilization, allocation, and robustness of the network is optimized. Until now, wireless network management has typically been dictated by in-band information, such as wireless measurements, client locations, or even device state. This position paper explores fundamental new ways to manage the network by utilizing out-of-band data provided by a rich deployment of sensors. Out-of-band data can capture information about the users, objects, or environments associated with a network device, meaning that richer contextual policies can be implemented in the network. We propose an architecture called SenseNet, which builds upon three recent trends: (1) the massive deployment of sensors today, (2) the existence of deep-learning algorithms to glean meaningful insights from the sensory data, and (3) the provisioning of edge computing resources to provide real-time actuation of new sensor-based policies. A brief evaluation shows the feasibility and motivates SenseNet and afterwards challenges towards practical deployment are discussed.

CCS CONCEPTS

- Networks → Network architectures.

KEYWORDS

Network management, wireless networking, edge computing, user context

ACM Reference Format:

Shazal Irshad, Eric Rozner and Apurv Bhartia, Bo Chen. 2020. Rethinking Wireless Network Management Through Sensor-driven Contextual Analysis. In *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications (HotMobile '20)*, March 3–4, 2020, Austin, TX, USA. ACM, New York, NY, USA, Article 4, 6 pages. <https://doi.org/10.1145/3376897.3377863>

1 INTRODUCTION

Today's environments are outfitted with more sensors than anytime in history. On the personal level, wearables and phones monitor user characteristics such as activity, health, and emotion. In users' homes or workplaces, cameras, proximity sensors, smart digital

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotMobile '20, March 3–4, 2020, Austin, TX, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7116-2/20/03...\$15.00

<https://doi.org/10.1145/3376897.3377863>

Apurv Bhartia, Bo Chen

Cisco Meraki

{apurv,bochen}@meraki.com

assistants, and other IoT devices are readily deployed. Still more sensors exist on the city-scale, monitoring weather, transportation, and implementing city-wide CCTV. A large variety and deployment of sensors enables technologies to adapt to user context or environmental factors. Until this time, however, most research has focused on adapting user applications or services based on context. In this paper, we investigate how a rich set of deployed sensors can allow administrators to more flexibly manage their network. We introduce fundamental new quality-of-service provisioning based on external sensor data, and outline many other mechanisms to integrate sensor readings directly into network management.

Consider the following motivating scenario. After the Boston Marathon bombing cellular service was quickly overwhelmed. Imagine video monitoring infrastructure deployed around the bombing area could be utilized to understand which users may actively be in distress, or simply infer which users are near the blast radius. Armed with this data, the cellular provider could then implement access control based on the urgency of its users, perhaps preventing an overloaded network from denying service to victims or responders.

The above example highlights the potential contributions of this work. Until now, most Quality-of-Service (QoS) decisions were a function of *in-band* network data, that is information obtained from wireless links (e.g., link quality, interference, etc) or devices attached to the network (e.g., application-level flows, device location, etc). And while management functionality has recently been adapted to certain classes of users (such as high-paying users or users that have exceeded data limits) or types of applications (e.g., bandwidth limiting certain apps), such data still resides within the network operator. In contrast, there has been no way to consider *out-of-band* information such as a current user's emotions, activities, context, or even dress in network management decisions. In other words, this paper aims to understand how information that cannot be gleaned from the network itself can be used to make better and more flexible network management decisions.

For the camera-based example to become reality, computer vision algorithms must be robust and mature enough to realize administrator demands and preferences. Luckily, recent trends in deep learning have provided a variety of algorithms that can achieve human-level accuracy on a variety of tasks such as image recognition, object detection, activity and facial recognition, and pose estimation. And while these algorithms have traditionally been computationally expensive, advances in both machine learning (building more lightweight models) and systems research (allocating resources and efficiently configuring computer vision architectures for specific tasks [13]) show promise in running a set of computer vision algorithms on a small edge cluster.

Running computer vision algorithms on a small edge cluster enables responsive allocation and adaptation of wireless resources.

Network decisions may require tight timescales, and thus round-trip times to the cloud might not be easily afforded. Today, there is significant momentum in deploying edge-based compute, which is typically co-located with access networks. For example, HP, Microsoft, and AT&T have recently invested in edge computing deployments. Edge computing also provides a natural home for network slicing, function virtualization, other trends such as 5G to enable more programmable networks. A challenge, however, is dealing with the limited capacities of edge environments, while still providing low-latency, high-quality sensor analysis over a large class of users or devices.

In this position paper, we describe a system called SenseNet to realize our goals. First, we define and explore fundamental new ways to manage wireless networks by analyzing out-of-band sensory data. We enumerate the benefits of this new technology by showing how it can increase spectrum efficiency, provide new ways to enhance public safety, and enable new revenue streams (Section 3.1). Next, we outline a programmable framework allowing administrators to configure this new type of network architecture (Section 3.2). Afterwards, we highlight some preliminary results to motivate our vision. Last, challenges to realize the full potential of SenseNet are discussed in Section 5.

2 BACKGROUND

This section shows how SenseNet differs from and can co-opt prior work. Please note a comprehensive listing of previous work is not possible due to space constraints.

Context-aware applications and systems Prior research in this area is vast ([3] and references therein) and much of the research aims to understand how applications and services can adapt to a user’s context. Going beyond application-based contextual awareness, some schemes propose resource-aware contextual awareness. These schemes vary from external resource discovery and usage [7], to smart environment adaptation, to anticipatory computation [23]. Largely different from the focus of previous works, this paper aims to extend the benefits of contextual awareness to a new domain: wireless network management.

Wireless network management Briefly, network management considers such problems as spectrum assignment [27], interference modeling [11, 16, 26], roaming [4], and scheduling [28]. Users or applications may be managed differently based on policy (e.g., application rate-limiting or user throttling). These previous works commonly manage the network based on in-band state: state easily gleaned from the network or applications running on the device. In this paper, we expand the notion of network management to out-of-band state, allowing for fundamental new ways to contextually manage the network by integrating sensory and other data external to the network or its devices. While preliminary works have studied context-aware network management [24], they infer contexts of the *application* residing on the device (e.g., whether an app is in the foreground or not) to make network decisions. Our paper investigates a fundamental new way to contextually manage the network by examining the *user, object, or environment associated with the wireless device*.

Edge-based video processing As cameras are a widely deployed sensor, video analysis approaches are relevant to SenseNet design. Chameleon [13] is an example of an edge-based video analytics framework, designed to adapt configurations (e.g., input video resolution, frame rate, and model selection) based on edge compute or network availability. Schemes such as Mainstream [12] conserve edge resources by performing computation sharing across vision models. Schemes such as the above give promise to edge-based video processing.

Optimized deep learning Early deep learning architectures contained many parameters and thus were expensive to train and run inference. And while advances in deep learning continue to provide state-of-the-art accuracy on a variety of vision-based tasks, significant progress is being made in making deep learning architectures more portable and light-weight through techniques like reducing precision, designing models with fewer parameters, and sparsifying networks. SenseNet can judiciously utilize advances in lightweight deep learning to sense more accurately from the edge.

3 SENSOR-BASED NETWORK MANAGEMENT

Below we show how utilizing out-of-band sensors to manage the wireless network opens many areas of study. First, we enumerate how new quality-of-service models can be supported in SenseNet. Then we propose a programmable framework to easily configure SenseNet environments.

3.1 New models of QoS

Below, we detail how utilizing out-of-band sensory information enables new QoS models that can make wireless utilization more efficient, contribute to social good, and provide new revenue streams.

More efficient wireless utilization As IoT, mobile health, cyber-physical systems, and smart environments proliferate, wireless spectrum demand continues to grow at a significant rate. And while techniques have been developed to increase capacities (e.g., full-duplex communication [8], small cells [21], advances in MIMO [10, 15, 25]), there exists an untapped potential to ensure increased spectrum demands can be supported in the future. The amount of *useful* data transmitted over the wireless medium can be increased by scheduling traffic in an intelligent and efficient manner. For example, consider a scenario where the administrator wants to provide more network bandwidth to users actively using their device, or simply limit the bandwidth of certain applications (like YouTube) when users aren’t actively engaged with their phone. Computer vision techniques such as gaze following or pose estimation can analyze video feeds to determine when users are interacting with their device. When a user is actively engaging their device (by perhaps looking at the screen), higher QoS can be provided than when the user is not actively engaging their device (not actively looking at the screen). Alternatively, traffic can be buffered while users are not actively using their devices and then sent in bursts before usage occurs. Bursting has been shown to increase efficiency as aggregation can amortize MAC layer overheads and can also preserve energy by avoiding frequent tail latencies.

A second way to drive more useful data across a set of customers is to enable better management policies. In traditional wireless networks, management is performed on a device-level, but in SenseNet several devices can seamlessly be abstracted into one logical entity. In most networks today, individual devices are given a fair share of network bandwidth, meaning a single user with lots of devices may obtain significantly more bandwidth than a co-located user with a single device. With SenseNet, devices can be associated to a specific individual, and then those devices can be roped into a single quality-of-service class that is defined on a per-person level. This allows administrators more control over bandwidth allocation and can perhaps preserve bandwidth across users in a more fair way, allowing more users to be satisfied within a given channel.

Out-of-band integration for social good Sensors are already commonly deployed to benefit society today. Facial recognition tracks terrorists, enhances airport security, and is used by law enforcement and border patrol. With SenseNet new mechanisms can be deployed to increase public safety. For example, an unattended bag at an army base can be correlated with the location of wireless devices to determine if any wireless device is likely to reside in the bag. If so, our scheme can integrate with Wi-Fi or cellular networks servicing the area to cut off all network traffic to the specific device, preventing wireless devices to act as detonators.

Additionally, SenseNet can utilize facial recognition to track or limit the operational efficiency of criminals or terrorists. SenseNet furthers the state-of-the art by managing the wireless devices associated with targeted individuals. With SenseNet, a suspect's temporary "burner" phone can be automatically and continuously identified by associating any new phone to out-of-band facial identification. Then, authorities can be notified or advanced traffic monitoring techniques can be initiated. Such approaches enhance tracking when the individual exits camera vantage points by continuing to track the temporary device via wireless location.

Last, SenseNet can also act as an additional deterrent against unauthorized access in enterprise networks by identifying and penalizing adversaries indulging in denial-of-service, evil-twin [2], and man-in-the-middle attacks.

Enabling new revenue streams By utilizing out-of-band information, wireless providers can create new revenue streams into their service offerings. In the simplest case, wireless providers may opt to enhance the wireless performance of certain customers, citizens, or devices. A clothing company may pay a wireless service provider to offer free wireless services, or higher-quality wireless service, as part of a high-profile marketing campaign in which users are rewarded for wearing certain brands. This idea can easily be generalized, where shops, hotels, or airlines provide customized, higher-end wireless experiences to a special set of valued customers. Such experience can easily be integrated into commonly deployed facial recognition algorithms that exist in retail today [29].

If personal-based services are less attractive, perhaps due to privacy concerns, then there exists many other revenue opportunities that preserve privacy. For example, consider a cellular network that provides better service to individuals within a certain class of vehicles. City municipalities and even some companies provide services for green vehicles (such as free parking or complimentary charging). With SenseNet, cellular network service can hinge on

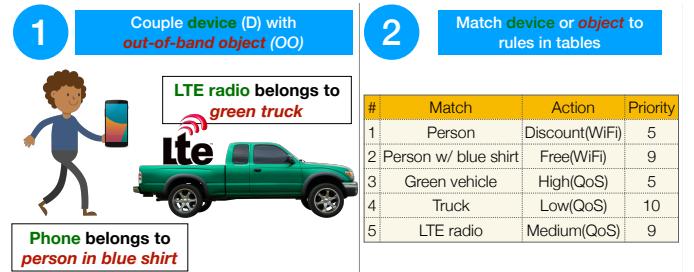


Figure 1: Example of match-action rule set

non-traditional entities like vehicle type. Cameras are already being deployed to count vehicles on roadways, and thus SenseNet can easily co-op this infrastructure. When companies or governments sponsor these type of promotions, service providers can obtain additional revenue.

3.2 A programmable framework to enable sensor-based network management

SenseNet provides a fundamental new way to manage wireless networks, and thus new mechanisms must be deployed to configure, program, and manage sensor-based policies.

Configurable rule set to institute policy Given SenseNet allows for out-of-band information to influence network management decisions, administrators must be provided with mechanisms to program and configure their networks in a seamless way. To facilitate such management, a configurable rule set can implement flexible policies. We take inspiration from software-defined networking technology like OpenFlow [20], and define policies by a set of match-action rules. Figure 1 shows an example. First, the wireless device (denoted D) must be coupled with its associated out-of-band object (OO). In the figure, the phone (D_1) is coupled with the person wearing a blue shirt (OO_1) and the LTE radio (D_2) is embedded in a green pickup truck (OO_2). The second part of the figure shows a simple rule table. In the table, the $\langle D_1, OO_1 \rangle$ tuple matches rules 1 and 2. Since rule 2 has higher priority, the user is granted free Wi-Fi. For the $\langle D_2, OO_2 \rangle$ tuple, rules 3, 4, and 5 match. Here, rules 3 and 4 match on the object (vehicle-based rules) and rule 5 matches on the device (LTE radio). Since rule 4 has the highest priority of the matched rules, the LTE radio in the green truck is put on a low quality-of-service class.

Action	Description
QoS Class	Move device to new QoS class
Bandwidth	Adjust bandwidth allocation to device
Associate	Force client to another AP; disassociate client
Billing	Assign specific billing profile to traffic
Alert	Monitor device traffic; alert administrator
PHY layer	Adjust PHY: rate, antenna configuration, etc.
MAC layer	Adjust MAC: aggregation levels, RTS/CTS, etc.
Scheduling	Buffer/burst traffic according to device state
Inspect	Start deep packet inspection (DPI) on device traffic

Table 1: Sample rule actions

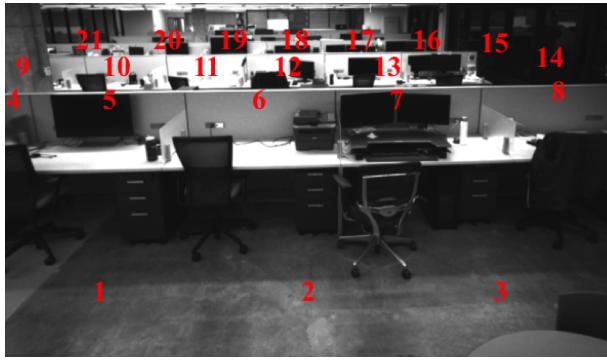


Figure 2: Lab setup with actual points

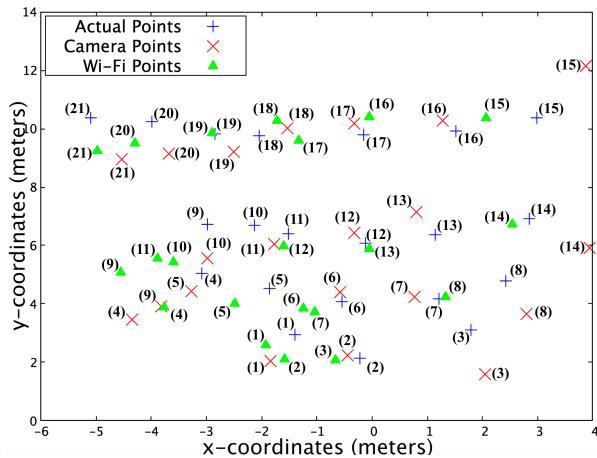


Figure 3: Feasibility study: Multi-domain localization

In practice, rule tables can be much more expressive, and like in OpenFlow, multiple tables can be chained together to form complex policies. The set of actions can vary significantly, and Table 1 outlines many possible examples. Just as the actions can vary significantly, so can the matching algorithms. Administrators can configure matches based on in-band data (IP addresses, customer IDs, etc), in-band metadata (device type, location), or characteristics of the out-of-band objects that the wireless device is coupled with.

4 EVALUATION

In this section we explore the feasibility of deploying a practical solution, discuss how SenseNet can improve localization, and highlight some initial results to motivate SenseNet.

Feasibility study An important first step to deploy SenseNet is aligning users in the sensory domain (*i.e.*, video domain) with the devices detected from the wireless domain. In the sensory domain, video will capture context about a specific user. In the latter, an individual may be identified by an address (*e.g.*, MAC, IP, or IMSI). While there has been some preliminary work on aligning users in the sensor domain and the wireless domain [6, 19, 22, 31, 34], these works typically require explicit client participation, additional

on-device sensors such as inertial sensors, movement or trajectory tracking, or additional infrastructure such as multiple cameras or BLE beacons. To understand the feasibility of aligning users in the wireless and sensory domain in a more practical manner, we analyze the performance of an architecture without the aforementioned constraints.

We conduct a simple experiment with three Google Wi-Fi APs and an Intel d435i depth camera deployed in a rectangular 1000 square foot lab space (see Figure 2). One of the APs (coordinates (0,0)) is co-located with the camera that looks into the room, and both are situated along the same wall. The other two APs (coordinates (-7.3,14.9) and (3.1,14.9)) are situated at the half-way point of the adjacent walls. A user moves around the room with a Google Pixel 3 phone, and at known points the user is localized in the Wi-Fi and camera domains. The Google WiFiRttScan [1] app¹ on the phone calculates time-of-flight to the APs using 802.11mc, which is run for about 5 seconds. Then, well-known trilateration techniques localize the user. The Intel depth camera outputs coordinates when selecting the user in the field-of-view. Figure 3 shows the camera and Wi-Fi localization techniques mapping to the same real-world coordinates, with the blue points representing the ground truth. For example, a user stands at position (1.21,4.19) for data point 7 (see Figure 2), the camera estimates the user's position to be (0.77,4.23), and the Wi-Fi estimates the position to be (-1.05,3.72). We find the error in Wi-Fi estimate for point 7 arises from an inaccurate time-of-flight estimate from the furthest AP.

Next, we understand how well this system can accurately pair the person's position in camera's view to the corresponding wireless location of the phone. Since the Wi-Fi and camera estimates are mapped to the same coordinate system, we perform pairing with the well-known Hungarian Method for assignment (as in [30]), which aims to minimize the total distance across all camera and Wi-Fi pairings. Over the 21 total data points, only 2 points are mismatched, indicating a fundamental primitive of SenseNet (marrying wireless and sensory data) is feasible without many extra constraints. Future work will need to investigate the performance of such pairings under more challenging conditions (*e.g.*, less AP coverage, less client traffic, larger areas, cellular, denser deployments, dynamic scenarios, non-depth camera, etc).

Improving localization We briefly provide insights into how wireless and camera location data can be combined to provide better overall localization accuracy. Compared to the ground truth, the average error in camera estimation is 1.08m, and the average error in Wi-Fi localization is 1.29m (graphs omitted due to space). We notice, however, the error of the camera-based technique increases at the boundaries of the camera's field-of-view (FoV), specifically at actual points 1, 3, 4, 8, 9, 14, 15, and 21 in Figure 2. In such cases, localization calculation based on combined wireless and camera info may increase the overall accuracy. For example, a simple combination scheme that uses Wi-Fi localization when objects are near the edge of the camera's FoV, and the camera localization otherwise, performs 15.2% better overall than the camera-only technique and 29% better than the Wi-Fi only technique. Future work is required to formalize the approach (a weighted average over camera and

¹Explicit client-side participation is not needed, as long as device supports 802.11mc. Ideally a software patch on the AP would suffice.

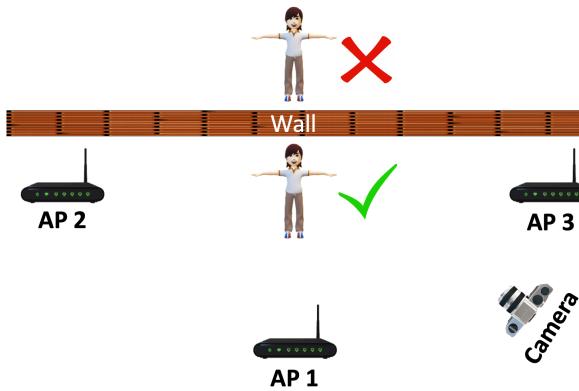


Figure 4: Geofencing experiment

Wi-Fi outputs seems promising), but the basic idea is wireless and camera information can be combined intelligently to provide more accurate localization estimates.

Geofencing To motivate a simple SenseNet use case, we perform a geofencing experiment as shown in Figure 4. Administrators may wish to define a geofencing policy in which users are able to access the Wi-Fi only when inside the building. In this case, traditional trilateration techniques can not guarantee the confidence of the result if a given user is close to the wall as time-of-flight information is similar. Other problematic cases may occur when APs are positioned along the same wall in a line or when not enough APs are deployed to perform accurate trilateration. With SenseNet, the camera can determine if a given user is within the geofence, and if so, the APs can grant wireless access. Our simple experiment was able to differentiate the indoor and outdoor cases, allowing for better wireless network management. Note that depending on camera arrangements, more fine-grained policies such as enabling access control based on specific room occupancy can be enabled. Such policies are useful in enterprises or other highly-controlled environments.

5 DISCUSSION AND CHALLENGES

While the piece-parts to enable SenseNet exist today, several challenges must be overcome to realize the full potential of sensor-based network management in practice.

Security and User privacy This paper does not explicitly consider user privacy in context-aware systems. Most sensing systems are deployed in public, which may limit concerns. To combat issues with privacy, SenseNet can rely on previous work. For example, intelligent brokers can give users ownership or preserve their context [3], or different contextual data can be shared based on its level of detail. More work is needed, however, to better understand the privacy ramifications of sensor-based network management. On the security side, work on trustworthy sensors [9, 14, 18, 32, 33] can ensure sensor data is valid.

Injecting and running models at the edge Section 3.2 outlines a programmable framework for SenseNet. Network administrators need not be experts in deep learning, and hence tools should be

provided to select and configure models that will implement a given rule set automatically. Functionality may be split between the edge and the cloud, and such partitioning should be mostly abstracted away from the administrator. Alternatively, tools can be developed that ensure rule sets are valid for a given set of deployed models at the edge or provide feedback on how expensive certain rule sets may be to deploy. In short, there exists many opportunities to providing tooling to ensure models are deployed automatically and efficiently at the edge.

Expanding the scope of sensors Future work can explore how different types of sensors can be integrated into network management. Video-based sensors are part of a larger class of *environmental sensors* such as microphones, occupancy sensors, or other smart infrastructure readings. In addition, *personal sensors* may provide health, emotional, or other contextual data such as activity recognition. Wireless providers could partner with wearable companies to allow users to opt-in to share their data in exchange for better or cheaper wireless service. There are many interesting opportunities to integrate extremely rich personal data into network management. For example, in an emergency scenario (say flooding or wildfire), resource assignment can be based on current stress or heart rate levels of its users. Alternatively, new ideas like determining user's emotions [17] and optimizing user's happiness within the network can be studied.

Mitigating sensor analysis costs Many deep learning techniques are computationally expensive, even when performing inference. With an expanded scope of sensory inputs, as well as proliferation of high-fidelity sensors (e.g., high-resolution, 360° videos with dense user populations), care must be taken to not overwhelm limited edge resources. Network management may require extremely tight timescales, which further intensifies requirements. As outlined in Section 2, progress is being made from the machine learning and systems community, but current techniques still generally suffer from low frame rates even on low resolution inputs.

Heterogeneous deployments Sensor and edge computing deployments are likely to be very heterogeneous in nature and hence solutions must generalize. Akin to the goal of P4 [5], ideally administrators can simply define a policy, and then a “compiler” can determine where and how functionality should be run. For example, some cameras may have on-board compute or built-in accelerators, which means dedicated edge servers may be relieved of some computation. As with P4, open standards must be created to allow components in the ecosystem to inter-operate, communicate, and work together in an effective and holistic manner.

6 CONCLUSION

This paper proposes a new way to manage wireless networks by analyzing widely deployed sensor data. Examples are given in which out-of-band sensor data can enrich network management, providing new techniques to increase usable bandwidth, implement new policies, and drive new revenue streams for wireless providers. Gaining meaningful insights from sensor data requires state-of-the-art machine learning algorithms, which must run at the network edge to ensure fine-grained network-based functionality can be enabled. Such algorithms, however, typically consume nontrivial

amounts of resources, and thus techniques must be devised to lessen their overheads. Finally, a brief evaluation shows the feasibility and further motivates our scheme.

7 ACKNOWLEDGEMENTS

We thank the anonymous HotMobile reviewers for their feedback. This work is partially funded by NSF-1908910.

REFERENCES

- [1] [n. d.]. WiFiRttScan App - Apps on Google Play.
- [2] Bandar Alotaibi and Khaled Elleithy. 2016. Rogue Access Point Detection: Taxonomy, Challenges, and Future Directions. *Wirel. Pers. Commun.* 90, 3 (Oct. 2016), 1261–1290.
- [3] Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. 2007. A Survey on Context-aware systems. *INTERNATIONAL JOURNAL OF AD HOC AND UBIQUITOUS COMPUTING* 2, 4 (2007), 2004.
- [4] Apurv Bhartia, Bo Chen, Derrick Pallas, and Waldin Stone. 2019. ClientMarshal: Regaining Control from Wireless Clients for Better Experience in Enterprise WLANs. In *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom '19)*. ACM.
- [5] Pat Bosshart, Dan Daly, Glen Gibb, Martin Izzard, Nick McKeown, Jennifer Rexford, Cole Schlesinger, Dan Talayco, Amin Vahdat, George Varghese, and David Walker. 2014. P4: Programming Protocol-independent Packet Processors. *SIGCOMM Comput. Commun. Rev.* 44, 3 (July 2014), 87–95.
- [6] Siyuan Cao and He Wang. 2018. Enabling Public Cameras to Talk to the Public. *IMWUT* 2, 2 (2018), 63:1–63:20. <https://doi.org/10.1145/3214266>
- [7] L. Capra, S. Zachariadis, and C. Mascolo. 2005. Q-CAD: QoS and context aware discovery framework for mobile systems. In *ICPS '05. Proceedings. International Conference on Pervasive Services, 2005*. 453–456. <https://doi.org/10.1109/PERSER.2005.1506566>
- [8] Jung Il Choi, Mayank Jain, Kannan Srinivasan, Phil Levis, and Sachin Katti. 2010. Achieving Single Channel, Full Duplex Wireless Communication. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking (MobiCom '10)*. ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/1859995.1859997>
- [9] M. H. Diallo, N. Panwar, S. Mehrotra, and A. A. Sani. 2018. Trustworthy Sensing in an Untrusted IoT Environment. In *2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 468–471. <https://doi.org/10.1109/PERCOMW.2018.8480384>
- [10] R. Irmer, H. Drost, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H. Mayer, L. Thiele, and V. Jungnickel. 2011. Coordinated multipoint: Concepts, performance, and field trial results. *IEEE Communications Magazine* 49, 2 (February 2011), 102–111. <https://doi.org/10.1109/MCOM.2011.5706317>
- [11] Kamal Jain, Jitendra Padhye, Venkata N. Padmanabhan, and Lili Qiu. 2003. Impact of Interference on Multi-Hop Wireless Network Performance. In *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking (MobiCom '03)*. Association for Computing Machinery, New York, NY, USA, 66–80. <https://doi.org/10.1145/938985.938993>
- [12] Angela H. Jiang, Daniel L.-K. Wong, Christopher Canel, Lilia Tang, Ishan Misra, Michael Kaminsky, Michael A. Kozuch, Padmanabhan Pillai, David G. Andersen, and Gregory R. Ganger. 2018. Mainstream: Dynamic Stem-Sharing for Multi-Tenant Video Processing. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. USENIX Association, Boston, MA, 29–42.
- [13] Junchen Jiang, Ganesh Ananthanarayanan, Peter Bodik, Siddhartha Sen, and Ion Stoica. 2018. Chameleon: Scalable Adaptation of Video Analytics. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18)*. ACM, New York, NY, USA, 253–266. <https://doi.org/10.1145/3230543.3230574>
- [14] B. Kantarci and H. T. Mouftah. 2014. Trustworthy Sensing for Public Safety in Cloud-Centric Internet of Things. *IEEE Internet of Things Journal* 1, 4 (Aug 2014), 360–368. <https://doi.org/10.1109/IJOT.2014.2373786>
- [15] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta. 2014. Massive MIMO for next generation wireless systems. *IEEE Communications Magazine* 52, 2 (February 2014), 186–195. <https://doi.org/10.1109/MCOM.2014.6736761>
- [16] Yi Li, Lili Qiu, Yin Zhang, Ratul Mahajan, and Eric Rozner. 2008. Predictable Performance Optimization for Wireless Networks. *SIGCOMM Comput. Commun. Rev.* 38, 4 (Aug 2008), 413–426. <https://doi.org/10.1145/1402946.1403005>
- [17] Robert LiKamWa, Yunxin Liu, Nicholas D. Lane, and Lin Zhong. 2013. MoodScope: Building a Mood Sensor from Smartphone Usage Patterns. In *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '13)*. ACM, New York, NY, USA, 389–402. <https://doi.org/10.1145/2462456.2464449>
- [18] He Liu, Stefan Saroiu, Alec Wolman, and Himanshu Raj. 2012. Software Abstractions for Trusted Sensors. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services (MobiSys '12)*. Association for Computing Machinery, New York, NY, USA, 365–378. <https://doi.org/10.1145/2307636.2307670>
- [19] Xiaochen Liu, Yurong Jiang, Purnet Jain, and Kyu-Han Kim. 2018. TAR: Enabling Fine-Grained Targeted Advertising in Retail Stores. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '18)*. ACM, New York, NY, USA, 323–336. <https://doi.org/10.1145/3210240.3210342>
- [20] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. 2008. OpenFlow: enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review* 38, 2 (2008), 69–74.
- [21] T. Nakamura, S. Nagata, A. Benjebbour, Y. Kishiyama, T. Hai, S. Xiaodong, Y. Ning, and L. Nan. 2013. Trends in small cell enhancements in LTE advanced. *IEEE Communications Magazine* 51, 2 (February 2013), 98–105. <https://doi.org/10.1109/MCOM.2013.6461192>
- [22] Le T. Nguyen, Yu Seung Kim, Patrick Tague, and Joy Zhang. 2014. IdentityLink: User-device Linking Through Visual and RF-signal Cues. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 529–539. <https://doi.org/10.1145/2632048.2636072>
- [23] Veljko Pejovic and Mirco Musolesi. 2015. Anticipatory Mobile Computing: A Survey of the State of the Art and Research Challenges. *ACM Comput. Surv.* 47, 3, Article 47 (April 2015), 29 pages.
- [24] Magnus Proebster, Matthias Kaschub, Thomas Werthmann, and Stefan Valentini. 2012. Context-aware resource allocation for cellular wireless networks. *EURASIP Journal on Wireless Communications and Networking* 2012, 1 (12 Jul 2012), 216.
- [25] Hariharan Rahul, Swaran Kumar, and Dina Katabi. 2012. MegaMIMO: Scaling Wireless Capacity with User Demands. In *ACM SIGCOMM 2012*. Helsinki, Finland.
- [26] Eric Rozner, Mi Kyung Han, Lili Qiu, and Yin Zhang. 2013. Model-Driven Optimization of Opportunistic Routing. *IEEE/ACM Trans. Netw.* 21, 2 (April 2013), 594–609. <https://doi.org/10.1109/TNET.2012.2205701>
- [27] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. 2007. Traffic-Aware Channel Assignment in Enterprise Wireless LANs. In *2007 IEEE International Conference on Network Protocols*. 133–143. <https://doi.org/10.1109/ICNP.2007.4375844>
- [28] Eric Rozner, Vishnu Navda, Ramachandran Ramjee, and Shravan Rayanchu. 2010. NAPMan: Network-Assisted Power Management for WiFi Devices. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys '10)*. Association for Computing Machinery, New York, NY, USA, 91–106. <https://doi.org/10.1145/1814433.1814445>
- [29] Nick Tabor. 2018. Smile! The Secretive Business of Facial-Recognition Software in Retail Stores. <http://nymag.com/intelligencer/2018/10/retailers-are-using-facial-recognition-technology-too.html>.
- [30] J. Teng, B. Zhang, J. Zhu, X. Li, D. Xuan, and Y. F. Zheng. 2014. EV-Loc: Integrating Electronic and Visual Signals for Accurate Localization. *IEEE/ACM Transactions on Networking* 22, 4 (Aug 2014), 1285–1296.
- [31] He Wang, Xuan Bao, Romit Roy Choudhury, and Srihari Nelakuditi. 2013. In-Sight: Recognizing Humans Without Face Recognition. In *Proceedings of the 14th Workshop on Mobile Computing Systems and Applications (HotMobile '13)*. ACM, New York, NY, USA, Article 7, 6 pages. <https://doi.org/10.1145/2444776.2444786>
- [32] Alec Wolman, Stefan Saroiu, Paramvir Bahl, and Victor Bahl. 2010. Using Trusted Sensors to Monitor Patients' Habits. In *1st USENIX Workshop on Health Security and Privacy (HealthSec) (1st usenix workshop on health security and privacy (healthsec) ed.)*. <https://www.microsoft.com/en-us/research/publication/using-trusted-sensors-to-monitor-patients-habits/>
- [33] D. Wu, P. Wang, X. Zhang, R. Yan, and R. X. Gao. 2016. A correlation-based approach to trustworthy sensing for cyber-physical systems. In *2016 IEEE International Instrumentation and Measurement Technology Conference Proceedings*. 1–6. <https://doi.org/10.1109/I2MTC.2016.7520438>
- [34] Jingao Xu, Hengjie Chen, Kun Qian, Erqun Dong, Min Sun, Chenshu Wu, Li Zhang, and Zheng Yang. 2019. iVR: Integrated Vision and Radio Localization with Zero Human Effort. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3, Article 114 (Sept. 2019), 22 pages.