Challenge: Unlicensed LPWANs Are Not Yet the Path to Ubiquitous Connectivity

Branden Ghena
University of California, Berkeley
brghena@berkeley.edu

Joshua Adkins
University of California, Berkeley
adkins@berkeley.edu

Longfei Shangguan
Microsoft Cloud & AI
longfei.shangguan@microsoft.com

Kyle Jamieson
Princeton University
kylej@cs.princeton.edu

Philip Levis
Stanford University
pal@cs.stanford.edu

Prabal Dutta
University of California, Berkeley
prabal@berkeley.edu

ABSTRACT

Low-power wide-area networks (LPWANs) are a compelling answer to the networking challenges faced by many Internet of Things devices. Their combination of low power, long range, and deployment ease has motivated a flurry of research, including exciting results on backscatter and interference cancellation that further lower power budgets and increase capacity. But despite the interest, we argue that unlicensed LPWAN technologies can only serve a narrow class of Internet of Things applications due to two principal challenges: capacity and coexistence.

We propose a metric, *bit flux*, to describe networks and applications in terms of throughput over a coverage area. Using bit flux, we find that the combination of low bit rate and long range restricts the use case of LPWANs to sparse sensing applications. Furthermore, this lack of capacity leads networks to use as much available bandwidth as possible, and a lack of coexistence mechanisms causes poor performance in the presence of multiple, independently-administered networks. We discuss a variety of techniques and approaches that could be used to address these two challenges and enable LPWANs to achieve the promise of ubiquitous connectivity.

CCS CONCEPTS

• **Networks** → **Network performance evaluation**; **Wireless access networks**.

ACM Reference Format:

Branden Ghena, Joshua Adkins, Longfei Shangguan, Kyle Jamieson, Philip Levis, and Prabal Dutta. 2019. Challenge: Unlicensed LPWANs Are Not Yet the Path to Ubiquitous Connectivity. In *The 25th Annual International Conference on Mobile Computing and Networking (MobiCom'19)*, October 21–25, 2019, Los Cabos, Mexico. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3300061.3345444>

1 INTRODUCTION

The Internet of Things (IoT) encompasses a broad array of technologies that connect the physical world with large-scale data processing and storage. Smart homes use machine learning to improve HVAC efficiency and occupant comfort [38], smart trucks provide real-time shipping updates and improve vehicle routing [43], and data fusion from soil sensors and drones conserves water and reduces chemical use [57].

Today, we can build ultra-low power devices that last for years on ambient energy or tiny batteries [20, 25]. We can also deploy highly scalable internet services that process massive streams of IoT data [26, 62]. The dominant remaining challenge for the Internet of Things is connecting these two. Energy efficient and robust networking beyond the range of personal area networks such as Bluetooth and 802.15.4 remains an ongoing technical challenge.

Over the past few years, a number of low-power wide-area networks (LPWANs) have gained popularity for their ability to fill this void. Their use of simple protocols and unlicensed bands allowed them to take a first-to-market approach, and their ability to transmit at ranges over a kilometer while drawing only a few hundred milliwatts enables exciting new applications. Numerous research results further improve their performance—backscatter drastically reduces the power needed for LoRa edge devices [15, 40, 51, 56], and coherent combining improves network range and throughput [9].

We argue, however, that connectivity for the Internet of Things remains an unsolved problem. The core claim of this paper is that unlicensed-band, LPWAN technologies as they exist today can serve only a narrow class of IoT applications.

This work is licensed under a Creative Commons Attribution International 4.0 License.

MobiCom'19, October 21–25, 2019, Los Cabos, Mexico

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6169-9/19/10.

<https://doi.org/10.1145/3300061.3345444>

Furthermore, it is unclear if even the improvements provided by recent research will be enough to expand these use cases. We find that two technical challenges remain for LPWAN protocols to be broadly useful: capacity and coexistence.

The first problem, **capacity**, is the amount of throughput shared by all devices on the network. LoRaWAN, one unlicensed-band LPWAN, provides 60 kbps of total throughput shared by devices over the range of several kilometers that a single base station can cover. Individually, low-bandwidth and long-range are not a problem; together, however, they prohibitively restrict the utility of the technology. We define a metric called *bit flux*, which measures the bit rate a protocol can provide over a unit area. Comparing the bit flux requirements of applications and the bit flux that LPWANs provide, we find that unlicensed LPWANs are only suitable for low-rate, sparse sensing applications.

The second problem faced by unlicensed LPWANs is **coexistence**. Even the limited capacity that LPWANs provide assumes that there is only a single network operating in a given area. The use of the unlicensed bands means this is unlikely to be true as the number of IoT applications and stakeholders using these applications grows over time. Unless coexistence between networks is addressed or the capacity of networks operating in the unlicensed band is increased to well above existing and future application needs, contention is likely to lead to poor performance and ultimately a lack of use by future deployments.

Recognizing both the problems and potential for solutions in this space, 3GPP has been developing cellular standards targeting LPWAN applications, with the most notable protocols, NB-IoT and LTE-M, beginning to operate in the US and abroad. While higher in cost, complexity, and power, our evaluation shows that these technologies meet many of the capacity needs that unlicensed LPWANs currently do not, and they avoid problems of coexistence by operating in licensed spectrum. They also provide coverage without requiring gateway deployments, a valuable consideration for many applications.

Our aim is not to propose any single solution to these challenges but rather to motivate future research in wide-area, unlicensed-band communications. With the release of these cellular technologies we are at a critical turning point in this space. We could drive to improve the protocols of unlicensed LPWANs so they are sufficient for application needs, and we could push for coexistence strategies in both protocol and regulation to ensure graceful degradation of applications. Or we could watch as dense and critical applications shift away from unlicensed LPWANs to cellular networks, taking with them the rich opportunity for future innovation and research that has traditionally followed the ubiquitous use of unlicensed bands.

Figure 1: Range and network throughput for several IoT network technologies. Maximum range is estimated from uplink path loss using the Hata model [13]. Network throughput is the uplink payload bitrate shared by all devices connected to a single gateway, accounting for access control overhead. While all emphasize long range and low throughput, each network technology has different capabilities based on its particular protocol choices.

2 NETWORK OVERVIEW

We begin by investigating the low-power, wide-area networks, which we split into two categories: unlicensed-band LPWANs and cellular technologies. For each network, we quantitatively describe range and throughput and also qualitatively explore additional aspects of the networks such as power and deployment that impact real-world use. Figure 1 plots maximum throughput versus maximum range for the networks we describe.

To determine maximum range, we start by determining the maximum path loss for each protocol given transmit power and receiver sensitivity for existing hardware. Transmit distance is then estimated using the Hata model [13] for protocols at 915 MHz and the Hata model PCS extension [35] for cellular protocols operating at 1900 MHz. In both cases the models are configured for medium-sized cities with end device and gateway heights of 1 m and 100 m respectively.

We determine network throughput as the total uplink payload bits per second provided across many devices connected to a single gateway. While the maximum throughput for a single device in a network is important, for the kinds of large-scale, machine-to-machine applications we describe individual devices do not stress the network. Instead it is the deployment of many devices, each individually with small throughput needs, that cumulatively can exceed the network capacity. Calculating network throughput begins with the maximum goodput for a single device, increased by the number of devices a single gateway can concurrently support and reduced by the cost of contention with other deployed devices in the same network.

2.1 Unlicensed LPWANs

Low-power, wide-area network protocols such as Sigfox [47] and LoRaWAN [27] utilize the unlicensed 915 MHz ISM band in the US to provide long-range communication. These unlicensed-band technologies were the first networks to directly target large-scale machine-to-machine communications.

LoRaWAN, a popular LPWAN protocol, is an open network standard built on top of the proprietary LoRa chirp-spread-spectrum physical layer. In LoRaWAN, rather than solely communicating with a single gateway, devices broadcast data in an ALOHA fashion which can be received by any gateway on the network. Each transmission is followed by two listening windows which can contain an acknowledgment or any other downlink destined for the device. While transmitting at 20 dBm, a LoRaWAN transceiver draws about 400 mW, but otherwise it can remain in an idle state indefinitely in which it draws 5 μ W [44]. Anyone who purchases a LoRaWAN gateway can operate their own network, but LoRaWAN is capable of managed network deployments as well, and several operators have deployed networks [16, 28].

In the US, LoRaWAN transmissions hop across 64 channels and additionally select a spreading factor from five possibilities, allowing devices to trade off range and throughput. A LoRaWAN device using the data rate with the most throughput (data rate three: 125 kHz bandwidth with spreading factor 7) is capable of about 5 kbps of goodput. Gateways are capable of receiving packets from different channels or spreading factors simultaneously if the gateway has a sufficient number of decoders. In practice, even if a gateway monitors all 64 channels it would only be capable of decoding at most 64 packets simultaneously. LoRaWAN's access control strategy of unslotted ALOHA reduces total capacity to 18% of maximum channel capacity in the optimal case [1], resulting in a little less than 60 kbps throughput for an entire network.

LoRaWAN is an open standard managed by a nonprofit association, the LoRa Alliance. Combined with accessible transceivers and unlicensed network deployment, this has made LoRaWAN particularly attractive for academic research. A number of papers have been published recently that improve LoRaWAN gateways [9, 10], propose alternative access control methods for LoRaWAN [41, 53], or simulate LoRaWAN networks [5, 55]. One particularly active area of research involves combining LoRa and backscatter techniques to improve range of backscatter technologies [15, 40, 51, 56].

Sigfox, another major player in the unlicensed LPWAN space, is a proprietary standard targeted at infrastructure monitoring. Similar to LoRaWAN, Sigfox utilizes ALOHA-style transmissions from devices in the network which can be received by any gateway. Radios are left off for the majority of the time, with messages sent to the device received during a listening window following each transmission. A

14 dBm Sigfox transmission draws 100 mW, which drops to 150 μ W in idle mode [2].

Unlike LoRaWAN, Sigfox utilizes narrowband transmissions that trade bitrate for distance to an even greater extent. The physical layer bitrate for Sigfox is 600 bps in the US, and the protocol retransmits each packet on two additional channels to increase reliability [47]. A single gateway is capable of receiving transmissions sent simultaneously by up to 270 devices [46], but in combination with an unslotted ALOHA MAC layer, Sigfox is only capable of about 4 kbps of goodput across an entire network.

Another difference from LoRaWAN is that Sigfox is a proprietary network. Sigfox does not support arbitrary users deploying their own networks. Instead partnered network operators deploy connectivity in various regions [48]. In turn, academic research focusing on the network has been limited, with almost no work that solely utilizes Sigfox.

2.2 Cellular Technologies

Cellular technologies are also important to the LPWAN story. GPRS, part of 2G GSM, was a popular network for machine-to-machine communications before reaching end-of-life in the US, and is still popular in the rest of the world. After a standards lag, new cellular networks are now available that target IoT use cases by design. These technologies offer an interesting comparison to unlicensed-band LPWANs and they have begun to see adoption by the research community. Due to their use of licensed frequency bands, cellular networks cannot be deployed by arbitrary users, but rather must be managed networks run by telecommunications companies.

GPRS serves as waypoint from which long-range machine-to-machine communications began. GPRS is capable of providing payload uplink of up to 80 kbps for common class 12 modems. Combined with a large number of channels (124 for the US 900 MHz band and 374 for the US 1900 MHz band) this allows it to theoretically provide a network throughput of 30 Mbps shared among all devices in a cell. GPRS does not provide the same kind of low-energy operation that modern LPWANs do, however, as it requires frequent paging responses to stay connected to a network, which results in high average power. As an example, the SIM800H GPRS modem draws about 4 mW in its lowest power connected mode and over 1 W while transmitting at maximum rate [49].

LTE-M is a recent LTE protocol that targets machine communications. LTE-M's maximum bandwidth is lower than traditional LTE protocols (1.4 MHz vs 5–20 MHz), but it can otherwise coexist in the same band as traditional LTE categories. In the US, Verizon and AT&T have deployed LTE-M support in their networks nationwide [3, 59]. Taking into account various protocol overheads, a single half-duplex LTE-M device is capable of sending payload data at a rate

up to 375 kbps. For a 20 MHz bandwidth network, up to 16 devices could transmit simultaneously for a total network throughput of 6 Mbps.

NB-IoT, another new LTE machine-to-machine protocol, uses even lower bandwidth at a cost of even lower throughput. A single device in a NB-IoT network is capable of a payload upload rate of 62.5 kbps while using 200 kHz of bandwidth. The advantage to operators is that NB-IoT is narrow enough to be deployed in the guard band at the edge of cellular channel allocations. This allows operators to support IoT needs without impacting their existing bandwidth allocations and is how T-Mobile has deployed NB-IoT throughout the US [50]. Assuming a single guard band deployment, the network throughput is equivalent to the single device throughput for NB-IoT.

Both LTE-M and NB-IoT address the traditional power issues of cellular communications through support for a power saving mode in which modems can disable their radio interfaces for extended periods of time (minutes to days) without being disconnected from the network. While the maximum power draw is still high, 1.4 W for LTE-M and 0.9 W for NB-IoT, their sleep mode power is only 30 μ W [54]. Both of the protocols also offer greater range compared to standard LTE communications.

In estimating the throughput and range of the LTE technologies we consider maximum throughput and range. In practice we expect range and throughput to trade off, but taking the maximum of both gives us an upper bound of their throughput, range, and subsequently bit flux.

2.3 Power Comparison

While our evaluation of networks focuses on throughput and range, power is a first-order concern for the networks and applications we consider. Many IoT devices are battery-operated, with communication a large drain on their limited energy supply. A typical measurement of communication energy use is bits per joule, but that metric functions particularly poorly for comparing protocols with very different amounts of overhead. While the cost of each bit transmitted over LTE protocols is quite low, the energy spent reestablishing a network connection upon wakeup from sleep needs to be accounted for as well for a fair comparison to protocols like LoRaWAN which do not have such connection overhead.

Instead, we compare protocols by presenting average power for a sample application workload: a 200 byte upload once per day. This data requirement is on the low side for LTE-M, which would be more efficient with larger payloads. Meanwhile, this requirement is high for Sigfox, which must fragment the payload across 17 packets. We believe this application workload is sufficient for at least providing a sense of the power tradeoff between these networks.

Network Technology	Average Power (μ W)			
	84 Bytes Per 1 Hour	84 Bytes Per 4 Hours	200 Bytes Per 24 Hours	1000 Bytes Per 24 Hours
Sigfox (155 dB)	110	29	11	56
LoRaWAN (143 dB)	12	3.0	1.1	5.1
LTE-M (144 dB)	50	25	12	13
LTE-M (164 dB)	2200	620	150	440
NB-IoT (144 dB)	62	22	13	15
NB-IoT (164 dB)	1800	520	100	240

Table 1: Average power for each network across example application demands. Expected power is presented for cellular protocols both with good connectivity (144 dB) and at maximum range (164 dB), while Sigfox and LoRaWAN are measured only at their maximum ranges. Application demands span from 84 Bytes each hour to 200 Bytes each day. LoRaWAN performs the best in all application cases, around an order of magnitude better than the cellular protocols in good connectivity. Sigfox must fragment payloads across many packets for all application examples, resulting in higher average power. The additional costs of more complicated physical layers and access control mechanisms lead to an increased power draw for the cellular protocols, particularly when at maximum range. NB-IoT performs better than LTE-M at maximum range, but both perform similarly otherwise.

Estimating average power allows us to account for both the transmission and the additional overhead of maintaining the communication protocol. Calculating average power for Sigfox and LoRaWAN is straightforward and we do so based on datasheet numbers for existing hardware [2, 44]. For LTE-M and NB-IoT, we present numbers based on prior literature, which models expected power draw based on expected latencies at multiple total path loss choices [21]. We present the numbers for both a relatively good connection (144 dB total path loss) and at maximum range (164 dB). For reference, the maximum path loss for Sigfox is 155 dB while LoRaWAN (at data rate three) has a maximum path loss of 143 dB. Table 1 presents the average power for several applications ranging from one 84 byte payload every hour to one 200 byte payload every day. LoRaWAN has the lowest overall average power for each application case, around an order of magnitude better than LTE-M or NB-IoT in a good connection and two orders of magnitude better than the cellular technologies at maximum range.

2.4 Studies of LPWANs

Several other works, like this one, explore the capabilities of LPWANs. An early measurement of machine-to-machine communication traffic was performed by Shafiq et al. [45], studying 2G and 3G networks in the US during 2010. One important finding was that while the traffic volume of such devices is individually low, the large population of devices is

what has to be managed successfully by networks. They also note that the devices they study have a higher ratio of uplink to downlink traffic, especially compared to smartphones. We follow the authors' cue and focus on aggregate network throughput for uplink traffic from deployed devices.

Other works explore issues with unlicensed LPWANs. Ismail et al. note that the large coverage area of wide-area networks results in an "unprecedented number of hidden terminals" [17]. Krupka et al. investigate cross-technology collisions between LoRa, Sigfox, and IQRF [23]. Vejlgard et al. measure ISM band interference in a city in Denmark and find that it results in a 50% packet error rate for indoor LoRa devices transmitting to an outdoor gateway [58]. Each of these works is part of the greater story that unlicensed-band LPWANs are not yet ready to connectivity for the Internet of Things. We continue down that path, first comparing the capabilities of LPWANs with application requirements and then presenting possible solutions to the problems we find as future research challenges for the networking community.

3 NETWORK BIT FLUX

To understand how well a network can support pervasively deployed applications, we develop a new metric, *bit flux*, which measures a network's throughput over its coverage area. Specifically, we measure bit flux in units of bit per hour per square meter.

$$\text{bit flux} = \frac{\text{network throughput}}{\text{coverage area}} = \frac{\text{bit/hour}}{\text{m}^2} \quad (1)$$

This measure, which is the two dimensional version of a metric first proposed by Mark Weiser [60], is valuable because it considers how much capacity an application would require from shared networking infrastructure over a large geographical region. Importantly, *this metric captures both the capabilities of networks and the requirements of applications*. A network that provides a higher bit flux than the application requires is capable of serving the connectivity needs of that deployment.

Bit flux is more useful than network throughput alone for deployments with multiple gateways. Just comparing network throughput to application data rate is sufficient for determining whether a single gateway can support a deployment within its coverage area. Once an application spans the deployment regions of multiple gateways, however, this sufficiency analysis becomes difficult because the capability of each gateway needs to be individually compared to the needs of the devices deployed within its coverage area. By looking at throughput capabilities and needs averaged over an area, bit flux can be used to compare the needs of applications of any size to the capabilities of networks with any number of gateways, as long as the deployments are relatively homogeneous in density.

Figure 2: Throughput per unit area (bit flux) as range is varied through power control. Plotted are the bit per hour per square meter for each of the unlicensed-band and cellular LPWANs we discuss. Using power control, networks can reduce their coverage area, increasing their bit flux and allowing them to satisfy the needs of more applications at the cost of the deployment of additional gateways. The minimum and maximum ranges are limited to the power options found in existing hardware for each technology.

Bit flux also accounts for networks that take advantage of spatial reuse. Reducing gateway range increases network capacity by allowing for more concurrent transmissions at the cost of more deployed gateways. This concept is a common method for increasing cellular network capacity, and many LPWANs have some capability for power control to support it. Because bit flux accounts for coverage area, networks with the ability to shrink coverage area can increase their bit flux accordingly. This means that when applied to networks, bit flux does not result in a single value but a function.

Figure 2 demonstrates the increase in bit flux for long-range networks as maximum range decreases. For each network, there is both a maximum and a minimum range that can be achieved based on the maximum and minimum transmit power of existing hardware. As shown, reducing range for a network has the capability of improving its bit flux by several orders of magnitude. Due to much higher throughput, LTE-M has a higher bit flux, even at maximum range, than Sigfox or LoRaWAN at their minimum range. Additionally, LoRaWAN offers a subset of the capabilities of NB-IoT, which has a larger range of power control configurations.

An important limitation of bit flux is that it only measures technical capability, rather than feasibility. While a LoRaWAN network has similar throughput to local area networks such as 802.15.4, LoRaWAN by default covers a much larger area and therefore would provide a significantly lower bit flux. While short-range networks like 802.15.4 could service the throughput needs of city-scale deployments, the need to deploy many gateways or use high-power transmitters would likely make them unrealistic networking choices.

Application	Single Location Throughput (bps)	Single Location Radius (m)	Pervasive Bit Flux ($\frac{bps}{m^2}$)
Zebrant [63]	53	75	0
Trash can monitoring [4]	0.38	370	0.003
Hospital clinic [6]	11	20	0.02
Volcano monitoring [61]	520	1,500	0.2
CitySee [30, 64]	20,400	5,700	1
Electricity metering [12, 39]	51,389	6,180	1.5
Habitat monitoring [29]	10	10	9
H1N1 [22]	18,000	60	43
IMT-2020 [18, 19]	35,556	564	128
MacroScope [52]	12	4	221
GreenOrbs [33, 64]	5,600	80	1,000

Table 2: Throughput, radius, and bit flux of sensing applications published in past sensor networking proceedings and the IMT-2020 standard [19]. The single location metrics show the requirements to deploy an instance of each application, while the pervasive metric assumes that the application is deployed at scale in its target environment. With throughput and bit flux spanning many orders of magnitude, these applications impose highly varying requirements on their underlying networks. While many networking technologies may meet the throughput requirements of a single application, they often do not have the capacity to support one or more of these applications at scale.

An additional problem lies in the amount of bandwidth required to support an application. A network may be able to support the throughput needs of an application but be doing so only by utilizing the entire bandwidth of the frequency band it occupies. In this case, no other networks could coexist within that same frequency band and geographic location. We explore network feasibility, in terms of number of gateways and bandwidth usage in Section 5.

4 PERVASIVE APPLICATIONS

The “Internet of Things” describes a wide and diverse range of applications. To understand and quantify their networking requirements, we survey notable application papers from the sensor networking literature, and consider their networking requirements in two deployment scenarios. The first, **single location** case, assumes the application is deployed to the fullest extent in a single location. We report the throughput and range required to support these deployments by multiplying the number of nodes in the deployment by the amount of data per measurement by the sampling interval.

A single instance of an application is often not consistent with the ubiquity targeted by the IoT. Therefore, we also consider the **pervasive** case for each application, which assumes that the application is scaled to be fully deployed in its target environment. For example, while a single location case may describe an application that monitors a single building, the pervasive case would include monitoring for all buildings

of that type throughout a city. The applications vary tremendously in deployment area, so we employ the bit flux metric to compare them in terms of bits per hour per square meter. The networking requirements for the eleven applications we survey are shown in Table 2, and are described below, along with the assumptions for their pervasive deployments.

Zebrant [63] is one of the earliest sensor network research deployments. It places GPS tracking collars on zebras that asynchronously send location data over a wide-area network. The incredibly low density of wild Grevy’s Zebras results in near zero bit flux over a wide area, with peak throughput coming from monitoring all zebras in a large herd [34].

Trash Can Monitoring [4] reports when trash cans are full in a deployment of 197 monitored trash cans throughout New York City’s Times Square. Each trash can reports approximately twice a day, and we assume the same frequency and density for a pervasive deployment.

Hospital clinic [6] measures patient vital signs in a 32 bed hospital clinic in St. Louis, USA. At scale all patients in the 2915 hospital beds in St. Louis would be monitored [32].

Volcano monitoring [61] senses seismic tremors across 16 devices on Reventador volcano in Ecuador, streaming data when an event is detected. The pervasive case covers a volcanic area at the same sensor density.

CitySee [30, 64] measures air quality from 1196 devices deployed in Wuxi, China, and we assume the same sensor density for a pervasive deployment.

Electricity metering [12, 39] in San Francisco, USA. Approximately 370,000 smart meters throughout the city report 250 byte readings once every four hours.

Habitat monitoring [29] measures microclimate and occupancy of bird burrows with 32 sensors on Great Duck Island off the coast of Maine, USA. A pervasive deployment would monitor the estimated 5000 Storm Petrel nests on Great Duck Island with 7500 sensors [7].

H1N1 [22] measures a single-day human contact graph of 850 people for modeling flu epidemiology in a school in San Francisco, USA. A full deployment would measure interactions for the 80,000 students in San Francisco [24].

IMT-2020 [18, 19] defines performance characteristics of 5G technologies. For machine-type communications it defines a connection density of one million devices per km^2 each transmitting a 32 byte packet every two hours.

MacroScope [52] monitors the microclimate of a redwood tree with 33 sensors on a tree in Sonoma, USA. A full deployment would place sensors on all trees in an old-growth forest, at a density of about 20 trees/acre [37].

GreenOrbs [33, 64] measures ecological data from 330 devices in a forest near Tianmu Mountain in China. We assume pervasive deployment at the same sensor density.

Application	Sigfox	LoRaWAN	NB-IoT	LTE-M	2G GPRS
Zebranet [63]	●	●	●	●	●
Trash can monitoring [4]	●	●	●	●	●
Hospital clinic [6]	○	●	●	●	●
Volcano monitoring [61]	○	●	○	●	●
CitySee [30, 64]	○	○	○	●	●
Electric metering [12, 39]	○	○	○	●	●
Habitat monitoring [29]	○	○	○	●	●
H1N1 [22]		○	○	○	●
IMT-2020 [18, 19]			○	○	●
Macroscopic [52]			○	○	●
GreenOrbs [33, 64]			○	○	●

Table 3: Sufficiency of a networking technologies to meet the pervasive bit flux requirements of each application. A circle indicates sufficiency, however an open circle indicates that range reduction is required for suitability, where suitability is defined as providing greater than five times the bit flux required for each application. The degrees of range reduction required to meet these cases varies significantly. For instance LTE-M can easily meet $5\times$ the capacity of the IMT-2020 standard with greater than 4000 m range, however NB-IoT must reduce its range to less than 1000 m to provide this same capacity.

The eleven applications differ by many orders of magnitude in their throughput and bit flux requirements. Applications that cover a dense phenomenon, such as the redwoods Macroscopic, have a relatively high bit flux even with low throughput, while applications that cover a large area (CitySee) or measure a sparse phenomenon (H1N1) have a low bit flux despite their high throughput. In terms of bit flux, the applications fall into two major categories. Sparse environmental or human monitoring require less than one bit per hour per square meter (Zebranet, trash can monitoring, volcano monitoring, CitySee, electricity metering, and hospital clinic). Denser monitoring (habitat monitoring, H1N1, Macroscopic, GreenOrbs, and IMT-2020) requires several orders of magnitude more bit flux.

5 NETWORK SUITABILITY

Now that we have described several pervasive applications, we investigate how well LPWANs meet their communication needs. To satisfy an application a network must provide equal or greater bit flux than the bit flux of the application in a pervasive deployment. This assumes a uniform distribution of application devices as well as a uniform distribution of gateways. If a network does not meet the bit flux needs at its maximum range, it can reduce range and increase the number of gateways deployed to increase capacity. For each

Figure 3: The proportion of the network capacity used by the H1N1 application for varying gateway density. As shown in Figure 2 and Table 3, networks can increase bit flux through power control to service certain applications at the cost of a decrease in range and a subsequent increase in gateway deployment density. LTE-M and 2G networks can service the application throughout San Francisco, USA (120 km²) with only a few gateways and a small proportion of their total network capacity. LoRaWAN and NB-IoT can also serve the application, but only by allocating a significant proportion of their capacity to it or deploying many gateways.

application, Table 3 displays which networks could serve its data needs for the pervasive case. For the Zebranet application, all networks we investigate suffice due to the low data rate over an extremely wide area. Similarly, all networks can handle the load of the trash can monitoring application because its data rate is so low. On the other side, GreenOrbs can only be satisfied by the cellular IoT solutions due to the density of deployed sensors.

However, the capability to serve the needs of an application does not mean that doing so would be reasonable. The open circles on Table 3 denote circumstances where over 20% of a network's bit flux capacity would be spent on a single application or where reduced range would be required to meet the bit flux needs. Using a majority of total network capacity means using the majority of the bandwidth in the frequency allocation that network occupies. For example, saturating the capabilities of a LoRaWAN network means saturating the throughput of all 64 LoRaWAN channels, a significant portion of the 915 MHz ISM band. This is not a realistic scenario for wide-area deployments in urban locales.

5.1 H1N1 Case Study

To put this idea into more concrete numbers, we take a deeper dive into the H1N1 application. In the pervasive example, we imagine the application deployed throughout the city of San Francisco. Figure 3 shows for each network the number of gateways that would be necessary to serve its throughput needs. The number of gateways is plotted as a line rather than a point because the network could vary its range through power control, changing the number of gateways necessary

for coverage. This is plotted against the proportion of total network capacity the application would be using for gateways deployed at that density with optimal power control. LoRaWAN, for instance, could serve the H1N1 application throughout all of San Francisco with only 24 gateways, but it would use all of its capacity for this application alone. Even with the greatest reduction in range, LoRaWAN would still use 50% of its capacity for the H1N1 application. In practice, this would dedicate a significant chunk of the ISM band towards this application alone.

Another case where capability does not equate to reasonableness is in terms of the number of gateways required to cover an application deployment. For example, NB-IoT could cover the H1N1 application case while only using 1% of its network bit flux, however, doing so would require the deployment of 2000 gateways throughout San Francisco. While this deployment size is not totally implausible based on new femtocell efforts [36], sufficient motivating applications would need to exist before a service provider would invest in such a dense deployment.

The resulting range of the network after power control should be used as a final consideration for a reasonable deployment. Several of the networks we describe are capable of reducing end device power until the resulting range is only several hundred meters. In these situations, a deployment of WLAN technologies, such as WiFi or 802.15.4 should be considered instead of an LPWAN as they can provide much greater throughput at lower energy budgets.

5.2 Are LPWANs Sufficient?

The first takeaway from this analysis is the relative success of cellular IoT technologies. NB-IoT and LTE-M could, at least conceptually, meet the needs of every application we describe. Their high throughput and wide coverage areas allow them to cover the needs of a low-throughput applications with a single deployment, but also affords them the opportunity to over-provision gateways to handle high-throughput applications. In the H1N1 case, LTE-M could cover all of San Francisco with a single gateway using 22% of the capacity of that entire band. Or, it could cover all of San Francisco with 24 gateways and only allocate 1% of the band capacity to that single application.

Even though their use for pervasive applications seems promising from a capacity standpoint, cellular technologies do have their own challenges, the most notable being fees to access the network and the higher average power requirements discussed in Section 2.3. Some IoT devices may have flexibility in their design to accommodate an increased energy demand in trade for network performance and reliability, but this will not be true for all applications.

For unlicensed LPWANs, we find two challenges in network suitability. The first is an issue of *capacity*. For several of the applications we describe, LoRaWAN deployments are unable to transport the data necessary even at minimum range. Even when LoRaWAN and Sigfox can meet the bit flux requirements of an application, they only do so with a dense deployment of gateways using a majority of the unlicensed bandwidth available. To handle pervasive application needs, unlicensed-band LPWANs will need to increase their capacity. This problem is primarily one of implementation. The selection of the Aloha access control mechanism, for instance, greatly reduces network throughput.

The second challenge, as demonstrated by the high percentage of bit flux needed to satisfy some applications, is one of *coexistence*. A network cannot assume that it is deployed in isolation. Especially in the context of city-scale deployments, many networks will be running in the same physical region. For long-term success, technologies making use of the unlicensed band are going to need to share it, either by using so little of the band that multiple networks can naturally coexist or by actively coexisting with other networks. This problem is a fundamental one for long-range networks in unlicensed bands.

These two problems, capacity and coexistence, do not mean LoRaWAN or other unlicensed-band, long-range technologies are unsuitable for all applications; with little contention, sensing deployments in remote areas with sufficiently low bit flux are well-provided for by unlicensed LPWANs. However, to succeed for pervasive applications in urban areas, solutions to these challenges will be necessary.

6 NETWORK SOLUTIONS

If capacity of the unlicensed band and the networks that use it is not sufficient to provide for the desired applications, contention, both within a network and between networks, can lead to poor throughput and unpredictable reliability. A number of solutions have been proposed to increase the capacity of individual LPWAN networks, some of which would also decrease the impact of collisions with coexisting networks. Researchers have also proposed active coordination between networks and widening the unlicensed band.

This section enumerates these techniques. We focus on LoRaWAN as the majority of research projects target it, but the resulting techniques are applicable to many unlicensed-band technologies.

6.1 Improving Transmission

Modifying LoRaWAN's access control mechanism could result in greatly increased capacity for a single network, although it would not greatly increase the ability of a network to coexist with other networks beyond the decrease in total

channel usage. Polonelli et al. describe a method for layering ALOHA slots on top of the existing LoRaWAN protocol [41]. They utilize acknowledgements for device synchronization with the gateway.

Alternatively, channel access can be explicitly scheduled. Trüb et al. demonstrate two possible TDMA schemes that could be employed for LoRaWAN systems which could improve network throughput to 60%, or three times that of unslotted ALOHA [53]. Neither work measures the increase in energy cost for implementing such schemes. For dense networks which would experience many packet collisions under ALOHA-style access control, however, the energy cost for scheduling may be lower than the cost of repeated transmissions. Access control changes would need modifications to software on both gateways and devices, a challenge for existing deployments.

6.2 Resilient Reception

Methods for better receiving packets in the presence of noise not only increase the capacity of a network, but also increase resiliency to the presence of coexisting networks.

DaRe [31] performs convolutional erasure coding on LoRaWAN application-layer data such that a lost packet can be recovered from other packets. Application layer coding may be one method for increasing data reception rates while requiring software changes to the gateways and end devices. Applying a code rate to payloads would increase energy use as on-air time increases.

Choir [10] leverages radio imperfections in frequency, time, and phase to simultaneously receive several transmissions. While Choir suggests that this could enable as much as a 30× increase in network capacity, NetScatter finds that this technique would enable no more than 5-10 simultaneous transmitters [15]. Charm [9] uses coherent combining to increase reception rates for transmissions weak in signal strength. This could also serve to increase resiliency to collisions, however we do not know the exact magnitude of this improvement. Deployment of systems like Choir and Charm require modifications to gateway hardware and software, but can be deployed without modifying existing devices and without any increase in energy.

We may also rely on the capture effect to improve reception rates. This causes a packet with stronger signal strength to be received despite a collision with a weaker signal, and can be achieved by densely deploying gateways without reducing the transmit power of end devices. To evaluate this technique, we simulate a deployment of LoRaWAN devices coexisting on a single channel using a modified version of LoRaSIM [5]. In our simulation, devices and gateways are deployed randomly across a 5 km region and devices send a 20 byte packet once per minute on average. As shown in

Figure 4: Increased deployment of gateways results in higher packet reception rate due to the capture effect. Shown is the reception rate for packets sent by 100 devices on the target network. As the total number of deployed devices, most not on the network, increases, collisions cause packets to be lost. Increasing the number of gateways deployed throughout the same area results in more packets received as some overcome collisions due to the capture effect.

Figure 4, when a network of a single gateway and 100 devices is deployed alongside 1000 devices on another network, the gateway receives 27% of transmissions, however when 10 gateways are deployed in the same scenario, the reception rate increases to a 57% due to the capture effect.

6.3 Increasing Bandwidth

An increase to the amount of bandwidth usable by unlicensed systems would result in increased capacity for all networks operating in the band. There has been progress towards making additional unlicensed spectrum available in the US, specifically TV bands around 600 MHz [11]. The amount of TV white space bandwidth available for unlicensed use depends on local channel usage, which varies widely based on location and population density. Still, one simulation finds that an average of 80 MHz could be available in cities in the US [14], triple the 26 MHz currently allocated to the 900 MHz ISM band.

While existing LPWAN transceivers have support for some of the TV whitespace frequencies, firmware changes on both devices and gateways would be necessary to exploit this hardware. Additionally, protocols would need to adopt the ability to determine which TV channels are available in a particular deployment area.

6.4 Coexisting through Coordination

Ultimately unlicensed-band collisions are inevitable between networks, and all the more inevitable due to the number of stakeholders present within the range of an LPWAN. If the underlying capacity of the band is not enough to meet the needs of its users, some form of coordination may be the

only hope of increasing capacity and creating predictable and fair performance.

This coordination could be done in both the frequency and time domains. Some protocols, such as WiFi, divide the frequency domain with protocol-specified limits of single-network bandwidth, ensuring some minimum number of networks can coexist. This technique is more difficult for unlicensed LPWANs, which may have to coexist with many more networks due to their range, and would additionally further limit their throughput.

Coordination becomes simpler for managed networks run by a select few service providers. A limited number of network operators could provide communication in the unlicensed bands, such as is occurring with Sigfox. A dominant regional provider could result in de facto ownership of the band, as other smaller network deployments would have to work around them. Removing the option for anyone to put up a gateway and create their own network would be a loss of some of the value of the unlicensed bands, however.

Techniques to enable coordination through inter-network communication have also been considered. De Poorter et al. present a design for an LPWAN management framework that includes cross-network and cross-technology optimization [8], however it requires synchronization of heterogeneous devices to enable layered TDMA schemes. WiSHFUL demonstrates a similar scheme, which changes access control mechanisms to reduce cross-technology interference [42].

To be successful, coordination needs buy-in from the majority users of network capacity. Regulations, such as the 1% per-device duty cycle limits in the EU are one mechanism to enforce this coordination, however per-network rather than per-device limits may be necessary to prevent takeover by a deployment of many devices.

6.5 Takeaways

These potential solutions beg the question: how far are unlicensed LPWANs from supporting the Internet of Things? Reviewing the techniques we have described, simple TDMA mechanisms could increase capacity by 3 $\times$, simultaneous reception by 5 $\times$, coherent combining by 1-2 $\times$, and increasing bandwidth to the predicted TV white space availability by 4 $\times$. Together these implementation changes might therefore generate a two order-of-magnitude improvement to LoRaWAN capacity, pushing it to be on par with LTE-M in capacity (and probably looking a lot like an LTE protocol in design).

Unfortunately, this alone would be insufficient to solve unlicensed LPWAN problems. Coexistence solutions, which licensed-band cellular technologies can ignore, will also be necessary. This problem is fundamental to long-range communication in unlicensed bands, and the coordination mechanisms necessary to overcome it are less clear and less studied in prior work.

Few of these techniques are free. Many, particularly modifications to access control mechanisms, would increase the energy cost of communication. One of the biggest strengths of existing unlicensed band networks like LoRaWAN is how low power they are, but the reality is that additional energy costs will need to be paid in order to provide connectivity for higher-throughput applications in urban areas. Exactly where in the tradeoff space between capability and energy cost future networking technologies ought to fall is still unclear, but we believe the question is worthy of exploration.

7 CONCLUSION

The unlicensed bands have long been the home to diverse, innovative, and incredibly impactful networking solutions, especially for the research community. As we investigate LPWAN technologies and pervasive applications, however, we find a disconnect between application needs and unlicensed LPWAN capabilities. We propose the bit flux metric for comparing wide-area application requirements to network capabilities in terms of throughput averaged over coverage area. Using this metric, we demonstrate that existing unlicensed band LPWANs can only serve a narrow class of IoT applications: low-rate, sparse sensing applications.

The problems of capacity and coexistence lead us to investigate possible methods of improving the networks. Improvements to implementation could provide two order-of-magnitude improvement in capacity, but fundamental problems remain in coexisting with other unlicensed band stakeholders that must be overcome as well. We hope that as a research community we can rise to this steep challenge and continue to support both pervasive applications on the unlicensed bands and the research opportunities provided by their ubiquitous use.

8 ACKNOWLEDGMENTS

We would like to thank our anonymous reviewers for their insightful feedback, and our shepherd, Swarun Kumar, for his guidance on how to improve this paper and clarify its arguments. We would also like to thank Basavaraj Tonshal and Hamed Asadi at Ford for many discussions about the capabilities of LPWANs.

This work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. Support was also provided by the Stanford Secure Internet of Things Project and the Stanford System X Alliance. Additionally, this material is based upon work supported by the National Science Foundation under grant numbers CNS-1824277, CNS-1617161, DGE-1256260, and DGE-1106400, and NSF/Intel CPS Security under grants CNS-1822332 and CNS-1505728.

REFERENCES

- [1] Norman Abramson. 1970. The ALOHA System: Another Alternative for Computer Communications. In *Proceedings of the November 17-19, 1970, Fall Joint Computer Conference*. ACM, 281–285.
- [2] Atmel. 2016. ATA8520E Datasheet.
- [3] AT&T Business. 2018. LTE-M Network Solutions. <https://www.business.att.com/solutions/Service/internet-of-things/networks/lte-m/>.
- [4] Bigbelly, Inc. 2019. New York City's Times Square Efficiently Manages 26,056 Gallons of Waste and Recycling Each Day with Bigbelly. <http://info.bigbelly.com/case-study/times-square-new-york-city>.
- [5] Martin C. Bor, Utz Roedig, Thiemo Voigt, and Juan M. Alonso. 2016. Do LoRa Low-Power Wide-Area Networks Scale?. In *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*. ACM, 59–67.
- [6] Octav Chipara, Chenyang Lu, Thomas C. Bailey, and Gruia-Catalin Roman. 2010. Reliable Clinical Monitoring Using Wireless Sensor Networks: Experiences in a Step-down Hospital Unit. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*. ACM, 155–168.
- [7] College of the Atlantic. 2018. Seabirds at Great Duck Island. <https://www.coa.edu/islands/great-duck-island/seabirds-at-gdi/>.
- [8] Eli De Poorter, Jeroen Hoebeke, Matthias Strobbe, Ingrid Moerman, Steven Latré, Maarten Weyn, Bart Lannoo, and Jeroen Famaey. 2017. Sub-GHz LPWAN Network Coexistence, Management and Virtualization: An Overview and Open Research Challenges. *Wireless Personal Communications* 95, 1 (2017), 187–213.
- [9] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. 2018. Charm: Exploiting Geographical Diversity Through Coherent Combining in Low-Power Wide-Area Networks. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 60–71.
- [10] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM'17)*. ACM, 309–321.
- [11] Federal Communications Commission. 2008. In the Matter of Unlicensed Operation in the TV Broadcast Bands, Additional Spectrum for Unlicensed Devices Below 900 MHz and in the 3 GHz Band. Report and Order, FCC-08-260.
- [12] Dominic Fracassa. 2019. CleanPowerSF tripling households served with municipal electricity. <https://www.sfchronicle.com/bayarea/article/CleanPowerSF-tripling-households-served-with-13618155.php>.
- [13] Masaharu Hata. 1980. Empirical Formula for Propagation Loss in Land Mobile Radio Services. *IEEE transactions on Vehicular Technology* 29, 3 (1980), 317–325.
- [14] Farzad Hessar and Sumit Roy. 2015. Capacity Considerations for Secondary Networks in TV White Space. *IEEE Transactions on Mobile Computing* 14, 9 (2015), 1780–1793.
- [15] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. 2019. Netscatter: Enabling Large-scale Backscatter Networks. In *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation (NSDI'19)*. USENIX Association, Berkeley, CA, USA, 271–283. <http://dl.acm.org/citation.cfm?id=3323234.3323258>
- [16] The Things Industries. 2019. The Things Network. <https://www.thethingsnetwork.org/>.
- [17] Dali Ismail, Mahbubur Rahman, and Abusayeed Saifullah. 2018. Low-power Wide-Area Networks: Opportunities, Challenges, and Directions. In *Proceedings of the Workshop Program of the 19th International Conference on Distributed Computing and Networking*. ACM, 8.
- [18] ITU-R. 2017. Guidelines for Evaluation of Radio Interface Technologies for IMT-2020.
- [19] ITU-R. 2017. Minimum Requirements Related to Technical Performance for IMT-2020 Radio Interface(s).
- [20] Vikram Iyer, Rajalakshmi Nandakumar, Anran Wang, Sawyer B. Fuller, and Shyamnath Gollakota. 2019. Living IoT: A Flying Wireless Platform on Live Insects. In *To appear in Mobicom*.
- [21] Pascal Jörke, Robert Falkenberg, and Christian Wietfeld. 2018. Power Consumption Analysis of NB-IoT and eMTC in Challenging Smart City Environments. In *2018 IEEE Globecom Workshops*. IEEE, 1–6.
- [22] Maria A. Kazandjieva, Jung Woo Lee, Marcel Salathé, Marcus W. Feldman, James H. Jones, and Philip Levis. 2010. Experiences in Measuring a Human Contact Network for Epidemiology Research. In *Proceedings of the 6th Workshop on Hot Topics in Embedded Networked Sensors*. ACM, 7.
- [23] Lukas Krupka, Lukas Vojtech, and Marek Neruda. 2016. The Issue of LPWAN Technology Coexistence in IoT Environment. In *2016 17th International Conference on Mechatronics-Mechatronika (ME)*. IEEE, 1–8.
- [24] Lapkoff and Gobalet Demographic Research Inc. 2010. Demographic Analyses and Enrollment Forecasts for the San Francisco Unified School District. <https://www.sfsud.edu/en/assets/sfusd-staff/enroll/files/DemographicReport3182010.pdf>.
- [25] Yoonmyung Lee, Suyoung Bang, Inhee Lee, Yejoong Kim, Gyouho Kim, Mohammad Hassan Ghaed, Pat Pannuto, Prabal Dutta, Dennis Sylvester, and David Blaauw. 2013. A Modular 1 mm³ Die-Stacked Sensing Platform with Low Power I²C Inter-die Communication and Multi-Modal Energy Harvesting. In *IEEE Journal of Solid-State Circuits*, Vol. 48. 229–243.
- [26] Wei Lin, Haochuan Fan, Zhengping Qian, Junwei Xu, Sen Yang, Jिंगren Zhou, and Lidong Zhou. 2016. StreamScope: Continuous Reliable Distributed Processing of Big Data Streams. In *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation (NSDI'16)*. USENIX Association, Santa Clara, CA, 439–453. <https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/lin>
- [27] LoRa Alliance. 2017. LoRaWAN 1.1 Specification.
- [28] MachineQ. 2019. Improving How the World Operates by Bridging the Physical and Digital Worlds - MachineQ. <https://machineq.com/>.
- [29] Alan Mainwaring, David Culler, Joseph Polastre, Robert Szewczyk, and John Anderson. 2002. Wireless Sensor Networks for Habitat Monitoring. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*. ACM, 88–97.
- [30] Xufei Mao, Xin Miao, Yuan He, Xiang-Yang Li, and Yunhao Liu. 2012. CitySee: Urban CO₂ Monitoring with Sensors. In *IEEE International Conference on Computer Communications (INFOCOM'12)*. 1611–1619.
- [31] Paul J. Marcelis, Vijay S. Rao, and R. Venkatesha Prasad. 2017. DaRe: Data Recovery Through Application Layer Coding for LoRaWAN. In *2017 IEEE/ACM Second International Conference on Internet-of-Things Design and Implementation (IoTDI'17)*. IEEE, 97–108.
- [32] Missouri Department of Health and Senior Services. 2018. MO Hospital Profiles By County. <https://health.mo.gov/safety/healthservregs/pdf/MOHospbyCounty.pdf>.
- [33] Lufeng Mo, Yuan He, Yunhao Liu, Jizhong Zhao, Shao-Jie Tang, Xiang-Yang Li, and Guojun Dai. 2009. Canopy Closure Estimates with GreenOrbs: Sustainable Sensing in the Forest. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys '09)*. 99–112.
- [34] P.D. Moehlman, D.I. Rubenstein, and F. Kebede. 2015. The IUCN Red List of Threatened Species. <http://www.iucnredlist.org/details/7950/0#sectionPopulation>.

- [35] Preben E. Mogensen and Jeroen Wigard. 1999. *COST Action 231: Digital Mobile Radio Towards Future Generation System, Final Report*. Chapter 5.2 On Antenna and Frequency Diversity in GSM; 5.3 Capacity Study of a Frequency Hopping GSM Network, 210–231.
- [36] Takehiro Nakamura, Satoshi Nagata, Anass Benjebbour, Yoshihisa Kishiyama, Tang Hai, Shen Xiaodong, Yang Ning, and Li Nan. 2013. Trends in Small Cell Enhancements in LTE Advanced. *IEEE Communications Magazine* 51, 2 (2013), 98–105.
- [37] National Park Service. 2015. Del Norte Coast Redwoods. <https://www.parks.ca.gov/pages/414/files/DelNorteSPFinalWebLayout2015.pdf>.
- [38] Nest. 2019. Nest Create a Connected Home.
- [39] Pacific Gas and Electric Company. 2016. EPIC 1.14 - Next Generation SmartMeter Telecom Network Functionalities.
- [40] Yao Peng, Longfei Shangguan, Yue Hu, Yujie Qian, Xianshang Lin, Xiaojiang Chen, Dingyi Fang, and Kyle Jamieson. 2018. PLoRa: A Passive Long-range Data Network from Ambient LoRa Transmissions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. ACM, 147–160.
- [41] Tommaso Polonelli, Davide Brunelli, and Luca Benini. 2018. Slotted ALOHA Overlay on LoRaWAN-A Distributed Synchronization Approach. In *2018 IEEE 16th International Conference on Embedded and Ubiquitous Computing (EUC)*. IEEE, 129–132.
- [42] Peter Ruckebusch, Spilios Giannoulis, Domenico Garlisi, Pierluigi Gallo, Piotr Gawlowicz, Anatolij Zubow, Mikoaj Chwalisz, Eli De Poorter, Ingrid Moerman, Ilenia Tinnirello, et al. 2017. Wishful: Enabling Coordination Solutions for Managing Heterogeneous Wireless Networks. *IEEE Communications Magazine* 55, 9 (2017), 118–125.
- [43] Samsara. 2019. Samsara for Construction.
- [44] Semtech. 2016. SX1276/77/78/79 Datasheet.
- [45] Muhammad Zubair Shafiq, Lusheng Ji, Alex X Liu, Jeffrey Pang, and Jia Wang. 2012. A First Look at Cellular Machine-to-Machine Traffic: Large Scale Measurement and Characterization. *ACM SIGMETRICS Performance Evaluation Review* 40, 1 (2012), 65–76.
- [46] Sigfox. 2017. Number of Orthogonal Sigfox Channels. <https://ask.sigfox.com/questions/3194/number-of-orthogonal-sigfox-channels.html>.
- [47] Sigfox. 2019. Sigfox Connected Objects: Radio Specifications. Revision 1.3.
- [48] Sigfox. 2019. Sigfox Operators. <https://partners.sigfox.com/companies/sigfox-operator>.
- [49] SIMCom. 2013. SIM800H Hardware Design.
- [50] T-Mobile Newsroom. 2018. T-Mobile Launches Nation’s First Plan for Narrowband IoT. <https://www.t-mobile.com/news/narrowband-iot>.
- [51] Vamsi Talla, Mehrdad Hesar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. 2017. LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3, Article 105 (9 2017), 24 pages. <https://doi.org/10.1145/3130970>
- [52] Gilman Tolle, Joseph Polastre, Robert Szewczyk, David Culler, Neil Turner, Kevin Tu, Stephen Burgess, Todd Dawson, Phil Buonadonna, David Gay, et al. 2005. A Macroscope in the Redwoods. In *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems*. ACM, 51–63.
- [53] Roman Trüb and Lothar Thiele. 2018. Increasing Throughput and Efficiency of LoRaWAN Class A. In *UBICOMM 2018*. ThinkMind, 54–64.
- [54] ublox. 2019. SARA-R4/N4 Series Datasheet.
- [55] Floris Van den Abeele, Jetmir Haxhibeqiri, Ingrid Moerman, and Jeroen Hoebeke. 2017. Scalability Analysis of Large-Scale LoRaWAN Networks in NS-3. *IEEE Internet of Things Journal* 4, 6 (2017), 2186–2198.
- [56] Ambuj Varshney, Oliver Harms, Carlos Pérez-Penichet, Christian Rohner, Frederik Hermans, and Thiemo Voigt. 2017. Lorea: A Backscatter Architecture that Achieves a Long Communication Range. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. ACM, 18.
- [57] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta N Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. 2017. FarmBeats: An IoT Platform for Data-Driven Agriculture. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI’17)*. 515–529.
- [58] Benny Vejlgaard, Mads Lauridsen, Huan Nguyen, István Z Kovács, Preben Mogensen, and Mads Sorensen. 2017. Interference Impact on Coverage and Capacity for Low Power Wide Area IoT Networks. In *2017 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.
- [59] Verizon Enterprise. 2018. Transitioning your IoT Environment to 4G LTE. <http://www.verizonenterprise.com/products/internet-of-things/4g-lte-iot/next-gen/#ps-overlay>.
- [60] Mark Weiser. 1993. Some Computer Science Issues in Ubiquitous Computing. *Commun. ACM* 36, 7 (7 1993), 75–84.
- [61] Geoff Werner-Allen, Konrad Lorincz, Jeff Johnson, Jonathan Lees, and Matt Welsh. 2006. Fidelity and Yield in a Volcano Monitoring Sensor Network. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*. USENIX Association, 381–396.
- [62] Mingming Zhang, Tianyu Wo, Tao Xie, Xuelian Lin, and Yaxiao Liu. 2017. CarStream: An Industrial System of Big Data Processing for Internet-of-vehicles. *Proceedings of the VLDB Endowment* 10, 12 (8 2017), 1766–1777. <https://doi.org/10.14778/3137765.3137781>
- [63] Pei Zhang, Christopher M. Sadler, Stephen A. Lyon, and Margaret Martonosi. 2004. Hardware Design Experiences in ZebraNet. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*. 227–238.
- [64] Xiaolong Zheng. 2018. Personal Communication.