LoRaX: Repurposing LoRa as a Low Data Rate Messaging System to Extend Internet Boundaries

MORGAN VIGIL-HAYES, Northern Arizona University, Flagstaff, AZ
MD NAZMUL HOSSAIN, Northern Arizona University, Flagstaff, AZ
ALEXANDER ELLIOTT, Georgia Institute of Technology, Atlanta, GA
ELIZABETH BELDING, University of California, Santa Barbara, Santa Barbara, CA
ELLEN ZEGURA, Georgia Institute of Technology, Atlanta, GA

Globally, 43% of households lack Internet access, primarily in regions where deployment and/or service costs are prohibitive, including in the least developed countries, rural locations, and regions with high concentrations of ethnic minorities and low-income populations. Unfortunately, this lack of Internet access increasingly equates to a lack of access to essential services, such as healthcare, education, and economic opportunities. In an environment of marginal economics, creative and varied approaches to obtaining access have flourished, including Internet kiosks long popular in the Global South, libraries as public access in the Global North, parking lot use of open WiFi access points, and spectrum-based solutions such as TV whitespace links and citizen band radio. In the near future, local 5G and the deployment of satellite constellations promise yet additional options in the price/performance space for access. In this context we are interested in the following research question: *How can the presence of multiple networks, with different price, performance, and geographic reach profiles, be best used in concert to improve access to critical services?* We propose that a robust answer to this question bears a holistic, cross-layer examination of new communication paradigms, network architecture innovation, and application design. We make this concrete by running to ground a specific case study of two networks, one high performance yet limited in geographic scope and the other low performance yet pervasive. Specifically our LoRaX (LoRa eXtends the Internet) system combines high bandwidth but non-pervasive Internet access with a low data rate, low power, yet ubiquitious network made possible by IoT developments. By focusing on two networks with extreme differences, we explore a design space that offers users new opportunities for participating in Internet-based services–even when high speed Internet connectivity is intermittent. We also reflect on the generality of the environment and our solution approach for future multi-network settings.

CCS Concepts: • **Networks** → **Network architectures**; **Network experimentation**; • **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**.

Additional Key Words and Phrases: digital divide, last mile connectivity, challenged networking environment, low data rate messaging, LoRa, LPWAN

ACM Reference Format:
Morgan Vigil-Hayes, Md Nazmul Hossain, Alexander Elliott, Elizabeth Belding, and Ellen Zegura. 2022. LoRaX: Repurposing LoRa as a Low Data Rate Messaging System to Extend Internet Boundaries. In *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS) (COMPASS '22), June 29-July 1, 2022, Seattle, WA, USA.* ACM, New York, NY, USA, 28 pages. https://doi.org/10.1145/3530190.3534807

COMPASS '22, June 29-July 1, 2022, Seattle, WA, USA

ACM ISBN 978-1-4503-9347-8/22/06…$15.00
https://doi.org/10.1145/3530190.3534807

1 INTRODUCTION

Despite significant efforts to deploy last-mile Internet connectivity over the past decade, only 57% of global households have Internet access [51]. The households most likely to lack access are those in the least developed countries [51], those in rural locations [76], and those belonging to ethnic minorities and low-income populations [26, 78]. Two main factors contribute to this digital inequality: cost of deployment and cost of service. For rural and remote communities, the cost to providers of deploying last-mile Internet connectivity can be difficult to offset based on the population size of the service area. Moreover, for communities with emergent economies and low-income demographics, monthly service costs associated with home Internet access or LTE subscriptions can be prohibitively expensive [78, 90]. Critically, a lack of Internet access increasingly equates to a lack of access to essential services, such as healthcare, education, and economic opportunities [83].

In an environment of marginal economics, creative and varied approaches to obtaining access have flourished. Internet kiosks have long been a part of the access landscape in the Global South [44], while libraries and community centers offer public access in the Global North and beyond [37]. The Internet access situation in marginalized communities became even more dire during the COVID-19 pandemic as nearly all aspects of school, work, and life moved online, producing examples of individuals traveling, in some cases significant time and distance, for Internet access [9, 56]. Looking ahead, the deployment of satellite constellations [53] and 5G networks promise yet additional options in the price/performance space for access. Significantly, now and into the future the Internet access landscape will consist of multiple technologies that individuals must navigate. Current network architectures and applications provide limited support for this navigation.

In this context we are interested in the following research question: *How can the presence of multiple access networks, with different price, performance, and geographic reach profiles, be best used in concert to improve access to critical services?* To make this concrete, we consider a specific instance made possible by recent advances in long distance, low power, but low bandwidth IoT technologies. Our two-network proof-of-concept system, LoRaX (LoRa eXtends the Internet)[1], supports access to services over a combination of low- and high-bandwidth network regimes. As illustrated by Figure 1, the key idea is to use the *pervasive*, low data rate network for messaging and service initiation, followed by service completion at a later time over the *geographically constrained* high bandwidth network. This novel *initiate-then-update (ITU)* communication paradigm is coupled to a user interface design that explicitly reveals to the user the difference in network capability to assist in effective use. Taken as a whole, this system produces a new point in the design space of cost-effective Internet access achieved by cross-layer considerations.

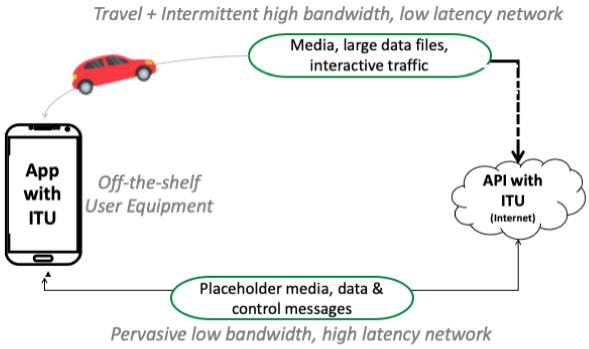

Fig. 1. Dual-regime ITU system.

Our work makes the following contributions:

- A new point in the design space of cost effective Internet access that re-purposes a newer link technology with the architectural innovation of a two-network system and tests the idea of involving the end user in navigating regimes of differing connectivity performance.

[1]LoRa is a popular IoT networking standard. There are alternatives to LoRa in the low power, long range IoT space, including SigFox [57, 64] and Narrow Band IoT [63, 81]. We selected LoRa for the availability of off-the-shelf equipment and the developer community. Our key contributions apply regardless of the specific low data rate technology chosen.

- The development of an *initiate-then-update* paradigm for applications operating across two network regimes. The key idea behind this paradigm is that a useful set of service calls can be partially completed with limited data rate transfers and then fully completed when high bandwidth access is available. An increasing number of services with RESTful APIs are compatible with this approach.
- A proof-of-concept implementation and performance evaluation that demonstrates the feasibility of using LoRaX as a cost-effective means for supporting economic development in contexts where Internet connectivity is not ubiquitous and thus the alternative is to drive for access.

After running our proof-of-concept to ground, we reflect on the generality of the multiple network paradigm and the approach we take to use these networks in concert, as well as the economic viability of a pervasive low cost, low power network. We anticipate that the design insights from LoRaX and the initiate-then-update paradigm may lead others to consider cross-layer and architectural approaches to bridging stubborn digital divides as well as provide ideas useful for future highly heterogeneous network environments, even when well-provisioned.

2 BACKGROUND & CONTEXT

Our work takes place in the context of two trends. The first is the considerable past effort to address digital divides with new technology-based solutions. The second is the increasing availability of multiple types of networks in a given region, beyond the common example of the cellular/WiFi dual network situation frequently encountered in well-provisioned areas today. To the first contextual trend, digital inequalities have proven to be persistent even in the face of efforts to deploy new infrastructure and increase accessibility [26, 51, 90]. In the research and development community, many technology advances to improve access naturally lie at the link layer and emphasize longer reach with moderate or higher bandwidth. Notable in this category are TV White Space (TVWS) links [19, 66, 71], Long-distance WiFi [7, 65, 73, 80], and LEO satellite constellations [8, 53]. These have one hop geographical coverage ranging from 10s of kilometers to 1000s for satellite, and data rates in the tens to low hundreds of Mbps. By comparison, low-power, low data rate technologies developed for IoT such as Low-powered Radio (LoRa) [62] have reach comparable to Long-distance WiFi and TVWS with much lower data rate, but notably lower subscriber and provider cost. Less common are architectural advances that use existing link technologies in novel ways. A notable example here is Disruption Tolerant Networks. With origins in delay tolerant networking for interplanetary communication [13], Disruption Tolerant Networks use a store-carry-forward paradigm to bridge regions of temporary disconnection [35]. DTNs were an area of active research and produced a bundle protocol for transmission across connectivity regions [48], numerous routing protocols for sparse and changing connectivity [1, 2, 28], and application prototyping [75].

To the second contextual trend, as network technologies become increasingly hyperlocal (e.g., 5G and CBRS) or regional (e.g., LoRa and Sigfox) in their service ranges, there is a general trend towards multiple network regimes overlapping in a given area. ICTD literature demonstrates that underserved communities tend to rely on these different network regimes for the distinct purposes of extending the reach of connectivity (regional networks) [73, 103] and providing robust local connectivity (hyperlocal networks) [46, 52]. However, given the continued trend of overlapping pervasive low-capacity networks and hyperlocal high capacity networks, we argue that there is a need for a paradigm that weaves these coinciding network regimes together rather than relegating services to one regime or another. To do this, we envision an *initiate-then-update (ITU) paradigm that supports data services over a combination of low- and high-bandwidth network regimes*. It may seem counter intuitive to draw on a low data rate technology in the context of improved access. Our key insight is that the reach and cost profile of increasingly pervasive IoT network coverage bears creative examination in the design space of cost-effective access.

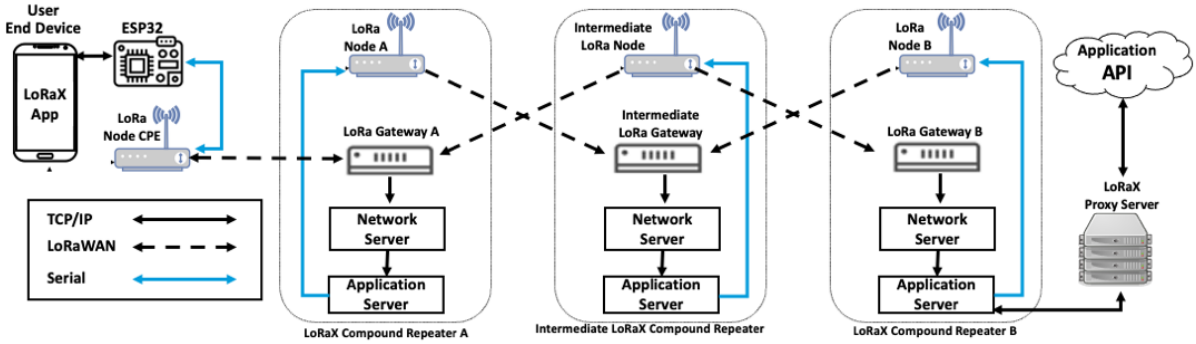

Fig. 2. LoRaX network architecture.

3 SYSTEM DESIGN & IMPLEMENTATION

We design and implement LoRaX as a dual-network instance of an ITU system that uses LoRa as the pervasive, low-capacity network and WiFi LANs as the high-capacity networks. LoRa provides an ideal test case for examining the potential uses of a pervasive, low-capacity network; it is highly economical (see Section 5.3) and has been growing in prominence as the de facto low power, wide area network (LPWAN) standard for supporting rural IoT [61]. It is critical to note that *any* pervasive, low-capacity network could be used in an implementation of an ITU system. For instance, one might envision an ITU system that uses 2G networks for the low-capacity regime and LTE/5G for the high capacity regimes. With LoRaX, the benefits of LoRa, namely low cost and wide coverage area can be integrated into a larger system that compensates for its low data rate capacity through the placement of a proxy that can coordinate service delivery over high-capacity and low-capacity network regimes on behalf of the user. In the context of Figure 1, LoRaX uses LoRa as the pervasive, low-capacity network and WiFi LAN as the intermittently available, high-capacity network. Notably, LoRaX extends the boundary of connectivity to Internet-based services through a proxy server that manages data transactions between low and high data rate network regions.

Figure 2 depicts the major components of the LoRaX system, with the protocols and information flow through the LoRa network to the proxy and on to the Internet. On the left, a User End (UE) Device, such as a smartphone or tablet, runs an application that interfaces with Internet-based services using LoRaX. A LoRa node located at the customer premise (CPE) acts as the access between an off-the-shelf, non-LoRa-enabled UE device and the LoRa network. In the middle, our LoRaX Compound Repeaters each combine a LoRa node with a LoRa gateway to support a multi-hop, bidirectional LoRa network from the UE to a LoRaX Proxy Server on the Internet. The LoRaX Proxy Server manages connectivity between the low rate LoRa network and the high rate Internet and accesses Application APIs on behalf of the user.

Since our usage of LoRa as part of LoRaX deviates from the conventional LoRa usage and architecture (which we describe in Appendix A), there are several key challenges that we needed to address. First, we needed to implement software support that would allow mobile apps to coordinate transactions executed via LoRaWAN and TCP/IP protocol stacks. Second, we needed to re-architect application services so that they were logically accessible through both high- and low-bandwidth regimes. We accomplished this by decomposing Internet-based APIs into smaller units of informational transactions that could be handled partially by LoRa and then providing a user interface that helped communicate with the user about partial and complete data transactions. As third challenge, we explored the possibility of increasing the coverage of LoRa even further by extending its reach through a multi-hop architecture. Since commercial LoRa only supports single-hop star topologies, extending the reach involved the implementation of a novel *compound repeater* architecture that allows off-the-shelf LoRa to

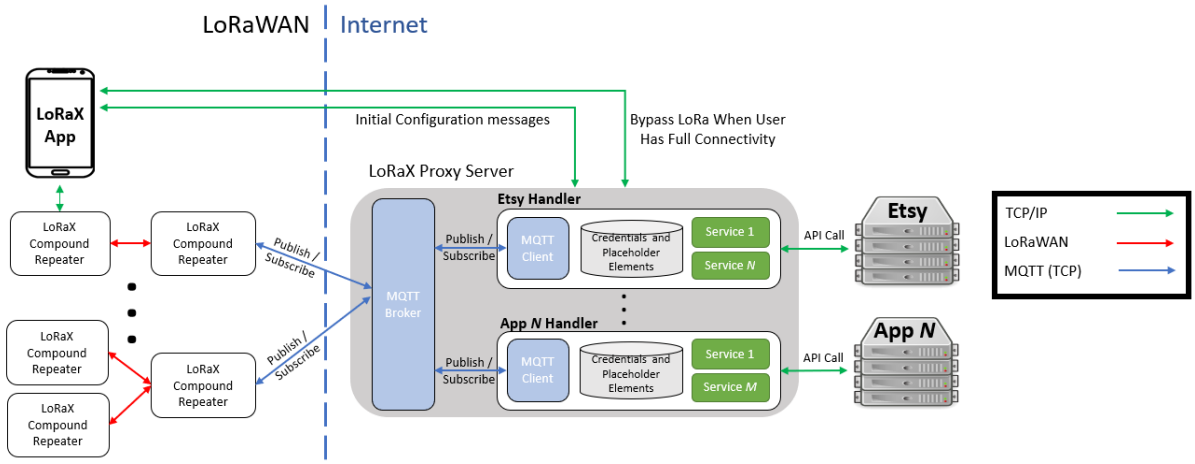

Fig. 3. LoRaX proxy server.

operate in a bi-directional, multi-hop configuration. We provide more extensive background detail about standard LoRa operation and configuration in Appendix A to help contextualize the details of our implementation. We organize descriptions of the LoRaX system design around the three challenges that our system addresses.

3.1 Challenge 1: Effectively Utilize and Integrate Two Regimes

Fundamental to the LoRaX system is effectively utilizing the low data rate LoRa network and integrating into the higher bandwidth Internet. Figure 3 depicts the LoRaX Proxy Server that orchestrates the dual network use. As shown, the proxy participates on the LoRa side in the publish/subscribe protocol that is used to distribute LoRa messages. The proxy participates on the Internet side by making API calls to services such as Etsy. To translate between low data rate, highly compact LoRa messages and full-fledged API calls, the proxy maintains state on a per-user and per-application service basis.

Use of LoRax requires an initial registration step with the proxy, as well as setup of information to support each desired service. This registration must be done over an Internet connection; the LoRa network is too slow to feasibly support setup. Initial setup involves supplying the user's account credentials, any API keys, and any service-specific details. As described in Section 3.2, a key principle of the LoRaX system is the ability to reference and provide placeholder elements until the user can travel for Internet access and update the placeholders with full resolution objects. The initial configuration includes creating and storing these placeholder elements at the proxy.

After setup, when a user has only LoRa access, their request is encoded into a LoRaX message and routed through multiple compound repeater hops to the LoRaX proxy server. When the message reaches a compound repeater that is connected to the Internet, it is forwarded by the Application Server to the LoRaX proxy server as an MQTT message that contains the formatted payload data that will be used as parameters for a RESTful API call. The MQTT message is published to an MQTT broker located on the LoRaX proxy server where it is then pushed to the proxy service that handles API calls for the application (e.g., Etsy). Once the appropriate proxy service receives the MQTT message, it then creates an HTTP request based on the REST API of the web-based service to which it corresponds. The LoRaX proxy server then handles the response from the RESTful API call, and forwards an encoded response back to the UE over the reverse pathway using LoRaX messages. For our

proof-of-concept, we designed a compact message representation that allowed us to extract the service, API call, and API parameters from a single LoRa message.

In the case where the user has a stable Internet connection, they will access the proxy over a TCP/IP connection bypassing the compound repeaters and supply the system with the listing details. That request is then forwarded from the LoRaX proxy server to the service (Etsy in this example) via its REST API just as in the case with no Internet connection. The response here is sent back to the UE through a TCP/IP connection.

3.2 Challenge 2: Re-architect Application Services

3.2.1 Initiate-then-Update. One of the contributions of our work is the reframing of networked service transactions through an *initiate-then-update* paradigm, which provides users with more pervasive access to services by allowing them to take advantage of multiple network regimes. The initiate-then-update paradigm leverages two common service trends in order to support transactions over a combination of low-bandwidth and high-bandwidth network regimes: (1) services based on transactions of well-defined, structured content objects; and (2) APIs based on CRUD (create, read, update, delete) database operations. *Our insight is that the creation of new content objects via APIs often requires a collection of individual information units, some viable for low data rate transmission and others requiring high bandwidth. Further, uploading partial information can preserve much of the value of full objects.* When information can be represented with few bits it can be effectively transmitted over low-bandwidth regimes; larger information requires high-bandwidth.

Since most APIs focus on content objects as the unit of transaction (not the individual information units that comprise a content object), typical API use occurs only when there is sufficient bandwidth to support the large information units that are part of the content object. The initiate-then-update paradigm disrupts typical client interactions with APIs by transmitting the small information units associated with a content object as soon as possible to a proxy (i.e., when pervasive low-bandwidth connectivity is available). To bridge the gap left by large information units required to create the content object, initiate-then-update relies on the proxy to provide "information placeholders"–stock data options that are used temporarily in the place of information units that could not be transmitted over the available regime. In the case of LoRaX, the LoRaX proxy server hosts a database that contains information placeholders for each service that it provides proxy access. Initiate-then-update next leverages the fact that most APIs enable a way to update previously created content objects. When a high-bandwidth regime becomes available, the initiate-then-update paradigm updates a content object containing information placeholders with the actual large information that was part of the content object formed by the application client.

3.2.2 UI Design and Etsy Prototype. A LoRaX-supported app running on a UE device has two possible connectivity options for communication with Internet-based services–a low data rate communication channel supported by LoRa and a comparatively high data rate channel supported by wireless broadband. While there have been several other platforms that use heterogeneous network channels for sending data based on availability or quality of the connection [12, 29, 32, 107], these approaches seek to make the app transition between different network interfaces invisible to the user. LoRaX takes a fundamentally different approach by making the communication channels (and their limitations) transparent to the user while also guiding the user through optimal use. We hypothesize that this transparency may be useful in areas with sparse or challenged broadband connectivity because information about *when* broadband connectivity is needed to achieve certain tasks can be used to help direct users to mobilize to places *where* broadband connectivity is available.

To investigate the potential and limitations of making network capabilities transparent to the user and supporting (at the application layer) transitions from low bandwidth to high bandwidth regimes, we designed a prototype version of a LoRaX UI to the Etsy app using the high fidelity prototyping tool Figma [36]. Doing so requires making design decisions about how to convey to the user which networks are currently available and

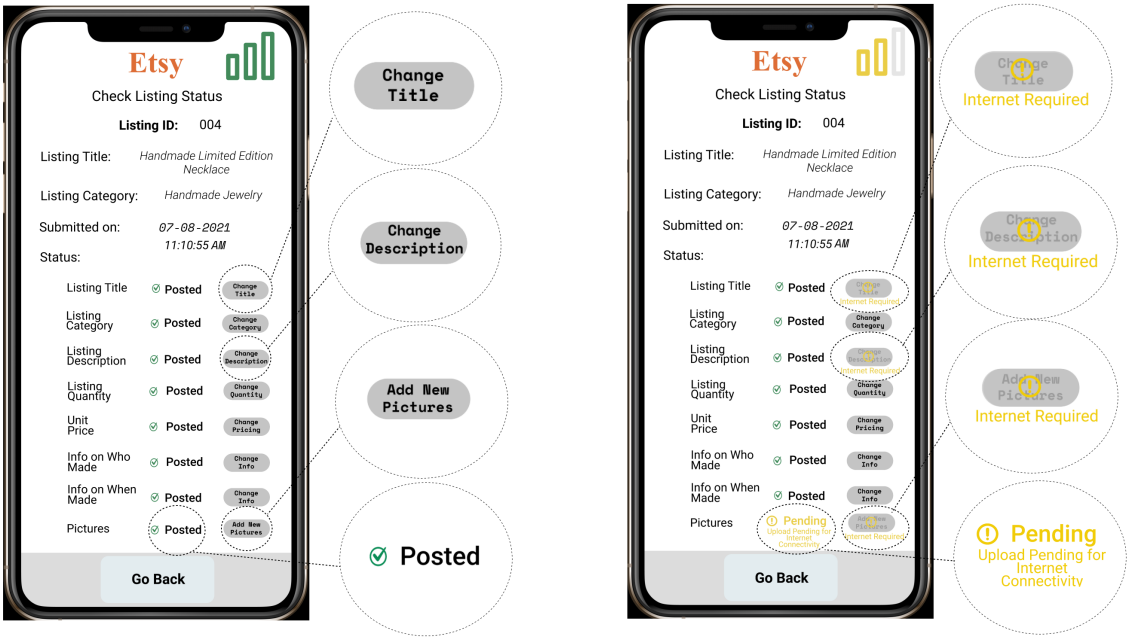

(a) Broadband connectivity

(b) LoRa connectivity

Fig. 4. Example of Etsy LoRaX UI in different connectivity statuses.

what interactivity options are currently possible, as well as carrying state for transactions that have been initiated and are awaiting update. We chose Figma because it allowed us to develop the prototype 'mock' user interface for the LoRaX smartphone app without developing an actual functioning LoRaX app. Figma was originally developed as a web-based tool for designers worldwide to collaborate on various design projects [36]. Figma supports creating interactive experiences on prototype user interfaces without requiring any under the hood back-end codes (e.g., codes for client-server communication) which also allowed us to simulate scenarios from different network regimes on the mock user interface for the LoRaX app.

As shown in Figure 4, the prototype simulated the user experience of accessing Etsy API calls with both (a) full broadband connectivity and (b) limited connectivity provided through LoRa. The prototype supports five common actions: linking an existing Etsy user account to the app, creating a new product listing, reviewing and modifying existing listing information, and checking notifications and alerts about listings. When broadband connectivity is available (Figure 4a), the user interface is designed to explicitly reveal which Etsy services are available in addition to providing information about the status of the UE's end-to-end connectivity to Etsy servers (green bars). Conversely, when only LoRaX connectivity is available (Figure 4b), the app "grays out" the services that cannot be supported with limited connectivity while also demonstrating the overall limited status of the end-to-end connection using LoRaX (yellow bars). Thus the design philosophy is four-fold: (1) keep a consistent look-and-feel between the two connectivity regimes to ease cognitive load and make the transition from one regime to another relatively seamless at the UI level, (2) provide multiple indications of the current connectivity; and (3) make explicit what operations can and cannot be performed, and (4) convey status for initiated operations that are pending update. In Sections 4.3 and 4.4 we describe the results of user testing with this prototype UI.

3.3 Challenge 3: Extend Reach

To be most cost effective, we want the LoRa portion of LoRaX to extend beyond the reach of a single gateway, by using LoRa in a multi-hop configuration. This represents a departure from the traditional data flow enabled by off-the-shelf LoRa, where a node connects to one or more gateways which in turn connect to a single network server, in a star-of-stars topology with the network server as the hub. While the star-of-stars topology maximizes data flow from many data collection points (sensors and IoT devices) to a single data sink (the network server), the multi-hop topology seeks to extend LoRa's reach beyond the transmission range of a single node-gateway link. Doing so is not straightforward because the LoRa protocol does not support direct node-to-node communication or direct gateway-to-gateway communication [61]. Mesh LoRa extends the reach of traditional LoRa by making every end node capable of acting as a router and able to directly communicate with other end nodes to forward their payloads to a gateway in a multi-hop fashion [49, 58]. However, implementing mesh LoRa requires developing custom hardware as well as protocol stacks to achieve direct end node to end node communication [49, 58]. This prevents mesh LoRa from taking immediate advantage of off-the-shelf hardware as well as the already available protocol suites. Because we do not require a mesh network, and we seek to use off-the-shelf hardware, we use the simpler approach of coupling nodes with gateways.

To implement a multi-hop topology, we designed *compound repeaters* (see Figure 2 and Figure 5) that combine a LoRa node and gateway together into one component. A compound repeater receives data over LoRa on its gateway, transfers the data over a serial connection to its node, and transmits from the node to the next compound repeater over LoRa. Figure 2 illustrates a multi-hop network, with one compound repeater functioning as a last hop that connects to UEs via a LoRa node that functions as customer premise equipment (CPE) and one compound repeater functioning as a connection point to the Internet via a LoRaX proxy server. The intermediate compound repeater illustrates the mid path configuration.

Compound repeaters address an additional logistical challenge presented by commercially-available LoRa devices in order to realize a multi-hop network, namely the limitations of the LoRa timing protocol as implemented by Class A LoRa devices (see Appendix A.1). In theory, multi-hop topology would be best implemented with a Class C device, which enables LoRa end devices to transmit and receive simultaneously; however, Class C devices are still not readily available as off-the-shelf development boards. To overcome the asymmetry between uplink and downlink transmission timing with Class A LoRa devices and the lack of end device-to-end device and gateway-to-gateway communication capabilities, the compound repeater couples LoRa nodes and LoRa gateways over a serial connection. This allows a compound repeater node to always be able to receive (by the attached gateway) from another compound repeater or transmit (by the attached node) to another compound repeater, thus achieving simultaneous transmit and receive capability without being restricted by the limitations that accompany the receive window timings of Class A LoRa devices.

Fig. 5. Implementation of LoRaX compound repeaters.

To accomplish simple routing over compound repeaters, we implemented an Application Server that determined whether packets needed to be forwarded to the next hop via LoRa or if they needed to be forwarded on to either the LoRaX proxy server or a LoRa CPE. Since the Conduit gateway device is powered by mLinux [95], we were able to implement an Application Server directly on the Conduit device. We implemented the Application Server logic with Node-RED [38], an event-driven programming language for wiring together hardware devices. The

Application Server is able to access packet payloads received by the gateway on the same compound repeater. If the payload is intended for a service for the LoRaX proxy server connected to the current compound repeater, then the packet is forwarded to that service via TCP/IP to the LoRaX proxy server. Conversely, if the payload is intended for a UE, then the packet is forwarded over LoRa by the gateway on the last-mile compound repeater and received by the LoRa CPE node. If a compound repeater is not connected to the intended packet destination (either LoRaX proxy server or LoRa CPE node), the Application Server will forward packets to the node on the current compound repeater via serial so it can be to be forwarded on to its final destination.

3.4 Additional Implementation Details

Our implementation uses off-the-shelf hardware. Specifically, we implement the UE-to-LoRa functionality using a LoRa node equipped with an ESP32 module that provides the node with TCP/IP communication capabilities over WiFi to the device. We use an ESP32 MCU developed by the Espressif Systems that comes equipped with a 32-bit microprocessor, integrated WiFi, Bluetooth, and support circuitry in a compact 39-pin module [94]. To support LoRa communications, the LoRa node uses The Things Uno (TTUno) [72], an off-the-shelf LoRa node development board based on Arduino Leonardo [5]. TTUno employs an 8-bit Atmel Atmega32u4 micro-controller paired with a Microchip RN2903 LoRa radio module, which handles the MAC and PHY layers of LoRa communications. The TTUno and ESP32 are connected to each other with a custom-designed Arduino shield that provides serial communication. Our implementation of the gateway uses a Multitech Conduit Programmable Gateway [96], which can be configured to function both as a LoRa Gateway and LoRa Network Server.

4 EVALUATION

We have three goals in our evaluation. The first is modest: a proof-of-concept that our prototype LoRaX system works end-to-end as intended with off-the-shelf LoRa devices plus our Arduino shield and software. This modest test allowed us to understand first-hand the barriers in using off-the-shelf and development platforms for novel LoRa uses; we reflect on this experience in Section 5.2. The second goal is to probe performance barriers to use of LoRaX as a solution for enabling Internet service use without broadband connectivity. The third goal is to assess usability of LoRaX, and in particular how well the user interface supports users in navigating regimes of different performance.

4.1 Measurement Methodology

Because we have repurposed LoRa for human-to-Internet service use, rather than its usual machine-to-machine use, the performance metrics of interest differ from those in prior work. (See Section 6 for further discussion of common metrics in prior work.) In particular, we focus on the end-to-end latency that a user experiences when initiating an Internet service action on an end device as well as the round trip time to receive a confirmation. Additionally, while telemetry applications can often tolerate data transmission loss, we require reliable transmission or notification of failure. Therefore we also measure packet loss that would lead to either a failure report to the user or retransmissions that increase effective latency.

The delays in the LoRa portion of the end-to-end path dominate performance, hence our measurement setup omits the proxy to Internet service component. In particular, as illustrated in Figure 6, our measurement setup utilizes two compound repeaters (Figure 5) with the ESP32 acting as a means for an end-user device to insert messages into the system. In these tests, a Raspberry Pi simulates an end-user device by sending LoRaX messages to the ESP32 on Compound Repeater A. The messages are then forwarded from the node over LoRa to Compound Repeater B. Compound Repeater B simulated a message-response cycle by returning LoRaX messages back to the end-user device via Compound Repeater A over LoRa. Compound Repeater A then forwarded the message back to the end-user device via TCP/IP from the ESP32. To evaluate the performance of the system, both LoRa gateways

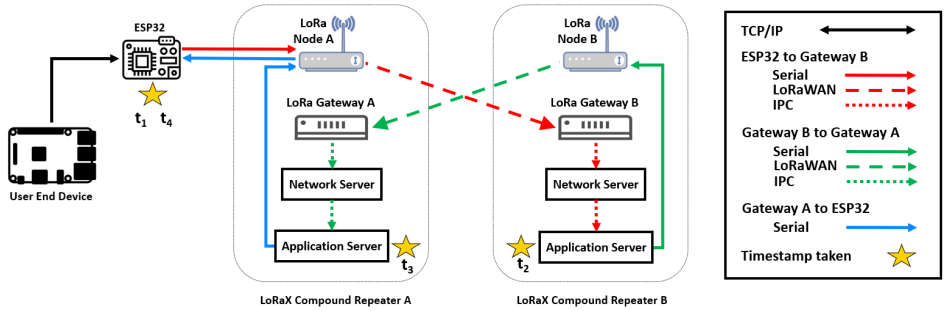

Fig. 6. Measurement configuration including points (stars) where NTP-synchronized timestamps were collected.

and the ESP32 were synchronized to the same NTP server. Timestamps were taken at four locations: at the ESP32 when the LoRaX message was received from the end-user device (t_1), at the LoRa gateway on Compound Repeater B (t_2), at the LoRa gateway on Compound Repeater A (t_3), and once again when the LoRaX message returned to the ESP (t_4). RTT was measured by $t_4 - t_1$, uplink delay by $t_2 - t_1$, and downlink delay by $t_4 - t_2$. For each test, we measured delays for 250 round trip messages at each distance and line of sight condition. When packets were dropped, we re-transmitted until we achieved success using 20 s as an arbitrary (conservative) timeout value.

For our test environment, we used The Things Network's adaptive data rate (ADR) mechanism, which dynamically optimizes the parameters used to determine the spreading factor (modulation rate), bandwidth, and transmission power used by the LoRa Microchip RN2903. ADR is recommended for devices that are generally static, which is the ideal deployment scenario for the compound repeaters. Notably, for all of our measurements, the spreading factor remained at a consistent SF=7, information we leverage in our trace-driven simulations. The LoRa gateways were configured to use an antenna gain of 3 dBi. All tests use a packet payload size of 13 bytes. We restricted the payload sizes to be 13 bytes as this was the minimum number of bytes required for encoding the necessary data to make a successful Etsy API call.

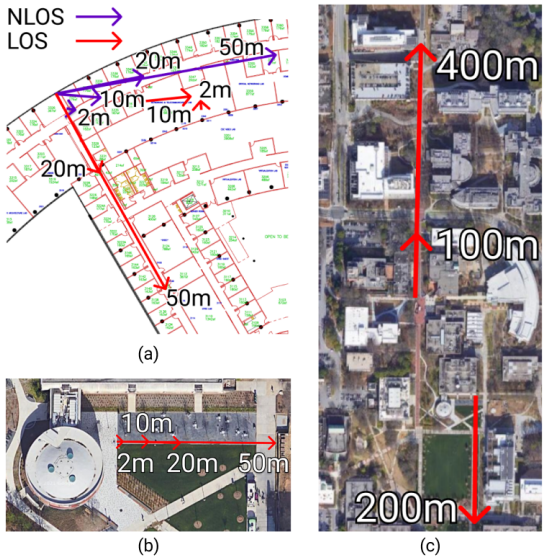

Fig. 7. Measurement locations for (a) indoor tests, (b) 2-20 m outdoor tests, and (c) 100-400 m outdoor tests.

4.2 Measurement Results

To establish a baseline on system performance, we measured end-to-end delay in line of sight (LOS) and non-line of sight (NLOS) indoor environments in a university building. For indoor NLOS tests, the compound repeaters were placed on the opposite sides of building corridors, forcing the LoRa transmissions to pass through multiple concrete walls as shown in Figure 7a. LOS measurements occurred in hallways.

We extended our measurements to a less-controlled outdoor environment depicted in Figure 7b and 7c to evaluate the robustness of LoRaX latency and PLR in the target setting and with longer distances between compound repeaters. While we were aware that our range would be limited due to the limitations of the PCB

.Location	Distance	LOS	Mean $Delay_{up}$ $t_2 - t_1$ (ms)	Std. dev. $Delay_{up}$ $t_2 - t_1$ (ms)	Mean $Delay_{down}$ $t_4 - t_2$ (ms)	Std. dev. $Delay_{down}$ $t_4 - t_2$ (ms)	Mean $Delay_{LoRa}$ $t_3 - t_2$ (ms)	Std. dev. $Delay_{LoRa}$ $t_3 - t_2$ (ms)	PLR (%)
Indoor	2m	Y	1,328.06	90.92	1,854.39	90.53	295.12	98.24	*0**
Indoor	10m	Y	1,336.9	62.54	1,845.36	62.37	291.24	106.81	1.96
Indoor	20m	Y	1,324.76	66.58	1,857.36	66.39	*281.18**	84.77	13.5
Indoor	50m	Y	1,313.11	78.03	1,869.18	77.86	298.19	102.52	1.96
Indoor	2m	N	1,320.86	84.95	1,860.56	84.7	*302.21***	110.96	0.4
Indoor	10m	N	*1,352.35***	98.7	*1,829.7**	98.31	283.84	75.93	9.42
Indoor	20m	N	1,334.16	72.74	1,846.98	72.69	288.88	114.9	0.79
Indoor	50m	N	1,342.19	81.52	1,839.31	81.72	293.72	*119.42***	1.57
Outdoor	2m	Y	1,337.86	86.1	1,843.9	86.24	293.83	110.34	1.57
Outdoor	10m	Y	1,304.14	96.88	1,877.16	97.03	294.91	78.47	0.8
Outdoor	20m	Y	1,300.01	69.64	1,881.78	69.57	288.55	*70.46**	*0**
Outdoor	50m	Y	*1,286.22**	*120.57***	*1,895.35***	*120.63***	298.05	86.93	1.19
Outdoor	100m	Y	1,322.71	83.74	1,858.82	83.61	286.66	77.83	0.8
Outdoor	200m	Y	1,296.15	*57.02**	1,885.05	*56.97**	293.00	110.57	4.58
Outdoor	400m	Y	1,307.22	63.15	1,874.27	63.11	291.25	98.93	*28.77***

Table 1. Overview of baseline indoor and outdoor measurements. (*Minimum observed values for each column; **Maximum observed values for each column.)

antennas onboard the LoRa nodes on the compound repeaters, our goal was to generally characterize the stability of performance with increases in distance so that we could extrapolate performance to more extended scenarios.

Our measurements are summarized in Table 1 where the contribution of each component of delay is isolated and the standard deviation is also provided, with the packet loss rate on the round trip path indicated in the last column. Four trends are notable. First, the mean delays on all components of the paths are remarkably consistent, with the difference between lowest to highest values no more than 10% of the mean. Second, the standard deviations are more substantial especially over the LoRa link, where we see standard deviation values as high as 50% of the mean. The uplink and downlink delay values are dominated by the delay between the ESP32 and the LoRa nodes, likely because it is serial, while the LoRa link is subject to much more environmental variation. Third, the indoor LOS and NLOS values show little difference indicating that for these distances and obstacles NLOS is not a problem. Finally, the outdoor data shows a generally increasing trend in packet loss rate as distance increases, with a notable increase going from 200m to 400m, suggesting that 400m is the practical limit for our off-the-shelf components.

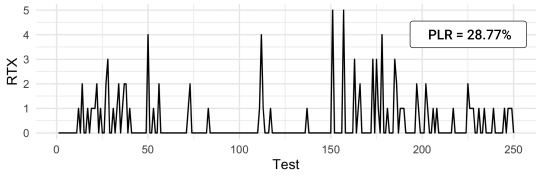

Fig. 8. Retransmission frequency for each test at 400 m.

Trace-driven Simulation. In order to understand how our results might translate to a more complex implementation with a greater number of hops, we used a trace-driven simulation approach based on our measurements at 400m. Specifically, we seek to understand how our measurements of LoRaX RTT and LoRa link delay translate to RTT values for a single Etsy API call made over LoRaX. We contextualize these RTT projections by comparing them to walking times (for hop distances up to 400 m) and driving times required to traverse the same distances. Our measurements and the measurements of others [59] indicate that the most significant impact on LoRa end-to-end delay is packet loss rate, which increases with larger packet payload sizes and NLOS links. Based on the consistency of our measured RTT and LoRa delays across increasing distances (Table 1) as well as other measurement-based studies of LoRa delay across more substantial distances (up to 9 km) [59, 60], we assume that

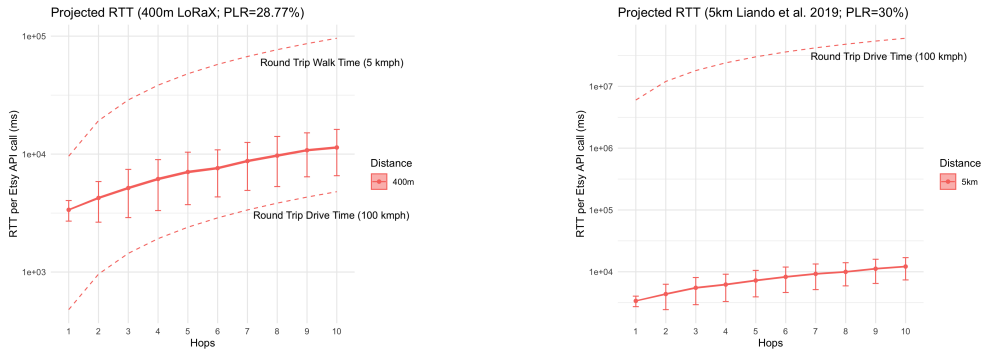

(a) (b)

Fig. 9. Projected RTT per API call relative to average drive time for the same distance to nearest Internet access point for (a) 400m hops and (b) 5km hops.

projected LoRa transmission delays will remain consistent at greater distances (up to 5 km) given our parameters of 13 byte payloads and Spreading Factor 7. We provide access to the software used to run our trace-driven simulation via GitHub[2].

For each multi-hop scenario that we simulate, we randomly select a delay value from our measured distribution of transmission delays from the ESP32 to the next hop compound repeater (ESP32 to Gateway B in Figure 6) to represent the delay of the first outgoing transmission from the LoRaX compound repeater attached to the UE to the next LoRaX compound repeater in the path (hop_1). We similarly select a value from our measured distribution of downlink transmission delays to represent the delay of the last transmission on the return pathway (hop_{2n}). For every other hop on the uplink pathway (hop_2 to hop_n) and downlink return pathway (hop_{n+1} to hop_{2n-1}), we randomly select a value from our measured distribution of LoRa transmission delay values (Gateway B to Gateway A in Figure 6). To account for packet loss, we use the PLR that we observed at 400 m to determine whether a transmission over a particular hop was successful or not. Until we are able to successfully transmit in our simulation, we add $2 * Mean\_LoRa\_TX\_Delay$ to our overall projected RTT, simulating the time it would take to detect a potential packet loss. The projected RTT represents the sum of delay values selected for each hop on the uplink and downlink path in addition to the projected amount of time it would take to detect all dropped packets. We run our simulation 1,000 times for each hop count.

We graph the mean projected RTT as a function hops in Figure 9, where the hop lengths are set to 400 m and 5 km. In each graph, the error bars represent the standard deviation of RTT values. For a simulation of 10 hops at 400 m, we project a mean RTT of 11.4 s (σ = 4.8 s) and a mean PLR of 26.5% (σ = 8.3%). We compare these projections to travel times associated with different modes of transport that a user might use to travel to the nearest point of Internet connectivity[3]. Given the relatively short hop length of 400 m, we compare our projections to walking and driving. For walking, we assume an average travel rate of 5 kmph [15]; for driving, we assume an average travel rate of 100 kmph. At this distance, driving times can be 2.4× faster than the RTT for a 10-hop configuration. However, we note that this may be a relatively short distance and users might be more apt to walk to the nearest point of Internet availability than drive. In this case, the projected RTT for a 10-hop configuration of LoRaX is 8.4× faster than the round-trip walk time. This means that in the time it takes a user to walk to a place of Internet availability, LoRaX could make over eight API calls on their behalf.

[2]https://github.com/CANIS-NAU/COMPASS2022.

[3]We assume that the nearest point of Internet connectivity is equivalent to the end-to-end path length.

We recognize that 400 m falls short of LoRa's promised transmission range of 15km in rural, outdoor settings [59, 61]. As we discuss in Section 5.2, a number of measurement studies have recognized the relative immaturity of off-the-shelf LoRa hardware [89, 100, 101]. In order to characterize how LoRaX might perform given improved hardware, we make some assumptions that allow us to project RTT values for hop lengths of 5 km. Based on an extensive measurement study of LoRaX performance using software defined radios (rather than off-the-shelf devices), Liando et al. were able to demonstrate that LoRa was able to transmit data across LOS links of up to 5 km with a PLR of 30% when the spreading factor was set to SF=7 and the payload length was 10 bytes. While the study did not include RTT or delay measurements, the authors noted that transmission delay was more dependent on payload size and spreading factor, rather than distance [59]. Based on the similarity of our own parameters (13 bytes and SF=7) to those of Liandro et al., we ran our simulation using the same delay distributions as those we measured for 400 m, but used a PLR of 30% to determine retransmission in the simulation of multiple hops at 5 km in length. We observed similar projected mean RTT values (μ = 12.2 s and σ = 4.8 s at 10 hops) as those projected for 400 m hops. However, the most notable distinction of the extended hop length is the relationship to round trip drive times across equivalent distances. Assuming that users would drive at a rate of 100 kmph to the nearest place of connectivity, LoRaX RTTs are 1,774-4,933× faster than the driving times, indicating that in the amount of time it would take for a user to drive to the nearest point of connectivity and back home again, thousands of Etsy API calls could be generated. For an Etsy merchant, this translates directly into potential revenue as they would be able to spend more time posting shop updates and advertising merchandise before having to travel to fulfill order requests.

4.3 User Testing Methodology

As mentioned earlier in Section 3.2.2, previous works that used heterogeneous network channels for communicating depending on the quality or availability of the connection made the transition between different network regimes invisible to the user [12, 29, 32, 107]. Because LoRaX takes a fundamentally different approach by making the states of the communication channels and their limitations transparent to users, we needed to determine effective ways to convey the available network states and how they might affect availability of application functionalities to the users. To do so, we designed and conducted preliminary user testings with a Wizard of Oz testing approach, using a prototype 'mock' user interface that simulates connections to Etsy using LoRaX as described in Section 3.2.2.

We conducted the preliminary user testing sessions over the Zoom video conferencing platform in accordance with Institutional Review Board (IRB) ethical and privacy guidelines and our institution's COVID-19 guidelines. Participation was voluntary and users were recruited using a snowball sampling method [41] by reaching out to local crafting communities and known Etsy users. Participants were not compensated for their participation. Zoom meeting audio and video were recorded after obtaining the participant's verbal permission. Overall, we were able to conduct the preliminary user testing sessions with eleven different users (n = 11) whose ages ranged from 20-45 years old. All of the users reported that they used smartphone apps multiple times a day and three of the users confirmed that they had used e-commerce platforms like Etsy. We provide background information about all the test users who participated in the preliminary user testing process in Table 2.

Each preliminary user testing session was conducted in two distinct phases. During the first phase our goal was to understand usability in each connectivity regime. Each user was given a list of tasks to complete including configuring the app for first-time use, creating and checking item listings on Etsy, modifying an already posted item listing, and reviewing Etsy

(a) (b) (c)

Fig. 10. Icons for representing connectivity status (LoRa connectivity shown).

Gender	Age (Years)	Education Level	Etsy Experience	Label
Male	23	Attending college for a 4-year degree	Not an Etsy user	U1
Non-binary	22	Attending college for a 4-year degree	Not an Etsy user	U2
Male	23	Attending college for a 4-year degree	Not an Etsy user	U3
Female	22	Attending college for a 4-year degree	Not an Etsy user	U4
Male	21	Attending college for a 4-year degree	Not an Etsy user	U5
Male	24	Enrolled in or completed a graduate degree	Not an Etsy user	U6
Male	23	BA or BS college degree	Not an Etsy user	U7
Female	23	Enrolled in or completed a graduate degree	Not an Etsy user	U8
Female	23	BA or BS college degree	Has been using Etsy as a seller for more than 6 months but less than a year. Uses Etsy 2-3 times a month.	U9
Female	45	AA degree from a vocational, technical, junior college, or community college	Has been using Etsy as a buyer for more than a year. Uses Etsy 2-3 times a year.	U10
Female	25	Enrolled in or completed a graduate degree	Has been using Etsy as a seller for more than a year. Uses Etsy 2-3 times a year.	U11

Table 2. Background information about the test users participating in the preliminary user testing.

notifications. Users were first asked to perform the tasks on a version of the user interface (UI) that assumed the high-bandwidth network regime. Then, users were asked to perform similar tasks with a version of the UI that represented the low-bandwidth network regime.

During the second phase of testing, our goal was to understand what connectivity icon designs were effective in conveying regime to a user. Users were shown three different sets of icon designs and asked to rank them according to their personal preferences. Icons included ascending bars (Figure 10a), a meter (Figure 10b), and a face (Figure 10c). Icons were color-coded to indicate whether the connectivity status included broadband connectivity (green), LoRa connectivity (yellow), or no connectivity (red). Icons were modified slightly in different connectivity statuses to provide additional indicators about connectivity (e.g., for broadband connectivity, the bars had three colored bars as opposed to only two colored bars for LoRa connectivity). To help account for order bias, we presented the options in a random order for each user. The users were then asked to provide a brief explanation of their ranking choices. Using methodologies from participatory design [69], sessions ended with asking the users for general feedback on the prototype UI, specifically what they liked, did not like, and what they would change about the interface based on their experience.

We analyzed the recordings of the preliminary user testing sessions and user feedback using a grounded theory approach that involves qualitatively identifying emergent themes ("open codes") in user's responses to questions about their perceptions of the interface [93]. As we collected more responses, we used the constant comparative method to categorize open codes into more general themes ("axial codes") that reflected some of the collective perspectives on the user interface [40].

4.4 Findings from Preliminary User Testing Sessions

During the first phase of all preliminary user testing sessions, the users were given a two part list of various tasks to complete using the prototype 'mock' LoRaX app user interface. Apart from minor navigational difficulties (e.g., clicking the wrong button or mistakenly choosing the wrong option), all of the test users (n = 11) could successfully complete all the tasks on the given list using the mock user interface. While analyzing the data obtained from this stage of user testing, we focused on finding what factors (e.g., icons, text labels, color coding, etc.) helped users to better understand about the different network regimes and how the available network regime affected the functionalities of the prototype LoRaX app. These analyses and implicit feedback from the users were complemented by the explicit user feedback that we obtained from the second phase of the user testing.

Two major themes emerged from user feedback. The first was about the iconography that should be used to denote information about current network regime. The second theme focused on the use of explicit labeling to provide information both about the network regime and the services available while connected through that regime.

Iconography. From our analysis of the data obtained from the user testing sessions, we found that when navigating between different network regimes (e.g., high-bandwidth and low-bandwidth) users preferred icons that carry clear meanings and are recognizable from similar settings. For example, one of the test users quoted why they chose network bar icons as their first preference for representing the current network regime: *"On my iPhone that's how the connectivity shows. That's the most intuitive"* (U11). Another user explicitly mentioned universality when explaining their preference for network bars as an icon indicating the network regime: *"I would just go with the network bars. That's a universal symbol for having connectivity"* (U7).

However, objects or symbols which the user comes across in everyday life might not always be obvious in carrying the message when they are not used in their typical contexts. For example, another test user explained a speedometer symbol might not be a good choice for representing connectivity status by stating: *"I understand that the speedometer is like Internet speed. But I don't know, I just haven't seen that before. I wouldn't say that's super intuitive"* (U11).

Explicit Labeling. Besides using appropriate symbols to make the user aware of the different connectivity statuses, we also investigated user perceptions of text labels to help clarify how different network regimes impacted the availability of services available through a LoRaX-enabled app. The general consensus found in the feedback from different test users is that legible text labels can help the user (especially if the user is not very tech savvy) make a connection between the network regime and the availability of services.

We also asked for users' opinions about attaching explicit text labels to icons used to display the current network regime. According to the users' feedback, such labels are not necessary if the icon symbol is universally understood. However, if less common symbols are used, such text labels will certainly help the user to understand the meaning behind the symbol more easily. According to one of the test users: *"If you're going to use the speedometer one or the smiley face then the labels are important. Because, I feel like people already kind of understand the bars mean connection"* (U4).

5 DISCUSSION

Our implementation and evaluation of LoRaX illuminates the possibilities and logistical challenges associated with pervasive low data rate messaging in the context of an ITU system. Here we discuss the implications of our work for enabling the ITU paradigm, challenges associated with limitations of off-the-shelf LoRa hardware, economic feasibility of a LoRaX network, and some of the limitations of our approach to design and evaluation.

5.1 General Applicability

While our work in this paper evaluates the feasibility of using LoRa as a pervasive low data rate channel that extends Internet services, it points to the possibility of a paradigm wherein pervasive low data rate networks are able to provide scaled-down interactions and services routinely offered over high capacity networks, but where these high capacity networks are prohibitive due to physical range, challenging terrain, and/or cost of deployment. As cyberphysical systems and smart community architectures expand beyond high-density urban environments [30, 55, 91] and seek to support highly mobile networked agents across dynamically networked environments [31, 87], it is useful to consider how pervasive low data rate networks might add value in regions where high capacity networks are not available. For example, in the case of designing smart and connected communities in more rural settings, a pervasive low data rate network channel might be used to continuously monitor the presence of mobile networked devices (e.g., a herd of livestock tagged with biosensors or an oncoming fleet of self-driving vehicles) and direct software-defined edge-based resources associated with high data rate network infrastructure to configure themselves to anticipate an increase in traffic load as part of smart infrastructure practices. Alternatively, in the case of disaster monitoring using unmanned mobile sensing units [87], a pervasive low data rate channel might be able to assist high capacity networks in coordinating and allocating dynamic spectrum resources based on agent mobility and data transmission needs [31, 88]. A non-trivial point to note here is that the integration and use of a pervasive low data rate network will naturally incur considerable delay on the communication process compared to its counterpart high capacity network. As demonstrated through the trace driven simulation results discussed in Section 4.2, in remote areas where high capacity networks are unavailable, people could easily face larger delays if they chose to walk or drive to the nearest point of high capacity network connectivity. We hypothesize that people would be tolerant of the lags stemming from the integration and use of pervasive low data rate networks when the alternative options (e.g., walking or driving to the nearest point of high capacity network connectivity) incur considerably higher lags.

Network Model. To begin to formalize the performance tradeoffs, consider an environment such as the one in Figure 1 where multiple networks with different geographic reach and different performance profiles co-exist. Imagine that one of these networks, N_1, has relatively low performance yet large geographic reach. The second network, N_2, has relatively high performance yet is geographically limited in reach. For simplicity, let the network performance be expressed as the average bandwidth b_1 for N_1 and b_2 for N_2, where $b_1 << b_2$. Suppose a user is in the region covered by N_1 but not N_2 and has a file of size s to upload. Further, suppose the travel time for this user to reach an access point of N_2 is T. In this simple model, the additional network N_1 offers benefit whenever $s/b_1 < 2T + s/b_2$, where the factor of 2 is to allow the user to travel to the access point and back home. If b_2 is very large then this is approximated by the condition $s/b_1 < 2T$.

For example, if $b_1 = 300kbps$ and (one way) travel time is 1 hour then files that are smaller than 270 Kbytes can be more efficiently transferred over the slower network N_1. If $b_1 = 37.5kbps$ and travel time is 1 hour then files that are smaller than 38 Mbytes are more efficient over N_1. (We use these bandwidth values for b_1 based on LoRa technology estimates.) The travel distance associated with 1 hour of travel depends, of course, on the travel modality. If a car is not available and the user must depend, for example, on a bicycle, then the user can cover about 25 kilometers in an hour. If the user is further constrained by schedule commitments, such as lack of access to childcare, the "travel" time could easily include realities such as waiting for another adult to come home before travel is possible. Thus even relatively short travel distances such as 5 km when accompanied by a 45 minute wait until travel is possible could easily tilt the preference to the slower network.

To incorporate the ITU paradigm in this simple model, suppose further that there are two sizes to the image, s_1 and s_2, where $s_1 << s_2$. The images are not equivalent in utility to the user, however the smaller image suffices as a temporary placeholder, and notably has much higher utility than no image. In this regime, the network N_1 offers benefit whenever $s_1/b_1 < 2T + s_2/b_2$. If we again assume that b_2 is large, this is approximated by $s_1/b_1 < 2T$.

With a factor of 10 compression from s_2 to s_1, the advantage of N_1 extends to original images of size between 2.7 and 380 Mbytes using the speed range and 1 hour travel time budget above. While simple, this model illustrates the interplay between an abstracted version of the application, the ITU paradigm, and the multi-network regime.

5.2 Maturity of OTS LoRa Development Platforms

As we implemented the LoRaX system using off-the-shelf (OTS) LoRa development hardware, we experienced several logistical challenges associated with the maturity and interoperability of different LoRa systems. At the time of our equipment purchase (2018), the two major commercial options for LoRa development boards were the TTN Uno (combines Arduino Leonardo with a Microchip RN2903 LoRa shield), which uses the TheThingsNetwork library for Arduino[4]; and the Dragino LoRa shield coupled with Arduino Leonardo, which uses the LMIC library[5]. Based on the quality of available documentation for both the hardware and the supporting libraries, we selected the TTN Uno; however, our actual experimentation with several of the boards revealed that the onboard PCB antenna was not powerful enough to transmit data more than half a kilometer–far shorter than the 10 km promised by the TTN Uno documentation [98]. Notably, we are not alone in facing these engineering challenges when exploring OTS hardware solutions for the development of system prototypes that leverage LoRa. Works investigating the use of OTS LoRa for a number of applications including rural fire monitoring [101] and urban healthcare IoT [100] have noted similar discrepancies in the achievable OTS LoRa node transmission ranges with both TTN Uno and Dragino development shields [89]. While several TTN Uno development forums provided suggestions for adding high-gain antennas to the TTN Uno, these instructions were all external to the formal TTN Uno documentation and required physical modifications of the board–far from the "plug-and-play" capabilities touted by OTS LoRa nodes. Although these types of engineering challenges are not completely insurmountable (as evidenced by in-depth evaluations of LoRa's transmission capabilities [17, 39, 59, 67, 77, 82], our own work and the work of others note that they are open challenges that can prevent novel experimentation and innovative deployments that seek to use OTS LoRa as a data transmission solution [89].

5.3 Economic Feasibility

A key consideration for any proposal to bridge digital divides is the economic cost of the proposed alternative. While technology costs are notoriously difficult for researchers to adequately estimate, we provide an analysis of feasibility based on costs for current off-the-shelf components. For the compound repeater design, the LoRa node is the cheapest component and costs approximately $60-$65. The LoRa gateway component can be found at costs ranging from approximately $400 to $800. It is noteworthy that when Class C LoRaWAN end devices become available, the compound repeater will be replaced by one of these devices and the cost will likely dramatically reduce. Each compound repeater can simultaneously service dozens of households up to 120 where the network becomes too saturated [11]. For the target users of this application, this limitation is unlikely to be realized.

The portion of the LoRaX architecture that is placed in the users' residences (CPE) is composed of much cheaper components. For our experimental setup, we use the same LoRa node from the compound repeater and pair it with an ESP32. With ESP32s being widely available and at a low cost of approximately $5, our experimental CPE comes in at approximately $70. Every household would need one CPE device that any number of smartphones would be able to connect to. When deploying the LoRaX system, there are a number of financing, management, and ownership models that are possible, especially since LoRa equipment is relatively inexpensive, small, and low-power (especially compared to alternatives like local cellular networks and TVWS). One common approach would be to place the burden of financing and installation of the compound repeaters onto the users of the system through monthly service fees [33]. In a "worst-case scenario" where a single compound repeater would be needed

[4] https://github.com/TheThingsNetwork/arduino-device-lib
[5] https://github.com/matthijskooijman/arduino-lmic

to provide service to a single, remote household, the cost of initial deployment could be completely amortized in a year with subscriber payments of $72.50/month. In other scenarios, community taxes or grant programs might fund the initial cost of compound repeater installation and they would simply charge a small fee for the low volume of traffic that would need to be routed out to the Internet. As for the CPE, we have seen multiple approaches to this in other community networks [14, 33]. Some communities choose to include the CPE with a service subscription. Other communities have chosen to make the users purchase the CPE outright, in which case the equipment is theirs to own even after terminating their service. The justification of this approach is that users will take better care of their equipment if they are the ones who own it. When considering long-term management of compound repeater infrastructure, we envision the LoRaX system being something that would be maintained by community network operators (such as those discussed by Potsch et al [79]) as a way to augment their community Internet backbone. However, when imagining a future where blending multiple, complementary network regimes in a single region can increase service access and coordination, we can also envision a model where larger national network service operators take on the task of deploying and maintaining the infrastructure necessary to support the ITU paradigm.

Based on these cost logistics, it is clear that as more users subscribe to a system like LoRaX, the more it drives down the cost of infrastructure deployment and maintenance. Moreover, there are network effects that are possible with increased adoption. As more users come to rely on services offered through an ITU model, there are increased market pressures to design services to be delivered through this paradigm of more gradual and granular service delivery, removing the need for extensive proxy services at the edge of the Internet. If ITU services are deployed through community network operators, this relieves the burden of those operators also needing to maintain extensive proxy services on their network.

5.4 Limitations

Performance evaluation. In theory, a single LoRa gateway is capable of supporting connections on the scale of 15 km [61, 82]. In our evaluation, we were only able to transmit at up to 400 m with LOS (PLR of 28.77%). This was largely due to the limitations of the transmission power of the LoRa module (Microchip RN2903) used by the LoRaX node in our implementation. While these hardware limitations prevented us from collecting measurements at a scale that is truly representative of the promises made by LoRa specifications, we were able to leverage our measurements and data provided through other measurement studies as part of trace-driven simulations that offered a more extended exploration of LoRaX's ability to provide access to Internet-based services.

We note that the indoor RTT's that we observed are about 3-4× as long as those that are observed in a study by Liang et al. focused on measuring RTT achieved by LoRa networks operating in indoor environments [60]. This is largely explained by the fact that the payload size that we supported in our measurements was almost twice as large as those evaluated in the Liang et al. study, which increases the transmission and propagation delays in both the uplink and downlink directions. We also note that the standard deviations of the delays and the PLRs in Table 1 vary non-intuitively with distance. The standard deviations of all the delays are either below 100 ms or stay very close to 100ms. We hypothesize that the standard deviations of the delays can be attributed to the accuracy of NTP over Internet as it can vary between 5–100 ms and is easily affected by delay scenarios on the network [104]. Additionally, we theorize that some portion of the PLRs that we observed are due to random instances of inter technology-interference on the sub-GHz/unlicensed ISM band [45] coupled with the limited transmission capabilities of the PCB antennas on the LoRa nodes in our measurement setup.

Despite the relatively long RTTs (3.1 seconds) that we measured over the LoRaX system, it is critical to evaluate the performance in the context of the application. Having a user experience a 3.1 second delay to post a new listing on Etsy or receive a notification about an item being purchased from their online store is negligible when compared to delays on the order of hours or days between opportunities to interact with Internet-based services.

User testing. There were several limitations in our preliminary user testing methodology. Our sample size of $n = 11$ was small and not necessarily representative of the demographics of users who would need to use a LoRaX-based application due to Internet connectivity challenges. However, given that our goal was to investigate the potential and limitations of making network capabilities transparent to the user, the sample was sufficient by demonstrating that users were able to successfully complete tasks under different network regimes and were able to reach a relative consensus about the how information about connectivity might be displayed visually and textually to provide appropriate indicators about service availability given network connectivity status.

6 RELATED WORK

Long-Range, Low-Bandwidth Messaging Services. There are a number of existing long-range, low-bandwidth systems that use the Short Message Service (SMS) communication service [27, 50, 103] as a control channel to facilitate client-server communication. Generally, these services use an SMS gateway and a GSM modem to deliver content between a user and a server using mobile devices over the existing GSM network infrastructure. Furthermore, these systems primarily target communication within a community, with application domains that include tracking health care information and supply chain management.

In contrast, LoRaX does not leverage any preexisting network infrastructure. By utilizing LoRaWAN [62], nodes can be added as needed, allowing a network to expand with its user base. Additionally, the only costs associated with LoRaX are those to purchase a gateway and nodes; there is no per-message charge, allowing a LoRaX network to scale with the number of messages sent. Instead of connecting users within a community, LoRaX is used to connect its users to Internet-based services.

For low data rate and high delay tolerance, the disruption tolerant networking research has produced routing algorithms (e.g., [42, 92]), architectural innovations (e.g., [35], and associated routing schemes to support data delivery in sparse networks. The three-tier Data MULEs project proposed the use of mobile entities to pick up data from stationary sensors and deliver it to access points [85]. Developed at roughly the same time, the Message Ferry approach similarly envisioned nodes with responsibility for moving messages around [108]. A key challenge in sparsely connected networks is knowing when there is data to send that might summon a MULE or a ferry. A low data rate control channel for this purpose is postulated in some prior work [108], long pre-dating our idea of adding a low data rate network as a complement to other networks.

Web Over Challenged Networking Environments. Since the release of the World Wide Web (WWW) in 1991 [106], providing access to diverse types of web contents over any challenged/heterogeneous network environments has been an active area of research. Inherent characteristics of any challenged networking environment, e.g., network heterogeneity, unstable/intermittent connection(s), high delays/jitters, bandwidth bottlenecks, often result in poor connectivity which unsurprisingly contributes to negative user experiences while browsing/searching on the web [21].

Approaches to improving performance have included techniques such as local proxy service and caching, hoarding, prefetching of web contents to mitigate the effects of a challenged networking environment [6, 86]. Another approach proposes integrating concurrently available multiple heterogeneous wireless networks with divergent capacities for transferring large data files to support interactive applications on the web [70]. This particular approach titled Integrating Multi-Path Data Transfer (IMPDT) utilized one of the concurrently available wireless networks which had a lower data rate but comparatively more stable connectivity as a control channel to initiate/set up the actual data transfer operation. In accessing the web, IMPDT utilized the stable, low data-rate control channel for transferring smaller-sized HTTP text files and unstable, high data-rate networks for transferring larger-sized data files. This work bears some relationship to our own in the selective use of a low bandwidth network, but differs in that the IMPDT networks were available everywhere. Our target environment must contend with different geographic reach for the low and high data rate networks.

Human-Network Interaction. Our inquiry into user perceptions of LoRaX is related to work that examines how users interact with information about networks and decision-making around networks [4, 22–25, 54, 99, 105]. A body of work by Chetty et al. investigates user perceptions of network usage and resources [22–24]. Their evaluation of uCap examined user perceptions of a system designed to help household network users monitor their network usage and manage network usage caps [25]. Field studies of uCap revealed that providing users with visibility into network performance caused users to take action and make decisions about their network activities based on information about the network. While uCap and related systems [54] examine how users adjust their network activities and resource allocation based on network information, a separate body of work provides frameworks for users to assign traffic from specific mobile applications to specific network channels (e.g., LTE and WiFi) [4, 99]. While our work also seeks to provide transparency about heterogeneous networks, we focus on users' perceptions about indicators of service availability rather than providing them with the capacity to select a particular network connection for service traffic.

LoRa Systems and Networking. Much of the published research regarding LoRa concerns measurements of physical and link layer performance under a variety of settings, and proposed improvements at these lower layers. For example, Carlsson et al. measure LoRa under different outdoor environmental conditions that include a dense forest, a city, and an open space [16]. They focus on physical layer measurements (e.g., RSSI, SNR) as a function of distance from the end device to the gateway, finding rapid falloff in the city setting and advantages to placing gateway antennas at higher elevations. Blenn and Kuipers similarly measure the popular LoRa Things network [10] for SNR and RSSI, as well as theoretical data rates. The effect of inter-network interference is considered for its potential to limit scalability [68], and mitigations are proposed by Voigt et al. [102] and in Choir [34], the latter of which improves the LoRa link layer for urban settings.

In work closer to our own, Lee and Ke investigate the feasibility of extending the reach of LoRa via a mesh structure that supports ad hoc sensor networking using the LoRa physical layer protocol and omitting the LoRaWAN MAC layer protocol [58]. Our work similarly investigates extending the reach of LoRa by facilitating multiple hops, but we have maintained support for the LoRaWAN MAC protocol. Chen et al. propose a comprehensive system architecture called TinyNet that unifies LoRa with other low-power radio standards such as Bluetooth low energy and 802.15.4 [20]. They measure multi-hop delay performance for a variety of protocol stack combinations. We avoid the complexities of Layer 2.5 in TinyNet by focusing on a single low power physical and link layer coupled logically with a single broadband link. Gu et al. propose a similar out of band control plane for a network with heterogeneous radios [43], using LoRa for the control plane and Zigbee for data plane. Finally, several recent systems exploit LoRa in uncoventional ways for long distance/wide area backscatter [74, 97] and propose alternatives that use chirp spread spectrum in novel ways [47].

7 CONCLUSION

As low-powered, low-bandwidth LPWAN technologies such as LoRa become more ubiquitous, it is valuable to consider how they might be leveraged in concert with high capacity networks with limited range to increase access to services. In this paper, we have described the design and implementation of LoRaX, a first-of-its-kind system that implements a novel initiate-then-update design paradigm to support API service access across a combination of low- and high-bandwidth network regimes. Our measurement evaluations of LoRaX revealed a high variability in the performance of off-the-shelf LoRA hardware with respect to RTT performance and PLR. Our trace-driven simulations demonstrated that with more reliable hardware that has more powerful transmission capabilities, LoRaX could enable UEs to make thousands of API calls in the time that it would take to drive to the nearest point of connectivity–highlighting the ability of a ubiquitous low data rate messaging system to provide meaningful access to Internet-based services. Our UI-design and preliminary user study demonstrated that users were able to successfully complete app-oriented tasks across multiple network regimes and demonstrated users'

acceptance of a UI-design that provided transparency about network availability and the resulting capacity for interaction with services. As we look forward to a future with increasingly multiple network environments that combine high-capacity networks with limited ranges and ubiquitous low-capacity networks, we anticipate that design insights from LoRaX and the initiate-then-update paradigm can help support a future of highly heterogeneous network environments.

ACKNOWLEDGMENTS

This work was funded by National Science Foundation Smart & Connected Communities award NSF-1831698. We would like to thank Md Ashiqul "Robin" Amin, Zane Fink, Rishabh Mahajani, Rafael Olievera, and Bill Eason for their help with engineering our implementation of the LoRaX architecture.

REFERENCES

[1] T. Abdelkader, K. Naik, A. Nayak, N. Goel, and V. Srivastava. 2013. SGBR: A Routing Protocol for Delay Tolerant Networks Using Social Grouping. *IEEE Transactions on Parallel and Distributed Systems* 24, 12 (2013), 2472–2481.

[2] T. Abdelkader, K. Naik, A. Nayak, N. Goel, and V. Srivastava. 2016. A Performance Comparison of Delay-tolerant Network Routing Protocols. *IEEE Network* 30, 2 (2016), 46–53.

[3] LoRa Alliance. 2021. *What is LoRaWAN® Specifications.* Retrieved June 19, 2021 from https://lora-alliance.org/about-lorawan/

[4] Gabriel T. Andreazi, Azzedine Boukerche, Luis H.V. Nakamura, Márcio A. Teixeira, Daniel C. Lobato, and Rodolfo I. Meneguette. 2016. A New Framework for Data Routing Management in Smartphones. In *Proceedings of the 14th ACM International Symposium on Mobility Management and Wireless Access* (Malta, Malta, November 13 - 17, 2016) *(MobiWac '16)*. Association for Computing Machinery, New York, NY, USA, 91–98. https://doi.org/10.1145/2989250.2989252

[5] Arduino. 2021. *Arduino Leonardo with Headers.* Retrieved June 19, 2021 from https://www.arduino.cc/en/Guide/ArduinoLeonardoMicro

[6] Carlos Baquero, Victor Fonte, Francisco Moura, and Rui Oliveira. 1995. MobiScape: WWW Browsing Under Disconnected and Semi-connected Operation. In *Proceedings of the First Portuguese WWW National Conference* (Braga, Portugal).

[7] Yahel Ben-David, Matthias Vallentin, Seth Fowler, and Eric Brewer. 2010. JaldiMAC: Taking the Distance Further. In *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions* (San Francisco, California, June 15, 2010) *(NSDR '10)*. Association for Computing Machinery, New York, NY, USA, Article 2, 6 pages. https://doi.org/10.1145/1836001.1836003

[8] Debopam Bhattacherjee, Simon Kassing, Melissa Licciardello, and Ankit Singla. 2020. In-Orbit Computing: An Outlandish Thought Experiment?. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks* (Virtual Event, USA, November 4 - 6, 2020) *(HotNets '20)*. Association for Computing Machinery, New York, NY, USA, 197–204. https://doi.org/10.1145/3422604.3425937

[9] Darrah Blackwater. 2020. *For tribal lands ravaged by COVID-19, broadband access is a matter of life and death.* Retrieved June 19, 2021 from https://www.azcentral.com/story/opinion/op-ed/2020/05/09/arizona-tribal-lands-need-broadband-access-help-fight-covid-19/5177693002/

[10] Norbert Blenn and Fernando Kuipers. 2017. LoRaWAN in the Wild: Measurements from The Things Network. (2017). arXiv:arXiv:1706.03086

[11] 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* (Malta, Malta, November 13 - 17, 2016) *(MSWiM '16)*. Association for Computing Machinery, New York, NY, USA, 59–67. https://doi.org/10.1145/2988287.2989163

[12] Waylon Brunette, Morgan Vigil, Fahad Pervaiz, Shahar Levari, Gaetano Borriello, and Richard Anderson. 2015. Optimizing Mobile Application Communication for Challenged Network Environments. In *Proceedings of the 2015 Annual Symposium on Computing for Development* (London, United Kingdom, December 1-2, 2015) *(DEV '15)*. Association for Computing Machinery, New York, NY, USA, 167–175. https://doi.org/10.1145/2830629.2830644

[13] Scott Burleigh, Vinton Cerf, Robert Durst, Kevin Fall, Adrian Hooke, Keith Scott, and Howard Weiss. 2003. The InterPlaNetary Internet: A Communications Infrastructure for Mars Exploration. *Acta Astronautica* 53, 4-10 (2003), 365–373.

[14] Richard Canevez, Carleen Maitland, and Matthew Rantanen. 2020. A Dynamic Perspective of Internet Service Provider Adoption of Emergent Network Technology: A Case Study of Tribal Digital Village. *Journal of Information Policy* 10, 1 (2020), 83–122.

[15] Nick Carey. 2005. Establishing Pedestrian Walking Speeds . https://www.westernite.org/datacollectionfund/2005/psu_ped_summary.pdf.

[16] Anders Carlsson, Ievgeniia Kuzminykh, Robin Franksson, and Alexander Liljegren. 2018. Measuring a LoRa Network: Performance, Possibilities and Limitations. In *Proceedings of the 18th International Conference, NEW2AN 2018, and 11th Conference, ruSMART* (St. Petersburg, Russia, August 27-29, 2018). Springer International Publishing, 116–128.

[17] Marco Cattani, Carlo Boano, and Kay Römer. 2017. An Experimental Evaluation of the Reliability of LoRa Long-range Low-power Wireless Communication. *Journal of Sensor and Actuator Networks* 6, 2 (2017), 7.

[18] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi. 2016. Long-range Communications in Unlicensed Bands: The Rising Stars in the IoT and Smart City Scenarios. *IEEE Wireless Communications* 23, 5 (October 2016), 60–67.

[19] Ranveer Chandra and Thomas Moscibroda. 2019. Perspective: White Space Networking with Wi-Fi Like Connectivity. *ACM SIGCOMM Computer Communication Review* 49, 5 (2019), 107–109.

[20] Gonglong Chen, Yihui Wang, Huikang Li, and Wei Dong. 2019. TinyNet: A Lightweight, Modular, and Unified Network Architecture for The Internet of Things. In *Proceedings of the ACM SIGCOMM 2019 Conference Posters and Demos* (Beijing, China, August 19-23, 2019) *(SIGCOMM Posters and Demos '19)*. Beijing, China, 9–11.

[21] Jay Chen, Lakshminarayanan Subramanian, and Kentaro Toyama. 2009. Web Search and Browsing Behavior Under Poor Connectivity. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems* (Boston, MA, USA, April 4-9, 2009) *(CHI EA '09)*. Association for Computing Machinery, New York, NY, USA, 3473–3478. https://doi.org/10.1145/1520340.1520505

[22] Marshini Chetty, Richard Banks, A.J. Brush, Jonathan Donner, and Rebecca Grinter. 2012. You're Capped: Understanding the Effects of Bandwidth Caps on Broadband Use in the Home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA, May 5 - 10, 2012) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 3021–3030. https://doi.org/10.1145/2207676.2208714

[23] Marshini Chetty, Richard Banks, Richard Harper, Tim Regan, Abigail Sellen, Christos Gkantsidis, Thomas Karagiannis, and Peter Key. 2010. Who's Hogging the Bandwidth: The Consequences of Revealing the Invisible in the Home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA, April 10 - 15, 2010) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 659–668. https://doi.org/10.1145/1753326.1753423

[24] Marshini Chetty, David Haslem, Andrew Baird, Ugochi Ofoha, Bethany Sumner, and Rebecca Grinter. 2011. Why is My Internet Slow? Making Network Speeds Visible. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada, May 7 - 12, 2011) *(CHI '11)*. Association for Computing Machinery, New York, NY, USA, 1889–1898. https://doi.org/10.1145/1978942.1979217

[25] Marshini Chetty, Hyojoon Kim, Srikanth Sundaresan, Sam Burnett, Nick Feamster, and W. Keith Edwards. 2015. UCap: An Internet Data Management Tool For The Home. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea, April 18 - 23, 2015) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3093–3102. https://doi.org/10.1145/2702123.2702218

[26] Federal Communications Commission. May 2019. Broadband Deployment Report. https://www.fcc.gov/reports-research/reports/broadband-progress-reports/2019-broadband-deployment-report.

[27] Mauro Ricardo da S. Teófilo, Daniel Risi, Vicente Ferreira Lucena Junior, Thomaz Philippe C. Silva, and Luiz Carlos A.M. Cavalcanti. 2013. A Mobile-Based System for Content Delivery Over SMS. In *2013 International Conference on MOBILe Wireless MiddleWARE, Operating Systems, and Applications* (Bologna, Italy, November 11 - 13, 2013). IEEE, 77–84.

[28] Michael Demmer and Kevin Fall. 2007. DTLSR: Delay Tolerant Routing for Developing Regions. In *Proceedings of the 2007 Workshop on Networked Systems for Developing Regions* (Kyoto, Japan, August 27, 2007) *(NSDR '07)*. Association for Computing Machinery, New York, NY, USA, Article 5, 6 pages. https://doi.org/10.1145/1326571.1326579

[29] Shuo Deng, Anirudh Sivaraman, and Hari Balakrishnan. 2014. All Your Network Are Belong to Us: A Transport Framework for Mobile Network Selection. In *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications* (Santa Barbara, California, February 26 - 27, 2014) *(HotMobile '14)*. Association for Computing Machinery, New York, NY, USA, Article 19, 6 pages. https://doi.org/10.1145/2565585.2565588

[30] Jonathan Desdemoustier, Nathalie Crutzen, and Rudolf Giffinger. 2019. Municipalities' Understanding of the Smart City Concept: An Exploratory Analysis in Belgium. *Technological Forecasting and Social Change* 142, C (2019), 129–141. https://doi.org/10.1016/j.techfore.2018.1

[31] Karyn Doke, Qianli Yuan, Mila Gasco-Hernandez, Megan Sutherland-Mitzer, J. Ramon Gil-Garcia, Petko Bogdanov, and Mariya Zheleva. 2020. Supporting Resilience in Rural Emergency Preparedness and Response Through Improved Information Access. *GetMobile: Mobile Computing and Communications* 24, 2 (September 2020), 5–11. https://doi.org/10.1145/3427384.3427386

[32] Pengyuan Du, Qi Zhao, and Mario Gerla. 2019. A Software Defined Multi-Path Traffic Offloading System for Heterogeneous LTE-WiFi Networks. In *2019 IEEE 20th International Symposium on "A World of Wireless, Mobile and Multimedia Networks"* (Washington DC, USA, August 12, 2019) *(WoWMoM)*. IEEE, 1–9. https://doi.org/10.1109/WoWMoM.2019.8793045

[33] Marisa Elena Duarte. 2017. *Network Sovereignty: Building the Internet across Indian country.* University of Washington Press, Seattle, WA, USA.

[34] 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* (Los Angeles, CA, USA, August 21 - 25, 2017) *(SIGCOMM '17)*. Association for Computing Machinery, New York, NY, USA, 309–321. https://doi.org/10.1145/3098822.3098845

[35] Kevin Fall. 2003. A Delay-Tolerant Network Architecture for Challenged Internets. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (Karlsruhe, Germany, August 25 - 29, 2003) *(SIGCOMM '03)*. Association for Computing Machinery, New York, NY, USA, 27–34. https://doi.org/10.1145/863955.863960

[36] Figma. 2021. *About Figma, the collaborative interface design tool.* Retrieved August 13, 2021 from https://www.figma.com/about/

[37] Karen E. Fisher and Eiad Yafi. 2018. Syrian Youth in Za'atari Refugee Camp as ICT Wayfarers: An Exploratory Study Using LEGO and Storytelling. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies* (Menlo Park and San Jose, CA, USA, June 20 - 22, 2018) *(COMPASS '18)*. Association for Computing Machinery, New York, NY, USA, Article 32, 12 pages. https://doi.org/10.1145/3209811.3209873

[38] OpenJS Foundation. 2021. *About: Node-RED.* Retrieved June 19, 2021 from https://nodered.org/about/

[39] Orestis Georgiou and Usman Raza. 2017. Low Power Wide Area Network Analysis: Can LoRa Scale? *IEEE Wireless Communications Letters* 6, 2 (2017), 162–165.

[40] Barney G. Glaser. 1965. The Constant Comparative Method of Qualitative Analysis. *Social Problems* 12, 4 (1965), 436–445.

[41] Leo A. Goodman. 1961. Snowball Sampling. *The Annals of Mathematical Statistics* 32, 1 (1961), 148–170. http://www.jstor.org/stable/2237615

[42] Samo Grasic, Elwyn Davies, Anders Lindgren, and Avri Doria. 2011. The Evolution of a DTN Routing Protocol - PRoPHETv2. In *Proceedings of the 6th ACM Workshop on Challenged Networks* (Las Vegas, Nevada, USA, September 23, 2011) *(CHANTS '11)*. Association for Computing Machinery, New York, NY, USA, 27–30. https://doi.org/10.1145/2030652.2030661

[43] Chaojie Gu, Rui Tan, and Xin Lou. 2019. One-Hop Out-of-Band Control Planes for Multi-Hop Wireless Sensor Networks. *ACM Transactions on Sensor Networks* 15, 4 (July 2019), 29 pages.

[44] S. Guo, M. H. Falaki, E. A. Oliver, S. Ur Rahman, A. Seth, M. A. Zaharia, and S. Keshav. 2007. Very Low-Cost Internet Access Using KioskNet. *SIGCOMM Computer Communication Review* 37, 5 (October 2007), 95–100. https://doi.org/10.1145/1290168.1290181

[45] Jetmir Haxhibeqiri, Adnan Shahid, Martijn Saelens, Jan Bauwens, Bart Jooris, Eli De Poorter, and Jeroen Hoebeke. 2018. Sub-gigahertz Inter-technology Interference. How Harmful is it for LoRa?. In *Proceedings of the 2018 IEEE International Smart Cities Conference* (Kansas City, MO, USA, September 16 - 19, 2018) *(ISC2)*. IEEE, 1–7.

[46] Kurtis Heimerl and Eric Brewer. 2010. The Village Base Station. In *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions* (San Francisco, California, June 15, 2010) *(NSDR '10)*. Association for Computing Machinery, New York, NY, USA, Article 14, 2 pages. https://doi.org/10.1145/1836001.1836015

[47] 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* (Boston, MA, USA, February 26 - 28, 2019) *(NSDI'19)*. USENIX Association, 2560 Ninth St. Suite 215 Berkeley, CA, United States, 271–283.

[48] Adrian Hooke. 2002. Towards an Interplanetary Internet: A Proposed Strategy for Standardization. In *Proceedings of the SpaceOps 2002 Conference* (Houston, Texas, October 10 - 19, 2002) *(SpaceOps Conferences)*. American Institute of Aeronautics and Astronautics, Reston, VA, USA. https://doi.org/10.2514/6.2002-T5-34

[49] Heon Huh and Jeong Yeol Kim. 2019. LoRa-based Mesh Network for IoT Applications. In *Proceedings of the 2019 IEEE 5th World Forum on Internet of Things* (Limerick, Ireland, April 15 - 18, 2019) *(WF-IoT)*. IEEE, 524–527. https://doi.org/WF-IoT.2019.8767242

[50] Occam Technologies Inc. 2021. *FrontlineSMS.* Retrieved June 19, 2021 from https://www.frontlinesms.com/

[51] International Telecommunication Union. 2019. *Measuring Digital Development: Facts and Figures 2019.* Technical Report. International Telecommunication Union, Geneva, Switzerland. 1–11 pages. https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2019.pdf

[52] David L. Johnson, Veljko Pejovic, Elizabeth M. Belding, and Gertjan van Stam. 2012. VillageShare: Facilitating Content Generation and Sharing in Rural Networks. In *Proceedings of the 2nd ACM Symposium on Computing for Development* (Atlanta, Georgia, March 11 - 12, 2012) *(ACM DEV '12)*. Association for Computing Machinery, New York, NY, USA, Article 7, 10 pages. https://doi.org/10.1145/2160601.2160611

[53] Imad Khan. 2021. *Starlink Internet Coverage, Cost, Speeds and the Latest News — What You Need to Know.* Retrieved June 19, 2021 from https://www.tomsguide.com/news/starlink-internet-coverage-speed-cost-satellites-ipo-and-latest-news

[54] Hyojoon Kim, Srikanth Sundaresan, Marshini Chetty, Nick Feamster, and W. Keith Edwards. 2011. Communicating with Caps: Managing Usage Caps in Home Networks. *SIGCOMM Computer Communication Review* 41, 4 (August 2011), 470–471. https://doi.org/10.1145/2043164.2018526

[55] Gerhard Kormann-Hainzl, Helena Lovasz-Bukvova, and Hölzl. 2021. Are Smart Villages Just Smaller Smart Cities? Call for a Region-Type-Specific Approach to the Smartification of Communities. In *Proceedings of the Central and Eastern European eDem and eGov Days Conference* (Budapest, Hungary, May 5 - 8, 2021) *(CEEE|Gov Days)*. Facultas Verlags-und Buchhandels, Vienna, Austria, 115–125. https://doi.org/10.24989/ocg.v341.8

[56] KTAR. 2020. *Navajo Nation President describes water, internet issues amid COVID-19.* Retrieved June 19, 2021 from https://ktar.com/story/3074025/navajo-nation-president-describes-water-internet-issues-amid-covid-19/

[57] Alexandru Lavric, Adrian I Petrariu, and Valentin Popa. 2019. Sigfox Communication Protocol: The New Era of IoT?. In *Proceedings of the 2019 International Conference on Sensing and Instrumentation in IoT Era* (Lisbon, Portugal, August 29 - 30, 2019) *(ISSI)*. IEEE, 1–4. https://doi.org/10.1109/ISSI47111.2019.9043727

[58] H. Lee and K. Ke. 2018. Monitoring of Large-Area IoT Sensors Using a LoRa Wireless Mesh Network System: Design and Evaluation. *IEEE Transactions on Instrumentation and Measurement* 67, 9 (2018), 2177–2187.

[59] Jansen C. Liando, Amalinda Gamage, Agustinus W. Tengourtius, and Mo Li. 2019. Known and Unknown Facts of LoRa: Experiences from a Large-Scale Measurement Study. *ACM Transactions on Sensor Networks* 15, 2, Article 16 (Feb. 2019), 35 pages. https://doi.org/10.1145/3293534

[60] Ruobing Liang, Liang Zhao, and Peng Wang. 2020. Performance Evaluations of LoRa Wireless Communication in Building Environments. *Sensors* 20, 14 (2020). https://doi.org/10.3390/s20143828

[61] LoRa Alliance. 2018. *LoRaWAN 1.0.3 Specification.* Retrieved March 31, 2021 from https://lora-alliance.org/sites/default/files/2018-07/lorawan1.0.3.pdf

[62] LoRa Alliance. 2019. *LoRaWAN Specification.* Retrieved March 31, 2019 from https://lora-alliance.org/about-lorawan

[63] Borja Martinez, Ferran Adelantado, Andrea Bartoli, and Xavier Vilajosana. 2019. Exploring the Performance Boundaries of NB-IoT. *IEEE Internet of Things Journal* 6, 3 (2019), 5702–5712.

[64] Kais Mekki, Eddy Bajic, Frederic Chaxel, and Fernand Meyer. 2019. A Comparative Study of LPWAN Technologies for Large-scale IoT Deployment. *ICT Express* 5, 1 (2019), 1–7.

[65] Alan Mickelson, Rick Wallace Kenyon, Bennett Miller, Heinz Ulrich Boehmer Fiehn, Mark Hinkle, Kaitlin Yahime Mazotti, Stefano Costa, Nicholas Bollen, and Chris Dizon. 2017. University of Colorado at Boulder WiLDNet Testbed. In *Proceedings of the 2017 IEEE Global Humanitarian Technology Conference* (San Jose, CA, USA, October 19 - 22, 2017) *(GHTC)*. IEEE, 1–9. https://doi.org/10.1109/GHTC.2017.8239268

[66] C. Mikeka, M. Thodi, J.S.P. Mlatho, J. Pinifolo, D. Kondwani, L. Momba, M. Zennaro, and A. Moret. 2014. Malawi Television White Spaces (TVWS) Pilot Network Performance Analysis. *Journal of Wireless Networking and Communications* 4, 1 (2014), 26–32.

[67] Konstantin Mikhaylov, Juha Petaejaejaervi, and Tuomo Haenninen. 2016. Analysis of Capacity and Scalability of the LoRa Low Power Wide Area Network Technology. In *Proceedings of the European Wireless 2016; 22th European Wireless Conference* (Oulu, Finland, May 18 - 20, 2016). VDE, 1–6. https://ieeexplore.ieee.org/document/7499263

[68] Konstantin Mikhaylov, Juha Petajajarvi, and Janne Janhunen. 2017. On LoRaWAN Scalability: Empirical Evaluation of Susceptibility to Inter-Network Interference. In *Proceedings of the 2017 European Conference on Networks and Communications* (Oulu, Finland, June 12 - 15, 2017) *(EuCNC)*. IEEE, 1–6. https://doi.org/EuCNC.2017.7980757

[69] Michael J. Muller and Sarah Kuhn. 1993. Participatory Design. *Communications of the ACM* 36, 6 (June 1993), 24–28. https://doi.org/10.1145/153571.255960

[70] Akira Nagata, Shinya Yamamura, and Masato Tsuru. 2011. Web Browsing Over Multiple Heterogeneous Challenged Networks. In *Proceedings of the 6th ACM Workshop on Challenged Networks* (Las Vegas, Nevada, USA, September 23, 2011) *(CHANTS '11)*. Association for Computing Machinery, New York, NY, USA, 57–60. https://doi.org/10.1145/2030652.2030669

[71] Kagiso Ndlovu, Zablon A. Mbero, Carrie L. Kovarik, and Arif Patel. 2017. Network Performance Analysis of the Television White Space (TVWS) Connectivity for Telemedicine: A Case for Botswana. In *Proceedings of the 2017 IEEE AFRICON* (Cape Town, South Africa, September 18 - 20, 2017). IEEE, 542–547. https://doi.org/10.1109/AFRCON.2017.8095539

[72] The Things Network. 2021. *The Things Uno.* Retrieved June 19, 2021 from https://www.thethingsnetwork.org/docs/devices/uno/

[73] Rabin Patra, Sergiu Nedevschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric Brewer. 2007. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *Proceedings of the 4th USENIX Conference on Networked Systems Design Implementation* (Cambridge, MA, April 11 - 13, 2007) *(NSDI'07)*. USENIX Association, Berkeley, CA, United States, 7.

[74] 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* (Budapest, Hungary, August 20 - 25, 2018) *(SIGCOMM '18)*. Association for Computing Machinery, New York, NY, USA, 147–160. https://doi.org/10.1145/3230543.3230567

[75] Alex Pentland, Richard Fletcher, and Amir Hasson. 2004. Daknet: Rethinking Connectivity in Developing Nations. *Computer* 37, 1 (2004), 78–83.

[76] Andrew Perrin. May 2019. Digital Gap Between Rural and Nonrural America Persists. https://www.pewresearch.org/fact-tank/2019/05/31/digital-gap-between-rural-and-nonrural-america-persists/

[77] Juha Petajajarvi, Konstantin Mikhaylov, Antti Roivainen, Tuomo Hanninen, and Marko Pettissalo. 2015. On the Coverage of LPWANs: Range Evaluation and Channel Attenuation Model for LoRa Technology. In *Proceedings of the 14th International Conference on ITS Telecommunications* (Copenhagen, Denmark, Decmber 2 - 4, 2015) *(ITST)*. IEEE, 55–59. https://doi.org/DOI:10.1109/ITST.2015.7377400

[78] Pew Research Center. June 2019. Mobile Fact Sheet. https://www.pewresearch.org/internet/fact-sheet/mobile/

[79] Thomas Pötsch, Paul Schmitt, Jay Chen, and Barath Raghavan. 2016. Helping the Lone Operator in the Vast Frontier. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks* (Atlanta, GA, USA, November 9 - 10, 2016) *(HotNets '16)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3005745.3005757

[80] Bhaskaran Raman and Kameswari Chebrolu. August 2005. Design and Evaluation of a New MAC Protocol for Long-Distance 802.11 Mesh Networks. In *Proceedings of the Eleventh Annual International Conference on Mobile Computing and Networking*. Cologne, Germany, 156–169.

[81] Rapeepat Ratasuk, Benny Vejlgaard, Nitin Mangalvedhe, and Amitava Ghosh. 2016. NB-IoT system for M2M communication. In *Proceedings of the 2016 IEEE Wireless Communications and Networking Conference* (Doha, Qatar, April 3 - 6, 2016). IEEE, 1–5. https://doi.org/10.1109/WCNCW.2016.7552737

[82] Usman Raza, Parag Kulkarni, and Mahesh Sooriyabandara. 2017. Low Power Wide Area Networks: An overview. *IEEE Communications Surveys & Tutorials* 19, 2 (2017), 855–873.

[83] Anique Scheerder, Alexander van Deursen, and Jan van Dijk. 2017. Determinants of Internet Skills, Uses and Outcomes. A Systematic Review of the Second-and Third-level Digital Divide. *Telematics and Informatics* 34, 8 (2017), 1607–1624.

[84] Semtech. 2021. *What is LoRa?* Retrieved June 19, 2021 from https://www.semtech.com/lora/what-is-lora

[85] Rahul C. Shah, Sumit Roy, Sushant Jain, and Waylon Brunette. 2003. Data MULEs: Modeling and Analysis of a Three-tier Architecture for Sparse Sensor Networks. *Ad Hoc Networks* 1, 2 (2003), 215–233.

[86] Saureen Shah and Apurva Joshi. 2010. COCO: A Web-Based Data Tracking Architecture for Challenged Network Environments. In *Proceedings of the First ACM Symposium on Computing for Development* (London, United Kingdom, December 17 - 18, 2010) *(ACM DEV '10)*. Association for Computing Machinery, New York, NY, USA, Article 7, 7 pages. https://doi.org/10.1145/1926180.1926189

[87] Alireza Shamsoshoara, Fatemeh Afghah, Abolfazl Razi, Sajad Mousavi, Jonathan Ashdown, and Kurt Turk. 2020. An Autonomous Spectrum Management Scheme for Unmanned Aerial Vehicle Networks in Disaster Relief Operations. *IEEE Access* 8 (2020), 58064–58079. https://doi.org/10.1109/ACCESS.2020.2982932

[88] Alireza Shamsoshoara, Mehrdad Khaledi, Fatemeh Afghah, Abolfazl Razi, and Jonathan Ashdown. 2019. Distributed Cooperative Spectrum Sharing in UAV Networks Using Multi-Agent Reinforcement Learning. In *2019 16th IEEE Annual Consumer Communications Networking Conference* (Las Vegas, NV, USA, January 11 - 14, 2019) *(CCNC)*. IEEE, 1–6. https://doi.org/10.1109/CCNC.2019.8651796

[89] Jothi Prasanna Shanmuga Sundaram, Wan Du, and Zhiwei Zhao. 2020. A Survey on LoRa Networking: Research Problems, Current Solutions, and Open Issues. *IEEE Communications Surveys Tutorials* 22, 1 (2020), 371–388. https://doi.org/10.1109/COMST.2019.2949598

[90] Laura Silver, Emily A. Vogels, Mara Mordecai, Jeremiah Cha, Raea Rasmussen, and Lee Rainie. November 2019. Mobile Divides in Emerging Countries. https://www.pewresearch.org/internet/2019/11/20/mobile-divides-in-emerging-economies/

[91] Zachary Spicer, Nicole Goodman, and Nathan Olmstead. 2021. The Frontier of Digital Opportunity: Smart City Implementation in Small, Rural and Remote Communities in Canada. *Urban Studies* 58, 3 (2021), 535–558. https://doi.org/10.1177/0042098019863666

[92] Thrasyvoulos Spyropoulos, Konstantinos Psounis, and Cauligi S. Raghavendra. 2005. Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks. In *Proceedings of the 2005 ACM SIGCOMM Workshop on Delay-Tolerant Networking* (Philadelphia, Pennsylvania, USA, August 26, 2005) *(WDTN '05)*. Association for Computing Machinery, New York, NY, USA, 252–259. https://doi.org/10.1145/1080139.1080143

[93] Anselm Strauss and Juliet Corbin. 1994. Grounded Theory Methodology. *Handbook of Qualitative Research* 17, 1 (1994), 273–285.

[94] Espressif Systems. 2021. *ESP32*. Retrieved June 19, 2021 from https://www.espressif.com/en/products/socs/esp32

[95] Multi-Tech Systems. 2021. *About mLinux*. Retrieved June 19, 2021 from https://www.multitech.net/developer/software/mlinux/about-mlinux/

[96] Multi-Tech Systems. 2021. *MultiTech Conduit® - Programmable Gateway for the Internet of Things (MTCDT Series)*. Retrieved June 19, 2021 from https://www.multitech.com/brands/multiconnect-conduit

[97] Vamsi Talla, Mehrdad Hessar, 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 (sep 2017), 24 pages. https://doi.org/10.1145/3130970

[98] The Things Network. 2020. *Network Architecture*. Retrieved June 19, 2021 from https://www.thethingsnetwork.org/docs/network/architecture.html

[99] Victor Tortorello, Luis H.V. Nakamura, Daniel C. Lobato, and Rodolfo I. Meneguette. 2015. Android Application to Management Multiple Networking Interfaces. In *2015 International Conference on Distributed Computing in Sensor Systems*. 201–202. https://doi.org/10.1109/DCOSS.2015.29

[100] A. Valach and D. Macko. 2018. Exploration of the LoRa Technology Utilization Possibilities in Healthcare IoT Devices. In *Proceedings of the 2018 16th International Conference on Emerging eLearning Technologies and Applications* (Stary Smokovec, Slovakia, November 15 - 16, 2018) *(ICETA)*. IEEE, 623–628. https://doi.org/10.1109/ICETA.2018.8572032

[101] Roberto Vega-Rodríguez, Sandra Sendra, Jaime Lloret, Pablo Romero-Díaz, and Jose Luis Garcia-Navas. 2019. Low Cost LoRa based Network for Forest Fire Detection. In *2019 Sixth International Conference on Internet of Things: Systems, Management and Security* (Granada, Spain, October 22 - 25, 2019) *(IOTSMS)*. IEEE, 177–184. https://doi.org/10.1109/IOTSMS48152.2019.8939193

[102] Thiemo Voigt, Martin Bor, Utz Roedig, and Juan Alonso. 2017. Mitigating Inter-Network Interference in LoRa Networks. In *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks* (Uppsala, Sweden, February 20 - 22, 2017) *(EWSN '17)*. Junction Publishing, USA, 323–328.

[103] Lu Wei-Chih, Matt Tierney, Jay Chen, Faiz Kazi, Alfredo Hubard, Jesus Garcia Pasquel, Lakshminarayanan Subramanian, and Bharat Rao. 2010. UjU: SMS-Based Applications Made Easy. In *Proceedings of the First ACM Symposium on Computing for Development* (London, United Kingdom, December 17 - 18, 2010) *(ACM DEV '10)*. Association for Computing Machinery, New York, NY, USA, Article 16, 11 pages. https://doi.org/10.1145/1926180.1926200

[104] Ulrich Windl, David Dalton, Marc Martinec, and Dale R. Worley. 2006. *The NTP FAQ and HOWTO : Understanding and using the Network Time Protocol (A first try on a non-technical Mini-HOWTO and FAQ on NTP)*. Retrieved June 19, 2021 from http://www.ntp.org/ntpfaq/NTP-a-faq.htm

[105] Susan P. Wyche, Thomas N. Smyth, Marshini Chetty, Paul M. Aoki, and Rebecca E. Grinter. 2010. Deliberate Interactions: Characterizing Technology Use in Nairobi, Kenya. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA, April 10 - 15, 2010) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 2593–2602. https://doi.org/10.1145/1753326.1753719

[106] Robert H. Zakon. 1997. *Hobbes' Internet Timeline*. Retrieved June 19, 2021 from https://rfc-editor.org/rfc/rfc2235.txt

[107] Di Zhang, Yuezhi Zhou, Xiang Lan, Yaoxue Zhang, and Xiaoming Fu. 2018. AHT: Application-Based Handover Triggering for Saving Energy in Cellular Networks. In *Proceedings of the 2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking* (Hong Kong, China, June 11 - 13, 2018) *(SECON)*. IEEE, 1–9. https://doi.org/10.1109/SAHCN.2018.8397106

[108] W. Zhao, M. Ammar, and E. Zegura. 2004. A Message Ferrying Approach for Data Delivery in Sparse Mobile Ad Hoc Networks. In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing* (Roppongi Hills, Tokyo, Japan, May 24 - 26, 2004) *(MobiHoc '04)*. Association for Computing Machinery, New York, NY, USA, 187–198. https://doi.org/10.1145/989459.989483

A LORA PRIMER

Our starting point for the LoRaX system is the low-powered wide area networks (LPWANs) that have emerged as a promising solution for data connectivity to support wide-area telemetry at large scale [18, 82]. By operating in the unlicensed industrial, science, medical (ISM) spectrum band with low power requirements and impressive advertised coverage footprints of up to 15km, Low-powered Radio (LoRa) has been touted as the preferred low-powered data transmission technology to support low data rate Internet of Things (IoT) data transmissions [82].

LoRa's primary use is the transmission of low data rate sensor values to applications in the cloud where the data is aggregated and analyzed as depicted in Figure 11. As illustrated, LoRa-enabled sensors on the left of the figure stream data at low rate to one or more LoRa Gateways in range. The LoRa Gateway then forwards packets to an Internet-connected LoRa Network Server, typically via cellular or WiFi access. The LoRa Network Server is responsible for sending acknowledgements, filtering data (e.g., if multiple gateways receive the same packet), and forwarding data to the appropriate Application Server for processing.

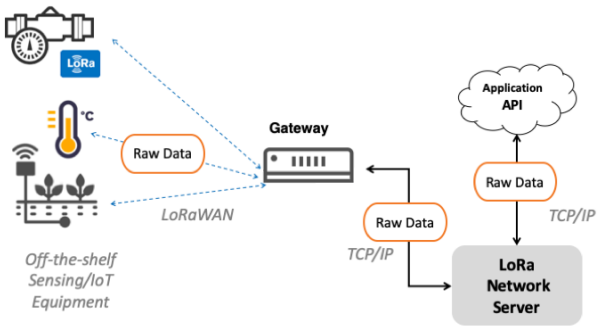

Fig. 11. Standard LoRa network architecture.

LoRaWAN is a combination of MAC and network layer protocols that manages communications between end devices and gateways by acting as a routing protocol while facilitating shared access to the physical layer [84]. LoRaWAN networks assume a star-of-stars network topology where gateways are connected to a central network server via a standard TCP/IP backbone and act as transparent bridges between end devices and central network servers by converting LoRaWAN packets to TCP/IP packets and vice versa [3, 61].

A.1 LoRa Device Classes

The LoRa specification defines three classes of end devices, each with different behavior for transmitting, listening/receiving, and idle time: Class A, Class B, and Class C. For our initial design, we considered several system architectures based on the characteristics of the different LoRa end device classes. Class A is the minimum supported behavior by LoRa-compatible end devices; some support Class C, which was the second class to be implemented, and currently only a handful support Class B, the last to be standardized. Because Class A is the most commonly available, we use it for our design and implementation.

In Class A mode, an end device's communication is always initiated by the device itself. The end device will only listen for a downlink message shortly *after* it sends an uplink message. Once the end device completes asynchronous transmission of a message, its radio is turned off, or made idle, for a configurable period of time T1, often 1 second. It then turns on the radio in receive mode and listens for the preamble of an incoming downlink message (from the gateway). This is the first receive window, and its duration must be long enough to detect a message's preamble. If it does detect the preamble of a message, the end device leaves the radio turned on to receive the entire message. If the message is received intact and its destination address is this end device, then the radio is turned off (set back to idle). A diagram for the typical Class A timing is shown in Figure 12.

Alternatively, if no preamble is detected during the first receive window, or if a message was received but its destination was a different end device, then the end device leaves the radio idle until the end of a second, configurable period of time T2, at which time it turns on the radio again in receive mode and listens for a downlink message preamble. This is the second receive window. As before, the radio is turned off at the end of the window if no preamble is detected. Otherwise, the radio is kept on until the entire message is received.

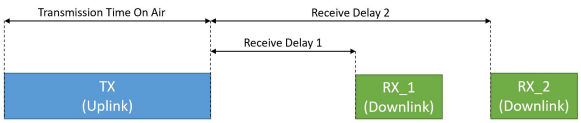

Fig. 12. LoRa Class A end device timing

As LoRa was originally developed to support IoT, a major benefit of using Class A communications is power savings; by limiting the time spent in receive mode, the overall power consumption of the device is minimized. . However, the main downside of using Class A is the increased difficulty of bidirectional communication. Since the end device will only listen for a downlink message immediately after sending an uplink message, there can be no asynchronous downlink communication to the end device. Further, the end device will only accept at most one downlink message for each uplink message that it sends. Because of these constraints, the typical use case for Class A devices is for battery-powered sensors that primarily only need to send data, and only occasionally need to receive messages in return.

A.2 Tuning Factors

Timing can be tuned using a parameter referred to as the Spreading Factor (SF), which determines the frequency at which data is sent. SF ranges from 7-12, where lower SF values correspond to higher transmission rates and higher SF values correspond to lower transmission rates. SF7 is considered to be the ideal setting for maximizing throughput while minimizing power consumption due to the shorter times required to transmit a single payload. However, higher SF values can be used to ensure greater reliability of transmission. To enable LoRa end-devices to be responsive to changing interference conditions, the gateway can communicate to devices which SF value it should use for best performance at any give time. This adaptive data rate (ADR) mechanism is an optional setting that we leverage in our experiments in Section 4.1.

Additionally, LoRa allows for three different channel widths: 125 kHz, 250 kHz, and 500 kHz [61, 98]. Notably, the ADR mechanism also moderates the channel bandwidth from transmission to transmission for achieving optimal throughput and power consumption.