

Lecture 15

LPWANs Challenges

CS433 – Wireless Protocols for IoT

Branden Ghena – Spring 2026

With images from Swarun Kumar (CMU)
and Dinesh Bharadia (UCSD)

Materials in collaboration with
Pat Pannuto (UCSD) and Brad Campbell (UVA)

Administrivia

- Lab: WiFi
 - Due tonight (can use slip days if necessary)
 - For throughput optimization two requirements
 1. Do at least *something* at the PHY layer
(there are multiple low-hanging fruit here)
 2. Do something for packet size
(definitely consider what the packet size ought to be)
 - You should be able to beat 100 kbps without much effort
 - And much faster is possible
- Lab: LoRa
 - Tomorrow! See you all in lab

Today's Goals

- Deep-dive into challenges LPWANs face
- Explore academic research improving LoRaWAN capabilities and exploring LoRaWAN deployments

Outline

- **LPWAN Challenges**
- Improving LoRaWAN
- Wide-area Network Analyses

Challenge for Unlicensed LPWANs Paper

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

Branden Ghena

University of California, Berkeley
brghena@berkeley.edu

Joshua Adkins

University of California, Berkeley
adkins@berkeley.edu

Longfei Shangguan

Microsoft Cloud & AI
longfei.shangguan@microsoft.com

Kyle Jamieson

Princeton University
kylej@cs.princeton.edu

Philip Levis

Stanford University
pal@cs.stanford.edu

Prabal Dutta

University of California, Berkeley
prabal@berkeley.edu

ABSTRACT

Low-power wide-area networks (LPWANs) are a compelling answer to the networking challenges faced by many Internet of Things devices. Their combination of low power, long range, and deployment ease has motivated a flurry of research, including exciting results on backscatter and inter-

ACM Reference Format:

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

- MobiCom 2019


Two major categories of low-power, wide-area network protocols

- **Unlicensed LPWANs**

- Sigfox, LoRaWAN, etc.
- Unlicensed band, 915 MHz (US)
- Managed or user deployable

- **Cellular IoT**

- LTE-M and NB-IoT
- Licensed cellular bands
- Managed networks
- Rolled out US-wide




Do novel networks meet application needs?

- How do we compare varied requirements and capabilities?
 - Networks have throughput per gateway and range of gateway.
 - Applications have throughput per device and deployment area.
- Each gateway must support throughput for all devices in its coverage area.
 - Deployment areas are often wider than a single gateway's range.
- Solution: compare the density of communication.
 - Data communication rate per unit area.

New metric for wide-area communication.

Our proposed metric: **bit flux**

- $bit\ flux = \frac{network\ throughput}{coverage\ area}$
- Units: bit per hour / m²




Bit flux can measure application needs.

For an application:

$$\text{bit flux} = \frac{\sum \text{each device's uplink}}{\text{deployment area}}$$

- Assumes a relatively homogeneous distribution.




Bit flux can measure network capabilities.

For a network:

$$\text{bit flux} = \frac{\text{gateway goodput}}{\text{gateway coverage area}}$$


- Assumes a non-overlapping deployment of gateways.
- Note that bit flux alone ignores the total number of gateways required.




Bit flux accounts for spatial reuse.

- Reducing coverage area and deploying additional gateways improves capacity.


- $$\text{bit flux} \uparrow = \frac{\text{gateway goodput}}{\text{coverage area} \downarrow}$$




Bit flux measurement for LoRaWAN.



Networks differ in capability by orders of magnitude.



Range reduction results in a bit flux curve for each network.



Let's compare network capabilities to a real-world application.

Smart household electric meters.

- ~250 bytes of data every 4 hours
- ~370000 electric customers in San Francisco



$$\frac{\frac{250 \text{ bytes}}{4 \text{ hours}} * 370000 \text{ devices}}{120 \text{ km}^2} \approx \frac{51000 \text{ bps}}{120 \text{ km}^2} \approx 1.5 \frac{\text{bph}}{\text{m}^2}$$


Bitflux for various applications

- See paper for full details on each application
- Creates a comparison point between applications and networks!


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

<https://brandenghena.com/projects/lpwan/ghena19lpwans.pdf>

All networks are capable of meeting the data needs of electricity metering.




Unlicensed LPWANs lag behind Cellular IoT in ability to support applications.



2G < 0.03% utilized 18

Sigfox requires range reduction to meet application needs.



Capacity solutions are relatively straightforward.


- Better access control mechanisms. Explore CSMA?
- Recover simultaneous transmissions (Choir and Charm).
- Increase bandwidth (TV white spaces).
- All likely come at the cost of increased energy usage...
 - Results in a protocol that looks pretty similar to cellular...

Adwait Dongare, et al. "Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks." *IPSN'18*

Rashad Eletreby, et al. "Empowering low-power wide area networks in urban settings." *SIGCOMM'17*


Abusayeed Saifullah, et al. "SNOW: Sensor network over white spaces." *SenSys'16*

LoRaWAN devotes most of its network capacity to a single application.



Coexistence is inevitable in urban areas.

- Urban environments and long range lead to many overlapping deployed networks.
- Capacity problems worsen coexistence by devoting more bandwidth to one application.
- It's not just electricity metering...




Coexistence in unlicensed bands is a more difficult problem.

- No methods for inter-network negotiation so far.
 - Without buy-in from most deployments, all access control becomes uncoordinated.
-
- **Cellular IoT does not have this problem**

Cellular may dominate future deployments.

- LTE-M and NB-IoT are now deployed in the US (and worldwide).
- Licensed bandwidth avoids the coexistence problem.
- Cellular may solve many applications but is not a perfect solution.
 - Still has higher energy and monetary costs for use.
 - Also limited to where service is already available.



Unlicensed LPWANs are still useful for some scenarios.

- Controlled or unoccupied regions have reduced coexistence concerns.
 - Industrial factories, farms, parks and forests.
- Unlicensed networks are very exciting for research.
 - Anyone can deploy a network wherever they want.
 - Much easier to explore protocol modifications and new technologies.
- Research suffers without real-world applications.
 - Problem areas are strong recommendations for new research.
 - New research is only useful if they will have real-world impacts.

Implications – Low-Power Wide-Area Networks.

- Existing unlicensed LPWANs face significant challenges in supporting urban applications.
 - Best suited for industrial or agricultural uses in controlled environments.
- Research directions for unlicensed LPWANs:
 - improve network capacity,
 - and enable coexistence.
- Cellular IoT networks (LTE-M and NB-IoT) are positioned to solve the needs of city-scale sensing.
 - If the money and energy costs are there.

Break + Question

- How important is a homogenous distribution?
- Can you come up with a scenario where this breaks?

For an application:

$$\textit{bit flux} = \frac{\sum \textit{each device's uplink}}{\textit{deployment area}}$$

- Assumes a relatively homogeneous distribution.

Break + Question

- How important is a homogenous distribution?
- Can you come up with a scenario where this breaks?
- Consider densely populated pockets of transmission over wide areas
 - Deployment area needs to correspond to pockets, not to entire range

For an application:

$$\textit{bit flux} = \frac{\sum \textit{each device's uplink}}{\textit{deployment area}}$$

- Assumes a relatively homogeneous distribution.

Outline

- LPWAN Challenges
- **Improving LoRaWAN**
- Wide-area Network Analyses

Unlicensed LPWANs are a very active research area

- Active work on:
 - Better access control mechanisms
 - Increasing bandwidth by utilizing other spectrums like TV white spaces
 - Recover simultaneous transmissions
 - Goal: reduce re-transmissions, increase end-device sleep and battery life
- Let's look at some research improvements that have been made to LoRa

Can we distinguish data LoRa chirps that have collided?

Empowering Low-Power Wide Area Networks in Urban Settings

Rashad Eletreby
Carnegie Mellon University
Pittsburgh, PA
reletreby@cmu.edu

Swarun Kumar
Carnegie Mellon University
Pittsburgh, PA
swarun@cmu.edu

Diana Zhang
Carnegie Mellon University
Pittsburgh, PA
dianaz1@andrew.cmu.edu

Osman Yağan
Carnegie Mellon University
Pittsburgh, PA
oyagan@andrew.cmu.edu

ABSTRACT

Low-Power Wide Area Networks (LP-WANs) are an attractive emerging platform to connect the Internet-of-things. LP-WANs enable low-cost devices with a 10-year battery to communicate

1 INTRODUCTION


Recent years have witnessed Low-Power Wide Area Networks (LP-WANs) emerge as an attractive communication platform for the Internet of Things (IoT) [37]. LP-WANs enable low-power devices

- SIGCOMM 2017


Choir concept

- Can we distinguish data from LoRa chirps that have collided?
 - Yes! By applying signal processing to the problem

Data in time domain




Data in frequency domain





What happens when LoRa chirps collide?

Different data





+



What happens when LoRa chirps collide?


Same data




Hardware offsets!

Imperfections in hardware create offsets


Carrier frequency offsets (CFO)



Timing offsets (TO)




Decoding colliding packets



Choir results

- Recovering collided packets resolves losses due to Aloha transmission!
- Increases maximum throughput on the network considerably!
- Downsides:
 - Requires hardware modifications on gateways
 - Computationally expensive



Can we recover weak LoRaWAN signals?

Charm: Exploiting Geographical Diversity Through Coherent Combining in Low-Power Wide-Area Networks

Adwait Dongare
Carnegie Mellon University
adongare@cmu.edu

Revathy Narayanan
Carnegie Mellon University
IIT Madras
revathyn@andrew.cmu.edu

Akshay Gadre
Carnegie Mellon University
agadre@andrew.cmu.edu

Anh Luong
Carnegie Mellon University
anhluong@cmu.edu

Artur Balanuta
Carnegie Mellon University
Instituto Superior Técnico, Lisbon
artur@cmu.edu

Swarun Kumar
Carnegie Mellon University
swarun@cmu.edu

Bob Iannucci
Carnegie Mellon University
bob.iannucci@west.cmu.edu

Anthony Rowe
Carnegie Mellon University
agr@ece.cmu.edu

ABSTRACT

Low-Power Wide-Area Networks (LPWANs) are an emerging wireless platform which can support battery-powered devices lasting


KEYWORDS

LPWAN, sensor networks, coherent combining, diversity

- IPSN 2018


Charm

- Take advantage of multiple gateways in range of a device
- Combine signals received at each gateway to recover packets that weren't received cleanly
- Particularly useful at decoding weak signals




Coherent combining

- Signal received at each receiver is a combination of
 - Signal
 - Wireless channel
 - Random noise
- Signal and wireless channel are similar at all receivers
- Noise is different though!
Possibly independent?
- When combined, signals are coherent (build in strength) whereas noise is incoherent (spreads out)



Charm uses coherent combining across gateways

- Gateways send signal data to the cloud
- The cloud can perform combining on the data and recover signals
- Challenges
 - Only send the data when it's needed
 - Tight time synchronization on the data
 - Gateway hardware changes



Can we recover packets that have bad CRCs?

Opportunistic Packet Recovery Paper

A Cloud-Optimized Link Layer for Low-Power Wide-Area Networks

Artur Balanuta
Carnegie Mellon University
Instituto Superior Técnico
artur@cmu.edu

Swarun Kumar
Carnegie Mellon University
swarun@cmu.edu

Nuno Pereira
Carnegie Mellon University
npereira@cmu.edu

Anthony Rowe
Carnegie Mellon University
agr@ece.cmu.edu

ABSTRACT

Conventional wireless communication systems are typically designed assuming a single transmitter-receiver pair for each link. In Low-Power Wide-Area Networks (LP-WANs), this one-to-one design paradigm is often overly pessimistic in terms of link budget because client packets are frequently detected by multiple gateways

ACM Reference Format:


Artur Balanuta, Nuno Pereira, Swarun Kumar, and Anthony Rowe. 2020. A Cloud-Optimized Link Layer for Low-Power Wide-Area Networks. In *The 18th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '20)*, June 15–19, 2020, Toronto, ON, Canada. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3386901.3388915>

- MobiSys 2020

Opportunistic Packet Recovery (OPR)

- Can we recover packets that have bad CRCs?
 - What if we have some information about where the interference might have occurred during the signal transmission?
 - OPR demonstrates that we can recover packets!
- Process
 1. Receive bits even for bad packets
 2. Measure RSSI for each bit along the way
 3. Look for changes in RSSI that signals interference
 4. Try different values for the effected bits until the CRC succeeds!

Detecting error bits in transmitted packets



Correctly identifies 83% of the corrupted bits
with 17% false positives


Multiple gateways do even better

- If multiple gateways receive the packet, they can compare their received data
- Common bits are likely correct, while differences are likely interference

Node	Payload Hex	Bit Errors
TX	00000000000000000000000000000000	-
RX1	000000000000 2008 0000000000000000	2
RX2	000000000 2001406 0000000000000 10	6
RX3	000000000 1002004 0000000000000000	3

OPR total design

- Bad packets are sent to the OPR Server for handling along with RSSI data
- OPR server attempts to reconstruct the packet



- In practice, system can correct up to 72% of CRC errors!
 - Also completes correction in time for the ACK response (within 1 second)
- Downsides: not really much, this should really be used in practice

Can LoRa use a CSMA approach?

Busy-Signal Multiple Access (BSMA) Paper

BSMA: Scalable LoRa networks using full duplex gateways


Raghav Subbaraman[†], Yeswanth Guntupalli[†], Shruti Jain[†], Rohit Kumar[†],
Krishna Chintalapudi[§], Dinesh Bharadia[†]

[†]University of California San Diego, [§]Microsoft Research
USA

{rsubbaraman,yguntupa,shjain,rokumar,dineshb}@eng.ucsd.edu,krchinta@microsoft.com

Abstract

With its ability to communicate long distances, LoRa promises city-scale IoT deployments for smart city applications. This long-range, however, also increases contention as many thousands of devices are connected. Recently, CSMA has been proposed as a viable MAC for resolving contention in LoRa networks. *In this paper, supported by measurements, we demonstrate that CSMA is ineffective in urban deployments.* While gateways stationed at rooftops enjoy a long communication range, 70% of the devices placed at street level fail to sense each others' transmissions and remain hidden, owing to obstructions by tall structures. *We present Busy Signal Multiple Access (BSMA), where the LoRa gateway transmits a downlink busy signal while receiving an uplink transmission.* The IoT devices defer up-




- MobiCom 2022

CSMA provides better throughput than Aloha

- Current LoRaWAN uses Aloha for uplink communication
 - So there's nothing stopping devices from colliding
- Switching to CSMA could provide significantly more throughput
 - Reduces time and energy wasted on packet collisions
- Concern: does CSMA for LoRa actually work in the real world?


Obstructions cause hidden terminal problem

- Hidden terminal problem is a concern here
 - Devices may not be able to hear each other
- Measured reality: devices fail to detect simultaneous transmissions 70% of the time!!




Hidden terminals greatly reduce CSMA throughput

- Perfect CSMA could provide twice the throughput of Aloha
- But, at 70% failure for CSMA, it's not really an improvement over Aloha




Busy-signal can alleviate this problem!

- Idea: have the Gateway transmit a “busy signal” on one of the downlink channels while an uplink is in progress
- Other devices can hear the gateway, even if they can’t hear each other




BSMA design



- End devices do not attempt to transmit while a busy signal is present
 - Instead, they defer transmission for a random duration, and go back to sleep
 - When waking up, they sense for a busy signal again before transmitting

Real-world improvements from BSMA application

- 50% improvement over Aloha when network is under load
 - Still some packets are lost due to collisions
- Could be combined with other approaches to still recover collided data
 - Maybe even intentionally send busy signal based on number of acceptable collisions
- Greatly reduced energy per packet as well
 - Since they're more likely to be received, and sensing doesn't cost too much
- Downside:
 - Must allocate channels for busy signal(s)



Break + Open Question

- What are the challenges in translating all of this research into real-world use?

Break + Open Question

- What are the challenges in translating all of this research into real-world use?
 - Improve gateway hardware: how do people buy/make it?
 - Needs a manufacturer to be interested
 - Multiple gateways are necessary
 - Not applicable to very small deployments
 - Network operators (TTN or Helium) would be good targets
 - They also have backend stuff running anyways, so adding OPR should be possible!

Outline

- LPWAN Challenges
- Improving LoRaWAN
- **Wide-area Network Analyses**

“The Helium Paper”

Federated Infrastructure: Usage, Patterns, and Insights from “The People’s Network”

Dhananjay Jagtap*
dtjagtap@eng.ucsd.edu
UC San Diego

Alex Yen*
alyen@ucsd.edu
UC San Diego

Huanlei Wu
huw012@ucsd.edu
UC San Diego

Aaron Schulman
schulman@cs.ucsd.edu
UC San Diego

Pat Pannuto
ppannuto@ucsd.edu
UC San Diego

ABSTRACT




In this paper, we provide the first broad measurement study of the operation, adoption, performance, and efficacy of Helium. The Helium network aims to provide low-power, wide-area network wireless coverage for Internet of Things-class devices. In contrast to traditional infrastructure, “hotspots” (base stations) are owned and operated by individuals who are paid by the network for providing

1 INTRODUCTION

Core to the success of the Smart City will be its supporting infrastructure. While the edge devices are ready, there is not yet a widely-deployed supporting communication infrastructure suited to their needs. Embedded systems are defined by the walls, sidewalks, and windows into which they are literally embedded. For this network, scale will come not from millions of devices sending


- IMC 2021

Cell providers put spectacular effort (\$\$) into rolling out new wireless technologies



T-Mobile 4G 2017

What if we could just get people to put mini towers up on their own?



These are Internet-connected hotspots deployed as part of a crowd-sourced wireless infrastructure called **Helium**


People will do a lot if you pay them
(Especially if you pay them in cryptocurrency)

- People are paying real USD to buy Helium access points (~\$400/ea)
- They put in their time and energy to install and maintain them
- ...all so that they can earn Helium Tokens (HNT) on the Helium Blockchain
 - The Helium blockchain pays users HNT for coverage and ferrying data
- **Fall 2021 Numbers:**
 - **250,000 Hotspots deployed**
 - **Adding 2,000 new hotspots / day**
 - (Helium's blockchain came online on July 29, 2019)


The study: What does it look like when a network goes from 0 to 250,000 in two years?

- Because of the blockchain, we can see how it happened
 - We can observe historical growth and current utilization
- Because it's a peer-to-peer network, we can observe active infrastructure
 - We can learn about the ISPs people are relying on for backhaul
- Because of the open nature, we can easily do field measurements
 - We can analyze performance, robustness, and coverage of Helium


The Helium Architecture




This study mostly thinks in terms of hotspots



Cumulative and Daily Growth of Helium




Once-deployed, most hotspots don't move



And if a hotspot is going to move, it moves early on

Takeaway: Coverage is reasonably stable



ISPs Become a Chokepoint

	ISP	Number of Hotspots
1	Spectrum	2497
2	Comcast	1922
3	Verizon	1590
4	Cablevision	450
5	AT&T	338

- 1,588 cities rely on only **1 ASN** (~an ISP)
 - 414 of these cities have at least 2 hotspots
 - (Palma, Spain has 76 hotspots)
- **2020 Spectrum outage in LA could have taken down 291 out of 333 hotspots (87%)**


What this study doesn't measure (well):
The end-user performance of the network



Simplistic coverage models to do not map to real-world performance



(Urban)



(Suburban
)

“The CityWAN Paper”

Citywide LoRa Network Deployment and Operation: Measurements, Analysis, and Implications

Shuai Tong¹, Jiliang Wang¹✉, Jing Yang¹, Yunhao Liu¹, Jun Zhang²

¹Tsinghua University ²ENNEW Technology Co., Ltd.

tl19@mails.tsinghua.edu.cn, jiliangwang@tsinghua.edu.cn, jing-yan18@mails.tsinghua.edu.cn
yunhao@tsinghua.edu.cn, junzhang@enn.com

ABSTRACT

LoRa, as a representative Low-Power Wide-Area Network (LPWAN) technology, holds tremendous potential for various city and industrial applications. However, as there are few real large-scale deployments, it is unclear whether and how well LoRa can eventually meet its prospects. In this paper, we demystify the real performance of LoRa by deploying and measuring a citywide LoRa network named **CityWAN**, which consists of 100 gateways and 19,821 LoRa end nodes, covering an area of 130 km² for 12 applications. Our measurement focuses on the following perspectives: (i) Performance of applications running on the citywide LoRa network; (ii) Infrastructure efficiency and deployment optimization; (iii) Physical layer signal features and link performance; (iv) Energy profiling and cost estimation for LoRa applications. The results reveal that LoRa performance in urban settings is bottlenecked by the prevalent blind spots, and there is a gap between the gateway efficiency and network coverage for the infrastructure deployment. Besides, we find that LoRa links at the physical layer are susceptible to envi-



Figure 1: The deployment environments of LoRa based municipal devices in the *CityWAN*.

large-scale Internet of Things (IoT). As a representative LPWAN

- SenSys 2023

What happens if we “really deploy” LoRa at-scale in a real city?

- Looks at realistic use cases, mostly utility metering

Table 1: Number of deployed devices in various smart city applications.

Device Name	Number	Device Name	Number
Gas meter	18765	Water meter	591
Heat meter	245	Gas stove	17
Gas alarm	143	Regulator valve	26
Electricity meter	19	Hygrothermograph	2
Intelligent module	8	Pressure sensor	3
SOS Emergency button	1	Gate magnetism	1



Location of physical infrastructure (in, deeply in, outside a building) makes a significant difference




Figure 3: CDF of packet loss rates for different applications.




Figure 4: SNR distribution of transmissions in different applications.

Coverage becomes less “simple sphere” the further you get from the gateway

- Also sometimes referred to as “urban canyon” effect




Figure 6: Fine-grained coverage measurement for the ESP map in a typical urban environment.

Emerging, non-traditional demodulators from research **do** help in real-world settings




Figure 8: Weak signal demodulation: SNR thresholds for different demodulation mechanisms.




Figure 9: Coverage areas of standard LoRa and the signal combination based demodulation mechanism.

Link failures tend to be coordinated in time




Figure 16: Bursty loss pattern of LoRa transmissions.

Competing users for ISM band demonstrate regular interference

- In deployment city, Digital Terrestrial Multimedia Broadcast (DTMB, a digital TV standard) partially overlaps with LoRa
 - Runs from the hours of 6:00-24:00
 - 5 SNR loss to LoRa transmissions while running




Figure 18: Fine-grained SNR measurements on a typical LoRa link, showing link fluctuations with different transmission times of a day.

Some channels are worse than others

- Though this varies with time...
 - Left: Over a 24h window
 - Right: Over a month-long view (of same link)




Figure 20: Link SNRs over channels show the trend of (a) aggregation for the short-term measurement in a day and (b) fluctuation for the long-term measurement in a month.

Being high up helps you see more gateways, and seeing more than “a few” gateways makes packet loss rare




Figure 21: Number of connected gateways at different floors.




Figure 22: Packet loss rates versus number of reached gateways.

Reliability doesn't come for free... long SFs take meaningfully longer to send (and thus consume much more energy/packet)




Figure 25: Node battery lifetime with different data rates and uplink transmission demands.

Table 2: Lifetime estimation for different network technologies across various application demands.

Network Technology	End Node Lifetime (years)			
	Heat Meter 90 Bytes Per Day	Gas Meter 256 Bytes Per Day	Hygrometer 64 Bytes Per Hour	Regulator 256 Bytes Per Hour
LoRa (143 dB)	9.5	6.7	3.8	1.6
NB-IoT (144 dB)	5.3	3.5	1.4	1.1
Sigfox (155 dB)	3.4	2.3	0.9	0.6
LoRa (157 dB)	4.5	2.9	1.7	0.8
NB-IoT (164 dB)	2.5	1.7	0.7	0.6
LTE-M (164 dB)	2.2	1.4	0.5	0.5

...but there is more to a deployment than just battery life for real world networks

Table 3: System Expenditures.

	Cost	LoRa	NB-IoT	Sigfox
Deployment Expenditure	User Equipment (\$)	4-6	6-12	4
	Site build (K\$)	2.1	21	10.5
	Site lease (K\$/year)	0.4-1.1	3.7-8.4	0.9-1.1
Operation Expenditure	Spectrum (K\$/kHz/site)	0	0.001	0
	Electricity (K\$/year)	0.1	1	1
Maintenance Expenditure	Battery replacement (\$/node/year)	0.8	1.4	1.1
	Facilities maintenance (relative to deployment expenditure/year)	20%	10%	15%




Figure 26: Cost evaluation for three LPWAN technologies under different scenarios and user densities.

- This is an excellent example of the type of technology comparison we expect for the final design report...
 - (though you have to also include the text that explains *how* you get your numbers!)

In sum: Does LoRa live up to its promises for Smart Cities?

It depends and work is ongoing.

Outline

- LPWAN Challenges
- Improving LoRaWAN
- Wide-area Network Analyses