Towards Next-Generation Global IoT: Empowering Massive Connectivity with Harmonious Multi-Network Coexistence

Ziyue Zhang, Xianjin Xia, Ruonan Li, Yuanqing Zheng
The Hong Kong Polytechnic University, Hong Kong, China
{ziyue.zhang,ruo-nan.li}@connect.polyu.hk,
{xianjin.xia,yuanqing.zheng}@polyu.edu.hk

ABSTRACT

LoRaWAN offers a compelling solution for delivering cost-effective network access to millions of IoT devices worldwide. However, operators face challenges in scaling their services to meet the growing demands of IoT connections. Moreover, current LoRaWANs foster competition rather than cooperation among coexisting networks, resulting in substantial capacity degradation as network density increases. To identify the root causes limiting LoRaWAN scalability and to enable harmonious coexistence among network operators, this paper conducts an in-depth investigation of operational LoRaWANs. For the first time, our study reveals that the capacity degradation in LoRaWAN is not due to the traditionally believed issues (such as interference or packet collisions in the wireless medium) but rather a newly identified issue termed *decoder contention problem*. This problem cannot be effectively resolved using conventional approaches and will significantly hinder the scaled deployment of LoRaWANs as a next-generation global infrastructure. Based on our new findings, we propose design principles that guide our exploration for effective strategies to address this emerging practical problem. We develop concrete deployable solutions to mitigate contention, optimize spectrum utilization, and promote spectrum sharing among network operators. Extensive evaluations demonstrate that our strategies effectively boost network capacity close to the theoretical bound, and support the coexistence of up to six networks with significant improvement in spectrum efficiency.

CCS CONCEPTS

• **Networks** → **Network protocols**; **Network performance analysis**; *Network management*; *Wireless access points, base stations and infrastructure.*

KEYWORDS

IoT connectivity, Network capacity, Concurrent communication, Spectrum sharing, Frequency planning, LoRaWAN, Coexistence

ACM Reference Format:
Ziyue Zhang, Xianjin Xia, Ruonan Li, Yuanqing Zheng. 2025. Towards Next-Generation Global IoT: Empowering Massive Connectivity with Harmonious Multi-Network Coexistence. In *ACM SIGCOMM 2025 Conference (SIGCOMM '25), September 8–11, 2025, Coimbra, Portugal.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3718958.3750504

SIGCOMM '25, September 8–11, 2025, Coimbra, Portugal

ACM ISBN 979-8-4007-1524-2/2025/09
https://doi.org/10.1145/3718958.3750504

1 INTRODUCTION

LoRaWAN for Connecting Global IoT. LoRaWAN has established itself as the *de facto* standard for providing cost-efficient, low-data-rate wireless connectivity to IoT devices [18, 35]. The EU countries and the US are expanding the provisions of LoRaWAN as a foundational infrastructure to facilitate city- and nation-scale IoT connections [19, 21, 23]. Unlike cellular networks, LoRaWAN uses unlicensed ISM bands and adopts optimized low-power designs for IoTs, providing communication services at low cost. Moreover, LoRaWAN allows private deployments to complement public infrastructures. Currently, more than 70 operators are providing LoRaWAN services across over 100 countries worldwide [20] (Table 2). With sustainable investments from governments and industries, LoRaWAN is envisioned to evolve into a global infrastructure for interconnecting millions of IoT devices in smart cities, utility management, agriculture, *etc.*

However, are current LoRaWANs capable of serving the ever-increasing demands of IoT connections at a global scale? Empirical studies with real-world operational LoRaWANs reveal that existing systems suffer from the following limitations (as detailed in § 2.2):

(1) Limited network capacity with significant performance gaps. Operational LoRaWANs support very limited user concurrency. For instance, TTN (a global LoRaWAN operator [18]) only supports a maximum of 16 users to actively communicate at the same time. Although LoRaWAN empowers massive connectivity by duty-cycling IoT devices with low active durations, the low concurrency fundamentally limits the network capacity (measured as *the maximum number of IoT connections*). Operators may resort to dense gateway deployment to increase network capacity. Our investigations, however, reveal that extra LoRaWAN gateways may not necessarily yield capacity improvements in practice. The achievable network capacity of real-world LoRaWANs falls significantly short of the theoretical capacity of LoRaWAN (*e.g.*, <35%) and cannot scale to meet the demands of fast-growing IoT connections worldwide.

(2) Capacity contention across networks. Differing from cellular networks, where each operator has exclusive access to licensed spectrum and fully utilizes the capacity, all LoRaWAN deployments (private and public) share the same unlicensed bands and must contend for communication. Our experiments with operational LoRaWANs reveal that the capacity limit (*i.e.*, 16 concurrent users) not only constrains individual network concurrency but also caps the total concurrency across all coexisting networks. Consequently, each network acquires only a fraction of the limited capacity, which

diminishes further as more networks coexist. This capacity degradation arising from multi-network coexistence poses a significant scalability challenge, hindering the widespread deployment of LoRaWAN as a global IoT infrastructure.

New Findings. Existing studies primarily attribute such capacity limitations to packet collisions and interference in the shared wireless medium [9, 12, 46, 49, 55]. However, our study indicates that even in the absence of packet collisions, LoRaWANs still experience significant capacity gaps. Our in-depth investigation with operational LoRaWANs (as detailed in § 3) identifies four primary factors that contribute to the practical capacity gaps:

(1) Limited capability per gateway. A single gateway monitors multiple channels but can only receive a small portion of packets limited by decoder resources (*e.g.*, 16 decoders per gateway). As incoming packets are processed in a First-Come-First-Served (FCFS) way, the later-arriving packets have to be discarded when decoders are fully occupied.

(2) Inefficiency with homogeneous reception. Operators typically configure their gateways with standard LoRaWAN channels. As a result, these gateways operating within the same spectrum observe identical packets in the same order, leading all gateways to receive the early packets while none captures the later ones. This homogeneous reception results in resource waste and forfeits the opportunity to utilize the additional decoders from multiple gateways to receive new packets and enhance overall network capacity.

(3) Non-optimal operational strategy. LoRaWAN does not associate users with dedicated gateways. Instead, all gateways within range receive and forward a user's packets. While this strategy enhances coverage and redundancy, it can also introduce inefficiencies—some users may unnecessarily occupy decoder resources across multiple gateways, while others may be left without access to any decoders.

(4) Inefficient spectrum sharing. Current LoRaWANs lack mechanisms for spectrum sharing among network operators. In multi-network coexistence scenarios, a gateway operating in the shared spectrum detects packets not only from its own network but also from other coexisting networks. Although LoRaWAN packets include network identifiers, these identifiers cannot be accessed until the packets are successfully decoded. As a result, packets from all coexisting networks compete for and consume decoder resources at each gateway prior to packet reception or rejection decisions.

Based on these new findings, this paper reveals that the practical capacity gaps observed in operational LoRaWANs are fundamentally caused by the contention for gateway decoder resources among both intra- and inter-network users, rather than by conventional issues such as network interference, packet collision, or spectrum access contention.

Our Proposal and Contributions. We propose four basic design principles and eight strategies to address the decoder contention problem, as summarized in Table 1. In-depth investigations are conducted to explore the suitability of proposed strategies for mitigating decoder contention in real-world LoRaWANs. Our explorations uncover the strengths and limitations of each strategy (§ 4.2). Finally, we select four strategies for implementation in our *AlphaWAN* system based on three criteria: (1) no modification to COTS hardware or the underlying protocol; (2) compatibility with legacy systems; and (3) compliance with LoRaWAN and ISM band regulations.

AlphaWAN integrates the strategies into the LoRaWAN stack as two new primitives: The *intra-network channel planning* involves the joint optimization of channel configurations for both gateways and end nodes. This primitive minimizes decoder contention among users within the same network and facilitates the efficient use of decoder resources across multiple gateways to enhance overall network capacity. The *inter-network channel planning* introduces a mechanism for efficient spectrum sharing among network operators. It assigns coexisting networks to frequency-misaligned channels, physically isolating their communications to mitigate inter-network contention.

We implement *AlphaWAN* based on an open-source LoRaWAN platform ChirpStack [1]. We evaluate *AlphaWAN* with real-world experiments and large-scale emulations. The results show that *AlphaWAN* supports 3× more IoT connections using the same spectrum and gateway resources. By deploying and leveraging extra gateways, *AlphaWAN* further increases network capacity and approaches the theoretical limit of LoRaWAN. Moreover, *AlphaWAN* enables harmonious coexistence of up to six networks, achieving a 778.1% improvement in spectrum utilization.

Our research makes the following contributions: (1) We demonstrate that today's operational LoRaWANs fall short of capacity to serve IoT connections at scale. We delve into the underlying gateway design and network operations, identifying a new issue termed *decoder contention problem* as the root cause of the practical capacity gaps. (2) We explore eight strategies to address the decoder contention problem and extensively investigate their feasibility and practicability for real-world LoRaWANs. (3) We present *AlphaWAN* to comprehensively manage intra- and inter-network contentions. The source codes of the project can be accessed with https://a1phawan.github.io.

Ethics: This work does not raise any ethical issues.

2 BACKGROUND AND MOTIVATION

2.1 LoRaWAN Based IoT Connectivity

The Internet of Things (IoT) landscape has been undergoing a paradigm shift from short-range connectivity, such as Wi-Fi, BLE, and ZigBee, toward Low-Power Wide-Area Networks (LPWANs). The prevalence of LPWANs is primarily owing to their long communication range and high energy efficiency, enabling a single gateway to cover vast areas (*e.g.*, >10 km in suburban and 5 km in urban [19]). LPWANs are particularly well-suited for low-data-rate applications that demand long-term operation of battery-powered IoT devices, often lasting several years without maintenance.

Among the LPWAN technologies, LoRaWAN, NB-IoT, and LTE-M account for over 90% of the global market share [21]. Unlike NB-IoT and LTE-M, which utilize licensed spectrum, LoRaWAN operates within unlicensed ISM bands at lower operational costs. Moreover, LoRaWAN supports both public and private network deployments, complementing existing technologies to provide flexible and cost-efficient connectivity for a wide range of IoT use cases.

Figure 1 illustrates the architecture of a LoRaWAN system, which comprises end nodes, gateways, network servers, and application

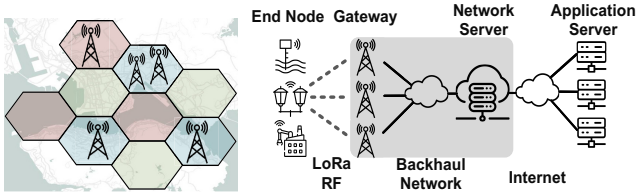

Figure 1: LoRaWAN architecture.

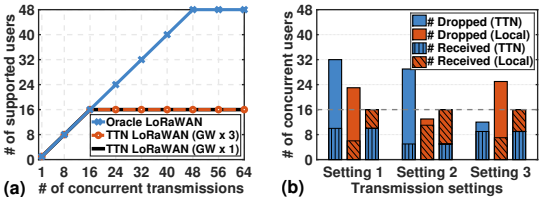

Figure 2: (a) Capacity gaps of TTN LoRaWAN, and (b) capacity when two LoRaWANs coexist.

servers. LoRa is used as an access technology, enabling long-range connectivity between end nodes and gateways. Gateways function as access points and forward packets to network servers via a backhaul link. Typically, the gateways and network servers (*i.e.*, *LoRaWAN infrastructure*) are deployed and maintained by network operators. IoT devices (termed *LoRaWAN users*) subscribe to network operators and access the infrastructure for communications with minimal configuration. LoRaWAN has been officially recognized by the International Telecommunication Union (ITU) as a standard technology for providing LPWAN connectivity to IoT devices [18, 35]. More than 70 operators worldwide offer commercial LoRaWAN services to the public [19, 22].

2.2 Practical Capacity Gaps of LoRaWAN

A central question facing the network operators is whether LoRaWAN can scale to meet the demand of massive IoT connectivity. Since all LoRaWAN networks share the same unlicensed ISM bands, it is common for multiple LoRaWANs from different network operators to coexist and compete. A more pressing concern arises: Can LoRaWAN deliver sufficient network capacity under such coexistence conditions? To explore this, we conduct preliminary studies using real-world LoRaWANs to assess their practical capacity.

In our context, network capacity is measured as the maximum number of IoT connections a network can support, rather than the traditional metric of throughput in cellular and Wi-Fi networks. LoRaWAN supports massive IoT connectivity through two key mechanisms: (1) concurrent communication across different channels and with orthogonal data rate settings; (2) duty cycling, which limits the active time of each LoRaWAN user and scatters user communications across time. Importantly, the number of concurrent communications a network can handle fundamentally bounds its maximum user connections. Based on this insight, our study adopts the maximum number of concurrent users as a core metric for evaluating LoRaWAN's network capacity.

We conduct experiments with two operational LoRaWANs: one globally-operated LoRaWAN with enterprise-grade gateways and network servers operated by TTN [18], and one local LoRaWAN with COTS gateways operated by a local service provider. For the TTN LoRaWAN, we subscribe to their services and connect our LoRa nodes to their gateways for experiments. For the local LoRaWAN, we build an experimental LoRaWAN system using the same hardware and spectrum settings. The network server is configured with an open-source ChirpStack server [1].

In our experiments, we schedule varying numbers of nodes to transmit concurrently using different sub-channels and data rates. Application servers record the number of successfully received packets from the two LoRaWAN networks. Figure 2a presents the reception results from the TTN LoRaWAN. For comparison, we also plot the theoretical maximum number of concurrent users that could be supported by LoRaWAN as Oracle. A significant gap is observed between the practical and theoretical capacities: the TTN LoRaWAN receives at most 16 concurrent packets, which is only one-third of the theoretical limit. To investigate whether additional gateways could increase overall network capacity, we deploy two more gateways operating on the same spectrum and repeat the experiments. Surprisingly, no capacity improvement is observed in either the TTN LoRaWAN (Figure 2a) or our local LoRaWAN (not shown).

Next, we study the impact of multi-network coexistence. We configure nodes from both the TTN LoRaWAN and our local LoRaWAN to transmit concurrently within the same frequency band (923–925 MHz). Figure 2b presents the number of received packets and the dropped ones of each network under various transmission settings. Despite variations in the reception performance of each LoRaWAN, we unexpectedly observe that the total number of successfully received packets from the two coexisting networks always adds up to 16 across all settings. This observation indicates that the practical capacity limit (*i.e.*, 16 concurrent users) indeed constrains the aggregate number of simultaneous transmissions across all co-located networks, with each network acquiring only a portion of the available capacity.

2.3 Research Goal

Our research aims to uncover the causes of LoRaWAN's practical capacity gaps and to explore effective solutions for mitigating these issues. The ultimate goal is to facilitate massive connectivity with LoRaWAN by enabling more IoT transmissions within the shared ISM bands, closing the gap between practical performance and theoretical potential. This vision advances LoRaWAN toward large-scale deployment and harmonious multi-network coexistence, better supporting the rapidly growing IoT connections worldwide.

3 DEMYSTIFYING LORAWAN CAPACITY

3.1 Understanding Gateway Receptions

Although LoRaWAN gateways can adopt various Semtech chipsets (*e.g.*, SX1301, SX1302, and SX1308), they share a common architecture with similar hardware and packet reception pipeline. We

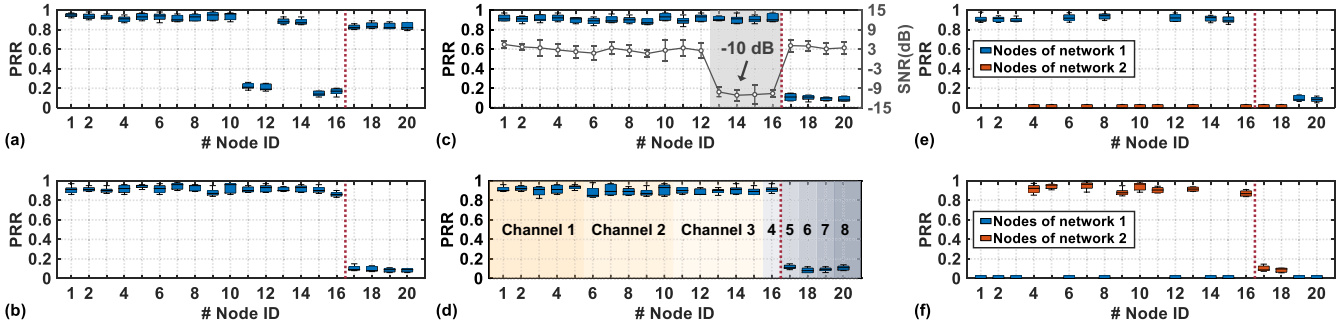

Figure 3: (a,b) Packet Reception Ratios (PRRs) of 20 concurrent nodes with Scheme (a) and Scheme (b); (c,d) impacts of packet SNRs and channel crowdedness; (e,f) PRRs at gateway 1 and gateway 2 of two coexisting LoRaWANs.

use a COTS gateway (WisGate RAK7268CV2) [42] with an embedded SX1302 radio for a case study. This COTS LoRa gateway can simultaneously monitor on eight channels with a frequency span of 1.6 MHz. However, this gateway only receives 16 concurrent packets at most, limited by its computation and storage resources (*e.g.*, only 16 decoders available [44]). This serves as a physical limit for the capacity of COTS gateways based on the Semtech chipsets.

To understand how a COTS gateway receives packets when the number of concurrent packets exceeds the gateway's capacity, we set up the gateway and control 20 nodes to transmit concurrently over different channels at different data rates without packet collisions among the nodes. We schedule the nodes to transmit in 20 micro time slots (*i.e.*, node i in the i^{th} slot) with two schemes: (a) ensuring that the leading preamble symbol of the 20 nodes arrives at the gateway in order, and (b) ensuring that the final preamble symbol of the 20 nodes arrives in order. Comparing the results of the two schemes in Figures 3a and 3b, we confirm that COTS gateways start receiving a packet after finishing receiving the packet's preamble, *i.e.*, termed *lock-on*. In particular, the incoming concurrent packets are received in the order of their lock-on time. Packets arriving late are dropped when the number of packets exceeds the gateway's capacity.

Using the settings of Scheme (b) as a basis, we alter the SNR conditions and channel distributions of the 20 nodes to study their impacts on packet receptions. The results, presented in Figure 3c and Figure 3d, demonstrate that the gateway does not prioritize high-SNR packets over low-SNR ones, as long as the packet SNRs meet the basic reception threshold. Packets from crowded channels (*e.g.*, Channels 1~3) and those from idle channels (*e.g.*, Channel 4) are treated fairly by the gateway. The decision to receive or drop a packet is purely based on the lock-on time of the packets, as long as their SNRs suffice for packet decoding.

To study how the transmissions from different coexisting networks impact a gateway's reception behaviors, we deploy another LoRaWAN within the same spectrum and disperse the 20 nodes across two LoRaWANs (10 nodes each). As advised by the LoRaWAN specification [2], the nodes of the two networks use different frame sync words in their transmitted packets. Figure 3e and Figure 3f plot the packet reception results of the two networks respectively. For Network 1, the gateway does not receive packets from Network 2. However, Network 2's packets still contend for and occupy Network 1's gateway resources, preventing the gateway from receiving some later packets from Network 1. The same phenomenon is observed in Network 2 as shown in Figure 3f. These results indicate that the gateways must first decode a packet to extract the embedded sync words and then perform packet filtering. Therefore, all packets from both networks compete for the decoder resources and capacity of every gateway.

By assembling the experimental observations, we can unveil the underlying reception mechanism of a COTS gateway (see details in § C). The key insights are: (1) LoRaWAN gateways process incoming concurrent packets in a First-Come-First-Served (FCFS) manner based on the lock-on time of the packets; (2) A gateway must first decode an incoming packet to determine whether it is destined to its network.

3.2 Capacity Gaps: Cause and Implications

Our studies have confirmed that a single LoRaWAN gateway suffers from capacity limitations mainly due to the constraint of decoder resources. We have examined other COTS gateways and summarized their capacity data (Table 4 in § C). None of these gateways has sufficient decoders to fully support the theoretical capacity of their operating channels.

Can we deploy more gateways to improve capacity? Our field studies observe that commercial LoRaWAN operators such as TTN [18], Senet [36], and ZENNER [60] typically run networks following the LoRaWAN standard. When multiple gateways operate within the same spectrum, they typically choose from the standard channel plans with homogeneous frequency settings. As a result, multiple co-located gateways observe the same incoming packets in the same order. As all gateways process concurrent packets following the FCFS mechanism, the gateways receive the same early packets and drop the same late packets. Consequently, LoRaWAN operators cannot improve the overall network capacity by deploying more gateways.

Coexisting LoRaWANs suffer capacity degradation. Due to the lack of spectrum sharing and coordination mechanisms, coexisting LoRaWANs often operate with the same standard channel plans. Although the LoRaWAN protocol introduces identifiers such as frame sync words, session keys, and node IDs to distinguish the

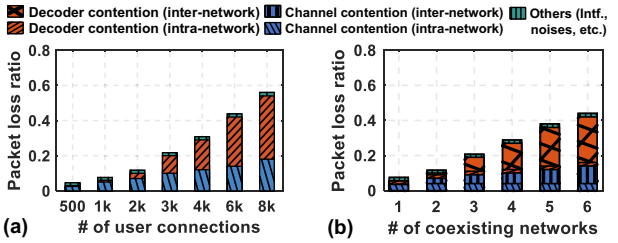

Figure 4: Decoder contention problem (a) under varying user scales of a single LoRaWAN, (b) when multiple LoRaWANs coexist (each network connects 1k users).

packets of different networks, these identifiers cannot be extracted until the packets are successfully decoded. As a result, packets from all co-located networks compete for the limited decoder resources of each gateway and may occupy gateway capacity without yielding any successful packet receptions. Consequently, the decoder resources at individual gateways become a bottleneck that constrains the aggregate capacity of all coexisting LoRaWANs. As more operators deploy LoRaWANs in the same area, the capacity available to each network diminishes due to increased contention for gateway decoder resources.

In summary, the capacity gaps observed in operational LoRaWANs are fundamentally caused by the *Decoder Contention Problem* as follows.

Decoder Contention Problem. The practical capacity of a LoRaWAN is limited by the decoder resources of gateways. Suboptimal LoRaWAN operating strategies, such as homogeneous channel configurations across gateways and coexisting networks, lead to intensive contentions for gateway decoder resources among both intra- and inter-network users. These contentions are further intensified by LoRaWAN's unique characteristics, including operation in unlicensed ISM bands, long communication range, and the absence of user-gateway association. Specifically, because LoRaWAN allows any gateway within range to receive and forward packets from any user, some users unnecessarily consume decoder resources at multiple gateways, while others fail to access any available decoders. These factors collectively contribute to the practical capacity gaps of LoRaWANs.

To investigate the practical implications of the decoder contention problem, we record and analyze packet loss in real-world LoRaWAN deployments, classifying the causes into three categories: decoder contention, channel contention (*i.e.*, multiple nodes using identical transmission settings), and other factors (interference, poor SNRs, *etc*). The results are presented in Figure 4. Notably, while channel contention is traditionally considered the dominant source of packet loss in large-scale LoRaWANs [9, 46, 55], our results reveal that decoder contention becomes more critical as the network size surpasses 3,000 users, as shown in Figure 4a. Even in smaller-scale deployments (*e.g.*, 1,000 users per network), cross-network decoder contention emerges as the leading cause of packet loss when three or more networks coexist (see Figure 4b). These results call for urgent and targeted solutions to address decoder contention, particularly in the context of large-scale and coexisting LoRaWAN deployments.

4 *ALPHAWAN*: NEXT-G GLOBAL IOT

4.1 Principles for Boosting Capacity

To enhance the capacity of LoRaWANs, our work aims to mitigate the bottlenecks by addressing various factors contributing to the decoder contention problem. We propose four design principles: (1) Optimize spectrum utilization, (2) Add more decoders and spectrum resources, (3) Manage user contention, and (4) Isolate coexisting networks.

In this section, we first present strategies following the above principles to address the decoder contention problem for capacity enhancement, then investigate their practicality for operational LoRaWANs, followed by the system implementation of *AlphaWAN* in § 4.3.

4.2 Strategies and Design Considerations

4.2.1 Optimizing spectrum utilization. This category of strategies focuses on optimizing channel configurations among gateways to enable them to more effectively accommodate a greater number of users within the same spectrum.

Strategy ① increases per-channel decoder resources to support more users over each channel. This can be realized by changing the number of channels and frequencies a gateway operates on, both of which are programmable on COTS gateways. Our rationale is that if a gateway operates on fewer channels, all decoder resources of the gateway are exclusively reserved for those channels. As a result, each channel has more resources to service more concurrent users.

We have practically tested this strategy with five gateways operating in a 1.6 MHz spectrum (see Figure 5a). The total capacity of the spectrum increases from 16 to 48 concurrent users as the number of channels per gateway reduces from 8 to 2. This strategy can work with all commercial LoRaWAN gateways including older models, enabling operators to expand network capacity without upgrading gateways. However, a drawback of this strategy is that it may under-utilize gateways' radio resources such as Rx chains.

Strategy ② addresses the decoder contention problem by proactively configuring LoRaWAN gateways with heterogeneous channels. Heterogeneous channel configuration enables co-located gateways to detect distinct sets of packets across different frequencies, with each gateway observing a unique subset of users contending for its decoder resources. This increases the likelihood that a packet arriving late at one gateway may be received earlier at another, thereby improving the chances of successful reception. Consequently, packets from delayed users, which would otherwise be dropped by all gateways in current LoRaWANs, gain a new opportunity for delivery. Importantly, this strategy allows the decoder resources of multiple gateways to be utilized collectively, enhancing the system's ability to receive additional packets.

A feasibility study is presented in Figure 5b. As the three gateways operate with heterogeneous channels, a capacity improvement from 16 to 24 concurrent users is observed. We expect higher gains by employing more gateways and optimizing their frequency plans.

Principles	Strategies	Implementation method	Practicability	Adopted in *AlphaWAN*?
Optimize spectrum utilization	① **Improve per-channel resource utilization**	Adjust the number of channels per GW	Programmable, supported by COTS GWs	**Yes**
	② **Heterogeneous channel configuration**	Diversify channel configurations of GWs	Supported by COTS GWs	**Yes**
Add extra resources	③ More decoders per GW	Upgrade to the newest GWs	Not supported by legacy GWs	No
	④ More spectrum resources	Expand to new frequency bands	Limited ISM bands for LoRaWAN	No
Manage user contention	⑤ Smaller cell with shortened transmit range	Adaptive Data Rate, transmit power control	Suboptimal spectrum utilization	No
	⑥ Divide large cells into sub-regions	Directional antennas	Less effective to LoRaWAN	No
	⑦ **Contention management for LoRaWAN**	Joint channel planning and ADR/TPC optimize	Supported by COTS GWs and end-nodes	**Yes**
Isolate coexisting networks	⑧ **Spectrum sharing across operators with misaligned channel plans**	Create channel plans per operator with optimal frequency misalignment	Supported by COTS GWs and the LoRaWAN standard	**Yes**

Table 1: Summary of proposed strategies to address the decoder contention problem, implementation methods, and their practicability for LoRaWAN.

4.2.2 Adding extra resources. **Strategy ③** addresses the decoder resource limitation by adding more decoders per gateway. For instance, the latest RAK gateway [41] employs two SX1303 radios with 32 decoders embedded, supporting up to 32 concurrent users (see Table 4). This strategy, however, requires hardware modifications to gateways. Operators may need to upgrade their infrastructure with the latest gateways to achieve higher capacity.

Strategy ④ aims to expand LoRaWAN operations to new spectrum, thereby reducing user contention in the crowded ISM bands and utilizing extra frequencies to serve more users. Operators can apply for new spectrum from local authorities (*e.g.*, FCC in the US). However, this strategy is less practical for operational LoRaWANs at present. Besides, though more frequencies can increase the total network capacity, the per-spectrum user capacity does not improve.

4.2.3 Managing user contention. **Strategy ⑤** aims to decrease the communication ranges of LoRaWAN devices to restrict decoder contention among users within smaller regions. The current LoRaWAN offers an Adaptive Data Rate (ADR) mechanism, which could be leveraged to control the cell size of gateways by adjusting transmission power and data rate on a per-user basis.

We have tested the effectiveness of ADR within real-world LoRaWANs. As shown in Figure 6, ADR can effectively adapt the cell size, both in coverage range and user connections of the gateways. Without ADR, each user connects to seven gateways on average, meaning that a single user transmission competes for resources at seven gateways. ADR can reduce contention from seven gateways to two gateways, freeing five gateways for serving additional users. However, a side effect of LoRaWAN ADR is that it aggressively reduces cell size using high data rates (*e.g.*, >90% in DR5), as shown in Figure 6d. Similar results are observed in the TTN LoRaWAN, with 53.7% in DR5 as shown in Figure 6e. Such *unbalanced data-rate usage* can lead to suboptimal spectrum utilization and compromise the maximum user capacity per cell.

Strategy ⑥ seeks to use directional antennas to separate the packets from different directions and thus reduce user contention. However, our practices with the RAKwireless 12 dBi directional antennas [43] show that although packets from non-steered directions are weakened by 14 - 40 dB in signal power (see Figure 7), they can still be received at the gateways and compete with users from the steered direction. This is due to the high sensitivity of LoRaWAN radios (*e.g.*, down to -148 dBm [44]). A LoRaWAN radio can reliably receive packets even when the signal is weaker than the noise. Thus, using directional antennas alone may not effectively reduce user contention.

Remarks: Conventional techniques, such as ADR, transmission power control, and directional antennas, can reduce contention to some extent by decreasing the physical signal strength of users' packets. However, the high sensitivity of LoRaWAN radios makes these techniques less effective. This necessitates addressing the decoder contention problem from a contention management perspective as follows.

Strategy ⑦ comprehensively manages user contention in LoRaWANs by collectively using ADR and frequency planning techniques. Unlike the current LoRaWAN ADR, which links users to the nearest gateways, we propose allowing some users to bypass congested nearby gateways and instead transmit to less crowded ones farther away. To achieve this, our strategy reallocates certain users from a congested nearby gateway to a less crowded one by switching users' frequencies to the channels operated by the target gateway. A joint optimization method can be employed to determine the best operating channels and user configurations (data rate and transmission power) for both gateways and end nodes. Although this strategy may increase power consumption for some users, it effectively balances workloads between nearby and distant gateways, alleviating decoder contention at congested ones. More importantly, it ensures full utilization of spectrum resources (high and low data rates) and the decoding capabilities of all gateways (nearby and distant), ultimately enhancing overall network capacity.

4.2.4 Isolating coexisting networks. **Strategy ⑧** isolates the users of coexisting networks by operating co-located LoRaWANs with distinctive frequency plans. As illustrated in Figure 8, the channels of coexisting networks should maintain proper frequency misalignment. This frequency misalignment could allow COTS LoRaWAN devices to exploit the inherent frequency selectivity of radio hardware to isolate packets from other co-located networks. Specifically, when a radio operates on a channel, the radio hardware truncates signals outside the frequency ranges of that intended channel (termed *frequency selectivity*). Packets not belonging to the intended network, due to frequency misalignment, would be truncated and dropped, not flowing into the following radio pipeline of packet decoding. Only packets intended for the network can flow

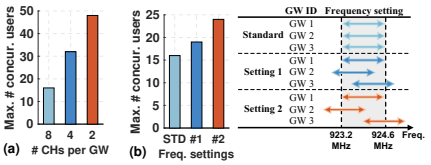

Figure 5: Capacity gains of (a) operating fewer channels per GW, and (b) heterogeneous channel adoptions.

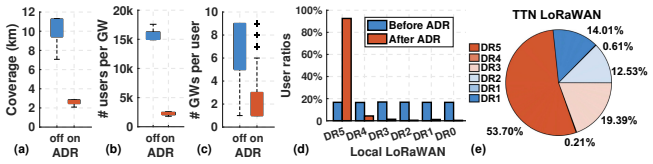

Figure 6: (a-c) Gateway cell size with and without ADR, (d,e) The current LoRaWAN ADR leads to unbalanced data-rate usage.

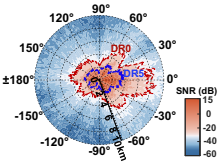

Figure 7: Directional antenna for LoRaWAN.

into the pipeline and consume decoder resources for packet reception. Packets from different co-located networks are thus naturally separated without incurring cross-network decoder contentions.

Notably, as the channels of coexisting networks overlap in frequency, there may be concerns about inter-channel interference. Thanks to the rich orthogonal data-rate settings of LoRaWAN, the implications of inter-channel interference can be negligible for practical LoRaWANs. Our experiments with COTS devices, both gateways and end nodes, show that overlapping channels with >40% frequency misalignment can ensure correct reception for over 80% of packets even using non-orthogonal data rates, as shown in Figure 8.

4.3 System Implementation

This section details the system implementation of *AlphaWAN*, which integrates **Strategies** ①, ②, ⑦, and ⑧ into the LoRaWAN stack, providing comprehensive deployable solutions for operational LoRaWANs. We choose these strategies for several reasons:

- The selected strategies comprehensively address the decoder contention problem from different design spaces: ① and ② focus on spectrum optimization, ⑦ targets contention management, and ⑧ facilitates spectrum sharing across coexisting networks. They work complementarily to enhance capacity and scalability for LoRaWANs.
- These strategies enhance the LoRaWAN capacity from new perspectives, differing from the conventional wisdom adopted in cellular networks and WLAN. These strategies are specifically designed to address the decoder contention problem in LoRaWANs and exploit the unique capabilities of LoRaWAN devices and network operation characteristics for capacity enhancement.
- These strategies are compatible with COTS gateways as well as millions of legacy LoRa nodes and do not require any modifications to the underlying protocol or system operations of LoRaWANs. They can be readily implemented at the application layer and seamlessly integrated to upgrade commercial operational LoRaWANs.

From an implementation standpoint, the selected strategies, when adopted together, incorporate comprehensive channel planning and parameter optimization for gateways, end nodes, and across coexisting networks. *AlphaWAN* implements these strategies as two LoRaWAN primitives: (1) *intra-network channel planning* for gateways and end nodes within the same network, and (2) *inter-network channel planning* for coexisting networks across different operators.

4.3.1 Intra-network channel planning. This subsection addresses Channel Planning (CP) for a single network. The CP problem adopts strategies ①, ②, and ⑦ collectively to decide optimal channel adoptions for both gateways and end-nodes, formulated as follows.

CP input. We represent a LoRaWAN network as a triplet (GW, ND, CH), where GW, ND, and CH denote the sets of gateways, end nodes, and frequency channels, respectively. To model the effects of LoRaWAN ADR and transmit power control, we simplify the communication ranges of end nodes into various discrete distances, denoted by a set DR. $R_{ND\times GW\times DR}$ records the coverage relationships between nodes and gateways, where $r_{ijl} = 1$ if the i^{th} node can physically reach the j^{th} gateway using the l^{th} transmission distance, and 0 otherwise. Let U^t_{ND} represent the traffic rates of end nodes within a time window t. We introduce three constants, *i.e.*, C_j, P_j, and B_j, to denote the maximal decoders, maximal operating channels, and maximal radio bandwidth of the j^{th} gateways, respectively. U^t_{ND} is derived from the history traffic records of gateways (detailed in § 4.3.3), while the other data are known a priori.

CP output. The decisions made by CP include the number of operating channels and their detailed frequency configurations for both gateways and end nodes, as well as the ADR and transmit power settings for end nodes. We use binary variables f_{ik} and h_{jk} to indicate if the i^{th} node and the j^{th} gateway operate on the k^{th} frequency channel, and variable d_{il} to indicate if the i^{th} node uses the l^{th} transmission distance. The specific data rate and transmit power settings for a node are derived from the required transmission distance (*e.g.*, using a mapping table).

LoRaWAN user contention model. Connectivity between the i^{th} node and the j^{th} gateway is established if (1) the gateway is within the node's communication range and (2) the gateway's operating channels support the node's frequencies. We use $link_{ij}$ to indicate whether the i^{th} node can connect to the j^{th} gateway: $link_{ij} = 1$ if $\sum_{k\in CH, l\in DR}(r_{ijl}\cdot h_{jk}\cdot f_{ik}\cdot d_{il}) > 0$; otherwise $link_{ij} = 0$.

Given the nodes' traffics U^t_{ND}, we can count the number of packets transmitting over the channels of the j^{th} gateway as $k_j = \sum_{i\in ND} link_{ij}\cdot u^t_i$. These packets compete for the gateway's decoder resources and can potentially suffer from packet losses if k_j exceeds the gateway's capacity C_j. The higher k_j, the greater the potential for packet loss. Thus, we indicate the risk of packet loss at the j^{th} gateway as $\varphi_j = k_j - C_j$ if $(k_j > C_j)$, and 0 otherwise. The *risk of packet loss for a node* is defined by the minimum risk across all

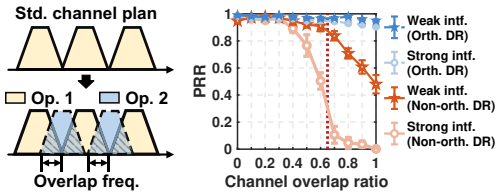

Figure 8: Overlapping channels and the packet performance.

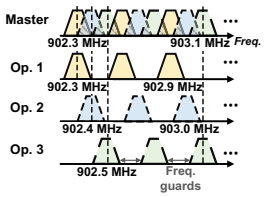

Figure 9: Spectrum sharing in *AlphaWAN*.

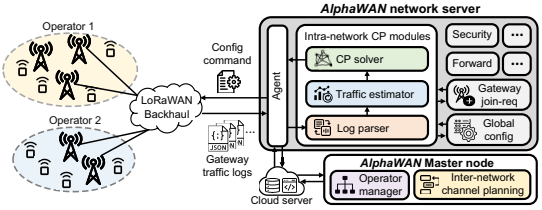

Figure 10: System architecture of *AlphaWAN*.

gateways serving the node, denoted as $\Phi_i = \min_{j \in GW}\{\varphi_j | link_{ij} = 1\}$.

Optimization goal and constraints. The optimization objective of the CP problem is to minimize packet loss risks for all nodes, which is formulated below.

$$\text{Minimize} \quad \sum_{i \in ND} \{\Phi_i\} \tag{1}$$

A feasible channel plan is subject to the following constraints:

Node connectivity constraint: Each node must connect to at least one gateway, *i.e.*, $\forall i \in ND : \sum_{j \in GW} link_{ij} > 0$.

Gateway radio constraints: The frequency span of channel settings for any gateway must not exceed the radio's maximal bandwidth B_j, expressed as $\forall j \in GW : \max\{k|h_{jk} = 1\} - \min\{k|h_{jk} = 1\} \leq B_j$. The number of operating channels is constrained to $\forall j \in GW : \sum_{k \in CH} h_{jk} \leq P_j$.

The formulated problem is a variant of the Knapsack Problem, which is NP-hard [25]. *AlphaWAN* runs an evolutionary algorithm [26] on a central server to search for approximate solutions to the problem. Notably, the 'goodness' of a computed solution depends heavily on the node traffic U^t_{ND}. A better solution can be expected if U^t_{ND} accurately represents the actual user traffic in a LoRaWAN. To improve the representativeness of U^t_{ND}, *AlphaWAN* measures user traffic using optimally selected time windows and aggressively uses samples with high capacity demand to train the problem solver. This ensures the computed channel plan can deliver satisfactory network capacity for the ever-increasing demands.

4.3.2 Spectrum sharing across networks. *AlphaWAN* facilitates spectrum sharing by implementing Strategy ⑧. This involves different operators defining and maintaining distinctive operating channels with suitable frequency misalignment. According to our empirical studies in Figure 8, coexisting channels with ≤70% overlapping ratios (*i.e.*, >30% frequency misalignment) give satisfactory reliability. As a larger overlapping ratio enables more networks to operate within the same spectrum, the optimal frequency misalignment should be determined based on the potential number of coexisting networks in a region.

To better coordinate spectrum usage, *AlphaWAN* shifts the responsibilities of channel division and maintenance from individual operators to a centralized Master node. The Master estimates the maximum number of networks coexisting in a region and selects a frequency misalignment to divide the LoRaWAN spectrum into frequency-overlapping sub-channels, as illustrated in Figure 9. Operators should register with the Master before deploying LoRaWAN infrastructure in a region. The *AlphaWAN* Master keeps an up-to-date record of channel occupancy in the area and assigns channels to operators based on their requests. As illustrated in Figure 9, different operators receive unique channel plans to minimize potential inter-network interference. Operators use the allocated channels to configure gateways and end nodes (*e.g.*, through intra-network channel planning), ensuring the smooth operation of their networks in the region.

4.3.3 Integration with LoRaWAN stack. Figure 10 presents the architecture of *AlphaWAN*. It operates over the backhaul of a LoRaWAN network, with core components running on the LoRaWAN network server and various application-layer agents running on gateways. The detailed implementations of *AlphaWAN* are described below.

Network server. We implement *AlphaWAN*'s channel planning functional components in an open-source LoRaWAN network server platform ChirpStack [1]. Specifically, three new modules are added to ChirpStack:

Log parser: Gateways send the data packets from end devices, along with metadata like receiving channel, timestamp, and SNR, to ChirpStack where the metadata is stored in operational logs. The log parser interprets the metadata from all gateways to extract information such as user traffic and user-gateway link profiles for the CP input.

Traffic estimator: This module combines data across gateways to restore the actual traffic patterns of end nodes. Representative traffic data from different time windows are selected as input for the CP problem solver.

CP solver: This module solves the Channel Planning (CP) problem and transforms the CP output into detailed channel configurations for gateways and end nodes.

Master node. We implement the *AlphaWAN* Master node as an independent process running on a cloud server. It accepts requests from operators (*i.e.*, the inter-network channel planning module on the network server) and responds with channel assignments, with data exchanges implemented via TCP. An operator bootstraps its network by updating the assigned operating channels in ChirpStack using the LoRaWAN channel creation commands [2].

Gateways. We implement the end-point agents of *AlphaWAN* at gateways, which receive commands from the server end of *AlphaWAN* and apply the updated channel configurations, if any, to gateways. These *AlphaWAN* agents are implemented using application-layer scripts that execute in a sandbox environment to configure gateway devices.

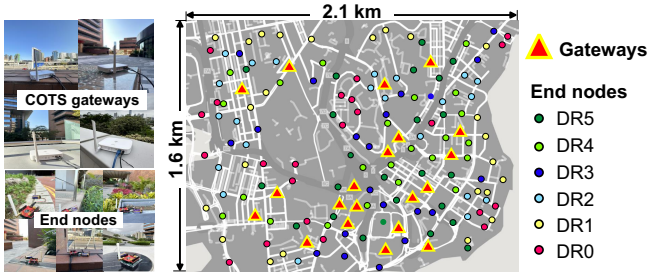

Figure 11: Devices and deployment map of the testbed.

End-devices. *AlphaWAN* exploits the LoRaWAN ADR commands [2] to configure frequency channels, data rates, and transmit power for end nodes. These commands are issued by the ChirpStack server and are supported by all COTS devices implementing the LoRaWAN protocol.

We note that *AlphaWAN* can be readily implemented with COTS gateways and does not require any hardware modifications to deployed LoRa nodes, which allows smooth upgrade of today's LoRaWAN infrastructure.

5 EVALUATION

We have implemented *AlphaWAN* and integrated it into ChirpStack [1]. A testbed system has been built using a ChirpStack network server and up to 20 RAKwireless gateways, providing LoRaWAN connectivity to a 2.1 km × 1.6 km urban area (see Figure 11) with diverse link conditions, including outdoor, indoor, building blockage, *etc.* Our experiments deploy up to 144 COTS LoRa nodes consisting of Dragino LoRa shields embedded with Semtech SX1276 radio. These LoRa shields are connected to Arduino Uno boards, which configure the radio chips to send LoRa packets with specific parameters. We conduct testbed experiments to evaluate the effectiveness of the various strategies adopted by *AlphaWAN* (§ 5.1). We study the applicability of *AlphaWAN* in addressing some real-world operational issues in § 5.2. Additional microbenchmark evaluations are presented in § 5.3.

5.1 Testbed Evaluation

5.1.1 More gateways, more gains! We first evaluate the capacity enhancement performance of *AlphaWAN*. We incrementally deploy 1~15 gateways and operate all gateways within the 916.8–921.6 MHz spectrum (4.8 MHz wide). This spectrum has 24 LoRaWAN channels, allowing 144 users to transmit concurrently. We set up 144 COTS LoRa nodes with different channels and orthogonal data rates. We schedule all nodes to transmit concurrently and count the number of packets received by the ChirpStack server.

We evaluate *AlphaWAN* using two configurations: a full version, which optimally adapts the number of operating channels per gateway, and a variant with Strategy ① disabled, where the number of operating channels per gateway is fixed at eight and does not reduce dynamically. We compare *AlphaWAN* against two baselines: (1) standard LoRaWAN, which uniformly configures gateways using three standard channel plans, and (2) a randomized channel planning strategy (Random CP), which adjusts the number of channels per gateway following Strategy ① but assigns channels to gateways at random.

Figure 12a presents the maximum number of concurrent users supported by different strategies. For reference, the theoretical bound of LoRaWAN is also plotted, labeled as LoRaWAN (Oracle). Our results show that both versions of *AlphaWAN* support more concurrent users than standard LoRaWAN, which supports 48 users. Without Strategy ①, *AlphaWAN* achieves a 143.3% improvement in capacity. When Strategy ① is enabled, *AlphaWAN* exhibits linearly increasing capacity as the number of gateways increases from 1 to 9, and eventually reaches the theoretical maximum, effectively mitigating the capacity gaps. Moreover, *AlphaWAN* achieves 132.3% higher capacity compared to Random CP.

5.1.2 Spectrum efficiency. This experiment evaluates the spectrum efficiency of *AlphaWAN*. We adopt the same experimental settings as described in § 5.1.1, with the number of gateways fixed at 15. We vary the total operating spectrum of the network from 1.6 MHz to 6.4 MHz and measure the maximum number of concurrent users supported under each spectrum setting. As shown in Figure 12b, standard LoRaWAN supports 16 concurrent users with 1.6 MHz and 64 users with 6.4 MHz. Both versions of *AlphaWAN* consistently achieve higher capacity than standard LoRaWAN under the same spectrum settings. The overall capacity of *AlphaWAN* proportionately scales with the number of gateways (Figure 12a) and spectrum bandwidth (Figure 12b). To enable a fair comparison of spectrum efficiency, we normalize the network capacity by spectrum bandwidth, reporting per-MHz user capacity, as indicated by the dashed lines in Figure 12b. *AlphaWAN* without Strategy ① supports more users per MHz than standard LoRaWAN, but fewer than Random CP. The full-version *AlphaWAN* achieves the highest spectrum efficiency, 292.2% higher than standard LoRaWAN and 130.7% higher than Random CP.

5.1.3 Contention management. This experiment examines the effectiveness of contention management in *AlphaWAN*. We employ 144 concurrent nodes and 15 gateways and operate them using the same experimental setup as described in § 5.1.1. We implement two variants of contention management: a full version of *AlphaWAN* that applies Strategy ⑦ for comprehensive contention management at both gateway and node sides, and another version that manages contention without cooperation from the node side. As shown in Figure 12c, *AlphaWAN* without cooperation from end nodes improves the mean capacity of LoRaWAN from 42 concurrent users to 57 users. The mean capacity further increases to 68 users when *AlphaWAN* comprehensively manages contention at both gateways and end nodes.

5.1.4 Spectrum sharing. This experiment evaluates the performance of the spectrum sharing mechanism in *AlphaWAN*. We deploy up to six networks within the same spectrum (1.6 MHz wide). Each network consists of 3 gateways and 24 end nodes belonging to different operators. The nodes of all networks are scheduled to concurrently send packets with a 10-byte payload. We operate the gateways with standard LoRaWAN and *AlphaWAN*, respectively, and measure the number of received packets for each network. For standard LoRaWAN, all networks use the same channel plans. For

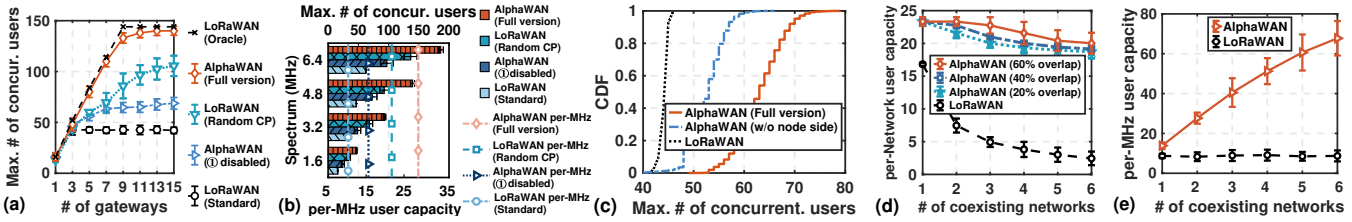

Figure 12: Capacity performance of *AlphaWAN* with (a) different numbers of gateways, (b) operating spectrum, (c) contention management, and (d,e) spectrum sharing among different numbers of coexisting networks.

AlphaWAN, we run a Master node to allocate coexisting networks with distinctive channel plans using frequency overlapping ratios of 20%, 40%, and 60%, respectively.

Figure 12d reports the user capacity of each network across different settings. As more networks are deployed together, the per-network capacity of standard LoRaWAN degrades dramatically. In contrast, *AlphaWAN* maintains a per-network capacity above 23 concurrent users, with only a slight reduction to 20 users when the number of coexisting networks exceeds three. Figure 12e compares *AlphaWAN* and standard LoRaWAN in terms of spectrum efficiency, measured as the total number of concurrent users per MHz (per-MHz user capacity). As the number of coexisting networks increases from one to six, *AlphaWAN* significantly improves spectrum utilization efficiency, achieving 158.9% to 778.1% higher per-MHz user capacity than standard LoRaWAN.

5.2 Study for Practical Network Operations

5.2.1 IoT connectivity at scale. This experiment evaluates the scalability of *AlphaWAN*. We deploy 15 gateways and 144 COTS LoRa nodes within the testbed area shown in Figure 11. The network operates across 24 LoRaWAN channels, spanning a total bandwidth of 4.8 MHz. Each node is configured to transmit packets containing a 10-byte payload. Instead of adhering to the standard 1% duty cycle specified by LoRaWAN, we configure each node with elevated duty cycles and schedule a node to transmit extra packets from different users in the extended active durations. This approach enables a single physical node to emulate the transmissions of up to ten LoRaWAN users across distinct time slots. The 144 LoRa nodes can collectively emulate the network activity of over 14,000 users operating under the standard 1% duty cycle. We run experiments with various strategies: (1) LoRaWAN with ADR disabled, (2) LoRaWAN with ADR enabled, (3) LMAC [12], the state-of-the-art (SOTA) MAC protocol for LoRaWAN to avoid packet collisions, (4) CIC [46], the SOTA collision resolving technique for LoRaWAN, (5) Random CP, a randomized channel planning strategy for LoRaWAN gateways, and (6) *AlphaWAN*. Note that CIC requires specialized gateway implementations, while the other strategies work with COTS gateways. For fairness, we only use CIC for resolving packet collisions and apply the same decoder resource constraints of COTS gateways (*i.e.*, 16 decoders per gateway) to CIC for packet decoding.

Figure 13 presents the network throughput and PRR performance of various strategies under different network scales. The throughput of LoRaWAN (w/o ADR), LMAC, and CIC saturates when the number of users reaches 6,000. LMAC and CIC outperform LoRaWAN (w/o ADR) thanks to their collision avoidance and resolving capabilities. However, as the user scale exceeds 6,000, the decoder contention problem emerges as the primary bottleneck (see Figure 13c). LoRaWAN ADR and Random CP partially mitigate this issue and continue to improve throughput beyond 6,000 users. Nevertheless, LoRaWAN ADR fails to fully utilize available spectrum resources, resulting in unbalanced data rate usage (see Figure 13d) and lower throughput compared to *AlphaWAN*. By leveraging all available data rates and spectrum resources, *AlphaWAN* effectively reduces decoder contention (see Figure 13c) and delivers scalable performance, maintaining over 85% PRR as the network size increases to 12,000 users. Notably, this performance is achieved using 15 gateways operating within a 4.8 MHz spectrum. Our solution is scalable to even larger network sizes by deploying additional gateways and expanding the spectrum.

5.2.2 Coexisting with legacy LoRaWANs. In practice, new network deployments adopting *AlphaWAN* may coexist with legacy LoRaWAN networks. This experiment evaluates the impacts of partial adoptions of *AlphaWAN* on the capacity performance of coexisting networks. We deploy four coexisting networks within the testbed area, using the same experimental settings described in § 5.1.4. We vary the number of networks participating in *AlphaWAN* for spectrum sharing. The remaining networks (not engaged) operate their gateways using the standard LoRaWAN channel plans.

Figure 14 presents the per-network user capacity across different settings. When none of the networks adopts *AlphaWAN*, all four networks compete for communication, resulting in the lowest capacity (averaging four users per network). When two networks (Network 3 and Network 4) adopt *AlphaWAN*, they benefit from spectrum sharing, each obtaining approximately 2× capacity gains. Moreover, the optimized frequency plans of these networks help reduce contention in the legacy LoRaWAN channels, allowing the coexisting legacy networks (Network 1 and Network 2) to achieve slightly improved capacity. However, due to the absence of efficient cooperation from legacy LoRaWANs, legacy networks may still introduce interference into some channels used by *AlphaWAN* networks, limiting the full potential of capacity improvements. As more networks adopt *AlphaWAN* and coordinate spectrum usage, the capacity of all coexisting networks progressively improves.

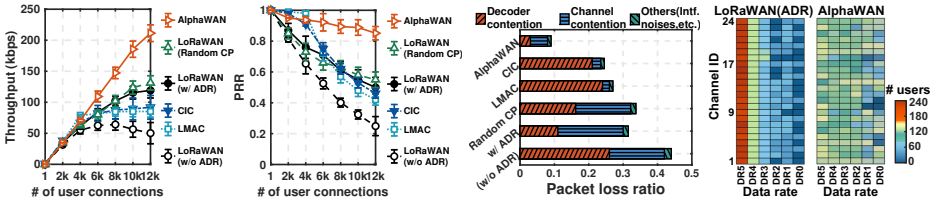

(a) Aggregated network throughput (b) Packet reception ratio (PRR) (c) Factors of packet loss at 6k user scale (d) Spectrum utilization

Figure 13: Scaled LoRaWAN operations: *AlphaWAN* vs. State-of-the-art.

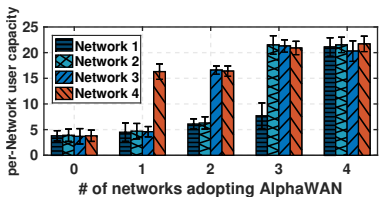

Figure 14: Capacity performance of four coexisting networks with partial adoptions of *AlphaWAN*.

5.3 Micro Benchmarks

5.3.1 Fairness among coexisting networks. This experiment evaluates whether *AlphaWAN* delivers fair capacity performance across coexisting networks. We deploy two LoRaWANs operating within the same 1.6 MHz spectrum, following the setups outlined in § 5.1.4. Each network consists of three gateways. The *AlphaWAN* Master assigns channel plans to the two networks using a 40% frequency overlap. For Network 1, we employ a fixed number of 48 users, scheduled to transmit concurrently across different sub-channels using orthogonal data rates. This setting represents the theoretical maximum capacity that LoRaWAN can support within the 1.6 MHz spectrum. Our experiment varies the number of concurrent users in Network 2 from 16 to 80 and measures the number of users each network can support under different user-load conditions.

Figure 15 reports the service ratios of each network, defined as the ratio of users whose packets are successfully received. The results show that both networks maintain user service ratios above 90% as the number of concurrent users in Network 2 increases from 16 to 48. Each network fairly obtains similar capacity performance (48 concurrent users) in the shared spectrum. When the user load in Network 2 exceeds 48, some users must use identical transmission settings, resulting in channel contention and packet failures. As shown in Figure 15, the service ratio of Network 2 drops significantly. In contrast, Network 1 maintains a relatively high service ratio (*e.g.*, >80%) with just a small drop due to increased interference in the shared spectrum. These results indicate that the communications of the two coexisting networks are effectively isolated, allowing each network to fairly utilize the full capacity of the spectrum.

5.3.2 Impact of spectrum sharing on packet performance. As the spectrum sharing mechanism of *AlphaWAN* operates coexisting networks with frequency-misaligned channels, it might introduce additional inter-channel interference (see § 4.2.4). This experiment evaluates the impact of spectrum sharing on LoRa packet reception.

We deploy two coexisting LoRaWAN networks, each establishing a node-to-gateway communication link. The channels for the two links are configured with 20% frequency overlap. The two links are scheduled to transmit simultaneously. Link 1 (*i.e.*, the link of Network 1) operates with a fixed data rate (DR4), while Link 2 is configured with varying data rates and transmission power levels to emulate different coexistence scenarios. We measure the packet reception performance of a COTS LoRaWAN gateway on Link 1 under varying SNR and coexistence conditions.

The results are presented in Figure 16. We observe that the reception SNR threshold of the gateway without the coexistence of Network 2 is approximately -13 dB. When Network 2 transmits on an overlapping channel with orthogonal data rates, the packet reception threshold does not change much compared to the case without Network 2 coexistence. A 3.3–3.7 dB increase in packet reception threshold is observed when Network 2 uses non-orthogonal data rates, due to the more pronounced inter-channel interference with non-orthogonal transmissions over the shared spectrum.

5.3.3 Latency analysis of AlphaWAN. A capacity upgrade operation in *AlphaWAN* comprises centralized computation (*i.e.*, solving the Channel Planning (CP) optimization problem), distribution of optimal channel configurations to gateways, and rebooting the gateways with the updated settings. When multiple networks coexist, an additional spectrum sharing procedure is required, involving message exchanges between operators and the *AlphaWAN* Master. This experiment evaluates the latency overheads associated with capacity upgrade using *AlphaWAN*.

We measure the end-to-end latency of a complete capacity upgrade operation in a single LoRaWAN network under varying user scales and for different numbers of coexisting networks (each with 3,000 users). The CP optimization problem is executed on a workstation with an Intel Core i5-13490F Processor. The backhaul connections between gateways and the centralized network server (running ChirpStack and the *AlphaWAN* Master) utilize a 2.5 Gbps Ethernet. We separately measure latency for each segment of the procedure. The sum latency corresponds to the time elapsed from the initiation of a capacity upgrade command to the point when the last gateway completes its reboot.

The results are presented in Figure 17. As shown in Figure 17a, the latencies are primarily dominated by CP solving and gateway rebooting. On average, gateway rebooting takes 4.62 seconds. The computational time for solving the CP problem increases from 0.45 seconds to 1.37 seconds as the network scales from 4,000 users (4 gateways) to 12,000 users (12 gateways). In scenarios with multiple coexisting networks, each network independently solves the CP problem for its own users in parallel. Although the spectrum sharing mechanism in *AlphaWAN* introduces additional 0.17–0.28 seconds spending on operator-to-Master communications, the overall capacity upgrade for 2 to 4 coexisting networks completes within 6 seconds, as shown in Figure 17b. This indicates that the capacity upgrade operation in *AlphaWAN* results in a system suspension

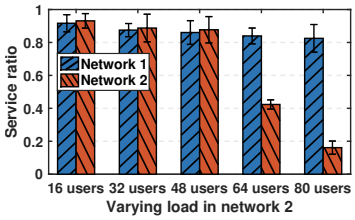

Figure 15: Fairness among coexisting networks under varying user loads.

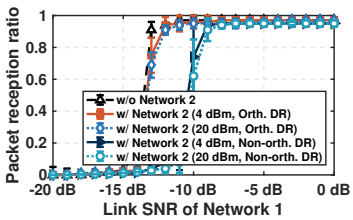

Figure 16: Impact of spectrum sharing on packet reception under different coexisting conditions.

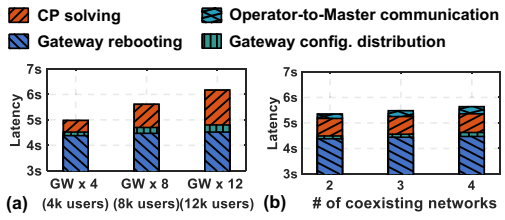

Figure 17: Latency of capacity upgrade with *AlphaWAN*: (a) a single network at different scales, (b) coexisting networks.

of less than 10 seconds. In practice, such upgrades can be scheduled during idle or designated maintenance periods, minimizing disruption to the normal system operation.

6 RELATED WORK

LoRaWAN In Status Quo. Pioneer measurement studies have been devoted to LoRaWAN to understand its performance and limitations. [33, 38] investigate LoRaWAN concurrent transmissions and varying link performance with a campus-scale testbed deployment. CityWAN [50] performs measurement studies on LoRaWAN for smart-city applications. [24] investigates the connectivity performance of a public operational LoRaWAN. Ghena *et al.* [14] identify capacity limitation and multi-network coexistence as two key challenges hindering the practical adoption of LoRaWAN and surveys techniques to address the challenges. To the best of our knowledge, we are the first to discover the decoder contention problem as a key bottleneck in large-scale LoRaWANs and present solutions to the problem.

Recent studies address architecture optimization for LoRaWAN. Joltik [57] designs universal sketches in LPWAN backhaul networks for efficient data analytics. HyperLoRa [16] integrates blockchain and edge computing into the LoRaWAN backhaul. QuAiL [11] focuses on sensory data processing in the LPWAN backhaul. [4, 34, 45] apply SDR-based C-RAN architecture to LoRaWAN, aiming to offload PHY layer signal processing from gateways to the cloud. In contrast, *AlphaWAN* does not change the current LoRaWAN architecture but exploits existing devices and network operation strategies to enhance LoRaWAN capacity.

Capacity enhancement. Many research efforts have been devoted to capacity enhancement for cellular networks [31, 39] and WLAN [37, 40]. For LoRaWAN, existing studies [5–7, 9, 30, 32, 46, 49, 51, 52, 55, 56, 59] mainly focus on PHY layer signal processing techniques to resolve packet collisions and support more concurrent transmissions over the same channel. Another category [12, 27, 28] studies MAC layer strategies and protocols for LoRaWAN to reduce packet collisions. XGate [58] proposes a new gateway design for LoRaWAN to enable packet reception across a broader spectrum. NELoRa [29], Falcon [48], and XCopy [54] improve gateway capability for receiving weak packets. However, none of these studies addresses the decoder contention problem.

Channel planning and optimization. Frequency planning and spectrum management have been studied in conventional cellular and Wi-Fi networks, but very few works have been proposed for LoRaWAN. Frequency planning in cellular networks mainly focuses on allocating frequency to various regions with the goal of minimizing inter-area interference and maximizing spectrum efficiency [3, 47]. For instance, [8] optimally assigns orthogonal sub-bands to neighboring gateways to enable bi-directional communication. [13, 15] classify the available spectrum into two groups: one serves edge regions in an orthogonal manner, the other serves all regions in a shared manner. Companies such as Federated Wireless [53] and Huawei [17] have introduced specialized mechanisms for Wi-Fi to facilitate spectrum management and sharing across coexisting networks. However, existing solutions designed for cellular and Wi-Fi are not directly applicable to LoRaWAN, due to significant differences in network architecture and operational characteristics. For LoRaWAN, Chime [10] optimizes reliability and power consumption by selecting the best frequency for each link. However, the capacity enhancement of LoRaWAN is not addressed.

7 CONCLUSION

This paper presents a systematic study of the practical capacity gaps observed in real-world operational LoRaWANs, identifying bottlenecks, unveiling the root causes, and providing comprehensive solutions. We delve into the underlying LoRaWAN system to uncover key performance bottlenecks inherent to resource-constrained gateways, shared ISM bands, and suboptimal network operations, which we conclude as the decoder contention problem. We propose strategies to address this problem from multiple design spaces, investigating their feasibility and practicality for LoRaWANs, and share our insights and lessons learned. We develop *AlphaWAN*, which integrates the proposed strategies into the LoRaWAN stack to effectively manage intra- and inter-network contentions for capacity enhancement. Extensive evaluations demonstrate that *AlphaWAN* can effectively improve the practical capacity of LoRaWANs. *AlphaWAN* can be directly deployed into today's operational LoRaWANs to meet the growing demands of global IoT applications.

ACKNOWLEDGEMENT

We sincerely thank our shepherd – Prof. Kyle Jamieson and the anonymous reviewers for their constructive comments and invaluable suggestions that helped improve this paper. This paper is supported by Hong Kong GRF under Grant No. 15218022, 15231424, 15211924, and 15216325. Xianjin Xia and Yuanqing Zheng are the corresponding authors.

REFERENCES

[1] 2024. *ChirpStack, Open Source LoRaWAN Network Server.* Retrieved June 15, 2024 from https://www.chirpstack.io/docs/

[2] LoRa Alliance. 2017. *LoRaWAN 1.1 Specification.* LoRa Alliance. https://lora-alliance.org/resource_hub/lorawan-specification-v1-1/.

[3] Ehsan Aryafar, Omer Gurewitz, and Edward W Knightly. 2008. Distance-1 constrained channel assignment in single radio wireless mesh networks. In *IEEE INFOCOM 2008-The 27th Conference on Computer Communications.* IEEE, 762–770.

[4] Artur Balanuta, Nuno Pereira, Swarun Kumar, and Anthony Rowe. 2020. A cloud-optimized link layer for low-power wide-area networks. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services* (Toronto, Ontario, Canada) *(MobiSys '20).* Association for Computing Machinery, New York, NY, USA, 247–259. https://doi.org/10.1145/3386901.3388915

[5] Qian Chen and Jiliang Wang. 2023. AlignTrack: Push the SNR Limit of LoRa Collision Decoding. *IEEE/ACM Transactions on Networking* 31, 5 (2023), 2070–2085. https://doi.org/10.1109/TNET.2023.3235041

[6] Weiwei Chen, Xianjin Xia, Shuai Wang, Tian He, Shuai Wang, Gang Liu, and Caishi Huang. 2025. Enabling Large Scale LoRa Parallel Decoding With High-Dimensional and High-Accuracy Features. *IEEE Transactions on Mobile Computing* 24, 5 (2025), 3520–3536. https://doi.org/10.1109/TMC.2024.3517343

[7] Weiwei Chen, Jiefeng Zhang, Xianjin Xia, Shuai Wang, Shuai Wang, and Tian He. 2024. Hitting the Sweet Spot: An SF-any Coding Paradigm for Empowering City-Wide LoRa Communications. In *2024 23rd ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN).* 225–236. https://doi.org/10.1109/IPSN61024.2024.00023

[8] P. Dutta, S. Jaiswal, D. Panigrahi, and R. Rastogi. 2008. A New Channel Assignment Mechanism for Rural Wireless Mesh Networks. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications.* 2261–2269. https://doi.org/10.1109/INFOCOM.2008.294

[9] R. Eletreby, D. Zhang, S. Kumar, and O. Yagan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM'17).* 309–321.

[10] Akshay Gadre, Revathy Narayanan, Anh Luong, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2020. Frequency Configuration for {Low-Power}{Wide-Area} Networks in a Heartbeat. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20).* 339–352.

[11] Akshay Gadre, Fan Yi, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2020. Quick (and dirty) aggregate queries on low-power WANs. In *2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN).* IEEE, 277–288.

[12] Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan, and Mo Li. 2020. *LMAC: Efficient Carrier-Sense Multiple Access for LoRa.* Association for Computing Machinery, New York, NY, USA.

[13] Rizwan Ghaffar and Raymond Knopp. 2010. Fractional frequency reuse and interference suppression for OFDMA networks. In *8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks.* IEEE, 273–277.

[14] 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* (Los Cabos, Mexico) *(MobiCom '19).* Association for Computing Machinery, New York, NY, USA, Article 43, 12 pages. https://doi.org/10.1145/3300061.3345444

[15] Naveed UL Hassan and Mohamad Assaad. 2009. Optimal fractional frequency reuse (FFR) and resource allocation in multiuser OFDMA system. In *2009 international conference on information and communication technologies.* IEEE, 88–92.

[16] Lu Hou, Kan Zheng, Zhiming Liu, Xiaojun Xu, and Tao Wu. 2021. Design and Prototype Implementation of a Blockchain-Enabled LoRa System With Edge Computing. *IEEE Internet of Things Journal* 8, 4 (2021), 2419–2430. https://doi.org/10.1109/JIOT.2020.3027713

[17] Huawei. 2025. *Huawei Redefines Mobile Access Networks with CloudAIR.* Retrieved June 24, 2025 from https://carrier.huawei.com/en/trends-and-insights/cloudran/huawei-redefines-mobile-access-networks-with-cloudair

[18] The Things Industries. 2024. *Where companies scale their LoRaWAN® solutions. Unlock the full potential of LoRaWAN with our secured, scalable and fully featured Network Server.* Retrieved November 20, 2024 from https://www.thethingsindustries.com/

[19] INTUZ. 2022. *LPWAN vs. LoRaWAN: The Better Technology For IoT Device Connectivity.* Retrieved Dec 18, 2024 from https://www.intuz.com/blog/lpwan-vs-lorawan-the-better-technology-for-iot-device

[20] IoT Analytics. 2018. *LPWAN technologies: How cellular MNOs are placing their bets.* Retrieved Oct 22, 2024 from https://iot-analytics.com/lpwan-technologies-cellular-mnos/

[21] IoT Analytics. 2024. *LPWAN market 2024: Licensed technologies boost their share among global 1.3 billion connections as LoRa leads outside China.* Retrieved Dec 18, 2024 from https://iot-analytics.com/lpwan-market/

[22] IoT for All. 2022. *LPWAN Market: Top trends propelling the industry demand through 2027.* Retrieved Dec 18, 2024 from https://www.iotforall.com/news/lpwan-market--top-trends-propelling-the-industry-demand-through-2027

[23] Jacob Schindler. 2024. *Four takeaways from the latest LPWAN market data.* Retrieved Dec 18, 2024 from https://www.sisvel.com/insights/four-takeaways-from-the-latest-lpwan-market-data/

[24] Dhananjay Jagtap, Alex Yen, Huanlei Wu, Aaron Schulman, and Pat Pannuto. 2021. Federated infrastructure: usage, patterns, and insights from "the people's network". In *Proceedings of the 21st ACM Internet Measurement Conference* (Virtual Event) *(IMC '21).* Association for Computing Machinery, New York, NY, USA, 22–36. https://doi.org/10.1145/3487552.3487846

[25] Richard M. Karp. 1972. Reducibility among Combinatorial Problems, Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (Eds.). Springer US, Boston, MA, 85–103. https://doi.org/10.1007/978-1-4684-2001-2_9

[26] Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. *Multidimensional Knapsack Problems.* Springer Berlin Heidelberg, Berlin, Heidelberg, 235–283. https://doi.org/10.1007/978-3-540-24777-7_9

[27] Nikolaos Kouvelas, R Venkatesha Prasad, Niloofar Yazdani, and Daniel E. Lucani. 2021. np-CECADA: Enhancing Ubiquitous Connectivity of LoRa Networks. In *2021 IEEE 18th International Conference on Mobile Ad Hoc and Smart Systems (MASS).* 374–382. https://doi.org/10.1109/MASS52906.2021.00054

[28] Nikolaos Kouvelas, Vijay S Rao, R. Venkatesha Prasad, Gauri Tawde, and Koen Langendoen. 2020. p-CARMA: Politely Scaling LoRaWAN. In *Proceedings of the 2020 International Conference on Embedded Wireless Systems and Networks* (Lyon, France) *(EWSN '20).* Junction Publishing, USA, 25–36.

[29] Chenning Li, Hanqing Guo, Shuai Tong, Xiao Zeng, Zhichao Cao, Mi Zhang, Qiben Yan, Li Xiao, Jiliang Wang, and Yunhao Liu. 2021. NELoRa: Towards Ultra-low SNR LoRa Communication with Neural-enhanced Demodulation. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems* (Coimbra, Portugal) *(SenSys '21).* Association for Computing Machinery, New York, NY, USA, 56–68. https://doi.org/10.1145/3485730.3485928

[30] Chenning Li, Xiuzhen Guo, Longfei Shangguan, Zhichao Cao, and Kyle Jamieson. 2022. CurvingLoRa to Boost LoRa Network Throughput via Concurrent Transmission. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22).* USENIX Association, Renton, WA, 879–895. https://www.usenix.org/conference/nsdi22/presentation/li-chenning

[31] Guoqing Li and Hui Liu. 2006. Downlink radio resource allocation for multi-cell OFDMA system. *IEEE Transactions on wireless communications* 5, 12 (2006), 3451–3459.

[32] Ruonan Li, Ziyue Zhang, Xianjin Xia, Ningning Hou, Wenchang Chai, Shiming Yu, Yuanqing Zheng, and Tao Gu. 2025. From Interference Mitigation to Toleration: Pathway to Practical Spatial Reuse in LPWANs. In *Proceedings of the 31th Annual International Conference on Mobile Computing and Networking* (Hong Kong, China) *(ACM MobiCom '25).* Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3680207.3723483

[33] 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 Trans. Sen. Netw.* 15, 2, Article 16 (feb 2019), 35 pages. https://doi.org/10.1145/3293534

[34] Jun Liu, Weitao Xu, Sanjay Jha, and Wen Hu. 2020. Nephalai: towards LPWAN C-RAN with physical layer compression. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) *(MobiCom '20).* Association for Computing Machinery, New York, NY, USA, Article 36, 12 pages. https://doi.org/10.1145/3372224.3419193

[35] LoRa Alliance. 2021. *LoRaWAN Formally Recognized as ITU International Standard for Low Power Wide Area Networking.* Retrieved Dec 18, 2024 from https://lora-alliance.org/lora-alliance-press-release/lorawan-formally-recognized-as-itu-international-standard-for-low-power-wide-area-networking/

[36] Netmore Senet. 2024. *Connecting the IoT Revolution.* Retrieved Jan 21, 2025 from https://senetco.com/

[37] Shuwei Qiu, Xiaowen Chu, Yiu-Wing Leung, and Joseph Kee Yin Ng. 2020. Joint Access Point Placement and Power-Channel-Resource-Unit Assignment for 802.11ax-Based Dense WiFi with QoS Requirements. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications.* 2569–2578. https://doi.org/10.1109/INFOCOM41043.2020.9155490

[38] Michael Rademacher, Hendrik Linka, Thorsten Horstmann, and Martin Henze. 2021. Path Loss in Urban LoRa Networks: A Large-Scale Measurement Study. In *2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall).* 1–6. https://doi.org/10.1109/VTC2021-Fall52928.2021.9625531

[39] Mahmudur Rahman and Halim Yanikomeroglu. 2010. Enhancing cell-edge performance: a downlink dynamic interference avoidance scheme with inter-cell coordination. *IEEE Transactions on Wireless Communications* 9, 4 (2010), 1414–1425.

[40] Hariharan Shankar Rahul, Swarun Kumar, and Dina Katabi. 2012. JMB: scaling wireless capacity with user demands. *SIGCOMM Comput. Commun. Rev.* 42, 4 (Aug. 2012), 235–246. https://doi.org/10.1145/2377677.2377722

[41] RAKwireless. [n. d.]. RAK7289 V2 Datasheet. https://docs.rakwireless.com/Product-Categories/WisGate/RAK7289-V2/Datasheet/, year = 2024, note = Accessed: 2024-08-14.
[42] RAKwireless. 2024. *Bringing WisGate Edge to Life (an OS that fits your needs).* Retrieved October 8, 2024 from https://www.rakwireless.com/en-us/products/wisgate-os2
[43] RAKWireless. 2025. *860-930MHz 12dBi Directional Antenna Datasheet.* Retrieved Jan 3, 2025 from https://docs.rakwireless.com/Product-Categories/Accessories/RAKARP01/Overview/#product-description
[44] Semtech. 2024. *SX1302 Datasheet.* Retrieved May 2, 2024 from https://www.semtech.com/products/wireless-rf/lora-core/sx1302
[45] Muhammad Osama Shahid, Daniel Koch, Jayaram Raghuram, Bhuvana Krishnaswamy, Krishna Chintalapudi, and Suman Banerjee. 2024. Cloud-LoRa: Enabling Cloud Radio Access LoRa Networks Using Reinforcement Learning Based Bandwidth-Adaptive Compression. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24).* USENIX Association, Santa Clara, CA, 1959–1976. https://www.usenix.org/conference/nsdi24/presentation/shahid
[46] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent Interference Cancellation: Decoding Multi-Packet Collisions in LoRa *(SIGCOMM '21).* Association for Computing Machinery, New York, NY, USA.
[47] Jungmin So and Nitin H Vaidya. 2004. Multi-channel MAC for ad hoc networks: Handling multi-channel hidden terminals using a single transceiver. In *Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing.* 222–233.
[48] Shuai Tong, Zilin Shen, Yunhao Liu, and Jiliang Wang. 2021. Combating link dynamics for reliable lora connection in urban settings. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* (New Orleans, Louisiana) *(MobiCom '21).* Association for Computing Machinery, New York, NY, USA, 642–655. https://doi.org/10.1145/3447993.3483250
[49] Shuai Tong, Jiliang Wang, and Yunhao Liu. 2020. Combating packet collisions using non-stationary signal scaling in LPWANs. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services.* 234–246.
[50] Shuai Tong, Jiliang Wang, Jing Yang, Yunhao Liu, and Jun Zhang. 2024. Citywide LoRa Network Deployment and Operation: Measurements, Analysis, and Implications. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems* (Istanbul, Turkiye) *(SenSys '23).* Association for Computing Machinery, New York, NY, USA, 362–375. https://doi.org/10.1145/3625687.3625796
[51] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. 2020. CoLoRa: Enabling Multi-Packet Reception in LoRa. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications.* 2303–2311.
[52] Jiliang Wang, Shuai Tong, Zhenqiang Xu, and Pengjin Xie. 2024. Real-Time Concurrent LoRa Transmissions Based on Peak Tracking. *IEEE Transactions on Mobile Computing* 23, 10 (2024), 9582–9594. https://doi.org/10.1109/TMC.2024.3365797
[53] Federated Wireless. 2025. *Empowering Connectivity with Shared Spectrum.* Retrieved June 23, 2025 from https://www.federatedwireless.com/
[54] Xianjin Xia, Qianwu Chen, Ningning Hou, Yuanqing Zheng, and Mo Li. 2023. XCopy: Boosting Weak Links for Reliable LoRa Communication. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking* (Madrid, Spain) *(ACM MobiCom '23).* Association for Computing Machinery, New York, NY, USA, Article 14, 15 pages. https://doi.org/10.1145/3570361.3592516
[55] Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2021. PCube: Scaling LoRa Concurrent Transmissions with Reception Diversities. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom'21).* ACM.
[56] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. FTrack: Parallel Decoding for LoRa Transmissions. *IEEE/ACM Transactions on Networking* 28, 6 (2020), 2573–2586. https://doi.org/10.1109/TNET.2020.3018020
[57] Mingran Yang, Junbo Zhang, Akshay Gadre, Zaoxing Liu, Swarun Kumar, and Vyas Sekar. 2020. Joltik: enabling energy-efficient" future-proof" analytics on low-power wide-area networks. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking.* 1–14.
[58] Shiming Yu, Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2024. Revolutionizing LoRa Gateway with XGate: Scalable Concurrent Transmission across Massive Logical Channels. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) *(ACM MobiCom '24).* Association for Computing Machinery, New York, NY, USA, 482–496. https://doi.org/10.1145/3636534.3649375
[59] Shiming Yu, Xianjin Xia, Ziyue Zhang, Ningning Hou, and Yuanqing Zheng. 2025. FDLoRa: Scaling Downlink Concurrent Transmissions With Full-Duplex LoRa Gateways. *IEEE Transactions on Mobile Computing* (2025), 1–14. https://doi.org/10.1109/TMC.2025.3572130
[60] ZENNER Connect. 2024. *Our LoRaWAN Network Availability.* Retrieved Dec 18, 2024 from https://zenner-connect.com/en/with-lorawan-into-the-internet-of-things/#availability

APPENDICES

Appendices are supporting material that has not been peer-reviewed.

A STATUS OF GLOBAL LORAWAN OPERATIONS

Table 2 lists some leading LoRaWAN service providers across various regions worldwide. With the number of IoT devices and applications continuously growing, more organizations have been embracing and benefiting from LoRaWAN with the help of professional support. Notably, as the authorized spectrum for LoRaWAN is limited to less than 6.5 MHz in over 70% of countries and regions (see Figure 18), achieving high-efficiency spectrum utilization becomes particularly critical for practical LoRaWAN operations.

LoRaWAN operator	Operational regions	Operation mode	# Gateways	# End nodes	User growth rate
The Things Industries	Global	Public	50K	1M	50%
Netmore Senet	EU/US/AU	Public	20K	2.3M	251%
Actility	EU/US/AS	Public	40K	4M	75%
ZENNER Connect	EU/US	Public	110K	8.9M	78%

Table 2: Status of some commercial operational LoRaWANs.

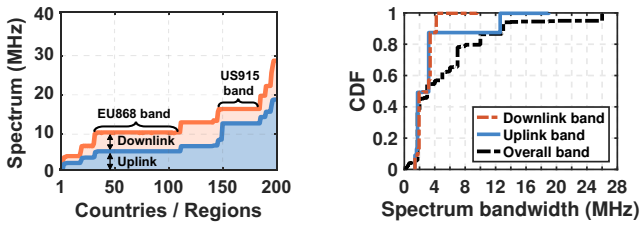

Figure 18: LoRaWAN spectrum across countries and regions.

B HOW IS LORAWAN OPERATED?

The operational strategies adopted by LoRaWAN differ drastically from those of traditional cellular networks, as summarized in Table 3. In cellular networks, each user is associated with a specific cell tower and allocated dedicated spectrum resources. This one-to-one association between users and resources allows cellular operators to flexibly scale capacity with the available resources. In contrast, LoRaWAN does not associate users with dedicated gateways or spectrum resources. Instead, the LoRaWAN spectrum is shared by all users across operators and accessed through competition. Specifically, a user's packet is received and relayed to the network server by all LoRaWAN gateways within the coverage area. LoRaWAN adopts such a strategy to make use of gateway redundancy for packet reliability. However, this can also lead to a single user occupying multiple gateways' resources, while many others may remain unable to use the network. Despite the differences in operational strategies, many LoRaWANs still follow the methods of cellular networks to allocate frequency resources and operate network gateways and user devices. This approach has been demonstrated to be unsuitable for LoRaWAN and can result in sub-optimal network capacity (refer to § 2.2).

LoRaWAN gateways and user devices operate in unlicensed ISM bands, with frequencies varying across regions (*e.g.*, 902-928 MHz

	Cellular network	LoRaWAN
User association	Associated with one cell tower	Not associated with any gateways
User-gateway connection	One user to one tower (One-to-One)	One user to multiple gateways (One-to-Many)
Operational spectrum	Licensed spectrum	Unlicensed ISM bands
Spectrum use strategy	Dedicated frequency resource; Allocated per-user basis	Shared spectrum; Contention based access

Table 3: Operational strategy differences from cellular networks and LoRaWANs.

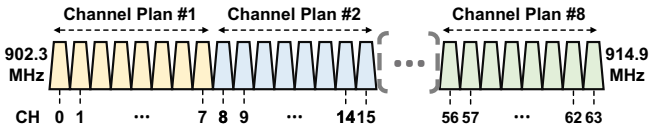

Figure 19: Channel Plans of Standard LoRaWAN.

in the US, 863-870 MHz in Europe). The LoRaWAN standard offers dynamic and fixed channel plans to operate the devices. LoRaWANs in the US915 band adopt fixed channel plans, where the majority of channel frequencies are statically defined. Operators select one of the predefined channel plans to configure a gateway for operation. In contrast, the dynamic channel plan (*e.g.*, the EU868 band) allows operators to define and alter channels dynamically for their networks. Meanwhile, several channels with fixed frequencies are reserved (*e.g.*, for joining requests and beacon broadcasting) and must be supported in all LoRaWANs.

Figure 19 demonstrates the channel plans in the US915 band. The LoRaWAN channels are numbered CH 0 from the minimum frequency to the maximum (CH 63). Starting with CH 0, every eight channels form a group termed a *channel plan*. For instance, channel plan #1 includes CH 0, CH 1, ···, CH 7, and CHs 8∼15 belong to channel plan #2. The frequency settings of a channel plan are used to configure the concurrent reception channels of a LoRaWAN gateway.

C RECEPTION PIPELINE AND RADIO ARCHITECTURE OF LORAWAN GATEWAY

The experimental investigations in § 3.1 uncover the underlying reception mechanism of a COTS gateway summarized below:

The reception pipeline of a LoRa packet consists of two phases: packet detection and decoding, as illustrated in Figure 20a. Packet detection involves identifying the preamble of a LoRa packet. Upon detecting a LoRa preamble, the gateway radio locks on the packet and initiates the decoding process, which includes detecting frame timing, verifying sync words, and decoding the payload. The RF radio of COTS gateways employs a parallel reception architecture, as shown in Figure 20b. The radio comprises multiple Rx chains, each capable of receiving different frequencies. A packet detector is applied to each Rx chain to detect incoming packets with varying data rates over a frequency. Once a packet is detected, it is passed to a dispatcher that allocates decoder resources from a pool to process the packet. The dispatcher handles packets from all Rx chains in a First-Come-First-Served (FCFS) manner based on the lock-on time of the packets. If all decoders in the pool are occupied, the dispatcher drops subsequent packets until any decoders become available. A

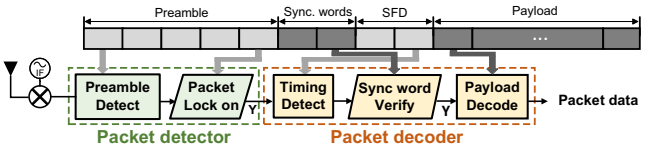

(a) Packet reception pipeline.

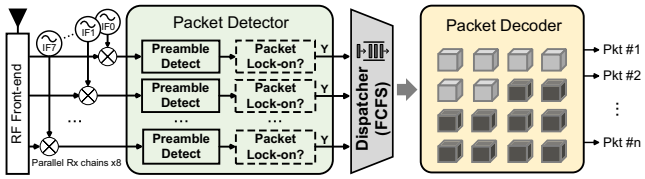

(b) The radio architecture and workflow of COTS gateway.

Figure 20: The workflow of concurrent packet reception of COTS LoRaWAN gateway: packet decoders are dispatched in a First-Come-First-Served manner.

Manufacturer	Product Model	Chipset	Rx Spectrum	# Rx chains	# Decoders	Theory Capacity	Practical Capacity
Dragino	LPS8N LPS8V2	SX1302	1.6MHz	8+1	16	**54**	**16**
RAKwireless	RAK7246G	SX1308	1.6MHz	8+1	8	**54**	**8**
	RAK7268CV2	SX1302	1.6MHz	8+1	16	**54**	**16**
	RAK7289CV2	SX1303	3.2MHz	16+2	32	**108**	**32**
Kerlink	Wirnet IBTS Wirnet iFemtoCell	SX1301	1.6MHz	8+1	8	**54**	**8**

Table 4: The maximum number of concurrent users supported by different commercial gateways.

decoder is released and returned to the pool after finishing the decoding of a packet.

We have empirically examined several mainstream COTS gateway products and summarized their capacity data in Table 4. None of these gateways have sufficient decoders to fully support the theoretical capacity of their Rx spectrum.

D SIMULATION STUDY ON USER EXPANSION OVER LONG-TERM OPERATIONS

This experiment evaluates the effectiveness of *AlphaWAN* in handling user expansions during the long-term operations of LoRaWAN infrastructure. We use a trace-driven method to simulate a network with 10 gateways running for one year (53 weeks). Over 100,000 packet traces are collected from 500 sites in our testbed, with packet SNRs ranging from -15 dB to 5 dB. We use the collected traces to synthesize node traffic across different frequency channels and simulate the communications of massive IoT nodes. Each node randomly sends packets at a 1% duty cycle following the LoRaWAN regulation [2]. The network uses a 4.8 MHz spectrum and starts with 1,180 users. New users connect to the network regularly (around 150 users join every week). Connected users send packets with random frequency channels and select data rates based on the link's SNR profiles. We plot the average packet reception ratio (PRR) of the entire network every week to examine the performance of *AlphaWAN*

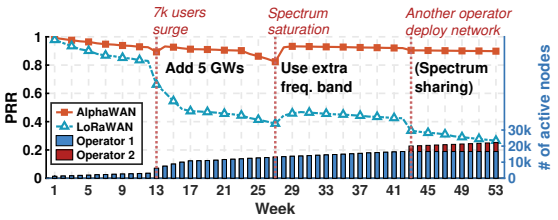

Figure 21: Simulation study of *AlphaWAN* on handling LoRaWAN user expansions over long-term operation.

as presented in Figure 21. The PRRs of standard LoRaWAN are also evaluated for baseline comparisons.

From week 1 to week 12, the number of connected users increases from 1,180 to 3,090. The PRRs of both strategies decrease slightly as more users join the network, with a faster PRR drop observed for standard LoRaWAN.

In week 13, we add 7,000 users to the network to simulate the deployment of a new IoT application. To respond to this event, *AlphaWAN* adds five more gateways to expand network capacity. For fairness, five gateways have also been added to the standard LoRaWAN. We observe that *AlphaWAN* successfully upgrades network capacity to accommodate the 7,000 new users with PRRs greater than 90%, while standard LoRaWAN fails to provide additional capacity for the users, experiencing a significant PRR drop.

By week 27, the overall number of network users increases to 13,190, and the spectrum saturates. We expand the network's operating spectrum with an additional 1.6 MHz (*i.e.*, eight more channels). This pulls up the PRRs of *AlphaWAN* from 80% to higher than 90%. A slight PRR increase is also observed for standard LoRaWAN, but the PRRs remain below 50% due to inefficient spectrum utilization. In contrast, *AlphaWAN* effectively uses the spectrum to deliver PRRs greater than 90%.

In week 43, another operator deploys its LoRaWAN network in the area, with five gateways and 3,430 connected users operating in the same spectrum. The PRRs of standard LoRaWAN drop rapidly due to increased contentions from the coexisting network of the new operator. In contrast, *AlphaWAN* handles this event with its spectrum-sharing mechanism and delivers high PRRs, even in week 53 when the number of users reaches 22,180.