Adapting LoRa Ground Stations for Low-latency Imaging and Inference from LoRa-enabled CubeSats

AKSHAY GADRE, Electrical and Computer Engineering, University of Washington, Seattle, United States
ZACHARY MACHESTER, Carnegie Mellon University, Pittsburgh, United States
SWARUN KUMAR, Carnegie Mellon University, Pittsburgh, United States

Recent years have seen the rapid deployment of low-cost CubeSats in low-Earth orbit, many of which experience significant latency (several hours) from the time information is gathered to the time it is communicated to the ground. This is primarily due to the limited availability of ground infrastructure that is bulky to deploy and expensive to rent. This article explores the opportunity in leveraging the extensive terrestrial LoRa infrastructure as a solution. However, the limited bandwidth and large amount of Doppler on CubeSats precludes these LoRa links to communicate rich satellite Earth images—instead, the CubeSats can at best send short messages. This article details our experience in designing LoRa-based satellite ground infrastructure that requires software-only modifications to receive packets from LoRa-enabled CubeSats recently launched by our team. We present Vista, a communication system that adapts encoding onboard the CubeSat and decoding configuration on commercial LoRa ground stations to allow images to be communicated. We perform a detailed evaluation of Vista by leveraging wireless channel measurements from a recent CubeSat (2021), and show that Vista can achieve 55.55% lower latency in retrieving data with 12.02 dB improvement in packet retrieval in the presence of terrestrial interference. We then evaluate Vista on a case study on land-use classification over images transmitted over the CubeSat link to further demonstrate a 4.56 dB improvement in image PSNR and 1.38× increase in classification accuracy over baseline approaches.

CCS Concepts: • **Networks** → **Sensor networks**; *Very long-range networks*; • **Hardware** → **Sensor applications and deployments**; • **Computer systems organization** → **Sensor networks;**

Additional Key Words and Phrases: Satellite imaging, Doppler compensation, automated error recovery, channel-aware computing

ACM Reference Format:
Akshay Gadre, Zachary Machester, and Swarun Kumar. 2024. Adapting LoRa Ground Stations for Low-latency Imaging and Inference from LoRa-enabled CubeSats. *ACM Trans. Sensor Netw.* 20, 5, Article 102 (July 2024), 30 pages. https://doi.org/10.1145/3675170

1 Introduction

The massive reduction in costs of launching satellite payloads [22] has led to a proliferation of CubeSats in Low Earth Orbit and truly democratized access to space for small players—researchers,

This work was supported by grants from NSF (Grants # 2106921, 2030154, 2007786, 1942902, 2111751), ONR, AFRETEC, MFI, CISCO, Safety21 and CyLab-Enterprise.
Authors' Contact Information: Akshay Gadre, Electrical and Computer Engineering, University of Washington, Seattle, Washington, United States; e-mail: gadre@uw.edu; Zachary Machester, Carnegie Mellon University, Pittsburgh, Pennsylvania, United States; e-mail: zacm@cmu.edu; Swarun Kumar, Carnegie Mellon University, Pittsburgh, Pennsylvania, United States; e-mail: swarun@cmu.edu.

ACM 1550-4859/2024/07-ART102
https://doi.org/10.1145/3675170

students, and satellite enthusiasts. Today, such CubeSats target varied geo-sensing and imaging applications and number in the thousands [37]. However, most CubeSats launched today fail to answer a simple question at low-latency: "Is my satellite online right now, and if so, what is it imaging?". Most current CubeSats rely on existing wireless technologies, developed for larger satellites, for downlink communication. However, ground stations for these technologies are extremely expensive to rent/install for small teams—with most teams able to downlink from CubeSats quite rarely—often once per few hours [86], i.e., whenever a ground station is in range. Imagine a time-critical event that occurs during this long window where the satellite is simply inaccessible. For instance, consider an interesting Earth image that needs to be sent downlink. Such images could be time-sensitive, capturing on-Earth emergencies such as flooding or weather events around specific locations on the ground observed by the CubeSat. Today, if the CubeSat happens to be in an area without this dedicated ground station infrastructure, it may take several hours to received the data downlink. In this article, we propose to break this infrastructure coverage bottleneck for low-latency access to satellites by adapting existing ground infrastructure that is widespread around the world—specifically, commercial LoRa base stations. Specifically, we propose and demonstrate how LoRa can act as a ***quick and dirty* complementary** modality for conventional CubeSat deployments to push time-sensitive information to the ground anywhere in the world. Specifically, we believe LoRa can be deployed (Section 5) on the next-generation of earth-observation satellites to retrieve a quick rough estimate of observed terrain to reduce the first-contact latency[1] of downlink communication.

The reasons behind why LoRa is an attractive wireless technology to enable this capability to communicate with CubeSats in a timely manner is three-fold: (i) There is a vast existing ground infrastructure that continuously listens for LoRa packets that can be used to receive packets around the globe. While these are not optimized for satellite communication, commercial LoRa base stations are capable of receiving communication signals from satellites, albeit at a lower data-rate. (ii) It's chirp-spread-spectrum narrowband modulation is immune to much of the on-ground communication in these bands. Furthermore, it has been shown that other technologies also are affected minimally due to LoRa traffic at these frequencies [39, 87]. (iii) It accommodates a large number of transmitters simultaneously. Our motivation study in Section 7.1 demonstrates the clear latency benefits of leveraging existing LoRa ground station infrastructure.

Yet, there are two key challenges in communicating data from LoRa-equipped CubeSats to existing ground stations—the limited bandwidth of LoRa and a highly lossy wireless medium. These stem from a mismatch between the typical LoRa client (a mostly static terrestrial device deployed a few kilometers away) and a CubeSat (a device in space moving at seven kilometers per second and hundreds or thousands of kilometers away). In fact, theoretical analysis of the impact of these Doppler effects on LoRa at high frequencies [46] has shown significant amount of packet drops and error in detection. Indeed, that is why most of the recently launched LoRa CubeSats [6, 7, 15, 19] operate at lower frequencies (137 MHz and 407–450 MHz) and leverage specialized receivers instead of utilizing the existing 915 MHz ISM band ground infrastructure.

In this article, we describe our experience in adapting commodity ground LoRa infrastructure to effectively communicate with LoRa-enabled CubeSats in low-earth orbit. We present Vista (see Figure 1), a CubeSat communication system that uses conventional LoRa modulation to directly deliver data downlink with high efficacy, by overcoming the challenges of limited bandwidth and infrastructure compatibility. We study the performance of LoRa communication in the ISM band

[1]Note, that first-contact latency is the the delay to reach the next zone of downlink connectivity, while traditional latency deals with the overall downlink of all captured data instead of just a snapshot. In this article, we primarily use latency to refer to this ***first-contact latency***.

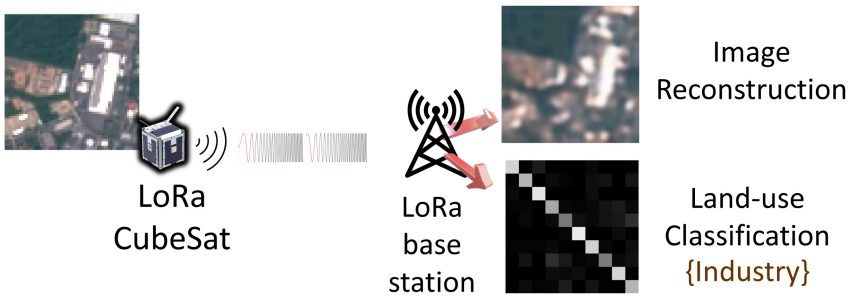

Fig. 1. Vista enables low-latency practical imaging and inference applications for LoRa CubeSats.

by collecting data from a launch of a LoRa CubeSat (launched by our team in 2021). We then use wireless channel measurements from this CubeSat deployed in space to demonstrate that Vista can provide an effective 55.55% lower latency and 12.02 dB improvement in packet detection SINR. We further perform a case study on land-use classification over images transmitted over this link to show 4.56 dB improvement in image retrieval SNR and 1.38× improvement in classification tasks.

The rest of this article details Vista's approach to address these challenges, namely: (1) Enabling existing LoRa ground station infrastructure to detect and decode CubeSat LoRa signals. (2) Designing channel-aware solutions for encoding images and performing inference on these images sent over the CubeSat LoRa link.

Connecting LoRa CubeSats with ground infrastructure: Engineering a terrestrial LP-WAN technology to communicate at a range fifty times larger than its conventional range is a challenging task. There are two primary challenges (apart from larger attenuation) that existing ground stations need to resolve for receiving packets from LoRa-enabled CubeSats: *(1) Larger Doppler Shifts* (by an order of magnitude) and *(2) Co-existence with terrestrial LoRa clients*. At the ground stations, the massive Doppler shift renders the packets from CubeSats undecodable at pre-existing LoRa base stations. Vista mitigates this by leveraging the fact that existing base stations typically have several hardware modules for packet detection which are typically multiplexed on the same receive chain that operate at various different bandwidths—LoRa traditionally provides these as part of its data-rate adaptation pipeline. We demonstrate how a careful selection of these pre-existing correlation modules can enable terrestrial ground stations to detect these missed packets with a simple software update. Next, we need to ensure that terrestrial clients' packets do not overwhelm packet transmissions from the CubeSats. Vista achieves this by understanding an unique behavior of CubeSat packets—the Doppler frequency shifts experienced by these packets are governed by their trajectory. Thus, Vista develops CubeSat-specific **Doppler-division multiple access (DDMA)** mechanism for LoRa that significantly improve the SINR (signal-to-interference-noise ratio) of packet detection for CubeSat packets in presence of interference from terrestrial clients and other CubeSats. Section 3 describes both of these solutions.

Channel-Aware Image Encoding: Given this bandwidth-starved CubeSat link, it is critical to maximize the utility of every byte to communicate data downlink. We specifically focus on communicating image data for the purpose of performing inference tasks. A näive solution would be to use traditional compression algorithms such as JPEG with appropriate compression ratio. Yet, there exists an exciting opportunity in the satellite image context, where rich prior information is available in how these images are structured, e.g., image data from NOAA LEO weather satellites [11].

Furthermore, compression of image data can also be actively guided by the nature of machine-learning tasks performed on these images. Most importantly, one can also account for the nature of the wireless channel between CubeSats and the LoRa ground stations, where factors such as intra-packet Doppler shift, noise, and satellite trajectories play a key role in the nature of image bit errors one can expect. In other words, rather than building a channel, data and task-agnostic image compression scheme, Vista designs an image encoding that is informed by the general nature of wireless channels, the structure of satellite imagery, and the tasks performed on this data.

Unlike prior work in joint source-channel coding for wireless image transmission [34, 35, 38, 42, 55] which deals with general uniform error distributions, operating LoRa (a technology that encodes data in frequency shifts) in 915 MHz ISM band (large amount of Doppler) leads to a non-uniform distribution of errors across bits. Section 4 describes our design of a data-task-driven autoencoder which is optimized for such space-to-ground wireless channels. We train this autoencoder with real satellite image data, our satellite trajectory, as well as our own LoRa channel measurements. We further make this encoding resilient to information loss due to discretization and wireless impairments. We provide various design points in our autoencoder that balances the diversity of tasks that can be performed with the image encoding on the ground and the amount of data that needs to be communicated.

We study Vista by developing a communication stack driven by channel-traces from a LoRa-enabled CubeSat in the ISM band launched in 2021 [21]. We study and report the observed Doppler effects and improved ground-station latency of this LoRa satellite. We then collect long-term wireless channel traces measured from the satellite in orbit across an year. We implement, evaluate and emulate various flavors of Vista on these channel measurements to showcase the potential benefits in terms of latency, packet detection and inference accuracy, demonstrating:

- 55.55% reduction in latency with a coverage across 44.92% of the CubeSat orbit vs. 33.27% of existing infrastructure.
- Up to 15 dB improvement in SNR for packet detection in high Doppler scenarios and 12.02 dB improvement in interference resilience towards terrestrial clients.
- 4.56 dB PSNR improvement in images retrieved and 38% improvement in accuracy of ML tasks over images communicated using Vista.

Contributions: Our contributions include:

- A novel CubeSat wireless communication system that uses conventional LoRa modulation to enable high-quality image inference over bandwidth starved links.
- A system that makes software modifications to commercial LoRa ground gateways to detect highly Doppler shifted packets from LoRa CubeSats in presence of intereference.
- Detailed implementation and evaluation through wireless channel measurements from a commercial LoRa enabled satellite launched by our team.

2 Overview

Vista is a CubeSat communication system that uses LoRa modulation at the satellite and conventional ground infrastructure, making software-only modifications at both ends to improve delivery and processing of images. Figure 2 depicts Vista's contributions. The rest of the article discusses the two key components of Vista's design.

Adapting LoRa CubeSat—Ground Station Links: We show how we can redesign existing commercial off-the-shelf LoRa base stations to detect and decode CubeSat packets accounting for the extremely high Doppler shifts these packets experience. Further, we leverage the special nature

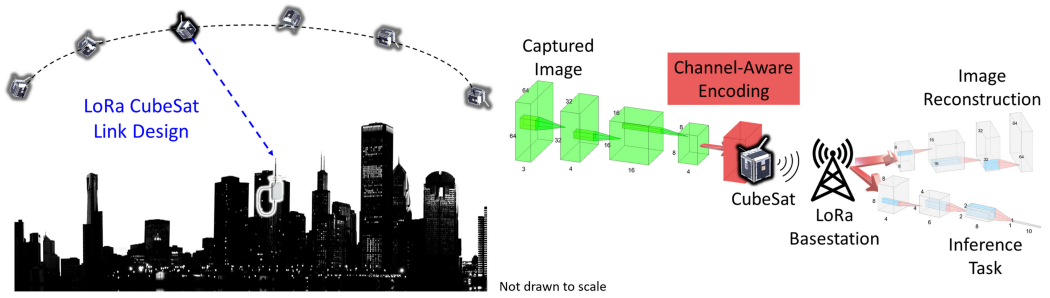

Fig. 2. Vista develops channel aware data-task driven encodings and Doppler-resilient packet detection to enable fine-grained imaging and low-latency inference from LoRa-enabled CubeSats using existing LoRa ground infrastructure.

of these satellite signals to develop trajectory-aware preambles for overcoming interference from other terrestrial LoRa sources. Section 3 describes our solution.

Channel-Aware CubeSat Image Encodings: At the satellite, Vista encodes data ensuring resilience to Doppler and other wireless impairments, while remaining cognizant of structure of the data and the applications ahead. Our case study evaluation of Vista decodes images and/or performs desired inference on the images, such as identifying land use patterns. Section 4 describes Vista's solutions at both the CubeSat and the ground infrastructure to enable this at low-latency.

3 Adapting LoRa Ground Stations

This section describes our solutions at the ground station for detecting LoRa packets being transmitted from CubeSats and improving SINR of detecting these packets in presence of terrestrial and CubeSat interference.

3.1 Detecting CubeSat Packets on LoRa Ground Station

Conventional LoRa ground infrastructure (operating at the 915 MHz ISM band) is not designed to receive LoRa signals from CubeSats. Furthermore, any approach that requires us to build custom ground infrastructure would negate the latency gains of leveraging existing infrastructure. Primarily, there are two key differences between the signals received at LoRa base stations from terrestrial sensors and those from the CubeSats. First, the satellites are traveling at speeds of several km/s (~7 km/s for our satellite), much larger than any terrestrial mobile LoRa clients. Second, is the fact that CubeSats operate in the low earth orbit (300–800 km altitude; 525 km for our CubeSat) which is a significantly larger distance. At the core, the base station needs to perform two tasks: (1) Detect and Decode the CubeSat packets, and (2) Know when the CubeSat is overhead.

Detecting and Decoding LoRa CubeSat packets: The attenuation of the transmitted signal at the base station in a good pass (~90^o elevation peak) is roughly 150 dB at an average. This leads to an average received power of −123 dBm significantly above the detection threshold of off-the-shelf LoRa base stations (−148 dBm). Fundamentally speaking, by leveraging the existing infrastructure (high gain antennas and LNAs), it is entirely feasible to detect and decode signals from the LoRa CubeSats. Indeed, prior deployments on hot-air balloons have demonstrated communication as far as 832 km using LoRa [83]. Thus, existing LoRa infrastructure with high-gain antennas can easily overcome its range bottleneck.

However, the massive amount of Doppler shifts in the CubeSat context present a difficult challenge. While LoRa was developed to be inherently resilient to terrestrial Doppler (a few Hz), the

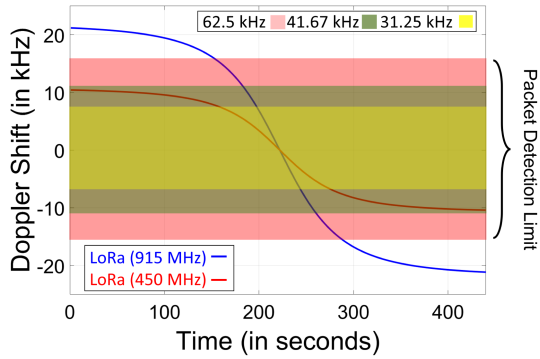

Fig. 3. Unlike CubeSats operating at lower frequencies (red), traditional ISM band LoRa base stations cannot detect LoRa CubeSat packets (blue) due to the large amount of Doppler frequency offset.

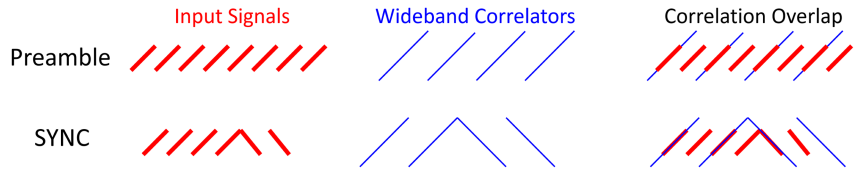

Fig. 4. Correlation of preambles and SYNC symbols with the already present wideband correlators in off-the-shelf LoRa base stations enables high Doppler packet detection and accurate Doppler estimation in Vista ground stations.

amount of Doppler seen from these CubeSats is on the order of several kHz (a significant proportion of the bandwidth). This means up to a third of the signal might be outside the receiver bandwidth. Along with the large communication range, this renders a large fraction of the packets undetectable (see Figure 3). If one were to use a software-defined radio with wide bandwidth, it would be easy to collect the raw signal data and post-process it to overcome this bottleneck. Instead, Vista takes the approach of re-configuring existing LoRa base stations with a simple software update to estimate and overcome Doppler.

Packet Detection using wideband correlators: The key reason for packet loss due to Doppler in LoRa base stations is because the correlation peak is far outside the bandwidth of the sampling unit at the receiver. To overcome this challenge, we tradeoff some noise resilience for improved Doppler robustness. Specifically, we set the receiver to the widest bandwidth correlators (250 kHz) and leverage the fact that even low-bandwidth LoRa preambles can correlate with these wideband correlators if configured properly. To see why this is possible, consider Figure 4: Notice that carefully configured correlators at 250 kHz share common signal fragments with those at lower bandwidths. However, only some fragments match (four out of eight preamble signals in Figure 4 (top)), meaning that this process results in a 3 dB reduction of detection SNR—trading off some SNR resilience for improved detectability. At this point, Vista can simply estimate Doppler using the SYNC symbols of a larger bandwidth (much like traditional LoRa) to obtain a coarse estimate of Doppler frequency shift. We illustrate this process in Figure 4 (bottom).

Adapting off-the-shelf hardware for wideband correlators: One crucial aspect of the above approach is the wide-band correlators that it requires for detecting the high-Doppler packets. Thus, to enable this approach on off-the-shelf ground stations it is critical to demonstrate that there exist correlators at higher bandwidths with the same frequency slope as that of narrow bandwidth

Preamble Frequency Slope (kHz/s)		Bandwidth (in kHz)				
		31.25	62.5	125	250	500
Spreading Factor	SF7	7.62	30.51	122.07	488.28	1953.12
	SF8	3.81	15.25	61.03	244.14	976.56
	SF9	1.90	7.62	30.51	122.07	488.28
	SF10	0.95	3.81	15.25	61.03	244.14
	SF11	0.47	1.90	7.62	30.51	122.07
	SF12	0.23	0.95	3.81	15.25	61.03

Fig. 5. Off-the-shelf ground stations are capable of receiving signals of various configurations across spreading factors (SF) and bandwidths (BW). The color-coded figure shows how each low bandwidth signal at a spreading factor between 7–10 has a higher bandwidth preamble with the same preamble slope (one such example has been marked in the figure).

transmissions. Most off-the-shelf ground stations [17, 76] support multiple spreading factors and bandwidths while CubeSat deployments [21] typically operate on the narrowest bandwidth configurations. As shown in Figure 5, all low-bandwidth configurations at lower spreading factors have an equivalent higher bandwidth correlator (preamble) that provides the same frequency slope as the original low-bandwidth transmission. In fact, this is not really that surprising as increasing bandwidth quadruples the slope while increasing spreading factor halves the slope. Thus, we can leverage these higher bandwidth configurations with the same frequency slope to detect the Doppler offset and use interrupts on the base station to decode the rest of the packet in the right configuration with the Doppler offset corrected.

Satellite trajectory to inform Doppler and vice-versa: As the above approach of estimating and correcting Doppler requires us to use wider preambles and accumulates more noise, it becomes critical for the base station to know when the CubeSat is overhead to switch into high Doppler-resilient mode—trading off SNR for detectability. However, most small CubeSats cannot be tracked by ground radars for the first few weeks (the most critical time for a CubeSat mission) [4]—owing to being highly clumped up in a volume with other entities launched with them. However, the few initial Doppler measurements obtained from our LoRa ground stations can help us indirectly inform satellite orbits, much earlier than ground radars can capture them. These orbital measurements can then directly inform future expected Doppler values for these same satellites in subsequent passes.

We use a simplified version of prior Doppler-to-orbit models [53, 54] to estimate the trajectory from our coarse Doppler measurements using the algorithm described in Algorithm 1. While the precise orbit of the CubeSat may not be known, one can still make very good predictions about its general structure and allowable shape. Mathematically, the variation of Doppler in an overhead pass from a CubeSat can be written in terms of two main parameters besides time (t) [28]: ϕ, the inclination angle and θ_{max}, which is the maximum elevation achieved by the satellite as viewed by the receiver:

$$\frac{\Delta f}{f} = D(t, \phi, \theta_{max}) = -\frac{1}{c} \frac{R_e r \sin(\omega_f t) \cos(\cos^{-1}(\frac{R_e \cos(\theta_{max})}{r}) - \theta_{max})\omega_f}{\sqrt{R_e^2 + r^2 - 2R_e r \cos(\omega_f t) \cos(\cos^{-1}(\frac{R_e \cos(\theta_{max})}{r}) - \theta_{max})}},$$

ALGORITHM 1: Trajectory Estimation using Doppler

Input: Measured Doppler $f_o^1, f_o^2 \ldots f_o^n$, **Period** t_p
Output: Trajectory parameters $\theta^*_{max}, \phi^*, t^*_{start}$
//Create Spline interpolated signal
//s(t)=Spline({$0, t_p, 2t_p, \ldots$},{$f_o^1, f_o^2 \ldots f_o^n$})
forall $\theta_{max} \in \{0 \rightarrow \frac{\pi}{2}\}, \phi \in \{0 \rightarrow \pi\}$ **do**
  $s^*(\phi, \theta_{max}) = D(t, \phi, \theta_{max})$
  //D is the emulated Doppler Curve $\Gamma(\phi, \theta_{max}) = max(s^*(t) * s(t))$
  //Cross Correlation $\tau(\phi, \theta_{max}) = \arg\max_t(s^*(t) * s(t))$
$\phi^*, \theta^*_{max} = \arg\max_{\phi, \theta_{max}} \Gamma(\phi, \theta_{max})$
$t^*_{start} = \tau(\phi^*, \theta^*_{max})$

where R_e is the radius of Earth, r is the orbital radius of the satellite, ω_f is the angular velocity of the satellite (function of spacecraft trajectory and Earth's rotation), and c is the speed of light. Given these parameters, we formulate our optimization as finding the best orbital parameters that fit our Doppler observations (f_o^i) across packets ($i = 1, \ldots, N$), specifically:

$$\underset{t_{start}, \phi, \theta_{max}}{\arg\min} \quad |\alpha|$$

$$\text{where } \alpha = \sum_{i=0}^{N} \left| f_o^i - D\left(t_{start} + it_{period}, \phi, \theta_{max}\right) \right|$$

The above objective function is not easy to optimize given the complex expression of the Doppler behavior. Specifically, while the expression is convex along the ϕ and θ_{max} dimensions, the randomness of t_{start} stops us from using gradient descent. However, since the possibilities of ϕ and θ_{max} are bounded, we can directly iterate on them as shown in Algorithm 1.

While there are many more complex and accurate models for estimating the trajectory from Doppler estimates, our results in Section 7.2 highlight the improvement in trajectory estimation due to better Doppler estimation by Vista. Upon estimation of this trajectory, we can leverage this information in future packets to get more accurate Doppler estimates of the incoming packets. Note that we can also use this information to reduce the bandwidth of correlator symbols to improve the SNR of reception. Finally, we can use this modified Doppler estimate to correctly decode the communicated information on the ground infrastructure and leverage it for decoding images and performing inference tasks.

3.2 Combating Interference from Other LoRa clients

As previously mentioned, the packets from LoRa-enabled LEO satellites are severely attenuated when they reach the ground station infrastructure. However, despite that stark 150dB of attenuation, these signals are still above the detection noise floor of the commercial off-the-shelf LoRa base stations in the absence of other radio interference. Yet, these ground stations are also designed to serve several terrestrial clients—their primary and original purpose. These clients are much closer to the base stations than the CubeSats meaning that their signals are usually larger in signal power compared with the CubeSat signals. Thus, the presence of terrestrial clients significantly reduces the capability of ground stations to detect packets from CubeSat clients.

One approach to overcome this problem of packet collisions is to leverage existing solutions for parallel decoding of collided LoRa packets such as Choir [48], MaLoRa [59], CoLoRa [84], and so on. However, these solutions are designed to decode packets that can be detected. Yet, our scenario, is plagued by an extreme near-far effect where the cubesat signal is undetectable owing

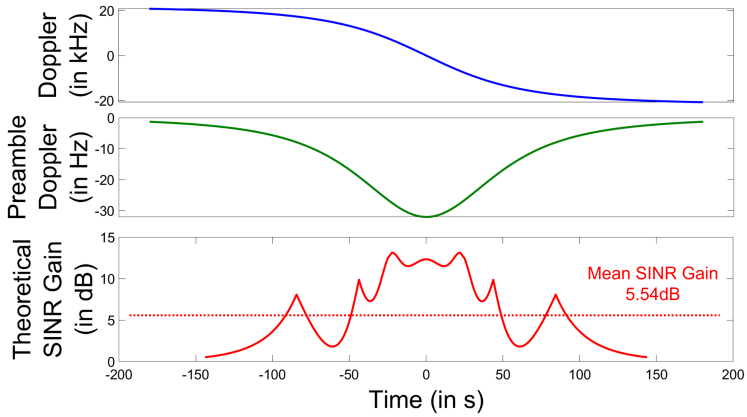

Fig. 6. Accumulation of Doppler across preamble of CubeSat LoRa packets leads to inefficient correlation with traditional preambles. Vista's trajectory-driven packet filters remove these offsets to achieve good packet detection in presence of terrestrial interference.

to being overwhelmed by ground clients. We, therefore, need a solution that somehow specializes in detecting packets from CubeSats and filters out packets from terrestrial clients.

At first glance, one might presume such filtering cannot be done given that CubeSat LoRa packets use the exact same modulation as ground clients. Yet, while the transmitted signal is the same, the channel between the satellite and the ground station is tightly coupled with the trajectory of the satellite leading to an important effect on every individual packet—intra-packet Doppler shifts (see Figure 8(a). Specifically, given that LoRa packets and symbols are several milliseconds long, the large amount of Doppler drift that affects these packets changes within a packet. This means instead being bound inside a bandwidth of the signal the signal may drift as much as several tens of Hz (see Figure 6). Thus, given the trajectory of the satellite (see Section 3.1), we can develop special packet detection filters that detect only the packets from a speci fic satellite instead of detecting packets from all LoRa clients.

Doppler-division Multiple Access (DDMA): The key idea behind DDMA is that while the packets face different Doppler frequency shifts, they are not random. Further, these shifts can be estimated across time using NORAD trajectory parameters or our own trajectory estimates above. Thus, the construction of these CubeSat packet detection filters given a Doppler trajectory Δf_t can be formulated as follows. Given a detection time t_0, we create the series of upchirps and downchirps in the preamble $x(t)$ similar to conventional LoRa packet detection. Given this basic preamble we multiply this vector by a satellite trajectory compensation factor described by $e^{-2\pi j(\Delta f_t - \Delta f_{t_0})t}$. We then use our new preamble $y(t) = x(t)e^{-2\pi j(\Delta f_t - \Delta f_{t_0})t}$ to detect packets from that specific satellite.

Why does this work? Similar to how normal preambles correlate well with standard preambles, this new preamble is designed to correlate perfectly with the packets of a specific CubeSat but partially with conventional LoRa packets. This improves the packet detection of CubeSat by reducing the correlation of noise/interference and improving the correlation with the desired signal.

Theorem 3.1. *The signal-to-interference-noise gain of Vista trajectory-driven packet detection filters given Doppler Δf_t, start time of packet t_0 and end time of preamble $t_{preamble}$ is given by:*

$$SINR_{gain} = 20\log_{10}\left(\frac{t_{preamble} - t_0}{\left|\int_{t_0}^{t_{preamble}} e^{-2\pi j(\Delta f_t - \Delta f_{t_0})t}\,dt\right|}\right)$$

Proof. Correlation of trajectory-driven preamble for CubeSat packets $\propto$

$$\left| \int_{t_0}^{t_{preamble}} e^{2\pi j(\Delta f_t - \Delta f_{t_0})t} e^{-2\pi j(\Delta f_t - \Delta f_{t_0})t} dt \right| = (t_{preamble} - t_0)$$

Correlation of trajectory-driven preamble for ground packets $\propto$

$$\left| \int_{t_0}^{t_{preamble}} e^{-2\pi j(\Delta f_t - \Delta f_{t_0})t} dt \right|$$

The rest follows. □

Combating interference from other LoRa CubeSats: Another source of interference would originate from other satellites using the same modality as Vista to communicate information downlink. However, the trajectories of these various LoRa satellites will always be different (otherwise, they would collide!) either in time, or the axis of trajectory. Thus, by using the unique trajectories of these satellites, we can detect and decode the packets from these satellites.

Theoretically, these trajectory-driven preambles can achieve up to 13.12 dB (mean 5.54 dB, see Figure 6) of SINR improvement per interfering client for packet detection in presence of terrestrial interference. We evaluate these preambles in Section 7.2 for various combination of terrestrial clients and LEO satellite trajectories.

4 Channel-Aware LoRa CubeSat Encoding for Images

In this section, we describe the key sources of information loss while decoding the detected packets – noise, discretization and intra-packet Doppler. We then showcase how we can co-optimize real world applications over this wireless channel for optimum performance via a case-study on LoRa CubeSat imaging and inference.

4.1 Sources of Information Loss

There are three key sources of information loss over this wireless link between the LoRa CubeSat and the ground station:

Noise causes LoRa chirps to be misclassified randomly across the frequency shifts. This means that each bit has an equal likelihood of bit-flip regardless of the location of the chirp. To estimate this error, we use the LEO satellite channel model as shown in Ref. [71]. We estimate the **symbol error rate** (**SER**) using the theoretical model in Ref. [26]. Based on measurements from our LoRa CubeSats launched in 2021 [21], our average SNR of −10 dB from these traces at LoRa ground stations correspond to roughly 10^{-4} SER. This directly corresponds to a BER of 10^{-4}.

Intra-packet Doppler affects bit errors asymmetrically. To understand how, we first need to understand how a LoRa chirp is decoded: The receiver typically deconvolves an incoming chirp with another upchirp to ascertain the frequency shift relative to the frequency shift of the preamble. Thus, generally much of the Doppler is overcome by the preamble and one would assume there would be no effect on the decoding process. While this is true for traditional wide-band satellite transmissions lasting a few millisecond, many LoRa packets last several hundreds of milliseconds. This means a significantly larger amount of Doppler accumulates over the duration of the packet, causing the relative offsets of chirps near the end of a packet to be severely deviated from the start of the packet (Figure 7(a)). Doppler-induced frequency shifts lead to the decoded chirps shifting into nearby bins. This phenomenon of misclassification into nearby bins instead of random frequency bins leads to the asymmetric effect across bits (Figure 7(c)). Remember that most adjacent bins only differ in the least significant bits, while the **most significant bits** (**MSB**) only change across massive frequency shifts (LSB). In fact, going from the MSB to the LSB, the likelihood of

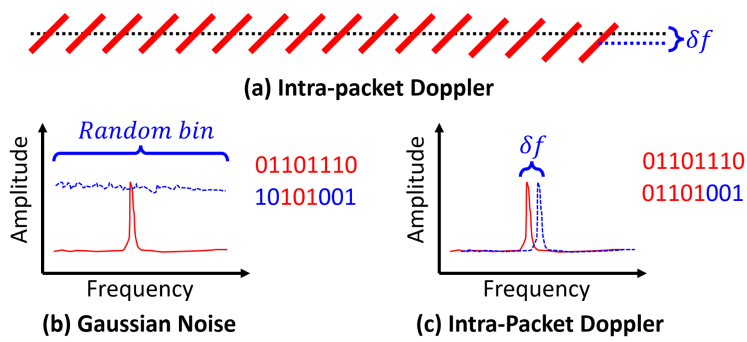

Fig. 7. LoRa symbols face an additional frequency shift due to accumulation of Doppler across the packet. Unlike gaussian noise which affects all bits of a symbol uniformly, intra-packet Doppler affects the LSB more than the MSB.

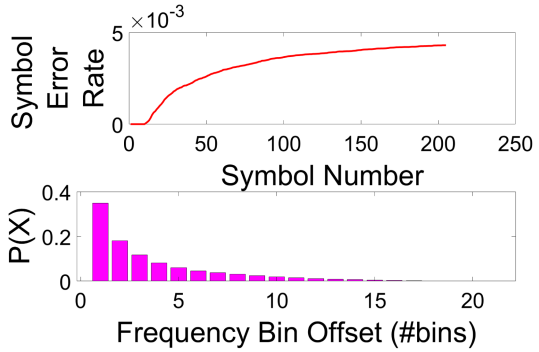

Fig. 8. Intra-packet Doppler affects symbols in a packet asymmetrically. The error likelihood in frequency bins offset follows an exponential distribution.

each bit getting affected increases exponentially. Figure 8 shows the distribution of symbol errors across symbols and the likelihood of amount of frequency bins shifted.

Another source of error is **discretization**. Remember that we are limited to communicating only 2048 discrete bits for every packet. Thus, we need to reduce the analog sensed information to a few bits causing information loss. This restricts any information sent over the wireless link to a very strict constraint.

While all satellite wireless links face the above challenges, these effects are magnified for LoRa CubeSats due to the small bandwidth of LoRa (125 kHz vs. several MHz) and limited size of a LoRa packet (256 bytes per packet). However, we show a unique opportunity that LoRa provides in optimizing applications over this wireless link due to the deterministic nature of information loss and improve data efficiency for inference over images.

4.2 LoRa CubeSat Imaging and Inference—A Case Study

Current approaches to communicate images: Traditional approaches to communicate images such as JPEG, TIFF and PNG are not designed for such low packet sizes (hundreds of bytes). This is due to the fact that much of the compression achieved using these approaches relies on carefully designed dictionaries, hash tables and encoding schemes, all of which take space to store for decoding the image. Thus, to achieve any coherent image communication, CubeSats typically

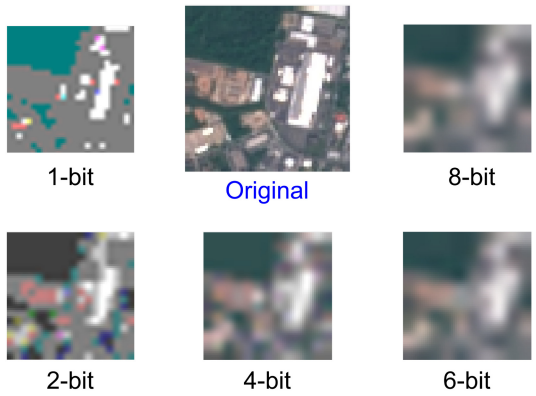

Fig. 9. Spatial and bit-level resolution tradeoff leads to large reduction in image quality of the baseline.

rely on two conventional compression approaches: subsampling and reducing bit resolution of every pixel. Subsampling reduces the spatial resolution of the image while the reduced bit resolution adds corruption due to discretization. Figure 9 demonstrates a few configurations that demonstrate significant loss of features, motivating the need for better solutions.

Data-Channel-Task Aware Approach: To ensure useful data compression within the constraints of LoRa packet size, Vista's key idea is to leverage important known information about the images collected by CubeSats, the wireless channel that the signals propagate over, and tasks performed on the ground. Specifically, we consider: (1) features extracted from a large dataset of typical satellite images, (2) a broad set of tasks that are typically performed on those images, and (3) a well-modelled LoRa-CubeSat channel.

Data Awareness: We first rely on the vast amount of Earth images taken by larger satellites in similar orbits (NOAA [11], Sentinel [18]) to learn how best to compress the images. One approach would be performing a statistical analysis and identifying the patterns across images. Vista instead chooses a more automated approach that learns these patterns using a convolutional autoencoder.

Task Awareness: Furthermore, this autoencoder can be optimized not only for reconstructing the images but also directed toward preserving maximal task-specific information. We design a common encoder for both image and task-recovery that is actively aware of the set of tasks that are intended to be performed at the ground. This loses a small amount of performance in both tasks but retains maximal information feasible for each task.

Channel Awareness: Finally, training and evaluating this autoencoder must account for bit-errors that these LoRa wireless packets typically encounter. We design a channel-task-aware image encoder that maximizes the output accuracy while being constrained to the limits of a LoRa packet. We do so by a two-stage process where the auto-encoder learns about *noise*, *intra-packet Doppler shift*, and *discretization* losses. As shown in Figure 10, the output of our solution preserves much of the useful information in the images for inference and imaging.

Vista Autoencoder Design: Our design of the autoencoder (shown in Figure 2) derives heavily from existing conventional wisdom in designing autoencoders over images for inference applications. We adopt a 3-layer convolutional neural network design on the satellite to maximize the accuracy-compute tradeoff. On the ground, we are relatively less resource constrained, and hence can have more layers. Our architecture co-optimizes for both tasks and image reconstruction using two parallel CNNs via a fully connected layer.

Fig. 10. Vista preserves much of the useful information in the images for inference and imaging.

Our autoencoder design builds on conventional wisdom from traditional convolutional autoencoder designs. As we go from the image towards the signature, we increase the number of convolution features captured by each layer and decrease the size of filters. This allows the Vista autoencoder to preserve the maximum amount of information while ensuring progressive compression. After each layer, we pass the output via a Rectified Linear Unit (ReLU layer) to capture non-linear behaviors which in turn maximizes the ability of Vista to improve over simple subsampling and other linear baseline approaches. Finally, we use a tanh layer at the output.

LoRa Payload and Channel: This output is then discretized to a low-bit representation and sent in a single packet with 256 byte payload. We note that encoder output size and discretization is primarily limited by available bandwidth. These tradeoffs can be task specific, as well as dictated by the wireless channel and we discuss these issues in Section 4.3. At the receiver, the LoRa base station decodes the packet and retrieves the communicated values.

Decoder Design: At the decoder, this signature is then passed on to a fully connected layer to distribute the input to two separate networks optimized for image reconstruction and labeling tasks, respectively. For image reconstruction, we have successive layers of upsampling, convolution, and ReLU to increase the spatial dimensions to that of the original image. For the labeling task, we have successive layers of convolution and ReLU layers with a final sigmoid layer to output the probability of belonging to a class.

Training the Network: We use the Mean Squared Error loss and Cross Entropy Loss functions for image reconstruction and multi-class classification, respectively:

$$l_n = (x_n - y_n)^2$$

$$\text{loss}(x, \text{class}) = \text{w}[\text{class}] \left(-x[\text{class}] + \log \left(\sum_j \exp(x[j] \right) \right)$$

We then train our network with five-fold cross validation on the training data to avoid overfitting while ensuring maximum accuracy. We have restricted our discussion in this section to primarily build a data-task-aware encoding for images based on our constraints. In the next section, we show how we can make this encoding robust to wireless impairments and improve inference performance.

4.3 Channel-Aware Vista Autoencoder Design

Given the above data-task-aware autoencoder design, we next endeavor to co-optimize the encoding with the wireless impairments described in Section 4.1. Broadly, our approach modifies the training of the autoencoder given the data, task and channel information by training it in two stages. As shown in Figure 11: In the ***first stage***, the network learns about the noise and Doppler related bit-errors. In the ***second stage***, it compensates for the discretization loss. Note that both

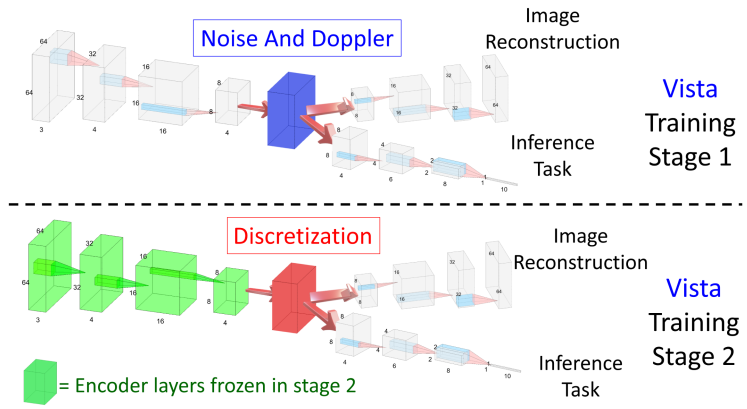

Fig. 11. Vista uses two-stage training to make the autoencoder channel-aware – **Stage 1**: learning noise and Doppler; **Stage 2:** calibration to discretization loss.

these stages of training will occur on the ground a priori and, hence, can be performed over a long time to converge.

Overcoming noise and intra-packet Doppler: To model the effect of noise and intra-packet Doppler, during the first stage of training, we add an additional layer that multiplies the input by a bit-flip operation layer. For a one bit output, we achieve this by creating random bit vectors of bit flips and then convert them to a (−1,1) vector (−1 signifying bit flip). Note that in a one-bit output scenario, the output of the tanh layer lies between (−1,1) which means a simple multiplication layer with our bit mask will work. Further, as this bit-mask remains constant during an epoch, the multiplication layer is both continuous and differentiable. Thus, during emulation of the bit-error, the gradient will also propagate over our custom layer and hence get corrected automatically. Note that this makes the network aware of the behavior of these errors and maximizes accuracy for the resulting image and task.

A new problem arises when each output is represented by multiple bits: The correct way to emulate bit flips would be to take our bit mask and XOR with the output of the encoder. However, it is well known that the XOR operation is not continuous or differentiable. To address this, we use a known approximation of the XOR operation [8] by a 2-layer neural network for the number of bits required to XOR. Hence, we add this XOR black box in place of the multiplication layer where the inputs to the XOR are given by our bitmask. This enables multi-bit outputs to learn the behavior of the bit error distributions caused by intra-packet Doppler.

Overcoming Discretization Loss: Discretization requires us to reduce these 4 byte float output elements to a few bits causing information loss. This means that if we use a 2×32×32 sized output, we can only send one bit of each floating point number to the actual decoder. On the other hand, using a 1×16×16 output for the encoder would allow us to communicate 8 bits of information to the output. As one can imagine, both of these scenarios severely affect the ability of the decoder to decode this input into coherent output.

To address the problem of discretization, we first attempt to understand how it affects the output of the decoder. We observe that the discretization loss at the input of the decoder propagates to roughly equal information loss in the output of a well-trained CNN. For example, the output of the 2×32×32 encoder becomes a single bit-mask on the output of the decoder. On the other hand, 1×16×16 image output suffers significantly less loss. However, this image output looks half the

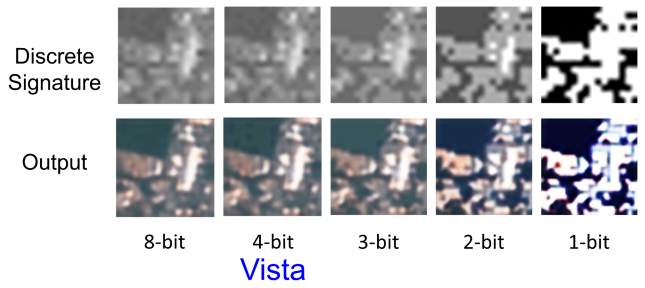

Fig. 12. Discretization of the encoding propagates to the output of the decoder.

resolution of the image output of 2×32×32 image, as each of the distorted pixels in the image is much larger, spanning double the number of underlying pixels.

The average information loss for k-bit discretization is equal to $\frac{1}{2^{k+1}}$. However, it is quite difficult to quantify the amount of information lost due to reduction in spatial resolution. The only potential way to solve this dilemma is to evaluate the loss on the prior available data across both axes and use the least error configuration. Our evaluation shown in Figure 12 shows that the optimum configuration näively operates at 4-bits.

However, we can do even better. We ask a simple question: "Can the decoder learn and specialize for this discretization to maximize the accuracy of the output of the image and tasks?" The first obviously direct approach would be to create a differentiable model of such discretization and add it as an additional layer to the neural network. However, such discretization is neither continuous nor differentiable. Furthermore, unlike XOR, there is no guarantee that it can be estimated with a small enough neural network. Instead, we take an indirect approach to teach this behavior to our decoder: We first freeze the encoder as it is already optimized to best represent the image in the limited set of bins. We affect the output of the encoder with the necessary discretization to create an *intermediate input*. We then relearn the decoder parameters for the image output on these affected intermediate inputs. Our results show how our approach minimizes information loss and maximizes the accuracy across configurations.

5 Discussion and Limitations

Impact of LoRa bandwidth on satellite imaging: While Vista achieves significant improvement in enabling low-latency downlink for quick information retrieval, an important concern is the LoRa's inherent bandwidth limitations as the primary downlink technology for an expensive commercial CubeSat deployment. In fact, our goal in this work is not to replace conventional downlinks, but instead to promote LoRa as a low latency complementary backhaul for specific applications. Furthermore, the contributions in this work are relevant broadly to the CubeSat community for the following reasons:

— There exist many smaller academic [3, 21] and industrial CubeSat deployments [10, 23] that can leverage the proposed mechanisms to improve downlink communication for free without significant cost investment quickly.
— LoRa has an extremely low power, weight and cost overhead [58] to be added as a secondary communication technology in larger CubeSats as a low-latency alternative for various applications that may require quick low-quality images, for instance, forest fire detection, oil fires in oceans.

— Many of the algorithms and solutions developed by Vista are also applicable for other low bandwidth technologies such as NB-IoT and FHSS transmitters traditionally used by other smaller CubeSats and hence present a broad applicability.

Open Challenges: Based on our experience launching the CubeSats, it is quite evident that there are several problems that need to solved to enable robust connectivity from CubeSats to the ground infrastructure. On the CubeSat, the list of open challenges range from enabling secure robust geo-location to new antenna designs and modulation schemes. On the ground station, it is critical to develop new hybrid base stations that can simultaneously serve terrestrial as well as CubeSat clients and accurately estimating the trajectory of these clients in the critical first month of their operation. Finally, there is a new opportunity for data scientists to maximize the utility of these bandwidth-starved CubeSat links by building better and more robust data-driven encodings.

Customizing Vista Training: A crucial component of Vista's benefits stem from both the data and application information that it presumes as well as the ability to operate across Doppler offset channels. Now the data and the application components will be mission specific and rely on having prior mission data from larger and higher bandwidth satellites. In our case, we use the application of land-use classification which has been explored quite extensively in prior work (See Section 8). However, the heterogeneous error distribution created by the Doppler curve can be computationally generated quite easily by knowing the altitude of the CubeSat (a factor that is known significantly prior to launch). Further even small inaccuracy in the altitude creates roughly similar Doppler curves creating similar error distributions for relevance to our requirements of training the data-application-channel-aware model. Vista shows that if we have been performing an application extensively using prior satellite missions, we can do so even at low-bandwidth with enough fidelity for practical use-cases via these encodings. All of the training will be performed a priori and programmed for the specific application. While the system will overcome channel-aware aspects of the training always, the success of data-application-aware component relies on the availability of high quality training data (which is true for all data-driven systems).

Spectrum Policy Concerns: There are two entities that govern the spectrum in the low-earth orbit. The first is the **Federal Communications Commission** (**FCC**) [40] that provides licenses to US-flagged spacecraft (i.e., any US company, university) for communication from satellites. The other entity is **National Telecommunications and Information Administration** (**NTIA**) [70] which provides special permissions for experiments on spacecraft backed by US government (i.e., NASA). The only policy-wise complicating factor in 915 MHz band [41] is that, it is specially allocated for federal civil and naval operators for radio-location services among others. However, recent meetings between these two entities point towards building a unified spectrum policy in the near future [69]. The CubeSat built by our team operates in the ISM band received permission to transmit a 915 MHz LoRa downlink communication at a bandwidth of less than 125 KHz.

Ground Station Feasibility Concerns: Another important feasibility concern for deploying Vista is that there is presumption that LoRa is allowed to operate at the same band across the world which is not true. However, SemTech has several transceivers that spans the complete spectrum utilized for LoRa across the world which can be multiplexed as well as specialized boards can be developed that can operate at multiple bands. There further remains the problem of having various antennas across these bands at the satellite which needs to be tackled in the future. Our evaluation specifically focused and operated under the US spectrum policy.

Furthermore, another complex topic of importance is how will ground stations serve both terrestrial and satellite clients. While most terrestrial base stations are today deployed with antennas with high lateral gains, most of the next generation of SemTech LoRa base stations have several

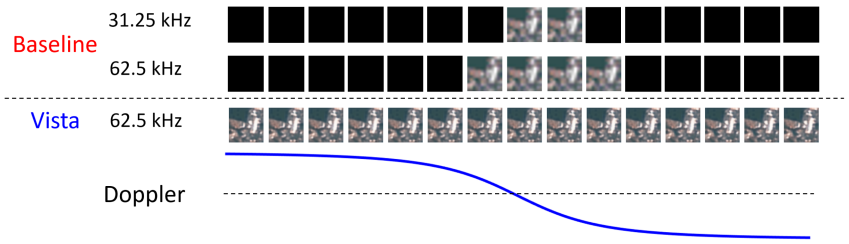

Fig. 13. Estimated image throughput over one pass of Vista CubeSat over an urban city with max elevation of 90° and inclination angle of 97.52°. Current terrestrial LoRa ground stations would be able to only detect and decode only 2–4 packets compared with the 15 encodings decoded via Vista.

RF chains that have separate ADCs, and so on. Thus, it is quite feasible to use an antenna with high longitudinal gain on one of the RF chains to receive only the satellite signals. However, given the current extensive deployment with purely omnidirectional antennas, we demonstrate how the new Doppler-correcting preambles suggested in Section 3.2 can serve as an intermediate solution for disambiguating the satellite and terrestrial clients.

Evaluation Limitations: Our evaluation is limited in the following aspects: (1) Actual end-to-end testing on the satellite is infeasible as ground-truth image information is unavailable due to limited satellite bandwidth. (2) Emulating high level Doppler offsets by actually moving a transceiver on the ground is infeasible as the orbital velocity of Vista satellite is 7 km/s. Thus, we use the signal measurements from our launched satellite to perform trace driven evaluation. (3) We use images from public LEO satellites as candidate packet payloads for our trace-driven emulation. Note that these are not from CubeSats due to limited public data of satellite imagery specifically from CubeSats.

6 Implementation and Evaluation

Our evaluation leverages signal and channel measurements from the CubeSat to inform Vista's approach to overcome the wireless challenges described in the article. We collected various channel traces from our own base station as well as other base stations around the world (via SATNOGSDB [16]). We use these channel traces to interpolate the channel metrics (such as attenuation, noise floor of the receiver, other interference metrics) that we use to evaluate our system. We implement a ground-based evaluation testbed that emulates CubeSats attenuation and Doppler behavior using SemTech transceivers SX 1257 and SX 1262. The CubeSat orbit parameters were as follows: altitude = 525 km, inclination angle (ϕ) = 97.52. The TLE information is described below:

$$1\ 70305C\ 21006E\ 21024.66625831\ .00096549\ 00000{-}0\ 57652{-}2\ 0\ 03$$

$$2\ 70305\ 97.5225\ 88.3417\ 0007037\ 262.1956\ 98.7323\ 15.11146297\ 16$$

The output power from the CubeSat was 27 dBm and the received SNR on the ground with conventional approaches varied between −10 to −30 dB. However, Vista solutions can improve the packet detection SNR significantly by approaches mentioned in Section 3.1 and evaluated in Section 7.2.

Baseline: As shown in Figure 13, the baseline of subsampling (Section 4.2) and communicating images will lose several packets without getting decoded. However, comparing our solution with pure noise would be unfair. Thus, we assume that packets were somehow detected for evaluation.

Training the channel-aware autoencoder: The training happens in two stages: (1) In the first stage during training we add the noise and Doppler layer (weights set to not change internally) and

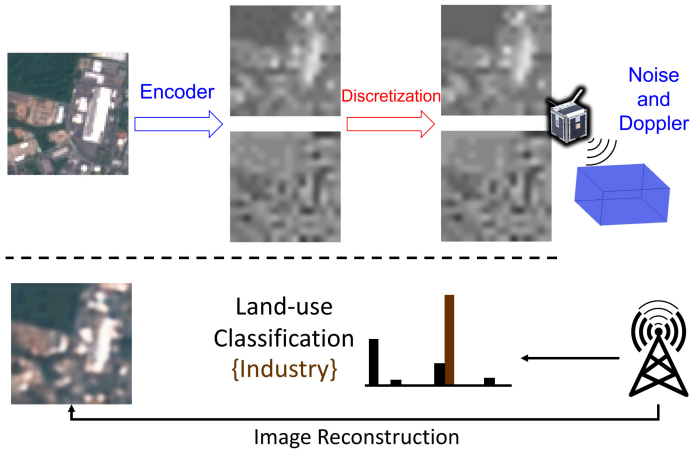

Fig. 14. Process of Vista evaluation.

train the Vista on training data. (2) In the second stage we freeze the weights of the encoder and first generate *intermediate inputs* as suggested in Section 4.3. We then discretize them to fit within our communication budget of 256 bytes. We train the decoder chain on this discretized input to calibrate for the discretization losses.

Evaluating Vista: Figure 14 shows an example run of how Vista evaluation is performed. The input image is encoded at the encoder to output the signature which is then discretized and modulated atop LoRa. It is then passed through the channel model informed by wireless channel measurements and analyzed noise and Doppler channel in Section 4.3. At the receiver, the signature is demodulated and passed to the decoder for outputting class labels and reconstructed image. For image reconstruction, we use the mean squared error loss and for classification, we estimate the confusion matrix. The baseline is subsampled to minimize information loss as shown in Figure 9. It is passed through the same architecture without the discretization, noise and Doppler effects to ensure fair comparison.

Dataset used: We use publicly available labeled dataset with images captured from Sentinel-2, EuroSat Land Use Dataset [56]: 27000 13-channel 64×64 images with 10 seperate land use labels. For our evaluation, we restrict ourselves to RGB channels as those are the ones that image captures.

7 Results

We evaluate the three aspects of Vista: Latency Benefits (Section 7.1), Doppler and Interference Resilient Packet Detection (Section 7.2) and a case-study on Imaging and Inference for LoRa CubeSats (Section 7.4).

7.1 Latency Benefits of Vista

Setup: We evaluate the potential latency benefits of using existing LoRa ground stations (operating on the 915 ISM band) over specialized ground infrastructure such as large generic satellite base stations or LoRa ground stations in lower UHF bands. Now, it is virtually impossible to know the locations of every possible ground station in all of these categories and thus we endeavor in procuring the publicly known locations of these ground stations on three platforms : The Things Network [24] (terrestrial LoRa network), T TinyGS [20] (LoRa grounds stations operating at lower bands), SatNogsDB [16] (Satellite Ground Stations). Assuming the TLE information above obtained

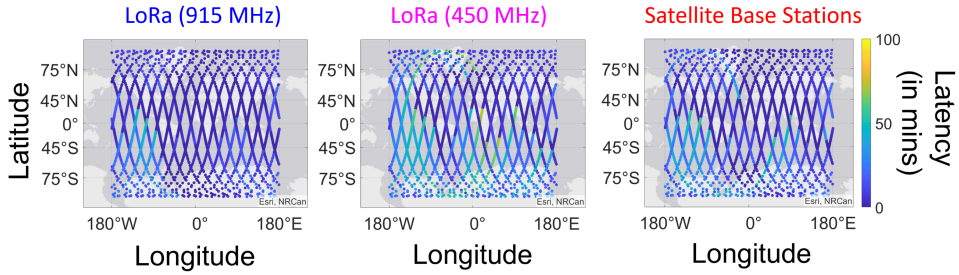

Fig. 15. **Vista Latency Benefits:** Demonstrates the potential latency using different ground station infrastructures over the trajectory of Vista satellite across Feb 1, 2022 to Feb 3, 2022. LoRa ground infrastructure operating in the ISM band can enable lower latency compared to other ground station infrastructures in developing regions of the world.

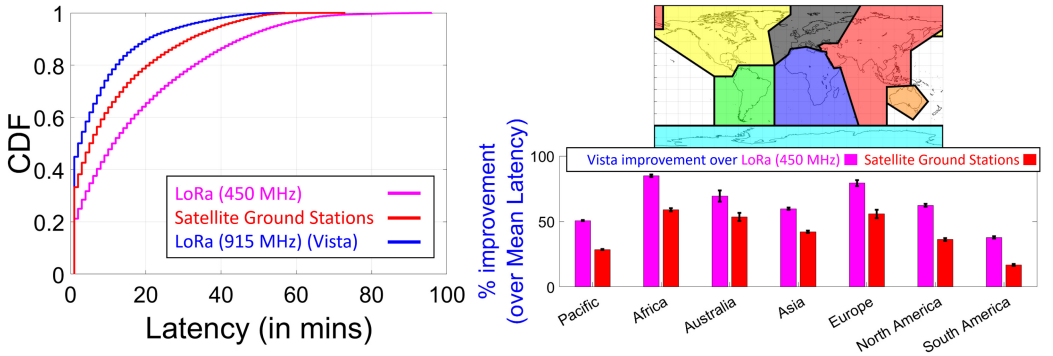

Fig. 16. **Vista Latency Benefits:** ISM-band LoRa CubeSats can enable 55.55% lower latency and 44.92% coverage compared to 33.27% due to the extensive availability of terrestrial ground infrastructure in the ISM band. These benefits are geographically distributed across the globe with significant improvement noted in developing countries around the world.

from NORAD, we use the location of our satellite across a year in space to evaluate the potential latency of retrieving information from any of these networks. We further analyze the benefits geographically from the satellite **point-of-view** (**POV**). Since, there is no clear definition of continental boundaries, we used the map described in (Figure 16, top right) to classify locations of the satellite across continents and geographically analyze the benefits provided by Vista.

Results: Figure 15 shows the latency across trajectory of the satellite using various ground infrastructure during two days of operation. It is quite evident that across many developing regions where it is difficult to develop and deploy expensive specialized ground infrastructure, leveraging existing LoRa deployments in ISM bands provide a clear win. Our evaluation across a year of Vista trajectory (Figure 16, left) demonstrates that that leveraging the widely deployed 915 MHz LoRa ground infrastructure will enable significantly lower latency of image downlink compared to using generic satellite ground stations or specialized LoRa infrastructure. Across the orbit of an year, LoRa infrastructure at 915 MHz provides a coverage of 44.92% compared to 21.23% of TinyGS and 33.27% of SatNogsDB. In terms of 90%ile latency, using Vista for LoRa CubeSats can potentially reduce latency by up to 55.55%. This benefit indeed is distributed across the globe with improvement across continents. As shown in Figure 16 (bottom right), the benefits are most pronounced in Africa where there is a lack of satellite reception infrastructure as well as Europe which has the

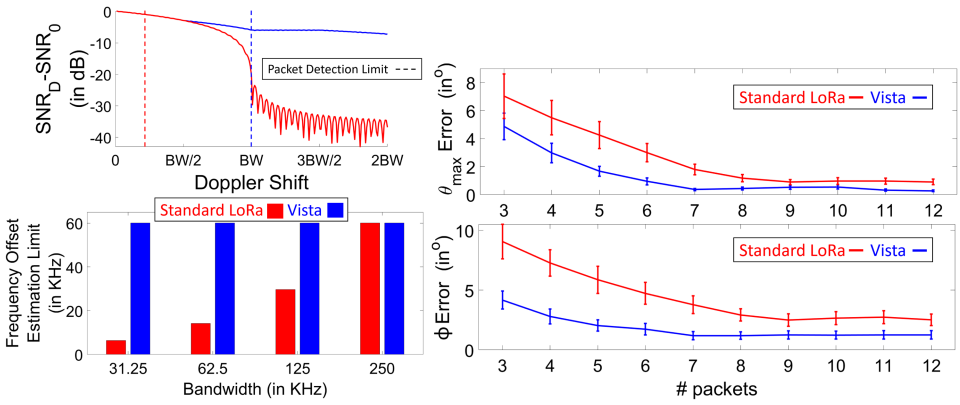

Fig. 17. **Vista Packet Detection Benefits:** (a) Demonstrates the ability of Vista to detect packets beyond Doppler limits of standard LoRa receivers; (b) Vista's accurate Doppler estimation leads to lower errors in predicting trajectory of the CubeSat; (c) Vista's trajectory-driven preambles improve resilience to interference from terrestrial clients and other CubeSats.

largest deployment of LoRa infrastructure. This demonstrates the potential of LoRa transceiver using Vista to reduce the latency of transmitting a quick snapshot to ground infrastructure.

7.2 Packet Detection Benefits of Vista

Setup: We evaluate the efficacy of Vista's wider correlators and trajectory-driven preambles in detecting Doppler offset packets. We use off-the-shelf Semtech SX1262 base stations to receive LoRa packets with increasing Doppler offsets and measure the limits of packet detection limits. We measure the received signal power across offsets using SemTech SX1257 base stations (with access to I/Q samples) for Vista's wideband correlators over standard narrowband decoding. We use the state-of-the-art orbit models [53, 54] to emulate the behavior for LEO trajectory by varying values of maximum elevation (θ_{max}), inclination angle (ϕ) and start time of transmission (t_{start}). We also add additional frequency noise as expected from a real satellite due to angular velocities, offset estimation resolution and other factors. We then use Vista's approach to estimate the trajectory using Algorithm 1. We then use this trajectory to test the efficacy of the trajectory-driven preambles.

Doppler-Resilience: Our measurement study in Figure 17(a) demonstrates that the frequency offset estimation limit of our wideband correlators is as good as the widest correlators supported by the hardware. Furthermore, using these correlators adaptively shows large gains in SNR reaching up to 15 dB when Doppler is as large as the bandwidth. Further, packet detection limits quadruple the resilience to Doppler shifts compared with baseline LoRa base station. As shown in Figure 13, this could lead to large increase in packet throughput when receiving signals from a single CubeSat pass.

Trajectory Optimization: Our results show that using the preamble along with the sync symbols achieves 4× better frequency offset estimation resolution compared with only two sync symbols. Further, due to the better resolution in measuring the offset estimate, we can estimate the trajectory with a resolution of 1.2° in inclination angle (ϕ) and 0.5° in max elevation (θ_{max}) which is comparable to the accuracy of available TLE files for much bigger satellites [4] (see Figure 17(b)).

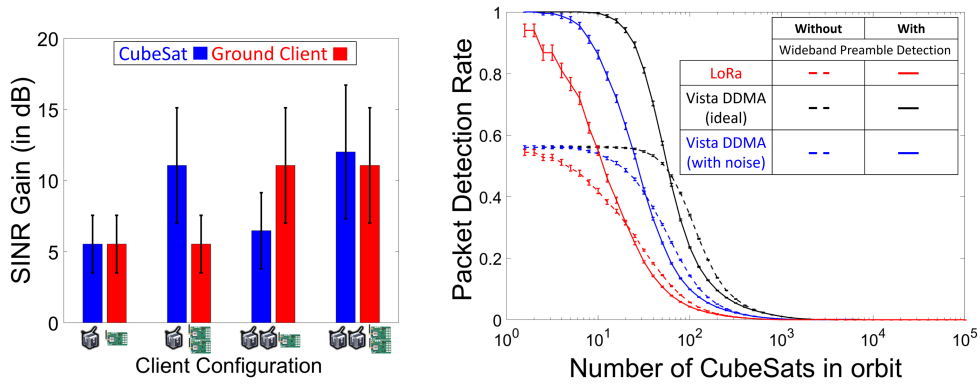

Fig. 18. **Vista Interference Resilience and Scalability:** (a) Vista's DDMA reduces interference significantly for incumbent terrestrial clients as well as other CubeSats. (b) DDMA also enables roughly 4× scalability of CubeSat clients operating within the same spectrum compared with traditional LoRa.

Furthermore, the inclination angle remains consistent within a margin of 1^{o} over the lifetime of a CubeSat, which can be further used to refine Doppler measurements.

7.3 Scalability and Interference Resilience of Vista

Setup: In this section, we evaluate the efficacy of our trajectory-dependent preambles' based Doppler-division multiple-access (DDMA, Section 3.2) to improve the interference resilience and scalability of LoRa-enabled CubeSats. We first demonstrate the signal-to-interference-noise (SINR) gain (in dB) by taking multiple 2 satellite configurations (by varying trajectory parameters) along with 2 terrestrial clients (one each deployed close and far from the ground station) and demonstrate the SINR benefits of leveraging a DDMA based packet reception. We take channel data for ground clients by choosing the clients from publicly available datasets [51] and take the channel data for clients in space by simulating various satellite trajectories. We then scale this evaluation to multiple (upto 10^5 satellite orbits and demonstrate how both Vista's packet detection solutions (namely, wideband correlators and DDMA) work together to improve packet reception. While this simulation assumes all the colliding packets are transmitted simulteneously, the true scalability of LoRa will further amplify the number of satellites that can be simulteneously serviced by a single ground station.

Interference Resilience: Our results in Figure 18(a) demonstrate the SINR gains of leveraging trajectory-driven preambles described in Section 3.2 for detecting packets in presence of interference. As we can clearly see, when there are two concurrent CubeSat transmitters along with two terrestrial clients, Vista preambles improves the interference resilience of packet detection by 12.02 dB for CubeSat clients and 11.08 dB for terrestrial clients at an average. Note that the baseline absolute SINR of configurations with more clients will be lower and thus this valuable gain would mean the difference between detecting a packet and missing it completely.

Scalability: Our results in Figure 18(b) show the key benefits of using the wideband preamble detection and DDMA for multiple access in improving the packet reception rate of LoRa-based CubeSats. First, as demonstrated in Figure 13, only 56% of packets transmitted across a single CubeSat pass can be detected using conventional off-the-shelf ground stations leading to the lower packet detection rate. Further, this improvement in packet detection SINR enables more packets to

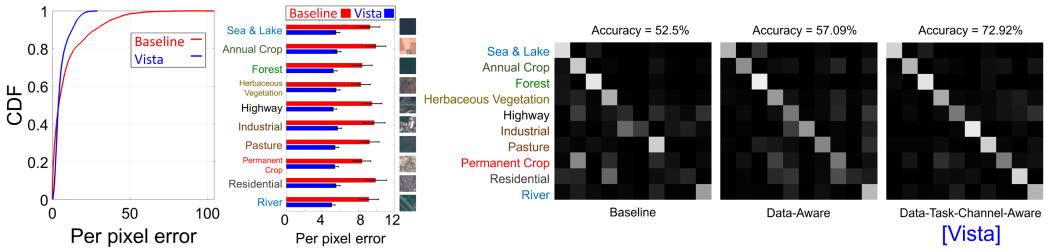

(a) (b)

Fig. 19. **Vista Channel-Aware Imaging and Inference:** (a) Vista produces significantly lower per-pixel error across 5120 RGB images captured by Sentinel-2 satellite leading to 4.56 dB improvement in PSNR (b) The data-task-channel aware approach of Vista achieves 1.38× higher accuracy in classifying 10 different land-use labels compared to baseline approaches.

be detected by the ground base station. Across the various packet detection rates, we can clearly see an average 4–4.5× increase in number of CubeSat clients that can be supported. Note that these numbers are bottlenecked by the noise floor which plays a crucial role in determining if the packets are detected. This means there is a new opportunity to further improve scalability by improving the noise resilience of the core LoRa technology (for instance, longer preambles and interspersing the preamble symbols between the packet).

7.4 LoRa CubeSat Imaging and Inference

Setup: Our model (Figure 2) uses a fully-connected layer to co-train two parallel CNNs optimized for image reconstruction and multi-class land-use classification with 50–50 (5,120 images each) train-test split. The dataset is taken from Ref. [56], which contains labeled images captured from Sentinel-2 [18] satellite from an altitude of 800km. We only choose the RGB spectrum inputs to classify the image into the 10 labeled land-use patterns: Sea & Lake, Annual Crop, Forest, Herbaceous Vegetation, Highway, Industrial, Pasture, Permanent Crop, Residential and River. For the ***baseline***, we leverage a subsampled image with 8 bit resolution as input. For the ***data-aware approach***, we specially train a Vista autoencoder specifically for reconstructing images. Finally, ***Vista's approach*** co-trains two parallel CNNs optimized for image reconstruction and multi-class land-use classification.

Improving Image Retrieval: Figure 19(a) shows that the baseline's per pixel error is significantly larger than Vista's approach. One of the crucial characteristics to observe is that the 90th percentile error is especially larger for the baseline which means while in some cases the baseline may do well, as more fine grained features get involved the reconstruction becomes much worse. For example, a lake may be easier to identify from even a low-resolution image of the baseline while an industrial complex is far more challenging to make out. This enables Vista to achieve a lower mean squared pixel error of 5.42 compared with 9.17 which is 40% lower error. This corresponds to an average 4.56 dB improvement in image SNR (note that the ideal image SNR is ∞). Note, however, that Vista's approach was not specialized for image reconstruction but also for accomplishing task using the same 256 bytes of information. A specialized network such as data-aware approach which is specialized for image reconstruction may perform even better at image reconstruction but will perform poorly at tasks. This is due to the fact that the latent space of compression is not necessarily tuned in the right dimensions for the task. Thus, we do not compare Vista with this method as a fair comparison.

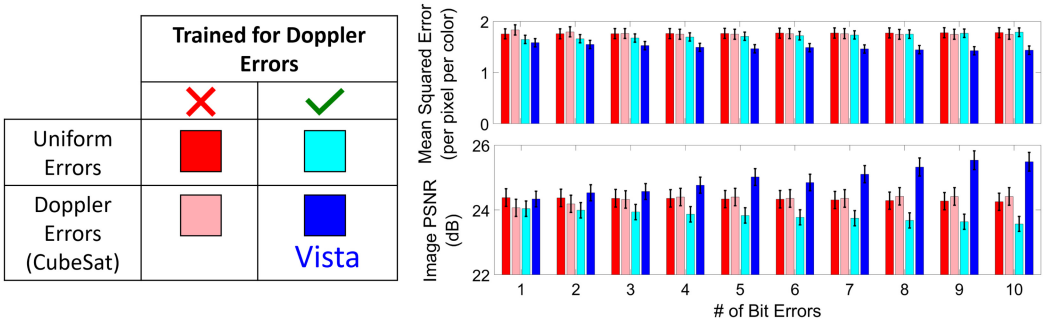

Fig. 20. **Vista Ablation Study:** Training Vista on the exponential error distribution produced by the effect of the Doppler on LoRa modulation makes Vista resilient to Doppler error distributions. As we can clearly see, while a normal training procedure produces similar PSNR and MSE metrics for both errors, the two-step training of Vista, significantly improves the PSNR and lowers the MSE in the presence of Doppler produced bit errors.

Improving Inference Quality: As shown in Figure 19(b), Vista's approach outperforms both the baseline approach as well as data-aware approach significantly leading to 72.92% accuracy in classifying the images compared to 52.5% and 57.09% accuracy of the former over 5120 test images (512 images for each class). The classification tasks where Vista's approach achieves much of its gains are those which require high resolution intricate details such as outlines of buildings for residential and industrial complexes (higher spatial resolution). Further, industries and highways are both made of concrete (grey colored) and face larger errors in the baseline. Note that prior work on the complete 13-channel dataset achieves 84.48% classification accuracy in the original work [56]. But since we use only RGB channels, our classification accuracy is lower.

7.5 Vista Ablation Study and Overhead

In this subsection, we present the impact of Vista on the CubeSat compute and memory resources as well as the efficacy of training an autoencoder aware of the channel error characteristics.

Channel-aware Autoencoder vs. Autoencoder Performance: We evaluate the efficacy of our training approach in improving the noise resilience of Vista's channel-aware encodings. Note that there have been several works that have developed and deployed better and more comprehensive data-and-application aware encodings on satellites such as [43], and thus we will primarily focus this evaluation on the channel-aware nature of our encoding. ***Setup:*** We train two autoencoders on the same training data as follows: one trained with uniformly distributed errors (AWGN) and one trained on Doppler noise errors as described in Section 4 and Figure 8. We then evaluate both of these autoencoders using separate test data on both of these error distributions across bit errors. ***Improvement of Channel-aware Autoencoder:*** As we see clearly see, an autoencoder trained on the right distribution of Doppler errors does perform better compared with the autoencoder that was not trained on the same distribution. In fact, this is not surprising as training on the right distribution allows the autoencoder to ensure that the least valuable information is encoded on those bits. As shown in Figure 20, we see that as the number of bit-errors reach 10 bit errors (out of 2048 bits), the channel-aware autoencoder has a lower image reconstruction error in the presence of Doppler errors.

Power, Compute and Memory Overheads: The three key resources on the CubeSat are power available to use, compute resources available for processing the data collected, and memory to

store the data prior to downlink transmission. (1) ***Power:*** It is a well known fact, that wireless communication dominates the power consumption of CubeSats as enough transmit power is required for the signal to reach the ground. Typically, a CubeSat generates energy with solar and consumes it when in eclipse. While the compute and sensors consume power on the order of a few μW for each image capture, the power required for transmission downlink by the RF frontend (even for a low-power technology like LoRa) is 1-2 W to enable efficient reception on the ground infrastructure. Thus, the extra compute overhead of performing Vista has ***negligible*** impact on the power budget of downlink transmission. Indeed, recent works [43] have deployed more complex designed data-driven algorithms to compress images. (2) ***Compute:*** At first glance, one might feel the extra compute required to implement Vista will be significant. However, remember that the training of Vista will be performed on the ground prior to the launch of the satellite. Thus, the primary impact on compute latency would be to run Vista pre-trained algorithm on images. Our evaluations on a replica of V-R3X [21] compute resources show that it requires an average of ***4.83 ms*** to encode one image.[2] (3) ***Memory:*** The final constraining factor on the CubeSat to deploy Vista is the amount of memory required to store the encoder model that creates the encodings. Our encoding model with 8-bit weight representations requires only ***36 kB*** to store.[2]

8 Related Work

Low-power wide-area networks (LP-WAN) based CubeSats: Past work has explored improving the range [45, 50], throughput [52, 81, 90], scalability [51, 52, 57, 77] and reliability [26] of LP-WANs. Prior systems [68] have also delivered images over LoRa by communicating images over several packets. However, CubeSats do not have the luxury of perennial coverage.

The idea of using LoRa for CubeSats has been explored by both, companies such as Swarm Technologies [19] and Lacuna Space [9] as well as hobbyists as a simple off-the-shelf technology for communication [62] over long ranges. Much of this prior work relies on dedicated ground stations [20] that operate at lower frequencies (such as 137 MHz and 450 MHz) instead of traditional ISM-band (915 MHz) LoRa base stations. Since the launch of the satellite, there have been several other significant launches by academic [31, 47, 66, 73, 91] and government entities [30, 65] that have explored this alternative technology. Further, while there is some theoretical analysis on understanding LoRa behavior in CubeSat scenarios, most of these miss important aspects that become apparent when the CubeSat is actually deployed such as Doppler [49], ignore PHY layer considerations [72]. There have been several works recently that have explored the theoretical problems face by LoRa transceivers and how at low-orbits, these channels face several connectivity issues [29, 46, 64]. With the publication of new cellular standards 5G New Radio Release 17 [1] and Release 18 [2], industries across areas spanning wildlife conservation [13], healthcare logistics [14], banking [5] to oil and gas [12] have also adopted the cellular LP-WAN alternatives such as NB-IoT [25] as the backhaul. Further academics have studied the impact of leveraging such cellular technologies for backhaul [36, 75]. While cellular backhauls also present an attractive low-latency backhaul opportunity, this work primarily focuses on LoRa due to availability of satellite deployment for evaluation. There are two key issues that particularly hinder enabling connectivity for LoRa satellites in the ISM band – (1) Packets go undetected due to limitations of LoRa ground infrastructure (Figure 3) [17, 29, 46, 64] and (2) Intra-packet Doppler generates

[2]Note that given the evaluation of Vista was performed using trace-driven simulation, the real compute time and memory required will also depend on the size of the image and the model being processed. However, we believe that it will remain significantly low as shown in other works [43].

heterogeneous errors that are atypical of wireless channels. Vista provides new solutions informed by our CubeSats channel traces in the ISM band (Section 3) to overcome these limitations.

CubeSat Communication and Satellite Imaging: There has been much work on development of satellites for various applications spanning agriculture, atmospheric monitoring and remote sensing. On one end, are publicly funded programs such as NOAA, Sentinel, LandSat with large historical data, where academics have recently shown new opportunities [79, 80] in developing novel ground architectures for receiving image data. On the other hand, are private enterprises entering this space such as Planet Labs, Microsoft [85, 86] and Boeing Aerospace. There has been also much work in characterizing the trajectory [33, 78] and channel [61, 71] of these LEO satellites. Further, there has been analysis [63] of how future satellite constellations can provide instantaneous latency of data transfer (caveat to throughput of inter-satellite links). However, inter-satellite links are highly difficult to create and all satellite constellations deployed for earth-observation remain incapable of communicating significant amount of data between the satellites. (instead they act like individual satellites working together to observe the earth in spatially diverse locations). All of these above deployments are big satellites or large CubeSats ($\geq$ 6 unit), and leverage high throughput wideband technologies to communicate to ground infrastructure. These technologies typically require specialized ground infrastructure which leads to less number of base stations and hence larger latencies. Vista attempts to complement these high-bandwidth technologies by leveraging a long-range technology with large terrestrial ground infrastructure. Vista addresses the unfortunate consequences of using this narrowband technology in LEO – the lower data rate as well as high Doppler-to-bandwidth ratio. Similar to L2D2 [86] which leverages existing satellite ground infrastructure for other satellites, Vista enables the ability to leverage existing LoRa ground infrastructure for communicating with LoRa enabled CubeSats.

There has also been recent work [82] that deals with reducing the latency of transmitting high fidelity data from satellites to distributed ground stations. While that work focuses on minimizing the *average overall latency* of downlink transmissions from satellites to distributed ground stations by developing a globally optimal schedule, this work primarily focuses on reducing the *first contact latency* of transmitting a quick and rough estimate of the image for low latency response. We believe LoRa along with Vista can act as a complementary communication mode in special circumstances (such as disaster management scenarios) to these solutions.

Error-Resilient Image Compression over wireless links: Imaging is one of the most important applications for today's CubeSats with usecases in weather monitoring, satellite imaging, cloud detection, agriculture, and so on. [27, 32, 56, 67, 93]. Thus, these images are typically compressed over the wireless link using autoencoders (deep neural networks with narrow waist) to create signature optimized for a set of tasks [74, 88, 92]. These narrow outputs are used to communicate between the sensor and the receiver as they maximize information density for a set of tasks [60]. Unfortunately, if autoencoders are trained unaware of the communication channel, there is large information loss due to discretization, and bit errors [89]. There is of course a lot of work in joint source-channel coding for wireless image transmission [34, 35, 38, 42, 55] which deal with general uniform error distributions seen in typical wireless channels. Further, there have been recent works [43, 44] that show the efficacy of data and application-aware encodings in reducing computational bottlenecks on the satellite. However, operating LoRa (a) technology that encodes data in frequency shifts) in 915 MHz ISM band (large amount of Doppler) leads to a non-uniform distribution of errors across bits as described in Section4.3. In this article, Vista shows how, by making these data and application driven encodings aware of this non-uniform distribution in the LoRa CubeSat context, we can design extremely small encodings robust to an adverse communication channel (Section 4).

9 Conclusion and Future Work

This article presents Vista, a LoRa-based CubeSat communication system that complements typical high-bandwidth satellite communication links to enable low-latency imaging and inference by using existing commodity LoRa ground stations. We develop custom software-only solutions to adapt existing LoRa ground infrastructure to receive packets from CubeSats. We then build a channel-aware image encoding scheme that is informed by the physical-layer channel characteristics of satellite signals. We present a detailed evaluation of Vista by leveraging wireless channel measurements from from a recent CubeSat launched by our team. Our results demonstrate 55.55% lower latency, a 4.56 dB improvement in recovered image quality and the 1.38× improvement in accuracy of land-use classification. We plan to deploy Vista end-to-end on a future CubeSat mission to estimate our system's ability for varied Earth sensing applications.

Acknowledgments

The authors would like to thank the anonymous reviewers for providing insightful reviews and suggestions to improve the manuscript. We would like to thank members of the WiTech and NEWT Lab for their feedback and advice on the manuscript.

References

[1] accessed Jan 01, 2024. 3GPP - Release 17. Retrieved from https://www.3gpp.org/specifications-technologies/releases/release-17
[2] accessed Jan 01, 2024. 3GPP - Release 18. Retrieved from https://www.3gpp.org/specifications-technologies/releases/release-18
[3] accessed Jan 01, 2024. Artemis CubeSat Kit and Curriculum . Retrieved from https://franceszhu.space/artemis-cubesat-kit
[4] accessed Jan 01, 2024. Assessment of TLE-based Orbit Determination and Prediction for Cubesats. Retrieved from https://ntrs.nasa.gov/api/citations/20190004996/downloads/20190004996.pdf
[5] accessed Jan 01, 2024. Banco Santander funds Sateliot with €6 Million through its High- Growth Enterprise Program. Retrieved from https://sateliot.space/en/news-sateliot-space/banco-santander-funds-sateliot-with-e6-million-through-its-high-growth-enterprise-program/
[6] accessed Jan 01, 2024. FossaSat-1, an Open Source Satellite for the Internet of Things. Retrieved from https://www.hackster.io/news/fossasat-1-an-open-source-satellite-for-the-internet-of-things-7f31cab00ef5
[7] accessed Jan 01, 2024. Gossamer Piccolomini (Gossamer 1). Retrieved from https://tinygs.com/satellite/Gossamer
[8] accessed Jan 01, 2024. How neural networks solve the XOR problem. Retrieved from https://towardsdatascience.com/how-neural-networks-solve-the-xor-problem-59763136bdd7
[9] accessed Jan 01, 2024. Lacuna Space. Retrieved from https://lacuna.space/
[10] accessed Jan 01, 2024. Nanosatellites - FOSSA Systems. Retrieved from https://fossa.systems/satellites/
[11] accessed Jan 01, 2024. NOAA Satellites. Retrieved from https://www.noaa.gov/satellites
[12] accessed Jan 01, 2024. Oil and Gas | 5G Satellite Connectivity's Impact on the Oil and Gas Value Chain. Retrieved from https://sateliot.space/en/news-sateliot-space/oil-and-gas-%7C-5g-satellite-connectivity%E2%80%99s-impact-on-the-oil-and-gas-value-chain/
[13] accessed Jan 01, 2024. Sateliot and Endangered Wildlife Trust (EWT) To Revolutionize Wildlife Conservation in Southern Africa with 5G IoT Satellites. Retrieved from https://sateliot.space/en/news-sateliot-space/sateliot-and-endangered-wildlife-trust-(ewt)-to-revolutionize-wildlife-conservation-in-southern-africa-with-5g-iot-satellites/
[14] accessed Jan 01, 2024. Sateliot Joins The VITAS Project to Research and Implement Healthcare Logistics Solutions with Drones. Retrieved from https://sateliot.space/en/news-sateliot-space/sateliot-joins-the-vitas-project-to-research-and-implement-healthcare-logistics-solutions-with-drones//
[15] accessed Jan 01, 2024. SATLLA-2B. Retrieved from https://tinygs.com/satellite/SATLLA-2B
[16] accessed Jan 01, 2024. SatNogs DB. Retrieved from https://db.satnogs.org/
[17] accessed Jan 01, 2024. SemTech SX1262. Retrieved from https://www.semtech.com/products/wireless-rf/lora-core/sx1262
[18] accessed Jan 01, 2024. Sentinel Missions. Retrieved from https://www.esa.int/Applications/Observing_the_Earth/Copernicus/The_Sentinel_missions

[19] accessed Jan 01, 2024. Swarm takes LoRa Sky High. Retrieved from https://spectrum.ieee.org/swarm-takes-lora-skyhigh
[20] accessed Jan 01, 2024. TinyGS, the Open Source Global Satellite Network. Retrieved from https://tinygs.com/.
[21] accessed Jan 01, 2024. What is V-R3x? Retrieved from https://www.nasa.gov/ames/v-r3x
[22] accessed Jan 01, 2024. Why Space Launches are getting cheaper? Retrieved from https://www.facebook.com/watch/?v=3421542817907662
[23] accessed Jan 01, 2024. Wyld announces global LoRa satellite deal. Retrieved from https://wyldnetworks.com/blog/wyld-announces-global-lora-satellite-deal
[24] accessed Jan 10, 2017. The Things Network. Retrieved from https://www.thethingsnetwork.org/
[25] 3GPP. 2018. NarrowBand - Internet of Things. Retrieved from https://www.gsma.com/iot/narrow-band-internet-of-things-nb-iot/. Accessed: 2019-01-30.
[26] Orion Afisiadis, Matthieu Cotting, Andreas Burg, and Alexios Balatsoukas-Stimming. 2019. On the error rate of the LoRa modulation with interference. *IEEE Transactions on Wireless Communications* 19, 2 (2019), 1292–1304.
[27] Shreya Agrawal, Luke Barrington, Carla Bromberg, John Burge, Cenk Gazen, and Jason Hickey. 2019. Machine learning for precipitation nowcasting from radar images. *Machine Learning and the Physical Sciences Workshop at the 33rd Conference on Neural Information Processing Systems (NeurIPS'19)*. 4.
[28] Irfan Ali, Naofal Al-Dhahir, and John E. Hershey. 1998. Doppler characterization for LEO satellites. *IEEE Transactions on Communications* 46, 3 (1998), 309–313.
[29] Thomas Ameloot, Marc Moeneclaey, Patrick Van Torre, and Hendrik Rogier. 2021. Characterizing the impact of doppler effects on body-centric lora links with SDR. *Sensors* 21, 12 (2021), 4049.
[30] Parinya Anantachaisilp, Marut Muangkham, Nuttawat Punpigul, and Mason Thammawichai. 2020. Store and forward CubeSat using LoRa technology and private LoRaWAN-Server. (2020).
[31] Rifat Shaarook and Srimathy Kesan. 2021. Low power LoRa transmission in low earth orbiting satellites. *Intelligent Environments* (2021), 233.
[32] Valentin Bickel, Lukas Mandrake, and Gary Doran. 2021. A labeled image dataset for deep learning-driven rockfall detection on the moon and mars. *Frontiers in Remote Sensing* 2, 2 (2021), 640034. https://www.frontiersin.org/journals/remote-sensing/articles/10.3389/frsen.2021.640034
[33] Nafeesa Bohra, Hermann De Meer, and Aftab A. Memon. 2009. Analysing the orbital movement and trajectory of LEO (Low Earth Orbit) satellite relative to earth rotation. In *International Conference on Personal Satellite Services*. Springer, 1–11.
[34] Eirina Bourtsoulatze, David Burth Kurka, and Deniz Gündüz. 2019. Deep joint source-channel coding for wireless image transmission. *IEEE Transactions on Cognitive Communications and Networking* 5, 3 (2019), 567–579.
[35] Philippe Burlina and Fady Alajaji. 1998. An error resilient scheme for image transmission over noisy channels with memory. *IEEE Transactions on Image Processing* 7, 4 (1998), 593–600.
[36] Gabriel Maiolini Capez, Santiago Henn, Juan A. Fraire, and Roberto Garello. 2022. Sparse satellite constellation design for global and regional direct-to-satellite IoT services. *IEEE Trans. Aerospace Electron. Systems* 58, 5 (2022), 3786–3801.
[37] Supriya Chakrabarti. accessed Jan 01, 2024. How Many Satellites are Orbiting Earth? Retrieved from https://astronomy.com/news/2021/09/how-many-satellites-are-orbiting-earth.
[38] Ramamurti Chandramouli, N. Ranganathan, and Shivaraman J. Ramadoss. 1998. Adaptive quantization and fast error-resilient entropy coding for image transmission. *IEEE Transactions on Circuits and Systems for Video Technology* 8, 4 (1998), 411–421.
[39] Gonglong Chen, Wei Dong, and Jiamei Lv. 2021. LoFi: Enabling 2.4 GHz LoRa and WiFi coexistence by detecting extremely weak signals. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.
[40] Federal Communications Commission. 2021. Federal Communications Commission. Retrieved from https://www.fcc.gov/
[41] Federal Communications Commission. 2021. Radio Spectrum Allocation. Retrieved from https://www.fcc.gov/engineering-technology/policy-and-rules-division/general/radio-spectrum-allocation
[42] Victor Debrunner, Linda DeBrunner, Longji Wang, and Sridhar Radhakrishnan. 2000. Error control and concealment for image transmission. *IEEE Communications Surveys & Tutorials* 3, 1 (2000), 2–9.
[43] Bradley Denby, Krishna Chintalapudi, Ranveer Chandra, Brandon Lucia, and Shadi Noghabi. 2023. Kodan: Addressing the computational bottleneck in space. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*. 392–403.
[44] Bradley Denby and Brandon Lucia. 2020. Orbital edge computing: Nanosatellite constellations as a new class of computer system. In *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems*. 939–954.
[45] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. 2018. Charm: Exploiting geographical diversity through coherent combining in low-power wide-area

networks. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 60–71.

[46] Alexander A. Doroshkin, Alexander M. Zadorozhny, Oleg N. Kus, Vitaliy Yu Prokopyev, and Yuri M. Prokopyev. 2019. Experimental study of LoRa modulation immunity to Doppler effect in CubeSat radio communications. *IEEE Access* 7, 7 (2019), 75721–75731.

[47] Ziad El Khatib, Adel Ben Mnaouer, Sherif Moussa, Mohd Azman Bin Abas, Nor Azman Ismail, Fuad Abdulgaleel, Ibrahim Elmasri, and Loay Ashraf. 2022. LoRa-enabled GPU-based CubeSat yolo object detection with hyperparameter optimization. In *2022 International Symposium on Networks, Computers and Communications (ISNCC)*. IEEE, 1–4.

[48] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. 2017. Empowering low-power wide area networks in urban settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 309–321.

[49] Lara Fernandez, Joan Adria Ruiz-De-Azua, Anna Calveras, and Adriano Camps. 2020. Assessing LoRa for satellite-to-earth communications considering the impact of ionospheric scintillation. *IEEE Access* 8 (2020), 165570–165582.

[50] 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.

[51] 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.

[52] Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan, and Mo Li. 2020. LMAC: Efficient carrier-sense multiple access for lora. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–13.

[53] William H. Guier and George C. Weiffenbach. 1959. *The Doppler determination of orbits*. Technical Report. JOHNS HOPKINS UNIV LAUREL MD APPLIED PHYSICS LAB.

[54] William H. Guier and George C. Weiffenbach. 1998. Genesis of satellite navigation. *Johns Hopkins APL Technical Digest* 19, 1 (1998), 15.

[55] Thomas Guionnet and Christine Guillemot. 2003. Soft decoding and synchronization of arithmetic codes: Application to image transmission over noisy channels. *IEEE Transactions on Image Processing* 12, 12 (2003), 1599–1609.

[56] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12, 7 (2019), 2217–2226.

[57] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. 2019. Netscatter: Enabling large-scale backscatter networks. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*. 271–284.

[58] Max Holliday, Kevin Tracy, Zachary Manchester, and Anh Nguyen. 2022. The V-R3X mission: Towards autonomous networking and navigation for Cubesat swarms. In *The 4S Symposium 2022*.

[59] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2023. Don't miss weak packets: Boosting LoRa reception with antenna diversities. *ACM Transactions on Sensor Networks* 19, 2 (2023), 1–25.

[60] Yihan Jiang, Hyeji Kim, Himanshu Asnani, Sreeram Kannan, Sewoong Oh, and Pramod Viswanath. 2019. Turbo autoencoder: Deep learning based channel codes for point-to-point communication channels. *Advances in Neural Information Processing Systems (NeurIPS'19)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc. Vol 32, 2758–2768. https://proceedings.neurips.cc/paper_files/paper/2019/file/228499b55310264a8ea0e27b6e7c6ab6-Paper.pdf

[61] Joe Khalife, Mohammad Neinavaie, and Zaher M. Kassas. 2020. Blind doppler estimation from LEO satellite signals: A case study with real 5G signals. In *Proceedings of the 33rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2020)*. 3046–3054.

[62] Bryan Klofas, Jason Anderson, and Kyle Leveque. 2008. A survey of cubesat communication systems. In *5th Annual CubeSat Developers' Workshop*. 1–36.

[63] Zeqi Lai, Qian Wu, Hewu Li, Mingyang Lv, and Jianping Wu. 2021. Orbitcast: Exploiting mega-constellations for low-latency earth observation. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. IEEE, 1–12.

[64] Intanon Lapapan, Wetchaphat Pa-in, Vachiravit Chotimanon, and Peeramed Chodkaveekityada. 2021. LoRa multichannel access to doppler effect in cubeSat radio communication. In *Proceedings of the Innovation Aviation & Aerospace Industry - International Conference (IAAI'21)*, Chiang Mai (Ed.). MDPI: Basel, Switzerland, 28–30. DOI: 10.3390/IAAI-2021-10588 https://sciforum.net/paper/view/10588

[65] Abhas Maskey, Pooja Lepcha, Hari Ram Shrestha, Withanage Dulani Chamika, Tharindu Lakmal Dayarathna Malmadayalage, Makiko Kishimoto, Yuta Kakimoto, Yuji Sasaki, Turtogtokh Tumenjargal, George Maeda, et al. 2022. One

year on-orbit results of improved bus, LoRa demonstration and novel backplane mission of a 1U CubeSat constellation. *Transactions of the Japan Society for Aeronautical and Space Sciences* 65, 5 (2022), 213–220.

[66] Syed Zafar Abbas Mehdi, Aiffah Mohd Ali, and Safiah Zulkifli. 2023. LoRaWAN CubeSat with an adaptive data rate: An experimental analysis of path loss link margin. *Aerospace* 10, 1 (2023), 53.

[67] Mohammad Modava and Gholamreza Akbarizadeh. 2017. Coastline extraction from SAR images using spatial fuzzy clustering and the active contour method. *International Journal of Remote Sensing* 38, 2 (2017), 355–370.

[68] Matteo Nardello, Harsh Desai, Davide Brunelli, and Brandon Lucia. 2019. Camaroptera: A batteryless long-range remote visual sensing system. In *Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems*. 8–14.

[69] National Telecommunications and Information Administration. 2021. FCC, NTIA Name Staff Representatives to Advisory Committees to Further Technical Collaboration. Retrieved from https://www.ntia.doc.gov/press-release/2022/fcc-ntia-name-staff-representatives-advisory-committees-further-technical.

[70] National Telecommunications and Information Administration. 2021. National Telecommunications and Information Administration. Retrieved from https://www.ntia.doc.gov/

[71] William P. Osborne and Yongjun Xie. 1999. Propagation characterization of LEO/MEO satellite systems at 900-2100 MHz. In *1999 IEEE Emerging Technologies Symposium. Wireless Communications and Systems (IEEE Cat. No. 99EX297)*. IEEE, 21–1.

[72] Maria Rita Palattella and Nicola Accettura. 2018. Enabling internet of everything everywhere: LPWAN with satellite backhaul. In *2018 Global Information Infrastructure and Networking Symposium (GIIS)*. IEEE, 1–5.

[73] V. Yu Prokopyev, S. S. Bakanov, V. K. Bodrov, E. N. Chernodarov, A. A. Doroshkin, V. N. Gorev, A. Yu Kolesnikova, A. S. Kozlov, O. N. Kus, A. V. Melkov, A. A. Mitrokhin, A. A. Morsin, A. E. Nazarenko, I. V. Neskorodev, A. V. Pelemeshko, Yu M. Prokopyev, D. A. Romanov, A. M. Shilov, M. V. Shirokih, A. A. Sidorchuk, A. S. Styuf, and A. M. Zadorozhny. 2021. NORBY CubeSat nanosatellite: design challenges and the first flight data. *Journal of Physics: Conference Series* 1867, 1 (2021), 012038. https://dx.doi.org/10.1088/1742-6596/1867/1/012038

[74] Seunghyoung Ryu, Hyungeun Choi, Hyoseop Lee, Hongseok Kim, and Vincent W. S. Wong. 2018. Residential load profile clustering via deep convolutional autoencoder. In *2018 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 1–6.

[75] Giancarlo Sciddurlo, Antonio Petrosino, Mattia Quadrini, Cesare Roseti, Domenico Striccoli, Francesco Zampognaro, Michele Luglio, Stefano Perticaroli, Antonio Mosca, Francesco Lombardi, et al. 2021. Looking at NB-IoT over LEO satellite systems: Design and evaluation of a service-oriented solution. *IEEE Internet of Things Journal* 9, 16 (2021), 14952–14964.

[76] Semtech. 2016. SX1301 datasheet. (2016). https://semtech.my.salesforce.com/sfc/p/E0000000JelG/a/44000000MDnR/Et1KWLCuNDI6MDagfSPAvqqp.Y869Flgs1LleWyfjDY

[77] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent interference cancellation: Decoding multi-packet collisions in LoRa. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 503–515.

[78] Xi Shen, Defeng David Huang, Boming Song, Claire Vincent, and Roberto Togneri. 2019. 3-D tomographic reconstruction of rain field using microwave signals from LEO satellites: Principle and simulation results. *IEEE Transactions on Geoscience and Remote Sensing* 57, 8 (2019), 5434–5446.

[79] Vaibhav Singh, Akarsh Prabhakara, Diana Zhang, Osman Yağan, and Swarun Kumar. 2021. A community-driven approach to democratize access to satellite ground stations. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 1–14.

[80] Vaibhav Singh, Osman Yağan, and Swarun Kumar. 2022. SelfieStick: Towards earth imaging from a low-cost ground module using LEO satellites. In *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*.

[81] Kai Sun, Zhimeng Yin, Weiwei Chen, Shuai Wang, Zeyu Zhang, and Tian He. 2021. Partial symbol recovery for interference resilience in low-power wide area networks. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*, IEEE, 1–11.

[82] Bill Tao, Maleeha Masood, Indranil Gupta, and Deepak Vasisht. 2023. Transmitting, fast and slow: Scheduling satellite traffic through space and time. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*. 1–15.

[83] The Things Network. 2020. LoRa World Record Broken: 832km/517mi using 25mW. Retrieved from https://www.thethingsnetwork.org/article/lorawan-world-record-broken-twice-in-single-experiment-1

[84] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. 2023. CoLoRa: Enabling multi-packet reception in LoRa networks. *IEEE Transactions on Mobile Computing* 22, 6 (2023), 3224–3240. DOI: https://doi.org/10.1109/TMC.2021.3138495

[85] Deepak Vasisht and Ranveer Chandra. 2020. A distributed and hybrid ground station network for low earth orbit satellites. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 190–196.

[86] Deepak Vasisht, Jayanth Shenoy, and Ranveer Chandra. 2021. L2D2: Low latency distributed downlink for LEO satellites. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 151–164.

[87] Thiemo Voigt, Martin Bor, Utz Roedig, and Juan Alonso. 2017. Mitigating inter-network interference in LoRa networks. In *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks*. 323–328.

[88] Haoyu Yang, Piyush Pathak, Frank Gennari, Ya-Chieh Lai, and Bei Yu. 2019. DeePattern: Layout pattern generation with transforming convolutional auto-encoder. In *Proceedings of the 56th Annual Design Automation Conference 2019*. 1–6.

[89] Lita Yang, Daniel Bankman, Bert Moons, Marian Verhelst, and Boris Murmann. 2018. Bit error tolerance of a CIFAR-10 binarized convolutional neural network processor. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 1–5.

[90] 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.

[91] Alexander M. Zadorozhny, Alexander A. Doroshkin, Vasily N. Gorev, Alexander V. Melkov, Anton A. Mitrokhin, Vitaliy Yu. Prokopyev, and Yuri M. Prokopyev. 2022. First flight-testing of LoRa modulation in satellite radio communications in low-earth Orbit. *IEEE Access* 10 (2022), 100006–100023. DOI: https://doi.org/10.1109/ACCESS.2022.3207762

[92] Lingjun Zhao, Huakun Huang, Xiang Li, Shuxue Ding, Haoli Zhao, and Zhaoyang Han. 2019. An accurate and robust approach of device-free localization with convolutional autoencoder. *IEEE Internet of Things Journal* 6, 3 (2019), 5825–5840.

[93] Mingmin Zhao, Peder Olsen, and Ranveer Chandra. 2023. Seeing through clouds in satellite images. *IEEE Transactions on Geoscience and Remote Sensing* 61 (2023), 1–16.

Received 7 September 2023; revised 23 February 2024; accepted 27 May 2024