

A Community-Driven Approach to Democratize Access to Satellite Ground Stations

Vaibhav Singh
Carnegie Mellon University
vaibhav3@andrew.cmu.edu

Akarsh Prabhakara
Carnegie Mellon University
aprabhak@andrew.cmu.edu

Diana Zhang
Carnegie Mellon University
dianaz1@andrew.cmu.edu

Osman Yağın
Carnegie Mellon University
oyagan@andrew.cmu.edu

Swarun Kumar
Carnegie Mellon University
swarun@cmu.edu

ABSTRACT

Should you decide to launch a nano-satellite today in Low-Earth Orbit (LEO), the cost of renting ground station communication infrastructure is likely to significantly exceed your launch costs. While space launch costs have lowered significantly with innovative launch vehicles, private players, and smaller payloads, access to ground infrastructure remains a luxury. This is especially true for smaller LEO satellites that are only visible at any location for a few tens of minutes a day and whose signals are extremely weak, necessitating bulky and expensive ground station infrastructure.

In this paper, we present a community-driven distributed reception paradigm for LEO satellite signals where signals received on many tiny handheld receivers (not necessarily deployed on rooftops but also indoors) are coherently combined to recover the desired signal. This is made possible by employing new synchronization and receiver orientation techniques that study satellite trajectories and leverage the presence of other ambient signals. We compare our results with a large commercial receiver deployed on a rooftop and show a 8 dB SNR increase both indoors and outdoors using 8 receivers, costing \$38 per RF frontend.

CCS CONCEPTS

• **Hardware** → **Signal processing systems; Wireless devices; • Computer systems organization** → **Cloud computing; Embedded and cyber-physical systems; • Networks** → **Wireless access points, base stations and infrastructure.**

KEYWORDS

LEO Satellites, Ground stations, Synchronization, Coherent Combination, low-cost, Distributed Systems

ACM Reference Format:

Vaibhav Singh, Akarsh Prabhakara, Diana Zhang, Osman Yağın, and Swarun Kumar. 2021. A Community-Driven Approach to Democratize Access to Satellite Ground Stations. In *The 27th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '21)*, October 25–29, 2021, New Orleans, LA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3447993.3448630>



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

ACM MobiCom '21, October 25–29, 2021, New Orleans, LA, USA
ACM ISBN 978-1-4503-8342-4/21/10.
<https://doi.org/10.1145/3447993.3448630>

1 INTRODUCTION

Today, if you launch a small satellite in space, you are likely able to communicate with it from the ground station only for about ten minutes *per day* [4, 5, 58]. Further and perhaps paradoxically, the rent one pays on today's satellite ground stations [6, 38, 85] for every minute of access could quickly add up to rival the actual launch costs of a small satellite (with recent nano-satellites costing as little as \$2000 to launch). Building or buying your own base station – including the few existing open-source designs – is an even larger investment [71, 89, 95]. These cost and communication bottlenecks mean that very little useful data can be retrieved from small satellites – often less than a few hundred kilobytes per day. More problematically, should the nanosatellite be unresponsive over the ten minute window or miss sending relevant data, an entire day's worth of data is lost. Given that many small satellite missions are only designed to last a few months [52] with mechanical or radiation-related failures being a constant threat, every day's data is precious. In other words, wireless communication to the ground remains the key bottleneck in ensuring the success of the small satellite revolution – and a problem that we in the MobiCom community are best suited to address.

There have been several admirable efforts that have sought to reduce or open source ground station infrastructure. Amazon's AWS ground stations [7], Stellarstation [2] and Planet Labs [69] are recent efforts that allow users to rent time on base stations at extremely fine granularity, providing cost savings. Various open source implementations, e.g., Satnogs [30] have built frameworks that simplify deploying your own base station. However, across these efforts a fundamental challenge remains – the cost of individual base stations is high [71, 89, 95], and these need to be placed in optimal pre-surveyed locations (e.g., rooftops) with significant installation and maintenance cost – both of which limit the scale and add to the cost of the network.

In this paper, we take the first steps towards addressing this problem by building Quasar, a satellite ground station that costs a few tens of dollars and can be placed anywhere (including indoors). Our system is designed so that while each single base station is unlikely to pick up even the strongest satellite signal, teams of wimpy base stations can collectively pool their data at the cloud to coherently combine and decode them. While there is rich literature on distributed coherent combining in wireless LANs, cellular networks, low-power IoT and even radio observatories [19], we believe the Low Earth Orbit (LEO) satellite context and our per-base station price point brings forth unique system design challenges that we

address. We implemented and evaluated our system by distributing Quasar’s inexpensive ground stations over a large U.S. city and obtained a distributed image of the city (and surrounding regions) captured from LEO satellites overhead in space that was not decodable (or even detectable) individually at any single ground station. Our satellite hardware system design will be made fully open source and has an associated online portal for contributing satellite data to create the framework for a community-driven satellite ground infrastructure.

The rest of this paper describes our solutions to the key-challenges in building a low-cost and distributed community-run satellite ground station infrastructure. We focus on three specific challenges:

(1) Low-Cost vs. Synchronization of Base Stations: To perform coherent combining accurately, ground stations must be perfectly synchronized in time, frequency and phase. Yet, precise and accurate clocks to synchronize base stations are expensive and often require deployment outdoors (e.g., GPS clocks, roof rent), which goes against Quasar’s accessible spirit and low-cost approach. Quasar addresses this challenge by using signals-of-opportunity from unrelated ground devices like amateur radio towers as effective beacons to synchronize distributed ground stations. We further show how we can use the known trajectory and velocity of satellites over time to account for Doppler shifts in frequency.

(2) Guidelines for Antenna Deployment: An important drawback of being able to deploy a node in any environment is significant signal attenuation, particularly in dense urban and indoor settings. Indeed, this is precisely why satellite ground station sites are typically chosen to ensure a clean line-of-sight path to satellites overhead. Quasar addresses this challenge by providing general guidelines on optimal indoor placement and orientation of nodes and their antennas, accounting for factors including cardinal directions, elevation, and proximity to adjacent surfaces using observations of trajectory and measured attenuation from satellite reception in different indoor scenarios.

(3) Noise Resilience vs. Limited Backhaul: To be truly resilient to noise, one would have to coherently combine signals from as many ground stations as possible – including those where signals are deeply buried under noise to the point of not even being detected. Yet, this means that ground stations would have to constantly stream all received signals to the base stations, even if they appear like complete noise. This is a significant strain on uplink bandwidth (particularly when this uplink is home broadband) and can disincentivize system adoption through a community-driven approach. To mitigate this challenge, Quasar uses the known trajectory of satellites as well as cues from ground stations that do receive less noisy versions of the satellite’s signal to decide, which of noisy received signals to push to the cloud, while other more noisy receptions can be locally discarded. We show that the resulting approach requires a very modest fraction of the total uplink bandwidth in a typical home broadband connection.

We implement and evaluate Quasar on an RTL-SDR (costing \$25) that can be connected to a Raspberry PI or the user’s laptop. These are connected to EXS136SMI Laird technologies antennas [28] allowing for a form factor of 9.95 by 1.20 centimeters and not requiring specialized installation. The Quasar toolkit is developed

in Matlab and Python with an associated visualization tool available via a web portal designed using ReactJS. All source code and hardware designs developed through this project have been made open source to encourage community adoption and can be accessed at <https://www.witechlab.com/quasar.html>. Our results reveal the following:

- Quasar achieves an SNR gain of 8 dB by aggregating signals at 8 ground stations, compared to the signals at individual ground stations.
- Quasar conserves backhaul bandwidth by 33% at the average user-deployed ground station by predicting when they need to relay data up-stream.
- Quasar achieves an average structural similarity index of 0.85 outdoors and 0.74 indoors in recovering weather satellite images using 8 ground stations; an improvement from 0.52 and 0.35 respectively, for a single ground station

Contributions: This paper’s technical contributions are:

- The design of an open-source, distributed and community-driven satellite ground infrastructure that effectively decodes LEO satellite signals from individually low-cost gateways that can be deployed anywhere.
- An approach to perform efficient synchronization, configuration and coherent combining of signals from low-cost gateways to recover signals individually undecodable at any single gateway.
- An end-to-end system evaluation demonstrating efficient satellite data decoding, with minimal strain on backhaul bandwidth.

2 RELATED WORK

Satellite Ground Infrastructure: Satellite signal access and specialized ground stations have been well studied since the time the first satellites were sent to space. Different designs have been explored based on specific applications [29, 54]. Prior work has also explored receive beamforming using large arrays with applications to radio astronomy [16, 39, 46] – different from our work. Recently, many non-government entities, e.g., universities, launching their own satellites and subsequently building a ground station [1, 21, 82], ranging about \$6000-10000 and countless volunteer student hours, or creating networks of receivers that have limited participation [50, 75, 78]. This effort and cost has limited the deployment of ground stations to only a few amateur radio enthusiasts who are willing to deploy huge antennas on their rooftops.

Recently there has been increasing interest in both academia [77, 90] and industry [2, 7, 55, 68, 69, 86] to design ground station network from which users can simultaneously access multi-satellite data using a subscription model. While these efforts have improved satellite data access by reducing latency using innovations in the cloud and ground station network, we differ in our approach where we try to receive single satellite data using distributed receivers. SatNogs[30] allows users to schedule receptions after deploying one’s own ground stations as well as access to processed data from other users; however, it doesn’t provide raw I/Q samples for processing. There have been many efforts towards miniaturizing parts of the ground station hardware [49, 66, 74, 76]. These works reduce




Figure 1: LEO satellites (a) operate in low altitudes and hence offer only a small angular view of the Earth, and (b) operate in polar orbits and complete one rotation in just 100 min.

the cost by innovating on the SDR or rotator front, but size of antenna and space required still remains the bottleneck. We strive to fill a gap in this space: cheaper receivers that do not require roof access and whose data can be shared among users. We argue that providing this will lower the barrier of entry to satellite receivers and encourage wider spread networks of ground stations, particularly in under-served geographical regions.

Satellites in Mobile Computing: In the mobile computing context, we've seen a budding interest in space, for applications including lower-power GPS solutions for ground and satellites [14, 57, 60], planetary network testbeds [10], and to supplement other sensing modalities [12, 51] as well as recent surge in interest in enabling satellite Internet [20, 37, 87] and satellite IoT [23, 63]. Since latency, bandwidth and inter-satellite links are the bottlenecks in enabling satellite internet, initial proposals in [13, 23, 36] have tried to suggest some approaches to overcome these. We distinguish ourselves from this body of work in that we are focused on applying mobile computing concepts to ground stations for LEO satellites.

Coherent combining in Wireless Networks: Enabling synchronous reception for wireless distributed systems have also been extensively studied [17, 26, 32, 53, 62, 79, 84] showing promising techniques for enhancing communication as well as sensing performance. Coherent combining has also been integral to varied wireless systems for enhanced communication [26, 32, 33, 70, 83], location tracking [31, 43, 44, 61, 92] and sensing applications [22, 31, 42, 99, 100]. Designs for synchronization of low cost receivers using GPS and custom hardware [3, 25, 35, 45, 48] have been studied, but we aim to reduce the complexity of the receiver by not adding additional hardware which goes against the spirit of our system. Coherent combining has also been used for enabling optical communication [11, 34]. In contrast to this body of work, our objective is to build on this work in the LEO satellite context.

3 PRIMER & MOTIVATION

In this section, we describe briefly why satellite base station infrastructure is expensive, particularly for LEO small satellites. We note that we focus on LEO satellites because they offer the cheapest launch costs for nano-satellites and form the vast majority of satellites launched by small players (hobbyists, schools, etc.).

LEO satellites offer a small angular view of the Earth: This is due to their proximity to the Earth – typically a few hundred kilometers away from its surface [4, 5, 97] (about 160 - 2000 km). While orbital parameters vary, this means that it takes approximately

15 LEO satellites to form a constellation that can cover the earth, compared to only 3 in MEO [15] (e.g. GPS). That is, a LEO satellite covers about 34 million sq.km. or approximately 7% of the Earth's total surface area. While governments, defense forces and large enterprises can afford to rent or install base stations that span the Earth for LEO, this becomes an expensive proposition for hobbyists, educational institutions or small firms launching cubesats.

LEO satellites are only overhead for a short time: A second related challenge that LEO satellites face is that they are visible over any given geographical location only for about a few tens of minutes per day [4, 5, 58, 97]. This is both due to the small angular view of the satellites as explained above and their high speeds – they revolve around the Earth in about 60-100 minutes typically. Note that orbital parameters vary every revolution around the Earth, so they often do not visit the same location of the Earth more than a few times a day. This means that should a satellite data not be received for any reason (e.g. faults) over these time windows, all data for the day is completely lost. Unfortunately, most hobbyists rent ground stations at a very small number of locations (typically just one [6, 38, 85]), due to high rent costs.

Nano-Satellites live short lives: The life of a nano-satellite is notoriously hard to predict given that it can be knocked offline due to multiple reasons such as radiation-related failure, battery failure or orbital decay. On aggregate however, small satellites live much shorter lives than large satellites at LEO (e.g. the International Space Station), given that satellite weight directly influences orbital decay. It is common for cubesats to last as little as days to as long as few years, but most typically a few months [52]. This means that most hobbyists only obtain about few tens to a few hundreds of days worth of communication traces from a cubesat before it goes offline. This makes it all the more critical to lower the cost and availability of satellite-to-ground communication links – there are only so many chances one gets to communicate before a satellite goes offline, and exactly when this might happen is hard to predict.

4 QUASAR DESIGN AND CHALLENGES

In this section, we discuss why we chose our particular design of coherent combining across low-cost ground stations. We follow up by explaining the research challenges in designing such a system and present an outline for the rest of the paper.

4.1 Design – Why Coherent Combining?

Quasar's approach to lower the cost of access to space places low-cost gateways that cost a few tens of dollars indoors or outdoors. We envision a community-operated design where satellite enthusiasts can place inexpensive receivers, wherever they have space available within their homes, with guidelines provided on what locations will be optimal. Signals from these gateways are then coherently combined at the cloud to recover weak signals that may not be individually recoverable at any single gateway.

Why not engineer cheaper ground stations? One might wonder why we opted for the coherent combining approach rather than geo-replicating existing ground station technology or engineer cheaper ground station hardware. Our choice stems from the fact that a major bottleneck in the cost of ground stations today is installation [71, 89, 94, 95]. Today, satellite ground stations are




Figure 2: System Overview: Quasar comprises of: (1) Synchronizing Receivers (2) Guidelines for placement of receivers (3) Coherent Combining in the cloud

bulky and deployed on specially chosen locations, typically on rooftops or hills. Installing a new base station at a carefully chosen location incurs a cost that might often rival hardware components.

Our choice: Instead, our approach opts for the opposite “deploy where you can” approach, allowing receivers to be placed indoors or outdoors and relies on inexpensive hardware with a small form factor. The result is significant improvements in ease of deployment, with the flip side being that the ground stations are unlikely to receive strong satellite signals. We address this problem through coherent combining – adding up weak receptions across ground stations that are geo-distributed to recover the satellite’s signal.

4.2 Challenges and System Outline

The rest of this paper discusses three key challenges in Quasar’s design. We illustrate how each of these challenges interplay to contribute to Quasar’s architecture in Fig. 2.

Low-Cost Synchronization: Given that our ground stations are inexpensive and often indoors, they lack the ability to natively receive time, frequency and phase synchronized signals across geo-distributed receiving devices. Sec. 5 discusses our approach to address this challenge that relies on signals-of-opportunity from other ground devices (e.g. Amateur Radio devices) to synchronize.

Optimal Ground Station Placement: While our system allows for ground stations to be deployed indoors or outdoors, we provide general guidelines to satellite enthusiasts on where ground stations should be deployed for optimal performance. Sec. 6 describes our analytical and data-driven approach.

Minimizing Strain on Backhaul: Having collected synchronized signals from distributed ground stations, we now need to collect satellite receptions without overwhelming backhaul bandwidths of home broadband connections. Sec. 7 discusses our approach to do this and to perform coherent combining effectively at the cloud.

We note that these techniques and engineering guidelines are designed keeping in mind the unique challenges in the context of distributed satellite reception while ensuring that the system can scale seamlessly with increasing number of satellites and users.

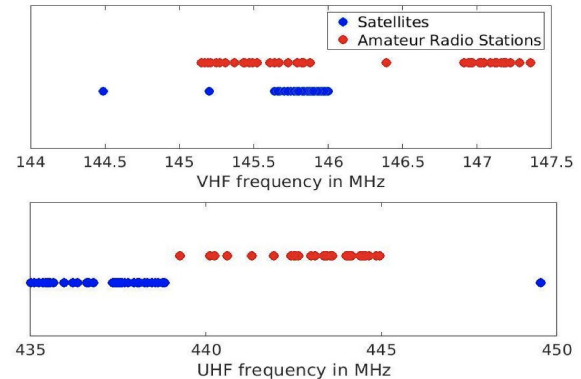


Figure 3: Overlap between LEO satellites and Amateur radio towers and repeaters in Pittsburgh city.

5 SYNCHRONIZATION

Coherent combining is essentially adding signals in a phase-coherent manner across multiple receiving ground stations to boost signal to noise ratio. This is why synchronizing distributed receivers is a key component of any coherent combining system. The need for this is further accentuated in our context because our system comprises several inexpensive radios serving as ground stations. Owing to the low cost of these radios, their hardware components are not designed to provide frequency stability for a long time [73] – particularly the long reception times of ten minutes, which is typical for Low Earth Orbit satellites. Moreover, because our radios are distributed over a wide geographical area, techniques that involve sharing a frequency stable clock are not a viable solution. To add to this, many of our radios are deployed indoors with weak or no access to popular time and frequency sharing systems like GPS.

In addition to these factors, there are certain challenges unique to the LEO satellite receiver context that makes the use of prior distributed receiver system techniques ineffective. We need to keep in mind that unlike most ground-based distributed systems, (1) the transmitter in the case of LEO satellites is almost never under the control of the users (except for the entity or organization owning the satellite). This means that receivers deployed on ground by general public cannot control and coordinate with satellite transmitters to aid synchronization by incorporating smart preamble design, Forward Error Correction, etc. (2) The high orbital speed of LEO satellites introduces frequency offsets called Doppler shifts of greater magnitude than seen in ground based systems. (3) The signals received from satellites are extremely weak due to the high path loss incurred across 1000s of kilometer separation. Therefore, traditional WiFi inspired phase-based synchronization techniques will not work because the phase of the received signal itself is noisy due to low SNR. An important difference here from GPS signals is that we target generic LEO satellites whose transmit signals might not be designed specifically for synchronization and localization.

Keeping these factors in mind, the design of our distributed receiver system should be such that it works without tweaking the satellite signals, can work even in low SNR scenarios and finally, should not impose additional burden on the users. In designing a solution to operate with these constraints, one can intuitively think

of using beacon systems akin to GPS which can provide strong signals indoors to lock on to. We note that many satellites today including Earth observation, weather monitoring, amateur and educational satellites operate in or near amateur radio bands in VHF (144-148 MHz) and UHF (420-450 MHz) amateur frequency bands [8, 59, 64]. Fig. 3 plots the LEO satellites [91] and Amateur radio stations operating in the VHF and UHF bands in Pittsburgh city region[18, 47]. It can be seen that the operating frequencies of satellite and amateur radio stations significantly overlap in the VHF band and lie in close proximity in the UHF band. Because of their low frequency of operation, amateur radio signals can travel long distances and be received indoors. Given their proximity to many LEO satellites' frequency of operation, it is lucrative to use these signals for synchronization inspired by other systems which use separate transmitters for synchronization in MIMO [79] context. We refer to these signals as signals of opportunity. The advantage of using signals of opportunity is that they are already available, thus preventing the need to use additional resources by the users for synchronization. In areas where these amateur radio signals are not available, we propose one of the users to produce an artificial signal of opportunity by installing a transmitter to serve the vicinity. When the received SNR is not poor, we can use preamble and frequency synchronization headers from the received satellite packet itself.

5.1 An Analysis of Hardware Offsets

Before describing our techniques in detail, let us first understand the different offsets which need to be synchronized especially in our context. To characterize these offsets succinctly, we develop a narrow band signal model as follows. Consider a receiver labelled as i . These receivers suffer three broad classes of offsets: (1) *Frequency Offsets*: when their carrier frequency (f_i) is offset from that of the satellite (f); (2) *Phase Offsets*: when the locking of their phase locked loops result in a phase (ϕ_i) that is offset from the satellite (ϕ_s); and (3) *Sampling Frequency Offsets*: when their sampling frequency (t_i) is offset from the satellite (t_s).

A challenge unique to satellites is introduced by the high orbital velocities of LEO satellites. Their motion causes significant Doppler shift (f_{d_i}) which can reach around 500-1000 Hz for a minute of reception. Mathematically, our channel model can be written as:

$$h_{s \rightarrow i} = \hat{h}_{s \rightarrow i} e^{j(2\pi f(t_i - t_s) + 2\pi(f_i + f_{d_i} - f)t + (\phi_i - \phi_s))} \quad (1)$$

Where $\hat{h}_{s \rightarrow i}$ is the true wireless channel in the absence of any hardware offsets and $h_{s \rightarrow i}$ is the measured wireless channel. The remainder of this section describes two solutions to nullify these offsets. Remember, our goal is to synchronize different receivers with each other, not necessarily synchronize each receiver to the transmitter. We use this approach in our solution using signals of opportunity. In the absence of any such signals, we synchronize each receiver to the transmitting satellite.

5.2 Using Signals of Opportunity

While receiving satellite signals, radios inevitably simultaneously listen to any signal of opportunity that falls within their bandwidth. As described earlier, the decision to use signals of opportunity for synchronization is driven by the following two factors: (1) lack of control of satellite transmitters and low SNR of received signals, and (2) overlap of satellite transmissions with amateur radio bands.




Figure 4: Quasar uses signals from nearby amateur bands opportunistically to synchronize different satellite receivers.

Our approach is to leverage channel measurements from frequency bands corresponding to the satellite and the signal of opportunity. Using these measurements, we synthesize an effective channel between each pair of receivers that is free of all time varying offsets as follows.

First, we deal with Doppler shifts using local measurements at each receiver. Each receiver experiences a different Doppler shift based on its geographical position and the satellite trajectory under consideration. Simplified Perturbation Models [40] - called SGP4 - model satellite trajectory mathematically and report Doppler velocities for a satellite with respect to a point on Earth at any time. Using the Doppler shift values from the SGP4 model, we can nullify the effect of Doppler for i^{th} receiver by multiplying measured $h_{s \rightarrow i}$ with $e^{-j2\pi \hat{f}_{d_i} t}$, where \hat{f}_{d_i} is the Doppler shift predicted by the model.[24, 67, 80, 96] report the accuracy of the SGP4 model for orbit determination with errors of the order of 1 km. However, Doppler shift for any satellite is calculated using the rate of change of relative range of satellite from earth based receiver (also called the range rate). Despite the large error in position, the error in range rate is not high due to the slow growth in position error over time, resulting in less than 150 Hz average doppler shift prediction error for VHF and UHF bands [98]. This error in Doppler estimation results in a small value of frequency offset that can be easily accounted for using phase correction on per packet basis using the preamble. Let the Doppler free channel after this be $\tilde{h}_{s \rightarrow i}$.

To remove other offsets, we make use of measurements from a signal of opportunity. Similar to Equation 1, the channel from a signal of opportunity involves all the offsets other than Doppler. We consider the case when the center frequency of signals of opportunity is (f_o) and that of satellite (f).Based on our model for the hardware offsets in Sec. 5.1 above, the channel model for receiving a signal of opportunity at the i^{th} receiver is:

$$h_{o \rightarrow i} = \hat{h}_{o \rightarrow i} e^{j(\theta_{o \rightarrow i} + 2\pi f_o(t_i - t_o) + 2\pi(f_i - f_o)t + (\phi_i - \phi_o))} \quad (2)$$

Our synchronization objective at this point can simply be thought of as removing, step by step, the various offset terms and obtaining a value that is only dependent on the \hat{h} terms. To do so, we realize that we can synchronize each pair of receivers (i, k) at any time. As a starting point, we first subtract channel phases of i and k for satellite and signal of opportunity separately using $\tilde{h}_{s \rightarrow i} \tilde{h}_{s \rightarrow k}^*$

and $h_{o \rightarrow i} \hat{h}_{o \rightarrow k}^*$ respectively. The idea here is to remove offset terms that correspond to the satellite and the source of the signal of opportunity individually that inherently appear in a similar fashion across both receivers from the same source. These terms individually compute to:

$$\begin{aligned} \tilde{h}_{s \rightarrow i} \tilde{h}_{s \rightarrow k}^* &= \hat{h}_{s \rightarrow i} \hat{h}_{s \rightarrow k}^* e^{j2\pi f(t_i - t_k) + 2\pi(f_i - f_k)t + \phi_i - \phi_k} \\ \tilde{h}_{o \rightarrow i} \tilde{h}_{o \rightarrow k}^* &= \hat{h}_{o \rightarrow i} \hat{h}_{o \rightarrow k}^* e^{j2\pi f_o(t_i - t_k) + 2\pi(f_i - f_k)t + \phi_i - \phi_k} \end{aligned} \quad (3)$$

(.)^{*} denotes complex conjugate. It is easy to see that the above two terms can be broken down into two components – offset free-channels (indicated by the \hat{h} terms) and hardware-offset related phases. Fortunately, the hardware-offset related terms are identical except for the center frequency of the satellite and the signal of opportunity transmitter, which tells us that we can get rid of most of them by dividing these two terms or equivalently, subtracting their phases as shown below:

$$h_{ik} = \tilde{h}_{s \rightarrow i} \tilde{h}_{s \rightarrow k}^* \hat{h}_{o \rightarrow i} \hat{h}_{o \rightarrow k}^* \quad (4)$$

The resulting term that we call h_{ik} is free of frequency and phase offsets of between receivers i and k but has a remaining phase offset term $2\pi(f - f_o)(t_i - t_k)$ as shown below:

$$h_{ik} = \hat{h}_{s \rightarrow i} \hat{h}_{s \rightarrow k}^* \hat{h}_{o \rightarrow i} \hat{h}_{o \rightarrow k}^* e^{j2\pi(f - f_o)(t_i - t_k)} \quad (5)$$

This term characterizes the frequency offset between the satellite and signal of opportunities transmitter as seen at the receivers i and k , arising due to different operating frequencies. Since both the satellite and signal of opportunity are not under the user's control, traditional distributed MIMO [72] and cellular [79] approaches cannot be used where the reference base station (in this case the signal of opportunity) transmits for a very short interval just before client transmission (in this case the satellite) on the same frequency to estimate the phase and cancel the additional offset term.

To remove this offset, we need to estimate the phase of the channel at each receiver at the same time and frequency. Inspired by [32], we extrapolate the averaged phase of the satellite and the signal of opportunity channel using piecewise cubic spline extrapolation to estimate the phase difference at the guardband between the two frequencies at each receiver. The difference in the extrapolated phase at the guardband corresponds to the term $2\pi(f - f_o)t_i$ for receiver i . The same process is repeated for receiver k and the difference of the 2 terms accounts for the additional term $2\pi(f - f_o)(t_i - t_k)$. Extrapolation works because the (1) Doppler spread leading to the fast fading channel is already accounted for using the satellite's trajectory in Doppler correction step and (2) similar narrowband transmission characteristics between satellite and LP-WANS [32] in the VHF and UHF bands. We note that the above formulation completely accounts for satellite signal channels (captured within the \hat{h} terms) as well as all hardware-related offsets. This approach also extends to any pair of receivers and can be replicated prior to coherent combination (described in Sec. 7).

5.3 Synchronization Using Satellite Signal

In the absence of signals of opportunity, we make use of the preamble and frequency synchronization markers which are already present in satellite packets. Unlike the previous approach, here we are synchronizing each receiver to the transmitted satellite signal. This approach requires that the signal strength of the received




Figure 5: SNR gain achieved using Signals of Opportunity (SOP) vs Preamble for synchronization.

satellite packet be sufficiently high so that the preamble can be detected as well as the phase of the signal in time and frequency is recoverable. We note that in low SNR scenarios there are situations when the data in the packet is undecodable, even if the preamble is detectable. This is because the preambles are long and structured sequences that can be detected at lower SNRs with correlation compared to data payload which can be random and buried under noise. This results in incorrect phase estimation for frequency and timing synchronization. Hence, a minimum SNR is needed to ensure synchronization with satellite signals. In cases where the individual SNR is low, our only option is to leverage signals of opportunity as in Sec. 5.2 above, while our approach in this section addresses the other cases where the SNR is above a minimum threshold.

With the received signal, our objective is to identify the different offsets in Equation 1. The carrier frequency offset and the Doppler shift manifest as a single term that linearly change the slope of the phase over time. By estimating the slope, the sum of these two terms is estimated. Similarly, the sampling frequency offset creates a slope over frequency when the signal is viewed in the frequency domain. Estimating slope in the frequency domain, we get the sampling frequency offset. Other offsets such as phase offsets, do not evolve over time and can be estimated using a simple one tap equalizer. Fig.5 compares the analytical SNR gain of the system after coherent combination (Sec.7.1) using the two synchronization methods described in Sec.5.2 and Sec.5.3 for 6 receivers at different minimum individual SNRs by decreasing the magnitude of AWGN noise in satellite signal. It can be seen that both the methods show similar trends in SNR gain, however, the method using signals of opportunity starts yielding gains at lower received SNR because it is less susceptible to errors in phase of satellite signal.

6 GUIDELINES FOR ANTENNA PLACEMENT

Flexibility in placement of ground stations is critical for adoption of Quasar. A significant barrier to entry for deploying a base station for the average consumer, beyond just expense, is access to a roof. Quasar allows deploying base stations at a much lower cost inside the home. When considering optimal placement of base stations in an environment, users should consider the following observations:

Orbital Parameters: Many LEO satellites are in polar orbits, to support coverage of the earth for applications like providing internet or monitoring weather. CubeSats are also often deployed from

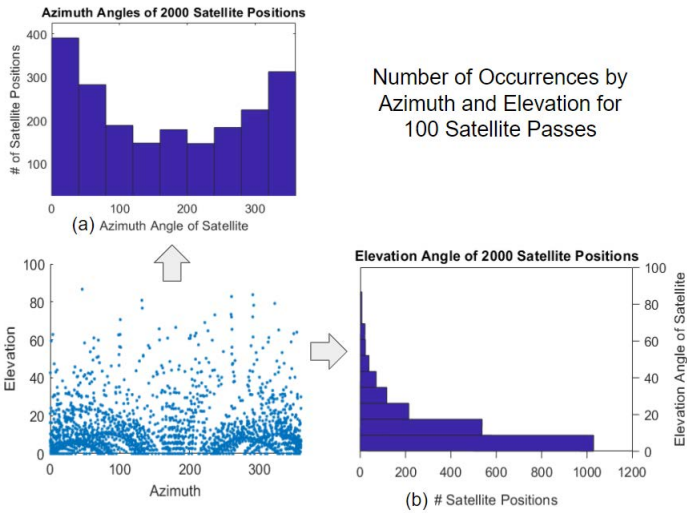


Figure 6: 100 satellite passes of 10 LEO satellites, showing (A) Azimuth Angles with local maxima at 0, 360 degrees (North) and 180 degrees (South). (B) Elevation angles primarily near the horizon (close to 0 degrees)

the International Space Station, at an orbital inclination of 51.64° . In practice, we see from the ground that most satellites pass roughly from north to south or south to north, as shown in Fig. 6(a). Furthermore, most buildings are rectangular and coarsely align with the cardinal directions. Thus, for the typical case, a base station deployment could be selected based on system requirements: East or West for fewer, more complete passes, or North or South for more, less complete passes.

Antenna Behavior: The monopole antennas we use (Laird EXS136) are vertically polarized duck antennas, as they co-optimize low cost and small form factor (4 in) even at meter-scale wavelengths. Duck antennas are omnidirectional and propagate roughly in a toroid, and thus are most receptive of signals near the horizon when they are vertically oriented and most receptive of signals overhead when they are horizontally oriented. As most satellite passes are relatively near the horizon as seen in 6(b), we recommend vertically orienting antennas within the home for typical usage.

We show SNR across time for a very good trajectory for three antenna directions in Fig. 7. The vertical antenna performs better in the beginning while the satellite is near the horizon, and the horizontal antenna performs better in the middle when the satellite is overhead, supporting our hypothesis. We note our antenna oriented at 45° performs poorly over the duration of the trace, which we surmise is due to reflectors, polarizers, and diffractors in the environment (e.g. buildings), which are typically built at right angles.

Obstructions: Obstructions to indoor settings can be quite problematic. Since wavelengths for LEO satellites are usually quite large (on the order of meters), they can typically propagate through most building materials. However, some building materials have quite a high permittivity and can result in significant attenuation, particularly moist masonry, which can attenuate a signal up to 26x




Figure 7: SNR over a satellite pass for different indoor antenna angles. Note the vertically oriented antenna works better in the beginning, when the satellite is near the horizon, and the horizontal antenna works better in the middle, when the satellite is overhead.




Figure 8: SNR over a satellite pass from windows at different indoor elevations. Note that the second and third floors behave comparably well, but the first floor suffers because of nearby ground obstructions.

as much as a vacuum [56]. RF waves cannot propagate through metal, either, so received signals might be poor near window frames and doors. Environmental factors such as leaves may attenuate the received signal as well, particularly after a rain.

Floor: Conventional wisdom for antenna deployment recommends deploying ground stations as high as possible [88]. In practice, we found that for signals from satellites, there is little difference in receiving data across floors when near a window as long as the antennas are high enough to avoid interference from ground reflections (i.e. a wavelength). In Fig. 8, we show SNR across time for three different floors with the same antenna configuration. We observe that the second and third floors behave similarly well, but interference from ground reflection results in the first floor performing relatively poorly.

Multipath: Multipath complicates the ideal placement of ground stations in an environment-specific way. Unlike the aforementioned considerations, devising universal placement advice that accounts for multipath in a new environment is challenging. We believe that a data-driven approach to placement guidance that accounts for multipath is possible, when Quasar is deployed at a truly dense and global scale. We leave this for future work.

7 COHERENT COMBINING

Upon deploying the receivers (following the guidance in Sec. 6) and synchronizing different receivers at the cloud (Section 5), in

this section, we describe how Quasar coherently combines received satellite signals. After synchronization, the signals no longer have any hardware specific offsets and are phase coherent. With this, our main objective is to combine weak signals from tiny receivers into high-quality received signals. We set out the exact formulation that we use for coherently combining such signals in Sec. 7.1 as well as the associated system design challenges in Sec. 7.2.

7.1 Coherent Combining Algorithm

Consider N receivers receiving a signal x and each with a different channel $h_{s \rightarrow i}$. Intuitively, it suffices to simply time, frequency and phase synchronize each satellite receiver to any single common one. Let us assume that ground station k is this common receiver to create the illusion that the entire network of ground stations is connected to receiver k 's clock. Following the operations described in Sec. 5.2 on $h_{s \rightarrow i}x$, we get $h_{ik}x$. With this, we have synchronized the i^{th} receiver to k . After synchronizing each stream, we perform maximal ratio coherent combining [27] as follows.

$$x_{\text{combined}} = \sum_{i=1}^N h_{ik}^* h_{ik} x \quad (6)$$

(\cdot) * denotes complex conjugate. In order to perform this, we need to estimate the channels h_{ik} after synchronization. This is done using a one tap equalizer and estimating this over the preamble portion of x . When h_{ik} is estimated accurately, each receiver's signal adds in phase and noise adds up incoherently, thus providing gain in SNR compared to a single receiver. This gain is termed as diversity gain.

It is important to note that irrespective of which synchronization approach is used, the algorithm for coherent combining resembles Equation 6. Specifically, in the case of synchronizing each receiver to the transmitted satellite signal (according to Sec. 5.3), $h_{s \rightarrow i}x$ now becomes, $\hat{h}_{s \rightarrow i}x$. Once again, $\hat{h}_{s \rightarrow i}$, is estimated for each receiver separately and we combine them as $\sum_{i=1}^N \hat{h}_{s \rightarrow i}^* \hat{h}_{s \rightarrow i} x$.

7.2 System Design

From a design perspective, obtaining a strong combined signal through a cloud architecture for coherent combining presents interesting tradeoffs. This is because in theory, the quality of the result of coherent combining will always improve as more and more streams are provided from receivers. But from a systems point of view, we note that there could be several reasons why it may be unnecessary to send raw I/Q streams to the cloud all the time. (1) Because LEO satellites are visible at any point on the Earth for only short periods of time during the day, signals received by a particular receiver is irrelevant for most of the day, except when the satellite is overhead. (2) As explained in Sec. 6, not all receivers have *good* signal to noise ratio for each pass of the satellite. On the other hand, in certain deployments, there may be too many *good* receivers. In both these cases, sending raw I/Q streams from all receivers is also not necessary and at some point results in diminishing returns.

Based on these observations, our solution in tackling these scenarios is to design the system such that the cloud is a central authority in deciding when and which receivers to combine and therefore, save vital upstream bandwidth at the distributed ground stations. We propose the receiver-cloud architecture encompassing the following features:

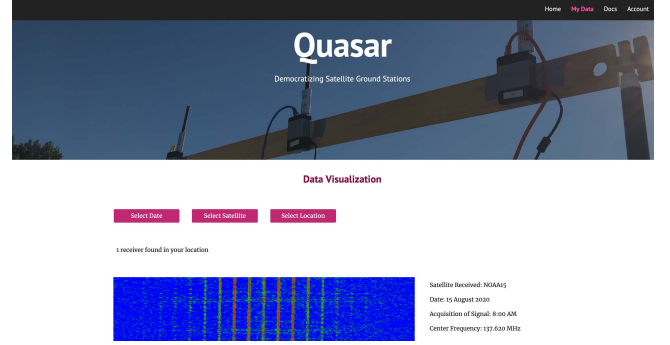


Figure 9: A screenshot of Quasar's web portal (all designs, code will be made open source). Quasar seeks to create an active community towards democratizing access to satellite reception.

Location aware satellite reception: LEO satellites pass over a location only for a small amount of time over the course of a day. This is precisely why receivers don't need to listen always. In our system, receivers make use of satellite trajectory model predictors (such as SGP4) to identify at what times different satellites pass overhead at its specific location. With this information, receivers turn on their radios only for this estimated duration. In the event of multiple satellite passes occurring over an overlapping time interval, users can decide which satellite to listen to. In this way, each receiver independently decides when to listen to a satellite and record I/Q samples. Given a sufficient population of radio receivers and common interest in listening to a particular satellite, this will result in successful recovery of data using coherent combination for each satellite.

Report recording statistics to cloud: Upon completing the recording, each receiver computes basic statistics such as Average SNR and Peak SNR. It then reports them to the cloud along with the initial recording time stamp at which satellite's signal was recorded and the center frequency that was locked to listen to it. The cloud receives these and performs an analysis to decide which receivers to combine (which we detail, next). Meanwhile, in case the receiver is expecting another satellite pass overhead, it pushes the recorded data into a circular buffer and proceeds with the new recording.

Cloud decides which streams to combine: The greater the number of streams coherently combined, the greater is the signal strength that is boosted. On one hand, combining as many signals as possible when the average SNR reported by many satellite receivers are poor makes intuitive sense. But on the other hand, combining fewer signals makes sense if each of them have very high average SNR. The quantity that drives the choice between the *aggregate all* and *aggregate what you need* strategies hard is the uplink bandwidth cost. To choose a suitable operating point along the bandwidth consumed vs. SNR gained curve, the cloud considers the satellite's data payload (image, voice, ...) and decides on an acceptable SNR (*SNR-Desired*) for the coherently combined signal that gives a good quality of decoded data. It then computes the maximum theoretical SNR that can be achieved by coherently combining different sets of receiver streams. If none of them are at least as high as *SNR-Desired*,



Figure 10: Receiver in indoor setting



Figure 11: Ground truth antenna installed on rooftop

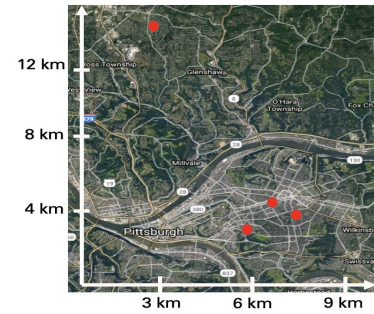


Figure 12: Testbed: red dots are buildings with receivers

the cloud decides to throttle down $SNR-Desired$ and repeat the process until at least one set of receiver streams matches $SNR-Desired$. If more than one of the sets exceed $SNR-Desired$, cloud decides to select the set with least cardinality, that is, minimum number of receiver streams. While our current implementation makes the choice based on cardinality, given that upstream capacity across Quasar’s deployment locations are relatively similar, one could alternatively factor in other elements such as the user’s available bandwidth, fairness considerations and user-set upload limits. Once the set of streams to aggregate is decided, the cloud instructs receivers to upload the recorded data corresponding to the timestamp of interest. Once the relevant data is uploaded to the cloud, initial preprocessing is done to ensure every data point corresponds to the same sampling rate (users can choose sampling rate depending on their hardware constraints and report the rate during cloud upload). If different sampling rates are used, the signal is resampled in software to bring each of them to the same rate, a step that works only if the original sampling rate is more than the bandwidth of the signal. After this pre-processing step, packet preamble is used to detect the beginning of the data in each signal followed by offset removal and combination as described in Sections 5 and 7.1 respectively.

Channel state fluctuations during data exchange period: We note that signal reception at each of the ground stations happens at the same time and the channel is estimated from these signals on a per packet basis before coherent combination using the preamble. The channel fluctuations during the cloud upload process is irrelevant for coherent combination since we estimate the channels for the period when the satellite is transmitting and not for the period when the data is uploaded to the cloud.

Open-Source Framework and Portal: One of the main thrusts of Quasar is to create a vibrant community to democratize the access to satellite receptions. To this end, we open source our hardware designs and code, including a web portal (see Figure 9) which represents the front end of different operations happening behind the screen. Users can create their account and upload multiple satellite data to the portal. Placement guidance discussed in Sec. 6 helps them choose an appropriate location for receiver placement. Once chosen, users add global coordinates of the receiver and meta tags describing this the date,time, SNR metrics and sampling rate of each reception. Based on the cloud’s instructions, selected data is uploaded to the backend, where receivers’ data is coherently combined following operations in Sections 5 and 7.1. Users can view the coherently combined data by clicking on the processed data

tab. To demonstrate the usefulness of coherent combining, we also display the raw data received from individual receivers.

8 IMPLEMENTATION

In order to allow more and more people to record and contribute signals for Quasar’s design to work, we seek to reduce the cost and effort of deployment by using cheap and small form-factor hardware that can be deployed from anywhere (indoors or outdoors). Below we describe Quasar’s hardware and software implementation.

Quasar Hardware: Quasar’s receiver hardware is composed of a computer, a software defined radio (SDR) and an antenna. We use the Laird Technologies EXS136SMI Tuf Duck and ICOM FAS24V half wave helical antennas. These antennas offer unity gain with toroidal radiation pattern and vertical polarization which makes it suitable for land radio in the VHF band. Both the antennas are rated to operate between 136 MHz and 150 MHz and designed to be used with handheld VHF radio receivers like ICOM F50,etc. While primarily designed for VHF terrestrial and amateur radio, the band of operation makes it suitable for satellite signal reception in the VHF band, particularly the US NOAA and Russian METEOR weather satellites transmitting at around 137 MHz. Any other VHF radio antenna operating in the frequency band of interest would be compatible with Quasar’s design as long as it can be interfaced with the SDR receiver RF port. At the SDR front, we use the popular low cost RTL-SDR dongle R820T2 RTL2832U offered by RTL-SDR blog. The ease of use of the plug-and-play dongle made it a natural choice which can be affordably scaled for a large scale distributed receiver system. The total bill of materials for our RF-frontend is \$38. The final part of the receiver chain is the computer which can be as simple as a Raspberry Pi system with USB port that can be ssh-ed into remotely or a laptop with an internet connection, to stream the received signal to the cloud.

Quasar Software: At the software front of the receiver, we use RTL-SDR software installed on the computer connected to the dongle to configure the receiver parameters. In order to aid visualization for the users, we also use GQRX software to view the real time waterfall of the received signal from the RTL-SDR. An important point to note is that we disable the Automatic Gain Control (AGC) setting in RTL-SDR receivers used in our system and operate at the lowest gain setting to prevent the degradation of the low power satellite signal in the presence of high power signal of opportunity. The option to disable the AGC is present both in GQRX as well as through a

Linux terminal command. The collected data is then streamed to the cloud over internet connection for further processing involving synchronization followed by coherent combination. We use Matlab and Python for Digital Signal Processing at the backend as well as generating visual results like waterfall and satellite data for the web portal front end.

Testbed: We evaluate Quasar on Carnegie Mellon University campus as well as a city scale testbed in Pittsburgh, where the receivers were distributed among eight locations to receive signals in urban and suburban neighborhoods. For the outdoor experiments, users were asked to collect data from rooftops of buildings while for the indoor experiments, users were asked to collect data from within their homes and other indoor locations in the university campus. The distribution of candidate receiver locations are shown in Fig.12. In order to verify the performance of our system, we install UC-1374-531R quadrifilar helix antenna [9] (see fig.11) connected to LNAU-0137-648 low noise amplifier and USRP N210 SDR to collect the ground truth signal.

Regulations: No special permits or licenses are needed to receive signals from Amateur radio frequency bands. We only focus our attention on satellites whose transmission characteristics are publicly available. This is true for most of the satellites transmitting in the Amateur radio bands including the NOAA weather satellites used for evaluation. For satellite reception and decoding in Quasar, the signal characteristics like preamble, modulation, frequency, bandwidth, reception time for both satellite and signal of opportunity are needed. In addition, for Doppler correction TLE files containing orbital parameters are needed to be given as input to the SGP4 model. All the above mentioned information is publicly available as mentioned in Section 5.

9 RESULTS

In this section, we evaluate the performance of our systems across various microbenchmarks and system parameters. For all parts of the evaluation, we tune our receivers to receive signals from NOAA 15, NOAA 18, NOAA 19 and METEOR M2 low earth orbit satellites with one of their transmitters transmitting at 137-138 MHz. NOAA satellites are operated by the National Oceanic and Atmospheric Administration, an American weather and scientific agency. They transmit frequency modulated AM data over a bandwidth of 34 kHz in the VHF band in Automatic Picture Transmission (APT) [65] format. The payload of the signal contains periodic snapshots of the region taken by the satellite as it moves overhead. METEOR M2 is a Russian LEO satellite transmitting similar earth images in the 137.100 MHz over a bandwidth of 120 KHz in Low Rate Picture Transmission(LRPT) [81] format. While our evaluation is limited to the VHF band, we believe that the techniques readily extend to other frequency bands with minor change in hardware (the receiver antenna). While signals of opportunity in these bands might not always be available, we use an external transmitter to emulate a signal of opportunity for the purpose of evaluation. These are narrowband (50 kHz) FM modulated signals generated in a band atleast 200 kHz adjacent to the satellite frequency. FM modulation is chosen to emulate the transmission characteristics of most amateur radio towers [47]. The transmit power is limited to 30 dBm because of limitations of the hardware used.




Figure 13: SNR gain with and without synchronization for 2 groundstations

9.1 Microbenchmarks

Synchronization: In order to perform coherent combination, we need to make sure that the channels of individual ground stations are offset free as discussed in Section 5. To characterize the performance of synchronization, we use two co-located receivers deployed on the rooftop of an academic campus building, tuned to the same frequency listening to NOAA weather satellite signals. We then measure the gain in SNR upon coherently combining the signals from the 2 receivers across packets with and without performing synchronization.

Fig.13 shows the CDF of SNR gain achieved for the 2 co-located receivers compared to a single receiver. We can see that for the case where coherent combination was done after synchronization of the receivers, the average gain achieved is 2.8 dB. Also plotted is the gain achieved when coherent combination is attempted without prior synchronization. We can see that the synchronized reception achieves 1.8 dB more gain on an average compared to the case without synchronization, showing the importance of the synchronization component.

SNR gain vs Bandwidth tradeoff: While coherent combination from multiple receivers boosts the signal power ensuring correct reception, not all the signals received from every receiver may be needed to achieve the desired gain as described in 7.2. We receive signals using 8 spatially distributed ground stations in both indoor and outdoor locations on an academic campus building and perform coherent combining opportunistically – We use the best channel with maximum absolute power coefficient as the baseline and coherently combine signals from channels in the decreasing order of channel powers for both indoor and outdoor locations.

Fig. 19 shows the SNR gain compared to the best receiver. While the diversity gain increases initially, it starts to saturate as more and more weaker channels are used to coherently combine and regenerate the signal. It can be seen that the gains begin to marginalize as weaker channels are used, so much so that one pays more for uplink bandwidth resources in exchange for negligible incremental gains. For achieving the same SNR gain with respect to the best channel, we can save 33.33% of uplink bandwidth and 50 % of uplink bandwidth can be saved by taking a 1 dB hit in SNR gain.

9.2 Diversity Gain metrics

Diversity Gain: In order to characterize the average diversity gain for Quasar, we deploy 8 ground stations at different indoor and




Figure 14: SNR gain with increasing number of groundstations




Figure 15: SNR gain across different satellites




Figure 16: SNR gain at different maximum elevation angles

outdoor locations in our city scale test bed shown in Fig. 12 and to receive signals from the aforementioned satellites and perform coherent combining at the cloud. In order to perform a fair analysis of the diversity gain, we use the average channel power of all ground stations as the baseline for calculating the diversity gain as the number of ground stations are increased.

Fig.14 plots the average diversity gain of Quasar in both indoor and outdoor scenarios. We can see that gain follows a logarithmic trend as the number of ground stations are increased and achieves an average gain of 8.1 dB for 8 ground stations. The gain in indoor scenario closely follows the gain achieved when ground stations are deployed outdoors, showing that Quasar’s coherent combining works equally well both indoors and outdoors. Note that diversity gain for 2 receiver case is different from that in Fig.13 because the baseline SNRs for single receiver in these two cases are different.

Effect of satellite diversity: Satellite signals are dependent on the orbital parameters of the individual satellite – path loss is dependent on the distance between satellite and the receiver, Doppler offset is dependent on the satellite’s velocity, etc. In this subsection, we try to characterize what effects satellite trajectories have on the coherent combination gain. NOAA 15, NOAA 18, NOAA 19 and METEOR M2 orbit the earth at an orbital altitude of 813 km, 856 km, 850 km and 825 km respectively and follow sun-synchronous polar orbits completing one rotation in about 100 minutes. We calculate the average SNR gain achieved for a fixed number of 6 ground stations deployed both indoors and outdoors with respect to the baseline SNR calculated in above paragraph by averaging across all base stations for different satellite passes. Fig.15 shows all the satellites performing equally well in terms of SNR gain both indoors and outdoors, which shows that Quasar’s gains are not significantly impacted by changes in orbital parameters.

Effect of satellite trajectory: Another aspect of the LEO satellite reception is the amount of visibility of the satellite from a particular region on earth. Not all passes of LEO satellite over a particular region are ideal – i.e. right overhead at the zenith at the 90 degrees elevation angle. If during a particular pass, the elevation angle available to the user is small, there is a high likelihood that the signal from the satellite will not be received correctly. There are two reasons for this: (1) signals are more likely to be attenuated by taller structures around the receiver (this is why traditional approach is to install receivers on rooftops or at an elevated place), and (2)

signals travel much larger distance from the satellite to the receiver at lower elevation angles compared to higher elevation angles.

To characterize the effect of elevation angle on Quasar, we collect signals across multiple passes covering a wide range of elevation angles. We categorize the maximum elevation angle during a satellite pass into 3 groups: poor (maximum elevation < 45 degree), moderate (45 degrees < maximum elevation < 70 degrees) and good (maximum elevation > 70 degrees). Fig.16 shows the variation of SNR gain with respect to the average baseline for 6 ground stations across the 3 groups of elevation angles. We can see that poor elevation angle offers lower gain and the gap between indoors and outdoors is also higher. This is due to the further attenuation of the already weak signal from walls and other obstructions. As expected, the SNR gain improves from poor to moderate to good passes as elevation angle increases.

9.3 Data Recovery metrics

While diversity gain of a multi antenna system provides good insights into the performance, we should also carefully examine if the gain achieved is good enough for data recovery. This is important since the received signal might still not be recoverable despite the diversity gain produced by coherent combination. Since the data transmitted by the LEO weather satellites are low resolution images of portions of earth, we use Structural Similarity Index Measure (SSIM) [93] and Peak Signal to Noise Ratio (PSNR) [41] metrics below to quantify image quality of recovered image to check if Quasar is capable of recovering the desired data. We calculate these metrics using the image generated by ground truth Weather antenna deployed on the academic campus building roof as described in Sec. 8.

SSIM Variation: SSIM is an image quality metric that perceives structural degradation between a test image and a reference image by taking into account factors like luminance, contrast and structure. It is an absolute metric with values close to 1 indicating good image quality. To see the variation of SSIM as we increase the number of ground stations (and hence expect to improve quality), we generate image after adding every additional receiver and measure its SSIM metric with the ground truth image generated by rooftop weather antenna. This metric is generated for different satellites and multiple passes. Fig.17 shows improvement in average SSIM from 0.52 to 0.85 outdoors and 0.35 to 0.74 indoors from single ground station to 8 ground stations.




Figure 17: SSIM variation with increasing number of groundstations




Figure 18: PSNR improvement with increasing number of groundstations




Figure 19: SNR gain - Bandwidth tradeoff




Figure 20: NOAA 18 image with one receiver (i.e. baseline) vs. Quasar with 4 and 8 receivers (left to right)

PSNR Variation: PSNR is another metric used to compare images – it measures the mean squared error (MSE) in test image from the reference image and reports the inverse of MSE as PSNR in dB. We calculate the average PSNR for increasing number of ground stations in a similar fashion to SSIM described in the above paragraph. Fig.18 shows an average 5 dB improvement in PSNR from one ground station to 8 ground stations. While our approach may not exactly match the reference image, we must keep in mind that these metrics are defined for computer vision and image reconstruction fields which have much more stringent requirements compared to wireless communication performance. We also present in Fig.20, how the recovered image evolves as more and more receivers are added in coherent combining. We believe further improvements in image quality are possible with better image coding available in other LEO satellites.

10 DISCUSSIONS AND LIMITATIONS

While our discussion thus far focuses on enabling Quasar using synchronization and coherent combining techniques for broadcast data, in this section we comment on the applicability to other types of data as well as touch upon some of limitations of our system.

Number of receivers and amount of data needed: An important question that arises is how many receivers are needed to get the same performance as a single big base station. The number of receivers needed depends on the gain of the single big antenna, the intermediate hardware filters used and the ADC resolution of the SDR used. Typically, the difference in gain between a big high gain directional yagi antenna and the small vertically polarized duck antenna used in this paper is of the order of 10-12 dB. From our evaluation in Section 9, we can see that the SNR gains scale almost

logarithmically with the number of receivers for 8 antennas. Extrapolating from that, we expect the number of receivers required will be of the order of 15-20 to achieve similar performance to a single big receiver. The amount of data needed scales linearly with the number of receivers. However, when it comes to the burden on the individual user, the amount of data that needs to be collected at the user end is the same as that in case of a big powerful receiver.

Applicability to Unicast satellite transmissions: In its current form, Quasar is applicable for broadcast data. However, the techniques are applicable to unicast data reception by using multiple colocated antennas - similar to the phased array dish antenna used by SpaceX's starlink internet receivers. Whether one would use multiple small antennas or a big antenna depends on multiple factors - space and ease of installation, cost, etc. We qualitatively motivate Quasar with ease of deployment and use, however, a quantitative comparison is beyond the scope of this work.

11 CONCLUSION AND FUTURE WORK

This paper presents Quasar, a community-driven network of inexpensive satellite ground stations that together enable low-cost and effective communication infrastructure for Low Earth Orbit satellites. Quasar uses coherent combining to phase coherently add together signals received across receivers to boost signal to noise ratio. As a result, Quasar supports inexpensive receivers that can be deployed indoors or outdoors without significant installation expense.

In the future, we plan to grow Quasar as a community resource and expand the scope and scale of the Quasar testbed. We believe a future global Quasar testbed will prove as a valuable resource in providing useful signal quality data on satellite communication – crucial to the development of future space-based communications infrastructure. We believe the availability of distributed I/Q samples over a large geographical deployments can also open up opportunities for wireless atmospheric and geo-sensing leveraging existing satellite signals.

Acknowledgements: We thank the shepherd and reviewers for their insightful feedback which helped improve this paper. We thank Weiwei Fan, Deeptha Anil Kumar and Milind Srivastava for helping run experiments during different phases of the project. We also thank all the members of the WiTech Lab at CMU for their valuable inputs. This research was supported in parts by NSF grants 2030154, 1942902 and 2007786, Kavčić-Moura grant and IoT@CyLab.

REFERENCES

- [1] 2013. TU-Delft Expanding in space. <https://www.tudelft.nl/en/ae/news/spotlight/tu-delft-expanding-in-space/>
- [2] 2020. Infostellar. <https://infostellar.net/>
- [3] O. Abari, H. Rahul, D. Katabi, and M. Pant. 2015. AirShare: Distributed coherent transmission made seamless. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. 1742–1750.
- [4] Australian Space Agency. 2020. ORBITAL PARAMETERS: LOW EARTH CIRCULAR ORBITS. <http://www.spaceacademy.net.au/watch/track/leopars.htm>
- [5] The European Space Agency. 2020. Types of orbits. Retrieved August 21, 2020 from https://www.esa.int/Enabling_Support/Space_Transportation/Types_of_orbits
- [6] amazon.com. 2020. AWS Ground Station Pricing. Retrieved August 21, 2020 from <https://aws.amazon.com/ground-station/pricing/>
- [7] amazon.com. 2020. AWS ground stations. Retrieved August 21, 2020 from <https://aws.amazon.com/ground-station/>
- [8] AMSAT-UK. 2011. CubeSat UHF Downlink Frequencies Chart. <https://amsat-uk.org/2011/11/06/cubesat-uhf-downlink-frequencies-chart/>
- [9] antennas.us. 2020. WXSAT-APT-KIT, VHF APT Satellite Weather Kit. Retrieved August 21, 2020 from <https://www.antennas.us/WXSAT-APT-KIT-Satellite-Weather-Kit.html>
- [10] Jinzhen Bao, Baokang Zhao, Wanrong Yu, Zhenqian Feng, Chunqing Wu, and Zhenggu Gong. 2014. OpenSAN: A Software-Defined Satellite Network Architecture. In *Proceedings of the 2014 ACM Conference on SIGCOMM* (Chicago, Illinois, USA) (SIGCOMM '14). Association for Computing Machinery, New York, NY, USA, 347–348. <https://doi.org/10.1145/2619239.2631454>
- [11] Aniceto Belmonte and Joseph M Kahn. 2009. Capacity of coherent free-space optical links using diversity-combining techniques. *Optics express* 17, 15 (2009), 12601–12611.
- [12] Matteo Berioli, Nicolas Courville, and Markus Werner. 2007. Integrating satellite and terrestrial technologies for emergency communications: the WISECOM project. In *QShine 2007 Workshop: Satellite/Terrestrial Interworking*. 1–6.
- [13] Debopam Bhattacharjee, Simon Kassing, Melissa Licciardello, and Ankit Singla. 2020. In-orbit Computing: An Outlandish thought Experiment?. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 197–204.
- [14] P. Bissig, M. Eichelberger, and R. Wattenhofer. 2017. Fast and Robust GPS Fix Using One Millisecond of Data. In *2017 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 223–234.
- [15] Vincent Bonneau, Bastile Carle, Laurent Probst, and Bertrand Pedersen. 2017. Low-Earth Orbit Satellites: Spectrum access. https://ec.europa.eu/growth/tools-databases/dem/monitor/sites/default/files/DTM_LEO%20-%20Spectrum%20access%20v1_0.pdf
- [16] Donald CD Chang. 2015. Coherent power combining via wavefront multiplexing on deep space spacecraft. US Patent 9,077,427.
- [17] Aleksandra Checko, Henrik L Christiansen, Ying Yan, Lara Scolari, Georgios Kardaras, Michael S Berger, and Lars Dittmann. 2014. Cloud RAN for mobile networks—A technology overview. *IEEE Communications surveys & tutorials* 17, 1 (2014), 405–426.
- [18] Richard E. Christian. 2019. Pittsburgh Area Amateur Radio Repeaters. Retrieved January 10, 2021 from <http://www.pgscanner.com/amat.php>
- [19] Cecilia Clivati, Roberto Ambrosini, Thomas Artz, Alessandra Bertarini, Claudio Bertolotti, Matteo Frittelli, Filippo Levi, Alberto Mura, Giuseppe Maccaferri, Mauro Nanni, and et al. 2017. A VLB experiment using a remote atomic clock via a coherent fibre link. *Scientific Reports* 7, 1 (Feb 2017). <https://doi.org/10.1038/srep40992>
- [20] Walter Colitti, Kris Steenhaut, Nicolas Descouvemont, and Adam Dunkels. 2008. Satellite based wireless sensor networks: global scale sensing with nano-and pico-satellites. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*. 445–446.
- [21] Charles Cooper, Ronald Fevig, and Jason Patel. 2002. The CubeSat Ground Station at the University of Arizona. ftp://pir1-ftp.d.lpl.arizona.edu/pub/cubesat/cubesat_papers/gspaper.pdf
- [22] Rajsheshkar Das, Akshay Gadre, Shanghang Zhang, Swarun Kumar, and Jose MF Moura. 2018. A deep learning approach to IoT authentication. In *2018 IEEE International Conference on Communications (ICC)*. IEEE, 1–6.
- [23] Bradley Denby and Brandon Lucia. 2020. Orbital Edge Computing: Nanosatellite Constellations as a New Class of Computer System. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 939–954.
- [24] Wei Dong and Zhao Chang-yin. 2010. An accuracy analysis of the SGP4/SDP4 model. *Chinese Astronomy and Astrophysics* 34, 1 (2010), 69–76.
- [25] Adwait Dongare, Patrick Lazik, Niranjini Rajagopal, and Anthony Rowe. 2017. Pulsar: A wireless propagation-aware clock synchronization platform. In *2017 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 283–292.
- [26] 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.
- [27] Thomas Eng, Ning Kong, and Laurence B Milstein. 1996. Comparison of diversity combining techniques for Rayleigh-fading channels. *IEEE Transactions on communications* 44, 9 (1996), 1117–1129.
- [28] The Antenna Farm. 2020. Laird Technologies EXS-136-SMI. Retrieved August 21, 2020 from <https://www.theantennafarm.com/catalog/laird-technologies-exs-136-smi-2288>
- [29] Michael Fischer and Arpad L Scholtz. 2010. Design of a multi-mission satellite ground station for education and research. In *2010 Second International Conference on Advances in Satellite and Space Communications*. IEEE, 58–63.
- [30] Libre Space Foundation. 2020. <https://satnogs.org/>
- [31] Akshay Gadre, Revathy Narayanan, and Swarun Kumar. 2018. Maintaining UAV Stability using Low-Power WANs. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 738–740.
- [32] 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.
- [33] 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.
- [34] David J Geisler, Timothy M Yarnall, Mark L Stevens, Curt M Schieler, Bryan S Robinson, and Scott A Hamilton. 2016. Multi-aperture digital coherent combining for free-space optical communication receivers. *Optics Express* 24, 12 (2016), 12661–12671.
- [35] Khondokar Fida Hasan, Charles Wang, Yanming Feng, and Yu-Chu Tian. 2018. Time synchronization in vehicular ad-hoc networks: A survey on theory and practice. *Vehicular communications* 14 (2018), 39–51.
- [36] Yannick Hauri, Debopam Bhattacharjee, Manuel Grossmann, and Ankit Singla. 2020. "Internet from Space" without Inter-satellite Links. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 205–211.
- [37] Wen He. 2018. The Analysis of RTS/CTS Mechanism Based on LEO Satellite Wi-Fi Network. In *Proceedings of the 2018 2nd International Conference on Algorithms, Computing and Systems*. 184–189.
- [38] Caleb Henry. 2017. Japanese startup raises \$7.3 million for smallsat antenna sharing service. Retrieved August 21, 2020 from <https://spacenews.com/japanese-startup-raises-7-3-million-for-satsat-antenna-sharing-service/>
- [39] RM Hjellming and RC Bignell. 1982. Radio Astronomy with Very Large Array. *Science* 216, 4552 (1982), 1279–1285.
- [40] Felix R Hoots. 1980. Spacetrack report no. 3, models for propagation of norad element sets. <http://www.itc.nl/~bakker/orbit.html> (1980).
- [41] National Instruments. 2020. Peak Signal-to-Noise Ratio as an Image Quality Metric. <https://www.ni.com/en-us/innovations/white-papers/11/peak-signal-to-noise-ratio-as-an-image-quality-metric.html>
- [42] Haojian Jin, Jingxian Wang, Swarun Kumar, and Jason Hong. 2019. Software-Defined Cooking using a Microwave Oven. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–16.
- [43] Haojian Jin, Jingxian Wang, Zhijian Yang, Swarun Kumar, and Jason Hong. 2018. Wish: Towards a wireless shape-aware world using passive rfid. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*. 428–441.
- [44] Haojian Jin, Zhijian Yang, Swarun Kumar, and Jason I Hong. 2018. Towards wearable everyday body-frame tracking using passive rfid. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 4 (2018), 1–23.
- [45] Jakub Kaderka and Tomáš Urbanec. 2020. Time and sample rate synchronization of RTL-SDR using a GPS receiver. In *2020 30th International Conference Radioelektronika (RADIOELEKTRONIKA)*. IEEE, 1–4.
- [46] NE Kassim, RA Perley, WC Erickson, and KS Dwarakanath. 1993. Subarcminute resolution imaging of radio sources at 74 MHz with the Very Large Array. *Astronomical Journal* 106 (1993), 2218–2228.
- [47] KD6KPC/WREQ745. 2021. Allegheny, Pennsylvania Amateur Radio Repeaters. Retrieved January 10, 2021 from https://www.repeaterbook.com/repeaters/location_search.php?state_id=42&loc=Allegheny&type=county
- [48] R Kim, T Nagayama, Hongki Jo, and BF Spencer Jr. 2012. Preliminary study of low-cost GPS receivers for time synchronization of wireless sensors. In *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2012*, Vol. 8345. International Society for Optics and Photonics, 83451A.
- [49] Public Lab. 2020. DIY Satellite Ground Station. Retrieved January 10, 2021 from <https://publiclab.org/notes/sashae/06-26-2020/diy-satellite-ground-station>
- [50] Zachary J Jeffke. 2014. *Distributed Ground Station Network for CubeSat Communications*. Master's thesis. <https://vtechworks.lib.vt.edu/handle/10919/25151>
- [51] Y. Lin, W. Dong, Y. Gao, and T. Gu. 2020. SatLoc: A Virtual Fingerprinting Approach to Outdoor LoRa Localization using Satellite Images. In *2020 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 13–24.

- [52] Pierre Lionnet. 2017. How much operational lifetime can we expect from cubesats? Retrieved August 21, 2020 from <https://www.linkedin.com/pulse/how-much-operational-lifetime-can-we-expect-from-cubesats-lionnet>
- [53] Xinghua Liu, Zhenhai Xu, Xiang Liu, Siwei Chen, and Shunping Xiao. 2018. A clean signal reconstruction approach for coherently combining multiple radars. *EURASIP Journal on Advances in Signal Processing* 2018, 1 (2018), 1–11.
- [54] Scott Madry, Joseph N Pelton, and Sergio Camacho-Lara. 2017. Ground Systems for Satellite Application Systems for Navigation, Remote Sensing, and Meteorology. *hsa* (2017), 1343.
- [55] Lockheed Martin. 2018. Amazon Web Services and Lockheed Martin Team to Make Downlinking Satellite Data Easier and Less Expensive. Retrieved January 10, 2021 from <https://news.lockheedmartin.com/news-releases?item=128622>
- [56] Kenzi Meyer, Ece Erdogan, George Morcoux, and Mary Naughtin. 2008. Use of Ground Penetrating Radar for Accurate Concrete Thickness Measurements. *Proceedings of the AEI 2008 Conference - AEI 2008: Building Integration Solutions* 328, 1–10. [https://doi.org/10.1061/41002\(328\)67](https://doi.org/10.1061/41002(328)67)
- [57] Prasant Kumar Misra, Wen Hu, Yuzhe Jin, Jie Liu, Amanda Souza de Paula, Niklas Wirstrom, and Thiemo Voigt. 2014. Energy Efficient GPS Acquisition with Sparse-Gps. In *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks* (Berlin, Germany) (IPSN '14). IEEE Press, 155–166.
- [58] N2YO. 2020. LEO satellite schedule. <https://www.n2yo.com/passes/?s=28654&a=1>
- [59] nanosats.eu. 2020. World's largest database of nanosatellites, over 2500 nanosats and CubeSats. <https://www.nanosats.eu/>
- [60] S. Narayana, R. V. Prasad, V. Rao, L. Mottola, and T. V. Prabhakar. 2020. Hummingbird: Energy Efficient GPS Receiver for Small Satellites. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) (MobiCom '20). Association for Computing Machinery, New York, NY, USA, Article 9, 13 pages. <https://doi.org/10.1145/3372224.3380886>
- [61] Revathy Narayanan and Swarun Kumar. 2018. Revisiting software defined radios in the IoT era. In *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*. 43–49.
- [62] Ali A Nasir, Salman Durrani, Hani Mehrpouyan, Steven D Blostein, and Rodney A Kennedy. 2016. Timing and carrier synchronization in wireless communication systems: a survey and classification of research in the last 5 years. *EURASIP Journal on Wireless Communications and Networking* 2016, 1 (2016), 180.
- [63] Wyld Networks. 2020. Wyld announces global LoRa satellite deal. Retrieved January 10, 2021 from <https://wyldnetworks.com/wyld-announces-global-lora-satellite-deal/>
- [64] Cristóbal Nieto-Peroy and M Reza Emami. 2019. CubeSat Mission: From Design to Operation. *Applied Sciences* 9, 15 (2019), 3110.
- [65] National Climatic Data Center NOAA. 2004. NOAA KLM User's Guide: Section 4.2 - APT System. <https://web.archive.org/web/20070316190349/http://www2.ncdc.noaa.gov/docs/klm/html/c4/sec4-2.htm>
- [66] Nyan-Sat. 2020. A down-to-earth satellite challenge... Retrieved January 10, 2021 from <https://nyan-sat.com/index.html>
- [67] Daniel Oltrogge, Ramrath AGI, and AGI Jens. 2014. Parametric Characterization of SGP4 Theory and TLE positional accuracy. (2014).
- [68] Atlas Space Operations. 2020. EMPOWERING GLOBAL ACCESS TO SPACE. Retrieved January 10, 2021 from <https://atlasground.com/>
- [69] planet.com. 2020. Planet Labs. Retrieved April 30, 2020 from <https://www.planet.com/>
- [70] Akarsh Prabhakara, Vaibhav Singh, Swarun Kumar, and Anthony Rowe. 2020. Osprey: A mmWave approach to tire wear sensing. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 28–41.
- [71] Narayan Prasad. 2020. Satellite ground stations: installing your own vs using an existing network. <https://blog.satsearch.co/2020-05-06-satellite-ground-stations-installing-your-own-vs-using-an-existing-network>
- [72] Hariharan Rahul, Swarun Suresh Kumar, and Dina Katabi. 2012. Megamimo: Scaling wireless capacity with user demand. In *Proc. ACM SIGCOMM*, Vol. 4. 1.
- [73] rtl-sdr blog. 2015. Detecting Pulsars (Rotating Neutron Stars) with an RTL-SDR. <https://www.rtl-sdr.com/tag/frequency-drift/>
- [74] rtl sdr.com. 2017. rtl-sdr with Rpi. Retrieved April 30, 2020 from <https://www.rtl-sdr.com/video-tutorial-installing-gqr-x-and-rtl-sdr-on-a-raspberry-pi/>
- [75] Y. Sakamoto. 2014. UNISEC GSN-WG. <http://www.astro.mech.tohoku.ac.jp/~gsn/en/>
- [76] Satnogs.org. 2020. Satnogs Wiki. Retrieved April 30, 2020 from https://wiki.satnogs.org/Main_Page
- [77] Marco Schmidt. 2020. *Ground station networks for efficient operation of distributed small satellite systems*. Gesellschaft für Informatik.
- [78] Marco Schmidt and Klaus Schilling. 2008. Internet-Based Ground Stations Networks for Pico Satellites. (2008). <https://arc.aiaa.org/doi/pdf/10.2514/6.2008-3205>
- [79] Clayton Shepard, Hang Yu, Narendra Anand, Erran Li, Thomas Marzetta, Richard Yang, and Lin Zhong. 2012. Argos: Practical many-antenna base stations. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. 53–64.
- [80] Simon P Shuster. 2017. A Survey and Performance Analysis of Orbit Propagators for LEO, GEO, and Highly Elliptical Orbits. (2017).
- [81] sigidwiki.com. 2020. Low Rate Picture Transmission. [https://www.sigidwiki.com/wiki/Low_Rate_Picture_Transmission_\(LRPT\)](https://www.sigidwiki.com/wiki/Low_Rate_Picture_Transmission_(LRPT))
- [82] Christopher Simpson, Andrew Burjek, William Patton, Edward G. Hackett, and Charles O'Neill. 2019. Ground Station and Infrastructure Development at the University of Alabama, Tuscaloosa. In *70th International Astronautical Congress*. International Astronautical Federation, 1–7. https://www.researchgate.net/profile/Christopher_Simpson17/publication/341494736_Ground_Station_and_Infrastructure_Development_at_the_University_of_Alabama_Tuscaloosa/links/5ec42700458515626cb82957/Ground-Station-and-Infrastructure-Development-at-the-University-of-Alabama-Tuscaloosa.pdf
- [83] Vaibhav Singh, Susnata Mondal, Akshay Gadre, Milind Srivastava, Jeyanandh Paramesh, and Swarun Kumar. 2020. Millimeter-wave full duplex radios. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.
- [84] Fikret Sivrikaya and Bülent Yener. 2004. Time synchronization in sensor networks: a survey. *IEEE network* 18, 4 (2004), 45–50.
- [85] Marco Sozzi, Francesco Marinello, Andrea Pezzuolo, and Luigi Sartori. 2018. Benchmark of satellites image services for precision agricultural use. In *Proceedings of the AgEng Conference, Wageningen, The Netherlands*. 8–11.
- [86] Leaf Space. 2020. Leaf Line: Shared Ground Station Network. Retrieved January 10, 2021 from <https://leaf.space/leaf-line/>
- [87] starlink.com. 2020. Starlink. Retrieved April 30, 2020 from <https://www.starlink.com/>
- [88] R D Straw and Gerald L Hall. 1999. Antenna Height and Communication Effectiveness.
- [89] M2 Antenna Systems. 2020. LEO-Pack, 436CP16/2MCP8A. Retrieved August 21, 2020 from <https://www.m2inc.com/FGLEOPACK>
- [90] 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.
- [91] Mineo Wakita. 2021. Satellite Frequency List. Retrieved January 10, 2021 from <http://www.ne.jp/asahi/hamradio/je9pel/satslist.htm>
- [92] Jingxian Wang, Junbo Zhang, Rajarshi Saha, Haojian Jin, and Swarun Kumar. 2019. Pushing the range limits of commercial passive rfids. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*. 301–316.
- [93] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [94] Satnogs Wiki. 2020. Ground Station Antennas. Retrieved August 21, 2020 from <https://wiki.satnogs.org/Antennas>
- [95] Satnogs Wiki. 2020. Rotators. <https://wiki.satnogs.org/Rotators>
- [96] Christopher F Wildt. 2017. *Accuracy in orbital propagation: A comparison of predictive software models*. Technical Report. Naval Postgraduate School Monterey United States.
- [97] Matt Williams. 2017. What is Low Earth Orbit? <https://www.universetoday.com/85322/what-is-low-earth-orbit/>
- [98] Moon-Hee You, Seong-Pal Lee, and Youngyeul Han. 2000. Adaptive compensation method using the prediction algorithm for the Doppler frequency shift in the LEO mobile satellite communication system. *ETRI journal* 22, 4 (2000), 32–39.
- [99] Diana Zhang, Jingxian Wang, Junsu Jang, Junbo Zhang, and Swarun Kumar. 2019. On the Feasibility of Wi-Fi Based Material Sensing. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–16.
- [100] Yang Zhang, Yasha Iravantchi, Haojian Jin, Swarun Kumar, and Chris Harrison. 2019. Sozu: Self-Powered Radio Tags for Building-Scale Activity Sensing. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 973–985.