



LAVA: Fine-Grained 3D Indoor Wireless Coverage for Small IoT Devices

R. Ivan Zelaya¹, William Sussman¹,
Jeremy Gummeson³, Kyle Jamieson², Wenjun Hu¹

Yale University¹, Princeton University², University of Massachusetts at Amherst³

ABSTRACT

Small IoT devices deployed in challenging locations suffer from uneven 3D coverage in complex environments. This work optimizes indoor coverage with LAVA, a Large Array of Vanilla Amplifiers. LAVA is a standard-agnostic cooperative mesh of *elements*, i.e., RF devices each consisting of several switched input and output antennas connected to fixed-gain amplifiers. Each LAVA element is further equipped with rudimentary power sensing to detect nearby transmissions. The elements report power readings to the LAVA control plane, which then infers active link sessions without explicitly interacting with the endpoint transmitter or receiver. With simple on-off control of amplifiers and antenna switching, LAVA boosts passing signals via multi-hop amplify-and-forward. LAVA explores a middle ground between smart surfaces and physical-layer relays. Multi-hopping over short inter-hop distances exerts more control over the end-to-end trajectory, supporting fine-grained coverage and spatial reuse. Ceiling testbed results show throughput improvements to individual Wi-Fi links by 50% on average and up to 100% at 15 dBm transmit power (193% on average, up to 8× at 0 dBm). ZigBee links see up to 17 dB power gain. For pairs of co-channel concurrent links, LAVA provides average per-link throughput improvements of 517% at 0 dBm and 80% at 15 dBm.

CCS CONCEPTS

• **Hardware** → **Wireless devices**; • **Networks** → **Physical links**; **Wireless local area networks**;

KEYWORDS

Smart Surfaces, Programmable Radio Environments, Non-Uniform 3D Coverage, Multi-Hop Amplify-and-Forward

ACM Reference Format:

R. Ivan Zelaya¹, William Sussman¹, Jeremy Gummeson³, Kyle Jamieson², Wenjun Hu¹ Yale University¹, Princeton University², University of Massachusetts at Amherst³. 2021. LAVA: Fine-Grained 3D Indoor Wireless Coverage for Small IoT Devices. In *ACM SIGCOMM 2021 Conference (SIGCOMM '21)*, August 23–28, 2021, Virtual Event, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3452296.3472890>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCOMM '21, August 23–28, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-8383-7/21/08.

<https://doi.org/10.1145/3452296.3472890>

1 INTRODUCTION

As the vision of IoT becomes a reality [68], off-the-shelf devices today crowd the unlicensed frequency bands in enterprises and homes, with new consumer gadgets deployed in an unplanned manner. Unlike a decade ago, wireless network optimization today is no longer concerned with only laptops, smartphones, or tablets with Wi-Fi connectivity on 2.4 and 5 GHz. More unlicensed bands are available today, such as 900 MHz and 6 GHz in the US. Devices are increasingly diversified, using multiple versions of Wi-Fi alongside other standards like ZigBee, ZWave, and Bluetooth. The hardware quality varies in terms of antenna and signal processing capability. Even on a single floor, they spread over more disparate and sometimes rather challenging locations, such as smart plugs that may be behind furniture, Roombas roaming on the floor, or smart bulbs and security cameras on the ceiling and potentially in the corners of a room. Some only require a single packet for signaling (e.g., flipping the light switch or circuit breaker), while others (e.g., smart cameras) require more bandwidth.

The fundamental challenge facing these technologies is uncontrolled signal propagation indoors, resulting in complex, uneven 3D coverage. Small IoT devices which use a single cheap and inefficient omnidirectional antenna each might easily get blocked by furniture like bookshelves and cabinets. Since enterprise Wi-Fi APs or ZigBee coordinators are usually on the ceiling, while home APs may be in arbitrary locations, it is crucial to provision for both vertical and horizontal coverage. Unplanned deployment of IoT devices also makes them susceptible to antenna polarization mismatch between the endpoints [14] and cross-technology interference [24].

The situation is poised to worsen amid the ever-expanding IoT landscape [33]. While shifts towards smarter homes, offices, and healthcare have been in progress for some time, measures to counter the COVID-19 pandemic have accelerated the transformation of remote work and telehealth. New gadgets showcased at the latest Consumer Electronics Show (CES 2021) aim to improve the well-being of the home occupants [32]. Remote monitoring for eldercare or assisted care adds even more devices [58]. Smarter offices incorporate specialized IoT applications but are increasingly hindered by a lack of standardization across the entire IoT ecosystem [22]. Given this IoT medley, deploying a dedicated wireless infrastructure per application scenario simply becomes unscalable.

To provision coverage for the myriad new devices, infrastructure support faces several new requirements (§2). First, complex geometries of 3D propagation environments and antenna beam patterns exacerbate non-uniform signal propagation already seen in 2D. This motivates power adjustment along the propagation paths to fine-tune energy distribution. Second, capability and hence performance metrics vary between devices, which nevertheless need to share

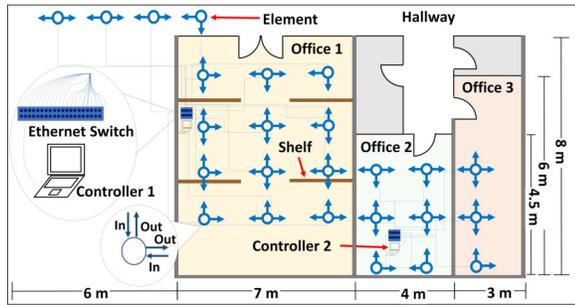


Figure 1: Ceiling testbed deployment for LAVA.

infrastructure support for a cost-effective large IoT deployment. Third, the above application trends highlight the need to consolidate and reuse infrastructure for multi-standard operations and co-existence. For wireless, this motivates a standard-agnostic approach, which then needs to operate at the signal level, even below the physical layer operations as specified by individual standards.

The closest approaches to signal-level enhancements are *smart surfaces* or *physical-layer relays*. Existing relay prototypes [2, 7, 29] extend the range with amplify-and-forward, but incur complex variable amplification gain control at each relay or even a full-duplex radio with substantial baseband circuitry to cancel self-interference. These relays are individually costly, limiting their numbers in a deployment. Further, they are currently customized for either Wi-Fi or 60 GHz and do not readily apply to other technologies.

Smart surfaces (e.g., [6, 30, 39, 54, 67, 70, 75]) are emerging concepts to *generate* favorable perceived channels with many passive radio frequency elements embedded into the environment. The early end-to-end prototypes [4, 14, 17, 38] have demonstrated the potential for beamforming or polarization rotation. LAIA [37, 38], RFocus [4], and ScatterMIMO [17] leverage *spatial diversity* between multiple paths between the transmitter and receiver endpoints and beamform towards the receiver, while LLAMA [14] rotates the polarization of signals to minimize polarization mismatch between endpoint antennas. These approaches, however, rely on pervasive deployment of the smart surfaces near the communication endpoints. Handling wideband, multi-link, and multi-technology co-existence scenarios in 3D would require a massive number of antenna elements. Further, these existing prototypes have so far focused on single-link optimizations, implicitly assuming both endpoints at the same vertical level. They do not yet handle co-channel interfering links or vertical offsets between endpoints.

This paper explores a middle ground between smart surfaces and physical-layer relays, combining the best of both worlds (§3). We design a mesh infrastructure that is amplified like a relay but considerably simpler, per relay unit, like a smart surface. Specifically, we present LAVA, a *Large Array of Vanilla Amplifiers*, as a cooperative mesh of amplifying *elements*, each consisting of several switched input and output antennas connected to fixed-gain amplifiers, a simple power meter, and an optional phase shifter. The elements are arranged in a simple grid topology and connected to a multi-level control architecture of Arduinos and PCs as controllers.

LAVA boosts passing signals primarily via successive amplify-and-forward, which is particularly beneficial to long links bottlenecked by pathloss. With fine-grained multi-hopping, LAVA aims

to specify a dominant end-to-end signal path composed of short segments, in order to overpower and supplant other multipath signal components. LAVA effectively distributes power more uniformly between the endpoints. Short links tend to be bottlenecked by multipath fading and benefit less from multi-hop amplification than phase alignment. In those cases, LAVA can adopt an optional phase adjustment stage, selecting from four settings on the last hop to improve performance at low complexity. Compared to signal alignment-based “smart surfaces,” active amplification in LAVA is more amenable to wideband signals and concurrent links. Compared to conventional multi-hop mesh networking, LAVA relays at the signal level, transparent to MAC layer operations and without halving the capacity per hop or necessitating full-duplex radios. The denser the LAVA topology, the more opportunities to construct multiple diverse and disjoint routes, each servicing a distinct endpoint communication link. A single LAVA deployment easily supports multiple wireless technologies in a standard-agnostic fashion.

To avoid coordination with endpoints, LAVA adopts rudimentary per-element power sensing to locate an active transmitter. Each power meter generates a voltage reading corresponding to the input signal, which is run through slow and fast exponentially weighted moving averages (EWMA) on the Arduino for coarse-grained frame detection. The control plane then infers active link sessions based on the power reports from the elements, and identifies a route through the LAVA mesh. Each element on the chosen route is activated simply with amplifier on-off control and antenna switching. LAVA does not change any MAC protocols and is only activated when the standard MAC protocol permits transmissions.

We envision that a fully integrated hardware realization of LAVA can become part of the building utility infrastructure, with elements plugged in to the regular power and network infrastructure analogous to how lighting fixtures are installed. For the current prototype, we deploy a testbed of distributed and collaborative LAVA instances above the secondary ceiling of real offices and corridors, covering an area of more than 80 square meters (Figure 1 and §4). Each LAVA instance occupies one room or area in the figure. This deployment strategy is minimally intrusive or visible. We also experiment with a floor deployment with a dense, hexagon-based topology in a large open lab.

Extensive evaluation (§5) shows LAVA is effective. Our power sensing mechanism is promising despite the low-quality part. Multi-hop amplify-and-forward in the ceiling testbed provides signal power and throughput improvements to individual Wi-Fi and ZigBee links alike, by 50% on average and up to more than 100% for Wi-Fi links at 15 dBm transmit power, 193% on average and up to more than 8× at 0 dBm (these transmit power levels capture the specification diversity of IoT devices (Appendix §B)), and up to 17 dB for ZigBee links. The performance at 0 dBm with LAVA, compared to using 15 dBm, shows the *power offloading* potential with LAVA. For pairs of co-channel concurrent links, LAVA improves the average throughput at 0 and 15 dBm transmit power by 517% and 80% respectively. A dense topology facilitates co-channel spatial reuse. Further, comparative results from the two testbeds highlight the curse of vertical offsets. LAVA can easily mitigate multipath fading without further phase adjustment if both endpoints are in the same plane as the elements, though this is more challenging in 3D.

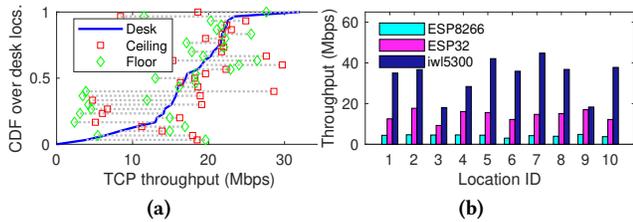


Figure 2: TCP throughput variation: (a) ESP32 across locations; (b) Three devices roughly co-located.

To summarize, this paper makes the following contributions: First, we highlight the effects of non-uniform 3D signal propagation in indoor IoT deployments. We expect this to be more complex in upcoming challenging outdoor scenarios. Second, we propose an infrastructure setup, LAVA, to provide multiple, fine-grained control points of signal propagation and energy distribution, achieved with simple, frequency-agnostic multi-hop amplify-and-forward. Adding an optional phase shifting stage further permits LAVA to operate in two modes: either as a relay mesh (using amplification), or for beamforming (using phase alignment). We believe LAVA represents another direction in the design space for reconfiguring the wireless environment. Third, we present a new network architecture to passively monitor communication between the endpoints and infer network conditions with a few assumptions. This takes a step towards a separate ambient infrastructure, decoupled from the standard wireless deployment. Fourth, we implement the first large-scale and distributed “smart surface,” by experimenting with a ceiling testbed in real offices and compare it to a floor testbed in a large open lab. Extensive experiments elucidate both the potential and challenges for such a ceiling deployment. This work does not raise any ethical issues.

2 PROVISIONING 3D COVERAGE

While indoor Wi-Fi coverage has received significant attention, the devices concerned used to revolve around laptops, smartphones, and tablets. These operate at or near the desk level, where the human user is. Therefore, the coverage considerations used to be approximately confined to a single plane. In contrast, IoT devices today are more pervasively deployed, from floor to ceiling, which exacerbates the coverage challenges previously described. The situation can be more complex for challenging outdoor scenarios, such as ocean IoT [11, 23, 35, 49, 62] and space IoT [27, 47, 48, 57, 61].

Non-uniform 3D coverage. Figure 2(a) illustrates the extent of non-uniform coverage horizontally and vertically. We place an ESP32 module in 30 different horizontal locations, at 3 vertical offsets (floor, desk, or ceiling level), in or just outside Office 1 (Figure 1). The ESP32 receives TCP traffic from an AP on the ceiling in Office 1, sent at 15 dBm transmit power. We can see that the link quality varies significantly across locations both vertically and horizontally. Most notably, there is no definitive ranking between the three vertical offsets. Non-uniform propagation behavior is already well documented for 2D coverage, and Figure 2(a) is simply a manifestation of this in 3D. The office furniture further interacts with signal propagation in complex ways. Perhaps less obviously, since most endpoints use linearly polarized antennas, vertical signal propagation is more affected by the beam pattern (mis)alignment.

Diverse device capability. To add to the inherent variability in the signal propagation behavior, multiplicity in the wireless standards, the signal processing algorithms, and hardware capability naturally diversifies performance capability. Figure 2(b) shows a simple example of TCP throughput variation across three Wi-Fi devices. An ESP32 module, an ESP8266, and an Intel iw15300 are placed roughly side by side around Office 1 and connected to the same AP. The ESP8266 is limited by its processing capability to about 5 Mbps, whereas the iw15300 can support 3×3 MIMO. ESP32 is used in some of the IoT devices we came across.

Multi-standard co-existence. The rise of smart-X scenarios has ushered in a vast number of often small IoT devices using various wireless standards (some listed in Appendix B). For examples, we found two smart cameras using Z-wave and Wi-Fi respectively, while smart light bulbs may use Bluetooth, ZigBee, or Wi-Fi.

2.1 Requirements for 3D Coverage

To cater to diverse link scenarios above, infrastructure support faces several new requirements.

First, we need to fine-tune energy distribution during signal propagation, at a finer granularity than radiating all power from the transmitter alone. This motivates power adjustment along the propagation paths, which should also mitigate SNR loss due to antenna beam pattern misalignment between endpoints in 3D. Second, performance metrics differ by device, but individualized infrastructure support is impractical for large-scale heterogeneous IoT deployment. This requires versatility in a shared infrastructure. While throughput improvement is a common performance optimization goal, low-end, battery-powered IoT devices benefit most from power provisioning from the infrastructure to prolong battery life, e.g., by reducing the transmit power or duration and/or increasing sleep time. This motivates dense deployment of an “access infrastructure” near the endpoints. Third, since networks operating over different wireless standards nevertheless face similar uneven signal propagation issues, reusing infrastructure support across standards is highly desirable for scalability. This motivates signal-level optimizations so as to be agnostic to even physical layer formats and operations specified in standards. Multi-standard co-existence issues, such as cross-technology interference, should be handled simply as concurrent link setups.

Collectively, these point to a solution that augments the propagation environment densely, capable of injecting power at the signal level, but untethered to any communication endpoints.

2.2 Design Considerations

Long or heavily attenuated links benefit from signal amplification *between* the communication endpoints. Conceptually, the simplest approach is amplify-and-forward [7]. Practically, however, this strategy forces a balancing act between signal and noise amplification. When applied to large networks, this necessitates multiple amplifying relays, with the concomitant problems of placement, coverage, and inter-relay interaction that require careful planning and/or real-time adaptation to dynamic link conditions. We next explore design considerations to balance effectiveness and complexity.

The case for reduced amplification. An amplify-and-forward relay amplifies both signal and noise. This increases the perceived

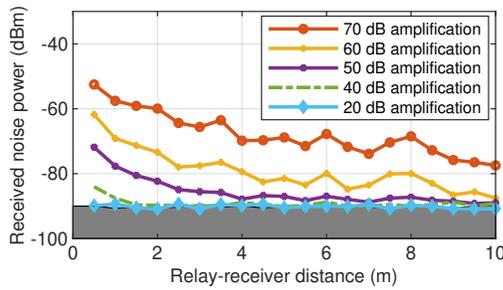


Figure 3: Noise amplification over distance.

noise at the receiver if the distance from the relay to the receiver is insufficient to attenuate the amplified noise to below the ambient noise floor. Based on channel measurements, prior work tunes down the relay’s amplification level to avoid this situation [7].

To study this relationship, we experiment in a large open lab environment (§5) in the 2.4 GHz band, connecting a directional patch antenna (§4) directly to an amplifier via a low-loss cable. A USRP receive antenna is placed some distance away to measure received background and relay noise. Figure 3 shows that once the relay’s amplification gain reaches 50 dB, the receiver must be at least 4 m away to reduce amplified noise to 2–3 dB above the noise floor. Thus, the higher the amplification gain, the larger the *exclusion area* needed around the amplifier to avoid significant noise amplification. Moreover, an *intra-relay* positive feedback loop forms when a relay’s amplifier gain exceeds the net path loss from its output to its input port (the relay’s *reverse path*). A 60 dB amplifier gain requires a 30 cm separation (between directional antennas) to avoid feedback, but a 1 cm separation suffices for 20 and 40 dB gains. Finally, any noise amplification above the noise floor is vulnerable to further amplification by nearby relays, which could eventually rise to reduce end-to-end SNR. Therefore, to support a *dense* coverage area, the amplification gain must be capped.

The case for a multi-hop mesh. Low-gain amplifiers have limited range. Taking a leaf out of wireless mesh networking, we organize the amplifiers into a multi-hop mesh. Each multi-hop chain can be viewed as an array *in series*, in contrast to multi-antenna arrays that have the antennas operate *in parallel*. A multi-hop relay also promotes a dominant direct path between successive hops. When the signal on that path overpowers any other multipath effects in the environment, per-hop amplification largely overcomes potential phase misalignment due to multipath propagation. This makes an active relay mesh topology less susceptible to channel fluctuations than prior approaches [3, 38, 73].

Incorporating phase alignment. When multi-hop amplification does not dominate over other signal components, phase alignment between multipath propagation can be helpful. Phase shifting has proven effective in previous surface designs [17, 38]. However, there are three main cons: First, phase alignment requires explicit channel information feedback from the endpoints to the control plane, which is not easily supported on off-the-shelf IoT devices; Second, with many phase settings per hop and several hops per LAVA route, the search complexity increases exponentially; Third, phase shifting effects are sensitive to the carrier frequency. We need a simple heuristic to balance decision complexity and performance.

Summary. Combining reduced relay amplification, a multi-hop relay mesh, and, optionally, phase-shifting, we have the key features of LAVA. The main challenges then are careful topology construction and a control algorithm to identify the most suitable routes through the array, ideally without coordinating with endpoints.

3 LAVA DESIGN

LAVA is a network of directional, multi-hop amplify-and-forward relays, which we call *elements*. This section describes the element hardware design, the topology construction, the control plane architecture and algorithms.

3.1 RF Element Hardware

Minimalist amplify-and-forward unit. For a minimalist design (called a *simple element*), the idea is to place an amplifier between two directional antennas to serve as a basic amplify-and-forward relay (Figure 4). We use a fixed and low-gain amplifier for this unit. A fixed gain simplifies the control to an on-off switch, and a low gain helps to avoid positive feedback and noise amplification.

We use directional antennas for better power gain and to enforce the signal direction. Since directional antennas have limited beamwidths, we need multiple directional antennas per element for signal coverage in all directions. Alternately, we can place multiple simple elements next to one another to cover all directions. The former requires a more sophisticated element *circuit*, whereas the latter shifts the complexity to the element *layout*.

Multidirectional element design. The idea behind a multidirectional element is to mimic the signal coverage, as if omni-directional antennas were used at both the input and the output. Therefore, this requires integrating multiple simple elements into one circuit and switching between the input directions and between the output directions. Our multidirectional element design (Figure 5) achieves this by enabling the RF path between any of up to four input antennas and any of up to four output antennas. All antennas are oriented towards different cardinal points to route signals.

Each multidirectional element has a basic power meter to detect input signal power, generating reports for the LAVA controllers (§3.3). We also include an optional phase shifter before the output RF switch (omitted in Figure 5).

3.2 LAVA Topology

Since individual elements alone are not sufficient to cover a broad area, we need multiple elements organized into a suitable topology to support multi-hop routes in a flexible way. There are two main considerations: the number of possible next hops for each element and the inter-hop spacing. The former is determined by the antenna configuration of each element. Given the multi-directional element design, we line up elements horizontally as shown in Figure 1. This planar, grid-like topology permits at least one route between any two antennas in the mesh, thus enabling signal flow through LAVA in any direction within its coverage area. Strictly speaking, this is only optimal if each directional antenna has a horizontal beamwidth of 90° . However, we find empirically that using four directional antennas, each with a 60° beamwidth, is already sufficient. Adding

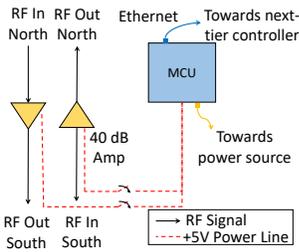


Figure 4: Simple element design.

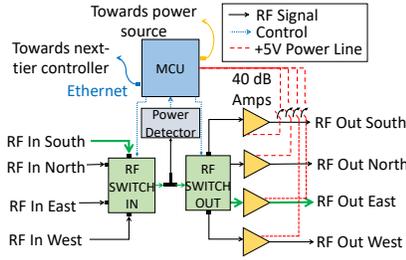


Figure 5: Multidirectional element.

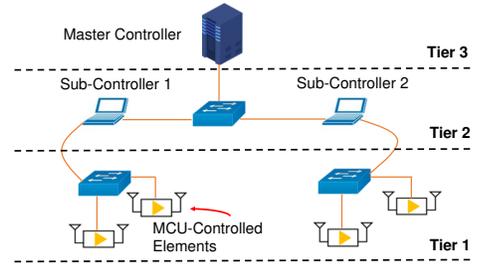


Figure 6: LAVA control architecture.

vertical hops can be more difficult in real environments; hence, we rely on the vertical antenna beamwidth for vertical coverage.

For inter-hop spacing, the goal is to maximize the signal received by the next hop such that each additional hop simply preserves the SNR seen at the first hop on the path, all while avoiding noise amplification. This corresponds to a distance over which the amplified noise power drops to the background noise level. Given our use of low-gain amplifiers, this distance is very short and can be empirically set to a fixed value based on the per-element amplification gain (§5.1). The inter-hop spacing defines the *spatial resolution* of our testbed for power detection.

Control architecture. LAVA follows a 3-tier hierarchical controller network architecture for scalability (Figure 6). The lowest tier consists of dedicated microcontroller units (MCUs) acting directly on the elements. There is one MCU per element, to set the *states* of its element (i.e., the input and output antennas and whether to turn on the amplifier) and report power measurements to the next tier up. The middle tier consists of sub-controllers, each connected to a disjoint subset of the MCUs based on physical proximity. The sub-controllers collect and filter power measurements from the MCUs and forward commands directly to the MCUs. The top tier is a single master controller whose role is to analyze all the power reports forwarded by the sub-controllers, compute the route through the LAVA mesh, identify the elements on the route to activate, and send state change commands only to the MCUs for those elements via the sub-controllers. The LAVA elements form a distributed monitoring network, while the master controller collects sensing measurements and runs most control plane algorithms.

3.3 Control Plane Action

The overall goal of the LAVA control plane is to (a) identify the elements that are “closest” to the communication endpoints in terms of the power received from or delivered to those endpoints, and (b) activate a route between them (Figure 7). We refer to the element that receives the most power from the transmitter endpoint as the *entry element*, and the one that is expected to deliver the most power to the receiver endpoint as the *exit element*. To identify these elements, we design a power-based frame detection mechanism for each element to detect an active signal power source nearby. This information is reported to the LAVA control plane to infer link-level information for the endpoints and activate the appropriate elements. Given that LAVA operates at a signal level, it is agnostic to the layers above the physical layer. The controller assumes some link relations, but does not establish or maintain them.

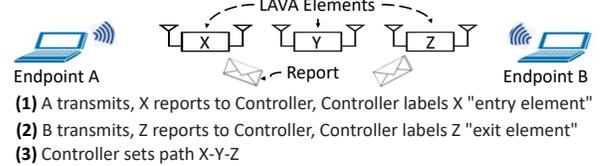


Figure 7: Identifying *entry* into and *exit* from LAVA.

Per-element power monitoring. Each element listens for an active signal source nearby and infers the presence of frames as an indication of an active traffic session. This implies an active endpoint transmitter within detection range, and lets us identify the entry and exit elements. The power detector in each element generates a voltage reading on a logarithmic scale. This voltage *decreases* as the input signal power *increases*. The MCU samples this voltage every 16 μ s for a duration of about 7.2 ms and creates a signal profile. For example, Figure 8(b) shows a detected signal profile, corresponding to the input signal shown in Figure 8(a), with the voltage drop in Figure 8(b) matching the input signal increase. This raw profile is then passed through slow and fast exponentially weighted moving averages (EWMA) for edge detection (Figures 8(c)(d)) Finally, we assess whether the gaps between adjacent edges are consistent with legitimate frames, as opposed to hardware glitches or other sources of noise. If there is a signal present, the MCU saves the *voltage index*, which is defined as the difference between the highest and lowest voltages in the signal profile (after filtering). The higher the input power, the higher the voltage index.

Each profile with a detected signal is then labeled with a *positive* flag, or *negative* otherwise. When an element observes 3 consecutive positive profiles, it begins reporting voltage index measurements to its direct subcontroller. The subcontroller then forwards these measurements to the master controller for analysis, unless the reports are from an intermediate element of an activated route. In that case, the power measurements will be unreliable and will be discarded by the subcontroller. Once reporting starts, if the MCU then observes 5 consecutive negative profiles (i.e., no signal present), it notifies the subcontroller, stops sending measurements, and resumes monitoring for an active signal. The numbers of consecutive profiles needed, 3 and 5, are chosen very conservatively due to a hardware artifact in our current prototype. The off-the-shelf power meters produce occasional hardware glitches. To guarantee reliable operations, we need to operate at a slower timescale (§5.2), which currently limits the *temporal resolution* of our testbed. This can be improved through better engineering and testing of the parts, and is not an inherent limitation of LAVA.

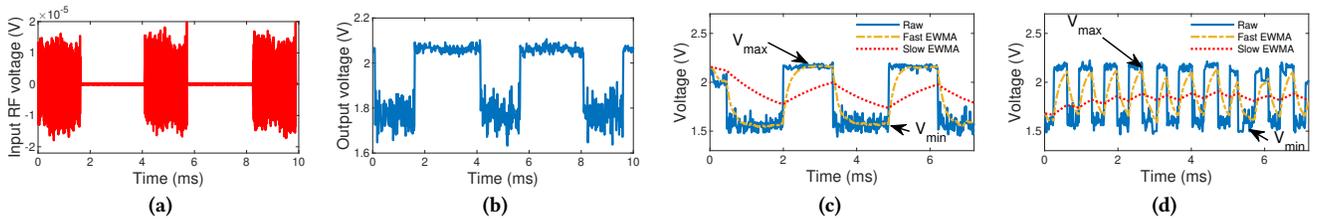


Figure 8: (a) Input signal and (b) the corresponding voltage measured by the logarithmic power detector during a long train of 1500-byte frames. (c), (d) Voltage index generation for 1500-byte and 300-byte OFDM frames, respectively.

“Entry” into and “exit” from LAVA. We cannot sense a completely passive receiver. Therefore, we assume bidirectional, half-duplex traffic. This is reasonable as TCP traffic is common and runs in both directions. Initially, neither the entry nor the exit element is identified (Figure 7). The master controller listens for reports, and when frames are detected and reported by an element, that element becomes the initial *entry* element. When another element sufficiently far away detects and reports frames, the far element becomes the initial *exit* element. Given nearby elements often detect signals from the same endpoint transmitter and their reports may conflict, we cannot reliably resolve different transmission sources beyond the temporal and spatial resolutions of our testbed. I.e., the two endpoints should be at least two hops (roughly 2 m) apart, and their transmissions separated by at least three positive power profiles (50 ms, numbers determined empirically in §5.2). This means the entry and exit elements cannot be the same (but can be adjacent), and LAVA does not support very short links this way. Short links benefit more from phase alignment along multiple paths than signal amplification (discussed below and in §5.2). Conversely, the endpoints need not be in range of each other, provided higher layers can establish and maintain links, e.g., via manual configuration, without relying on physical or link layer proximity detection.

Typically, multiple elements will detect the same transmission, so the master controller needs to identify the best entry element. Initially the first reporting element is regarded as the entry element candidate, because at this point the controller has no way of knowing if another report is going to arrive. If a subsequent report arrives within a short time, showing a sufficiently “higher” voltage index, the new reporting element will replace the current entry element. We discard additional reports based on the inter-report interval across elements, derived empirically in §5.2. Any reports within this interval are assumed to detect the same transmitter.

Finally, to support concurrent transmissions, our algorithm requires the same temporal and spatial separations between transmissions for distinct links. Once the entry/exit elements of the first link are identified, LAVA monitors the environment for reports sufficiently far away from the activated elements and repeats the identification process for the second link.

LAVA route computation. After identifying a pair of entry and exit elements, the master controller finds the shortest route between these two. Given the small number of elements in our testbeds, our current implementation simply enumerates all viable routes and selects the shortest. A larger testbed can employ standard algorithms like Dijkstra. The master controller then maps the route to suitable states for each element on the route. This process happens once for every active session between the endpoints. In our design, each hop

roughly compensates for the signal power lost from the previous hop, so more hops would offer little benefit. This is another manifestation of the inter-hop separation; we need consecutive elements turned on to counter multipath effects by amplifying a dominant path. This in effect supplants signals along other paths.

We determine the input antenna on the entry element and output antenna on the exit element heuristically. To activate the route, the master controller instructs each element along the route (via the subcontrollers) to receive and relay on the appropriate input/output antennas. Hence, the entry element reports on whichever antenna is set for the route. When neither the entry nor the exit detects a signal source, the LAVA route is deactivated, and the LAVA system returns to its initial passive monitoring state. For concurrent links with disjoint endpoints, frames detected at least two hops away from active routes are considered to be from new endpoints, and a disjoint path is enabled for each new pair.

Phase shifts. Sometimes enabling a LAVA route does not provide the best SNR improvement, typically if only two hops are used. In these cases, multi-hop signal amplification alone is not sufficient to supplant other propagation paths.

Hence, for a short LAVA route, a final, optional stage of our algorithm involves a short explicit-feedback session between the receiving endpoint and the master controller. The controller iterates over four phase shifts (every 90°) at the exit element. The endpoint receiver reports the corresponding SNRs to the controller, which then selects the best phase shift. This strategy is empirically shown to achieve near-optimal performance at low decision complexity (§5.2).

4 IMPLEMENTATION

LAVA elements. For easy testing and flexibility, all LAVA elements are built with loose parts. The circuitry can be integrated into a printed circuit board (PCB) to ease future deployments. We use 37 dB ZX60-2531MA+ amplifiers for 2.4 GHz from Mini Circuits, each requiring a 5 V power source and between 102 and 120 mA of current. We also use directional 2.4 GHz/5 GHz dual-band indoor patch antennas [51], with 8 dBi gain at 2.4 GHz, linear (vertical) polarization, an azimuthal beamwidth of approximately 66° and a vertical beamwidth of around 16° . Since our amplifiers support a much wider frequency band than the antennas are optimized for, we add a bandpass filter [5] to filter out spurious signals above or below 2.4 GHz.

For a *simple element*, each amplifier is only connected to two directional antennas and a bandpass filter. For a *multidirectional elements*, recall that we add two RF switching stages to determine the incoming and the outgoing signal directions respectively (Figure 5).

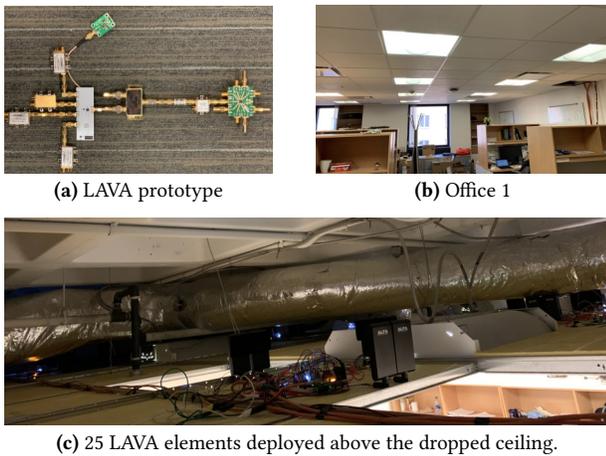


Figure 9: (a) Multidirectional LAVA element prototype (RF circuitry); (b) the large office space below the LAVA testbed; and (c) the LAVA elements above the dropped ceiling.



Figure 10: A floor testbed of 38 simple elements.

Figure 9(a) shows an assembled prototype. To achieve this, we use a non-reflective SP8T RF switch with a 2 dB insertion loss [60], with directional antennas [51] on four of its ports, each oriented in a different direction. The output of this switch is then amplified by a 20 dB amplifier to compensate for signal loss incurred in subsequent stages of the circuit (3 dB loss from a power splitter and 7 dB insertion loss from our custom board). The amplified signal is then passed through the bandpass filter and bifurcated with a power splitter. One branch flows to a logarithmic detector [44] to compute the voltage index, and the other to the output switch.

For the output stage, we use a custom-made board integrating an SP4T RF switch [59] and a 6-bit digital phase shifter [52]. The board first phase-shifts the signal and then routes it to one of the four output ports of the switch. Each is then connected to an amplifier and a directional antenna. The amplifiers have to be connected at each output port of the RF switch; if fed to the input port instead, the amplification would overcome the isolation provided by the RF switch (the lowest being 23 dB, less than the 37 dB amplification gain), causing signal leakage on all ports.

Per-element control. Each element is controlled and powered by an Arduino MCU with an Ethernet shield [18] and a TFTP bootloader to allow remote upgrade of the Arduino code. The Arduino MCU’s I/O pins can provide up to 40 mA of current, not enough to power the 37 dB amplifiers. Instead, we use the SPI interface in

the Arduino MCU to drive the highs and lows of an 8-bit serial-in, parallel-out shift register [56]. Its pins in turn drive a CMOS transistor array with 8 circuits [65], each delivering up to 500 mA. These circuits serve as power supplies to amplifiers, and the state of the shift register (set by the MCU) determines which amplifier is on (if any) at any given moment. Finally, the digital I/O pins in the MCU control the phase shifter and RF switches. An analog pin in the Arduino MCU reads the output voltage of the logarithmic detector [44]. Internally, the MCU computes the voltage index of the raw voltage measurements and reports back to the controller.

Control infrastructure. The elements are managed by two sub-controllers, both connected to a master controller. All sub-controllers are shuttle PCs DH110 with an Intel Core I7-7700 Quad Core processor and 32 GB RAM. The controllers and the Ethernet-enabled Arduino MCUs are connected to a 48-port 100 Mbps switch. This wiring is largely hidden from view above the dropped ceiling, and is seamlessly integrated with the infrastructure already in place for regular network operations. Note that we could easily integrate the LAVA network into existing production networks, but we decide instead to keep a separate subnet for experimentation purposes.

Power distribution. To power the elements, we use 12 V, 0-30 A DC power supplies, connected to fuse blocks [21] and all elements via 18 AWG gauge electrical wires. We could use the power outlets in the ceiling space, but decide instead to keep our testbed separate from the regular production infrastructure. Power over Ethernet (PoE) would be another option to power the elements and simplify future deployment.

Testbed deployment. Since each LAVA element embeds power sensing capability, the overall mesh needs sufficient spatial coverage to ensure both adequate detection of active signal sources and signal power amplification. A ceiling deployment meets both requirements, and can also route signals around furniture. This is in contrast and complementary to existing smart surface testbeds [4, 14, 15, 17, 38], large or small, where the surface acts like a side wall, with implicit antenna polarization and orientation.

We deploy 25 multidirectional LAVA elements above the secondary, dropped ceiling spanning three actively used offices and the corridor (Figure 1), logically forming two sets of collaborative LAVA arrays. This ensures the testbed is largely hidden from view (Figure 9(b)), with roughly one element next to each set of lighting fixtures (Figure 9(c)). Further, the ceiling space is mostly empty other than housing pipes and cables, and imposes few constraints to constructing the ideal grid topology. This is our default testbed, mimicking a real deployment. We also set up a floor testbed of simple elements in a more complex, hexagonal topology (Figure 10). Routes through the testbed are manually set or automatically selected from a pre-configured list.

5 EVALUATION

We validate LAVA design decisions with microbenchmarks and then study large-scale link-level performance. Note that we cannot directly compare LAVA to existing surface prototypes. They are designed for different operational settings, and do not fully support concurrent links. Instead, we compare the relevant signal manipulation *strategies*, and defer a qualitative comparison to §6.

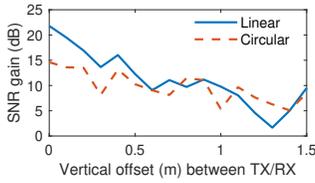


Figure 11: Vertical coverage by different antennas.

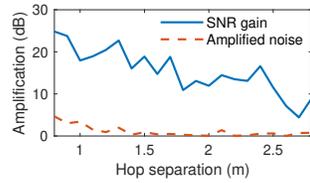


Figure 12: Amplification performance over inter-hop distance.

Setup. We deploy Wi-Fi (USRPN210s and WARPv3) and ZigBee nodes (XBee S1 motes) around the LAVA testbeds. Unless otherwise stated, all the microbenchmarks are run with two WARPv3 forming a single-antenna Wi-Fi OFDM link, and all experiments in §5.3 involving WARPv3 boards adopt the phase-shifting stage of LAVA. This is the only endpoint platform that can communicate with the LAVA control plane for phase selection. To measure throughput, we use Intel 5300 Wi-Fi cards, ESP32, and ESP8266 modules. Their transmit power levels are set at 0 dBm and 15 dBm to sample varied IoT device specifications (Appendix B). The former matches battery-operated devices, and the latter typical Wi-Fi clients.

5.1 LAVA Element and Topology Design

Element-wise vertical coverage. With a planar deployment, LAVA’s vertical coverage depends on the vertical beamwidth of individual antennas. However, there is a tradeoff between the beamwidth and power density, given the same amount of power radiated from an antenna. We consider two types of antennas: the (linearly polarized) patch antenna we use, and a circularly polarized patch antenna designed for 2.4 GHz, with 8 dBi gain, 65° vertical and azimuthal beamwidths. We place two endpoints at either side of a simple LAVA element, whose input antenna is 0.3 meter from the transmitter and whose output antenna is 1.5 meters from the receiver. The transmitter and the element are both on the floor, but the receiver is placed at vertical offsets ranging from 0 to 1.5 meters from the floor. For each receiver elevation and each type of antenna, we measure the SNR gain from this element at the receiver (Figure 11). The linearly polarized antenna (15° vertical beamwidth) outperforms the circularly polarized antenna at lower vertical offsets and both achieve comparable performance for higher offsets. The circularly polarized antenna induces more performance variability due to constant polarization rotation [9, 50]. Therefore, we use the linearly polarized antenna despite a smaller vertical beamwidth.

Hop separation. Next, we estimate the optimal LAVA inter-hop distance mentioned in §3.2. We line up two LAVA elements between the transmitter and the receiver, all four in a straight line of sight. The endpoint transmitter and the first element are placed in fixed locations, 30 cm apart, ensuring roughly constant received power at the first hop. The distance between the second element and the receiver is also fixed at 1.5 m. We then vary the distance between the two elements (i.e., moving the second element and the endpoint receiver together). For each hop separation, we measure the SNR gain at the receiver with the two-hop relaying over without (Figure 12). Although seen even at 2.8 m between the two LAVA elements, the SNR gain becomes less predictable, because a larger

hop spacing is more susceptible to multipath effects while reducing the signal power collected by the second hop.

We repeat this experiment for other setups and reach the same qualitative conclusion, i.e., *the optimal hop spacing appears to be where the signal amplification at each hop compensates for the path loss from the previous hop*. As noise degrades most slowly in the short line of sight setup reflected in Figure 12, these results suggest an inter-hop spacing of 1.5 to 2 m in our prototype construction.

5.2 Control Plane Microbenchmarks

The ability of LAVA to detect and aid an ongoing transmission depends on the accuracy of the voltage index reports and link inference mechanism. We next evaluate the sensitivity and effectiveness of these decisions.

Detection sensitivity to inter-frame interval. We first study an element’s ability to detect a valid frame (i.e., structured signals) for different inter-frame intervals. A signal source transmitting 1500-byte OFDM frames is connected via an attenuator to a random input port in an element, with an incoming SNR of 30 dB. We then vary the inter-frame time interval and capture 100 consecutive voltage profiles, each lasting 7 ms, and compute the voltage indices. A measurement is deemed successful if it generates an accurate voltage index. Figure 13(a) shows the accuracy of LAVA element reports for an ongoing transmission as a function of the inter-frame time interval. There is a sudden drop in detection accuracy when the inter-frame interval exceeds 2 ms, i.e., when fewer than two frames can be unambiguously discerned. Conversely, our voltage indices are reliable provided at least two frames are found within the 7 ms window.

False positives (FP) and false negatives (FN). We next assess the accuracy of *individual* voltage index reports. A false positive occurs when a non-zero report is received (i.e., signal detected) despite the absence of an active transmission. Similarly, a false negative occurs when a report of 0 is received (i.e., no signal detected) when an active transmission is occurring. *FP* and *FN* reports occur due to background noise and hardware glitches.

We place the transmitter near 10 different elements, and verified the rate of false positives and false negatives when the transmission was turned on and off, respectively. The element closest to the transmitter classifies 200 consecutive voltage profiles, each covering 7 ms. Figure 13(b) shows a maximum *FP* rate of 11.5% and a maximum *FN* rate of 4%. To compensate for the high *FP* rate per voltage profile, therefore, an element only reports to the controller on a perceived active transmission after capturing 3 consecutive non-zero voltage index profiles. Similarly, it reports on the completion of a transmission only after capturing 5 consecutive negative voltage index profiles. This conservative reporting strategy reduces the false positive and negative rates from any element to zero. More importantly, by requiring 3 consecutive positive reports, each highly suggestive of the presence of a structured signal, the control plane is unlikely to react to sporadic signals (like Wi-Fi beacons) or sustained, unstructured transmissions (like microwave blasts).

Sensitivity to input SNR. We next evaluate the sensitivity of voltage index measurements to the input SNR on each element and the detection consistency across multiple elements. Variations in these measurements can arise from hardware and operational

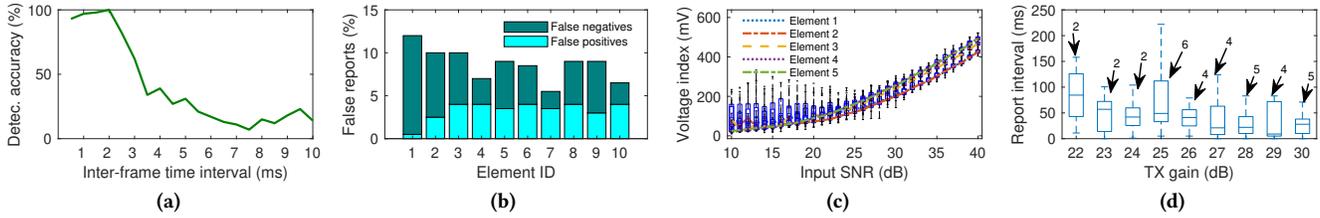


Figure 13: Power sensing microbenchmarks. (a) Detection accuracy vs. inter-frame time interval. (b) Power detection false positives and false negatives over some time window on each element. Due to noise and hardware glitches, individual reports can be inaccurate; Therefore, an element only sends power reports after detecting 3 frames, eliminating false reports. (c) Detection consistency across input SNRs and elements. (d) Measuring the inter-report interval to filter conflicting reports; The number of reporting elements is shown above each arrow.

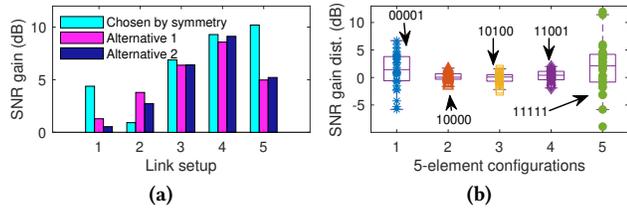


Figure 14: (a) Validating the symmetry of the LAVA routes for 5 links. (b) 3D multi-hopping dynamics for 30 links; each boxplot shows the SNR gain across 30 links for a particular 5-element on-off configuration represented by the adjacent binary number, with the “1” bits showing the “on” elements in the 5-element chain.

differences in the circuit parts. In this experiment, a signal source transmitting 1500-byte OFDM frames is connected via an attenuator to a random input port in an element. We then vary the TX gain of the signal source to generate input SNRs from 10 to 40 dB for the LAVA element. For each SNR, we collect 100 voltage index reports, each captures signal profiles for 7 ms. We then repeat this process for 4 other randomly selected elements. Figure 13(c) shows reported voltage index distribution. The spread of the measured voltage corresponding to 10 and 40 dB input SNRs is 492 mV. But for any given SNR level, the median voltage index reported by any two elements varies by at most 68 mV, or 14% of the spread, validating the consistency of the measurements across elements.

Inter-report arrival time. Multiple elements may detect the same transmitter, and more than one element reporting on the same transmitter might mislead the control algorithm. Therefore, we investigate how to filter out spurious reports. We place a USRP in the center of Office 1, vary its transmit power level (TX gain), and count voltage index reports from any elements in a 5 s window. We typically set TX gains between 20 and 30 dB, producing 5 to 50 dB per-element input SNRs in our ceiling testbed. A reporting element might send multiple reports during this period. A higher TX gain solicits more reporting elements. Figure 13(d) shows the distribution of inter-report arrival times for each transmit power level and the number of reporting elements (indicated by arrows). Two elements report at 22 dB, compared to five at 30 dB. For any TX gain over 23 dB, the median inter-report interval is below 50 ms. Therefore, 50 ms indicates the temporal resolution of power detection of distinct transmitters. Successive reports arriving within this window are deemed to indicate the same active transmitter.

LAVA route symmetry. Our control algorithm assumes that the entry element for a link should also be the exit element for the same

link in the reverse direction. To validate this, we fix one endpoint location, move the other to 5 different locations, and identify the closest LAVA element for the moving endpoint. Then, we enable the shortest route (*ideal*) between this element and the element closest to the fixed endpoint. Separately, we try *alternative* routes identified through visual inspection, and observe that in 4 out of 5 links, the *ideal* route outperforms the alternatives (Figure 14(a)).

Multi-hop dynamics. To understand additional element selection strategies to route signals through LAVA, we next investigate multi-hop dynamics. 5 LAVA elements are placed on stands to mimic the ceiling deployment, and the receiver endpoint is placed on the floor in an area of roughly 4 by 6 m near the 4th and 5th elements in the chain. The transmitter is beyond the first element on the same level as the elements. We then turn on different combinations of elements in the 5-element chain (Bit 1 in the binary numbers in Figure 14(b) corresponds to an element being on), and measure the perceived received signal quality at the endpoint receiver for 30 different receiver locations.

This setup generates both a strong, direct line-of-sight path between the endpoints, as well as multiple amplified paths through one or more activated LAVA elements. LAVA shows most benefit when an amplified path dominates over the others and may hurt the performance otherwise. A complete *chain* (all elements on) enabled between endpoints outperforms incomplete chains (one or more elements off), seen in 8.9 dB gain on average in the endpoint receive SNR. All elements on corresponds to the optimal inter-hop spacing, maximizing the amplification benefit while minimizing multipath effects. This translates to activating all successive elements on a route. Conversely, single-hop amplification (the first two boxplots) has little effect unless that hop is very close to the transmitter. Short links spanning one or two elements suffer from multi-path effects more often than benefit from the amplified path.

We also repeat the experiment with endpoints and LAVA elements all on the same level (mimicking the floor testbed). Any amplified path dominates over the direct path between endpoints easily now. While the best performance is still seen with all elements on, turning off any element rarely hurts.

Adding phase shifts. Since the ceiling deployment may not completely counter multipath effects, especially for short links, we explore adding phase alignment between multipath propagation (§3.3). We enable 3-element routes in the ceiling testbed for 25 different links. For each element on a route, we first consider 16 phase settings (every 22.5°); hence, there are 16³ (= 4096) possible phase combinations for each 3-element route, and we find the best

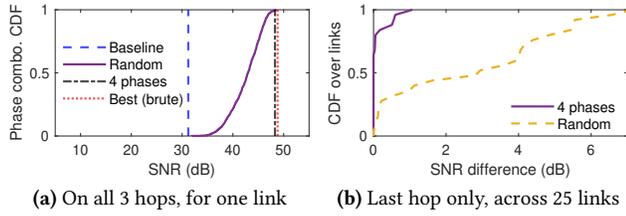


Figure 15: Effects of phase shifting. Choosing the best out of 4 phase settings on the last hop is within 1 dB of the optimal.

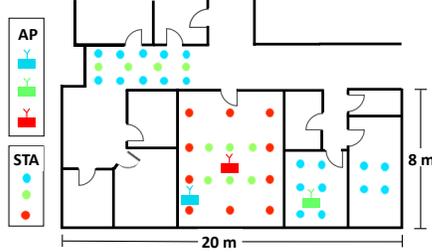


Figure 16: AP and client (STA) locations around the ceiling testbed.

and worst received SNRs after enumerating these 4096 settings. Second, we randomly pick a 3-element phase combination at a time (equivalent to using no phase shifters), repeating 4096 times. Third, we only adjust the phase shift on the last hop of the 3-element route, iterating over 4 phase settings only, and selecting the best.

In the link shown in Figure 15(a), choosing random phase shift settings on all three elements can provide anywhere between 1 and 20 dB gain. Choosing the best out of 4 settings (0° , 90° , 180° and 270°) achieves within 1 dB of the optimal gain. The same observations apply across all links (Figure 15(b)). Therefore, selecting one phase setting out of 4 possibilities and only on the last hop of a route can deliver additional benefit at low complexity.

5.3 Large-Scale Experiments

Our large-scale experiments are set up according to the spatial and temporal resolutions in our testbeds (explained in §3 and empirically identified above). We use *iperf* to generate UDP and TCP traffic for the Intel 5300 cards, ESP32 and ESP8266 and use the WARPLab environment for the WARPv3 boards. Due to the hardware quality limitation of the power meters in our prototype, we set the inter-frame time to around 2 ms. For comparison between UDP and TCP, we generate bidirectional traffic sessions for both, with the reverse direction starting 50 ms after the forward direction.

Individual Wi-Fi and ZigBee links. We deploy 40 WARPv3 Wi-Fi links around the LAVA ceiling testbed area. To capture 3D coverage effects, one endpoint is placed on the ceiling (AP), and the other ~ 2.5 m below it at the locations shown in Figure 16 (STA). We also set up links in the floor testbed at 30 different locations with the USRP N210s. Figure 17(a) shows a median per-link SNR gain of 14 dB from our floor testbed, but 11 dB from the ceiling testbed. The maximum gains are 25 dB and 18 dB, respectively. Next, we place ZigBee nodes at 20 different locations (blue dots in Figure 16). Figure 17(b) shows LAVA provides a median RSSI gain of 6 dB and up to 17 dB. In contrast, across 30 different ZigBee links, the floor testbed provides a median gain of 14 dB and up to 31 dB. These results show both (SISO) Wi-Fi and ZigBee benefited similarly when

signals mainly travel in one plane, since LAVA is standard-agnostic. However, the vertical offsets for the ceiling testbed and experiments present more challenges, and exacerbate antenna quality difference between the Wi-Fi and ZigBee nodes.

Throughput across Wi-Fi devices. We deploy 40 Wi-Fi links with the Intel 5300 cards, as shown in Figure 16, and measure TCP and UDP throughput with and without LAVA at 15 and 0 dBm transmit power. The median and maximum throughput gains are 25% and 192%, respectively, at 15 dBm, and $2\times$ and $8.7\times$ at 0 dBm (Figure 17(c)). Unsurprisingly, LAVA performs particularly well for low-power communications. Figure 17(d) further shows the percentage of *throughput realized*, i.e., the throughput using 0 dBm transmit power with LAVA, as a percentage of the throughput at 15 dBm without LAVA. Half the links achieve 75% or more of the throughput at 15 dBm. This suggests that LAVA can provide *power offloading* support.

Finally, we place ESP8266, ESP32 and iw15300 clients side by side at 10 locations (different from those in Figure 16) and a single iw15300 endpoint on the ceiling to which they connect. The median throughput increase for each is 6.5%, 34% and 16% (Figure 18(a)). With its rate capped at 7 Mbps, the ESP8266 receives the least throughput benefit from LAVA, but can leverage reduced transmit power. Devices like the ESP32 (single-antenna, 40 MHz channels), widely adopted in commodity IoT solutions, can benefit the most.

Mobility support. We place a WARPv3 receiver on a trolley in the hallway, and move it at 1 m/s. The WARPv3 transmitter is fixed at the center of Office 1, and we disable phase-shifting. The average SNR over a 30 s period increases by 4.65 dB (Figure 18(b)), showing that LAVA can provide reasonable SNR gains without requiring fine-grained details about fluctuating channel conditions.

Multi-link scenarios. We next examine how well our power sensing mechanism handles multiple links. We set up an iw15300 AP and 5 different iw15300 Wi-Fi client pairs, allow each AP-client link to operate at different time instances, and measure per-link throughput with and without LAVA (Figure 19(a)). LAVA increases the average throughput for either client by 113% and 386%, and by up to 293% and 814%. I.e., all links benefit from LAVA this way.

LAVA vs multi-AP scenarios. We also compare utilizing LAVA to having access to additional APs, e.g., in enterprise settings or with Google Home, where a single client can connect to either of two APs. A suitable LAVA route is enabled for each iw15300 client-AP link, and the client transmits at 0 dBm. Figure 19(b) shows that LAVA increases the average throughput for either client-AP link by 95% and 547%, respectively. This shows that LAVA is orthogonal and complementary to multi-AP setups. Moreover, LAVA can be better than switching to another AP, even though the AP density in Figure 16 is higher than in most deployments, including the production Wi-Fi network in our office. In setup 4, switching from AP2 to AP1 increases the throughput from 1.4 Mbps to 11.3 Mbps, or $8\times$, whereas LAVA increases the throughput to 18.8 Mbps, or $13\times$.

Concurrent links. We next investigate whether two links can benefit simultaneously from LAVA. Given the temporal and spatial resolutions of our testbeds, we can only effectively evaluate two concurrent links that are at least two hops (around 4 m) apart. We

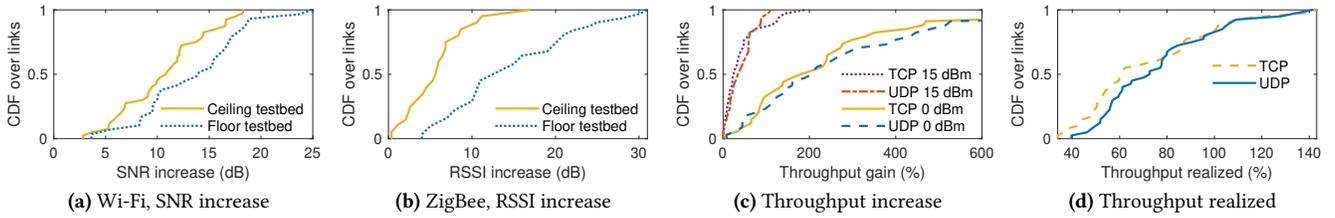


Figure 17: LAVA improves both Wi-Fi and ZigBee links in terms of SNR or throughput. Further, a significant portion of the TCP/UDP throughput for 15 dBm can be realized at 0 dBm when aided by LAVA (shown in (d)).

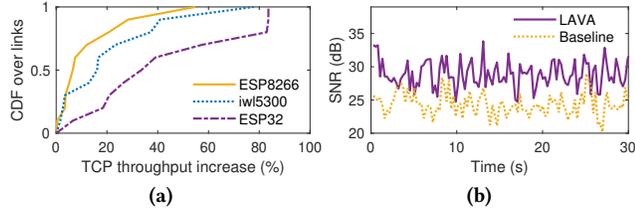


Figure 18: (a) Throughput increase for different devices: ESP8266, ESP32 and Intel iw15300. (b) Snapshot of the SNR when moving the receiver at 1 m/s.

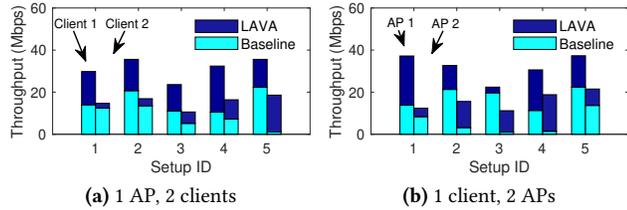


Figure 19: LAVA can improve the throughput for both multi-client and multi-AP scenarios.

set up 5 different co-channel WiFi-WiFi link pairs. LAVA is transparent to regular Wi-Fi MAC operations and the gains provided by LAVA in this setup will inevitably be affected by MAC-layer carrier sense. We use the ESP32 modules as the endpoints and not the iw15300 stations, because the former appears to back off less than the latter while providing more reliable throughput measurements than WARPLab (or XBee). We then measure *per-link* TCP throughput for four cases (Figures 20(a)(b)): (i) when both links are active; (ii) like (i) but with LAVA; (iii) when only one link is active; and (iv), like (iii) but with LAVA. Instead, LAVA improves the per-link throughput during concurrent transmissions by 5× on average at 0 dBm (by 80% at 15 dBm). More importantly, by introducing spatial segregation, LAVA sometimes enables the independent links to communicate concurrently with the same performance as if they were operating in isolation.

While phase alignment, used explicitly or implicitly in most previous smart surface prototypes, is frequency *selective* and rarely optimizes for multiple frequencies simultaneously, signal amplification, used primarily in LAVA, is frequency *agnostic* over a wide range. Therefore, we next set up 5 Wi-Fi link pairs concurrently on *different* frequencies (Channels 1 and 11 respectively). Figures 20(c)(d) show that LAVA increases the per-link throughput at 0 and 15 dBm by 55% and 68% on average, respectively.

To summarize, LAVA can improve concurrent wideband links, co-channel or otherwise, just as for single-link scenarios, provided the routes for distinct links do not share elements.

Energy propagation and spatial reuse. Next, we assess the effects of a dense topology of low-gain amplifiers on energy distribution, using the floor testbed. A transmitter is placed in an empty area (roughly at coordinates (1,1) in Figure 21(a)), whose initial transmit power was set to achieve 23 dB receive SNR at an arbitrary location, (x_{rx}, y_{rx}) . We then reduce the transmit power by 7 dB and find a 4-hop relay setting, each hop providing 37 dB amplification, that yields the same SNR at (x_{rx}, y_{rx}) again. We measure the signal strength (Figures 21(a)(b)) for both scenarios at evenly-spaced locations around the Wi-Fi link. Figure 21(b) shows smaller interference ranges around the link than Figure 21(a), suggesting this multi-hop amplify-and-forward mechanism can improve spatial reuse, when combined with reduced endpoint transmit power. With this, we demonstrate LAVA can support concurrent WiFi-ZigBee transmissions on the *same* frequency, thus providing a mechanism to mitigate cross-technology interference. Two ZigBee motes (tuned to 2.480 GHz) and two Wi-Fi nodes (on 2.484 GHz) are placed on the floor, around the floor testbed, such that neither receiver can directly “hear” from either transmitter. With manually set routes, LAVA can provide working links for both concurrently, whose SNRs are shown in Figure 21(c).

6 RELATED WORK

Smart surfaces. Antenna array based proposals [8, 17, 37, 38, 67, 75] advocating programming the radio environments mainly consider passive phase alignment between multipath propagation, whereas metasurface based “smart surface” ideas [6, 10, 14, 16, 20, 25, 28, 31, 36, 39–41, 43, 54, 66, 70, 74] further consider polarization rotation and absorption (attenuation). The end-to-end prototypes to date, LAIA [38], RFocus [4], ScatterMIMO [17], and LLAMA [14], have sampled the design space for certain use cases. Compared to endpoint-centric schemes, LAIA, RFocus, and ScatterMIMO can leverage more antennas over a large space and hence more degrees of freedom [34]. They are ideally suited to avoiding a single dead spot in space or frequency, or reducing antenna correlation at either endpoint for individual links. ScatterMIMO builds on LAIA and RFocus to reduce the search space of the surface configuration and aims to improve MIMO spatial multiplexing efficiency. LLAMA rotates the polarization of signals to minimize antenna polarization mismatch between endpoints. These prototypes are not yet optimized for co-channel concurrent links or 3D coverage, and require explicit per-frame feedback from the endpoint receiver to the surface control plane. LAIA and RFocus require pervasive deployment to increase beamforming gain, while ScatterMIMO and LLAMA are yet to be extended to large scale operations. Only LAIA considers concurrent links, but on different frequencies.

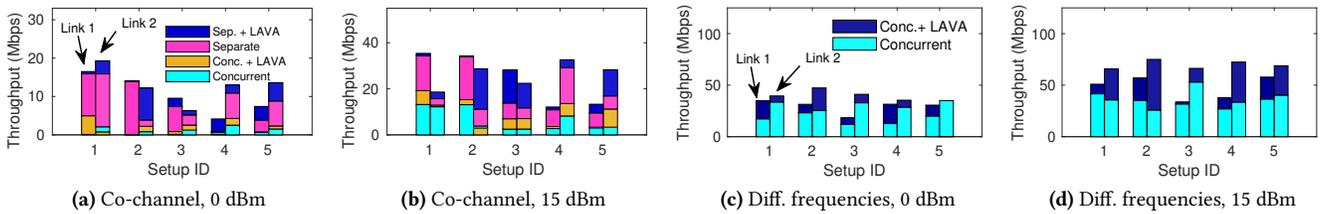


Figure 20: TCP throughput for concurrent links on the same or different frequencies at 0 or 15 dBm transmit power.

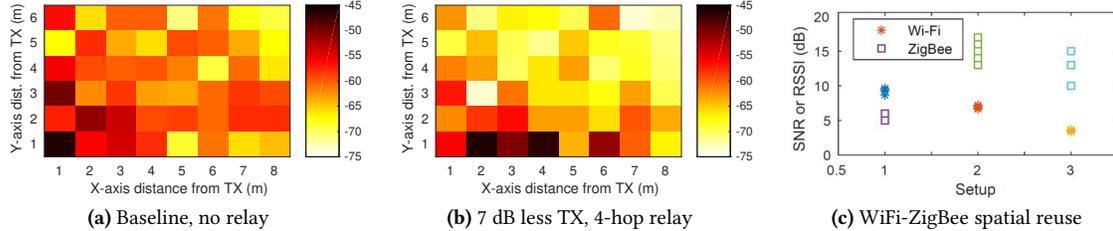


Figure 21: (a), (b) LAVA improves spatial reuse by reducing the interference range; (c) WiFi-ZigBee co-channel concurrent links.

Instead, LAVA injects power into the environment to combat path loss, thus more suited to small single-antenna or battery-powered devices instead of MIMO optimizations. With fine-grained amplification, LAVA handles wideband and concurrent links more easily and can operate without endpoint feedback. Our ceiling deployment also contrasts with and complements previous approaches (§4) while accounting for 3D coverage. Unlike LLAMA, LAVA does not counter signal polarization, but can mitigate power loss from endpoint beam pattern misalignment. LAVA routing signals within the array may resemble SurfaceMIMO [12]. However, SurfaceMIMO is not a reconfigurable surface, while LAVA is not designed to generate new spatial paths for MIMO.

Amplify-and-forward relays. MoVR [2] for 60 GHz and Fast-Forward [7] and DelayForward [29] for Wi-Fi OFDM dynamically adapt the amplification gain. LAVA builds on these with fixed-gain, multi-hop relaying, minimalist hardware and control logic, for a standard-agnostic, cost-effective solution.

Multi-AP deployments and mesh networking. To improve wireless coverage, multi-AP deployment [45, 46, 53, 69] is the standard approach for enterprises, while Wi-Fi repeaters [19, 63, 64] or Google Home employ mesh networking in homes. However, the typical enterprise AP density is one per 100 m², still insufficient to handle complex environments at a sub-room level (Figure 2). Further, multiple APs or mesh nodes require more frequencies or cause interference, and are technology-dependent. A client within range of multiple APs or Google Home nodes can frequently oscillate between associating with different APs. Alternatively, *distributed antenna systems* (DAS) use coaxial or fiber cables to introduce a separation between the access point and its antennas. This can improve coverage and network capacity [71, 77] via more efficient MIMO spatial multiplexing. For these approaches, upgrades to new protocol versions and hardware are cumbersome.

LAVA adds power at the signal level without halving capacity per hop, while using less components than full-fledged APs or mesh nodes. LAVA is transparent to the standard version or hardware implementation at endpoints, and is orthogonal and complementary to DAS. Fine-grained multi-hopping allows LAVA to exert more

control over the end-to-end signal trajectory and offer more disjoint paths for concurrent links.

Static signal shapers. Static signal shapers [13, 72] or reflectors [1, 2, 26, 76] are other takes on improving signal propagation behavior, but lack configurability. Directional antennas have been leveraged on communication endpoints to *statically and strategically restrict* signal propagation to avoid interference between networks [42]) or limit network coverage [55]. LAVA leverages directional antennas to route and confine the dominant signal trajectory.

7 CONCLUSION

In this paper, we highlight the need to re-examine non-uniform wireless signal reception in 3D given the rise of diverse, small IoT devices and more challenging and unplanned deployment locations. We propose a standard-agnostic infrastructure solution, a *Large Array of Vanilla Amplifiers* (LAVA), to programmatically *prescribe* the dominant signal propagation paths in the environment with multi-hop amplify-and-forward. This is combined with a passive power monitoring approach to detect active endpoint transmissions.

We design and implement a LAVA prototype in two testbed setups, above the secondary ceiling in real offices and on an empty lab floor. Comparison between the ceiling and floor testbeds further highlights the effect of 3D signal propagation. Extensive evaluation shows that our LAVA prototype can strengthen a range of link setups, including single-antenna Wi-Fi and ZigBee links and co-channel concurrent links. The power detection mechanism opens a new direction for inferring the network conditions. LAVA is therefore a promising wideband framework towards building an infrastructure decoupled from any existing wireless networks.

8 ACKNOWLEDGMENTS

We thank Josh Chavez, Bo Hu, Zhuqi Li, Yao Peng, Kevin Ryan, and Yaxiong Xie for helping with early iterations of LAVA element testing and testbed assembly. We also thank the anonymous reviewers and our shepherd, Rajalakshmi Nandakumar, for insightful comments. This work is partially supported by the National Science Foundation under Grant Nos. CNS-1763212 and CNS-1763309.

REFERENCES

- [1] O. Abari, D. Bharadia, A. Duffield, D. Katabi. Cutting the Cord in Virtual Reality. *Proceedings of ACM HotNets*, 2016.
- [2] O. Abari, D. Bharadia, A. Duffield, D. Katabi. Enabling High-quality Untethered Virtual Reality. *Proceedings of USENIX NSDI*, 2017.
- [3] F. Adib, S. Kumar, O. Aryan, S. Gollakota, D. Katabi. Interference Alignment by Motion. *Proceedings of ACM MobiCom*, 2013.
- [4] V. Arun, H. Balakrishnan. RFocus: Practical Beamforming for Small Devices. *Proceedings of USENIX NSDI*, 2020.
- [5] Taoglas 2.4GHz Bandpass Filter. <https://www.digikey.com/product-detail/en/taoglas-limited/BPF.24.01/931-1467-ND/6362804>.
- [6] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M. Alouini, R. Zhang. Wireless Communications Through Reconfigurable Intelligent Surfaces. *IEEE Access*, 7, 2019.
- [7] D. Bharadia, S. Katti. FastForward: Fast and Constructive Full Duplex Relays. *Proceedings of ACM SIGCOMM*, 2014.
- [8] E. Björnson, L. Sanguinetti, H. Wymeersch, J. Hoydis, T. L. Marzetta. Massive MIMO is a reality—What is next?: Five promising research directions for antenna arrays. *Digital Signal Processing*, 94, 2019.
- [9] S. M. Bowers, A. Safaripour, A. Hajimiri. Dynamic polarization control. *IEEE Journal of Solid-State Circuits*, 50(5), 2015.
- [10] X. Cao, B. Yang, H. Zhang, C. Huang, C. Yuen, Z. Han. Reconfigurable intelligent surface-assisted MAC for wireless networks: Protocol design, analysis, and optimization. *IEEE Internet of Things Journal*, 2021.
- [11] C. J. Carver, Z. Tian, H. Zhang, K. M. Odame, A. Q. Li, X. Zhou. AmphiLight: Direct air-water communication with laser light. *Proceedings of USENIX NSDI*, 2020.
- [12] J. Chan, A. Wang, V. Iyer, S. Gollakota. Surface MIMO: Using Conductive Surfaces For MIMO Between Small Devices. *Proceedings of ACM MobiCom*, 2018.
- [13] J. Chan, C. Zheng, X. Zhou. 3D Printing Your Wireless Coverage. *Proceedings of ACM HotWireless Workshop*, 2015.
- [14] L. Chen, W. Hu, K. Jamieson, X. Chen, D. Fang, J. Gummesson. Pushing the Physical Limits of IoT Devices with Programmable Metasurfaces. *Proceedings of USENIX NSDI*, 2021.
- [15] K. W. Cho, M. H. Mazaheri, J. Gummesson, O. Abari, K. Jamieson. mmWall: A Reconfigurable Metamaterial Surface for mmWave Networks. *Proceedings of ACM HotMobile Workshop*, 2021.
- [16] B. Di, H. Zhang, L. Song, Y. Li, Z. Han, H. V. Poor. Hybrid beamforming for reconfigurable intelligent surface based multi-user communications: Achievable rates with limited discrete phase shifts. *IEEE Journal on Selected Areas in Communications*, 38(8), 1809–1822, 2020.
- [17] M. Dunna, C. Zhang, D. Sievenpiper, D. Bharadia. ScatterMIMO: Enabling virtual MIMO with smart surfaces. *Proceedings of ACM MobiCom*, 2020.
- [18] SunFounder Ethernet Shield W5100. <https://www.sunfounder.com/ethernet-shield-w5100-for-arduino.html>.
- [19] NETGEAR N300 Wi-Fi Range Extender. <https://www.netgear.com/home/products/networking/wifi-range-extenders/WN3000RP.aspx>.
- [20] R. Fara, P. Ratajczak, D.-T. P. Huy, A. Ourir, M. Di Renzo, J. De Rosny. A Prototype of Reconfigurable Intelligent Surface with Continuous Control of the Reflection Phase. *arXiv preprint arXiv:2105.11862*, 2021.
- [21] Blue Sea Systems Blade Fuse Block. https://www.blueseas.com/products/5029/ST_Blade_Fuse_Block_-_12_Circuits_with_Cover.
- [22] The Future of IoT and the Digital Workplace. <https://www.cmswire.com/digital-workplace/the-future-of-iot-and-the-digital-workplace/>.
- [23] R. Ghaffarivardavagh, S. S. Afzal, O. Rodriguez, F. Adib. Underwater Backscatter Localization: Toward a Battery-Free Underwater GPS. *Proceedings of ACM HotNets*, 2020.
- [24] S. Gollakota, F. Adib, D. Katabi, S. Seshan. Clearing the RF smog: making 802.11n robust to cross-technology interference. *Proceedings of the ACM SIGCOMM*, 2011.
- [25] H. Guo, Y.-C. Liang, J. Chen, E. G. Larsson. Weighted sum-rate maximization for reconfigurable intelligent surface aided wireless networks. *IEEE Transactions on Wireless Communications*, 19(5), 3064–3076, 2020.
- [26] S. Han, K. Shin. Enhancing Wireless Performance Using Reflectors. *Proceedings of IEEE INFOCOM*, 2017.
- [27] Y. Hauri, D. Bhattacherjee, M. Grossmann, A. Singla. "Internet from Space" without Inter-Satellite Links. *Proceedings of ACM HotNets*, 2020.
- [28] P. del Hougne, M. Fink, G. Lerosey. Optimally diverse communication channels in disordered environments with tuned randomness. *Nature Electronics*, 2(1), 36–41, 2019.
- [29] K.-C. Hsu, K. C.-J. Lin, H.-Y. Wei. Full-duplex Delay-and-forward Relaying. *Proceedings of ACM MobiHoc*, 2016.
- [30] C. Huang, S. Hu, G. C. Alexandropoulos, A. Zappone, C. Yuen, R. Zhang, M. Di Renzo, M. Debbah. Holographic MIMO surfaces for 6G wireless networks: Opportunities, challenges, and trends. *IEEE Wireless Communications*, 27(5), 118–125, 2020.
- [31] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah, C. Yuen. Reconfigurable intelligent surfaces for energy efficiency in wireless communication. *IEEE Transactions on Wireless Communications*, 18(8), 4157–4170, 2019.
- [32] Healthy living with IoT. <https://www.iotforall.com/healthy-living-with-iot-ces>.
- [33] Five Emerging IoT Trends for 2021. <https://www.forbes.com/sites/theyec/2021/12/28/five-emerging-internet-of-things-trends-for-2021/?sh=6a96fb40205d>.
- [34] S. A. Jafar, M. J. Fakhereddin. Degrees of freedom for the MIMO interference channel. *IEEE Transactions on Information Theory*, 53(7), 2007.
- [35] J. Jang, F. Adib. Underwater backscatter networking. *Proceedings of ACM SIGCOMM*, 2019.
- [36] G. R. Keiser, N. Karl, S. R. U. Haque, I. Brener, D. M. Mittleman, R. D. Averitt. Structurally Tunable Nonlinear Terahertz Metamaterials using Broadside Coupled Split Ring Resonators. *arXiv preprint arXiv:2104.05757*, 2021.
- [37] Z. Li, Y. Xie, L. Shangquan, R. I. Zelaya, J. Gummesson, W. Hu, K. Jamieson. Programmable Radio Environments with Large Arrays of Inexpensive Antennas. *GetMobile: Mobile Computing and Communications*, 23(3), 23–27, 2019.
- [38] Z. Li, Y. Xie, L. Shangquan, R. I. Zelaya, J. Gummesson, W. Hu, K. Jamieson. Towards Programming the Radio Environment with Large Arrays of Inexpensive Antennas. *Proceedings of USENIX NSDI*, 2019.
- [39] C. Liaskos, S. Nie, A. Tsioliaridou, A. Pitsillides, S. Ioannidis, I. Akyildiz. A New Wireless Communication Paradigm through Software-Controlled Metasurfaces. *IEEE Communications Magazine*, 56(9), 2018.
- [40] C. Liaskos, S. Nie, A. Tsioliaridou, A. Pitsillides, S. Ioannidis, I. Akyildiz. Realizing wireless communication through software-defined hypersurface environments. *Proceedings of IEEE WoWMoM*, 2018.
- [41] X. Liu, Y. Liu, Y. Chen, H. V. Poor. RIS enhanced massive non-orthogonal multiple access networks: Deployment and passive beamforming design. *IEEE Journal on Selected Areas in Communications*, 39(4), 1057–1071, 2020.
- [42] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, P. Steenkiste. DIRC: Increasing Indoor Wireless Capacity Using Directional Antennas. *Proceedings of ACM SIGCOMM*, 2009.
- [43] Y. Liu, X. Liu, X. Mu, T. Hou, J. Xu, M. Di Renzo, N. Al-Dhahir. Reconfigurable intelligent surfaces: Principles and opportunities. *IEEE Communications Surveys & Tutorials*, 2021.
- [44] Logarithmic Detector AD8318. <https://www.analog.com/media/en/technical-documentation/data-sheets/AD8318.pdf>.
- [45] A. Miu, H. Balakrishnan, C. E. Koksal. Improving Loss Resilience with Multi-radio Diversity in Wireless Networks. *Proceedings of ACM MobiCom*, 2005.
- [46] R. Murty, J. Padhye, R. Chandra, A. Wolman, B. Zill. Designing High Performance Enterprise Wi-Fi Networks. *Proceedings of USENIX NSDI*, 2008.
- [47] S. Narayana, R. V. Prasad, T. Prabhakar. SOS: Isolated health monitoring system to save our satellites. *Proceedings of ACM MobiSys*, 2021.
- [48] S. Narayana, R. Venkatesha Prasad, V. S. Rao, L. Mottola, T. Venkata Prabhakar. A Hummingbird in Space: An energy-efficient GPS receiver for small satellites. *GetMobile: Mobile Computing and Communications*, 25(1), 2021.
- [49] Scalable and Ultra-Low Power Ocean IoT. <https://www.media.mit.edu/projects/ultra-wideband-underwater-backscatter-via-piezoelectric-metamaterials/overview/>.
- [50] M. Okatan, J. Mantese, S. Alpay. Polarization coupling in ferroelectric multilayers. *Physical Review B*, 79(17), 2009.
- [51] Alpha Network Antenna APA-M25. https://www.alfa.com.tw/products_detail/234.htm.
- [52] Macom Digital Phase Shifter MAPS-010164. <https://www.macom.com/products/product-detail/MAPS-010164>.
- [53] H. Rahul, H. Hassanieh, D. Katabi. SourceSync: A Distributed Wireless Architecture for Exploiting Sender Diversity. *Proceedings of ACM SIGCOMM*, 2010.
- [54] M. D. Renzo, M. Debbah, D.-T. Phan-Huy, A. Zappone, M.-S. Alouini, C. Yuen, V. Sciancalepore, G. C. Alexandropoulos, J. Hoydis, H. Gacanin, J. de Rosny, A. Bounceur, G. Lerosey, M. Fink. Smart radio environments empowered by reconfigurable AI meta-surfaces: An idea whose time has come. *EURASIP Journal on Wireless Communications and Networking* volume, 2019.
- [55] A. Sheth, S. Seshan, D. Wetherall. Geofencing 802.11 coverage areas to physical boundaries. *Proceedings of Pervasive*, 2009.
- [56] Texas Instruments 8-Bit Shift Register. <http://www.ti.com/lit/ds/symlink/sn74hc595.pdf>.
- [57] V. Singh, A. Prabhakara, D. Zhang, O. Yağan, S. Kumar. A community-driven approach to democratize access to satellite ground stations. *Proceedings of ACM MobiCom*, 2021.
- [58] Smart hospital. <https://www.techrepublic.com/article/new-smart-hospital-platform-could-be-the-digital-transformation-tool-healthcare-needs/>.
- [59] Skyworks SP4T Switch SKY13414-485LF. [https://www.skyworksinc.com/-/media/SkyWorks/Documents/Products/701-800/SKY13414_485LF_201689\].pdf](https://www.skyworksinc.com/-/media/SkyWorks/Documents/Products/701-800/SKY13414_485LF_201689].pdf).
- [60] Analog Devices SP8T Switch HMC321ALP4E. <https://www.analog.com/media/en/technical-documentation/data-sheets/hmc321a.pdf>.
- [61] Ho IoT in Space is Taking Off. <https://www.iotechexpo.com/2019/12/iot/iot-in-space/>.
- [62] F. Tonolini, F. Adib. Networking across boundaries: Enabling wireless communication through the water-air interface. *Proceedings of ACM SIGCOMM*, 2018.
- [63] Linksys - AC750 Boost Range Extender. <http://www.linksys.com/us/p/P-RE6300/>.

- [64] TP-Link RE450 1750 Wi-Fi Range Extender. https://www.tp-link.com/us/products/details/cat-5508_RE450.html.
- [65] Toshiba DMOS Transistor Array. <https://www.digikey.com/product-detail/en/toshiba-semiconductor-and-storage/TBD62781APG/TBD62781APG-ND/8570316>.
- [66] S. Venkatesh, X. Lu, H. Saeidi, K. Sengupta. A high-speed programmable and scalable terahertz holographic metasurface based on tiled CMOS chips. *Nature Electronics*, 3(12), 785–793, 2020.
- [67] A. Welkie, L. Shangguan, J. Gummesson, W. Hu, K. Jamieson. Programmable Radio Environments for Smart Spaces. *Proceedings of ACM HotNets*, 2017.
- [68] What is the Internet of Things. <https://www.scientificamerican.com/article/what-is-the-internet-of-things/>.
- [69] G. Woo, P. Kheradpour, D. Shen, D. Katabi. Beyond the bits: Cooperative packet recovery using physical layer information. *Proceedings of ACM MobiCom*, 2007.
- [70] Q. Wu, R. Zhang. Towards Smart and Reconfigurable Environment: Intelligent Reflecting Surface Aided Wireless Network. *IEEE Communications Magazine*, 58(1), 2020.
- [71] J. Xiong, K. Sundaresan, K. Jamieson, M. A. Khojastepour, S. Rangarajan. MIDAS: Empowering 802.11ac Networks with Multiple-Input Distributed Antenna Systems. *Proceedings of ACM CoNEXT*, 2014.
- [72] X. Xiong, J. Chan, E. Yu, N. Kumari, A. A. Sani, C. Zheng, X. Zhou. Customizing indoor wireless coverage via 3D-fabricated reflectors. *Proceedings of ACM BuildSys Workshop*, 2017.
- [73] Q. Yang, X. Li, H. Yao, J. Fang, K. Tan, W. Hu, J. Zhang, Y. Zhang. BigStation: Enabling Scalable Real-time Signal Processing in Large MU-MIMO systems. *Proceedings of ACM SIGCOMM*, 2013.
- [74] I. Yoo, M. F. Imani, T. Slesman, H. D. Pfister, D. R. Smith. Enhancing capacity of spatial multiplexing systems using reconfigurable cavity-backed metasurface antennas in clustered MIMO channels. *IEEE Transactions on Communications*, 67(2), 1070–1084, 2018.
- [75] Q. Zhang, Y.-C. Liang, H. V. Poor. Large Intelligent Surface/Antennas (LISA) Assisted Symbiotic Radio for IoT Communications, 2020.
- [76] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Zhao, H. Zheng. Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers. *Proceedings of ACM SIGCOMM*, 2012.
- [77] H. Zhu, J. Wang. Radio Resource Allocation in Multiuser Distributed Antenna Systems. *IEEE JSAC*, 2013.

APPENDIX

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

A LAVA DOCUMENTATION

The source code and design schematics of our LAVA ceiling testbed are now available at: https://github.com/riz333/lava_sigcomm 21. This directory contains four subdirectories:

- **control/**: contains the LAVA control plane algorithm.
- **misc/**: contains miscellaneous auxiliary functions.
- **controlCodeMCU/**: contains the Arduino microcontroller control code.
- **circuit_diagrams/**: contains detailed circuit diagrams of the LAVA multidirectional prototype.

control/ contains a topology file, *finaltop.txt*, which describes the testbed topology, a Makefile that generates executables and three code files:

- *conclinks.c*: implements concurrent links at master.
- *newalg.c*: implements single links at master.
- *newfwd.c*: bridges master to Arduinos at subcontroller.

misc/ contains a Makefile that generates executables and three code files:

- *getv.c*: queries voltage at an Arduino.
- *setnode.c*: sends commands to an Arduino.
- *setpath.c*: single-link route selection.

controlCodeMCU/ contains the code implemented in the Arduino microcontrollers to control the elements.

circuit_diagrams/ contains four circuit schematics (pertaining to the LAVA multidirectional element design):

- *amplifier_selection.pdf*: describes how the LAVA element activates an amplifier.
- *element_control_and_power_schematic.pdf*: shows all connections between the element’s hardware and the microcontroller unit.
- *element_rf_schematic.pdf*: shows the element’s RF connections.
- *pcb_prototype.pdf*: schematic of the element’s custom PCB board (RF switching and phase shifting).

B IOT DEVICES AND CAPABILITIES

Table 1: Example Wireless IoT devices

Product	Battery Powered?	Wireless IC	Standard	Tx Power (dBm)
August lock	Yes	CC2541	Bluetooth 4.0	-23 to 0
Microsoft Band watch	Yes	AR3002	Bluetooth 4.0	-20 to +20
Sketchers GOwalk watch	Yes	nRF8001	Bluetooth 4.0	-18 to +4
Garmin fenix 5 Plus watch	Yes	CYW20719	Bluetooth 5.1	0 to +5
		ATWILC1000	802.11 b/g/n	+12.2 to +18.9
		nRF52832	ANT	-20 to +4
iControl Piper camera	Yes	ZW0301	Z-Wave	-20 to 0
Blink XT camera	Yes	CC3100	802.11 b/g/n	+13.0 to +18.3
FunAce Robot camera	No	MT7601U	802.11 b/g/n	+15 to +19
GE dimmer	No	CSR1010	Bluetooth 4.1	Up to +9
Kasa bulb	No	LBM100	802.11 b/g/n	Up to +18
LIFX Mini bulb	No	ESP32	Bluetooth 4.2	-12 to +9
			802.11 b/g/n	+12 to +20.5
WiZ bulb	No	ESP-WROOM-02	802.11 b/g/n	+13 to +20.5
Teckin SP23 plug	No	ESP8266EX	802.11 b/g/n	+15.5 to +21.5
BlitzWolf BW-SHP4 plug	No	ESP8285	802.11 b/g/n	+13 to +20
TP-Link HS100 plug	No	AR9331	802.11 b/g/n	+16 to +20
Beam Alert range extender	No	QCA9531	802.11 a/b/g/n	+20 to +26 (2 chains)