Programmable Radio Environments for Smart Spaces

Allen Welkie
Princeton University

Longfei Shangguan
Princeton University

Jeremy Gummeson
University of Massachusetts
Amherst

Wenjun Hu
Yale University

Kyle Jamieson
Princeton University

ABSTRACT

Smart spaces, such as smart homes and smart offices, are common Internet of Things (IoT) scenarios for building automation with networked sensors. In this paper, we suggest a different notion of smart spaces, where the radio environment is programmable to achieve desirable link quality within the space. We envision deploying low-cost devices embedded in the walls of a building to passively reflect or actively transmit radio signals. This is a significant departure from typical approaches to optimizing endpoint radios and individual links to improve performance. In contrast to previous work combating or leveraging per-link multipath fading, we actively *reconfigure* the multipath propagation. We sketch design and implementation directions for such a programmable radio environment, highlighting the computational and operational challenges our architecture faces. Preliminary experiments demonstrate the efficacy of using passive elements to change the wireless channel, shifting frequency "nulls" by nine Wi-Fi subcarriers, changing the 2×2 MIMO channel condition number by 1.5 dB, and attenuating or enhancing signal strength by up to 26 dB.

1 INTRODUCTION

The vision of the Internet of Things (IoT) is commonly described in terms of "smart" scenarios. For example, smart homes and offices are already commercially viable solutions for building automation, leveraging networked sensors that interact with the environment. As more and more wirelessly connected IoT devices are deployed, however, the wireless channels are increasingly strained. In this context, unplanned wireless networks such as Wi-Fi, Zigbee, and LoRa have proved their utility, but are also known to degrade in performance when users move into a "dead zone," when they must scale to more users, or when many operate in close proximity.

Figure 1: A passive type of PRESS-instrumented room's walls might contain an array of low-cost antennas which either reflect or attenuate wireless signals incident on the wall, configuring the environment to improve wireless networks operating nearby.

These important problems lie at the heart of the overall effort to scale up the overall wireless network bandwidth to satisfy an ever-growing user demand. Specifically:

(1) How best to eliminate dead zones in the presence of the vagaries of multipath propagation?
(2) How best to leverage *spatial multiplexing* in the multiuser MIMO channel, to simultaneously move packets to or from multiple clients?
(3) How best can a network mitigate interference from other, nearby networks?

At the physical layer, the questions revolve around how to combat multipath fading or suppress interference across antennas, for the same receiver or across receivers.

Most—if not all—of the solutions thus far formulated involve enhancing the *endpoint* radios involved in communication. For example, techniques to align and/or null interference between and among different wireless networks [2, 8, 16, 18, 20, 22] hinge on multi-antenna transmitters' ability to adjust their transmissions at the transmission endpoint. Indeed, the advent of "massive MIMO" techniques in next-generation cellular networks applies this conventional technique in the extreme, introducing many more (switched) antennas than both radio chains and users at an access point (AP), so that the AP may search for a set of antennas that forms a "well-conditioned" MIMO channel to those users, enhancing the overall network throughput [10, 27, 28, 30, 37].

In millimeter-wave networks, recent work has explored using knowledge of reflector locations to inform AP placement [34], placing static reflectors to *replace* direct point-to-point links [38], and using programmable phased-array reflectors [1] to create alternate links that circumvent obstacles. For

frequencies below 10 GHz, on the other hand, the number of reflectors, diffractors, and absorbers in the environment potentially dwarfs the number of antennas at the endpoints of wireless communication. Hence there is reason to speculate that there may be even more degrees of freedom in these networks to change the *environment* itself rather than the wireless communication endpoints (*i.e.*, users and APs).

This paper asks whether it might be possible to build a smarter *environment* by electronically modulating the environment itself. Our discussion in this paper points the way to a multitude of techniques that will eventually actively *reconfigure* indoor multipath propagation: a *Programmable Radio Environment for Smart Spaces* (PRESS). One realization of a potential PRESS design envisions a matrix of low-cost antenna elements connected to passive loads and embedded in the walls of a building, as shown in Figure 1. In this way, one can think of a PRESS-instrumented building as a massive, programmable antenna. There are at least three ways such an environment can facilitate better communication between wireless endpoints:

Enhancing individual wireless links. The destructive effect of multiple spatial paths arriving at a receive antenna out of phase results in a frequency "null", as shown, for example, in the dashed grey curve representing wireless channel H_{11} in Figure 2. But if the PRESS elements in the walls of a building could shift the phase of the wireless signals the walls reflect, we would be able to eliminate the null at the frequency of interest. The OFDM modulation and channel coding operating on each link would then see a "flatter" channel, and could offer a greater bit rate, and hence throughput, to higher layers. Spatial "dead spots" that Wi-Fi networks anecdotally suffer from are often the result of this problem.

Improving Large MIMO performance. While Massive MIMO systems sometimes enjoy the luxury of many more antennas than users, this is not guaranteed at peak demand and reduces the systems to *large MIMO* instead. MIMO techniques' gains depend on the availability of spatially-uncorrelated paths, which can be quantified with the *condition number* of the channel matrix. In large MIMO systems the channel frequently becomes *poorly conditioned* [15, 21, 37], where similar MIMO channels between different users degrade the throughput of conventional MIMO algorithms. But by manipulating the wireless channel itself, PRESS may be able to improve the large MIMO conditioning, restoring performance without additional AP processing complexity.

Network "harmonization" and spatial partitioning. Dynamic frequency-division schemes split the wireless medium in frequency between clients [7, 26, 29] or wireless networks [6, 16, 22], with attendant network efficiency benefits. But they benefit most when the communication channels allocated to the respective parties (*i.e.* the left-hand side of the spectra in Figure 2 for AP and Client 1, and the right-hand side for AP and Client 2) are strongest, while interference channels (those between interfering senders and "bystanders" shown at the bottom of Figure 2) are weakest. A first possibility is
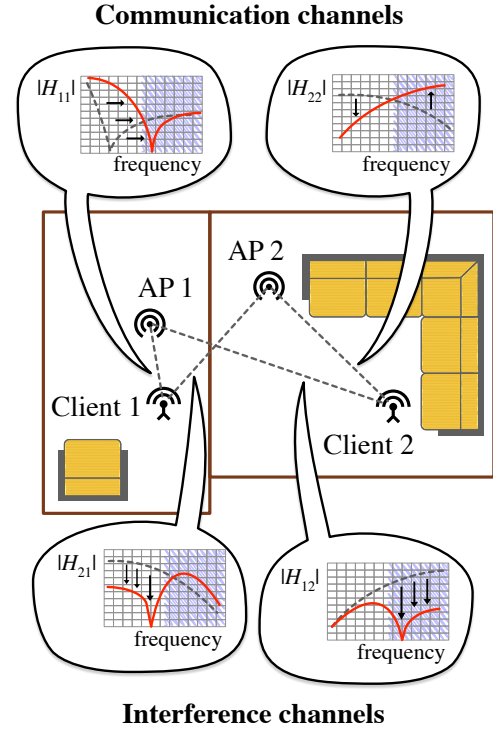
Communication channels

Figure 2: Partitioning and harmonizing two wireless networks with PRESS. Dashed grey lines in the plots show the communication (AP 1 to Client 1 and AP 2 to Client 2) channels and interference (AP 1 to Client 2 and AP 2 to Client 1) channels *before* PRESS, solid red lines represent the same channels *after* PRESS' effect.

therefore to *harmonize* the communication and interference channels between two nearby networks by moving the location of frequency nulls, attenuating channel strength on certain sub-bands, and enhancing it on others, as shown in Figure 2. Another instance of network harmonization is *interference alignment* [4, 13]: aligning the interference that two networks cause at a receiver in a third network, so that that receiver may remove the interference from both interfering networks in a single nulling step [8, 18]. A third possibility is simply to reduce interference between different pairs of wireless conversations, spatially partitioning the space. These techniques stand in contrast to previous proposals of *physically* separating transmitted signals via sectorization in a cellular network, or 3D-printing Wi-Fi AP antenna reflector hardware [5], because they work flexibly, in the environment itself, and do not constrain the transmitters.

2 DESIGN CHALLENGES

A series of inter-dependent challenges present in realizing the above vision makes configuring the radio environment for the foregoing applications perhaps even more challenging than optimizing individual links from their endpoints.

The first challenge we encounter is how to choose radio

hardware that can induce power and phase changes significant with respect to the active links, balancing and trading off efficacy with associated cost. As the authors of the *Braidio* [12] design observe, *active* radio architectures consist of receive and transmit subsystems, which in turn are composed of power-hungry mixers and power amplifiers to quadrature-modulate/demodulate and amplify the communication signal, respectively. In contrast, *passive* backscatter designs [11, 14, 19] contain neither, simply reflecting energy using passive components such as capacitors, inductors, and switches. Considering efficacy, the *PhyCloak* [25] system has demonstrated that a full-duplex, active "obfuscator" radio operating in close proximity with a transmitting radio can alter the wireless channel amplitudes, delays, and Doppler shifts between the transmitter and a nearby receiver. But active radios (not to mention full-duplex active radios) are relatively expensive and power-hungry, and so are unlikely to scale to deployment as envisioned in Figure 1 across an entire building. On the other hand, passive PRESS elements have a cost advantage, so can scale to a relatively large and dense array of small antennas, but may not reflect enough power over enough range to make enough of a difference to the wireless channel. We anticipate that our eventual design will involve a mixture of both active and passive elements, with the latter significantly outnumbering the former.

Secondly, a PRESS deployment wholly- or partially-comprised of a large number of densely-deployed passive elements entails a non-trivial "inverse problem." A standard signal model [31] comprised of one or more paths, each of which is characterized by its angle of departure from the sender ϕ_l, propagation delay over the air τ_l, Doppler shift γ_i, and angle of arrival at the receiver θ_l (for path l), has been shown to accurately predict the wireless channel between a sender and receiver [32]. But PRESS demands the inverse direction of this calculation: given the existing wireless channel between sender and receiver, or sender and "bystander" radio experiencing interference from that transmission, we seek to compute the signal path parameters $\{\phi_m, \tau_m, \gamma_m, \theta_m, \ldots\}$ for an existing or additional path or paths such that the *superposition* of the existing, modified, and additional paths yields the desired wireless channel, for example to make one of the changes to the wireless channel.

The foregoing leads us to a third broad set of challenges: how to rapidly actuate the PRESS array to produce the desired wireless channel. This entails several tasks:

(1) PRESS must gather all the required wireless channel information between transmitters, interferers, "bystanders," and intended receivers.
(2) The system must quickly navigate through an enormous search space of channel parameters to identify the configuration that produces the desired effect.
(3) The system must then apply the chosen configuration to the PRESS array, possibly controlling each element in the vicinity individually to achieve the configuration intended. If the current communication patterns involve


Figure 3: PRESS element configuration—three SP4T switches that can independently switch between three different reflective path lengths (0, $\lambda/4$, and $\lambda/2$ additional path length) and one absorptive load.

multiple wireless links operating over different time or frequency slots, we would like the system to attempt to optimize them jointly and simultaneously, if possible.

In order for ongoing communication to reap the benefits of the PRESS array, the latter must perform the above all during the channel *coherence time*, the amount of time over which the wireless channel remains constant. Typical values of the channel coherence time at 2.4 GHz range from *ca.* 80 milliseconds while almost stationary (0.5 mph movement) down to *ca.* six milliseconds at running speed (6 mph) [31], placing upper bounds on the time for PRESS to act. But furthermore, depending on traffic patterns, PRESS will very likely reap additional performance benefits from switching strategies on packet-level timescales of one to two milliseconds, as the set of senders and receivers changes. Note here that a trade-off exists between agility and optimization: one might jointly optimize over a large set of likely communication links, obviating the need to change the PRESS array for each link's communication, but possibly complicating the optimization problem. On the other end of the design space, one might optimize solely over a single communication link, or two individual networks, but hard-forcing the above timing constraints. One can imagine hybrid tradeoffs and dynamic strategies that leverage these extreme positions as well as points on the spectrum in between. These timing considerations lead us to the last of our system design challenges.

Finally, we face a number of more practical operational challenges. These include how to deploy, power, and maintain the PRESS array itself. Communication between a centralized or semi-centralized PRESS controller and the array elements is a key design issue—we seek some combination of wired and wireless communication that is practical and at the same time achieves low-latency control over the PRESS array. We furthermore seek a control plane design that does not interfere with communication in the wireless data plane: further thoughts on possible solutions can be found in §4.

3 AN EXPLORATORY STUDY

The aim of this experimental study is to gauge both potential gains and challenges involved in programming the radio
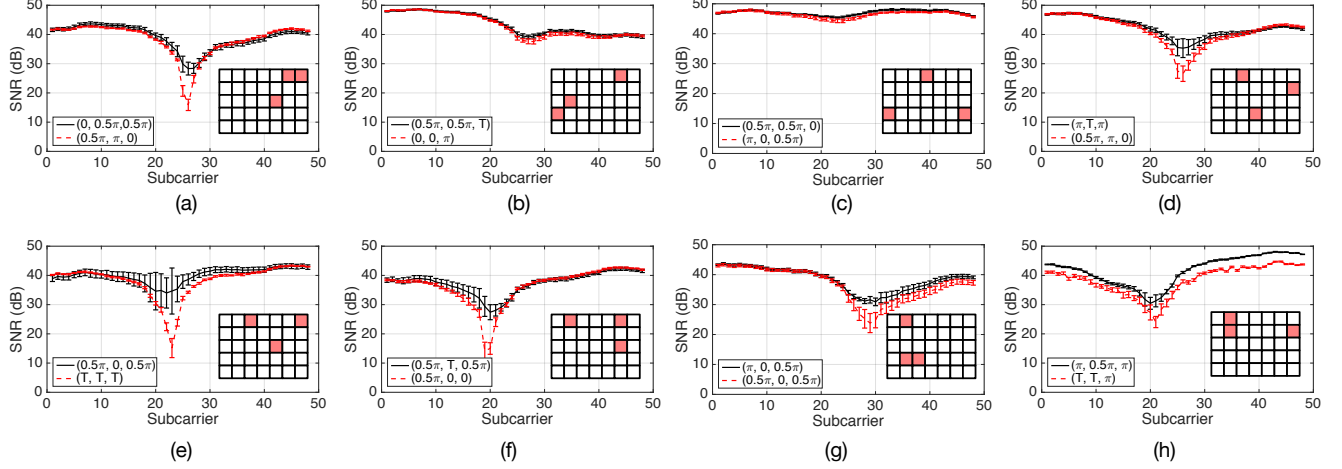
Figure 4: The measured per-subcarrier SNR for two PRESS configurations for each of eight randomly generated PR-ESS element locations (a) through (h). Locations are shown in the grid in each plot. In each figure, we plot the two configurations that give the largest single-subcarrier SNR difference. (The reflection coefficient configurations featured are different across figures. "T" means the element is terminated.)

environment. As a proof of concept, we experiment with a small passive PRESS array in a controlled indoor setting. We first run experiments involving transmitter and receiver in line of sight. In these scenarios, the effect of the PRESS element configurations on the per-subcarrier SNR is limited to less than 2 dB. This small change is expected, as the line-of-sight signal dominates over the reflection of much lower strength from the passive PRESS elements. This suggests that a *passive* PRESS array is best suited to improving non-line-of-sight links, whereas line-of-sight links require some *active* PRESS elements. Therefore, the following passive-element experiments are all done in a non-line-of-sight setup.

3.1 Prototype Implementation

Unless otherwise stated, our endpoints are WARP software-defined radios [33], transmitting Wi-Fi-like OFDM signals comprised of 64 subcarriers over 20 MHz on channel 11 of the ISM band (2.462 GHz). The transmit and receive endpoints use 2 dBi-gain omni-directional antennas [23] as could be expected in an operational setting. Our initial prototype PRESS element is comprised of a 14 dBi, 21° azimuthal beamwidth parabolic antenna [17] or omnidirectional antenna attached to three single-pole, four-throw (SP4T) RF switches [24] as shown in the configuration of Figure 3. We connect each endpoint of each SP4T switch to a RF waveguide of a different length comprised of cable adapters left open, thus reflecting radio energy. We control the RF switch through a micro-controller time synchronized with the WARP radios' transmissions.

3.2 Experimental Results

In the following experiments, the transmitter sends one frame comprised of multiple OFDM symbols and the receiver estimates the channel state information from the training sequences in the frame. Between successive frames, the RF

waveguide attached to each antenna is changed using the SP4T RF switch. Three of the four waveguides attached to each antenna are left open and the lengths differ by a quarter of a wavelength which changes the phase of the reflection from each antenna by $\pi/2$. The fourth waveguide is terminated with an absorptive load which effectively eliminates any reflection. Three antennas are used, which means there are 64 different PRESS antenna configurations. Because of the latency in our experimental setup, the channel for these 64 different combinations cannot be measured within channel coherence time (it takes about 5 seconds to measure all of the combinations). To compensate, we iterate through the 64 combinations 10 times and calculate statistics on the SNR for each PRESS antenna configuration.

To obtain a wireless channel with significant reflected components, we use an experimental setup that blocks the direct path between the transmitter and receiver. As can be expected, this channel demonstrates much more frequency selectivity than the line-of-sight experimental setup. We place the PRESS antennas in eight randomly generated locations in a grid 1–2 meters from both the transmitting and receiving antennas. Each antenna placement results in a different scattering environment due to the movement of our experiment equipment. At each of these antenna placement configurations we measure the wireless channel for the 64 different antenna reflection coefficient configurations.

3.2.1 Link Enhancement. We seek to measure the range of induced SNR changes one can expect for the different PRESS configurations. To measure this, we calculate which two configurations give the largest difference in subcarrier SNR across all subcarriers. The results are plotted in Figure 4. The biggest change in SNR occurs when one of the PRESS reflection coefficient configurations produces a "null" on one of the measured subcarriers. These nulls occur when

Figure 5: Complementary CDF of the change in null location (subcarrier index) between pairs of PRESS element configurations, among configurations that exhibit a null. Each curve contains data from a separate experimental repetition.

Figure 6: *Left:* Complementary CDF of the change in minimum SNR among subcarriers between pairs of PRESS element configurations. The probability axis is cropped to show more detail. *Right:* Complementary CDF of the minimum SNR among subcarriers for all 64 PRESS element configurations. Each trace is one of the 10 trials of the 64 configurations. The probability axis is cropped to show more detail.

Figure 7: Two configurations demonstrating control over frequency selectivity.

the multipath components—from PRESS and elsewhere—destructively interfere at the receiver. By changing the PRESS reflection coefficients, we can change how the multipath components superpose at the receiver and avoid or move the null. In these eight experiments, the largest change in the mean SNR on any given subcarrier is 18.6 dB, and the largest change in the SNR within one experimental repetition is 26 dB. Some nulls occur in configurations where the PRESS antennas are all terminated with an absorptive load and are not contributing reflection paths. For example, the data point with the lowest SNR in Figure 4(e) happens to be the configuration where all three antennas were terminated with an absorptive load. This demonstrates that PRESS configurations can remove nulls caused only by environmental reflections.

We can get an understanding of how the results vary across PRESS element configurations by investigating one of the PRESS element positions. In Figures 5 and 6 we plot statistics on the data from Figure 4(e). In particular, we would like to assess how PRESS can affect the frequency selectivity of a channel. To that end, we calculate statistics on how far the nulls in the channel move as a function of PRESS element configurations. In Figure 5, we plot the complementary CDF of the difference (measured in number of subcarriers) of the location of the most significant null in all of the 64^2 pairs of PRESS element configurations. The location of the most significant null is the subcarrier number corresponding to the minimum SNR value for a given configuration, and we only consider configurations that have a subcarrier SNR that is at least 5 dB less than the median subcarrier SNR. Of these pairs, most show either no change in null location or a change of only one subcarrier, but a few show changes of over three subcarriers (1 MHz).

We would also like to understand how the channel gain profile across subcarriers varies with the PRESS configuration. Do only a few of the configurations demonstrate these nulls? If we are experiencing a bad channel with a certain configuration, how likely will switching configurations improve the channel significantly? This informs the level of control precision over the PRESS system required to reap significant gains.

Figure 6 shows two complementary CDFs: the complementary CDF of the difference in dB of the minimum SNR across subcarriers for pairs of PRESS element configurations on the left, and the complementary CDF of those minimum SNRs for the 64 different configurations on the right. Around 38% of the configuration changes cause a 10 dB SNR change on at least one subcarrier, and less than 9% of the configurations show a worst subcarrier channel gain below 20 dB.

3.2.2 Network Harmonization. We would like to demonstrate the feasibility of PRESS regarding the spatial partitioning goal outlined in §1. In this experiment, we use two USRP N210 radios with only two PRESS elements, each of which is attached to four different reflective cable lengths and no absorptive load, to decrease the reflected phase granularity. Also, instead of randomly generated element placement, the elements and the surrounding environment were manipulated until a frequency-selective channel was found, to emulate a finer and more numerous PRESS array. Figure 7 shows that two of the PRESS element configurations exhibit clear and opposite frequency selectivity; each one favors its own half of the band.

3.2.3 Large MIMO Performance. The last remaining goal outlined in §1 is for PRESS to improve MIMO performance. To measure this effect, we replace the transceivers with a 2×2 MIMO transceiver pair in a non-line-of-sight

Figure 8: The distribution of MIMO channel condition number across subcarriers and experimental repetitions. Each curve on the CDF is a separate PRESS phase setting, with the phase settings demonstrating the best (lowest) and worst (highest) condition numbers appearing thicker and in color.

configuration (using a USRP X310 and two UBX 160 daughterboards) and measure the 2×2 channel matrix for each of the 64 PRESS configurations. Omnidirectional PRESS elements are deployed co-linear to the transmit antenna pair with λ spacing between the PRESS antenna elements. In MIMO systems, the condition number of the channel matrix is critically important to the channel capacity. In Figure 8, we plot a CDF of the channel matrix condition number across subcarriers for each PRESS configuration. Each CDF was computed from the mean of 50 successive channel measurements. This data shows that particular PRESS configurations (exemplars highlighted in red and blue) have a substantial impact on the MIMO matrix condition number and hence the maximum obtainable channel capacity when the number of client antennas approaches the number of AP antenna elements. Furthermore, we anticipate the impact of the PRESS elements to increase as the MIMO channel dimension increases past 2×2, as previously shown [21, 37].

4 DESIGN SPACE EXPLORATION

In light of the results from the foregoing exploratory study, we now identify and discuss some promising regions in the rich programmable radio environment design space. Our discussion is in two parts: first, we discuss design possibilities for the substrate itself (§4.1); then we consider the important issue of how to *control* the PRESS substrate (§4.2).

4.1 Substrate Choices

Element antenna design. Array elements may exhibit varying levels of directionality. Directionality can be achieved with either phased array antennas, as MoVR uses for millimeter-wave transmissions [1], directional antenna designs themselves (as demonstrated above in §3), or a combination of the two. More directional antennas would have a larger effect on a given link, but are more selective and may be more difficult to embed into the environment (the parabolic antennas used in §3 would be difficult to embed into walls but log-periodic or custom PCB antennas could be used and still offer directionality). Note that through phase-coherent signal

combining [9] a large number of less directional antennas could emulate a single highly directional antenna, so PRESS could use either few well-placed directional antennas or many randomly placed but less directional antennas, or anything in-between.

Passive-active hybrid, and multi-tier designs. A small number of active PRESS elements might replace several more passive elements. As noted in §3, these active elements can help effect changes on line-of-sight links as well as reducing the overall PRESS array size. Power issues for the active elements could be addressed with energy harvesting devices. Further, we might divide the elements into groups, to harness diversity or power gains within each group and multiplex across groups, analogous to how *Hekaton* [36] groups antennas for analog and digital beamforming respectively.

Number of reflection coefficients per element. As we discovered experimentally, allowing each PRESS element to be tuned to different, finely-spaced phases increases the likelihood that the sum of reflected signals will constructively interfere, or cancel, at the receiver, thus increasing the desirable system effect. We conjecture that around eight phase values along with the off state may provide sufficient resolution, but plan on testing with continuously-variable phase shifting hardware.

4.2 Control Plane Choices

Mechanism. As in wired software-defined networks, there are benefits associated with a centralized or semi-centralized controller that has a more global view of the network conversations. This necessitates a mechanism by which the controller can actuate all the array elements rapidly. Likely wireless control plane candidates are low-frequency, low-rate bands (perhaps ISM or whitespace frequencies) that penetrate walls well and travel long distances. Other candidates include ultrasound in order to easily scope the control to a single indoor room, as well as wires between some subsets of the array elements.

Navigating the search space. With N PRESS elements, each having M possible reflection coefficients, enumerating the M^N possibilities in the search space for the optimal configuration becomes impractical. We will focus the search in the vicinity of intended receivers, and apply heuristics to prune the space. Other likely possibilities include the application of convex optimization [3] or machine learning techniques, as Remy [35] has used in congestion control.

5 CONCLUSION AND FUTURE WORK

Our immediate next steps include prototyping and experimenting with larger arrays of smaller antennas, to bring us closer experimentally to the vision outlined in §1. We anticipate that as long as the amount of antenna material in a PRESS array is electrically significant relative to the (likely non-metallic) interior wall construction, that PRESS array will have a significant effect on the radio environment.

REFERENCES

[1] O. Abari, D. Bharadia, A. Duffield, and D. Katabi. Cutting the cord in virtual reality. In *HotNets*, 2016.

[2] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly. Design and experimental evaluation of multi-user beamforming in wireless LANs. In *MobiCom*, 2010.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[4] V. R. Cadambe and S. A. Jafar. Interference alignment and degrees of freedom for the *K*-user interference channel. *IEEE Trans. on Info. Theory*, 54(8):3425–3441, Aug. 2008.

[5] J. Chan, C. Zheng, and X. Zhou. 3D Printing Your Wireless Coverage. In *HotWireless*, 2015.

[6] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A case for adapting channel width in wireless networks. In *SIGCOMM*, 2008.

[7] K. Chintalapudi, B. Radunovic, V. Balan, M. Buettener, S. Yerramalli, V. Navda, and R. Ramjee. WiFi-NC: Wi-Fi over narrow channels. In *NSDI*, 2012.

[8] S. Gollakota, S. Perli, and D. Katabi. Interference alignment and cancellation. In *SIGCOMM*, 2009.

[9] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. 802.11 with multiple antennas for dummies. *SIGCOMM CCR*, 40(1):19–25, Jan. 2010.

[10] E. Hamed, H. Rahul, M. A. Abdelghany, and D. Katabi. Real-time Distributed MIMO Systems. In *SIGCOMM*, 2016.

[11] P. Hu, P. Zhang, and D. Ganesan. Laissez-Faire: Fully asymmetric backscatter communication. In *SIGCOMM*, 2015.

[12] P. Hu, P. Zhang, M. Rostami, and D. Ganesan. Braidio: An integrated active-passive radio for mobile devices with asymmetric energy budgets. In *SIGCOMM*, 2016.

[13] S. A. Jafar and M. J. Fakhereddin. Degrees of freedom for the MIMO interference channel. *IEEE Trans. on Info. Theory*, 53(7):2637–2642, July 2007.

[14] B. Kellog, A. Parks, S. Gollakota, D. Wetherall, and J. Smith. Wi-Fi backscatter: Internet connectivity for RF-powered devices. In *SIGCOMM*, 2014.

[15] N. Kita, W. Yamada, A. Sato, D. Mori, and S. Uwano. Measurement of Demel condition number for 2×2 MIMO-OFDM channels. In *IEEE VTC*, 2004.

[16] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi. Bringing cross-layer MIMO to today's wireless LANs. In *SIGCOMM*, 2013.

[17] Laird GD24BP Parabolic Antenna. Datasheet.

[18] K. Lin, S. Gollakota, and D. Katabi. Random access heterogeneous MIMO networks. In *SIGCOMM*, 2011.

[19] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. Smith. Ambient backscatter: Wireless communication out of thin air. In *SIGCOMM*, 2013.

[20] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. DIRC: Increasing indoor wireless capacity using directional antennas. In *SIGCOMM*, 2009.

[21] K. Nikitopoulos, J. Zhou, B. Congdon, and K. Jamieson. Geosphere: Consistently turning MIMO capacity into throughput. In *SIGCOMM*, 2014.

[22] G. Nikolaidis, M. Handley, K. Jamieson, and B. Karp. COPA: Cooperative power allocation for interfering wireless networks. In *CoNEXT*, 2015.

[23] PulseLarsen W1030 Omni Antenna. Datasheet.

[24] Peregrine Semiconductor PE42441 UltraCMOS SP4T RF Switch (10 MHz–8 GHz). Datasheet.

[25] Y. Qiao, O. Zhang, W. Zhou, K. Srinivasan, and A. Arora. PhyCloak: Obfuscating Sensing from Communication Signals. In *NSDI*, 2016.

[26] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-aware rate adaptation and MAC protocols. In *MobiCom*, 2009.

[27] H. Rahul, S. Kumar, and D. Katabi. JMB: Scaling Wireless Capacity with User Demands. In *SIGCOMM*, 2012.

[28] C. Shepard, H. Yu, N. Anand, L. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: Practical many-antenna base stations. In *MobiCom*, 2012.

[29] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang. Fine-grained channel access in wireless LAN. In *SIGCOMM*, 2010.

[30] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker. SAM: Enabling Practical Spatial Multiple Access in Wireless LAN. In *MobiCom*, 2009.

[31] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[32] D. Vasisht, S. Kumar, H. Rahul, and D. Katabi. Eliminating channel feedback in next-generation cellular networks. In *SIGCOMM*, 2016.

[33] Rice Univ. WARP platform (v. 3). Website.

[34] T. Wei, A. Zhou, and X. Zhang. Facilitating robust 60 ghz network deployment by sensing ambient reflectors. In *NSDI*, 2017.

[35] K. Winstein and H. Balakrishnan. TCP ex machina: Computer-generated congestion control. In *SIGCOMM*, 2013.

[36] X. Xie, E. Chai, X. Zhang, K. Sundaresan, A. Khojastepour, and S. Rangarajan. Hekaton: Efficient and practical large-scale MIMO. In *MobiCom*, 2015.

[37] Q. Yang, X. Li, H. Yao, J. Fang, K. Tan, W. Hu, J. Zhang, and Y. Zhang. BigStation: Enabling scalable real-time signal processing in large MU-MIMO systems. In *SIGCOMM*, 2013.

[38] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Zhao, and H. Zheng. Mirror mirror on the ceiling: Flexible wireless links for data centers. In *SIGCOMM*, 2012.