

Good Times For Wireless Research

Rahman Doost-Mohammady, Oscar Bejarano, and Ashutosh Sabharwal
Rice University

ABSTRACT

Wireless communications has been one of the major success stories of engineering, and as a field of research, its future might be even brighter than its past. The field will need many breakthroughs to achieve the grand vision of next-generation networks, and hence, it is important to empower thousands of researchers. One of the challenges is how do we empower experiment-based wireless research at a scale and speed not possible today. In this paper, we briefly discuss the challenges faced by the experiment-based wireless research and ongoing efforts to address those challenges.

CCS CONCEPTS

• **Hardware** → **Beamforming**.

KEYWORDS

Wireless Testbeds, Massive MIMO, Beamforming

ACM Reference Format:

Rahman Doost-Mohammady, Oscar Bejarano, and Ashutosh Sabharwal . 2020. Good Times For Wireless Research. In *Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH '20)*, September 21, 2020, London, United Kingdom. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3411276.3414557>

1 INTRODUCTION

The first demonstration of wireless communication was performed by Marconi in 1896 over a distance of 3 Km, and by 1901, he had demonstrated successful communication over a distance of 3000 Km [10]. Marconi's seminal work led to a Nobel Prize in 1909. In 1948, Shannon published his landmark paper [20] to establish the conceptual and theoretical foundations for reliable communications. The two breakthroughs have been an inspiration for many generations of engineers, and foundational to today's wireless networks.

The results by Marconi and Shannon also serve as a reminder that breakthroughs come in many forms – from disruptive experimental demonstrators to disruptive theoretical contributions. Both have the power to expand what we consider possible. At the very core, all scientists and engineers are inspired to achieve similar feats in their careers - to move the field of their research forward, either in our fundamental understanding or by designing novel systems, or preferably both. In this paper, we focus on experiment-based wireless research.

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 ACM 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.

WiNTECH '20, September 21, 2020, London, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8082-9/20/09.

<https://doi.org/10.1145/3411276.3414557>

Experiment-based wireless research, which often requires extensive laboratory facilities, tends to face two key barriers. The first one is *access* - experimental research laboratories often require significant financial resources and hence can be out-of-reach for many research groups. This lack of research infrastructure also impacts training and education, as laboratory-based pedagogy needs laboratory resources and local expertise to teach them. The second challenge is that of *replication*. Since each lab ends up selecting their special mix of hardware and experimental environments, replicating experiments and reproducing experiments faces a large barrier. As a result, building upon results of prior research becomes a slow process, leading to what we label as “amplification loss” in experimental research. In contrast, an analysis result (e.g., a theorem) enjoys a natural advantage - fewer external resources (like a laboratory) are required to build upon past research.

In this paper, we outline our personal perspective on why and how we are contributing to solving the two above-mentioned challenges in conducting novel experimental research. The paper thus serves as an update for the research community on our ongoing work as part of the POWDER-RENEW platform, which consists of developing and deploying a highly capable software-defined massive MIMO (mMIMO) wireless *network*.

2 SHOULD WE EXPECT MORE BREAKTHROUGHS?

All successful fields have to continually contend with the question - *are there any big breakthroughs on the horizon?* There is no doubt connectivity, wired and wireless combined, has fundamentally reshaped our lives. Wireless connectivity alone is more than a trillion-dollar industry and connects billions of people. In the late 2000s, the combination of 4G/LTE and smartphones brought video streaming to mobile clients, which was often stated as the grand challenge for wireless in the 1990s. Once achieved, the wireless research community started to wonder if there was a future for wireless research [4]. Fortunately, the demand for wireless data continued to explode and the conversation shifted from not just connecting people but also devices, with the introduction of broad concepts like Internet-of-Things and machine-to-machine communications, giving birth to 5G and beyond.

Simultaneously, academic research kept its pace and the last decade saw many new research ideas take hold, like (i) mMIMO [12] that is now an integral part of cellular systems, (ii) communications in mmWave bands [16] that is being adopted in both cellular and Wi-Fi, and (iii) in-band full-duplex, that is now part of wireless products and in DOCSIS 4.0 cable modem standards [3, 5]. We note that the three examples, mMIMO, mmWave, and full-duplex, followed different routes. mMIMO was an analytical demonstration of a way to scale network capacity. mmWave adoption relied on measurement studies that demonstrated its potential as a viable band

for cellular communications. In-band full-duplex was an experimental demonstration that was followed by analytical and systems research.

In our opinion, the right question to ask is “what is the driving vision?” and once answered, meeting the grand challenge almost always inspires breakthroughs to achieve the grand challenge. For example, the vision of 5G to connect *everything* to the Internet at sub-millisecond latency and Gigabits-level throughput is an excellent example of a grand driving vision. Such a vision statement can focus the global research community on first agreeing on a common “Mount Everest” to climb, and then developing novel ways of reaching the peak. We can be confident that this purpose-driven research will inspire new breakthroughs.

None of us can predict the form of the breakthrough (else we would invent it ourselves!) or who will be responsible for it. Instead, we ponder a specific question related to empowering breakthroughs, in the context of experiment-based research.

What research resources could accelerate the innovation process by enabling “amplification gain” while ensuring scientific rigor?

3 SHARED PLATFORMS AND TESTBEDS

In this section, we start by reviewing the 3Rs of scientific research and then discuss how shared research testbeds can address these.

3.1 3Rs of Scientific Research

The 3Rs of a scientific research process are *repeatability*, *replicability*, and *reproducibility* [13]. Repeatability relates to the ability of the original team to replicate their own experiments using their own experimental setup. Repeatability in the originating lab is a crucial step in validating that the observed results are not outliers due to random effects and more importantly, can be explained if reliably repeated. Replicability relates to a different team being able to use the same apparatus as the original team to arrive at the same results as the original team. One can view this step as being crucial for review purposes and for ensuring a specification of the experiment that removes any special expertise of the original team to successfully conduct the experiment. Reproducibility represents arriving at the same results using potentially different but equivalent apparatus, conducted by another team. This last R elevates the original concept into a scientific principle.

The wireless research community has the luxury to use theory, simulations, and experiments to establish new results. For example, we are blessed with strong theoretical foundations, e.g., information, communications, and networking theory. For theoretical contributions, the 3Rs have always been straightforward. Disclosing all the steps of an analytical result is a requirement for publishing, and hence for ensuring replicability. Expressed in the universally accepted language of mathematics, the only “hardware” apparatus needed are a pen and paper, and we know that a theorem is the same theorem on any paper inked with any pen (digital or otherwise). Simulations-based results faced the replicability and reproducibility challenge in the early days. However, with the advent of scientific software packages, personal computing, and open-source movement (e.g., a network simulator project,

<https://www.isi.edu/nsnam/ns/>), meeting the 3R’s of simulation-based results has become straightforward. Since simulations are a computer program, they are (in concept) a mathematical description like in an analysis paper.

However, meeting the 3Rs for experiment-based research remains challenging even today. Repeatability may be the easiest R, but replicability and reproducibility remain challenging. Building a lab capable of reproducing another result is a resource-intensive endeavor, and thus, comparing against past results is often a very challenging step. Reproducibility is also severely hampered by the fact that papers describe experiments in a haphazard manner - *there is no rigorous language for describing experiments*. Unfortunately, this situation is not specific to wireless research, it is a challenge for all sciences and a topic of many discussions [13].

One could argue that searching for a solution may not only be hard (if it was not hard, the scientific community would have had solved it already) but also unneeded. The progress in science and engineering in the last two centuries has been fantastic by any measure, and hence whatever we are doing as a community is working. However, we will contend that our inability to achieve the 3Rs in experimental research creates two challenges, as described below.

The first concern is a well-discussed challenge – we are publishing potentially irreproducible science and engineering. Imagine, trying to leverage an innovative idea for ultra-low latency method to control a mission-critical robot in a nuclear plant, and finding that there were holes in the original research design and reporting, rendering the whole idea incorrect.

The second concern is that experiment-based results and research is suffering from “amplification loss.” If the process of innovation had measurable speed or intensity, then the presence/absence of 3Rs will be found to directly contribute to that speed. For example, if researcher A publishes a theorem, researcher B can understand the details and build on that result, thereby “amplifying” the results of researcher A. However, since experiments are challenging to replicate and reproduce, we contend that experiments-based results have a very small amplification factor.

3.2 Shared Testbeds As One Answer

A shared platform allows us to go beyond repeatability by encouraging the entire research community to replicate and build upon the work of others. Certain shared platform models can even enable reproducibility by permitting users to bring their own devices into an existing infrastructure. The best example of a shared platform is in computing - a general-purpose computer combined with an operating system supporting many applications allows sharing, reproduction, and extension by many researchers. The same concept was applied to wireless research with two popular platforms in the late 2000s, namely GNURadio [2] built on USRP [19] and WARP [23]. The two hardware platforms were the first two widely available software-defined radios, which allowed researchers to buy a general-purpose wireless research “computer.” The open-source frameworks GNURadio and WARP allowed many researchers to rapidly build new ideas; collectively the two platforms were used by hundreds of research groups worldwide, leading to thousands of papers using it for their research. In concept, the two platforms

allowed researchers to partially share an experiment by publishing their code, as the experimental conditions cannot be shared. However, the ORBIT shared testbed [17] made that possible, as researchers could replicate both the code and conditions used by other researchers, thereby proving that

Shared wireless research testbeds can enable a common language both for the description of the experiments and experimental conditions.

Now imagine that every paper that uses a shared testbed for their experiments had to publish their (i) code used to conduct the experiment and details of the experimental conditions, (ii) the data collected during that experiment, and (iii) the code used to convert the data into results. Since programs are nothing but mathematical steps, disclosing all three parts make the disclosure of an experimental paper equivalent to providing proof in an analytical paper, even the experimental conditions (at least statistically). It is true, that checking the programs is non-trivial but at least, the details are available for others to verify and more importantly, to build upon. The end result is the tantalizing possibility where

Shared testbeds and open-code, data, analysis can enable the much needed research “amplification” gain currently lacking in wireless experimental research.

For shared testbeds to fulfill the above promise, it is important that they are *capable* and *extensible*. Both elements are crucial to enable researchers to explore ambitious ideas.

There are many shared, capable, and flexible wireless testbeds, both as part of the NSF PAWR program (COSMOS [18] and AER-PAW [15]) and as part of the global effort, e.g. 5TONIC [9] in Europe and 5G Testbed Project in India [14]. The emergence of multiple large-scale shared wireless testbeds focused on different research sub-themes is evidence that the research community sees significant value in the model. In the next section, we describe the POWDER-RENEW testbed.

4 POWDER-RENEW TESTBED

The POWDER-RENEW testbed is an at-scale, diverse spectrum, software-defined *radio network* deployment. Among many unique elements, the Rice RENEW part of the testbed relates to mMIMO technology, which is the primary focus of this paper.

4.1 It’s a Synchronized mMIMO Network!

The POWDER testbed will deploy multiple 64-antenna mMIMO RENEW base stations, operating in two key 5G bands, i.e. 2.5 GHz and 3.6 GHz, along with many fixed and mobile client nodes across the University of Utah campus in Salt Lake City. The base stations’ coverage area will overlap while mobile clients can move in and out of individual or overlapped cell coverage areas.

Collectively, the mMIMO testbed will include hundreds of base-station antennas. To enable coherent receive processing and joint beamforming across *all* antennas, all base stations will be tightly synchronized both in time and frequency through an optical fiber network. Thus, it will allow at-scale experiments that are currently not possible in any shared network and will provide the community with a unique capability for diverse research areas.

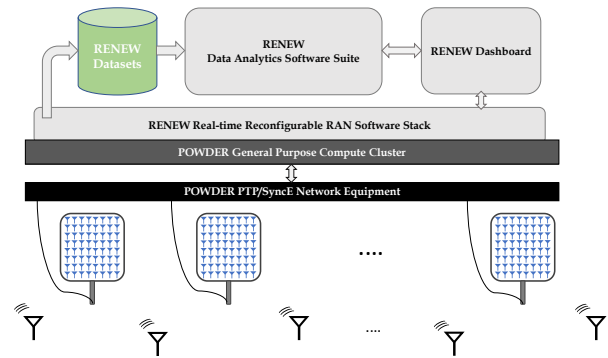


Figure 1: POWDER-RENEW mMIMO Testbed Architecture.

To achieve tight synchronization, the project is considering the use of IEEE 1588 or Precision Time Protocol (PTP) and Synchronous Ethernet (SyncE); these protocols allow synchronization of many distributed devices with sub-nanosecond accuracy. This level of synchronization using PTP and SyncE was first developed and demonstrated at CERN’s Large Hadron Collider [11] where many physical observations performed kilometers apart had to be time-synchronized. The synchronization technology is now commercially available under the name *White Rabbit* from Seven Solutions [21]. Moreover, network product vendors today offer commercial switches that support both SyncE and PTP protocols to be used in many domains, e.g. in datacenter and Telecom applications. Both solutions are currently being considered for the implementation of network synchronization in POWDER-RENEW.

4.2 It’s a Computing mMIMO Network!

With 64 antennas per base-station, the amount of data to be processed at each base station can become a heavy burden on the system. Even if sampling at only 5 MHz, the amount of data to be processed reaches up to 10.2 Gbps. As a result, some of the most pressing computing challenges in realizing mMIMO networks are in (i) meeting rigid timing constraints, and (ii) sharing large amounts of data between base stations and also with other edge cloud compute resources.

The RENEW platform architecture features many processing elements that are interconnected through both dedicated and general-purpose buses. Within each RENEW mMIMO base station, there are 32 software-defined radios, each equipped with System-on-Chip (SoC) components that include both FPGA fabric and ARM processors. All SoCs are interconnected through high-speed (12.8 Gbps) low-latency dedicated buses that could exchange data through Ethernet frames. Similarly, within each SoC, Ethernet frames can flow between ARM processors and FPGA fabric. Overall, through Ethernet and IP routing, data can flow to each element within an mMIMO base station, thereby creating a highly flexible heterogeneous computing array.

Additionally, RENEW base stations are interconnected with each other as well as with other edge compute elements composed of general-purpose servers, GPUs, and FPGA cards through fiber. Each

base station supports network speeds of up to 40 Gbps. This architecture permits the distribution of computation among all these elements, consequently enabling research on distributed computing solutions to meet the ultra-low latency goals of next-generation networks.

4.3 It's *Bare-metal* Open-Source!

The key feature of the POWDER-RENEW deployment is that it is accompanied with active development of an open-source software/firmware with the overarching goal of a fully-functional, field-operational, research radio access network (RAN) system that can be modified at every layer to enable research at all layers of the network stack. Thus, the researchers can have *bare-metal* access to the platform. We believe that the complete openness of the platform is a crucial capability for innovative cross-layer research.

Given the complexity of the system, the RENEW software suite provides a fast-track development framework, called RENEWLab [22], with the aim to kickstart mMIMO research and experimentation. RENEWLab provides many SDR applications that run in diverse development environments, specifically Matlab, Python, and C++, targeting various research needs. These applications range from test scripts for generating and transmitting signals over the air, to a more complex real-time channel measurement framework. In addition, it provides a *Data Analytics* framework for post-processing datasets that are recorded during live experiments. A *Dashboard* software is used to provide better user experience for reconfiguration and data management.

4.4 An Example RENEW Application

The ability to characterize the wireless channel is key in developing many innovative ideas. Conversely, channel traces are sufficient to validate the feasibility of many novel ideas. For instance, to validate the performance of a new precoding technique in mobile environments without having to build a fully-functional system, researchers can collect channel traces from various mobility scenarios and evaluate performance in an emulated environment; examples of past novel research contributions using real channel traces are discussed in §5.

The RENEWLab *sounder* suite supports real-time mMIMO channel measurement with the ability to capture and record live channel traces through active interaction of clients and base station antennas. Through a reconfigurable software/hardware framework, the user can configure client devices to send uplink pilot signals of arbitrary length in arbitrary times during recurring frames, where the length of the frame is also configurable and can be as low as a few microseconds. An example of such a frame is shown in Fig. 2.

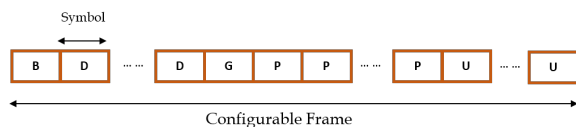


Figure 2: Example of a configurable frame for channel measurement. Symbol letters stand for Beacon, Guard, Pilot, Uplink, and Downlink, respectively.

To define a particular scenario for channel measurement, the user can specify the frame structure for base station antennas and clients through certain strings, composed of five letters, "B", "G", "P", "U", and "D", which stand for Beacon, Guard, Pilot, Uplink and Downlink, respectively. Each letter represents one or more OFDM symbols (subframe). Beacon signals are specific to base station antennas. A beacon is used to trigger over-the-air time synchronization at the clients. During a pilot symbol, the respective client transmits a user-defined pilot signal while all base station antennas receive it. During a guard symbol, the respective antenna is idle (no transmit or receive). For example, to collect time-orthogonal pilots from four clients and two uplink symbols, the base station can be assigned BGPUPU and while the four clients are given, DGPUGGUU, DGGPGGUU, DGGGPGUU, and DGGGGPUU, respectively. Note that, frame timekeeping and synchronization has been implemented in software and firmware (FPGA). As previously mentioned, FPGA reference design for base station and client radio elements are also made open-source, enabling researchers to build on the existing platform and innovate in real-time algorithm development.

Raw data (IQ samples) collected by the base station antennas, along with metadata specifying the experiment scenario and parameters used, are recorded in HDF5 file format [8] and can be post-processed using RENEW Data Analytics software suite, as shown in Fig. 1.

4.5 3Rs and POWDER-RENEW

RF hardware platforms are not only expensive but also difficult to set up and maintain. The increase from a few radios in a typical legacy wireless system to dozens of them in a massive MIMO system only exacerbates these problems. For the large fraction of wireless researchers, the costs of acquisition, deployment, and maintenance of a testbed can be prohibitive. POWDER-RENEW provides researchers with access to deployed state of the art resources for wireless experimentation thereby eliminating the cost and hardware management barriers.

In addition, the shared nature of the platform provides an opportunity to go beyond repeatability and enable both replicability and reproducibility. The platform is comprised of both fixed and mobile clients. The mobile clients are installed in university shuttles with a fixed route to allow for *repeatability* in multiple types of scenarios. In each experiment, users can record any information about the experimental setup (i.e., metadata) for future use, thus facilitating repetition. The POWDER-RENEW team encourages the release of both code and datasets by researchers, e.g. as part of their published manuscripts. The practice of sharing all POWDER-RENEW codebase and experimental scenarios means that other research teams can statistically *replicate* results and leverage past work without significant effort. Finally, the platform has adopted a BYOD (Bring Your Own Device) model which provides researchers with the flexibility to leverage a well-established deployment and use their own devices for their experiments. For instance, a team could use a different SDR to *reproduce* the work shown by other teams on the already deployed hardware.

5 RESEARCH EXAMPLES

In this section, we provide two examples to demonstrate how highly capable experimental platforms can enable novel research directions.

5.1 mMIMO Full-duplex

In-band full-duplex wireless is perhaps one of the best examples of research facilitated by the availability of software-defined radios. Two teams published their results in 2010, using the two available software-radio platforms: [3] used USRP and [5] used WARP. The experimental demonstration was essential as in-band full-duplex wireless transmissions were considered impossible – not because of a theoretical limitation but because of practical limitations of finite-resolution analog-to-digital converters. The two demonstrations used different ideas to achieve self-interference suppression in the *analog domain* before the received signal reaches the analog-to-digital converter. The two demonstrations spurred a significant research activity in subsequent years. Around the same time, 5G was considering the adoption of mMIMO and hence a natural question was on how to enable mMIMO full-duplex. All proposed techniques in the literature used additional analog hardware to achieve self-interference reduction, an idea that was not desirable for large arrays.

In [6], we demonstrated an all-digital approach called SoftNull, to enable many-antenna full-duplex with only digital-domain modifications. In the SoftNull design, the array is partitioned into a set of transmit antennas and a set of receive antennas, and self-interference from the transmit antennas to the receive antennas is reduced by transmit beamforming. The method was designed to convert existing mMIMO systems into full-duplex systems without modification; see Figure 3. SoftNull was proposed as a layer below the physical layer, tasked to only reduce self-interference, and agnostic to the upper layer processing. Thus SoftNull can operate on the output of algorithms for downlink MU-MIMO (such as zero-forcing beamforming) without modifying their operation.

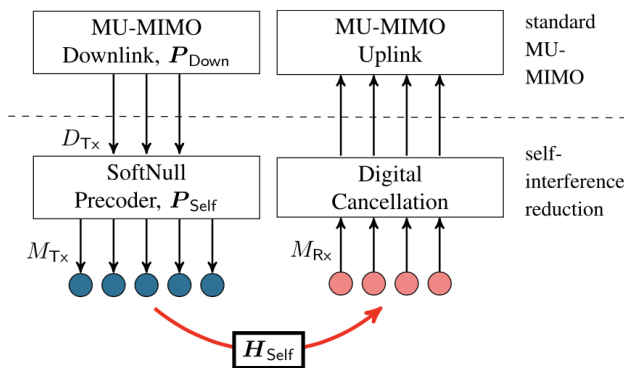


Figure 3: Softnull Architecture [6].

We collected channel measurements using a 72-element two-dimensional planar antenna array (see Figure 4), with mobile nodes placed in many different locations, measuring self-interference channels and uplink/downlink channels both outdoors, indoors,

and in an anechoic chamber; see Figure 5. The platform operates in the 2.4 GHz ISM band, with 20 MHz bandwidth. We use these real over-the-air channel measurements to simulate SoftNull and evaluate its performance extensively. The essence of the experimental results can be captured by the following two measurement-based conclusions.




Figure 4: The 72-element Argos platform (version 2) used for collection of Softnull dataset [6].




Figure 5: Different channel measurements scenarios [6].

SoftNull enables a large reduction in self-interference while sacrificing relatively few effective antennas. However, the amount of reduction depends on the environment: more scattering results in less suppression. In an outdoor low-scattering environment, SoftNull provides sufficient self-interference reduction while sacrificing only a few effective antennas. This fine understanding, based on channel environment, can also be explained using analysis. In [7], we extended the technique to design JointNull which considers joint beamforming and self-interference suppression and outperforms SoftNull significantly. Both SoftNull and JointNull datasets are open-source and available at <https://www.renew-wireless.org>.

Status report: In-band full-duplex has now been ratified in DOCSIS 4.0, the wireline cable modem standards, and is appearing in 5G base station products. The experimental demonstration was crucial to change the conversations around full-duplex; the availability of open-source platforms and public datasets is now crucial to move the concept from in-lab demonstrations to practical systems.

5.2 FDD mMIMO

A key challenge for frequency-division duplexing (FDD) mMIMO is the large overhead in acquiring channel state information for transmit beamforming. As a result, FDD operation of mMIMO was dubbed as the critical open question [1].

In [24], we conducted extensive channel measurements employing a 64-antenna base station as shown in Figure 6. The measure-



(a)

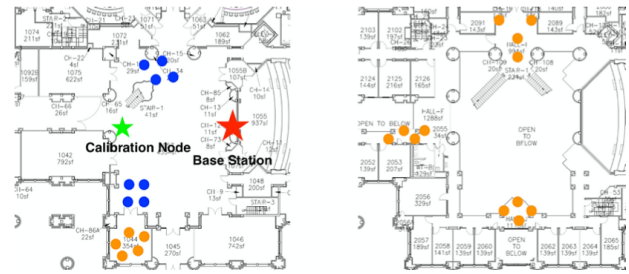


(b)


Figure 6: 64 antenna array used for collecting FDD measurements [24].

ments were performed over two 2.4 GHz ISM bands, separated by 72 MHz, for both indoor and outdoor environments; see Figure 7.

Using the measurement data, we discovered that only 4 downlink



(a)



(b)

Figure 7: (a): Indoor environment with the base station located indoor: locations of the base station and the mobiles in a typical office indoor environment. (b): Outdoor environment with the base station located outdoor: locations of the base station and the mobiles in an outdoor area of Rice University campus [24].

Angle-of-departures (AoDs) are sufficient to closely approximate the actual downlink channel, with an average correlation as high as 0.85. That is, the channel between one mobile antenna and the base station array can be characterized with much fewer parameters in the angular domain (4 complex channel coefficients) than antenna space (64 complex channel coefficients), thereby reducing the overall measurement dimension. Additionally, and more importantly, we show that at the base station, the uplink Angle-of-arrival (AoA) set has a strong correlation with the downlink AoD set, in that uplink AoAs are very close in number and magnitude to the downlink AoDs. Therefore, the estimated uplink AoA set can be directly applied as an estimated downlink AoD set. We then leveraged the insights from measurements to devise a new downlink training method, called *directional training*; shown in Figure 8. In directional training, the training symbols are sent in downlink AoDs that are estimated based on uplink AoAs. Since the number of dominant paths is very small, e.g. 4 paths, only a few downlink training symbols beamformed in specific AoDs are sufficient to estimate the gains and phases of the downlink AoDs.

Again, establishing that FDD mMIMO might be possible requires two important steps to happen. First, the measurements have to reveal some structure that could shed light on a solution. In our

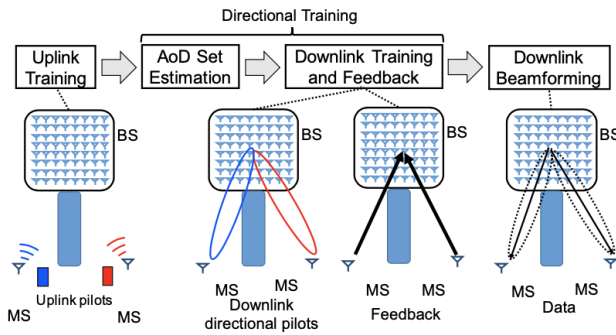


Figure 8: Steps in directional training [24].

case, that was the low-dimensionality of the channel in the angle-domain, compared to the number of antennas. And second, using that structure to derive a scalable channel estimation method that achieves good performance at a fraction of the overhead of the optimal method. Both steps require a flexible testbed; in this case, we leveraged the wideband nature of the testbed to study FDD systems.

6 CONCLUSIONS

Overall, we see a bright future for both the wireless research community and wireless technologies. Considering the number of ongoing shared testbed efforts, testbed design and use itself is becoming a mini-research area, and hence there is an opportunity to continue learning as the community gathers more experience in the coming years.

ACKNOWLEDGMENTS

This work was partially supported by NSF PAWR Grant 1827940 for the POWDER-RENEW project.

REFERENCES

- [1] E. Björnson, E. G. Larsson, and T. L. Marzetta. 2016. Massive MIMO: ten myths and one critical question. In *IEEE Communications Magazine*, Vol. 54. 114 – 123.
- [2] E. Blossom. 2004. GNU Radio: Tools for Exploring the Radio Frequency Spectrum. *Linux J.* 2004, 122 (2004), 4.
- [3] J. I. Choi, M. Jain, K. Srinivasan, P. A. Levis, and S. Katti. 2010. Achieving single channel, full duplex wireless communication. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*.
- [4] M. Dohler, R. Heath, A. Lozano, C. Papadias, and R. Valenzuela. 2011. Is the PHY layer dead? *IEEE Communications Magazine* 49 (2011), 159 – 165.
- [5] M. Duarte and A. Sabharwal. 2010. Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results. In *Proceedings of the Forty Fourth Asilomar Conference on Signals, Systems and Computers*.
- [6] E. Everett, C. Shepard, L. Zhong, and A. Sabharwal. 2016. Softnull: Many-antenna full-duplex wireless via digital beamforming. In *IEEE Transactions on Wireless Communications*, Vol. 15. 8077–8092.
- [7] N. M. Gowda and A. Sabharwal. 2018. JointNull: Combining partial analog cancellation with transmit beamforming for large-antenna full-duplex wireless systems. In *IEEE Transactions on Wireless Communications*, Vol. 17. 2094–2108.
- [8] The HDF5 Group. 2006. HDF5 Library and File Format. <https://www.hdfgroup.org/solutions/hdf5/>.
- [9] IMDEA Networks Institute. 2016. 5TONIC: Open 5G Research and Innovation Laboratory. <https://www.5tonic.org/>.
- [10] B. Jabbari. 1997. Introduction to the Classic Paper by Marconi. In *Proceedings of the IEEE*, Vol. 85. 1523–1525.
- [11] P. Loschmidt, G. Gaderer, N. Simanic, A. Hussain, and P. Moreira. 2009. White rabbit - sensor/actuator protocol for the CERN LHC particle accelerator. In *SENSORS, 2009 IEEE*. 781–786.

- [12] T. L. Marzetta. 2010. Noncooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas. *IEEE Transactions on Wireless Communications* 9, 11 (2010), 3590–3600.
- [13] S. L. McArthur. 2019. Repeatability, Reproducibility, and Replicability: Tackling the 3R challenge in biointerface science and engineering. In *Biointerphases*, Vol. 14.
- [14] Indian Institute of Science. 2018. 5G Testbed Project. <https://ece.iisc.ac.in/~5G-Testbed/>.
- [15] K. Powell, A.S. Abdalla, D. Brennan, V. Marojevic, R.M. Barts, A. Panicker, O. Ozdemir, and I. Guvenc. 2020. Software Radios for Unmanned Aerial Systems. In *Proceedings of the 1st International Workshop on Open Software Defined Wireless Networks* (Toronto, ON, Canada).
- [16] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez. 2013. Millimeter Wave Mobile Communications for 5G Cellular: It Will Work! *IEEE Access* 1 (2013), 335–349.
- [17] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh. 2005. Overview of the ORBIT Radio Grid Testbed for Evaluation of Next-Generation Wireless Network Protocols. In *Proceedings of WCNC*.
- [18] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejcki, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Maheshwari, P. Skrimponis, and C. Gutterman. 2020. Challenge: COSMOS: A City-Scale Programmable Testbed for Experimentation with Advanced Wireless. In *Mobicom* (London, United Kingdom).
- [19] Ettus Research. 2020. Universal Software-defined Radio Platform (USRP). <http://www.ettus.com>.
- [20] C. E. Shannon. 1948. A mathematical theory of communication. In *The Bell System Technical Journal*, Vol. 27. 379–423.
- [21] Seven Solutions. 2020. White Rabbit Synchronization and Deterministic Timing Solutions. <https://sevensols.com/>.
- [22] Rice University. 2020. RENEW Software Git Repository. <https://gitlab.renew-wireless.org/renew/renew-software>.
- [23] Rice University. 2020. Wireless Open Access Research Platform. warp.rice.edu.
- [24] X. Zhang, L. Zhong, and A. Sabharwal. 2018. Directional Training for FDD Massive MIMO. In *IEEE Transactions on Wireless Communications*, Vol. 17. 5183–5197.