Computer Networks 188 (2021) 107870

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Good times for wireless research

Rahman Doost-Mohammady[1,*], Oscar Bejarano[1], Ashutosh Sabharwal[1]

Rice University, United States of America

ARTICLE INFO

Keywords:
Wireless testbeds
Massive MIMO
Beamforming

ABSTRACT

Wireless communications has been one of the major success stories of engineering, and as a field of research, it's future maybe 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 discuss the challenges faced by the experiment-based wireless research and ongoing efforts to address those challenges. We discuss the RENEW platform as part of the POWDER-RENEW city-scale testbed in Salt Lake City, which is dedicated to the deployment of massive MIMO technology and the tools that are developed for experimentation.

1. Introduction

The first demonstration of a 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 [1]. Marconi's seminal work led to a Nobel Prize in 1909. In 1948, Shannon published his landmark paper [2] to establish the conceptual and theoretical foundations for reliable communications. The two breakthroughs have been 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.

Experiment-based wireless research, which often requires extensive laboratory facilities, tends to face two key barriers. 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*.

We also highlight how capable testbeds allow asking and answering important research questions. In Section 6, we describe our recent research results on mMIMO full-duplex [3] and frequency-division duplex (FDD) mMIMO [4]. In [3], we demonstrated that a precoding-based transmitter can (largely) eliminate the need for analog cancellation; analog cancellation is essential to achieve full-duplex operation with small number of transmit antennas. The result has strong practical implication, in that mMIMO arrays developed using *traditional transceivers* can achieve full-duplex operation. In [4], the mMIMO testbed allowed us to investigate channel reciprocity in FDD systems. Discovering a strong angle-reciprocity from measured data, we proposed a scalable channel estimation method that allowed us to reduce channel estimation overhead by an order of magnitude. This in turn can help address the biggest open challenges for achieving FDD mMIMO [5].

* Corresponding author.
E-mail address: doost@rice.edu (R. Doost-Mohammady).
[1] Equally contributed toward producing this work.

https://doi.org/10.1016/j.comnet.2021.107870
Received 16 November 2020; Received in revised form 17 January 2021; Accepted 25 January 2021
Available online 1 February 2021
1389-1286/

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 streaming video 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 [6]. 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 [7] that is now an integral part of cellular systems, (ii) communications in mmWave bands [8] 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 [9,10]. 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* [11]. 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 ensuring replicability. Expressed in the universally accepted language of mathematics, the only "hardware" apparatus needed is 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 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 simulations-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 [11].

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 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 late 2000s, namely GNURadio [12] built on USRP [13] and WARP [14]. 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. The ORBIT shared testbed [15] 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 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 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 explore ambitious ideas.

There are many shared, capable, and flexible wireless testbeds, both as part of NSF PAWR program (COSMOS [16], AERPAW [17], and Colosseum [18]) and as part of the global effort, e.g. Arena [19], 5TONIC [20], 5G Testbed Project in India [21]. The emergence of multiple large-scale shared wireless testbeds, focused on different research sub-themes, is evidence that the research community sees a 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 RENEW part of the testbed relates to mMIMO technology, which is the primary focus of this paper. Fig. 1 illustrates all the hardware and software elements of the POWDER-RENEW testbed related to mMIMO. In this section, we will review all these elements and describe their status and current development efforts as well as future directions for the testbed.

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 experiments at scale that are currently not possible in any shared network and provide the community a unique capability for diverse research areas.

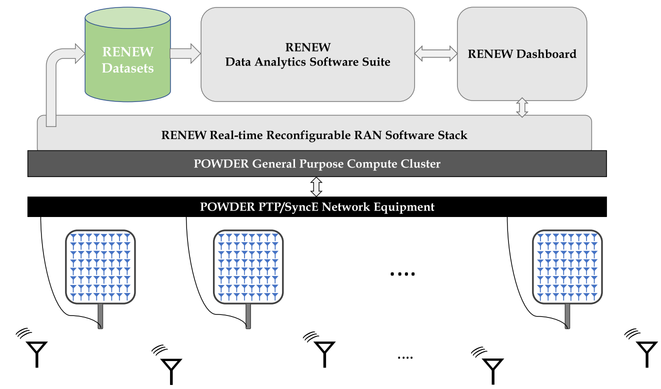

Fig. 1. POWDER-RENEW mMIMO Testbed Architecture.

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 [22] 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 [23]. Moreover, network product vendors today offer commercial switches that support both SyncE and PTP protocols to be used in many domains, e.g. in data center 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 is cumbersome. Even with 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 amount of data between base stations and also with other edge 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 radio (SDR) modules, each equipped with System-on-Chip (SoC) devices that include both FPGA fabric and ARM processors. With a combination of daisy-chaining and tree topologies, all SoCs are connected to a central *hub* through high-speed (12.8 Gbps) low-latency dedicated buses that could exchange data with Ethernet frame format. 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 a 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 Ethernet fiber where connectivity for each base station is enabled by up to 4x 10 Gbps Ethernet links. The architecture is shown in Fig. 2. Earlier works [24,25] provide extensive detail on the hardware design and architecture of RENEW base stations.

4.3. It's bare-metal open-source!

The key feature of POWDER-RENEW deployment is that it is accompanied with active development of an open-source software/firmware

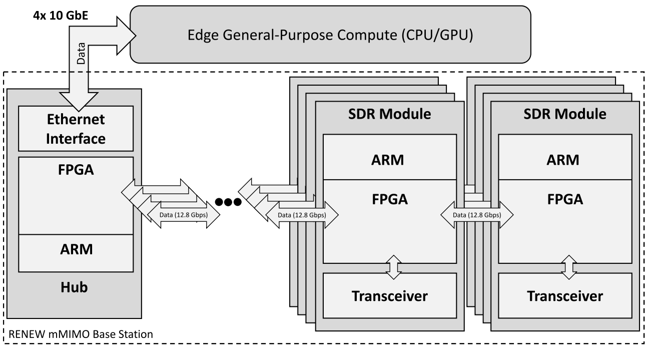

Fig. 2. RENEW mMIMO base station hardware architecture.

with the overarching goal of a fully-functional, field-operational, research radio access network (RAN) system. The system can be modified at every layer and thus, provides the researchers *bare-metal* access to the platform. We believe that the complete openness of the platform is a crucial capability for innovative cross-layer research.

RENEW software suite provides a fast-track development framework, called RENEWLab [26], 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 more complex real-time channel measurement framework. The latter is a real-time channel measurement and data recording software, called *sounder* which enables collecting large-scale mMIMO channel datasets.

In addition, RENEWLab provides a Python-based *Data Analytics* framework for post-processing datasets that are recorded during live experiments. A *Dashboard* software is used to provide hardware health monitoring and better user experience for reconfiguration and data management.

Most importantly, RENEW provides a highly configurable real-time mMIMO baseband processing software, named *Agora* [27]. Agora supports up to 64-antennas on the base station and is capable of zero-forcing beamforming to up to 16 spatial streams in both downlink and uplink directions. Agora is written in C++ and being further developed to support layer-2 features, such as HARQ and multi-user resource-block scheduling.

In Section 5, we elaborate on the features and use cases of the RENEW software.

4.4. 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 latter 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 own experiments. For instance, a team could use a different SDR to *reproduce* the work shown by other teams on the already deployed hardware.

5. RENEW software

The RENEW platform offers a suite of software tools that users can leverage for mMIMO research and development. Currently, the two major tools available consist of a configurable real-time channel sounder and a real-time baseband processing software.

5.1. Sounder: Real-time channel measurement

The ability to characterize the wireless channel is key in developing many innovative ideas in wireless communications. 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 Section 6.

The RENEWLab *Channel Sounder* supports real-time mMIMO channel measurement with the ability to capture and record live channel traces through active interaction between clients and base station antennas in a Time Division Duplex (TDD) mode. The software, as depicted in Fig. 3 includes two main components, (i) a multi-threaded network I/O component that receives pilot and uplink time-domain IQ samples of active mMIMO clients from all base station antennas, and (ii) a multi-threaded disk I/O that records the received samples in a multi-dimensional dataset with HDF5 file format [28]. Metadata, specifying the experiment scenario and parameters used, are also recorded along with the IQ samples. RENEW Data Analytics software provides a baseline tool for post-processing the datasets and plotting various metrics such as channel coherency and achievable rates for the recorded channels.

Both software and hardware can be configured through JSON files where users specify parameters such as the radio sample rate and carrier frequency, transceiver gains, TDD frame duration, OFDM symbol size, type of pilot signals, among others. Additionally, the user can configure the schedule for the TDD frames, e.g., the time slot for the time-orthogonal pilot signals from various users and the uplink and downlink data slots where all users or base station antennas transmit simultaneously. The length of each slot can be configured to accommodate an arbitrary repetition of OFDM symbols. Such flexibility allows frames as short as a few microseconds, thus enabling experimentation under very low channel coherence time conditions. An example of such a frame is shown in Fig. 4.

To define a particular scenario for channel measurement, the user can specify the transmission structure for base station antennas and clients, via a character string. Each string consists of a combination of five predefined letters, “B”, “P”, “U”, “G”, and “D”, which stand for, Beacon, Pilot, Uplink, Guard, and Downlink, respectively. Each letter represents a slot containing one or more OFDM symbols. 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 slot, the respective client transmits a user-defined pilot signal while all base station antennas simultaneously receive it. During a guard slot, the respective antenna is idle (no transmit or receive). For example, to collect time-orthogonal pilots from four clients and two uplink data slots, the base station can be assigned BGPPPPUU, while the four clients

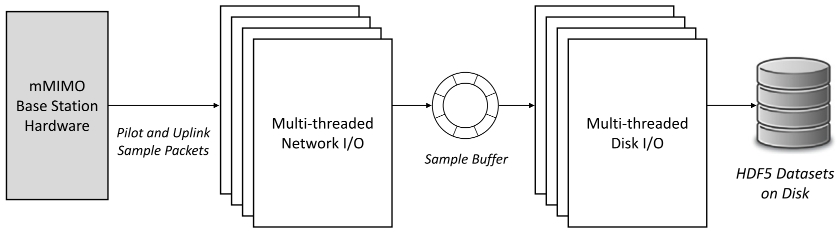

Fig. 3. Architecture of the Sounder Software.

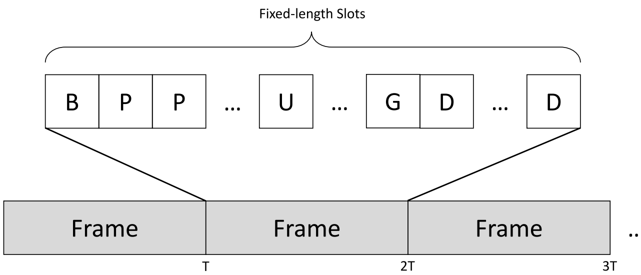

Fig. 4. Example of a configurable frame for channel measurement. Slot letters stand for Beacon, Pilot, Uplink, Guard, and Downlink, respectively.

are given, BGPGGGUU, BGGPGGUU, BGGGPGUU, and BGGGGPUU, respectively. Note that, frame-time keeping and synchronization has been implemented in both software and firmware (FPGA). An FPGA-based synchronization is helpful in high sample rate scenarios where high computational overhead may overwhelm the host machine. As previously mentioned, FPGA reference design for base station and client radio elements are also made open-source, enabling researchers to build on top of the existing platform and innovate on the development of real-time algorithms.

5.2. Agora: mMIMO baseband processing

mMIMO promises many-fold improvement in spectral efficiency through many-antenna multi-user beamforming in comparison with signal-antenna or small-scale MIMO, e.g. 4 × 4 MIMO, which was adopted in the fourth generation of cellular network technology. It is known that the spectral efficiency improvement in mMIMO scales with the minimum of the number of base station antennas and the number of spatial streams. However, with this improvement comes the cost of much larger computational complexity that also scales proportionally with the number of base station antennas. Such complexity puts mMIMO technology at odds with the ongoing trend of RAN virtualization and the adoption of cloud-RAN (C-RAN) in 5G cellular network deployment and operation. In C-RAN, major processing blocks in the radio access network, and the PHY layer in particular, are assumed to be running on general-purpose processors. Therefore, there is a big question in C-RAN standardization groups such as O-RAN [29], as to whether the current processors can handle the large computational cost associated with mMIMO, especially as the number of base station antennas are predicted to reach hundreds of antennas if not more, in the next decade [30].

Agora [27] addresses this problem by exploiting the rich data parallelism in mMIMO baseband processing and designs a threading model that scales the processing to all cores on a single many-core commodity server. A block diagram for the mMIMO baseband processing implemented by Agora is shown in Fig. 5. The base station software in Agora performs zero-forcing beamforming in both uplink and downlink directions. Additionally, it uses Low Density Parity Check (LDPC) for channel coding in both directions, simultaneously for all spatial streams. Agora relies on the Intel FlexRAN [31] open-source LDPC library for both downlink encoding and uplink decoding.

Table 1
Median and 99.9th Percentile Baseband Processing Latency in Agora [27].

Server Type	Number of cores	Median (ms)	99.9th (ms)
2 × Xeon E5-2697 v4, 2.3 GHz	32	1.34	1.38
4 × Xeon Gold 6130, 2.1 GHz	26	1.19	1.29
4 × Xeon Gold 6252N, 2.3 GHz	23	1.13	1.19
2 × Xeon Gold 6240, 2.6 GHz	23	1.12	1.15

Through extensive testing, Agora's real-time processing capability at 20 MHz bandwidth is demonstrated on several types of Intel many-core servers using Intel AVX-2 and AVX-512 extensions. Agora is shown to achieve a processing latency that is well below the 4 ms latency requirement for Enhanced Mobile Broadband (eMBB) use case of 5G-NR. Table 1 shows Agora's median and 99.9th percentile latency for processing 1 ms frames at 20 MHz in a 64 × 16 mMIMO setup, as reported by [27].

Agora also implements baseband processing for clients emulated as SDRs. Baseband processing software for multiple clients can be run on the same host machine while it is also possible to run each client independently on different computers. In addition, Agora includes an mMIMO channel simulator that allows an end-to-end emulation of the mMIMO systems between base station software and client software. While the simulator currently implements simple Rayleigh fading mMIMO channels, ongoing work includes linking the datasets generated by the Sounder to the simulator to enable realistic testing and development of mMIMO algorithms by the researchers without the need to access the mMIMO hardware.

Configuration of Agora is almost identical to the Sounder. JSON scripts are used to set many operational parameters for the software, including the frame schedule as it was described in Section 5.1. In addition to the parameters that were listed for the RENEWLab Sounder, the Agora user can set other parameters such as the modulation order and LDPC coding parameters to be used in the mMIMO baseband processing.

Fig. 6(a) shows a lab test bench of a 64-antenna RENEW mMIMO base station and eight mobile clients that is used to evaluate Agora. In the setup, 4 ms frames are used with each of the eight clients sending time-orthogonal full-band pilots based on Zadoff–Chu sequence as well as randomly-generated uplink data. Each slot includes 1 OFDM symbol of size 512 with 300 data subcarriers at the center and 64-QAM (6-bit) modulation for each subcarrier, corresponding to 1800 coded bits per symbol per user. LDPC coding with 1/3 rate is used, indicating 600 uncoded bits per OFDM symbol.

As reported in [27], Agora easily handles the 64 × 8 MIMO processing in this setup in real-time with block error rate (BLER) below 10%, the target rate defined by the 5G NR standard. Fig. 6(b) illustrates the worst BLER among beamformed users as the number of users increase from 1 to 8. With 8 users, BLER is still below the target which, considering 64-QAM modulation and 1/3 LDPC coding, demonstrates a nominal spectral efficiency of 16 bps/Hz is achieved.

Agora is an ongoing project that will continue to improve on these results and also add additional features, in particular, medium access and scheduling layers. Optimal multi-user scheduling in mMIMO is an active research area by itself and it is an important part of the network

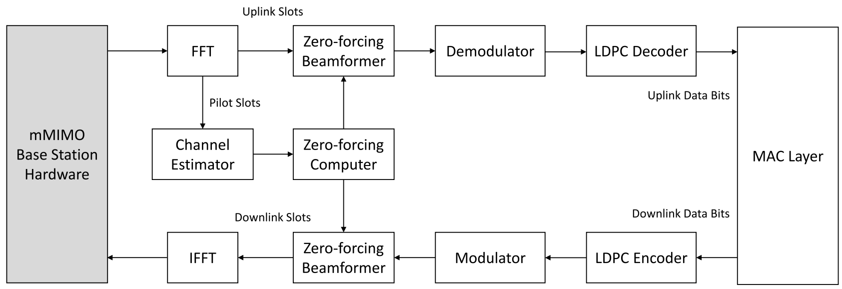

Fig. 5. mMIMO baseband processing block diagram implemented by Agora [27].

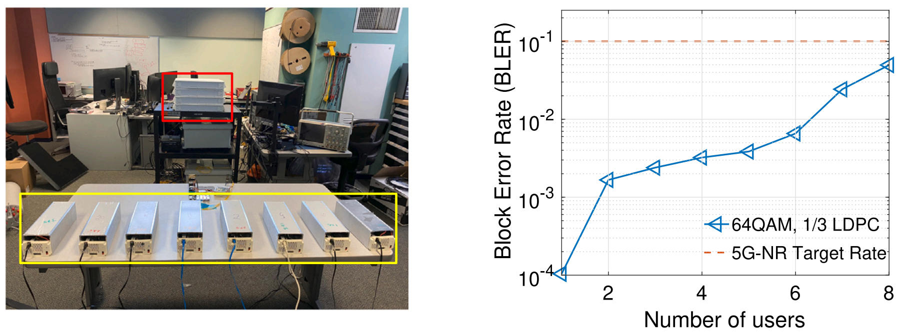

(a) Hardware setup including a 64-antenna RENEW base station, shown in a red box, and eight clients, shown in yellow box [27].

(b) Worst-user block error rate (BLER) vs. number of client uplink streams with Agora and a 64-antenna mMIMO base station [27].

Fig. 6. Performance evaluation of Agora on RENEW hardware [27].

stack that could bring about the real potential of mMIMO in terms of spectral efficiency. A baseline real-time implementation of users scheduling within Agora could open up tremendous opportunities for innovation by the user community.

6. Research examples

In this section, we provide two examples that demonstrate how a highly capable experimental platform can enable novel research directions.

6.1. 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: [10] used USRP and [9] used WARP. 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 to 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 [3], 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

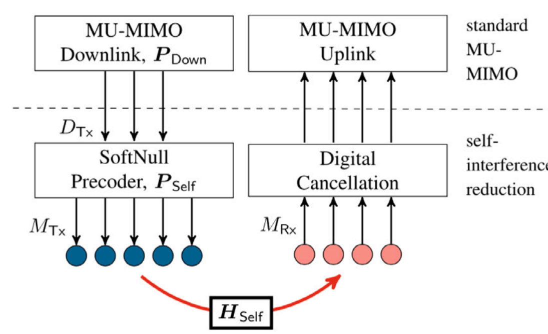

Fig. 7. Softnull Architecture [3].

from the transmit antennas to the receive antennas is reduced by transmit beamforming. In fact, the method was designed to convert existing mMIMO systems into full-duplex systems without modification; see Fig. 7. SoftNull was proposed as a layer below 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.

We collected channel measurements using the 72-element two-dimensional planar antenna array shown in Fig. 8, with mobile nodes placed in many different locations, measuring self-interference channels and uplink/downlink channels outdoors, indoors, and in an anechoic chamber; see Fig. 9. The platform operates in the 2.4 GHz ISM band, with 20 MHz bandwidth. We use these real over-the-air channel measurements to emulate SoftNull and evaluate its performance extensively. The essence of the experimental results can be captured by the following two measurement-based conclusions.

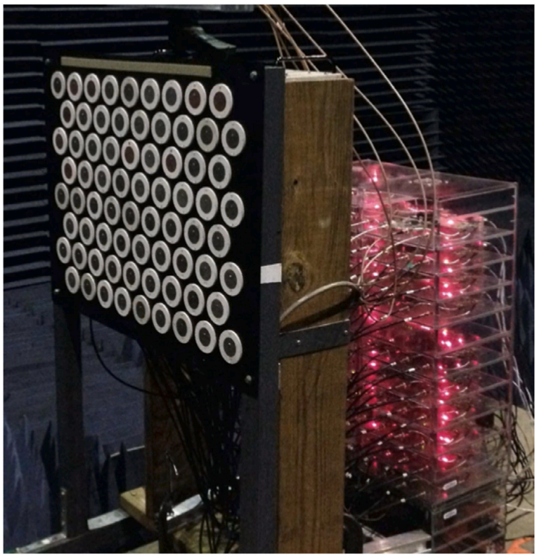

Fig. 8. The 72-element Argos Version 2 used for collection of Softnull dataset [3].

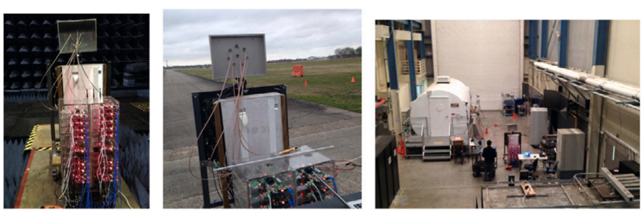

Fig. 9. Different channel measurements scenarios [3].

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 [32], 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 appearing in 5G base station products. The experimental demonstration was crucial to change the conversations around full-duplex. Similarly, the availability of open-source platforms and datasets is now crucial to move the concept from in-lab demonstrations to practical systems.

6.2. 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 [5].

In [4], we conducted extensive channel measurements employing a 64-antenna base station as shown in Fig. 10. The measurements were performed over two 2.4 GHz ISM bands, separated by 72 MHz, for both indoor and outdoor environments; see Fig. 11. Using the measurement data, we discovered that only 4 downlink Angle-of-departures (AoDs) are sufficient to closely approximate the actual downlink channel, with average correlation as high as 0.85. That is, the channel between

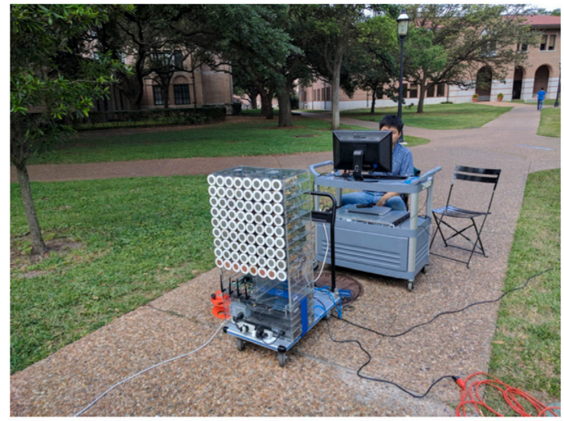

(a)

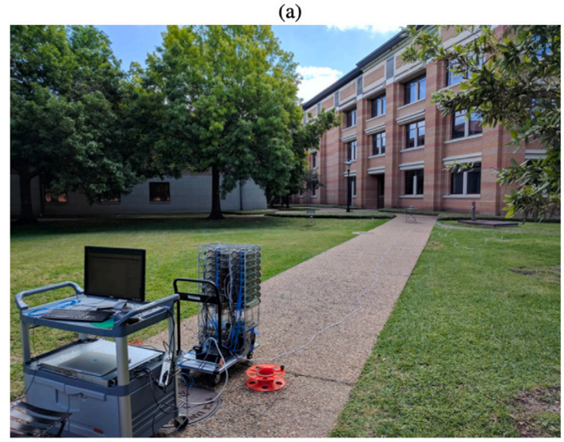

(b)

Fig. 10. 64 antenna array used for collecting FDD measurements [4].

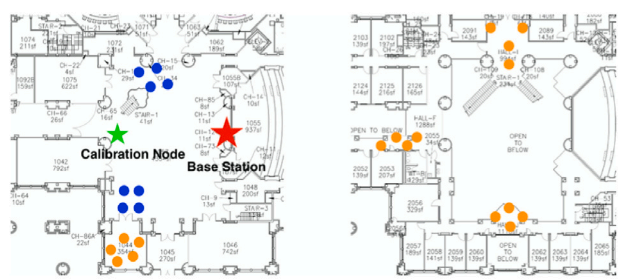

(a)

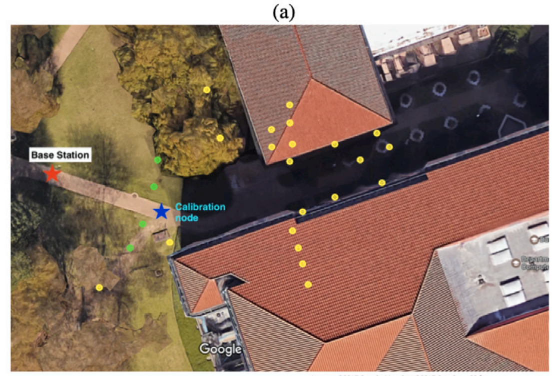

(b)

Fig. 11. (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 outdoor area of Rice University campus [4].

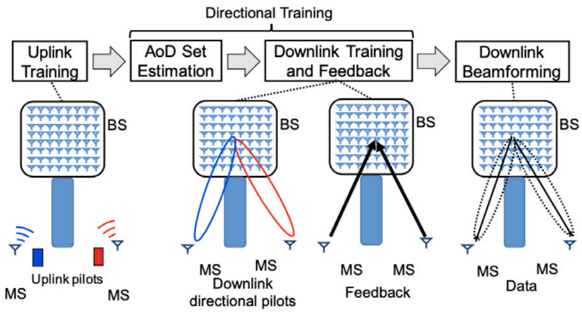

Fig. 12. Steps in directional training [4].

one mobile antenna and the base station array can be characterized with much fewer parameters in 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 estimated downlink AoD set. We then leveraged the insights from measurements to devise a new downlink training method, called *directional training*; shown in Fig. 12. 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 which could shed light toward a solution. In our 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.

7. 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 coming years.

Acknowledgments

This work was partially supported by NSF PAWR, United States of America Grant 1827940 for the POWDER-RENEW project.

References

[1] B. Jabbari, Introduction to the classic paper by marconi, Proc. IEEE 85 (10) (1997) 1523–1525.
[2] C.E. Shannon, A mathematical theory of communication, Bell Syst. Tech. J. 27 (3) (1948) 379–423.
[3] E. Everett, C. Shepard, L. Zhong, A. Sabharwal, Softnull: Many-antenna full-duplex wireless via digital beamforming, IEEE Trans. Wirel. Commun. 15 (12) (2016) 8077–8092.
[4] X. Zhang, L. Zhong, A. Sabharwal, Directional training for FDD massive MIMO, IEEE Trans. Wirel. Commun. 17 (8) (2018) 5183–5197.
[5] E. Björnson, E.G. Larsson, T.L. Marzetta, Massive MIMO: ten myths and one critical question, IEEE Commun. Mag. 54 (2) (2016) 114–123.
[6] M. Dohler, R. Heath, A. Lozano, C. Papadias, R. Valenzuela, Is the PHY layer dead? IEEE Commun. Mag. 49 (2011) 159–165.
[7] T.L. Marzetta, Noncooperative cellular wireless with unlimited numbers of base station antennas, IEEE Trans. Wireless Commun. 9 (11) (2010) 3590–3600.
[8] T.S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G.N. Wong, J.K. Schulz, M. Samimi, F. Gutierrez, Millimeter wave mobile communications for 5g cellular: It will work!, IEEE Access 1 (2013) 335–349.
[9] M. Duarte, A. Sabharwal, 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, 2010.
[10] J.I. Choi, M. Jain, K. Srinivasan, P.A. Levis, S. Katti, Achieving single channel, full duplex wireless communication, in: Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking, 2010.
[11] S.L. McArthur, Repeatability, Reproducibility, and Replicability: Tackling the 3R challenge in biointerface science and engineering, Biointerphases 14 (2019).
[12] E. Blossom, GNU Radio: Tools for exploring the radio frequency spectrum, Linux J. 2004 (122) (2004) 4.
[13] Ettus Research, Universal software-defined radio platform (USRP), 2020, http://www.ettus.com.
[14] Rice University, Wireless open access research platform, 2020, warp.rice.edu.
[15] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, M. Singh, Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols, in: Proceedings of WCNC, 2005.
[16] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejski, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Maheshwari, P. Skrimponis, C. Gutterman, Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless, in: Mobicom, 2020.
[17] K. Powell, A. Abdalla, D. Brennan, V. Marojevic, R. Barts, A. Panicker, O. Ozdemir, I. Guvenc, Software radios for unmanned aerial systems, in: Proceedings of the 1st International Workshop on Open Software Defined Wireless Networks, 2020.
[18] Northeastern University, Colosseum: The world's most powerful wireless network emulator, 2020, https://www.northeastern.edu/colosseum/.
[19] L. Bertizzolo, L. Bonati, E. Demirors, A. Al-Shawabka, S. D'Oro, F. Restuccia, T. Melodia, Arena: A 64-antenna SDR-based ceiling grid testing platform for sub-6 GHz 5g-and-beyond radio spectrum research, Comput. Netw. 181 (12) (2020).
[20] IMDEA Networks Institute, 5TONIC: Open 5g research and innovation laboratory, 2016, https://www.5tonic.org/.
[21] Indian Institute of Sciences, 5g testbed project, 2018, https://ece.iisc.ac.in/~5G-Testbed/.
[22] P. Loschmidt, G. Gaderer, N. Simanic, A. Hussain, P. Moreira, White rabbit - sensor/actuator protocol for the CERN LHC particle accelerator, in: SENSORS, 2009 IEEE, 2009, pp. 781–786.
[23] Seven Solutions, White rabbit synchronization and deterministic timing solutions, 2020, https://sevensols.com/.
[24] C. Shepard, R. Doost-Mohammady, J. Ding, R.E. Guerra, L. Zhong, ArgosNet: A multi-cell many-antenna MU-MIMO platform, 2017.
[25] C. Shepard, J. Blum, R.E. Guerra, R. Doost-Mohammady, L. Zhong, Design and implementation of scalable massive-MIMO networks, in: Proceedings of the 1st International Workshop on Open Software Defined Wireless Networks, 2020, pp. 7–13.
[26] Rice University, RENEW software git repository, 2020, https://gitlab.renew-wireless.org/renew/renew-software.
[27] J. Ding, R. Doost-Mohammady, A. Kalia, L. Zhong, Massive MIMO is a reality—What is next?: Five promising research directions for antenna arrays, Digit. Signal Process. 94 (2019) 3–20.
[28] The HDF5 Group, HDF5 library and file format, 2006, https://www.hdfgroup.org/solutions/hdf5/.
[29] O-RAN Alliance, Operator defined open and intelligent radio access networks, 2019, https://www.o-ran.org/.
[30] E. Björnson, L. Sanguinetti, H. Wymeersch, J. Hoydis, T. Marzetta, Massive MIMO is a reality—What is next?: Five promising research directions for antenna arrays, Digit. Signal Process. 94 (2019) 3–20.
[31] Intel, FlexRAN LTE and 5G NR FEC software development kit, 2019, https://software.intel.com/content/www/us/en/develop/articles/flexran-lte-and-5g-nr-fec-software-development-kit-modules.html.
[32] N.M. Gowda, A. Sabharwal, JointNull: Combining partial analog cancellation with transmit beamforming for large-antenna full-duplex wireless systems, IEEE Trans. Wirel. Commun. 17 (3) (2018) 2094–2108.

Rahman Doost-Mohammady is an assistant research professor at Rice department of Electrical and Computer Engineering since January 2020. He received his Ph.D., M.S. and B.S., all computer engineering from Northeastern University, Delft University of Technology and Sharing University of Technology in 2015, 2009, and 2006 respectively. His broad research interests include wireless systems and networking and embedded reconfigurable computing. He is also the technical lead for the RENEW project at Rice University aimed at building an open-source reconfigurable research testbed for massive MIMO.

Oscar Bejarano received his Ph.D. in Electrical and Computer Engineering from Rice University in 2015. He worked as an RF Systems Engineer at Cisco Systems until 2018. Since 2018, he has been a research engineer at Rice University for the RENEW project. His research interests lie in the field of wireless networking systems.

Ashutosh Sabharwal (Fellow, IEEE) received the Ph.D. degree from The Ohio State University in 1999. He then joined Rice University. He is the Founder of the WARP Project, an open-source project that was used at more than 125 research groups worldwide and used by more than 500 research articles. He demonstrated full-duplex wireless, a concept that has now been adopted in DOCSIS 4.0 cable modem standards. His research interests include wireless theory, protocols, and open-source platforms. He received the 2017 Jack Neubauer Memorial Award, the 2018 IEEE Communications Society Award for Advances in Communication, the 2019 ACM Sigmobile Test-of-Time Award, and the 2019 Mobicom Best Community Contribution Paper Award.