

# MillimeTera: Toward A Large-Scale Open-Source mmWave and Terahertz Experimental Testbed

Michele Polese<sup>1</sup>, Francesco Restuccia<sup>2</sup>, Abhimanyu Gosain<sup>2</sup>, Josep Jornet<sup>2</sup>, Shubhendu Bhardwaj<sup>3</sup>, Viduneth Ariyarathna<sup>3</sup>, Soumyajit Mandal<sup>4</sup>, Kai Zheng<sup>5</sup>, Aditya Dhananjay<sup>5</sup>, Marco Mezzavilla<sup>5</sup>, James Buckwalter<sup>6</sup>, Mark Rodwell<sup>6</sup>, Xin Wang<sup>7</sup>, Michele Zorzi<sup>1</sup>, Arjuna Madanayake<sup>3</sup>, Tommaso Melodia<sup>2</sup>

<sup>1</sup>University of Padova, Padova, Italy - <sup>2</sup>Northeastern University, Boston, MA USA - <sup>3</sup>Florida International University, Miami, FL, USA - <sup>4</sup>Case Western Reserve University, Cleveland, OH, USA - <sup>5</sup>Pi-Radio Inc., New York, NY, USA

<sup>6</sup>University of California, Santa Barbara, Santa Barbara, CA - <sup>7</sup>Stony Brook University, Stony Brook, NY, USA

## ABSTRACT

The promise of widespread 5th generation (5G) and beyond wireless systems can only be fulfilled through extensive experimental campaigns aimed at validating the large body of theoretical findings on millimeter wave (mmWave) and Terahertz (THz) frequencies. However, experimental research efforts in this field are often stymied by the lack of open hardware, open-source software, and affordable testbeds accessible by the research community at large, who is now forced to perform simulation-based research or – if at all possible – small-scale, *ad hoc* experiments. After discussing existing research challenges in mmWave and THz testbeds, in this paper we propose *MillimeTera*, a *vision* for a new generation of disruptive experimental platforms that will radically transform the *status quo* in mmWave and THz research. We next discuss our preliminary hardware and software efforts, and finally provide a roadmap of our main design and development goals in the years to come.

## CCS CONCEPTS

• Networks → Wireless access points, base stations and infrastructure; Network performance evaluation;

## KEYWORDS

5G, mmWave, Terahertz, Testbed, Experiments, Large-scale

## ACM Reference Format:

Michele Polese, Francesco Restuccia, Abhimanyu Gosain, Josep Jor-

net, Shubhendu Bhardwaj, Viduneth Ariyarathna, Soumyajit Mandal, Kai Zheng, Aditya Dhananjay, Marco Mezzavilla, James Buckwalter, Mark Rodwell, Xin Wang, Michele Zorzi, Arjuna Madanayake, Tommaso Melodia. 2019. *MillimeTera: Toward A Large-Scale Open-Source mmWave and Terahertz Experimental Testbed*. In *3rd ACM Workshop on Millimeter-wave Networks and Sensing Systems (mmNets'19), October 25, 2019, Los Cabos, Mexico*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3349624.3356764>

## 1 INTRODUCTION

Thanks to new applications in fields such as autonomous cars, augmented and virtual reality (AR/VR), sensing, space-borne Internet, and next-generation cellular networks, the need for wireless capacity is exploding, and has already increased by 100-1000x over the last decade [4]. To fulfill this unprecedented demand, the severe *spectrum crunch* in the traditional sub-6 GHz radio bands calls for the use of the millimeter wave (mmWave) and Terahertz (THz) frequencies, which will dominate the wireless world for the next 10-15 years [3, 8]. These fundamental advances will not come for free, and will necessarily force an end-to-end re-design of *every layer* of the wireless protocol stack [24].

However, despite the recent surge of mmWave and THz research [22], state-of-the-art advances still remain mostly theoretical in nature, with much emphasis on propagation and physical layer issues [6, 10, 11]. Today, experimental mmWave and THz wireless research is indeed monopolized by a small set of companies that spend tens of millions of dollars to build bleeding-edge prototyping platforms in-house, yet deliver mostly closed-source software and hardware architectures. In this harsh environment, the academic community is often unable to compete with industry to perform high-impact experimental research.

It is painful to see that our scientific community possesses enormous intellectual potential, but finds itself shut out of the real world of experimentation due to lack of access to versatile and powerful prototyping platforms. Thus, it is clear that the current state of affairs calls for immediate action at the forefront of experimental research. *This paper describes an*

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*mmNets'19, October 25, 2019, Los Cabos, Mexico*

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6932-9/19/10...\$15.00

<https://doi.org/10.1145/3349624.3356764>




Figure 1: In this effort, a number of nodes will be distributed to partnering universities, with either pairs of SDR units (SDR nodes, as shown in (a)) or mini-hubs (with multiple SDR nodes, as shown in (b)). The mini-hubs can be used for online and physical experiments (c).

effort to democratize mmWave and THz wireless research (60-240 GHz) by implementing a scalable, versatile, affordable, and bleeding-edge SDR prototyping platform named *MillimeTera*. Our goal is to build a significant number (close to 150) of mmWave to THz enabled nodes that will be deployed in a distributed hub-and-spoke model in universities across the United States and made accessible to the international research community.

Figure 1 shows an overview of our effort, which is expected to finally provide an experimental complement to purely simulation-based studies in the mmWave and THz spectrum. The community will thus receive a *scalable, versatile, programmable, and affordable SDR-based testbed that will enable cross-disciplinary and convergent experimental research that goes beyond computer-based mathematical models*. This vision of democratized access is bolstered by two key features: (a) all the hardware schematics and software will be open-source; and (b) the nodes will be locally or remotely accessible to all academic institutions. We will thus break the barrier to entry for experimental mmWave and THz research by jump-starting academic experimental research that will address end-to-end and system level challenges in these emergent frequency bands.

The remainder of the paper is organized as follows. In Section 2, we discuss existing open research challenges in mmWave and THz testbeds. We postulate our vision for the large-scale distributed testbed in Section 3, in which we discuss preliminary and future work in both hardware (Section 3.1) and software components (Section 3.2). Finally, Section 4 concludes the paper.

## 2 STATE OF THE ART OF MMWAVE AND THZ TESTBEDS

The mmWave band has been recognized as a cornerstone for 5th generation (5G) cellular networks and next-generation wireless local area networks. So far, the research efforts in this area have addressed plenty of pivotal questions, from propagation [10] to the Medium Access Control (MAC) layer and above [16]. The progress has been so substantial that the first mmWave-enabled commercial networks have been deployed, and a discussion on the support of THz communications for beyond 5G networks is already taking place [2, 11].

Despite the above silver linings, our research community is now painfully aware that existing simulation- and analysis-based research – as well as small-scale testbed validation – is not enough to produce groundbreaking advances beyond what is available today. Regrettably, existing research testbeds for mmWave and THz communications are either closed – because of commercial reasons or because they are based on Commercial Off-the-Shelf (COTS) devices – or extremely expensive. Specifically, testbeds with COTS equipment are presented in [13, 14, 18, 20]. These platforms have been instrumental in studying and understanding end-to-end behaviors of 60 GHz IEEE 802.11ad systems, for example with Multi-path TCP (MPTCP) [13]. Yet, the control of the protocol stack is inevitably limited to the reprogrammability of the devices' firmware. Hacking of low-level code has led to the possibility of controlling the beams [18], but the overall programmability is limited to a handful of operations.

Custom-built testbeds, on the other hand, allow the control of both the communication stack and the hardware. Critically, this capability requires a significant effort in terms of implementation and design, and often leads to *ad hoc* setups for small-scale experiments. The most significant effort in this area is *OpenMili*, presented in [23]. This testbed, characterized by open source software and hardware, features a 1 GHz bandwidth and a programmable protocol stack, and uses commercially-available Field Programmable Gate Arrays (FPGAs). However, it features a custom-designed phased array, which might be a show-stopper for networking research groups with limited expertise on RF design. This setup has been used, for example, to evaluate the performance of a possible approach for sub-6 GHz and 60 GHz integration [19]. Another small-scale testbed with a custom phased array and integration with traditional USRPs is presented in [1]. The authors, however, do not provide details on the bandwidth that is supported by the system, which is likely limited by the capabilities of the host processor. Finally, x60 [15], built on top of the National Instrument (NI) mmWave SDR platform [7], is able to achieve a bandwidth up to 2 GHz and has a configurable protocol stack. However, its cost (hundreds of thousands of dollars as of today) could ultimately prevent its widespread adoption by the research community.

At the THz frequencies, the access to experimental testbeds is even more challenging. For the time being, there are only a few experimental platforms in Asia, Europe and the USA, consisting of COTS or custom-designed up- & down-converters to THz frequencies connected to measurement equipment for channel characterization [12] or, in the better case, to high-performance arbitrary waveform generators and digital storage oscilloscopes for off-line signal processing and testing of physical-layer solutions. In any case, there is no control of the analog front-ends, which rely on fixed directional antenna systems, or real-time operation of the physical and higher layers.

## Research Challenges

We now discuss some of the research challenges (and possible solutions) in the mmWave domain that lack an experimental evaluation. They have been identified through an analysis of the available literature in the mmWave and THz areas, and through a survey directed to a group of 50 researchers participating in the NSF-sponsored mmWave Research Coordination Network Research Coordination Network (RCN).<sup>1</sup>

At the physical layer, experimental validation would address questions related to (i) waveform design (especially when considering carrier frequencies in the THz region); (ii) the trade-offs between performance and complexity for modulation and coding schemes, when aiming for data-rates

<sup>1</sup>[https://mmwrcn.ece.wisc.edu/?page\\_id=1594](https://mmwrcn.ece.wisc.edu/?page_id=1594)


in excess of 100 Gb/s; (iii) the design and practical implementation of channel estimation techniques tailored to ultra-broadband channels, i.e., 10 GHz or more of consecutive bandwidth; and (iv) appropriate beam design for mmWave and THz in mobile scenarios and with large antenna arrays. A testbed with a modular nature will also enable the large scale integration and testing of RF integrated circuits (RFICs) and new antenna designs.

At the MAC layer, instead, the main issues are related to (i) support of directionality, from initial access to beam tracking in mobile scenarios, and (ii) optimal scheduling strategies accounting for the characteristics of the channel. When considering end-to-end scenarios, a crucial aspect is understanding which capabilities should be supported by the lower layers of the protocol stack to enable efficient higher layer operations (*i.e.*, network, transport and application), and which new paradigms should be implemented in the whole protocol stack (*e.g.*, to natively support multi-connectivity, or the exchange of cross-layer information). From a system perspective, the availability of a large number of nodes will enable unprecedented experiments on network architectures and efficient mobility procedures that target reliability at such high frequencies. For example, multi-hop mesh networks could be studied, evaluating the trade-offs between spatial multiplexing and interference, resource allocation, and dynamic topology updates.

Large-scale, low-cost, open testbeds can benefit the research on applications of the mmWave and THz spectrum. Indeed, the integration of communication and sensing, considering the large amount of available bandwidth, would enable breakthroughs in vehicular, drone networks, and indoor positioning (*e.g.*, for e-Health applications). Additionally, cross-layer and high-rate multimedia solutions could be tested on a real and open system. Moreover, we could address secure communication schemes that rely on the characteristics of the mmWave and THz propagation, for example, to address eavesdropping, jamming, and other kinds of attacks. Finally, when it comes to THz applications, the light-matter interactions of THz signals also enable new imaging and spectroscopy applications, ranging from 3D scanning with sub-mm resolution to non-invasive chemical detection and characterization of products.

## 3 OUR VISION

We now describe our proposed *MillimeTera* infrastructure, which has been designed based on detailed feedback from the community on their urgent needs. The testbed has been designed to be flexible, extensible and responsive to the research challenges and experimental needs over the next 5 years and beyond. We describe the proposed SDR unit (and its phased evolution) in Section 3.1, and the full-stack and



**Figure 2: An early prototype of the 60 GHz 4-channel fully-digital MillimeTera SDR, with (1) a Xilinx KC705 FPGA board; (2) the transceiver board with various mixed-signal and RF chips from Texas Instruments and Analog Devices; and (3) a Linux host computer.**


control software in Section 3.2. *MillimeTera* aims at providing community-sourced designs and technologies that will be integrated into each SDR unit, going beyond what a single research group is capable of implementing. Furthermore, in Section 3.2 we describe how these SDR units will be deployed in a *geographically-distributed, flexible, and versatile testbed*.

### 3.1 SDR Units: Hardware

We will first design and develop novel 60 GHz SDR units, which will feature a *fully-digital* transceiver with 4 independent streams. Notice that [23] is only capable of analog beamforming. An early version is shown in Figure 2. Each SDR (Figure 2) will comprise the following modules:

**FPGA board for Baseband processing:** We will use the Xilinx KC705, which is a COTS board with a low-cost Kintex KC7K325T FPGA device. Featuring over 320k logic cells and 840 digital signal processing (DSP) blocks, this FPGA is ideal to perform the computations needed for a 4-channel system. The board connects 4 high-speed serial gigabit-transceiver (GT) ports to the FPGA mezzanine card (FMC) connector, enabling the FPGA to communicate with the digital-to-analog converters (DACs) and analog-to-digital converters (ADCs) on the transceiver mezzanine board. The board supports 1 Gbps and 10 Gbps Ethernet connections with the host computer, allowing for efficient integration with SDR software. The board also has i) 1 GB of DDR3 memory that can be used to buffer transmitted and received waveforms, ii) flash memory, iii) a USB programming interface, iv) a memory-card slot, and v) multiple clock sources.


**Transceiver mezzanine board:** This board will interface with the FPGA through an FMC connector. On the transmitter side, the time-domain digital I-Q data is sent from the FPGA to a bank of Texas Instruments DAC chips using JESD204B (subclass 1), which is a high-speed serial interface



**Figure 3: Design and simulation of lens for plug and play deployment at 60 GHz. (a) Simulation model used for semi-spherical lens; (b) Demonstration of beam-steering with single element switching.**

capable of multi-channel synchronization and 10 Gbps connectivity per lane. Each DAC chip has 4 converters, thereby supporting two complex baseband channels. The resulting analog signals are sent to a bank of Analog Devices HMC6300 mmWave up-converters, which produce the resulting 60 GHz RF signals. All the HMC6300 chips are provided with a reference local oscillator (LO) signal, so that the multiple channels are phase-locked. The resulting 60 GHz signals are fed to a bank of patch antennas, arranged in a  $4 \times 1$  or a  $2 \times 2$  configuration. The receiver chain is symmetrical to the transmitter chain, with mmWave down-converters and ADCs used instead of DACs and mmWave up-converters. A clocking chip is used for i) synchronization and DAC/ADC device locking, and ii) providing the core and GT reference clocks to the FPGA.

**THz and Lens-based Beamforming Evolutions:** A critical need exists in the academic research community to create the physical hardware that enables beamforming at THz ( $> 100$  GHz) frequency bands. Due to the relatively small separation between antennas at these frequency bands (on the order of millimeters), a transmit or receive beamforming device demands a monolithic integration of multiple antenna elements in a single chip. These beamforming chips require significant engineering effort and are highly proprietary solutions; they are hence extremely costly and unavailable to the broader academic community. Nonetheless, given the importance of this element when considering mmWave and THz directional communication, in a second phase of the testbed development, *MillimeTera* will provide THz digital beamforming modules to be integrated in the SDR unit. The target frequencies will be 120 and 240 GHz, due to Federal Communications Commission (FCC) regulations and the possibility of comparing bands with and without oxygen absorption. Additionally, the flexible and modular architecture of *MillimeTera* will allow the integration of alternative solutions, such as those described in [17].



**Figure 4: SDR unit and interactions between the FPGA board (which performs baseband processing) and the user host implementing the software protocol stack.**

An alternative beamforming solution is represented by *plug and play* pre-fabricated lenses, to improve system gain, link distance or beam-steering capabilities without affecting the power budget or adding active noise sources to the system. Full-wave simulation results for such a case are reported in Figure 3, which shows that a 25 dBi gain is possible with a lens of diameter  $2r \approx 4$  cm, dielectric constant  $\epsilon_r = 11$  and height  $H = 2.8$  cm. Beam-steering is also demonstrated using excitation of increasingly off-axis elements (1) to (5).

*Radio Frequency System-on-Chip (RF-SoC) for high performance THz baseband:* Besides the need for reconfigurable directional antenna systems, another key aspect of THz communications, and perhaps the main reason to justify its adoption, is the much larger consecutive bandwidth available beyond 100 GHz. Currently, the limitations in CMOS-based ADCs and DACs hamper its exploitation. For example, the fastest ADCs and DACs can support sampling frequencies in the order of 100 Giga-samples-per-second (Gsp), limiting the bandwidth to 50 GHz. Alternatively, multiple relatively narrowband (e.g., 2 GHz bandwidth) DACs and ADCs can be utilized to multiplex channels at THz frequencies. MilliMeTera will leverage the state of the art ADC/DAC arrays (up to 16 parallel converters) available in modern Xilinx RF-SoC chips to implement highly parallelized systems that can make the most of this bandwidth.

### 3.2 SDR Units: Software

The hardware platform presented in Section 3.1 needs to be complemented by the firmware, drivers, and software necessary to implement a wireless protocol stack and the necessary signal processing.

Specifically, our core SDR system will consist of (i) PHY layer processing in the FPGA of the board described in

Sec. 3.1; and (ii) a networking protocol stack that runs in a general-purpose processor of the host computer. This simple yet very effective design produces a flexible software-based radio, implementing widely different wireless protocols.

*Embedded software and Host Computer:* Figure 4 depicts a (simplified) end-to-end architecture of an SDR unit, with the bottom portion implementing the hardware-based logic, and the top portion describing the software-based protocol stack that will run on the host computer. Time-critical functionalities at the physical and MAC layer will be preferably executed in the FPGA, while pure software implementations may be run on the host. The implementation effort will be divided into (i) baseband processing and (ii) kernel-space Linux drivers, to interface the FPGA to the outside world. This strategy will allow us to keep the RF-related processing separated from the outside world, yet accessible through easy-to-use high-speed Linux drivers. We will make both a universal hardware driver (UHD) and a customized driver leveraging TCP sockets available to the community for ease of implementation and interfacing.

The high-level integration between hardware and software will be managed through the Platform for Advanced Wireless Research (PAWR) control framework (CF) software [5, 9, 21], which will handle the deployment of the relevant software blocks and experiment execution. The end-to-end system will thus rely on a robust, tested, and widely-used experimentation framework that can be run by both experts and new users. The testbed software blocks include a variety of customizable SDR stacks (such as GNU Radio, OpenAirInterface, and OpenLTE), core mobile networking stacks (such as M-CORD, and OpenAirInterface Core), as well as general purpose network virtualization and cloud computing stacks (such as OpenStack, OpenDaylight, ONAP, and XOS/CORD). As a result, the entire system (i.e., radio transceivers, data converters, baseband unit, computing and network infrastructure) will be software-defined and controlled through the CF, thus enabling a wide range of test cases, including both standard and custom, user-defined experiments. In addition, *the aforementioned software will be extended to provide novel functionalities that address mmWave/THz-specific needs and use cases*, for example beam management.

The fully-open software-based network stack, shown in the top part of Figure 4, will address existing development issues at all network layers and provide a common benchmark for researchers in mmWave and THz technologies. Specifically, this platform will help researchers design, prototype, and test solutions for a wide variety of problem areas, as discussed in Section 1. We will also include, among others: (a) efficient beam alignment; (b) beam tracking and recovery from blockage; (c) compressive channel sensing; (d) support

of multi-beam multi-user transmissions; (e) joint beam training and transmission scheduling; (f) simple imaging. The development of these functionalities, either in the host or in the FPGA board, will be extensively documented, so that researchers can easily understand how to modify the baselines to test alternative solutions.

*Geographically-distributed Local Testbed Model:* Unlike the centralized models of PAWR or the ORBIT testbed, the key defining feature of this effort is that a majority of the SDR nodes will be distributed to various universities across the country. There will be three mini-hubs for development of educational modules and fundamental research, each hosting a shared remotely-accessible testbed of about 5 nodes each. The remaining nodes will be given to academic research groups in various locations, based on an application and approval process. The advantages of this approach are: (a) researchers get to deploy the nodes in mobile/dynamic scenarios of their choice, as opposed to relying on static placement in a remote testbed; (b) researchers can choose to combine the SDR nodes with other hardware such as co-processors or GPU/CPU racks in their labs; (c) researchers can exploit the modularity of the MillimeTera design to integrate custom antenna arrays or lens-based solutions; and (d) SDR-based systems can be characterized and measured using anechoic chambers.

## 4 CONCLUSIONS

This paper discussed *MillimeTera*, a distributed testbed for end-to-end experimentation at mmWave and THz frequencies. We first presented today's research challenges, then reviewed the state-of-the-art testbeds available in the research community and highlighted the absence of a cost-effective, open and easy-to-configure testbed. We then described the main features of the *MillimeTera* testbed, which will be based on a large number of SDR nodes, with mmWave (first phase) and THz (second phase) front-ends and customized software. We hope that *MillimeTera* will foster innovative research in next-generation wireless networks.

## REFERENCES

- [1] O. Abari, H. Hassanieh, M. Rodreguez, and D. Katabi. 2016. Poster: A Millimeter Wave Software Defined Radio Platform with Phased Arrays. In *Proc. of ACM MobiCom*. New York City, New York, 419–420.
- [2] I. F. Akyildiz, J. M. Jornet, and C. Han. 2014. Terahertz band: Next frontier for wireless communications. *Physical Communication* 12 (Sep. 2014), 16–32.
- [3] F. Boccardi, R. W. Heath Jr, A. Lozano, T. L. Marzetta, and P. Popovski. 2014. Five Disruptive Technology Directions for 5G. *IEEE Commun. Mag.* 52, 2 (Feb. 2014), 74–80.
- [4] Cisco. 2017. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update. *White Paper* (March 2017).
- [5] COSMOS. 2019. PAWR COSMOS Platform. <http://cosmos-lab.org>
- [6] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi. 2019. A Tutorial on Beam Management for 3GPP NR at mmWave Frequencies. *IEEE Commun. Surveys Tuts.* 21, 1 (First quarter 2019), 173–196.
- [7] National Instruments. 2019. Introduction to the NI mmWave Transceiver System Hardware. White Paper, <http://www.ni.com/product-documentation/53095/en/>.
- [8] F. Khan and Z. Pi. 2011. An introduction to millimeter-wave mobile broadband systems. *IEEE Commun. Mag.* 49, 6 (June 2011), 101 – 107.
- [9] Univ. of Utah. 2019. POWDER Platform. <http://powderwireless.net>
- [10] 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.
- [11] T. S. Rappaport, Y. Xing, O. Kanhere, S. Ju, A. Madanayake, S. Mandal, A. Alkhateeb, and G. C. Trichopoulos. 2019. Wireless Communications and Applications Above 100 GHz: Opportunities and Challenges for 6G and Beyond. *IEEE Access* 7 (2019), 78729–78757.
- [12] S. Rey, J. M. Eckhardt, B. Peng, K. Guan, and T. Kürner. 2017. Channel sounding techniques for applications in THz communications: A first correlation based channel sounder for ultra-wideband dynamic channel measurements at 300 GHz. In *Proc. of 9th IEEE Int'l. Conf. on Ultra Modern Telecommunications and Control Systems (ICUMT)*. 449–453.
- [13] S. K. Saha, S. Aggarwal, D. Koutsikolas, and J. Widmer. 2018. AMuSe: An Agile Multipath TCP Scheduler for Dual-Band 802.11ad/ac Wireless LANs. In *Proc. of ACM MobiCom*. New Delhi, India, 705–707.
- [14] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetz, D. Koutsikolas, J. Widmer, and M. Hollick. 2018. Fast and Infuriating: Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware. In *Proc. of IEEE Int'l. Conf. on Sensing, Communication, and Networking (SECON)*.
- [15] S. K. Saha, Y. Ghasempour, M. K. Haider, T. Siddiqui, P. De Melo, N. Somanchi, L. Zakajsek, A. Singh, R. Shyamsunder, O. Torres, et al. 2019. X60: A programmable testbed for wideband 60 GHz WLANs with phased arrays. *Computer Communications* 133 (Jan. 2019), 77–88.
- [16] H. Shokri-Ghadikolaei, C. Fischione, G. Fodor, P. Popovski, and M. Zorzi. 2015. Millimeter Wave Cellular Networks: A MAC Layer Perspective. *IEEE Trans. Commun.* 63, 10 (Oct 2015), 3437–3458.
- [17] A. Simsek, S. Kim, and M. J. W. Rodwell. 2018. A 140 GHz MIMO Transceiver in 45 nm SOI CMOS. In *Proc. of IEEE BiCMOS and Compound Semiconductor Integrated Circuits and Technology Symp. (BCI-CTS)*. 231–234.
- [18] D. Steinmetz, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices. In *Proc. of ACM CoNEXT*. Incheon, Republic of Korea, 414–425.
- [19] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom*. Snowbird, Utah, USA, 28–41.
- [20] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim. 2018. Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks. In *Proc. of ACM MobiCom*. New Delhi, India, 257–271.
- [21] US-Ignite. 2019. Platforms for Advanced Wireless Research. <http://advancedwireless.org>
- [22] M. Xiao, S. Mumtaz, Y. Huang, L. Dai, Y. Li, M. Matthaiou, G. K. Karagiannidis, E. Bjornson, K. Yang, C. I, and A. Ghosh. 2017. Millimeter Wave Communications for Future Mobile Networks. *IEEE J. Sel. Areas Commun.* 35, 9 (Sep. 2017), 1909–1935.
- [23] J. Zhang, X. Zhang, P. Kulkarni, and P. Ramanathan. 2016. OpenMili: A 60 GHz Software Radio Platform with a Reconfigurable Phased-array Antenna. In *Proc. of ACM MobiCom*. New York City, New York, 162–175.
- [24] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. Panwar, and M. Zorzi. 2019. Will TCP Work in mmWave 5G Cellular Networks? *IEEE Commun. Mag.* 57, 1 (Jan. 2019), 65–71.