

Software-defined Radios to Accelerate mmWave Wireless Innovation

Kai Zheng^{1,10} Aditya Dhananjay^{1,10} Marco Mezzavilla^{1,10} Arjuna Madanayake²
 Shubhendu Bharadwaj² Viduneth Ariyaratna² Abhimanyu Gosain³ Tommaso Melodia³
 Francesco Restuccia³ Josep Jornet³ Michele Polese⁴ Michele Zorzi⁴ Jim Buckwalter⁵
 Mark Rodwell⁵ Soumyajit Mandal⁶ Xin Wang⁷ Jaakko Haarla⁸ Vasilii Semkin⁹
¹New York University ²Florida International University ³Northeastern University ⁴University of Padova
⁵UC Santa Barbara ⁶Case Western Reserve University ⁷SUNY Stony Brook ⁸Aalto University
⁹VTT Research Center ¹⁰Pi-Radio Inc.

1 Motivation and Problem

Imagine a dystopian world in which pianists did not have access to pianos. They would dream up sheet music, and have a computer simulate a recital. They would then go to their premier conference *PianoCom* to present a paper on their novel sheet music; and perhaps even spend the evening arguing about whose sheet music was more pleasing to the ears.

This is the same dystopian world that, as systems-level wireless researchers, we find ourselves in today; a world in which meaningful and real-world experimental work is out of reach for many of us. The reason is simple: *existing wireless prototyping platforms are either prohibitively expensive, or utilize transceiver technologies that are about 10 years behind the bleeding-edge; for many use-cases, such platforms do not even exist at any price.* This is driving up the cost and time needed to perform basic millimeter-wave (mmWave) wireless R&D, as well as its translation into emergent use-cases that include cellular, self-driving vehicles, radar, satellites, industrial automation, augmented/virtual reality, entertainment, and public-safety networks. These problems are most severe in the mmWave and sub-terahertz (sub-THz) bands, as explained below.

The severe *spectrum crunch* in the legacy (sub-6 GHz) radio bands has necessitated the use of the mmWave and sub-THz frequencies to fulfill the growing wireless data demands, simply due to the enormous amounts of underutilized spectrum in these emergent bands [1; 2; 3]. However, leveraging mmWave spectrum and getting *real-world* systems to work remains a challenging open problem because the radio frequency (RF) propagation behavior as well as the characteristics of wireless devices that operate at these frequencies tend to be different from today's wireless systems in the sub-6 GHz range [4]. The

free-space path losses in the mmWave bands are 1-2 orders of magnitude larger due to shorter wavelengths. This can be mitigated through beamforming gain which results in ray-like *beams*. The sharp beams are required to overcome path loss, but lead to problems such as blockages and sharp reflections. The steerability and compact size of the beamformers allow the combination of multi-input multi-output (MIMO) techniques with beamforming gain to enable operation in these scenarios [4]. Despite the weaker signal strength, reflections are critical to maintaining connections. The ability to rapidly steer multiple independent and sharp (i.e., high-gain) beams is critical for network robustness in such mmWave systems.

These fundamental differences have forced an end-to-end re-design of *every layer* of the wireless protocol stack, thus opening up a rich set of exciting research problems. Through over 50 *customer-discovery* interviews (using the I-Corps model), we have identified a rich set of scientific problems that the community desires to solve experimentally:

Physical layer (see [5]): What are the best waveforms to use, given the different device characteristics in mmWave bands? What beamforming and nulling approaches are needed? What modulation and coding schemes should be used to trade performance with computational complexity and power? How should interference be managed and mitigated in dense multi-user settings? What are the fundamental wireless propagation characteristics? How should channel-estimation be performed? What are the ramifications of compressed-sensing-based techniques, given sparse mmWave and sub-THz channels? How should beam-width adaptation and user-tracking work for mobile devices?

Media Access Control (MAC) layer (see [6]): In highly directional and dynamic links, how should

fundamental operations like initial access, beam-tracking, and synchronization be performed? How should the power-performance-robustness trade-offs be managed? In multi-user cellular settings, how can users be scheduled to provide quality of service (QoS) guarantees while maintaining fairness? How many beams are required, and how should these beams be managed? Can we leverage machine-learning to predict the evolution of the required beams in dynamic scenarios? How should transmission scheduling be combined with beam-training operations? Fundamentally, should MAC protocols be based on contention-based or grant-based mechanisms? How can anomaly detection techniques be used to detect and recover from transient link blockages?

Network and Transport layer (see [7]): How should end-to-end transport protocols be re-engineered to account for highly variable local links? How should network-layer protocols leverage multi-beam-connectivity? What *atomic* operations should be supported by the lower-layers of the protocol stack in order to enable efficient higher-layer operation?

Applications layer (see [8; 9]): Can we implement joint communication and sensing (i.e., radar) to enable use-cases in vehicular or unmanned aerial vehicles (UAV) networks? How should spectrum be managed while sensing? How can anomaly detection, matrix completion, and tensor completion techniques be used to improve the quality? How should applications (such as high-definition video streaming and gaming) adapt in the context of communication links with high bandwidth but high variability?

Devices and Materials (see [10]): Researchers who design RF integrated circuits (RFICs) currently lack a *vehicle* for their state-of-art chips to be used in a real-world testbed. They desire a modular platform that offers them a way to integrate their chips into a testbed that is used by the larger community. Furthermore, researchers working on smart materials (such as those used for passive lenses) also demand such a modular platform for integration.

Multi-hop Mesh Networks (see [11]): Having pencil-thin transmit and receive beams enables dense spatial reuse of spectral resources. How should resource allocation algorithms be designed? How should channel assignment and routing protocols be re-designed to leverage the spectrum? What are the trade-offs between performance and overall network robustness? How should interference mitigation be performed, while providing self-healing capabilities?

Security (see [12]): Given highly directional

mmWave links, how can we guard against eavesdropping, jamming, and other kinds of attacks? How can multi-connectivity be leveraged to accomplish the goals of improved privacy and security? When passive lenses are placed on top of a planar antenna array, how can the additional spatial filtering be leveraged to provide better security?

The aforementioned list is just a sampling of the open problems that wireless engineers desire to solve and demonstrate experimentally, but are held back due to the lack of affordable and bleeding-edge mmWave prototyping platforms. This pain-point is felt not only in academia, but also in government research laboratories and industry. In summary: wireless technologies are not developed and deployed with the speed and quality that could have been achieved if the community indeed had access to affordable and bleeding-edge wireless prototyping platforms. *We need to enable more experimentation!* [13]

2 The Solution

Hardware: We are working on a scalable and affordable software-defined radio (SDR) platform, featuring a state-of-the-art fully-digital mmWave wireless transceiver. The system will consist of three main parts (see Fig. 1): a) an off-the-shelf Xilinx ZCU111 RFSoc-based evaluation board; b) a host computer that runs Linux and GNU Radio; and c) the Pi-Radio transceiver board featuring all the RF components. The system will feature four channels operating in the unlicensed 57-64 GHz frequency band, with 2 GHz bandwidth. On the transmitter side, the PHY-layer digital baseband IQs can be generated at the host or on the FPGA. The RFSoc board has eight high-speed DACs that convert these IQs into the four analog baseband channels (each channel needs two DACs: one for I and one for Q), which are then routed through to the transceiver board using the FMC connector on the ZCU111 FPGA board. On the transceiver board, the analog streams corresponding to the four channels are fed to a bank of four Analog Devices HMC6300 mmWave up-converters (output power of 15 dBm per channel). The resulting RF signals are routed to the antennas, arranged in a 1×4 linear configuration. Multi-channel phase coherent mmWave up-conversion is performed through a circuit that consists of the Texas Instruments LMX2595 RF synthesizer, followed by an amplification and splitter network. The receiver side is symmetrical. We note that the same phase-coherent LO is used for the TX and RX sides, so as to facilitate RF self-calibration. The HMC6301 mmWave

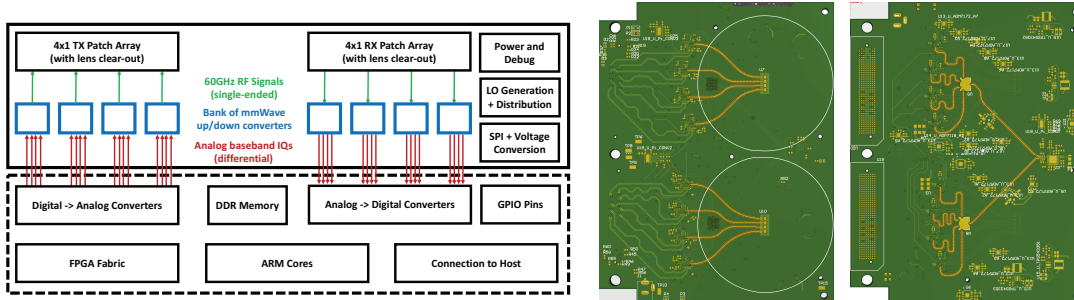


Figure 1: The block diagram of the SDR system under development, consisting of two main boards: a) the Xilinx ZCU111 evaluation board (dotted rectangle); and b) the Pi-Radio 60 GHz 4-channel transceiver board (solid rectangle); these two boards are mated using FMC connectors. Also shown are renderings of the transceiver board. The large circular keep-out areas and flange holes around the antennas are for the lens add-ons. The next board version could feature larger channel counts, as well as 2D beam-steering. Antenna design by Aalto University, Finland. All hardware design schematics and source code will be released using the *free and open-source* MIT license.

down-converter has a typical noise figure of 8 dB. The other components on the transceiver board include circuitry for the power tree and the SPI interface. We note that an early prototype of the 4-channel fully-digital transceiver board was developed at NYU (see Fig. 2). This system was based on a completely different baseband platform (Kintex), and the transceiver board had a different architecture that included DACs, ADCs, and clocking. The design of the NYU transceiver board and the LO generation prototype has been released publicly at [14].

We note two important points: a) the *fully-digital* transceiver represents the bleeding-edge in mmWave technology, previously unavailable to wireless researchers; and b) the system will be plainly affordable, as a matter of design and fundamental ethos.

Software: The FPGA can be programmed through Vivado and Simulink (System Generator). We use GNU Radio for programming the system on the host; this is the world’s most widely used SDR programming platform, and is supported by a large and active community of developers and users. While the FPGA is best suited for highly parallel and computationally intensive PHY-layer processing, the upper-layer (MAC and above) processing is better suited to be run on the ARM cores (on the RFSoc) or on the host computer. The split between the FPGA, ARM, and Host is completely configurable to allow for maximum flexibility and scalability.

Lens Add-on: The 1×4 antenna array provides only limited gain. To increase the gain, passive dielectric lenses will be placed on top of the antennas (see Fig. 3). There are two approaches to such lenses: a) a single lens placed over the array such that the




Figure 2: An early prototype of the NYU 60 GHz 4-channel fully-digital SDR system. The antennas have been designed at Aalto University, Finland. The SDR designs have been released using the MIT license.

elements are in the focal plane of the lens; and b) a *lenslet* array, wherein each antenna element has a lens element placed directly above it. Such lenses [15] can be made at $\leq \$10$ per unit, and can easily provide up to 25 dBi gain with a lens that is just a few centimeters in diameter. Flange holes on the transceiver printed-circuit board (PCB) can be used by the user to pop in and out their own lens systems.

GPU Acceleration: Graphics processing units (GPUs) are useful for highly parallel and numerically intensive processing tasks like coding/decoding. Besides being very powerful, GPUs offer a flexible and widely-used software stack to program them (for example, CUDA by NVIDIA). CUDA gives direct access to the GPU’s instruction set for the execution of compute kernels. Community feedback indicates that GPUs might be preferable to FPGAs for the higher-layer processing stack due to the fact that C/C++ can be used instead of register-transfer level (RTL) languages like Verilog/VHDL. This would also en-




Figure 3: Left: Single lens placed with the antenna elements in the focal plane; Right: Lenslet system with each antenna element having its own lens.

courage testbed use among the computer science (CS) community, and allow for more rapid translation from ideas to *over-the-air* implementations.

Migration to Upper mmWave Frequencies: Research groups around the world tape out impressive mmWave and sub-THz RFICs; such efforts involve significant effort and cost. However, there lacks a *translation* effort wherein these RFICs are integrated into SDR systems, and then used by the larger community. We plan to leverage these community-developed RFIC designs, refine the designs, and use these chips in our SDR platform. Critically, such efforts will involve packaging the bare die and integrating antennas onto the package. Such an effort will ensure that the community has access to bleeding-edge RFICs, leap-frogging even what the industry currently has access to. The FCC has opened up several GHz of unlicensed bandwidth in the 120 GHz and 240 GHz bands; these are therefore the most likely candidate frequencies for our effort.

Deployment Model (contingent on funding): We envision a geographically distributed hub-and-spoke model for deploying these SDR nodes in the real world. We plan to leverage existing publicly accessible testbeds (such as the ORBIT testbed at Rutgers University) to deploy these SDR nodes, such that users can log in remotely and run experiments. This will ensure a maintained and updated testbed with containerized software. However, users might want to run experiments in their own labs, with custom topologies and dynamic blockage and mobility scenarios. We therefore plan to make available several of these SDR nodes to researchers for use in their own labs, provided they agree to open-source the resulting work products. This will help with minimizing the replication of engineering effort across groups, but also help with a long-standing goal in the community: *being able to reproduce the results of others.*

3 Conclusion

We firmly believe that wireless research, especially in the mmWave (and beyond) frequencies, needs to have a strong experimental component for it to be meaningful and impactful. The community has been desperately searching for such platforms, and it is our hope to be able to provide such bleeding-edge and affordable SDR systems to the community.

References

- [1] F. Khan and Z. Pi, "An introduction to millimeter-wave mobile broadband systems," *IEEE Commun. Mag.*, vol. 49, no. 6, pp. 101–107, Jun. 2011.
- [2] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349.
- [3] 3GPP, "Study on scenarios and requirements for next generation access technologies," TR 38.913, 2016.
- [4] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, 2014.
- [5] Theodore Rappaport and Robert Heath and Robert Daniels and James Murdock, *Millimeter Wave Wireless Communications (Chapter 7)*. Prentice Hall, 2014.
- [6] H. Shokri-Ghadikolaei, C. Fischione, G. Fodor, P. Popovski, and M. Zorzi, "Millimeter Wave Cellular Networks: A MAC Layer Perspective," *IEEE Transactions on Communications*, vol. 63, no. 10, pp. 3437–3458, Oct 2015.
- [7] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. S. Panwar, and M. Zorzi, "Will TCP work in mmWave 5G Cellular Networks?" *CoRR*, vol. abs/1806.05783, 2018. [Online]. Available: <http://arxiv.org/abs/1806.05783>
- [8] A. Arbabian, S. Callender, S. Kang, M. Rangwala, and A. M. Niknejad, "A 94 GHz mm-Wave-to-Baseband Pulsed-Radar Transceiver with Applications in Imaging and Gesture Recognition," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 4, pp. 1055–1071, April 2013.
- [9] H. Singh, J. Oh, C. Kweon, X. Qin, H.-R. Shao, and C. Ngo, "A 60 GHz wireless network for enabling uncompressed video communication," *Articles*, p. 8, 2008.
- [10] J. F. Buckwalter, A. Babakhani, A. Komijani, and A. Hajimiri, "An Integrated Subharmonic Coupled-Oscillator Scheme for a 60-GHz Phased-Array Transmitter," *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 12, pp. 4271–4280, Dec 2006.
- [11] S. Singh, R. Mudumbai, and U. Madhow, "Interference analysis for highly directional 60-GHz mesh networks: The case for rethinking medium access control," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 5, pp. 1513–1527.
- [12] Y. Wu, A. Khisti, C. Xiao, G. Caire, K.-K. Wong, and X. Gao, "A survey of physical layer security techniques for 5G wireless networks and challenges ahead," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 4, pp. 679–695, 2018.
- [13] M. Polese, F. Restuccia, A. Gosain, J. Jornet, S. Bhardwaj, V. Ariyaratna, S. Mandal, K. Zheng, A. Dhananjay, M. Mezzavilla, J. Buckwalter, M. Rodwell, X. Wang, M. Zorzi, A. Madanayake, and T. Melodia, "MillimeTera: Toward A Large-Scale Open-Source mmWave and Terahertz Experimental Testbed," in *Proceedings of the 3rd ACM Workshop on Millimeter-wave Networks and Sensing Systems*, ser. mm-Nets'19. ACM, 2019, pp. 27–32.
- [14] [Online]. Available: <https://github.com/adityadhananjay/fully-digital>
- [15] J. Ala-Laurinaho, J. Aurinsalo, A. Karttunen, M. Kaunisto, A. Lamminen, J. Nurmiharju, A. V. Räsänen, J. Säily, and P. Wainio, "2-D beam-steerable integrated lens antenna system for 5G E-band access and backhaul," *IEEE Transactions on Microwave Theory and Techniques*, vol. 64, no. 7, pp. 2244–2255, 2016.