# RENEW: Programmable and Observable Massive MIMO Networks

Rahman Doost-Mohammady*, Oscar Bejarano*, Lin Zhong*, Joseph R. Cavallaro*,
Edward Knightly*, Z. Morley Mao†, Wei Wayne Li‡, Xuemin Chen‡, Ashutosh Sabharwal*
*Rice University, United States
{doost, obejarano, lzhong, cavallar, knightly, ashu}@rice.edu
†University of Michigan, United States
zmao@umich.edu
‡Texas Southern University, United States
{LiWW, Xuemin.Chen}@tsu.edu

***Abstract*—Massive MIMO is one of the key technologies in 5G wireless broadband, capable of delivering substantial improvements in capacity of next-generation wireless networks. However, due to its inherent complexity, its operation, reconfiguration, and enhancement present significant challenges and risks. In this paper we present RENEW, a fully programmable and observable massive MIMO network. We present the architectural design for full programmability at every layer of the wireless stack, from the radio hardware, including PHY and MAC layer configurations, all the way up to the network core functionality using network function virtualization. We also present mechanisms to enable observability at every layer of the stack. These include various indicators in the radio and core access network, hence enabling effective monitoring, troubleshooting, and performance evaluation of the network at large.**

***Index Terms*—Massive MIMO, 5G, open-source, software defined radio**

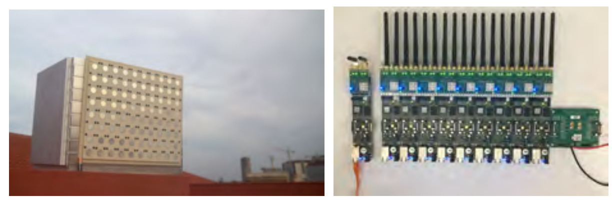

*Fig. 1:* (Left) *Third-generation Argos base station on rooftop as part of ArgosNet/RENEW along with an AA battery (in front of the base-station) to illustrate its size. The current design (power, antennas, networking) can scale to 160 antennas.* (Right) *Inside the base station are daisy chains of Iris software-defined radio modules; a nine-module chain shown here.*

## I. Introduction

The wireless industry rightfully boasts amazing statistics: a $1+ trillion market with nearly 7 billion devices, and each 10 MHz spectrum adds $3 billion to GDP & over 200,000 jobs in the workforce [1]. However, all predictions of *future* wireless usage make these numbers look insignificant, as we aspire to wirelessly connect nearly everything to the Internet. With augmented and virtual reality (AR/VR), Internet of Things (IoT), manufacturing, connected cars, drones and yet to-be-invented applications, the number of wireless-empowered applications is about to explode. In short, the wireless industry will soon have to become *everything to everyone*. For each innovative application domain, wireless connectivity will have to achieve a previously impossible combinations of bandwidth, reliability, power efficiency, complexity and coverage requirements. Unfortunately, today researchers have neither the foundational principles nor the experimental tools to systematically explore the large design space to realize this vision of future wireless.

The wireless networks that we have today are clearly not the wireless networks that we need for the future. Significant advancements cannot be attained by tinkering around the edges of existing networks. To drive forward the science of wireless networking, we need innovative researchers to build their own networks on a large scale and in real environments, with control and visibility from the lowest layers of the radio up to the top of the application stack. Since wireless devices are diverse and mobile, the testbed must be too; since technologies change rapidly (and sometimes unpredictably) at all layers of the stack, the platform must likewise be able to adapt to community needs to stay relevant. This living laboratory needs to be built with the precision of a scientific instrument so that experimenters can have confidence in the accuracy and reproducibility of their results, and must be built from the ground up to support the scientific process. It must support not only competition in the race for cutting edge technologies, but also cooperation and collaboration that enables researchers and industrial users to build on each other's work.

To address the above challenge, in 2018 we have begun the design of Reconfigurable Eco-system for Next-generation End-to-end Wireless (RENEW) – a large-scale massive MIMO scientific instrument and eco-system that will enable research at the forefront of the wireless revolution for years to come. RENEW is designed to enable a broad segment of the research community to address fundamental research questions in wireless communications and wireless-and-mobile networking in order to have a broad impact in accelerating wireless technology. In addition, RENEW will enable research on a

RENEW is a technology partner to POWDER. RENEW/POWDER is part of the National Science Foundation's PAWR program and is funded in part by NSF award CNS-1827940 and the PAWR Industry Consortium.

broad range of related applications and services that depend on wireless networks. RENEW leverages Rice Argos [2] to provide a fully open-source massive MIMO framework. Each many-antenna Argos base station employs multiple custom programmable Software Defined Radio (SDR) modules called Iris (from the company Skylark Wireless [3]), as their building blocks. Similarly, clients are based on Iris modules.

RENEW will be deployed as part of the POWDER project at the University of Utah campus. The network will largely consist of outdoor, high-powered transmitters as well as client nodes with high mobility (installed on shuttle buses) and a mix of indoor and outdoor fixed nodes deployed across the university campus [4].

## II. RENEW Technical Innovations

The key technical innovations of RENEW are:

**End-to-end Programmable**: Leveraging Rice Argos [2], RENEW will support programmability at all network layers and end-to-end, including programmable diverse spectrum access from 50 MHz to 3.8 GHz. Additionally, RENEW will have three tiers of computing resources adequate for cutting-edge research into software realization of network functions including the physical layer: (*i*) a hybrid edge computer, at each base-station, composed of FPGA and GPU/CPU-based processing, and (*ii*) small data center at the deployment site, with very low-latency communication across multiple base-stations.

**End-to-end Observable**: RENEW will provide end-to-end programmability along side end-to-end observability, a key ingredient for research advancement and reproducibility. RENEW observability will encompass infrastructure and programmable clients, resulting in extensive and unique data-sets, spanning many layers in the network, including apps.

**Infrastructure as Distributed System**: A testbed like RENEW is a heterogeneous distributed system for which application development is notoriously difficult. To address this problem, we will combine distributed systems technologies and the special properties of wireless infrastructure to build a RENEW resource manager. Akin to a distributed operating system, the resource manager presents each experiment a virtual testbed and coordinates multiple virtual testbeds. By leveraging the event-driven and data-driven nature of wireless network infrastructure, the resource manager further allows experimenters to record behavior traces, at various points in the testbed and across the network stack, and replay them to understand their designs under investigation.

**Ahead-of-standards Open-source Network Stacks**: RENEW will provide an open-access growing research repository with fully operational network stacks aiming to stay a generation or more *ahead* of wireless standards, in addition to supporting current standards (LTE and Wi-Fi) for baselines. The initial research repository will support recent advances such as Massive MU-MIMO, full-duplex, coordinated multi-point, diverse-spectrum PHY and access mechanisms, and software-defined networking (SDN) and network function virtualization (NFV). The repository will enable the community to research large system-level innovations against existing concepts throughout the network stack.

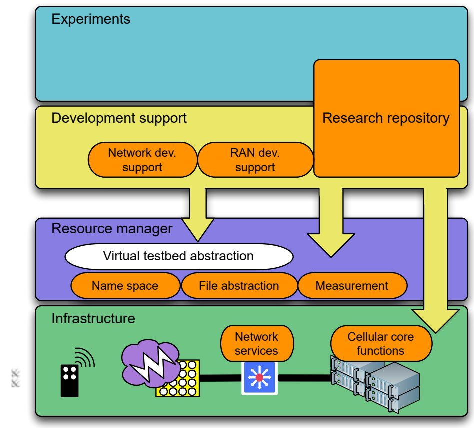


*Fig. 2: RENEW architecture layers.*

## III. Platform Technical Architecture Design

The RENEW end-to-end programmable architecture consists of three layers: infrastructure, resource manager, and development support for experimenters as summarized by Figure 2.

The *infrastructure*, described in §III-A, constitutes all the resources, both hardware and software, necessary for assembly of network functions at all layers. The *resource manager*, akin to a distributed operating system for the infrastructure, provides a *virtual testbed* abstraction that guarantees availability of specified resources to experimenters, creating an illusion to them that they are the sole users of these resources. Under the hood, the virtual testbed relies on a hierarchical name space and a virtual filesystem to address and virtualize infrastructure resources. The *development support* provides interfaces and libraries for users from diverse research communities to develop and debug experiments that may access infrastructure resources at different layers of the infrastructure. Included in the development support will be an open-access research repository of experiment implementations and building blocks to seed the research and education use for each research community. The collaborative development of the research repository will both seed and accelerate the use of the testbed by the broader research community.

### A. Infrastructure

Below we will describe the infrastructure for the RENEW platform.

*1) Programmable Radio Access:* **Iris SDR**: RENEW employs a programmable SDR module, called Iris from Skylark Wireless [3], as the building block for its many-antenna base stations and clients. Iris was developed for ArgosNet and has been used for both its base stations and clients. Iris combines an LMS7002M SDR transceiver [5] with the Xilinx 7z030 system-on-a-chip (SoC). LMS7002M is a frequency-agile 2×2 transceiver that can be tuned from 50 MHz to 3.8 GHz,

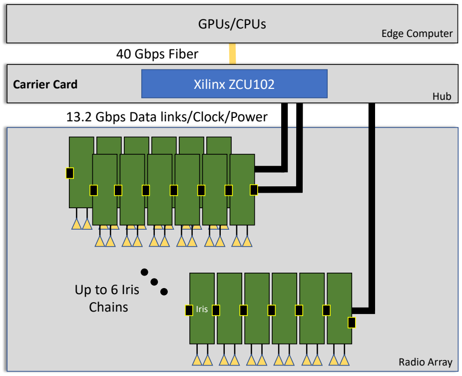


*Fig. 3: Architectural view of RENEW base station which is an upgrade of Argos base station.*

sampling a maximum of 56 MHz of contiguous bandwidth on each radio interface with 12-bit ADC. Xilinx 7z030 SoC includes dual-core ARM Cortex-A9 and medium-sized FPGA fabric. On Iris, the 7z030 SoC communicates with LMS7002M through high-speed I/Q lines from the FPGA fabric for data and through the SPI bus for control.

Iris carries an interchangeable RF frontend to support high RF performance at various bands. There are three RF frontends commercially available for Iris for use in RENEW. The first has no filtering or amplification but supports the full frequency range, ie. 50 MHz to 3.8 GHz, by feeding LMS transceiver output to antennas with 0 dBm transmit power. The other two frontends have filters & amplifiers and therefore are band-specific: one supports 470-700 MHz UHF with 30 dBm transmit power and the other supports both 2.4-2.7 GHz and 3.5-3.8 GHz with 28 dBm transmit power, respectively. Notably, each Iris has a width of 1.55" with two TX/RX antenna ports at the front-end. Such form factor allows half-wavelength separation at 3.8GHz, provided that dual-polarization antennas are used in the system.

**Base station**: The RENEW base station is based on Argos base station (shown in Figure 1) but with important updates. It consists of daisy chains of Iris modules as depicted by Figure 3, using two specialized bus connectors on each side of Iris modules. The bus connectors carry reference clock and power as well as 4 GPIO lines and 13.2 Gbps serial link routed to the FPGA fabric, thereby eliminating the need for cabling between Iris modules while locking all daisy-chained Iris modules in frequency and enabling phase-coherent MIMO operation. A *Hub board* like the Xilinx ZCU102 development board combines multiple daisy chains coherently through a customized backplane that provides high-precision clocking, power, and bi-directional connectivity. The Hub also aggregates/distributes streams of samples from/to all radios.

**Network clients**: Iris-based clients will also be deployed throughout the infrastructure. Because 7z030 SoC, the heart of Iris, has dual ARM Cortex A9 cores, the client is capable of running a full stack Linux to drive the lower network layers implemented in the FPGA fabrics. We will develop a traffic generator on top of the Linux stack. The generator behaves like a mobile app and sends/receives packets through the Iris radio front-end. A research user can control this generator in two ways. First, she can offline download a database to the client for the generator to use. More importantly, she can use the Ethernet or Wi-Fi connectivity to the client to drive its traffic generator in real-time, e.g., routing the traffic from a real smartphone to the client.

*2) Low-latency Network Infrastructure:* RENEW/POWDER will have a high-speed, ultra-low latency network connecting the multiple massive MIMO base stations relying on redundant 40 Gbps circuits which will be deployed over existing dark fiber resources. The network will provide commonly needed capabilities such as DNS, DHCP, SMTP, authentication, and logging services, and portions of them will be available as resources for experiments.

*3) Distributed Programmable Computing Resources:* The RENEW/POWDER platform will have programmable computing resources distributed across the infrastructure that are adequate for real-time implementation of physical layer innovations, especially software realization of radio-access and network functions. In addition, we will provision abundant and heterogeneous processing capability at a data center in the deployment site.

*4) Functions for Cellular Core:* On top of the flexible network infrastructure and data centers described above, RENEW/POWDER will support a key set of cellular network functions for traffic between the clients and the Internet. Instead of using expensive, non-reconfigurable specialized equipment like in today's LTE networks, RENEW/POWDER will realize these functions in software, leveraging SDN and NFV, for greater flexibility, efficiency, and scalability.

### B. Resource Manager

On top of the physical infrastructure described above, RENEW/POWDER employs a resource manager to manage and monitor its resources, provide arbitrated access to them, and enforce isolation between different users. The resource manager does so with an abstraction called *virtual testbed*. A virtual testbed consists of a set of resources defined and guaranteed by the manager. Each experiment runs inside its dedicated virtual testbed and is isolated from other concurrently running experiments. The virtual testbed abstraction leverages many resources' innate support for concurrent users and capacity provisioning, e.g., servers, fibers and network switches. For these resources, the resource manager admits users based on their capacity to ensure that concurrent running experiments do not interfere with each other. On the other hand, resources like antennas, FPGAs, and even the wireless channels, are difficult, if possible at all, to share by concurrent users. For these resources, the resource manager will allow time-division multiple access by default and cannot allocate them to two concurrently running virtual testbeds. The virtual testbed abstraction will build on top a hierarchical name space and virtual filesystem.

### C. Development Support

Given RENEW's complexity, we will provide a core set of frameworks to facilitate experiment development. **PHY and MAC experiment development**: RENEW will support both software-based and hardware (FPGA)-based development flows, allowing I/Q samples to be processed by various computing resources, from FPGA within the Iris module module to the server colocated with the base station and to the data center. For both flows, we will provide open-source examples. For real-time experiments that require the computational power of the FPGA inside the Iris module, RENEW supports the Xilinx FPGA development flow and provides scripts to automate much of it. Because RENEW abstracts each FPGA as a virtual file, configuring the FPGA is as easy as writing into the corresponding file.

For software-based development flow, RENEW abstracts Iris modules with SoapySDR [6], an open-source SDR abstraction framework. SoapySDR integrates multiple available SDR drivers in the market, such as UHD by Ettus Research. As a result, experiments implemented for the popular GNURadio framework can be readily run on RENEW. Additionally, RENEW supports development flows using multiple programming languages, including MATLAB, C, and Python.

**Network experiment development**: By supporting OpenAirInterface and OpenEPC, RENEW immediately supports their associated development frameworks. The orchestration layer on POWDER fully manages all of the resources, including compute and radio nodes. The experiments will be fully describable (and therefore repeatable) using JSON configuration files.

**Open-access research repository**: We will develop a large open-access research repository spanning all layers in the network stacks. While many of the simpler blocks are available as part of open-source development frameworks such as GNU radio, there are many advanced wireless algorithms that are not yet available. The repository will include:

(*i*) *Physical layer blocks* implemented in both hardware (Verilog/VHDL) and software (C/C++ and Python). These include centralized and decentralized algorithms for downlink precoding/beamforming and uplink equalization and detection including maximum-ratio combining, zero-forcing, and minimum mean-square error methods for massive MIMO. Also channel coding and decoding including turbo, LDPC, and Polar code families will be provided. Additionally, we will release blocks for modulation waveforms such as OFDM and filter bank multi-carrier waveforms for 5G research.

(*ii*) *Standard and experimental stacks* that include IEEE 802.11g/n/ac FPGA implementation, LTE Release 10 based on OpenAirInterface software framework with experimental modules such as massive MIMO with many client support. For massive MIMO research, we will provide a custom experimental stack developed in both hardware and software [7] with support for channel sounding, calibration, control channels, beamforming, as well as massive-MIMO full-duplex using the recent scheme SoftNull [8] which scales gracefully to massive MIMO systems. For cloud radio-access network (RAN) and coordinated multi-point (CoMP) research, we will provide C- and FPGA- based implementations of multi-cell time and frequency synchronization, automatic gain controls, and joint scheduling and beamforming.

(*iii*) To demonstrate the flexibility of RENEW's network and cellular core support, we will provide a *low-latency cellular core realization* [9] that uses a dedicated container to host necessary network functions for each client, in contrast to today's LTE evolved packet core (EPC) architecture where clients share instances of network functions. By locating the client container, and required network functions, in the edge computer close to the client, this design can significantly reduce cellular core latency: in today's LTE EPC, the client traffic has to be routed through shared network function instances that are often many hops removed from the client.

## IV. Supported Research Experiments

RENEW/POWDER will support core wireless research communities while also creating an environment that welcomes and supports new emerging and synergistic areas.

**Wireless Architectures**: With each generation, it has become increasingly harder to improve the *spectral efficiency* of wireless networks. It is not surprising then to see the pattern of innovations in the last decade that have been challenging traditional paradigms, e.g. massive MIMO [10]–[12], full-duplex [13], [14], Polar codes [15], and Coordinated Multi-point [16], [17]. Each of these ideas was so radical at their introduction that pure theory was not considered acceptable and experimental validations were considered crucial. Extrapolating forward, we expect that the next round of ideas will further challenge conventional wisdom, and hence will need large-scale testbed instantiations. RENEW/POWDER's computing and fiber networking architecture will achieve near-best-possible latencies across the whole network with massive MIMO, full-duplex and diverse spectrum support, and thereby allow exploration of designs that can explore the tradeoff, e.g. between throughput, latency, antennas per base-station, bandwidth, network core design and geographical coverage.

**Distributed and Network Systems**: RENEW/POWDER will open the network and computing infrastructures that are traditionally under tight control by carriers, with an expected impact similar to what software-defined network (SDN) technology has had on the Internet and data centers. Existing cellular networks consist of two distinct subsystems before interfacing with the Internet: the RAN and the evolved packet core (EPC) network. While RAN testbeds like WARP [18], NI and USRP [19], [20], BigStation [21] and ArgosNet [2] have allowed cutting-edge research about the last wireless hop, the researchers do not have access to a combined open-access RAN and EPC network. RENEW/POWDER will provide the first such end-to-end programmable wireless network testbed. It will leverage general-purpose systems based on an open-source software stack, and off-the-shelf connectivity, i.e., Ethernet over fiber.

**Mobile & Ubiquitous Computing**: The high throughput, low latency supported by RENEW/POWDER and the distributed computing resources in the testbed provide an opportunity to investigate technical challenges facing novel applications that serve mobile users. For instance, AV/VR applications are well-known for the demanding requirements in computational power and latency [22]. As a result, there is significant research on off-loading computation to edge computers or cloudlets that rely on high-speed, low-latency links [23]. Several emerging applications of AR/VR such as education, training of medical and military personnel, and entertainment will only be unlocked when such links are wireless [24]. Our platform will provide the infrastructure to investigate novel methods in multiple frontiers, from programming models to data privacy and system security.

**Security and Privacy**: Security in current and emerging mobile/wireless applications, such as Autonomous Vehicles or self-driving cars or IoT-enabled cyber-manufacturing, will be critical to operational safety [25]. While security has been explored at each of the individual layers, there has not been a systematic investigation into multi-layer network security – RENEW will provide a testbed to experimentally evaluate cross-layer security architectures, e.g., what security primitives should be built into the lower network layers, and across networked systems, from mobile devices to base station to gateway and eventually to cloud services, in order to enable strong security guarantees.

**Machine Learning/Artificial Intelligence**: Machine Learning is expected to adapt the system dynamically to its most efficient state, and to ensure and monitor its security and safety. Inferring and detecting anomalies over networks is an active area of research using FPGA and GPU computing systems [26]. Yet existing machine learning technologies are incapable of dealing with the unprecedented volume and rate of data required to protect future wireless communications. RENEW will provide an ideal testbed with real-time high-volume and high-speed wireless data; it will be invaluable for deploying and testing next generation machine learning algorithms for ultra-high speed data mining research.

## V. Conclusion

In this paper we introduce the open-source massive MIMO RENEW platform for enabling wireless research. The platform includes user-programmable software defined radios, plus extended functionality such as specialized antennas, radios, and hubs. It will also have modest local compute and storage capabilities (i.e., edge compute with sub-ms latency), and the ability to access large amounts of cloud computing capacity. On top of this infrastructure, we are planning to run a sophisticated testbed control framework that builds in support for complex device provisioning and a set of tools for scientific workflow management, collaboration, and artifact sharing. This platform is part of the RENEW/POWDER partnership created to design, build, and operate a new city-scale living laboratory for accelerating wireless technology innovation.

## References


[1] Statistics from CTIA reports. [Online]. Available: http://www.ctia.org

[2] Argos. [Online]. Available: http://argos.rice.edu

[3] Skylark Wireless. [Online]. Available: http://www.skylarkwireless.com/

[4] POWDER. [Online]. Available: https://powderwireless.net/

[5] Lime Microsystems. [Online]. Available: http://www.limemicro.com/products/field-programmable-rf-ics-lms7002m

[6] SoapySDR Project. [Online]. Available: https://github.com/pothosware/SoapySDR/wiki

[7] C. Shepard, A. Javed, and L. Zhong, "Control channel design for many-antenna MU-MIMO," in *Proc. ACM Int. Conf. on Mobile Computing & Networking (MobiCom)*, 2015, pp. 578–591.

[8] E. Everett, C. Shepard, L. Zhong, and A. Sabharwal, "SoftNull: Many-Antenna Full-Duplex Wireless via Digital Beamforming," *IEEE Transactions on Wireless Communications*, vol. 15, no. 12, pp. 8077–8092, December 2016.

[9] M. Moradi, Y. Lin, Z. M. Mao, S. Sen, and O. Spatscheck, "SoftBox: A customizable, low-latency, and scalable 5G core network architecture," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 3, pp. 438–456, 2018.

[10] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 186–195, February 2014.

[11] E. Bjornson, E. G. Larsson, and T. L. Marzetta, "Massive MIMO: ten myths and one critical question," *IEEE Communications Magazine*, vol. 54, no. 2, pp. 114–123, February 2016.

[12] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: practical many-antenna base stations," in *Proc. ACM Int. Conf. on Mobile Computing & Networking (MobiCom)*, 2012, pp. 53–64.

[13] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," in *Proc. ACM Int. Conf. on Mobile Computing & Networking (MobiCom)*, 2010, pp. 1–12.

[14] M. Duarte and A. Sabharwal, "Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results," in *Asilomar Conference on Signals, Systems and Computers*, 2010, pp. 1558–1562.

[15] J. Lin, Z. Yan, and Z. Wang, "Efficient Soft Cancelation Decoder Architectures for Polar Codes," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 1, pp. 87–99, Jan 2017.

[16] P. Marsch and G. Fettweis, *Coordinated Multi-Point in Mobile Communications: From Theory to Practice*. Cambridge, 2011.

[17] M. Sawahashi, Y. Kishiyama, A. Morimoto, D. Nishikawa, and M. Tanno, "Coordinated multipoint transmission/reception techniques for LTE-advanced," *IEEE Wireless Communications*, vol. 17, no. 3, pp. 26–34, June 2010.

[18] Wireless open-Access Research Platform. [Online]. Available: http://warp.rice.edu

[19] USRP. [Online]. Available: https://www.ettus.com/

[20] E. Luther, *National Instruments 5G Massive MIMO Testbed: From Theory to Reality*, September 2015, http://www.ni.com/white-paper/52382/en/.

[21] Q. Yang, X. Li, H. Yao, J. Fang, K. Tan, W. Hu, J. Zhang, and Y. Zhang, "BigStation: Enabling Scalable Real-Time Signal Processing in Large MU-MIMO Systems," in *Proc. ACM SIGCOMM*, 2013, pp. 399–410.

[22] P. Jain, J. Manweiler, and R. R. Choudhury, "OverLay: Practical Mobile Augmented Reality," in *Proc. ACM Int. Conf. on Mobile Systems, Applications, and Services (MobiSys)*, 2015, pp. 331–344.

[23] M. Gordon, D. K. Hong, P. Chen, J. Flinn, S. Mahlke, and Z. M. Mao, "Tango: Accelerating Mobile Applications through Flip-Flop Replication," in *Proc. ACM Int. Conf. on Mobile Systems, Applications, and Services (MobiSys)*, 2015, pp. 10–13.

[24] T. Chen, Z. Chen, Q. Shi, and X. Huang, "Road marking detection and classification using machine learning algorithms," in *2015 IEEE Intelligent Vehicles Symposium (IV)*, June 2015, pp. 617–621.

[25] A. Vosoughi, J. Cavallaro, and A. Marshall, "A Context-aware Trust Framework for Resilient Distributed Cooperative Spectrum Sensing in Dynamic Settings," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 10, pp. 9177–9191, Oct. 2017.

[26] D. Luebke, "CUDA: Scalable parallel programming for high-performance scientific computing," in *2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, May 2008, pp. 836–838.