Aerial Experimentation and Research Platform for Mobile Communications and Computing

Vuk Marojevic[1], Ismail Guvenc[2], Rudra Dutta[3], and Mihail Sichitiu[2]
[1]Department of Electrical and Computer Engineering, Mississippi State University, Mississippi State, MS
[2]Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC
[3]Department of Computer Science, North Carolina State University, Raleigh, NC
vuk.marojevic@msstate.edu, {iguvenc, rdutta, mlsichit}@ncsu.edu

Abstract—**The Third Generation Partnership Project (3GPP) released its first 5G specifications in December 2017. 3GPP Release 15 defines a new radio access network, new security architecture, services and applications, among others. These specifications will be refined with new releases that will include enhancements to cellular vehicular to everything (C-V2X), first introduced in LTE in Release 14, as well as the integration of unmanned aerial systems (UASs). This paper proposes a city-scale R&D platform enabling advanced mobile communications and computing experiments with 5G technology and UASs. The large-scale testbed is proposed by an academic-industry consortium to serve the needs of global researchers, developers, practitioners, regulators, and educators, among others. We present the unique aspects of this platform, its architecture and user services, for enabling experiments in controlled, yet production-like environments. Sample research projects, including 5G security and distributed wireless and computing resource management, indicate the broad R&D scope and the impact that such a platform can have on technology evolution, regulation, standardization and new computing and communications services.**

Keywords–5G, UAS, open research platform, testbed architecture.

I. INTRODUCTION

The Third Generation Partnership Project (3GPP) published its fifteenth release of mobile communication system specifications in December 2017, setting the foundations for 5G. With groundbreaking upgrades at the radio layer, the New Radio (NR) standard implements an advanced physical layer that supports millimeter wave communications and massive antenna arrays for full-dimensional multiple-input, multiple-output (MIMO) antenna systems. The 5G core network (5GC) has been redesigned for enhanced flexibility and service versatility with respect to the evolved packet core (EPC) developed for the long-term evolution (LTE).

The 5G RAN will be supported by a 5G service-based architecture (SBA), which defines a set of network features and mechanisms for network element registration, discovery and authorization, among others. It allows 5GC functions to be implemented as virtual network functions, and their dynamic integration and management for building and operating custom 5G networks on general purpose computing and networking hardware.

The goal of 5G networks is to provide ubiquitous, high-speed, low-latency connectivity for enhanced mobile broadband, massive machine type communication and real-time control in highly mobile environments. 5G will enable the tactile Internet, untethered augmented and virtual reality, and smart and connected vehicles, among others [1].

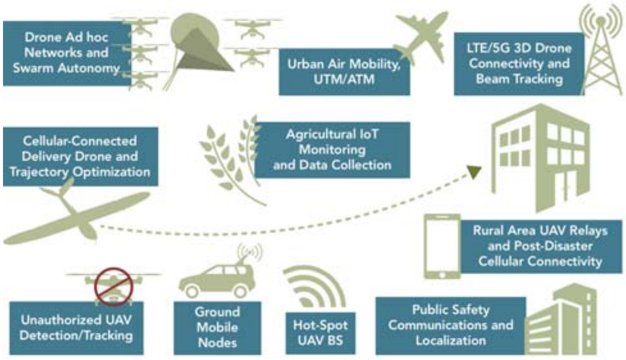

Fig. 1. 5G aerial networks and example AERPAW use cases and applications.

5G communications are largely driven by the need to be "untethered", and mobility is the root driver of many of the most difficult intellectual challenges in wireless theory and technology, from coarse-grained concerns such as handover to fine-grained ones such as beam-steering for tracking mobile nodes. Unmanned aerial systems (UASs) will leverage 5G services as well as extend the scope of terrestrial 5G networks (Fig. 1). These systems can be of a variety of types and carry different payloads to offer user services or support commercial, military, humanitarian, or public safety missions, among others [2].

Through prospective user and technical literature surveys, we have identified a broad range of specific UAS experiment scenarios, and derived common testbed requirements and the technical architecture to serve these. We propose the aerial experimentation and research platform for advanced wireless (AERPAW) [3], which is unique in its combination of being embedded in real outdoor spaces, of a variety of radio terrain ranging from agricultural to urban, and of native support of wireless mobility in experiments. In addition to production-like 5G networking equipment, the platform will feature software radios for both the fixed and mobile nodes. These are programmable by the experimenter, who can define custom experiments with commercial off-the-shelf or custom software and hardware that are properly isolated using virtualization technology and experiment verification and certification with the help of emulators and sandboxes.

978-1-7281-0960-2/19/$31.00 ©2019 IEEE

The remainder of this paper is organized as follows. Section II defines the requirements for the experimentation and research platform. Section III describes the architecture to fulfill these requirements. Experimenter services follows in Section IV. We discuss sample research projects in Section V, before concluding with a discussion of the scope and potential impact of the proposed platform.

II. REQUIREMENTS

AERPAW is meant to serve global researchers and provide them with a platform for experimentation in support of advanced wireless and computing R&D with terrestrial and aerial nodes. The main requirements for supporting user experiments are to provide remote access to advanced and programmable radios of various types, together with other hardware and software resources required to support sophisticated experimentation, including drivers, software development kits (SDKs), experiment templates, compute and storage resources, connectivity, and a variety of real-world spatial characteristics.

A unique requirement for AERPAW is its support for mobile radio endpoints operating in controlled, yet production-like environments. The testbed should enable at scale experiments that require several terrestrial radios and multiple mobile platforms to be operated in a wide geographical region of connected airspace.

Further aspects of the technical architecture are driven by the need for operational safety of airborne mobile nodes, the reproducibility of experiments, and the requirement to on-board users easily to take full advantage of this unique and evolving platform on an ongoing basis.

We conducted a user survey of the UAS communications and computing R&D community. Out of 71 responses, 51 are from professors, 6 are from graduate students, and the rest include people from industry, Idaho National Labs, NASA, and German Aerospace Center. Survey outcomes show that (1) having access to training tutorials/videos, (2) the ability to reserve the testbed, (3) the ability to access the platform through multiple means (web, ssh, etc.), and (4) full remote access to SDRs are the key features desired by the survey respondents. There is significant interest by the research community to use SDRs, both at cell towers and, more critically, at the UAS. Having access to mmWave SDRs and 5G NR mobile phones carried at a UAS are also highly demanded. For UAS applications, search/rescue with UAS, UAS detection/tracking, and the use of UAS as a hot spot base stations (BSs) received the highest interest in our survey.

III. OVERVIEW OF THE AERPAW ARCHITECTURE

The architecture describes the framework, logical components and tools to meet the requirement. Given the scope of this project, the chosen architecture is critical and merits discussion before detailing the platform design and subsequent development, deployment and operation plans.

Fig. 2 shows the overall architecture of AERPAW. We introduce the *AERPAW node* as the broadest generalization of the testbed's physical nodes. Each such node is a computer with varying degrees of CPU, storage, memory, and networking capabilities; the only defining characteristic of an AERPAW node is that it runs an OS capable of virtualization, contains one *Control VM (C-VM)*, running the AERPAW control and management software, and has at least one commodity networking interface that serves the C-VM. Although we refer to it as a VM, this *C-VM* has superuser privileges in the hypervisor (or host OS), and indeed may be implemented in the hypervisor.

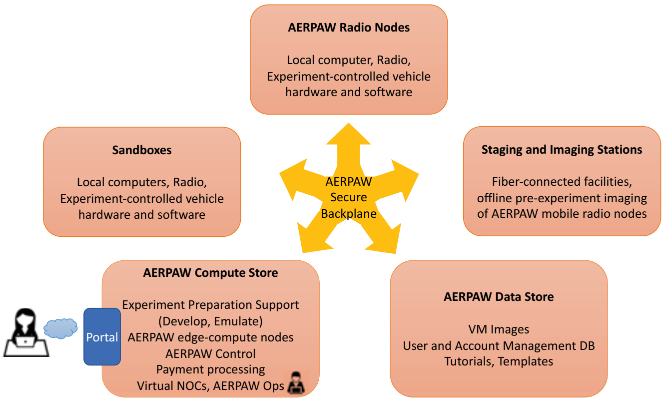

Fig. 2. High-level overview of the AERPAW architecture.

The *AERPAW control center* (conceptually central, but physically instantiated in a few distinct distributed locations, at least one each in North Carolina and Mississippi) will manage and orchestrate all such nodes. The commodity networking interface for this communication will be restricted only to the C-VM. Each AERPAW node will also run at least one *Experimenter VM (E-VM)*. Each E-VM will have a virtual network interface to the C-VM, which will appear to the E-VM to connect to a commodity network containing all other E-VMs of the same experiment. The C-VM will serve this out of the available commodity network interfaces. Security protocols will be realized by all AERPAW system components to provide an environment similar to a VPLS over these commodity networking channels, enabling access control and authorization, so we refer to this as the *Secure AERPAW backplane*. This backplane connects additional AERPAW system components, as described in the next section.

An *AERPAW radio node* will also be equipped with one or more experimental wireless equipment, available to the E-VM (to distinguish it from the special case that an AERPAW node is used as an pure edge-computing resource).

AERPAW will have a variety of advanced wireless equipment available for its nodes, including software-defined radios (SDRs), commercial 5G radios (mmWave and sub-6), LoRa and Sigfox IoT, and ultra-wideband (UWB) radios. Most nodes will be portable and node assignments will change over time in normal testbed operation, and AERPAW's technical architecture is designed to be general enough to handle any assignment. Specifically, an AERPAW (radio) node can installed and used at different locations and for different purposes in the testbed, subject to size, weight, power, and deployment constraints.

To realize the mobility we consider vital for AERPAW experiments, we use “vehicles” on which AERPAW radio nodes are installed during experiments. Some vehicles allow programmatic control of location and trajectory during the experiment, providing *experiment-controlled mobility*. Others represent piggybacking opportunity for AERPAW radio nodes on vehicles that describe trajectories according to their own real-world logic, creating *vehicle-controlled mobility*. The AERPAW control plane provides location sensing for all mobile nodes, and provides location control interfaces to nodes with experiment-controlled mobility.

Note that all node-vehicle combinations and associations are possible, subject to physical constraints. For example, a balloon or kiteoon can carry an aerial fixed or mobile node, whereas mobile aerial nodes may or may not provide experiment-controlled mobility. Similarly, there is no correlation implied with the functionality provided by the advanced wireless equipment; that is, any AERPAW radio node may, depending on the hardware and software capabilities, serve protocol roles such as base stations, UEs, or simply wireless link endpoints.

Most AERPAW fixed ground nodes will be deployed as semi-permanent on towers, poles, rooftops, or other available or purpose-built structures. In general terms, we call these locations *towers*. A physical vicinity in which multiple towers exist, and which offers a radio environment of specific characteristics, is called a *site*. Thus we can speak of sites representing agricultural land, urban areas, or residential suburban areas.

On the software side, we note that orchestration technology has been well explored and tested in the community, and we will make extensive use of the same or similar technologies, learning from the GENI (especially ExoGENI) and ORBIT/COSMOS, among others, and leveraging the well established code bases and communities around technologies such as OpenStack, Kubernetis, and OpenMANO. Similarly, to control advanced wireless devices, we will integrate tools and software that accompany such devices, and also open-source SDR implementations of advanced wireless protocols, such as srsLTE and open air interface (OAI), into AERPAW. Finally, autopilots (using open source software and hardware) such as Ardupilot and Pixhawk II will be integrated into AERPAW to provide control on unmanned aerial vehicle (UAV) trajectories, using open interfaces like MAVLink.

The AEPAW execution will involve the design, development, deployment, and operation of the platform. An overview of the platform is provided in Fig. 3, which illustrates a typical AERPAW fixed tower, a mobile node that can be an aerial or ground vehicle with one or more radios, a control center, AERPAW backplane, data store, and sandbox. Either the fixed tower and/or the mobile nodes can have one or more of the SDR equipment, 5G base station, RF sensor, IoT access station, among other possible equipment. In this section we will provide a simplified project management plan and high level explanation for each of the four stages towards operationalizing such a testbed, along with a description of the key deliverables during the different project phases.

IV. Experimenter Services

The purpose of AERPAW is to serve the global R&D community, as well as local communities, educators, regulators, and other users. Researchers and developers may use AERPAW to, for example, test new research results or technology prototypes in a production-like environment, where system assumptions can be verified in real operational conditions that are to some extent controllable.

A. Canonical Experiment Scenario

We describe below a small experiment (depicted in Fig. 4) that exercises the core functionality of the testbed and that can be used to derive the basic requirements for its design and operation. In designing this experiment, the user (researcher) plans to focus on a single wireless link and has chosen an available fixed location offered by AERPAW for one endpoint and an aerial mobile node for the other, both using SDR technology. The user intends to study the link characteristics as the mobile endpoint of the link traverses a specific 3D trajectory over the surrounding geography. We assume that the exact space-time trajectory of the UAS, as well as the exact sequence of SDR transmissions are decided at runtime by the experimental code implemented on the two nodes. For example, the research might address the fit of modulation schemes to mobile node velocity; thus, waveform configuration and UAS velocity may affect one another at runtime.

From the above, we see that the basic need is for users to be able to program computers with advanced wireless interfaces attached to them. Naturally, an experimenter may choose multiple fixed and mobile nodes and radios of various types. We also allow the user to include pure compute nodes in an experiment, such as to represent edge-compute nodes or the Cloud, or both. All nodes can run user-developed code, and it is a legitimate expectation that these distributed pieces of code will be able to coordinate their execution at runtime, and do so without compromising the integrity of the spectrum being used by the experimental wireless channels. In other words, the testbed needs to provide a back-channel to let the various pieces of experimenter code coordinate “out-of-band”.

We will extensively use virtualization techniques, but by the very nature of the testbed, sharing space or spectrum to collocate multiple experiments in the same location at the same time is only possible to a certain extent without curtailing experimenter freedom to write adaptive code in terms of positioning mobile nodes, varying radio frequencies, and iterating over scenarios. It follows then that experiments must have boundaries in time, space, and spectrum. Violating any of these boundaries must cause the experiment to be terminated in an automatic and responsive manner. The services to enable all these are described in continuation.

B. Experiment Preparation

In designing a community testbed, we are conscious of the need for actively and creatively enabling easy on-boarding, as

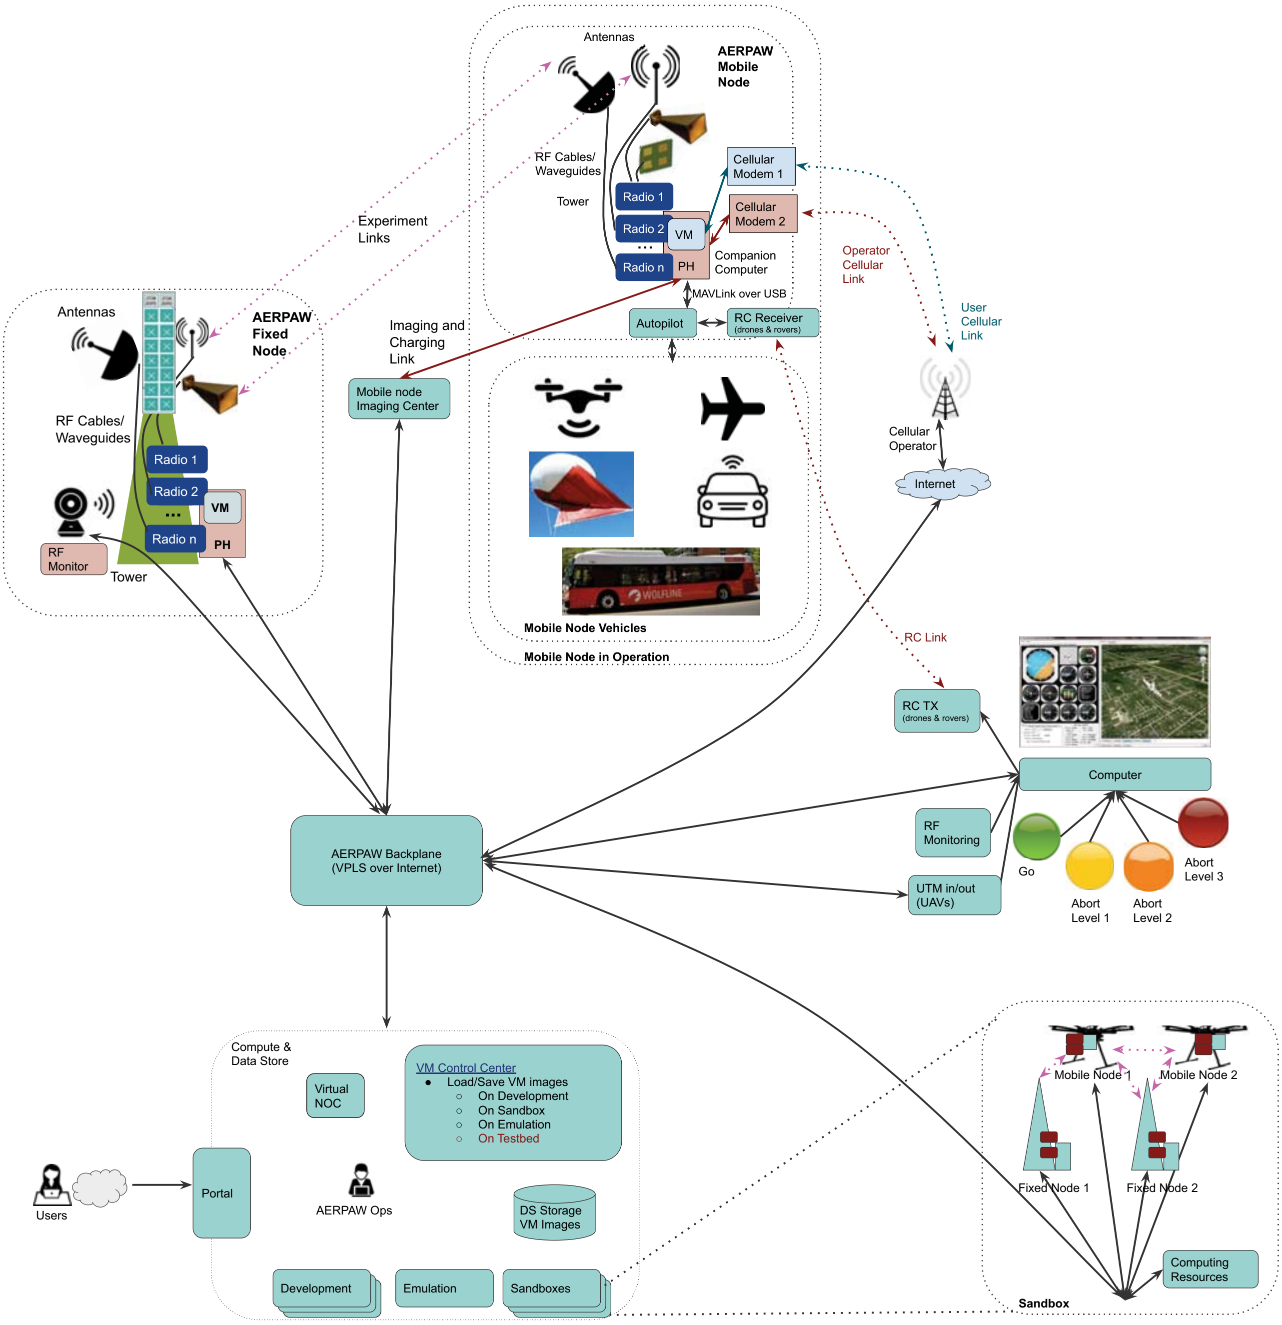

Fig. 3. AERPAW architecture at a glance, illustrating the fixed and mobile nodes, the sandbox, the control center, and the data store.

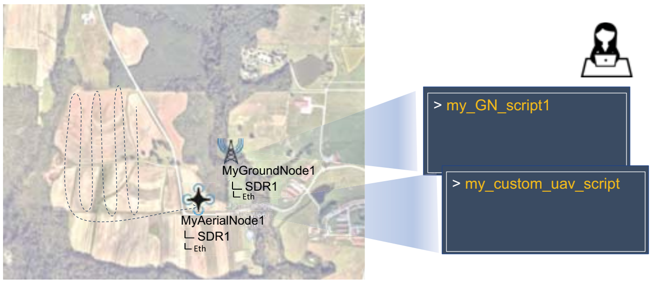

```
> my_GN_script1
```

```
> my_custom_uav_script
```

Fig. 4. Depiction of a canonical AERPAW experiment.

well as low-cost and self-guided training of testbed resources available to users, and have built these considerations into the very process of experiment preparation. There are multiple phases or stages an AERPAW experimenter will go through from inception to the final run of the experiment on the testbed (see Fig. 5). These are described in continuation.

(1) **Development.** In this first phase, the experimenter is required to start designing the experiment by defining the number of AERPAW radio nodes and edge-computing nodes, choosing sites, time, and spectrum. Note that these (and other user actions as below) will be possible to perform on

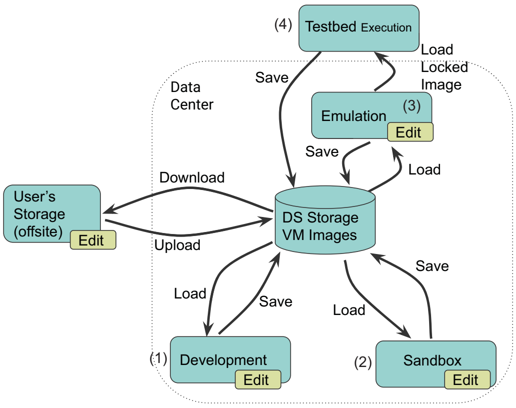

Fig. 5. VM life cycle during experiment preparation and execution.

AERPAW by various means; internally, all such actions reduce to API calls, and sufficiently advanced users may make use of these APIs directly, while others may make use of more intuitive interfaces, such as GUIs, to perform the same tasks.

After specifying the AERPAW nodes, the experimenter receives access to a number of VMs, one per AERPAW node, with certain characteristics (CPU, default drivers, etc.). This allows the experimenter to implement custom experimental logic, sequence, algorithms, etc. From here on, these same VM images are used for successive stages including actual testbed deployment.

(2) **Sandboxing.** AERPAW will maintain a sandbox environment (possibly distributed), containing computing resources to run VMs, and in addition radio and UAV hardware that mirror those deployed in the testbed. Sandboxes are safe execution spaces using measures such as attenuating cables in place of antennas for radios, and elimination of propellers from UASs. The sandbox environment allows the user to gain familiarity with the actual hardware, and ensure that the experimental code at the end of this stage is unlikely to produce unforseen problems when running on the actual testbed. The value of sandboxing is in ironing out the interaction of the experimenter code with the actual drivers and hardware. After such errors are eliminated, the experimenter is ready to submit the experiment.

(3) **Emulation.** At this stage, AERPAW checks the experiment on behalf of the experimenter, where a complete model of the entire experiment is run in software, with simulated models of hardware components, and actual components for every software component. The software simulations of the physical behavior of the UAVs and the radio links will be realistic enough to allow experimenters to gain a realistic idea of how the experiment will proceed in the actual testbed; for example, whether a planned UAV trajectory is likely to send it out of range of another node. When an emulation completes without any obvious errors or violations, the set of VM images that make up the experiment are marked as acceptable to run on the actual testbed. Upon seeing the results of the emulation, the experimenter can choose to modify their experiment; they must then be re-emulated successfully before they can be deployed on the testbed.

(4) **Execution.** In the final stage, the experimenter will deploy their experiment on the AERPAW testbed, where it is executed by the AERPAW operator.

While we described the preparation phases in canonical order above, an AERPAW experimenter will be able to iterate and go back-and-forth over them. Together, the preparation stages serve not only for self-guided learning, but saves time and effort for both the operator and the experimenter, and maximizes actual testbed use, since development, sandboxing, and emulation can be scaled as needed. The sandboxes, and the compute/data store that provides VMs for development and emulation, are connected to the rest of AERPAW by the secure backplane we mentioned earlier.

C. *Safety in Execution*

Although the emulation stage provides an advance check on the likely safety of the experiment, monitoring during execution is indispensable. The main checks are whether any violations are occurring for the predetermined boundaries on the spectrum (radiating in spectra outside those permitted for the experiment), space (experiment-controlled mobile nodes traveling outside the permitted boundaries), and time (experiment not terminated by the end of the scheduled period). AERPAW will use both automated and human monitoring, and will use a variety of different mechanisms. For example, a UAS encroaching beyond the experiment boundaries can be detected both by geofencing techniques as well as a human observer (pilot), but the latter is considered more serious, since only a more gross violation is likely to be observed by the human observer. Spectrum violations will be monitored by a frequency monitor that is installed in strategic locations.

When a violation is detected, AERPAW control will abort the experiment. Multiple abort modes exist for abort requirements of varying urgency. Every component of the AERPAW execution environment will recognize at least two levels of abort. The softer, or lower severity abort (“orange”) indicates that experiment execution should stop, and the state of the components be gracefully recovered. The higher severity (“red”) abort signifies immediate emergency shutdown of each component, potentially sacrificing state information.

V. ENABLED RESEARCH SAMPLES

A. *5G Security*

5G proposes a new security architecture that has yet to be tested, but initial studies already indicate several potential issues [4]. For LTE, the main security vulnerabilities were not identified and reported in open literature until years after the standards have been finalized and networks deployed. One of the reasons for this was the lack of available and affordable tools for LTE security research. LTE open-source software libraries that can run on personal computers and use commercial off-the-shelf software-defined radio (SDR) peripherals did not reach a sufficient level of maturity until recent years. Once software-defined LTE testbeds became available, many security problems were validated and practical solutions proposed [5]. Security and privacy in wireless networks has been a research topic for decades and is receiving a big push now because of the critical services that 5G networks will offer. User privacy, such as user identity and location, can be compromised because of how the 5G security specifications are defined [4]. The platform will allow testing the severity of the different types of 5G protocol exploits and implement and test countermeasures in a realistically deployed 5G experimental network.

In order to study the security weaknesses of next generation mobile communication systems, open-source implementations of base station and UE protocol stacks will be leveraged and a interference waveform library developed to launch controlled attacks and measure their effects. For example, if future 5G networks are used for command and control (C&C) of UAVs, these signals must be protected from interference and spoofing. Otherwise, an attacker may able to send valid C&C signals to one or multiple UAVs to take control of them.

B. Distributed Resource Management

Another sample research project is distributed scheduling and resource management in a D2D or V2V scenario. For example, 3GPP Release 14 specifies sensing-based semi-persistent scheduling (SPS) for LTE-based C-V2X. Despite the fact that testing has started on US roads, research has shown that SPS leads to many packet collisions [6], [7]. Hence, new scheduling approaches and congestion control mechanisms are needed, and researchers and developers expect important opportunities for R&D enabled by the flexibility and versatility of 5G systems. As opposed to using network simulators, such as ns-3 used in above papers, rapid prototyping and experimental verification of new technologies and procedures is invaluable especially for critical systems such as C-V2X and UAV communications. Such experiments will provide the real data that is needed to calibrate network simulators of emerging wireless technologies.

C. AI for Advanced Wireless Networks

A third sample research project is about resource and network management, parting from manual or well-defined management to learning-based optimization of network parameter settings, waveforms, hardware configurations, user services, and schedulers. Artificial intelligence (AI) and machine learning (ML) tools are gaining momentum, also for 5G [8]. AERPAW will expose access to all system parameters and educate researchers what their roles are and what the recommended ranges are. Researchers can push the limits and develop AI mechanisms on top of wireless systems for effective system or subsystem management and generate training data with AERPAW and without the restrictions of the data collected by commercial carriers. Ongoing research, such as [9], suggests that contributions using AI/ML for optimizing different parts of 5G or 5G-like systems will proliferate and all of them will need a platform to test these potential breakthroughs in wireless resource, network and service management.

UAS are complex systems that enable many new applications. Apart from wireless communications and mobile edge computing, there are several other technological and non-technological issues related to the deployment of UAS. AERPAW will also enable the testing of other related technology, systems and policies/regulations, including control systems, sensors, computer vision, compression, privacy, embedded systems, new energy sources and charging technology, sense&avoid, and transportation.

VI. CONCLUSIONS

Traditionally used mostly for non-critical voice communication, many of the current and emerging data and control communication systems that leverage cellular access networks have stringent requirements in terms of availability, reliability, integrity and privacy of services and user data. Applications include tactical communication, first responder ad-hoc networks, and mission-critical Internet of Things.

Because of the SBA concept that 5G introduces and the different commercial, military, and industrial verticals that 5G will enable, extensive and continuous testing of new protocol features, services and deployments in production-like environments is necessary. Specifically, new communications and computing services leveraging UASs and advanced wireless protocols, require an scalable testbed with configurable computing and networking nodes where standardization bodies, researchers, regulators, and industry can work together to test and evolve integrated terrestrial-aerial communications architectures and applications. We therefore propose a platform for experimental research in wireless communications and computing that deploys SDRs and 5G network infrastructure on ground and aerial nodes. The AERPAW team will collaborate with industry and users to deploy and operate such a network and support a variety of experiments.

ACKNOWLEDGMENT

The authors would like to thank the NSF PAWR program.

REFERENCES

[1] F. Boccardi, R. W. Heath, A. Lozano, T. L. Marzetta, and P. Popovski, “Five disruptive technology directions for 5G,” *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 74–80, 2014.

[2] J. Kakar and V. Marojevic, “Waveform and spectrum management for unmanned aerial systems beyond 2025,” in *Proc. IEEE PIMRC*, 2017, pp. 1–5.

[3] “AERPAW website,” https://aerpaw.org/, [accessed July 2019].

[4] R. P. Jover and V. Marojevic, “Security and protocol exploit analysis of the 5g specifications,” *IEEE Access*, vol. 7, pp. 24 956–24 963, 2019.

[5] R. P. Jover, “The impact of open source on mobile security research,” Tech. Rep., May 2016, https://goo.gl/hi4ukn, [accessed July 2019].

[6] R. Molina-Masegosa and J. Gozalvez, “LTE-V for sidelink 5G V2X vehicular communications: A new 5G technology for short-range vehicle-to-everything communications,” *IEEE Veh. Tec. Mag.*, vol. 12, no. 4, pp. 30–39, 2017.

[7] A. Nabil, K. Kaur, C. Dietrich, and V. Marojevic, “Performance analysis of sensing-based semi-persistent scheduling in C-V2X networks,” in *Proc. IEEE VTC-Fall*, 2018, pp. 1–5.

[8] M. Yao, M. M. Sohul, V. Marojevic, and J. H. Reed, “Artificial intelligence defined 5G radio access networks,” *IEEE Commun. Mag.*, vol. 57, no. 3, March 2019.

[9] M. Chen, W. Saad, C. Yin, and M. Debbah, “Data correlation-aware resource management in wireless virtual reality (VR): an echo state transfer learning approach,” *IEEE Trans. Commun.*, vol. 67, no. 6, pp. 4267–4280, June 2019.