Final Report from the

NSF Workshop on Future Research Infrastructure for the Wireless Edge

Washington D.C. on November 13-14, 2014

Workshop Chairs

Suman Banerjee (UW-Madison) NSF PI, Mark Berman (BBN Technologies), Abhimanyu Gosain (BBN Technologies), Ivan Seskar (Rutgers)

Workshop Participants

- Abhimanyu Gosain, GENI Project Office
- Alex C. Snoeren, UC San Diego
- Ashutosh Sabharwal, Rice University
- Athina Markopoulou, University of California, Irvine
- Bob Iannucci, Carnegie Mellon University
- Bryan Lyles, NSF
- Carl Kutsche, Idaho National Labs
- Changbo Wen, T-Mobile
- Darleen Fisher, NSF
- Geoffrey Challen, University at Buffalo
- George F. Riley, Georgia Tech
- Gil Zussman, Columbia University
- Grant Miller, NITRD
- Hongwei Zhang, Wayne State University
- Houbing Song, West Virginia University
- Ian Wong, National Instruments
- Ivan Seskar, Rutgers University
- Jack Brassil, HP Laboratories
- Jim Martin, Clemson University
- Joseph Soriaga, Qualcomm
- Kobus Van der Merwe, University of Utah
- Lakshminarayanan Subramanian, NYU
- Larry Landweber, UW-Madison
- Mark Berman, GENI Project Office
- Mario Gerla, UCLA
- Michael Ha, FCC
- Murat Yuksel, U of Nevada - Reno
- Nada Golmie, NIST
- Rangam Subramanian, NTIA
- Preston Marshall, Google
- Romit Roy Choudhury, UIUC
- Sampath Rangarajan, NEC Laboratories America
- Suman Banerjee, UW-Madison
- Thanassis Korakis, NYU
- Timothy J. Salo, Salo IT Solutions, Inc. / University of Minnesota
- Vanu Bose, Vanu Inc
- Wenye Wang, NC State University
- Xia Zhou, Dartmouth College
- Yunsheng Wang, Kettering University
- Z. Morley Mao, University of Michigan

Contents

1. Introduction

Mobile and wireless-based access has transformed the nature of computing, communication and access in the last decade. The number of mobile wireless network-connected devices (i.e. smartphones, wearable computers, Internet of things, etc.) is anticipated to grow exponentially. As these devices become pervasive and the expectations from users continue to grow, researchers and practitioners are experimenting with new mechanisms to deliver at-scale services connecting the distributed cloud directly to end users.

Today the most significant infrastructure for wireless experimentation at-scale is available through NSF programs such as GENI WiMAX, ORBIT, Emulab etc. Researchers have found wireless access to GENI's distributed cloud infrastructure to be a very fruitful environment for experiments and service deployments in a wide range of disciplines, including:

- Novel, non---IP mobile internetworking protocols
- Vehicular networking and applications
- Emergency, public safety, and healthcare IT communications

However, the diversity of wireless technologies (from 802.11x to WiMAX to LTE), each network's dramatically different coverage areas (from personal- to wide-area), and lack of a large 'representative' pool of mobile client devices places limitations on what researchers can learn from today's experimental infrastructure.

Hence, it is important for the community and the National Science Foundation to have a broad vision of the large scale future infrastructure needs to address key research challenges facing the wireless networking community as well as applications that can be supported in the next decade and beyond. As such, the focus of the NSF workshop on Future Research Infrastructure for the Wireless Edge held in Washington DC from November 13-14 was to bring together experts from academia, industry and government, to plan future research infrastructure that supports such experiments and service deployments.

2. Workshop Objectives and Structure

The goal of the workshop was to help chart a broad vision for the future wireless infrastructure needs for the research and academic communities. At the workshop, attendees discussed and debated the scale, operation, size and technological make-up of this infrastructure, and articulated a vision for the infrastructure and the associated research and applications that would be enabled.

This report has been written with the intention of (i) identifying the scale, scope and size of future wireless infrastructure; (ii) exposing the new and exciting research and applications issues; and (iii) stimulating academia and industry collaborations that would help along the evolution and operations of this wireless infrastructure.

The above questions were addressed by breakouts and were broadly partitioned into a number of technical areas, which include:

- Research
- Applications
- Industry Partnerships
- Outreach and Education
- Operations
- Spectrum Policy, Privacy and Security
- Open Source Research Platforms

To discuss these questions, a total of 40 participants were selected for this workshop comprising of some invited attendees and some from an open participation call issued to get broad participation from the wireless research community. The participants were divided into various breakout discussion groups throughout the day and a half event to allow more focused discussion on individual topics. The event was kicked off by a keynote address from Preston Marshall, Google and on day 2 a "panel" of Joseph Soriaga (Qualcomm), Michael Ha (FCC), Nada Golmie (NIST), Rangam Subramanian (NTIA) and Vanu Bose (Vanu,Inc.) addressed spectrum policy issues and alluded to other government agency efforts in this space. Various other individuals were given specific responsibilities in the workshop, e.g., to serve as leaders and scribes of specific discussion groups. This report incorporates input from various breakout groups and summarizes the discussions across the different groups throughout the workshop.

The detailed agenda of the workshop is presented in the Appendix.

3. Future Wireless Infrastructure Requirements

In the context of this workshop, the following seven topic areas are divided into the following sections that provide insight into requirements of future wireless testbeds.

- What are the research limitations of Today's Experimental Infrastructure?
- What are the recommendations and next steps?

3.1 Research

A crucial set of research questions we now face concern the architecture and design of a wireless edge network that will allow us to scale to a much larger number of connected devices. Our current network densities -- in client devices per unit area -- are expected to grow in the next decade by two orders of magnitude or more. That level of scaling threatens to overwhelm or break our current wireless networks, and motivates us to revisit conventional design principles, and explore new wireless network technologies capable of supporting dramatically higher client node densities.

3.1.1 What is not possible with today's Infrastructure

Our understanding of wireless networks has been well-served by deployment of experimental infrastructure (e.g., Orbit, GENI WiMAX, DOME, etc). Furthermore, recent efforts toward federated infrastructure have eased access for the research community. But the diversity of wireless technologies (from BLE to 802.11x to LTE), each network's dramatically different coverage areas (from personal- to wide-area), and lack of a large 'representative' pool of mobile client devices each place limitations on what researchers can learn from today's experimental infrastructure.

- Inability to study existing networks at scale

The research community needs a deeper understanding of how current operational architectures, protocols and algorithms scale. For example, if an existing network suffers performance degradation as traffic increases, we seek to identify its key bottlenecks. Otherwise, we risk encountering similar poorly understood emergent behaviors in the next generation of networks we design. Considering the huge complexity of an operational network with its many inter-related components (e.g., physical layer, access layer, network and mobility management, resource allocation), it is quite possible that performance degradation is not because of only one layer but a function of complex interactions between different protocol layers. Thus, it is desirable that the community has open access to experimental infrastructure that supports a deeper understanding of current operational networks at their highest scales in terms of user density, traffic and mobility.

- Limited access to cutting-edge networked devices and settings

The last decade has seen the rise of an amazing new number of *connected* applications - self-driving cars, infrastructure sensing, healthcare, environmental sensing, crowd-sourced information gathering, drones, wireless power meters and so on. Many new classes of applications have significantly diverse design objectives, and hence network and computational needs. To root our discussion in concrete terms, some of the growing applications represent significant departures from the conventional networks (and associated traffic) that we are familiar with today. These emerging applications include (i) connected vehicles (cars, fleets of trucks, automation for), (ii) health-care beyond traditional venues, i.e., clinics & hospitals, (iii) smart-grids, (iv) intelligent buildings, (v) smart cities, and (vi) disaster connectivity. Many of the new applications are mission-critical and hence the data traffic has to be delivered with high reliability, very low delay and extremely small jitter. To take one example, safe control of self-driving cars and drones will warrant an extremely high level of reliable connectivity. In addition, many mission-critical applications will require that network elements are also computational units, that is, they allow local computations done close to the source of data, with subsequent actuation.

In addition to these new network settings, the research community has limited access to emerging client devices (e.g., smartphones, wearable devices) that will be key attached components of future network infrastructure.

- Limited access to emerging large-scale operational networks

As we learn more of how to design next-generation networks from smaller testbeds, it is crucial to validate and optimize behavior at giga-scales. That is, while small scale specialized testbeds provide excellent control and initial validation; the promising concepts have to be tested at scale before they can be deployed globally. For example, existing testbed infrastructure was not intended for the study of the large-scale machine-to-machine networks, and the traffic types and behaviors we anticipate in that operational network. Similarly, we anticipate the need for large networks to operate in new and potentially shared radio spectrum, which is difficult or impossible to study with existing experimental resources.

3.1.2 Recommendations

- Understanding behavior of our existing networks at scale

A possible methodology to examine existing operational networks at scale will be to “tap” into those networks with specially designed “probes” or interfaces that capture data at multiple levels

- radio, signaling, packets, flows and apps. Collecting such a massive, rich, open and well-documented data-set -- even if for a small duration (say a few days to up to a few months) -- could truly open up new research directions.

a. **Quantification of Bottlenecks:** Two common approaches to increasing network capacity have been to 1) increase physical layer transmission rates and 2) increase network density. However, the increase in actual delivered throughput appears to be growing at a slower pace than the increases in physical layers and network densification. That is, we are losing throughput in our overall network deployments, and the causes of this loss are poorly understood. An immediate outcome of the massive, well-documented and open data-set is our ability to answer the following research questions
    i. What are the major causes of reduced capacity scaling in the current networks?
    ii. What cross-layer optimizations are possible today on current networks which will increase their capacity?
    iii. What are the fundamental limits in scalability of current network architectures even with optimizations?

b. **System Models for Future Analysis and Simulations:** An immediate outcome of visibility into a large-scale network operation will be development of network models that capture effects which are typically *not* studied together but have a direct impact on each other. For example, resource allocation per cell and across cells is directly impacted by mobility management techniques. However, we have no data-driven multi-cell models that capture physical layer conditions, user mobility, traffic matrices and user density. While understandably complex, having models driven by actual large-scale measurements enable us to answer to following broad set of research questions:
    i. If we replaced all or part of the network with a new access, network and mobility management technology, how would it impact the scaling behavior of the network?
    ii. If the cloud architecture was moved closer to network elements (as in `fog computing'), how would it impact the performance of specific traffic types (e.g. video)?

c. **Development of Massive-scale Management:** A key challenge in scaling network user density is network resource management at multiple time-scales, from short-term packet time-scale management to large-scale resource allocation and load balancing. Future networks are expected to be more heterogenous than today's, i.e. it is quite likely they will simultaneously leverage many links across different networks (multi-homing). The big management research challenges are:

- **i.** What network management fundamentals will enable us to scale optimally to expected user densities?
- **ii.** What type of practical architectures will reach those fundamental limits? Centralized, distributed, self-organizing, a hybrid of all?
- **iii.** In a landscape enabling open competition in deploying infrastructure, how can network operators build interference management as a fundamental construct, without having to "wire" all networks together to one centralized controller?

- Increasing access to cutting-edge networked devices and settings

For the development of testbeds deploying next-generation network technologies, it is occasionally desirable to develop dedicated smaller testbeds which allow researchers to control all aspects of the testbed. The increased level of control is important when the research is in such early stages that all aspects of network and computing are open. With such an open testbed, which is highly configurable, flexible and programmable, the following research directions will be enabled:

1. **Networks Designed for Low-latency Mission-critical Applications**: Control of many mobile entities (cars, drones) over lossy, unreliable wireless networks is relatively immature. For instance, what are the right network and computing design principles for controlling a large number of mobile entities where the cost of a mistake is extremely high? For example, an error in control of a connected car or drone is a potential accident that can be life-threatening. Can such mission-critical data supporting networks be designed in conjunction with other traffic in the network, which does not have similar requirements?

2. **Computing at the Edge**: When we co-locate computing and communications in one box, it is possible to perform optimizations which were previously not possible. The resulting design space is much larger than what is available in the current networks. The larger research questions relate to how network control and computation should be partitioned across different devices in the network - the mobile entity, network edge infrastructure, and the network core.

3. **Control of Massive Spatial and Spectral Resources**: With increased density of infrastructure and mobile nodes, the resulting networks have very high number of spatial resources (e.g., antennas, and number of antennas per device), and a large amount of

power and potential to access to diverse spectral (multiple bands) resources. The result is very resource rich, and at the same time, highly interference-rich environment. The open research questions relate to foundations of signaling, management and control of such massive amount of resources in a scalable manner, in both amount of computation and control overhead per node. An open question for study is the tradeoffs between fully distributed and fully centralized control.

.

- Designing for access to future large-scale operational networks

An interesting idea to ensure that researchers have access to future, large-scale operational networks is to ensure that they support virtualization. In such an environment a researcher could obtain a slice of a deployed wireless network and test novel research ideas with real deployed users in a live network. This methodology would enable potentially ground-breaking research in the design of **clean-slate network architectures**. A recurring theme in the above discussion is the architecture of the future networks. Some of the key aspects managing mobility and resource allocation across the network, while meeting all traffic demands for the types of traffic we have not experienced yet (e.g. self-driving cars, drones, Internet of Things). Thus, “convincing”-scale experiments can be performed on real networks but with the safety of not harming anyone; all test users in these experiments will be recruited with full consent and tests will only be deployed after significant testing on small scale networks.

To enable “convincing”-scale experiments, it is clear we need **research client devices** that permit much more access to low level interfaces than today’s mobile operating systems. That is, there is a pressing need to develop one or more types of fully programmable user devices to facilitate research. One device might serve as a “research phone” to be used in conjunction with testbed infrastructure, offering researchers unprecedented visibility to user and network behaviors (for those users electing to use the device). Other research clients might address other emerging applications, such as small form factor devices for investigation of IoT or machine-to-machine communications. We envision such a device to be entirely non-proprietary, permitting researchers to modify hardware and software spanning all protocol layers. In addition, the device might support multiple interfaces and multi-network connectivity (i.e., multi-homing at the interface and network levels). The device would be modular, permitting the inclusion of new interfaces, perhaps via a SIM or USB card slot. An example of a similar proposed device is the modular phone proposed in the Google Project Ara. The existence of targeted ‘Research clients’ will permit the design of new applications which leverage the devices’ unique functionality.

There are at least three fundamental research problems related to the creation of such a device. The first question is the electronic/mechanical/software design and architecture of a device that would permit all of the desired characteristics described above.

The second fundamental problem addresses the human-computer interface. To obtain representative data from a wireless testbed, it is critical that a large group of users *elects* to use the research device, and further uses it in a fashion similar to use of a personal device. The research community recognizes this to be a significant challenge – even minor changes (e.g., adding a dongle, changing a personal phone number (IMSI), lack of applications, etc) are enough to cause users to balk or renege on use, or simply produce non-representative data. Finding designs that minimize the friction of using a 'research client' is a critical challenge.

A third potential investigation is the joint optimization of infrastructure and mobile devices. The conventional assumption is that both network infrastructure and devices are essentially static and well known. But the network could potentially offer different services to different devices, or different services on-demand over time to a client, and understanding the tradeoffs associated with such dynamics, as well as programming interfaces and protocols, is currently not well understood.

3.2 Applications

3.2.1 Introduction and Background

The last decade has seen the rise of an amazing new number of *connected* applications - self-driving cars, infrastructure sensing, healthcare, environmental sensing, crowd-sourced information gathering, drones, wireless power meters and so on. Many new classes of applications have significantly diverse design objectives, and hence network and computational needs. The group identified applications that tackle emerging technical or societal trends, are reliant on characteristics *missing* from mobile networks, and need 10^3, 10^4, 10^5 endpoints.

3.2.2 What is not possible with today's Infrastructure

Today's wireless networks have enabled a computing revolution – moving many applications off of the desktop or laptop and making them mobile. The programmability of smartphones has been a significant factor in bringing about this revolution. Networks have played an important, but self-limiting role in the revolution. Broadly available wireless coverage and growing capacity have been significant factors. But the inability to closely intertwine computing on-phone (or in a wireless sensor or actuator), computing in the cloud, and communication has limited the class of possible applications.

We attempt to identify broad classes of applications that might be possible and to group them according to similar characteristics. These are

- Networked Outdoors
- Networked Cities
- Networked Vehicles
- Networked Me, and
- Networked Network

For each of these, we identify specific applications and their inherent value to society, their networking infrastructure requirements, and the characteristics or capabilities that are missing in today's networks.

- Networked Outdoors

The implicit promise of research in cyber-physical systems is the ability to link our physical world with the increasingly-rich world of information. Wireless networks facilitate accessing and coalescing fragments of information to support decision systems on a global scale. Earthquake early warning systems offer the hope of detecting earthquakes as close to their epicenter as possible and accelerating the propagation of this information from the *speed of sound* to the *speed of light*. Aggregation of sensor information on a wide scale can identify trends that impact health and safety – for example, the evolution of a plume of airborne contaminants.

The challenge in some of these applications, beyond providing network coverage in potentially remote locations, is implementing sense-compute-actuate processes in real time. Today's networks don't support the embedding of arbitrary application code and compel data to traverse potentially highly latent networks. Our view of networks in support of cyber-physical systems is stunted to the extent that the networks can't serve as general, flexible real-time computing platforms.

Applications	Infrastructure Requirements
• Earthquakes, Floods and other hazard sensing • Sensing of man-made hazards (radiation, chemical, biological) • Water quality monitoring • Environmental sensing • Network-supported UAS, multispectral imaging	• 100% coverage – including vertically • In-network computing and storage
Value	**Missing**
• Connectedness saves lives and property • Alerting: Milliseconds matter • Cyber-agriculture	• Low latency sense-compute-actuate • Ability to push computation into network • Ability to program multiple network elements as a single system (sensors, phones, SDN, in-network virtualized computing instances) • Mechanisms for enforcing deadlines (real-time computing)

Networked Cities

The costs of building and maintaining infrastructures in today's cities as well as the need to significantly increase the resilience of urban communities suggest that pervasive use of information technologies in the implementation of city infrastructure can yield significant improvements. Physical mechanisms for re-routing traffic based on real-time models, tools for observing and supporting public activity without having to be physically present, improved abilities to respond in times of emergency, and smart grids all put strain on today's wireless networks.

While prototypes and small-scale deployments of such futuristic city infrastructure capabilities are in the news, broad deployment will demand re-thinking. Coverage will need to withstand potentially significant physical disruption. Capacity will need to scale with demand on day-level, hour-level, and potentially minute-level scales. Location-based services will need to work in every corner of a city – above ground, at ground, and below ground. Networks will need to be as reliable as paint – and just as easy to deploy and re-deploy.

Applications	**Infrastructure Requirements**
• Software-Defined Roads • Street Video (not just Street View) • Surveillance-as-a-Service • Street Light Control • Emergency Response • Smart Grid	• High capacity, high bandwidth • Highly reliable wireless connections (security, robustness) • Location-aware passive infrastructure
Value	**Missing**
• Life, property, convenience	• Ability to co-scale capacity with demand • Jamming resistance • Reliability on par with comparable city services

Networked Vehicles

Vehicles have been late to the party in leveraging wireless networks, but that gap will close, and vehicles may well take a leadership role. This statement derives from the fact that today, a mobile-terminal-equipped car, for instance, is little more than a car with a phone under the hood. But phones, as terminals, were designed subject to extreme space and power limitations, neither of which cars have. One can easily imagine a car as a terminal that computes more, communicates more, and imposes much greater demand on the network (self-driving, telemetry-rich, entertainment-rich, work-environment-rich). But by the same logic, cars and other vehicles have the potential of themselves becoming network elements.

As in other use cases, car-as-network would require a kind of cooperative relationship with the network that is foreign to current network architecture. As in the case of networked cities, robustness of coverage will exceed today's realities. Real-time interactions with roadside services necessitate reduced latencies. The ability to program sets of vehicles and segments of network elements as a whole necessitates a fresh look at the mobile app paradigms.

Applications	**Infrastructure Requirements**
• Software-defined cars • Self-driving cars and car swarms • Collision avoidance systems • Work-while-commuting environments • Nanny telemetry (for insurers, parents) • Fleet telemetry (for auto manufacturers) • Car-as-network, Quadcopter-as-network	• Robust coverage – as reliable as asphalt • Road-localized high bandwidth • Low-latency near-to-near comms (~msec) • Computing-follows-vehicle • Driving-tailored CDN • High-throughput rest stops

Value	**Missing**
• Road capacity • Safety • Convenience • Reliability • Improved performance	• All the “-ilities” (reliability, scalability, …) • Open network architecture, ability to establish trust with cars, quads, phones, … • Virtualized computing in the network • Distributed programming • Mechanisms for enforcing deadlines (real-time computing) • Low latency sense-compute-alert

Networked Me

Consumers have been the primary beneficiaries of smartphone and wireless network technologies, and the pace of innovation will accelerate. On-phone apps have evolved from novelties to necessities. With the increasing integration of in-phone and wirelessly-connected-to-phone sensors, we can anticipate an explosion of applications that connect people to their places of work, homes, health care providers and others. Phones are a natural point of connection, computing and storage.

Mobile health applications and the network services behind them, when they enjoy the success that many anticipate, will evolve to demand a level of auditable reliability that is presently required of other life-support systems. But wireless networks will need to do so in the presence of massive data streams and ongoing, sophisticated attack.

Application	**Infrastructure Requirements**
• Mobile Health • Home automation	• Integration of sensors and network – cooperative power management • Phone is both local signal processor and gateway – deep network programming • Cloud connectivity for analytics, storage, alerting
Value	**Missing**
• Life, property	• Ability to handle massive data streams and data sets • End-to-end security • Reliability to the level of life-critical systems

Networked Network

Just as the network provides value in interconnecting entities in the cyber and physical worlds, it also provides value in interconnecting its own elements beyond basic functionality. By taking a more expansive and less-traditional view, the networking functionality on smartphones and smart devices as well as network-related functionality in apps can be considered part of the network. This makes sense because the overall behavior of wireless networks depends to a very significant degree on the aggregate behavior of these non-traditional elements.

If we were to adopt this view, however, we would face a conundrum. Today's networks are designed with a more-or-less adversarial view of smartphones, smart devices, and apps. To have the network embrace these elements would necessitate an interface across which phones, devices, and apps could, in a trustworthy way, exchange information with the network for mutual benefit. This would allow, for instance, the ability to have networks move and cache content based on what phones, devices, and apps might predict about their future state (where they will be) and needs (what they will use the network to access). This, in turn, could lead to a significantly improved quality-of-experience.

Application	**Infrastructure Requirements**
• Networks and Apps Co-operate • Content pre-caching • Local CDN (*e.g.*, scene-of-the-accident) • Network traffic coalescing – auto-broadcast	• Apps share context, plans w/network • Network shares state (e.g., radio shadow) w/apps • Both cooperate, co-optimize • Cloud comes down to earth (fog computing)
Value	**Missing**
• Improved resource utilization • Quality of experience	• Means to establish mutual trust between phone and network • Means to aggregate instrumentation across network elements + apps on the network and to reason about it • Research testbed requires humans-in-the-loop to generate interesting behaviors

Common Themes

In reflecting on these exciting future applications, several common themes emerge. First, we observe that the future network must provide communication **as well as** computation and storage services. This is a sea-change from the past and must include

- Open, flexible, deep (sensors, phones, network, app) programmability
- Ability to move computation and data toward the endpoints
- Support for programming many elements and the network as a whole

The challenge is not in offering the potential to move some *specific* computations into the network but rather the potential to move *any* computation into the network – under an open programming and provisioning model.

Second, analysis of these future applications suggests that capacity is not just a radio access network issue. Shifts in usage patterns impact the loads imposed at higher levels of the network and may necessitate on-the-fly reconfiguration in network topology. Importantly, research in areas that are subject to such usage pattern dependence will benefit from real test networks with real users making real choices day-to-day that surface such real fluctuations that *simulation and modeling alone cannot properly capture*.

Third, with our current perspective informed by the well-known evolution of smartphone usage, we can only assume that the Internet of Things, with 10's to 100's of devices for every smart phone, will place even heavier demands on the network:

Network research for IoT lags app/sensor/low power device research
Scaling-up yields increasingly-challenging research questions
Deadline-based (real-time) computing
Scalable maintenance and management

Emergent Focal Point Concepts

As we discussed these applications, the domains in which these applications will emerge took on their own identities:

- Cyber-GeoPhysical Systems and the Internet of Natural Things,
- Cyber-Aeronautical Systems, and
- Cyber-cities, -homes, -health

3.2.3 Recommendations and Next Steps

Out of these observations, three testbed concepts emerged:

- The Cyber-Physically Enabled Network Testbed
- The Cyber-Aeronautical Network Testbed for Civilian Applications
- The Multi-Campus Research MVNO

- The Cyber-Physically Enabled Network Testbed

A common element in many of the applications we considered is the ability to bridge the "cyber" to the "physical" via an inherently-programmable network. In Figure 1, we depict such a network testbed with the presumed ability to marshal computing (and storage) resources in the network toward the network endpoints that use them with the objective of minimizing latency in order to meet real-time computing constraints. To achieve this so-called cyber-physical enablement requires new degrees of freedom in creating and operating interconnected network elements; new means of creating "apps" that engage both the network and the devices and that include timing as an integral aspect of correct functionality; new tools for managing computing, communication and storage resources in the face of conflicting requirements; and a blurring of the boundaries between network, computer, sensor, and actuator.

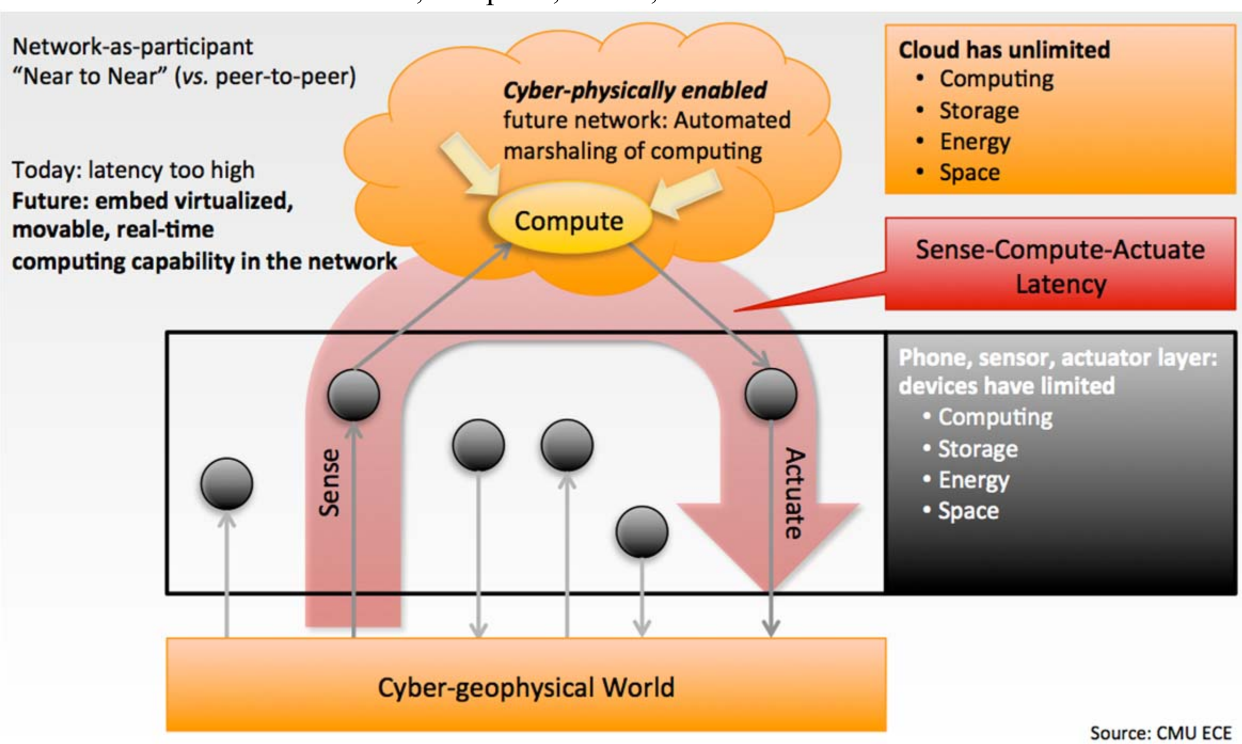

FIGURE 1

- The Cyber-Aeronautical Network Testbed for Civilian Applications

A specific example of a cyber-physically enabled network is a "cyber-aeronautical" system, which has its *unique opportunities and challenges* emerging for network-connected unmanned aerial systems (UAS). Such systems are poised for dramatic growth (pending the loosening of FAA regulation). Already, in their pre-commercial state, we foresee expectations being imposed on cellular networks for telemetry, flight planning, swarm control, and other aspects of wireless operation.

This testbed concept is built on the presumption that UAS devices and UAS-supporting wireless networks should be co-researched, and that there is a potentially very significant advantage to creating physical facilities that enable the US research community to *work well ahead of emerging commercial regulations*.

One can therefore envision a testbed across a number of large (many square miles) areas around the country coupled with appropriate wireless infrastructure equipped with the means to carry out cyber-aeronautical research. Few, if any, institutions can muster the necessary lab facilities, and not all can secure the right-to-fly under current regulations. The installations would each have

The right-to-fly under FAA regulations
"Interesting" terrain and simulated environments, similar in principle to military Combined Arms Collective Training Facilities (CACTF)
Purpose-built, open, multi-node 4G/5G network
Collaborative workspace, equipped labs, expertise
Flight schools, workshops, shared flight software

- The Multi-Campus Research MVNO

Agreeing that "scale" is a desirable attribute of a research network, the committee proposed the creation of a full-function cellular network under the control of researchers, but built by leveraging existing commercial cellular networks. The fundamental premise is that of a single NSF-operated network built out as a set of software-defined network elements (RAN, SDR, NFV) on the campuses of participating institutions, and federated as a Mobile Virtual Network

Operator (MVNO). This on-campus functionality would be backed up by an agreement with one or more commercial carriers to allow participating smartphones and devices to "roam" to commercial networks off-campus. Students, staff and faculty would be able to use this network via smartphones and other wireless devices just like a commercial network. But unlike a commercial network, the NSF wireless network would support instrumentation and reprogrammability at every level, at least for the on-campus elements. In essence, this would be a mobile incarnation of the NSFNET and would offer

- Scale and stress through real usage
- Human-in-the-loop
- Federation across campuses but with "one network" architecture
- Partnership with a carrier for off-campus coverage
- Full integration with the public switched telephone network (together with roaming capability, this would allow users to stay with one phone at all times, improving the scale of use)

3.3 Industry and Academia Partnerships

3.3.1 Motivation

A city scale wireless network deployment involves coordination of multiple partners. These are academic institutions, city governments, spectrum agencies, network operators and equipment vendors (Base station and User equipment). A fundamental aspect of a city scale testbed deployment is a strong partnership between all interested parties. Building these partnerships requires synergetic collaboration between academia and industry. Industry focus has recently shifted from hardware centric, highly constrained, and disaggregated control plane based network management testbed towards a cloud based software centric, highly dynamic & configurable, virtualized customer experience management testbed. The academic research community has already made significant contribution to software defined networks (SDN) in the data center and IT provider context.

The fundamental issue network operators care about is to increase network access and capacity. They are not keen on adopting current generation of wireless research technology due to the inflexibility of their current hardware based infrastructure and the high startup spectrum costs associated with new standards recommended frequency.

The research community is hampered by lack of access to low level network and MAC layer information that makes the cellular network frequency research hard to publish and they tend to revert to using open source 802.11X technologies.

The equipment manufacturers are always ahead of the curve to find the latest technology to commercialize and integrate with User Equipment. The recent drive in 5G standards has been led by Qualcomm and Ericsson.

3.3.2 Recommendations

Universities and cities are rich environments for researchers to test out new architectures, protocols and applications. This report hopes to foster collaboration between industry and academia by proposing solutions to the network operator issues of access and capacity. We propose the use of University WiFi networks as network offload points to improve user experience when operator network is not available. We propose industry partnerships with universities and cities funded by CC-NIE and CC-IIE grants to leverage the gigabit fiber backhaul networks, cloud infrastructure and compute clusters for distributed control plane network function virtualization (NFV).

To deploy large scale testbeds, we propose the starting point or epicenter as universities, which would be classified as containment access facilities (CAF), where a large deployment of programmable RF Base stations connected via gigabit campus network allows industry and researchers to have complete access to a managed and virtualized network deployment. These CAF's can be dedicated to different NFV infrastructure from cloud based core network infrastructure to distributed OpenFlow based control plane controllers at different universities/cities.

However, this testbed might be used only after initial experiments have been conducted at existing isolated facilities such as Idaho wireless range and ORBIT lab. The investigation of such collaborative mechanisms to mutually benefit both parties is the main recommendation.

3.4 Outreach and Education

3.4.1 Introduction and background

Penetrating into every corner of our society, wireless networks are foundational elements of the IT nervous systems today, and they are expected to be a basic enabler of transformative solutions of many societal challenges in the coming decades. For instance, wirelessly connecting and coordinating road vehicles can eliminate up to 90% of the accidents today, which implies more than one million people's lives saved per year across the world as well as significant reduction in accidents-induced congestion and economic cost; wirelessly connecting first-responders with one another, the environment, and control centers also enable effective, quick response to emergencies.

Enabling a wide range of innovation activities from model building, algorithm design, system prototyping and deployment, as well as user adoption, research infrastructures serve as a basic tool for engaging different stakeholders of wireless networking, including government agencies, industry, researchers, application developers, end-users, and the general public. While making sure that wireless networking research infrastructures enable transformative wireless networking research, we need to make sure that we fully leverage the catalyzing roles that research infrastructures play in helping transition basic research to real-world adoption and impact, including impact on education which will enable continuous innovation in wireless networking and the associated application domains. Given the broad society impact of wireless networks, different communities have developed a wide range of wireless networking or networked infrastructures, and we shall fully leverage the synergy between different communities in fostering an infrastructure ecosystem for accelerating innovations in a wide range of research, engineering, and business activities, especially in those areas (e.g., transportation, energy, homeland security) facing urgent grand challenges today. In transportation, for instance, every day we saved in realizing the vision of connected and automated vehicles implies over 100 people's lives saved per day in USA alone!

3.4.2 What is not possible with today's infrastructure

Considering the broad society impact of wireless networking, the inter-disciplinary nature of emerging wireless networked applications (e.g., those in wireless networked cyber-physical systems), and the urgency of realizing the many envisioned wireless networked solutions to societal challenges such as those in safety and sustainability, it is important for different communities of next-generation wireless networked solutions to work together in tackling the challenges in research and the transition of research results to practice and in fully leveraging the enabled opportunities of understanding networked systems in real-world settings. In particular, it is important to have research platforms and infrastructures that can enable different stakeholders to innovate on the same research infrastructure, achieving their individual as well as collaborative research missions.

3.4.3 Recommendations and next steps

It is recommended that the NSF community reach out and explore synergy with other agencies and organizations in domains of the national priorities, including but not limited to the following:

- Transportation
  - DOT/Volpe labs
  - DOT Safety Pilot deployment in Ann Arbor, Michigan
  - DOT connected vehicle testbeds in Michigan
  - Communities of (high-speed) trains
- Disaster response
  - DHS: OEC, OCIA, MPPD, FEMA, border patrol
- Public safety
  - FirstNet: DOC, NIST (Boulder)

It is recommended that we encourage different modes of interaction between communities of wireless networking infrastructures:

- With government agencies & industry
  - Explore how NSF networking infrastructures can contribute to missions of other agencies, their associated communities, and the industry
  - Explore how to leverage existing non-NSF infrastructures (e.g., those by DOT and the automotive industry) to support basic research by the NSF community
  - Explore mechanisms for different communities to coordinate research infrastructure initiatives for concerted progress across multiple disciplines in domains of national priority
- With research communities
  - Organize training/outreach workshops to the broad research communities, as exemplified by activities of GENI and NSFCloud
- With developer communities & end-users
  - Organize community-driven “challenges”, as exemplified by initiatives of US IGNITE and NIST Global City Teams Challenge
  - Develop infrastructures that engage developers and end-users in creative manners

For establishing research infrastructures that are sharable across different communities (who might be traditionally segregated in the innovation process), it is recommended that cross-community steering committees (e.g., NITRD at the federal level) coordinate long-term visions about inter-disciplinary initiatives that foster cross-community innovation and collaboration, and it is recommended that we investigate mechanisms for slicing/virtualizing research infrastructures.

It is recommended we pursue a wide range of education activities that are enabled by cutting-edge research infrastructures. Valuable activities include, but are not limited to, hands-on labs that leverage large scale, heterogeneous networking resources (wired, wireless), running MOOCs over infrastructures, and using catching platforms such as UAVs, robots, and smart phones to reach out to K-12 students and teachers.

3.5 Operations

This group was tasked with looking into operational and management issues for the wireless edge community resources under the assumption that these are reasonably large and complex and utilize variety of equipment. Some of the aspects that affect operations and management of the community resources are: scale, complexity and variability of equipment, types of users (both wireless users and experimenters) and types of services and experiments that are to be provided and managed by the future wireless edge testbeds with special emphasis on city-scale cellular testbed deployments.

3.5.1 Introduction

During the discussion, the group identified the following issues related to operations:

1. **Budgeting lifecycle**: Most grants are focused on designing and building testbeds with very little focus on operations and maintenance. This is especially true for large testbeds that have serious maintenance issues once they are fully operational. Also, shorter duration grants like CRI typically have very short production lifecycle given that most of the time is spent on development and initial deployment.
2. **Governance policies**: The issues of access control and scheduling can cause serious conflicts in the community. This is potentially even bigger problem if testbeds are jointly operated by wireless operators and academic institution which is highly likely to be the case in certain cases (e.g. city-scale deployments).
3. **Operations and management**: Typical academic institutions are not well suited to manage a large number of resources since such operation requires skilled and dedicated professional staff. Costs increase super-linearly when labor hours are added (support persons, maintenance and repair, etc). One of the questions related to operations of city-scale experimental wireless networks was whether there are incentives for existing operators to operate such a network in parallel to their production networks.
4. **Devices**: The number and types of existing opt-in devices is fairly and prohibitive for any city-scale testbed to support. Most existing testbeds use a small number of selected devices in order to minimize development efforts (applies mostly to phones but is true for other devices that people might use in city-scale deployment like modems used in laptop computers). The one trivial solution of asking users to carry dedicated experimental devices was seen as to restrictive and unlikely to be effective given that most real users insists on using their personal device. Effective solutions should also support multiple uses (i.e. be usable in individual testbeds, city-scale deployments and even larger MVNO environments).
5. **Users**: The group discussed the problem of how to incentivize real users to use the experimental infrastructure. Conflict between getting "real users" needs and experimental platforms needs to be addressed (i.e. possible reboot or reconfiguration of infrastructure while in use by opt-in users).

6. **Experimental Platforms**: The number of different platforms and corresponding frameworks is overwhelming typical users. Similarly, new testbeds are quite often allocating significant development efforts to introduce new management framework/platform without clear differentiation when compared with the existing ones.

3.5.2 Recommendations

- The operations of existing and future experimental resources need explicit funding line items for all phases of the project and need to clearly identify ate least these three phases and their duration: a.) design and deployment phase, b.) operational phase and c.) winding down ("sunset") phase.
- Certain awards (an in particular CRI) need to be extended beyond the typical 3 year cycle as it takes at least 1.5 to 2 year to develop and deploy a community resource (testbeds, software, etc) which leaves very short time for the actual operations.
- Each award needs to show that the resource is producing research. The following performance metrics were identified as relevant: number of conference and journal paper citing the resource usage, number of users and number of hours/experiments.
- Current NSF GENI usage policy that allows anyone with a current Shibboleth academic account from a participating university should be adopted as a default governance policy for the US academic institutions.
- Clear guidelines from NSF are needed regarding international access (especially since there might be export restrictions involved for certain resources).
- Proposals with industrial participation in deployment, operations and maintenance (e.g. with participation of equipment vendors or network operators) need to have explicit academic/research access and governance policies outlined.
- Large non-federated community resources need centralized, professional management and staffing for continued smooth operation especially if it involves regulatory and legal compliance. This is essential for city-scale operations and, in addition to managing complexity and sophistication of the deployment, saves costs and leads to consistent support.
- Proposals need to set expectations up front and clearly define interactions between the actors in development and deployment (government partners, companies and academia) as well as interactions between "ordinary" users and experimenters.
- One of the main requirements for a city-scale cellular experimental infrastructure is the seamless integration and support for dual use and most importantly support for the mandatory functions like 911, voice calling, etc. One possible solution is in

customization of dual SIM devices that are becoming available and could also be used as part of MVNO arrangement to extend coverage and incentive for regular opt-in users.

It is clear that there is no one-size-fits-all solution to the experimental testbed issue; there will always be one or more features a researcher needs that one testbed will support that others may not. The community needs to agree on a common baseline framework for testbed deployment and thus reduce user training and support costs as well as simplify development and deployment of custom testbeds.

3.6 Spectrum Policy, Privacy and Security

3.6.1 Background

SPECTRUM POLICY

Spectrum is a valuable shared resource and it is important that it is used efficiently and allocated in a way that fosters competition. However, this is not the case today – access to spectrum poses a huge barrier to entry. Furthermore, the licensed spectrum is inefficiently used.

The group discussion revolved around the use of licensed spectrum, new bands of unlicensed spectrum, pricing, and policy vs. technical solutions.

PRIVACY

The discussion focused on anonymization, and privacy vs. utility of datasets.

This is a bigger issue in the measurement community. What is the information we should collect to do research while not violating privacy? If we use a testbed, what are the rules for collecting information about users, anonymizing and sharing the data?

SECURITY

Security is important for any infrastructure. What is specific to the wireless testbed we envision? The consensus was that there will be higher likelihood of bugs and vulnerabilities being introduced, for the following reasons.

- Monitoring real mobile devices and infrastructure in urban areas is more open to threats (e.g., a rogue base station or malware on the phone) than previous closed wireless testbeds.
- The infrastructure will be programmable and thus more vulnerable. Programming will be typically performed by graduate students, which makes it more likely to be buggy.

3.6.2 What is not possible with today's infrastructure

SPECTRUM POLICY

In the past, experimental licenses have made spectrum access possible but painful. Participants reported their own experience with such licenses. Many times incumbents will shut down use because they can. FCC has simplified experimental license process, but the burden is on the user to comply with the legacy owner.

There are new bands of unlicensed spectrum coming up but a major challenge is that equipment is not yet available for non-commercial bands.

Pricing is a major barrier today. Participants debated whether spectrum should be expensive (after all, it is an important shared resource) and whether a recurrent/per-use fee would be appropriate on top of the one-time fee. Spectrum has often been bought cheap only to be resold later without actually being used. In summary, the issue seems to be more that of a fair allocation to foster competition and efficient use, rather than high price itself.

3.6.3. Recommendations and Next steps

SPECTRUM POLICY

The future seems to hold easier access to spectrum:

- Reuse of TV Whitespaces through simple database access
- New band at 3.5 GHz. Proposed rules have significant flexibility. International licensing of band for LTE makes equipment available
- Unlicensed LTE: Potential for easy access and devices, but limited to LTE.

The main recommendation was to foster competition:

- Enabling innovation and new entrants
- Need to protect incumbents
- Some policies implemented through SAS
- Allocation can be on short time scales

PRIVACY

Good practices for the use of the wireless testbed include the following:

- Anonymize datasets before making them publicly available. This is a common sense recommendation, but also a technically challenging one. Data anonymization is a research area on its own.
- Provide a clear separation between those who collect the data and those who analyze the data, for example via a limited API. This is similar to current practices at Google: a software engineer does not have direct access to data, but needs to use a particular API; for example, the engineer cannot narrow down the exact location of a user, only within some coarse distance.
- The collected data should stay within the testbed, while processed/sanitized versions of the data should be approved before release. This is similar to standard practices followed by research labs releasing their data. For example Sprint Labs and CAIDA collected detailed packet traces but made available anonymized/aggregated datasets.
- Provide incentives for users to opt-in the data collection.

What can NSF do to facilitate the process of data collection?

- Maybe NSF can help establish a template IRB for the use of the testbed. The participants discussed their experience with dealing with IRB. There is no uniformity across institutions. In order for IRB to be efficient, there should be one master agreement for a set of experiments and not one IRB per individual experiment. IRB is needed for those collecting the data and making them available. IRB may not be needed for those analyzing the data.
- Is a carrier exemption possible?

The participants tried to identify what privacy issues are unique or amplified in the case of the new wireless edge testbed:

- The wireless testbed we envision should involve real users using their personal devices in urban environments. The data collected on people's phones (location being only one of many) are much more sensitive than data that the networking community has been traditionally collecting so far (typically not directly linked to an individual). We may also monitor other personal devices beyond phones, such as RFIDs or sensors attached to a human, in which case the information is even more sensitive.
- When measurements are collected in the wild (in the city, in cars etc), non-testbed users will be implicitly or explicitly monitored as well. We should collect data for users who

opt-in and we need to drop data for users who did not agree to participate. However, this may not be sufficient, as information can be inferred even for users that are not explicitly monitored. In fact, we do not know a-priori the privacy implications of our data collection.

SECURITY

One recommendation was to look into the lessons learnt from DETER, to understand how to limit the effect of security threats within the testbed.

Another recommendation was to try to understand what operations users would be comfortable with on their own devices. For example, would users prefer having their data monitored or having an app run on their phone? Having multiple OSes or slices on a phone might help in separating what is acceptable for users to monitor from the rest of their activity.

3.7 Open Source Research Platforms

3.7.1 Motivation

The group was assigned the task of assessing the role of "open" platforms in facilitating next-generation research at the wireless edge. To begin, we define several types of openness that we use to categorize existing platforms:

- Open Access: you can easily program the platform.
- Open Source: source code for the platform is available. Frequently required by open access.A well-documented interface to the platform is available.
- Open APIs: are facilitated by open source.

As an example applying these categories to an existing component of today's wireless infrastructure, the Wifi access points (APs) provided by a wireless testbed would be:

Open access if the APs themselves could be reprogrammed by the experimenter, allowing all functionality provided by the device including its API to be altered;

Open source if the source code for the AP was available and provided to experimenters;

Open API if the APs provided a well-documented API allowing certain parameters or features to be changed.

Dependencies exist between the various types of openness but it is usually up to those maintaining the infrastructure to determine what types of openness will be provided to experimenters. For example, a wireless testbed could provide the source code for its routers to help experimenters understand their behavior (open source) but not allow experimenters to reprogram the routers (not open access) and only provide an API to control parameters such as channel assignment during the experiment (open API).

Current and Next-Generation Community Infrastructure Capabilities

The discussion of open platforms focused on how to utilize them to create community infrastructure testbeds enabling research at the wireless edge at scales currently not possible. Multiple components to constitute the wireless edge were considered, including:

- Mobile clients: Including both smartphones and other kinds of mobile devices that would utilize mobile data networks.

- Access points: We label any device providing connectivity to mobile clients as an access point, regardless of the particular protocol stack it supports. An example would be a 4G~LTE eNodeB.

- Backbone or core network: We consider the private network connecting access points to the global Internet to be part of the wireless edge.

- Cloud computation: Finally, we consider flexible storage and computation resources required by access point protocols or services running within the wireless edge as part of the wireless edge. An example would be computational resources required to implement packet filtering.

The discussion of current and next-generation community infrastructure is divided into three categories: what is already available today, what is possible given a modest amount of infrastructure investment, and what would require a great deal of additional investment. In many cases, today's experimenters lack access to capabilities already implemented by open platforms. For example, while hardware platforms supporting software-defined radio have been available for many years, we are not aware of an existing community testbed allowing researchers to utilize this capability. Deploying these existing open platforms at larger scale will improve on today's capabilities without requiring significant additional investment. Improvements require more substantial investment usually require both developing and deploying new open platforms.

Current Testbed Capabilities

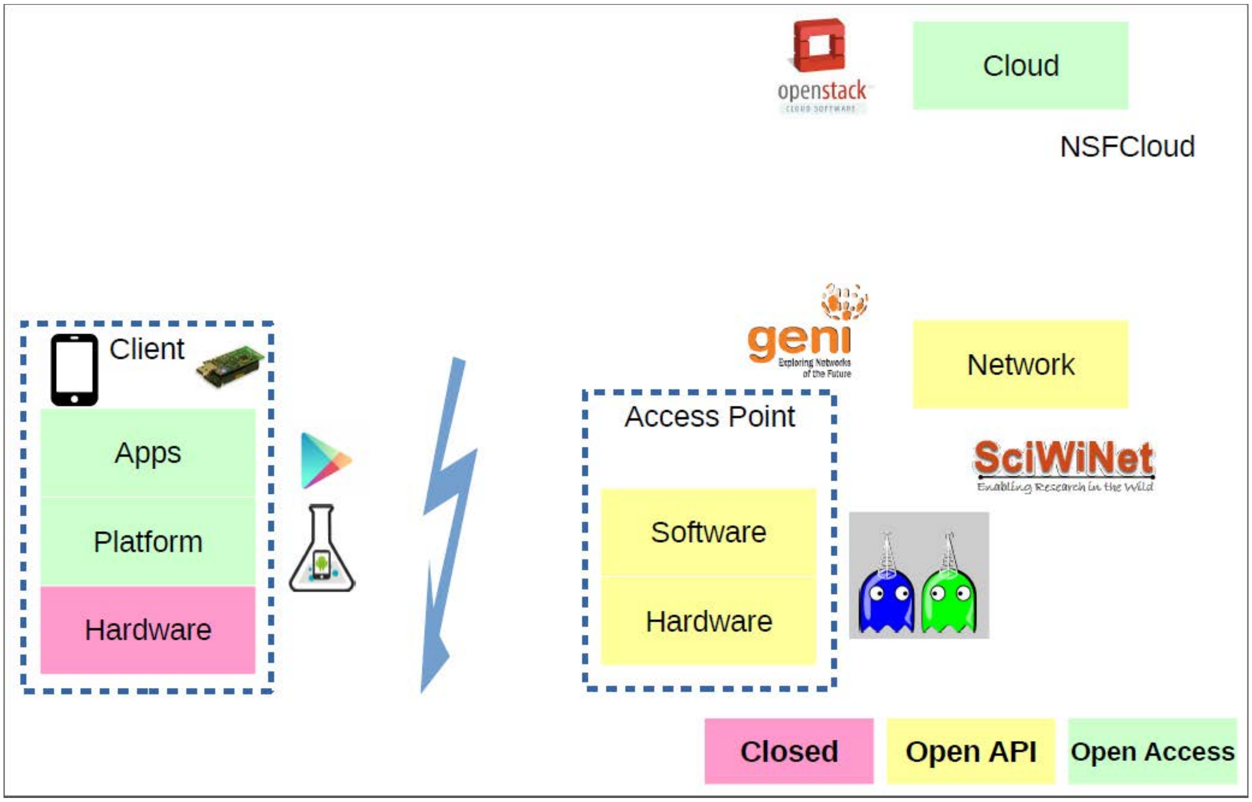

FIGURE 2

Open Access. We believe that open access platforms already exist for the cloud---including the OpenStack project and the NSF Cloud initiative---and that this problem has already been solved by other communities. On mobile smartphones, existing app marketplaces make open access available for smartphone apps, and the PhoneLab smartphone platform testbed represents community infrastructure providing open access to smartphone platforms as well.

Open API. Multiple network testbeds including GENI, PhantomNet/EmuLab and SciWiNet provide APIs allowing researchers to access networking components, including reconfiguring private network resources and access point software and hardware.

Closed. The only component completely off-limits to today's researchers is the low-level smartphone hardware. However, this has a significant impact on the ability to perform research on the wireless edge. For example, deploying experiments that alter low-level aspects of communication between mobile smartphones and wireless access points requires the ability to modify smartphone hardware or device drivers.

To summarize, while current community research infrastructure is promising, there are multiple limitations that could be addressed given a modest amount of additional investment.

Future Testbed Capabilities Requiring Limited Additional Investment

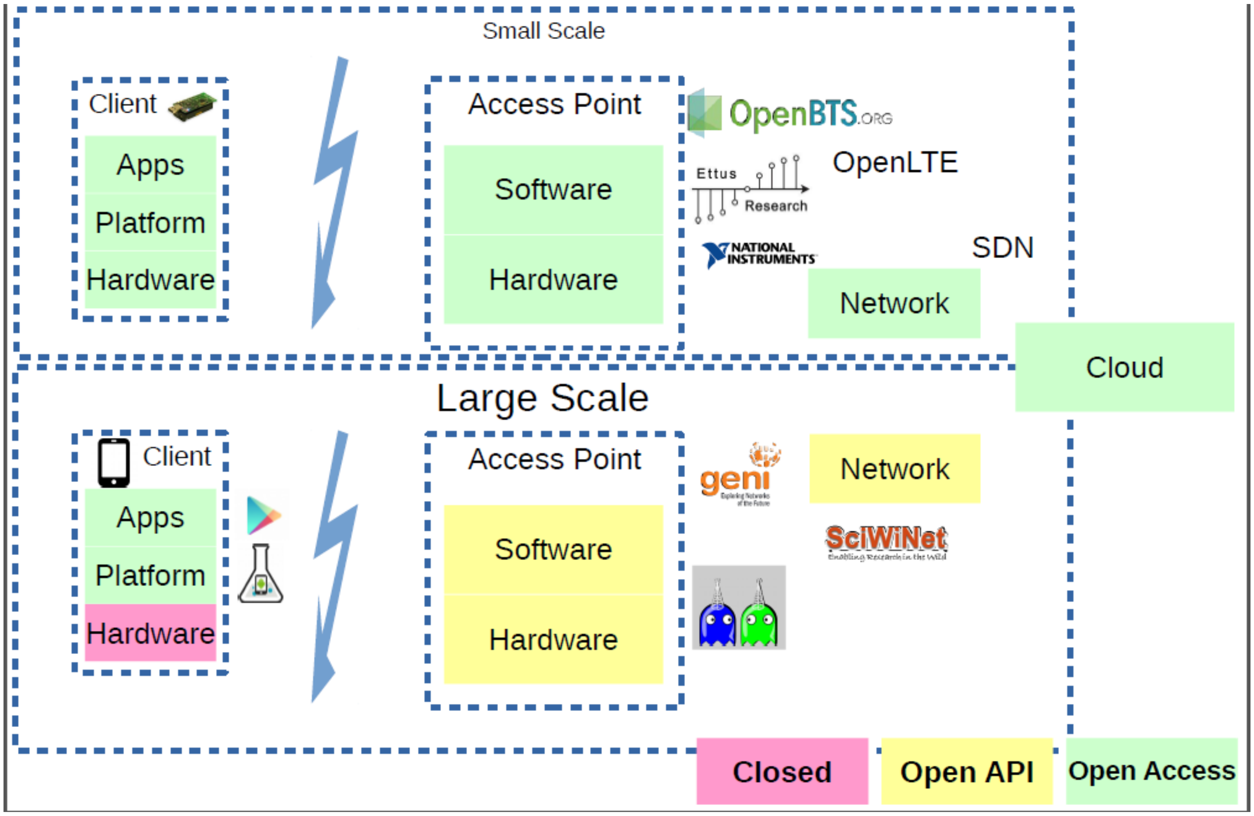

FIGURE 3

Figure 3 provides an overview of the community infrastructure possible given a small amount of additional investment, primarily through deployment of existing open platforms as part of new community testbeds. Note that we separately consider what is possible at both small (city) and large (nationwide) scale. This is because certain levels of openness require deploying and operating research infrastructure, and there are certain types of research infrastructure that will never be deployed at a nationwide scale. Specifically, wireless edge researchers will never deploy a mobile data network at the scale of those deployed by existing cellular providers. Complementary small- and large-scale deployments also helps satisfy users who might want to participate in small-scale experimental networks while not losing wide-area network coverage. For example, smartphone users who are willing to connect to experimental APs while on campus will still want to use the same device when they return home or while traveling.

At large scale the situation is similar to what is already deployed today, with additional investment serving to further deploy existing testbeds. For example, the SciWiNet model of using a Mobile Virtual Network Operator (MVNO) to provide more detailed access to internal mobile data network operations scales easily as it relies on partnering with existing cellular operators. In contrast, at a small scale more is possible. In particular, software-defined networking (SDN) and software-defined radio (SDR) can be used to provide open access to backhaul networks and access points (OpenBTS, OpenLTE, USRP), respectively. However, programming tools simplifying the process of programming SDNs and SDRs remains an open challenge and may be facilitated by the inclusion of these devices into public testbeds.

With mobile devices, we distinguish between commodity devices such as smartphones with energy and form-factor constraints and other types of devices using the mobile data network such as car-mounted devices. Experimentation with devices without energy or size constraints is much easier and facilitated by software-defined radio and the ability to deployed experimental hardware at small scales. However, access to low-level features of commodity smartphones remains a challenge.

Future Testbed Capabilities Requiring Significant Additional Investment

The main challenge to wireless edge research that cannot be solved through additional deployment of existing open platforms is access to smartphone hardware. This could take multiple forms, including access to source code for the mobile data radio device drivers for existing smartphone designs or the ability to produce research smartphones incorporating novel radio technologies or other new features. We suspect that addressing this challenge will be expensive but believe it may be worth determining the cost to design and manufacture a reasonable number of smartphones with a research-oriented design. It may also be possible to leverage emerging vendor interest in more modular smartphone platforms which would allow designs integrating novel features.

3.7.2 Summary and Recommendations

To summarize, many of the open platforms needed to create testbeds enabling research at the wireless edge already exist today but have not been deployed at appropriate scale and made available to researchers. Remaining blockers to completely open-access research infrastructure include scale and form factor. To address the scale issue we suggest complementary deployment of small-scale completely-open community research infrastructure alongside large-scale ``as open as possible'' research infrastructure. However, hardware enabling fully-programmable user-facing mobile systems remains a problem. Until this challenge is addressed the ability to experiment at both ends of the wireless edge commonly connecting people to the wireless data network will be limited.

4. Recommendations and Summary

All of the recommendations put forth in the sections above can be summarized into the following set of broad observations and recommendations.

Only testbeds operating at scale can expose challenges that occur at scale: Numerous design challenges of end-to-end systems manifest themselves when they get deployed at their true scale. Our current network densities -- in client devices per unit area -- are expected to grow in the next decade by two orders of magnitude or more. That level of scaling threatens to overwhelm or break our current wireless networks, and motivates us to revisit conventional design principles, and explore new wireless network technologies capable of supporting dramatically higher client node densities. The research community needs a deeper understanding of how current operational architectures, protocols and algorithms scale. For example, how mobile networks congest in presence of large volumes of video flows demanded by users, and how such congestion can be mitigated and managed, can only be researched with an infrastructure that supports such occurrences.

Adopt large user communities: A potential path to creating infrastructure at scale is to adopt large communities, e.g., a large campus with tens of thousands of users, expanding to the greater campus area and its users, and eventually to an entire city. A potential focus of such adoption would be to seek and identify compelling needs and challenges that could be local to that user community and evaluate how an infrastructure and its experiments could provide novel solutions to that community. Often university campuses with its complex set of user requirements, the need to support diverse users and usage scenarios, and with their great synergies between research and IT personnel, could provide a compelling microcosm to seed such communities.

Experimentation with usage scenarios at scale requires inter-operation with common user equipment, devices, and applications: At scale experimentation is not just about the infrastructure. It requires users and usage scenarios as well. To allow such experimentation, it is important to ensure that common end user devices, and their applications can work with experimental components. Hence, it is important to consider edge technologies and APIs that allow easy inter-operation of experimental concepts with common end user devices.

Engagement with industrial partners: A large scale infrastructure is not feasible without active engagement of industrial partners of different types --- operators with deep knowledge of operational aspects of such infrastructure, equipment vendors that are responsible for commercial grade hardware and software systems that form the infrastructure, and even content providers that generate most of the content transported over these infrastructure. Engagement with industrial partners, however, needs to be symbiotic. There should be alignment between the objectives of industrial partners and goals of the research community interested in the infrastructure. Industrial partners could potentially contribute hardware and software resources, personnel time, and know-how in operating such infrastructure. In return, they could benefit from new technologies that emerge from the research and in actively participating in evaluation of new concepts and technologies.

Engagement of local agencies: Local agencies, e.g., city governments, campus administrations, local utilities, and so on, can also be a valuable participant in the construction of such infrastructure. Wide area wireless infrastructure involves many logistical components. One example is access to sites, such as rooftops, roadsides, and utility poles, where equipment needs to be mounted. Having active involvement of local agencies would be critical to the success of such infrastructure. Beyond basic support, it would be beneficial if the local agencies also find benefit in participating in and in using such infrastructure to meeting their own goals.

Engagement with broader set of federal agencies: Many federal agencies that are focused on delivering new and improved services in multiple areas of national priority --- homeland security, healthcare, education, transportation, effective broadband, and much more. An infrastructure of the proposed form can potentially be leveraged in both experimenting with and in delivering applications in many of these domains. Partnership with such a broader set of federal agencies in creating such an infrastructure and finding mechanisms to share it also has great merit.

Regulatory support: Over the years, federal agencies such as the Federal Communications Commission (FCC) have provided regulatory mechanisms by which researchers can seek experimental spectrum licenses to experiment with new technologies under specific guidelines. Experimentation using large-scale wireless infrastructure as proposed, that engages local communities, industrial partners, and others, may benefit from additional regulatory support to promote such engagement. For instance, regulatory mechanisms that would encourage various forms of physical sharing the proposed wireless infrastructure --- for experimental use, to deliver commercial-grade services, and to support civic applications --- could be highly beneficial (through techniques such as slicing, virtualization and spectrum sharing techniques).

A strawman approach and potential next steps

Creating large-scale wireless experimentation infrastructure is a challenge and requires synergistic efforts from multiple partners. Based on our discussions in the workshop above, one could approach this effort in multiple stages. At the seed level, we could create “model campuses” where we adopt an entire campus community and its needs and build a wireless experimentation infrastructure replete with end-to-end components --- base stations, programmable edge computing platforms, dedicated backhauls, and cloud services. It needs effective partnership between a host campus, campus IT personnel, operators and equipment vendors, and researchers. The proposed infrastructure could be designed inherently to be virtualizable into multiple networks --- one for experimental use, one for delivering commercial-grade services, and a third and a fourth for civic applications, with users being allowed to move across these virtual networks. If multiple such model campuses can be instantiated, each with potentially unique requirements and challenges, the best practices of addressing these challenges could be exposed to the broad community. From these seeds, it is possible to organically expand deployments and communities to a greater campus area, and eventually to a “model city”; the latter requiring partnership and assistance of the local and federal agencies. Each campus or city could be encouraged to focus on a different set of application domains, e.g., in healthcare, transportation, public safety, homeland security and more, and the suitable partnerships would be meaningful in each local context. Similarly, each campus and city could also be encouraged to pick one single operator as a partner to create this shared infrastructure.

Finally, this first workshop explored some of the major issues around the needs, challenges, and opportunities of large-scale wide-area wireless infrastructure for experimentation; it appears that further discussions are necessary to help create a more concrete plan. Critical topics for further discussion include incentives for partnerships between research institutions, industrial partnerships, and local agencies, example architectures of shared infrastructure, and synergies with efforts of other federal agencies.

APPENDIX

<FILL AGENDA>