Final report from the

NSF Workshop on Future Directions in Wireless Networking

held in Arlington VA on November 4-5, 2013

Workshop chairs

Suman Banerjee (UW-Madison) and Dapeng Oliver Wu (U Florida)

Workshop scribes

Xiaojun Lin (Purdue U) and Xinyu Zhang (UW-Madison)

Workshop participants

Tarek Abdelzaher, University of Illinois, Urbana-Champaign
Salman Avestimehr, Cornell University
Victor Bahl, Microsoft Research
Suman Banerjee, University of Wisconsin-Madison
Stefano Basagni, Northeastern University
Doug Blough, Georgia Institute of Technology
Robert Brodersen, University of California Berkeley
Milind Buddhikot, Alcatel-Lucent Bell Laboratories
Joseph Camp, Southern Methodist University
Andrew Campbell, Dartmouth College
Guohong Cao, Pennsylvania State University
Alberto Cerpa, University of California Merced
John Chapin, DARPA
Yingying Chen, Stevens Institute of Technology
Kaushik Chowdhury, Northeastern University
Andrew Clegg, NSF
Joseph Evans, University of Kansas
Yuguang (Michael) Fang, University of Florida
Darleen Fisher, NSF
Deepak Ganesan, University of Massachusetts Amherst
Jie Gao, StonyBrook University
Mario Gerla, University of California Los Angeles

Erwin Gianchandani, NSF
Shyamnath Gollakota, University of Washington
Abhimanyu Gosain, Raytheon-BBN
Dirk Grunwald, University of Colorado
Marco Gruteser, Rutgers University
Qi Han, Colorado School of Mines
Tian He, Univeristy of Minnesota
Thomas Hou, Virginia Tech
Rose Qingyang Hu, Utah State University
Farnam Jahanian, NSF
Sneha Kasera, University of Utah
George Kesidis, Pennsylvania State University
Edward Knightly, Rice University
David Kotz, Dartmouth College
Robin Kravets, University of Illinois Urbana-Champaign
Bhaskar Krishnamachari, University of Southern California
P.R. Kumar, Texas A&M University
Li Erran Li, Alcatel-Lucent Bell Laboratories
Xiang-Yang Li, Illinois Institute of Technology
Xiaojun Lin, Purdue University
Xin Liu, University of California Davis
Chenyang Lu, Washington University in St. Louis
J. Bryan Lyles, NSF
Keith Marzullo, NSF
Tommaso Melodia, State University of New York, Buffalo
Eytan Modiano, Massachussetts Institute of Technology
Prasant Mohapatra, University of California Davis
Aarti Munjal, University of Colorado Denver
Matt Mutka, Michigan State University
Thyaga Nandagopal, NSF
Srihari Nelakuditi, U South Carolina
Guevara Noubir, Northeastern University
Katia Obraczka, University of California Santa Cruz
Joseph Paradiso, Massachussetts Institute of Technology
Craig Partridge, Raytheon-BBN
Lili Qiu, University Texas at Austin
Ted Rappaport, Polytechnic Insitute of New York University
Dipankar Raychaudhuri, Rutgers University
Phil Regalia, NSF
Kui Ren, State University of New York, Buffalo

Dennis Roberson, Illinois Institute of Technology
Romit Roy Choudhury, University of Illinois Urbana-Champaign
Sumit Roy, University of Washington
Anant Sahai, University of California Berkeley
Mahadev Satyanarayanan, Carnegie Mellon University
Anna Scaglione, University of California Davis
Henning Schulzrinne, FCC
Sanjay Shakkottai, University of Texas at Austin
Kang Shin, University of Michigan
John Smee, Qualcomm
Min Song, NSF
Oliver Spatscheck, AT&T Labs
Aravind Srinivasan, University of Maryland, College Park
Peter Steekiste, Carnegie Mellon University
Lakshmi Subramanian, New York University
Karthik Sundaresan, NEC Labs
Zhi (Gerry) Tian, NSF
Nitin Vaidya, University of Illinois Urbana-Champaign
ArunVenkataramani, University of Massachusetts Amherst
Bing Wang, University of Connecticut
Wenye Wang, North Carolina State University
Steven Weber, Drexel University
Kamin Whitehouse, University of Virginia
Dapeng Oliver Wu, University of Florida
Jie Wu, Temple University
Jiang (Linda) Xie, University of North Carolina at Charlotte
Guoliang Xing, Michigan State University
Yaling Yang, Virginia Tech
Michael Zavlanos, Duke University
Hongwei Zhang, Wayne State University
Junshan Zhang, Arizona State University
Xinyu Zhang, University of Wisconsin-Madison
Younguang Zhang, Microsoft Research
Zhi-li Zhang, University of Minnesota
Xia Zhou, Dartmouth College
Ty Znati, University of Pittsburgh
Gil Zussman, Columbia University

1. Introduction

Wireless communication networks are closely intertwined with the very fabric of both civilian and military environments, and their continued growth and well-being is critical to the success and health of our nation. Today's wireless communication industry worldwide is a trillion-dollar business that represents a substantial fraction of the global economy. Over the past two decades, it has had a transformative impact on our society and has revolutionized almost all aspects of human interaction. Through the use of smartphones, personal tablets and increasing number of wireless-based sensors in our environments, wireless networking systems touch every aspect of human lives today. Nonetheless, with exponential spread of mobile wireless devices, continued improvements in spectrum utilization, resource allocation, hardware miniaturization, and development of new renewable energy sources, we find ourselves still very much in the infancy of the development of such networks, which have the potential to grow and become orders-of-magnitude more sophisticated.

While past successes in this domain have been quite significant, wireless-based technologies can prove to be astonishingly transformative in future applications. The following are a few examples:

Operating room of the future. Critical surgeries in the future will be done by robotic surgeons. Human doctors in the emergency rooms of the future may monitor progress and may communicate, using wearable wireless devices like Google Glass, to other medical experts in remote locations. Continuous, robust, and high-bandwidth wireless access to streamed content will be crucial for the success of such operations.

Zero-accident transportation systems. In the US today, more than 30,000 people are killed every year in traffic accidents. Envisage a smart transportation infrastructure that will prevent accidents and save lives. With the emergence of autonomous vehicles, tomorrow's transportation systems will themselves form automated networks for moving people and goods from place to place. Sensors in these autonomous vehicles can observe the conditions of the road and surrounding vehicles, and can exchange information with other vehicles. Trustworthy wireless connections between vehicles and with the surrounding infrastructure will be needed to enable real-time collision avoidance and congestion management. Those wireless links must operate at high mobility (i.e., highway speed), with low latency, high availability, and high resistance to adversarial attack.

Cognitive assistance. Emerging wearable wireless devices may assist cognitively challenging tasks, such as recognizing faces, navigating through a building complex, locating misplaced objects, learning how to use new tools or instruments, or changing behaviors to improve personal health. Such a device is not only helpful for individuals with cognitive deficiency, but also can benefit everyone. It would be able to sense the current context, through built-in sensors as well as other sensors and devices that are wirelessly-connected around the human body and the neighboring environment. The intense information and computational requirements also require such a device to have a wireless connection to the cloud, where external information can be integrated and where more complex analysis can be accomplished. This, in turn, desires extremely low-power and low-latency communications, and in some cases high-throughput connections, for sharing rich local sensor data with the cloud and for streaming content to overlay local reality with augmented information.

Disaster response and recovery. Mobile phones, robots, and other electronic devices can be extremely useful to facilitate response and rescue operations in disaster areas. Note that in a disaster scenario, these devices must be able to operate even when the infrastructure is disabled or under heavy network demand. Such application scenarios challenge the notion that the network and the cloud are always available and demand new system designs, possibly including rapidly-deployable replacement infrastructure or even infrastructure-less operation.

Cyber warfare. Current trends in network design leverage heavily on data centers for data storage and analysis. While this trend is driven by many factors including cost, efficiency, and performance, data centers also become prime targets in warfare and the incapacitation of data centers would cripple many computing services. Similarly, network points of presence, transoceanic cables, and other aspects of peacetime system designs are vulnerable during wartime. Wireless-based ubiquitous access mechanisms are making such attacks easier to mount while allowing attackers to hide their tracks. At the same time, it appears feasible to design new defense mechanisms that leverage the power of distributed, crowd-sourced observations across millions of mobile devices and users.

Ubiquitous imaging. Imaging devices are becoming higher resolution, lower power and increasingly ubiquitous. It will soon be possible to find a camera on every human body, in every room, on every street, and in every vehicle. Applications of these cameras include safety, smart-building operations, and pervasive computing. Uploading images and videos from these cameras requires high-capacity wireless links, which will challenge the assumption that underlies nearly all last-mile wireless network technologies: that upstream bandwidth is typically only a fraction of downstream bandwidth. New wireless network architecture and design principles will be required to support a large fraction of upstream traffic in an efficient and scalable manner.

Industrial and equipment monitoring. Sensors are increasingly being used for real time monitoring of equipment and industrial processes and can generate more data than current wireless networks can support. For example, the jet engines on a modern airplane generate over 2 TB of data on a single transatlantic flight. This data contains important information about preventative maintenance, but not all of the data can be streamed in real time to the manufacturer for analysis. The transmission and analysis of such data, and the extension of data-centric applications to other vehicles and equipment worldwide, challenges the emerging trend that data analysis be performed in the cloud.

In-body sensors that predict epidemic outbreaks: In-body sensors will continuously monitor health statistics of each individual and alert physicians when needed. In addition, aggregated data from the entire population can predict outbreaks of epidemics even before they occur, leading to improved prevention and management of such conditions. Such applications require in-body sensors that not only have robust wireless connectivity, but also are highly energy-efficient so that sensors can stay powered and active for the entire human lifetime.

Several other compelling applications that bear similar challenges include: (i) improving lifestyle through better behavioral health, (ii) providing robust access to record, retrieve, and analyze multi-modal content in an untethered fashion, and (iii) enabling smarter and more sustainable homes, buildings, and cities that ensure greater energy efficiency and public safety.

Clearly, wireless networking is a central cog at the heart of many of these applications and systems. Hence, it is important for the community and the National Science Foundation to have a broad vision of the key research challenges facing the wireless networking community in the next decade and beyond. Towards that end, the NSF Workshop on Future Directions of Wireless Networking was held on November 4-5, 2013 in Arlington, Virginia. Critical challenges in wireless networking were discussed, and a vision for conducting fundamental and inter-disciplinary research in this area was articulated.

2. Workshop Objectives and Structure

The overarching goal of the workshop was to help chart a broad vision for the future of wireless networking research. At the workshop, attendees discussed and debated the critical challenges in wireless networking, and articulated a vision for conducting both fundamental and inter-disciplinary research in the area. The workshop participants were clearly briefed on the goal of the workshop. They were explicitly informed that the goal of the workshop was not to end with intellectual discussions, but rather to work as a team to create a common vision of the area that is presented to the community in the form of a report.

This report has been written with the intention of (i) identifying major open issues for future wireless research; (ii) exposing the research community to new and exciting inter-disciplinary problems; and (iii) stimulating far-reaching future research initiatives as well as collaborations that would help along the evolution of the wireless research community.

The area was broadly partitioned into a number of technical areas, which include:

- Wireless systems
- Wireless applications
- Wireless networking architectures and testbeds
- Controls and algorithms
- Emerging spectrum access technologies
- Metrics, measurements, and management
- Security and Privacy

In addition, the workshop also discussed various issues related to improving the quality of education in the domain, and achieving a better engagement between industrial activities and academic research.

To discuss all of these issues, a total of 105 participants were invited to the workshop, representing a broad spectrum of expertise in the wireless research community. The workshop was co-chaired by Suman Banerjee (UW-Madison) and Dapeng Oliver Wu (U Florida) who were assisted in report-writing activities by Xiaojun Lin (Purdue U) and Xinyu Zhang (UW-Madison). The participants were divided into various breakout discussion groups throughout the two-day event to allow more focused discussion on individual topics. Five experts from the community (Victor Bahl from Microsoft Research, Bob Brodersen from UC Berkeley, PR Kumar from Texas A&M University, Craig Partridge from Raytheon-BBN, and Mahadev Satyanaryanan from CMU) were requested to form a "critics' panel" to provide feedback to the participants throughout the workshop. Various other individuals were given specific responsibilities in the workshop, e.g., to serve as leaders of specific discussion groups and as scribes. 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. Research Thrust Areas

In the context of this workshop, the domain of wireless networking was partitioned into different topic areas. The topic areas are complementary and synergistic to each other.

3.1 Wireless Systems

3.1.1 Background and Prior Successes

The astronomical growth of wireless networks that we have witnessed in the past decade is a result of the synergistic advances across all elements of the wireless systems, including radio transmission technologies, network protocols, middleware, mobile devices, and applications. The notable prior successes include:

- *PHY layer*: New PHY transmission mechanisms that greatly increase bit rate and reduce interference: new multi-user and distributed MIMO techniques, interference cancellation and alignment, full-duplex radios, 60GHz channels and radios, wideband radios, new high-bandwidth point-to-point and point-to-multipoint backhaul technologies.
- *MAC layer*: New MAC protocols that challenge conventional contention/interference assumptions and interplay with the gamut of new PHY layer solutions.
- *Network layer*: New naming, addressing and routing protocols for cellular networks, mesh networks, backhaul networks and sensor networks, secure identity solutions, IP and cellular convergence, fast cellular handoffs, software defined cellular networks.
- *Transport layer*: New transport and congestion control protocols for cellular networks, sensor networks, mesh networks and wireless backhaul networks.
- *Middle-box innovations*: New cellular SDN frameworks, middle-box enhanced network services, new queuing algorithms, mobile caching protocols, mobile CDNs.
- *Device innovations*: Customizable mobile communication and sensing devices, low-power computing devices, energy scavenging devices and trusted computing platforms.
- *Mobile applications*: Explosive growth in new mobile applications for different platforms, new distributed mobile sensing applications and participatory sensing frameworks, cloud-powered mobile applications, new mobile user-interfaces through advances in speech, language, vision and gesture technologies and wearable computing.

3.1.2 Challenges and Opportunities

Despite these prior successes, the expectation for future network demand of mobile users, devices and applications is bound to increase astronomically, which will create a broad range of challenges for future wireless systems, including:

- Need for high-throughput: The growth of future wireless systems will require ubiquitous high-throughput mobile connectivity under a range of challenging network settings (urban versus rural, high device density, mobility etc.)
- Content-connectivity gap: Content growth outpaces connectivity growth creating serious resource availability problems in wireless networks in the future.
- High end-to-end latency: Network-intensive mobile applications experience high end-to-end latency that directly impacts usability and interactivity.
- Need for energy awareness: Energy remains a relatively scarce resource for mobile devices due to growing energy needs of applications and systems, intermittent power availability and limited advances in energy storage capacities.
- (Over)-dependence on the cloud: Current mobile applications heavily rely on the cloud infrastructure that any availability and fault-tolerance problems experienced in the network or the cloud directly impacts millions of end-users.
- Need for sustainable system design: future wireless systems will have to pay special attention to the growing carbon footprint of wireless networks and the energy constraints of mobile devices.

To be more specific, our discussion below will focus on two highly contrasting grand-challenge themes: *the 1000x Game* and *Stress-Free networks*. The 1000x game addresses the design of future wireless networks that are 1000x better than today across a variety of performance metrics. The Stress-free networks discussion focuses on future wireless networks that can seamlessly operate under extreme stress conditions, e.g., disaster relief scenario with increased demands on a weakened infrastructure. We next elaborate on these two themes.

3.1.2.1 The 1000x Game: designing networks that are 1000x better

How do we design future wireless systems that can provide 1000x improvement in end-to-end metrics such as throughput, latency reduction, energy reduction over current generation systems? Clearly, achieving such a level of improvement again requires advances across all elements of the wireless system. In particular, future PHY layers should at least provide 1000x capacity. Future MAC layers should intelligently handle 1000x increase in contention. Future network layers should intelligently handle 4-5 orders of increase in communications including accounting for device and service mobility. Finally, higher layer protocols need to be redesigned to be more energy and traffic demand aware. We summarize the five key underlying challenges below:

- Role of new low-layer technologies: How far can new PHY layer and MAC layer enhancements, such as multi-user MIMO, network MIMO and full-duplex, help in pushing the performance?
- Cross-layer inefficiencies: Existing systems experience large-amounts of cross-layer inefficiencies thereby negating many of the gains observed in individual layers. Can we obtain 10-100x gains purely from reducing cross-layer inefficiencies?
- Contention inefficiencies: Many wireless networks suffer from performance and predictability problems due to high levels of contention and interference, especially in large-scale, high-density environment.
- Role of SDN: Can we leverage software-defined middle-boxes that can understand the requirements of user applications and provide significant end-to-end gains by bringing content and services close to the mobile devices?
- Network-state aware design: If future wireless applications are completely aware of the underlying state of the network, how far can they adapt and optimize their behavior to observe end-to-end gains in the face of changing network conditions? How can we make applications network-aware without overwhelming them with updates regarding the network's conditions?

Next, we discuss the potential opportunities that could address the above challenges.

1000x end-to-end throughput: Achieving such improvement in throughput will require a concerted effort across PHY, MAC, and higher network layers. Promising solutions at the PHY layer include: (a) Full duplex radios that offer a possible gain of 2x; (b) distributed and network MIMO solutions that can eliminate contention and fully exploit MIMO advantage, enabling linear increase of capacity with number of antennas; (c) Highly directional transmissions that can significantly enhance throughput and spatial reuse; (d) mmWave technologies which operate on GHz bandwidth to offer substantially higher throughput. (e) Need for better cooperative spectrum sharing using cognitive radios given that spectrum is a scarce resource. Further, at the MAC layer, each PHY layer enhancement needs to be coupled with efficient protocols and primitives at the MAC layer to support intelligent contention management and interference handling. Moreover, promising directions at higher layers (above MAC) include: (a) Using a high density of small cells to increase end-to-end gains by a factor of 10x; (b) Designing future wireless routing protocols that are more channel-state aware and that can efficiently handle device mobility, provide efficient naming/addressing substrates and also intelligently move services in close proximity to devices and applications; (c)

Rethinking the design of end-to-end congestion control protocols that can achieve 1000x throughput over highly variable mobile communication channels.

1000x latency gains and the role of Software-Defined Networks (SDNs): Can we design future SDN-aware wireless network architecture that can provide a 1000x reduction in latency by intelligently moving network services closer to devices and users? Designing such systems raises several interesting questions: (a) Centralized SDN infrastructures incur high latency but is easy to maintain, while designing a completely distributed SDN that operates in close proximity to the devices raises significant distributed coordination and consistency semantics problems; (b) supporting extremely high levels of service mobility requires accurate prediction of end-user application requirements, learning of device mobility patterns and widespread replication of network services.

1000x energy reduction: How do we design future generation wireless systems that consume 1000x lesser energy, but yet meet the users' service requirements? Specific promising directions in this space include: (a) Design of domain-specific extremely low power devices (e.g., for environment monitoring and health); (b) Battery-less devices; (c) Low power communications using highly directional beams and mmWave technologies; (d) Use of small cells with dramatic reductions in power consumption levels; (e) New energy-aware protocols (e.g., medium access protocols) and OS level primitives for modular smart energy management.

Sustainability challenges: Closely tied to the energy challenges, future wireless systems are bound to face several sustainability challenges, many of which currently remain unaddressed. First, there is a need for sustainable off-grid wireless solutions given that stable grid power is largely absent in many parts of the world. Second, providing a sustainable model for universal access is difficult, given that most rural areas with limited populations and limited purchasing power have highly intermittent or no connectivity due to limited economic incentives for cellular providers.

3.1.2.2 Stress-free Networks: Systems that can seamlessly handle extreme loads

The stress-free networks discussion focused on the following theme*: Can we design stress-free wireless networks that can seamlessly operate and handle extreme network scenarios where the network load is significantly higher than the underlying network capacity*? Motivating scenarios include disaster or emergency response settings, highly crowded contexts or futuristic high-demand distributed mobile applications such as Internet of things, high-quality mobile video, mobile voting, mobile health, distributed smart energy and transportation systems. The design for stress-free networks would push wireless network design to radically new architectures that challenge many of the conventional assumptions. The key challenge is that end-users, devices and applications should gracefully adapt to changing network conditions and achieve "useful end-to-end guarantees" under extreme conditions. This question raises new, radically different problems:

Plan B networks: Under extreme load conditions, can the existing wireless network infrastructure have an automatic "plan B" mode of operation that provides completely different operational modes and guarantees for network services and end-to-end applications? A "plan B network" would automatically call for a new morphable end-to-end protocol stack and challenge many of the traditional network assumptions. For example, it would relax the security and reliability assumptions, but still provide useful end-to-end connectivity for certain applications.

Prioritization according to the value of information*:* A critical design metric in stress-free networks from an application standpoint is to characterize the "value of information" transmitted or received by an application. Stress-free networks should provide guaranteed service delivery or high priority packet delivery for critical applications and use the value of information to prioritize non-critical applications. Stress-free networks can have a direct impact on human behavior and human expectations from mobile applications, thereby providing new "social utility" models for communications.

Network agility, robustness, diversity*:* Stress-free wireless networks require efficient naming, addressing, routing, spectrum sharing and wireless cognition capabilities to quickly detect and respond failures. It should be robust in the face of changing network and link characteristics and intelligently leverage network diversity across different interfaces, service providers, topologies and link qualities.

Cloud support in stress-free networks: When mobile applications become increasing integrated with the cloud, disruption to wireless networks could potentially affect the experience of a large number of users. Can we design stress-free network that can still provide a cloud-like experience under extreme stress? Such a network would require a rethinking of the conventional mobile-cloud application model with a scalable resolution infrastructure, device-aware service mobility, and device-aware replication.

3.1.3 Recommendations

In summary, future research activities for wireless systems should focus on the following directions:

1. New PHY, MAC, and spectrum sharing technologies: We require new PHY layer technologies such as mmWave (60GHz) radios, advanced MIMO and full duplex technologies coupled with new MAC layer protocols to increase network throughput by a few orders of magnitude. Further, new models for spectrum access are needed to significant increase the spectrum reuse efficiency.

2. Design for scale: Future generation wireless systems and network infrastructure need to be designed to better scale with the potential astronomical growth in mobile devices and mobile traffic. This requires clean-slate technology solutions at the network and transport layers including scalable network architectures, scalable naming and addressing protocols, scalable routing protocols that are agile to fast mobility and robust to network changes and failures, and redesign of transport protocols that can quickly adapt to variable network conditions. Other potentially dramatic changes include: revisiting end-to-end metrics; revisiting end-to-end network layering concepts with the need for more cross-layer design; rethinking applications to be network-aware and to be adaptive to different network conditions.

3. Mobile-Cloud Integration: It is easy to forget the importance of the cloud while thinking about next-generation mobile and wireless networks. However, the rapid growth of cloud computing services in the last decade and the compelling manageability, security, and cost arguments for the cloud suggest that they are here to stay. As argued above, mobile and wireless networks can significantly enhance end-to-end user-perceived experience by intelligently leveraging the cloud. Thus, research in mobile networks must be accompanied by broader research in building massively scalable distributed cloud systems that enhance the mobile experience. In addition, we should identify the right balance of on-device/in-cloud computation that gives the maximum resilience and network usage efficiency.

4. Nanoscale networks: As wireless devices get increasingly smaller (e.g, ingestible devices used for intra-body sensing, nano-scale particle sensors that can be painted on a surface), it becomes increasingly difficult to address communication challenges and design appropriate wireless protocols for these small-sized devices.. Addressing systems challenges in this domain requires interdisciplinary research across various areas of not just computer science but also medicine, material sciences, and various engineering disciplines.

5. Design for sustainability and energy efficiency: A key consideration in system design is to incorporate sustainability as a first order design principle. Note that communication and computation already contribute to a significant part of the worldwide energy consumption. Hence, we believe future wireless system architectures should consider sustainability as a primary design metric to treat energy/power as a precious resource, to reduce carbon footprint and energy footprint, and to enhance technology adoption in the face of high variability in end-user purchasing power and connectivity.

6. Design for simplicity: There is a need for greatly simplified network architectures for mobile systems (possibly based on SDN) that are unlike what we have today, namely, a cellular network that interconnects with the TCP/IP Internet via a mind-numbing number of gateways and protocol translators. In the longer term, this multitude of protocol stacks impedes growth and innovation and increases management complexity.

3.2 Wireless Applications

3.2.1 Background and Prior Successes

As we discussed in Section 1, wireless technology has been a key enabler for an expanding range of mobile devices, including not only smartphones and tablets, but also wearable sensors, industrial machinery and health-monitoring instruments. Meanwhile, wireless devices are playing a key role in many new application domains, including personal health and wellness, real-time cognitive assistance, smart buildings, smart roads, autonomous vehicles, free-range assistive robots, and environmental monitoring -- all of which will push wireless networks to improve dramatically in throughput, latency, reliability, security, privacy, usability, and efficiency. In the past several years, real applications running on mobile devices have been driving the evolution of wireless networks, shaping the design decisions of every successful wireless technology available today, and constantly imposing research challenges at all layers of the network stack. For example:

- *Mobile computing*: The mobility needs of laptop computing helped to shape many design decisions of Wi-Fi including medium access control, network-association protocols, and dynamic frequency selection. The higher mobility levels of mobile telephones drove a very different set of design decisions for cellular communications technology, reflected not only in communication algorithms, but also in network architecture and protocols.
- *Personal peripherals*: The design of Bluetooth technology, e.g., the master-slave architecture and low-power design, was driven by the needs of peripherals such as mobile phones and ear-pieces.

- *Wireless sensing*: The limited energy budget of wireless sensors revealed many shortcomings of Bluetooth and drove the design of new low-power MAC protocols and low-power radio designs, eventually leading to the 802.15.4 and ZigBee standards.
- *Wireless positioning*: The need to determine the location of emergency callers and to present users with relevant location information has driven the design of wireless positioning systems. In particular, the need for a short time-to-first-fix led to Assisted GPS and the need for indoor localization led to Wi-Fi positioning services.

Currently, several technologies and trends are emerging and are likely to drive the next generation of wireless technologies. In particular,

- *Device diversity*: Device technology is shifting from portables (e.g., laptops and smartphones) to wearables (e.g., glasses and watches). Some predict that the smartphone will disappear within five years. Meanwhile, devices will shift from general-purpose computing tools to special-purpose devices (e.g., FitBit). As a result, application designers will need to cope with a growing diversity of devices.
- *Scale and density*: The number of wireless devices per person will grow tremendously, soon to reach 1.5 devices per person on average in the world, and later to reach 10 or even 100 devices per person. Current wireless network architectures lack scalability --- network performance degrades super-linearly with device density. Counter-intuitively, adding more infrastructure nodes, e.g., WiFi Access Points, may worsen the situation due to mutual interference.
- *Networking paradigms*: In many applications, there will be a near-permanent uplink to the cloud (or cloudlet) for pushing data from a person to the cloud for analytics. At the same time, wireless networks will increasingly be used locally, e.g. intra-building, intra-vehicle, inter-vehicle, and personal body-area networks. Data generated locally will increasingly be consumed locally (albeit with a subset uploaded as-possible to the cloud for archive or audit purposes). This trend will be particularly salient in rural, disconnected, and “hostile” network environments where general Internet connectivity is not possible.
- *Demanding applications*: Wireless networks will increasingly be used in high-confidence systems that are mission critical. They will need to become as reliable if not more reliable than wires, and they will need to provide performance guarantees. Also, many wireless networks are used in an application-specific fashion and multi-disciplinary expertise will be required for vertical system designs.

- *Machine networking*: Cameras, microphones, sensors, and crowd-sourced information will increasingly be used to drive automation and control systems, instead of just for human consumption. Sensors in such systems will generate massive data inflows, and will produce as much if not more data and network traffic than the World Wide Web. This will reverse current loads, where most data is produced in the cloud and consumed at the edge.
- *Context-awareness*: Radios are increasingly used as sensors to detect the geolocation of people, hand gestures, to differentiate between indoor/outdoor environment, and other aspects of device context.

3.2.2 Opportunities and Challenges

Future wireless networks will be challenged by a myriad of emerging applications, many of which will require us to rethink the design of wireless network architectures and protocols. A number of them were discussed in Section 1. The wireless research community should foresee ambitious applications that would bring powerful societal benefits and that specifically would highlight the pressing challenges to wireless networking. Many future wireless applications demand current assumptions and design principles to be revisited. As outlined below, they present interesting opportunities as well as challenges.

Sustaining uplink data flows. One challenge stems from the original assumption that mobile devices were expected to be data consumers. Contrarily, the deployment of participatory sensing systems, context-aware social applications and fleets of intelligent cars all result in the production of significant amounts of data at the end host. While much of this data ultimately can be collected in a cloud, some can also be immediately consumed by the end hosts, resulting in new patterns of communication. This throws doubts on the current architectural assumption that downlink traffic far overwhelms uplink.

Sustaining large scale and high density. Trends predict that there will be at least a 1000x more end users in the future, whether these devices are on a user's body or on the environment around them. Such massive scale requires new solutions for energy management (i.e., batteries of devices cannot be charged all the time), spectrum management (all of these devices must share the spectrum), real-time monitoring and feedback (these devices are collecting data that need to be processed and analyzed in the face of unreliable and unpredictable networks), as well as addressing and naming (it may not be necessary to know the names of all devices). Finally, such scale will necessitate diverse communication paradigms (user-centric versus data-centric).

Supporting real-time communication. As emerging applications are more closely linked with control systems (e.g., cars) and human perception (e.g., cognitive assistance), they also impose stricter latency requirements onto the network. Specifically, the network needs to support end-to-end latencies on the order of a few milliseconds for many interactions. For example, the network needs to be able to transfer camera images from a mobile device into the cloud, process them, and return the results in time for real-time display. This requires wireless access technologies that can establish rapid connections without compromising the energy savings from the current use of deep sleep modes. This may also require fundamental architectural changes to embed processing in the network fabric, such as supporting cloudlets that offer processing near the edge of the network.

Providing ubiquitous location information. Localization is the key to many mobile applications. Despite many years' research, accurate indoor localization, e.g., for wireless 911 calls or emergency responders, remains a challenge. Today's applications only have a simple concept of location that they expect to be accurate, leading to the need for better interactions between applications, mobile systems, mapping and location/localization services. On the other hand, some applications may be willing to trade off location accuracy for reduced energy consumption, with feedback about the expected accuracy.

Providing context-awareness without compromising security/privacy. Given the expected dynamics in future wireless networks, in terms of availability of bandwidth/spectrum, demands of the application, number of participating devices, it will be necessary to design context aware networks that can adapt to these environment and social changes. However, it is important to understand the security and privacy implications of collecting and using such context to drive the operations of a mobile system.

3.2.3 Recommendations

Applications are essential to the health of the networking community because, as they have done in the past, they shape assumptions and design principles for system design. Historically, application research was not given enough attention by conferences, journals, and funding agencies. We recommend three actions to strengthen the ties between technology development and driving applications.

Benchmarks. The community needs to stimulate the creation of application benchmarks that demand a paradigm shift in the field of wireless networking. Applications must be motivated by real needs in external disciplines, such as medicine, transportation, or computer vision. Benchmark studies must be replicable or reusable so that independent researchers can iteratively test and compare new network architectures with the workload from these benchmarks. The application benchmarks should do more than simply drive demand on throughput, latency, or energy -- they should demand that researchers revisit the basic tenets, principles, and conventions in network design.

Integration. There is explicit value in system integration as well as application support. Good research needs to find compelling research problems that explore not only theoretical limits, but also how to approach them in real systems, including both an understanding of current mobile system designs and expectations of future leading-edge systems. It is equally important to ensure that these research problems address needs and requirements of real-world applications, aiming for finding a healthy balance between research and practice. Instead of designing systems and algorithms for indeterminant users, there is a strong need to match application requirements with system and network capabilities, which ultimately can indicate which capabilities need to be the focus of future research.

Interdisciplinary. The most successful wireless and mobile applications today result from a thoughtful integration of expertise from many disciplines, both within computer science (such as machine learning, computer vision, image search) and beyond (behavioral health, transportation, control systems). All too often, however, developers writing applications for current wireless device make invalid assumptions about the behavior, performance and cost of the underlying wireless network -- in some cases, because of limited or outdated abstractions provided by the network stack API. Conversely, wireless networking researchers may be designing new network protocols and algorithms without an understanding of the needs of applications in other fields. It is of utmost importance to encourage and support interdisciplinary research on mobile and wireless applications, including efforts to prototype those applications in the context of novel wireless network technologies. The effort for cross-disciplinary collaboration and integration will bring new insights to wireless network research -- just what the community need to effectively identify future directions in wireless networks, and develop effective solutions for tomorrow’s applications.

3.3 Wireless Networking Architectures and Testbeds

3.3.1 Background and Prior Successes

The research agenda for wireless networking architecture is fairly intertwined with the types of testbeds suitable for research. We discuss these two components in turn.

Wireless Network Architectures. Wireless networks have become increasingly important to society and are now viewed as critical to the nation's communication infrastructure. With the rapid adoption of mobile data services, wireless/mobile networks now generate the majority of Internet traffic with the smartphone replacing the PC as the primary computing platform. As a result, wireless network architecture is a high-impact area of research where previous systems (for example, cellular and WiFi) that have been in operation for over 20 years are overdue for radical change and innovation.

NSF has invested in novel wireless architectures in the past, directly helping to foster innovations such as wireless mesh/ad-hoc/sensor networks, cognitive wireless networks, cross-layer protocol design, cooperative radio, multi-user MIMO, network coding and so on. Many of these innovations have been adopted by industry and have led to significant improvements in wireless services, with others currently in the pipeline.

One of the major problems with today's wireless systems is the "vertical silo" approach to building solutions involving specific PHY/MAC/network solutions that do not carry over between different technologies or applications. This has led to a proliferation of wireless protocols that prevent seamless integration of heterogeneous technologies into a single network. Consequently, there is an urgent need for a unified wireless protocol architecture which offers the features of high capacity, ultra-low latency, cross-layer visibility, unified API's across multiple technologies and seamless integration of optical and wireless and with the global Internet, to name a few requirements.

In addition, the protocol architecture needs to preserve some degree of modularity in order to constrain complexity, ease deployment, support sustainability, and achieve good security properties. This would be greatly facilitated by designing an architectural framework that supports on-the-fly fully programmable modification of MAC and PHY protocols throughout the network, made possible by emerging software defined radio technology. Security and privacy are also important areas of research for such an architecture, given the fact that wireless systems are susceptible to jamming, interception and traffic/location analysis attacks.

Wireless Testbeds. Research testbeds are of particular importance to the wireless field given the high cost of setting up individual experiments involving more than a few radio nodes. In addition, as programmable radio devices become ubiquitous, it is increasingly important that technology platforms and remote testbed access be available to a wide range of researchers and students, including those from the more general computer science community. With a new generation of wireless systems around the corner, it is ever more important for the research community to have access to state-of-the-art platforms and testbeds for realistic and at-scale experimentation.

In the past, NSF has supported a number of relatively successful wireless testbed and platform development projects including:

- Software-Defined Radio (SDR) Platforms: GNU/USRP, WARP, WiSER --- hardware/software kits distributed to the research community
- ORBIT (large-scale emulation and outdoor deployment of WiFi, SDR and cellular test equipment), operated as an open-access community resource since 2005
- Specialized testbeds such as CMU's Wireless Emulator, MoteLab (sensor network lab), Phonelab (emulation of mobile phone experiments), etc.
- Outdoor mobility testbeds such as DieselNet, Wisconsin's vehicular networking testbed, etc.

These radio platforms and wireless testbeds have enabled a significant number of experiments across the community and can be credited with a steady shift to experimental methodologies during the past decade. Another NSF-funded major effort is GENI, which has focused on remote testbed access and federation, both of which are relevant to wireless research.

3.3.2 Challenges and Opportunities

3.3.2.1 Wireless Architecture

Grand challenge. To make the research agenda for a unified protocol architecture more concrete, we introduce a "grand challenge" in wireless network architecture that tackles some of the key challenges facing wireless networks and access today, as follows: Design and validation of a future wireless network architecture that achieves the high-level performance goals of ~1000x capacity, 1 ms latency, while offering a future-proof protocol framework with features such as cross-layer visibility, unified API and SDR capability.

Towards this end, the following specific challenges need to be addressed:

- Integrating advanced communications technologies such as distributed MIMO and interference alignment, and translating the physical-layer capacity gain into network throughput improvement and latency reduction
- Incorporating technologies at higher frequencies (e.g., 60 GHz) using new network protocols, e.g., antenna beam steering

- Dynamic spectrum access capabilities including data plane features such as adaptive PHY/MAC and control plane features such as spectrum information exchange
- Effective integration with emerging optical and software defined network technologies likely to be used in access networks
- The ability to leverage in-network storage & computing (i.e. edge cloud or “cloudlets”) to significantly reduce network latency and improve service quality
- The architecture should work across multiple radio technologies, encouraging spectrum etiquette and promoting co-existence
- The architecture should provide APIs for fine-grained control and cross-layer algorithms. Meanwhile, it should maintain a degree of modularity necessary to manage complexity, security and privacy
- Incorporating features such as robustness, multi-homing, resource control, energy efficiency, light-weight protocols, etc.

Opportunities and connection to emerging areas. The design and evolvement of future wireless network architecture is inseparable from the following key research areas:

- Emerging cloud and software-defined network (SDN) technologies as important enablers for wireless networks
- Emerging content, service, and mobility-centric future Internet architectures that link seamlessly to wireless network architectures at lower layers of the protocol stack
- Internet-of-Things (eHealth, vehicular, underwater acoustic networks, environmental applications, and so on)
- Smart grid -- communications architectures supporting real-time, ultra-large scale smart grid sensor networks
- Big data including development of a national spectrum database and privacy-preserving approaches to storing mobility/usage data from more than 6 billion personal mobile devices
- Machine learning concepts related to cognitive radio networking

3.3.2.1 Wireless Testbeds

Grand challenge. To support the development of next-generation wireless architectures, the wireless research community should investigate large-scale real-world wireless testbeds capable of evaluating new designs to support ~1000x capacity and 1ms latency. The following testbeds are particularly important for the research community.

- Large-scale mobility testbed with real-users, something difficult for individual researchers to setup, motivating the need for a common resource
- Real-world large-scale spectrum observatory/testbed which provides a measurement basis for current work on dynamic spectrum
- Experimental platform for next-generation cellular networks ("5G"/60 GHz, cloud RAN, optical backhaul, etc.), which has historically been developed in a proprietary manner by industry
- Integration of optical, wireless and wired Internet.

Key features of the proposed testbeds should include:

- National-scale deployment with multiple urban/suburban/rural coverage areas
- Open access to research community with high degree of programmability
- Fine-grained control of radio MAC and PHY
- Programmable dynamic spectrum access capabilities
- Spectrum and traffic measurements at scale
- Real-world mobile end-users and outdoor spectrum use
- New technology platforms such as SDR, SDN, cloud, radio cloudlet, etc.

While the actual testbed implementation may be physically distributed for practical reason, it is important that the unified capabilities of the testbed(s) can support the broad research vision outlined above. This is clearly an ambitious undertaking, but one that is necessary to adequately address research issues related to the emerging heterogeneous networks and dynamic spectrum environment.

Other opportunities and connections: There are a number of important links to industry, government agencies, and other research communities in relation to the proposed future wireless testbeds. In particular, it would be useful if the community is better connected to organizations such as the FCC and NTIA. Furthermore, international collaborations in this domain are particularly important. Federation of testbeds with European (e.g., Fed4Fire) and Asian (NiCT/JGN) partners would be quite valuable as it would provide the community with a holistic view of testbed design. Furthermore, the community should explore synergistic connections to complementary activities in related topics, e.g., SDN architectures, applications of big data systems, and so on.

3.3.3 Recommendations

In summary, research activities in this topic area could focus around the following example issues:

1. **Future research programs in wireless should be aligned to address the grand challenge of "1000x capacity – 1 ms latency" in future wireless networks.** Achieving this goal will require architectural research on a number of topics including definition of a unified protocol stack for wireless systems, scalable network architecture, cross-layer visibility, efficient spectrum use (DSA), adaptive PHY/MAC capability reflecting recent progress with SDN technology, use of edge clouds, and others.

2. Given the above architectural grand challenge, **an ambitious community project ("mid-scale infrastructure" by NSF's definition) aimed construction of a large-scale future wireless network testbed is recommended.** Features of this testbed include national scale deployment, open access and programmability with fine-grained control of radio MAC/PHY, integration of software-defined radio (SDR) and software-defined networking (SDN) technologies, programmable dynamic spectrum capabilities, spectrum and traffic measurements at scale, and real-world mobile end-users.

3.4 Controls and Algorithms

3.4.1 Background and Prior Successes

Network control and algorithms have been at the heart of the Internet from its inception. More recently, with the emergence of wireless technology, network control and algorithms play a critical role for providing wireless access that makes efficient use of scarce wireless resources. Looking forward, it is clear that effective network control and the associated algorithms will play an increasingly important role in providing network services in future wireless and heterogeneous networks.

From its inception, the Internet has depended on network control algorithms for its operation. In fact, the Transmission Control Protocol (TCP) has control theory as its underpinning, as shown by Jacobson in his classical 1988 paper. Later, researchers in our community have shown that in fact TCP and its variants can be used to maximize network utility (e.g., see work on network utility maximization (NUM) by Kelly, Low and others). Further, the Internet is largely a distributed system, and distributed control algorithms, such as distributed Bellman-Ford routing, have been the engine that makes the internet work.

Moreover, nearly everyone in the world uses wireless local area networks to access the internet, especially with mobile wireless devices such as tablets and laptops. At the heart of wireless LANs is the underlying channel access scheme that enables users to share the scarce wireless medium. WiFi networks are based on carrier sensing and binary exponential backoff mechanisms for medium access control, both mechanisms developed by the network control algorithms research community over the past three decades.

A feature that distinguishes wireless networks from wireline networks is that wireless resources can sometimes be scarce, due to bandwidth, power and rate limitations. Significant advances in control theory and algorithms have been made to address this fundamental problem. In 1992 Tassiulas and Ephremides developed a throughput-optimal algorithm that makes routing and scheduling decisions based on queue backlog information. The max-weight framework has acquired fundamental importance in network control, as it provides an explicit mechanism for throughput-optimal control of multi-hop networks. Since then, it has been adopted into various networking contexts, including switches (switch scheduling in internet routers is based on the max-weight scheduling principle), satellite networks, wireless networks with power control and time-varying channel conditions, and optical networks. In fact, opportunistic scheduling, that takes advantage of the time-varying nature of wireless channel conditions, is an outgrowth of this research, and has now been adopted by nearly all next-generation wireless systems.

The field of control theory and algorithms has seen significant success in translation from analysis and theory to widespread incorporation into commercial wireless systems. This is evidenced by the algorithms in 3G and LTE wide area networks to support mobility and optimization of coverage, capacity and quality of service. Some examples of control algorithm application in these wide area network systems include power control through closed loop feedback, and dynamic rate adaptation and scheduling that incorporates efficient channel quality feedback to maximize performance while maintaining stability and robustness across a network of cells with embedded smaller cells. It is fair to say that many who use their mobile devices so frequently likely take for granted the large number of underlying control theory algorithms at work enabling the high level of mobility and efficiency in today's wireless communication networks.

3.4.2 Challenges and Opportunities

Decentralized network control and algorithms. With the ongoing and steady increases in the needed capacity of wireless communication networks, the future need for advancing the theory of network control and algorithms will continue unabated. In particular, there is a need for advancing the theory of decentralized and distributed control algorithms which achieve a network-wide utility within and across networks with minimal information exchange. Future network control algorithm and network analysis would be particularly useful if it is decomposable and scalable across heterogeneous networks each with heterogeneous elements -- a network of networks. Distributed and efficient algorithms and minimization of required side-information transfer will continue to be important aspects of efficient networks as the scale of the networks grows.

As practically achievable algorithms continue with advancing technology, there will be a need for intuition and comparison of alternate approaches against bounds from deep analysis and theoretical understanding. Similarly, while the sophistication of complex wireless network simulations can increase with technology and code development, there is a need for the control theory underpinning architectural and algorithmic tradeoffs to have a tantamount increase.

Integration of control, communication, computation and storage. In many emerging applications, where wireless networks are playing a central role -- from mobile computing to distributed robotics, cyber-physical systems and the Internet of thing, a common theme is that "control has gone wireless." Thus, algorithmic advances are needed not only to control the information flow in wireless networks in order to achieve traditional metrics such as throughput and capacity improvements, but also to enable the stability of other control loops that rely on the network, imposing higher requirements on latency. Other performance metrics of interest in these applications are energy efficiency, robustness, and privacy. With increasing heterogeneity of nodes (social, cyber, physical) and the linkages between them, these types of applications will also require improvements in our understanding of network geometry and its use in network control.

Another key challenge is the increasing blurring of boundaries between communication, computation, and storage. While recent years have seen a move towards centralized cloud computation, the increase of smart devices at the edge and the resultant taxing of scarce bandwidth creates an increasing need for algorithms that can decompose, virtualize, and allocate computation, communication and storage resources in a more flexible manner.

As a new economy arises around new applications such as cyber-physical systems and the Internet of Things, there are also new opportunities for online and adaptive network control and resource allocation algorithms in novel products. Indeed we have been seeing the rise of start-up companies in these areas that leverage these advances and we expect this impact to grow.

Dealing with highly-dynamic environments. New challenges arise also because of the increasing levels of dynamics in wireless networks, for instance, as cell-sizes become smaller in order to increase capacity, there will be a concomitant increase in inter-cell mobility, and in the number of femto-cells and access points. Similarly, new Internet of Things and CPS applications are likely to generate more small bursts of data with stringent requirements than long flows, requiring greater adaptation of the network to application traffic. But, on the flip side, new opportunities present themselves, such as the ability to control the location of network resources in order to dynamically adapt network performance using robotic relays. Together, these motivate the need for network control and resource allocation algorithms with greater flexibility and adaptation to changes, for instance through built-in learning mechanisms.

The I^2 Intelligent-Information Fabric: In summary, a grand challenge will be how to revolutionarily transform the wireless network from a "bit-pipe" to an intelligent information fabric. Traditionally, the wireless network has been a connectivity and bandwidth provider (i.e., a "bit-pipe") that connects the user with the data/service (e.g., in the cloud). With increased demands from the users (e.g., small latency, real-time control, real-time user behavior adaptation, user data privacy where key data is on the device and not in the cloud), we need to move to an intelligent fabric, where the network architecture evolves from {dumb bit-pipes + smart cloud computation} to {computation, control and humans embedded and distributed within the communication fabric}. This "smart fabric" will combine computation and storage within the wireless network to provide enhanced services to users. The research agenda is to determine the architecture for such a smart fabric (what, if any, are the design rules for layering and partitioning of functionality within the network), as well as the unified learning, computation, communication and complexity-management algorithms for both wireless/mobile users and nodes within this smart fabric.

3.4.3 Recommendations

This working group recommends that future research in this domain be focused along the following pathways.

1. **Abstractions that go between theory and practice:** Over the past 20 years the research community has developed a theoretical framework for optimal control of wireless networks. This theory, which has had influence in our community and beyond (e.g., transportation systems), was based on idealized modeling abstractions for wireless networks. As the practice of wireless networks matured, it is becoming clear that some of the modeling assumption made by the theoretical framework fail to capture some important aspects of practical systems, e.g., for highly dynamic scenarios. An important challenge to this community is to develop unified abstractions that bridge the gap between theory and practice, and to develop network control and algorithms that can have significant impact on practical wireless systems. Such abstractions will enable distributed and localized decision making to achieve globally optimal objectives with minimal message exchanges.

2. **Robust and Distributed Resource Allocation:** Wireless networks have emerged as a promising solution for providing last-mile Internet access for communities that are beyond the reach of the wired infrastructure. These networks will be tightly coupled with the wired Internet to provide Internet services to end-users. While users of wired networks are accustomed to highly dependable network services, wireless networks have to contend with interference, obstructions, jamming, and mobility, and are notoriously unreliable. A grand challenge to this community is to provide dependable communications in this challenging environment.

   Distributed network control and resource allocation algorithms are needed that provide reliable communications in multi-hop networks. With the ever increasing dependence of Cyber Physical Systems on communications for their control ("control has gone wireless"), there is an increasing need for network algorithms that can provide latency guarantees for real-time applications. Similarly, as we come to depend on wireless networks for critical applications, network architectures and algorithms must be designed to provide reliability guarantees (e.g., connectivity, probability of message delivery). Finally, network control and resource allocation algorithms must be secure in the face of both intentional attack (denial of service, jamming, etc.) as well as random failures.

3. **Theory and models for socio-technical systems:** TheI^2 fabric brings new challenges in network models. With humans and other non-traditional elements in the network, new models and tools are necessary to describe and characterize network behavior and performance.

   First, we need to develop data-driven models that characterize the network elements as well as the interactions between elements. This requires enhanced tools from machine learning and statistical inference that synthesize with models for communication and computing elements. These new models must be rigorously validated, with analytical guarantees of correctness. Specific applications range from distributed learning and optimization for 5G wireless systems, to networks that combine human elements and their cognitive responses within a strong mathematical framework.

   Second, these models will enable resource allocation and influence algorithms interpreted broadly -- resources include sensitive data, network resources (wireless spectrum, backhaul), computation resources, and human elements. Indeed human elements could serve in multiple roles including controllers, actuators, sensors and plants (cyber-human systems).

3.5 Emerging Spectrum Access Technologies

3.5.1 Background and Prior Successes

The past decade has witnessed a skyrocketing demand for commercial spectrum. Similarly, demand for unlicensed bandwidth keeps increasing, due to the continuing growth of WiFi deployment and the emergence of additional applications. Such trends will continue in the near future. For instance, Cisco projects a18 times increase in data usage from 2011 to 2016.

Meeting this huge demand for bandwidth is a challenge since most easily usable spectrum bands have been allocated. CTIA estimates that U.S. cellular companies will face a spectrum deficit of 275 MHz by 2014. On the other hand, many studies have shown that more than 90% of the allocated spectrum is unused or under-utilized. In other words, many of the frequency bands are only used in certain areas and/or only part of the time. This suggests that meeting future demands for wireless bandwidth will not only require more efficient communication and networking technologies, but also new techniques for increasing spectrum utilization. To address this issue, Dynamic Spectrum Access (DSA) can be developed to flexibly allocate spectrum to users and systems in response to system demands for bandwidth. Realizing this strategy through the use of cognitive radios requires not only advances in state-of-the-art radio technology but also new policies and economic models for spectrum use.

A significant body of work has been devoted to developing innovative DSA techniques that enable wireless devices to dynamically sense the communication environment and adapt their transmission schemes, in terms of waveform, spectrum access method, and/or networking protocols to meet the quality-of-service requirements of wireless applications. Significant progress has been made in many specific areas, including information theoretic analysis of dynamic spectrum access, modulation and waveform design, spectrum sensing, multiple access, power control, routing, cross-layer optimization, spectrum mobility and handoff, economics of cognitive radio systems, small-scale testbeds, software and hardware prototypes, and standardization.

3.5.2 Challenges and Opportunities

Most of the existing work focused on specific spectrum access functions and small-scale radio design. However, what is lacking is a new spectrum architecture and clean slate design that can multiply the effective capacity of the spectrum by a factor of 1,000 and enable spectrum sharing at scale across many dimensions, in terms of number of nodes, channels per radio, geographical extent, etc. Besides scalability, the new spectrum architecture must take into account emerging wireless applications that are likely to disrupt multiple industries and trigger a new surge of spectrum demands in home-automation, gaming, and sensor networking.

New models are needed to describe spatio-temporal spectrum availability and characteristics over a broad range compactly. For example, the FCC database of station licenses has been used to predict spectrum availability in the context of TV white space. New database models (e.g., a closed-loop feedback system) are needed for resource management to deal with fast changing spectrum usage environments, and should enable ubiquitous and distributed spectrum sensing and sharing. Another equally important domain is the 7 GHz of unlicensed millimeter-wave spectrum around the 60 GHz frequencies. Accurate channel models and efficient medium access algorithms are needed in order to achieve high spectrum utilization and reach gigabit-per-second data rates as anticipated.

Another challenge in spectrum access is to build a systematic understanding of the regulation side. A regulatory framework is key to tackle the challenges for the certification of cognitive radio devices and the enforcement of regulation polices. The regulatory concerns themselves need to be better analyzed, e.g., through a new theoretical framework that models the fundamental impact of regulatory constraints on wireless system performance. This will likely involve an integration of both economic and social issues with technical solutions.

There is an urgent need to develop fundamental theories that characterize the tradeoffs between capacity and latency, between spectrum sharing and security/privacy, and between power and bandwidth. A consensus is that stochastic, real-time optimization of spectrum allocation will play a key role in enabling effective spectrum sharing. It is envisioned that future generation wireless networks comprise systems of interdependent networks, such as social-tech networks and tech-economics networks, so new spectrum access techniques must take into account (and leverage) the interplay across different networks in multiple domains, including social networks, economics, policy, etc. Furthermore, spectrum access is fundamentally centered on dynamic sharing among heterogeneous uses. Compatible models are needed to be able to theoretically evaluate sharing models as well as to guide the design of actual system architectures for sharing. Regulators, engineers, and the public need to be able to understand the tradeoffs in a clear and transparent manner.

The evaluation of new spectrum access technologies remains open and challenging. Significant efforts are needed for developing open spectrum access testbeds, including both hardware and software components, that is able to test innovative ideas in a complete, realistic, working system at scale.

3.5.3 Recommendations

We summarize below some recommendations for emerging spectrum access techniques that have great potential to offer efficiency, robustness, security and resilience, in order to meet the needs and requirements of a variety of wireless applications.

1. **Clean Slate Design and Fundamental Theories.** New spectrum access techniques should not be constrained by the deployment in current networks, nor by current technologies. New architecture design should be from a holistic perspective and be targeted at developing comprehensive end-to-end solutions for reliable, scalable wireless networking. It is envisaged that future generation wireless networks comprise systems of interdependent networks, such as social-tech networks and tech-economics networks, so clean slate design should be based on an integration of economic and social concerns with technical issues.

2. **Millimeter Wave Technology.** Accurate channel and propagation models are needed for millimeter wave communications. It is crucial to develop a fundamental understanding of tradeoffs between power and bandwidth of millimeter wave technology. It remains largely open to devise scalable millimeter wave beamforming algorithms and highly directional communications protocols for spectrum access.

3. **Database Resource Design and Management.** New database models (e.g., a closed-loop feedback system) are needed for resource management to deal with fast changing spectrum usage environments, and to enable ubiquitous and distributed spectrum sensing and sharing.

4. **Cognition Capabilities going beyond Dynamic Spectrum Sharing.** It remains open how to operate with imperfect cognition and what can be achieved. Another challenge is how to harness a myriad of heterogeneous communications resources and capabilities efficiently, and what are the technical and regulatory implications.

3.6 Metrics, Measurements, and Management

3.6.1 Background and Prior Successes

A fundamental part of evaluating systems is to know what to measure, how to measure it and how to report those measurements. Those measurements can be used to create *models* that are used to describe a range of conditions in which wireless networks may operate. Proper models are critical to the evaluation and development of wireless networks.

The mobile and wireless community has developed models that both integrate and go beyond the prior contributions, including cross-channel models, occupancy models and mobility models. For instance, the fact that much spectrum remains unused for long durations (as indicated by measurements) could be converted into a model of spectrum occupancy that was then used to design different dynamic spectrum access policies.

Testbed experimentation, a cornerstone of recent wireless research, is often used to demonstrate sufficiency of the models used to evaluate a wireless system in certain analytical or simulated environment. Conducting an “in the wild” experiment demonstrates how the system can respond to unforeseen conditions, events and inputs not captured by those models. Test measurement can also be used for model discovery, for model validation, for exploring robustness, as well as for run-time calibration and fine-tuning of models and operating parameters.

Currently, the experimental methodology used by researchers and taught to students is often uninformed about methods such as Design of Experiments (or other multi-factorial analysis techniques) that could be used to conduct measurements over many factors in an efficient manner. These techniques could be used to determine the significance or importance of different factors in an experiment, and are much more efficient than factor-at-a-time experimentation.

This brought forward the notion that there should be greater emphasis on data-driven model development. Many existing models are based on model-fitting of data (*e.g.* propagation models that determine an exponent by fitting measured data) and others are derived from hypothesized behavior (*e.g.* mobility models or Gilbert-Elliot channel models). Because these models start with a preconceived behavior, they do not fully capture actions that can impact system performance. A sound model should be able to extract insightful patterns from data that are directly collected in the wild.

3.6.2 Challenges and Opportunities

The major challenge facing wireless measurement is *diverse data-driven models that can be used in closed loop to guide designs and architectures at all different layers, including the regulatory "layer".*

Learning models from diverse measurements. Today's wireless models are largely based on apriori assumptions, and measurements are usually gathered from carefully planned homogenous experiments to set the parameters within those models. There is a need to complement them with models automatically synthesized from diverse data sources ranging from actively controlled experiments to passive measurements. Machine learning techniques should be leveraged to fuse big data from potentially wildly dissimilar data sources (e.g. user tweets, WiFi measurements, ubiquitous photography, earth-sensing satellite measurements, and active perturbation of a subset of users' network experiences.).Dimension-reducing mechanisms such as PCA could be used to develop a set of parameters that could be used to configure experiments.

The automatically discovered and tuned models should be rich enough to allow simulation/emulation to evaluate the typical performance of wireless systems, so that designers can automatically discover unknown fragilities in their designs. Typical methods include "fuzz-testing" with random input, and deliberately breaking system state to trigger error-handling. The development of such a system to automatically characterize a model and reproduce the full range of behaviors is a worthy challenge.

Widening scope of models and measurements. An important aspect of this challenge is that the scope of the needed models (as well as the simulation/emulation tools) is much larger than it was in the past. For example, traditional omnidirectional propagation models and interference models are no longer sufficient. Higher frequencies (e.g., 60 GHz) and multiple-antenna systems demand a fresh look. Massively multiuser wireless systems demand models of the spatial correlations in link qualities across users.

Wireless measurements and models must go beyond RF itself. As we enter a world of shared spectrum and start to take account of regulatory interest in enhancing freedom, increasing innovation, and meeting other social needs, we need models that can approximately capture new effects (e.g. the economic choices of individuals and the investment decisions of firms) as well as new metrics (e.g. quantifying "innovation friendliness" for spectrum rules and "friendliness/neighborliness" for wireless systems in a shared environment).

Engaging other communities. For some of these models and measurements, an interdisciplinary effort is needed that appropriately engages with other communities like economics, business, consumer psychology, and possibly even law. For some others, a totally fresh approach might be needed that combines machine learning with human preferences to infer the underlying metrics. For yet others, it might be necessary to aggregate many different performance metrics developed independently for various previously isolated applications into new unified forms.

Beyond the designer-centric questions above, there are also important research opportunities to help those actually deploying, running, or maintaining wireless systems. For example, modern regulators need to be provided with an appropriately real-time picture of current spectrum use on a nationwide scale as well as broken down by specific regions and applications. Measurement doctrines are needed that are general enough to encompass the needs of monitoring and fine-tuning, as well as anomaly-detection and enforcement.

3.6.3 Recommendations

In summary, the following research thrusts should be promoted to amplify the effectiveness of wireless measurements and models:

1. Research on data-driven model development should be encouraged
2. Industry should be encouraged to shared activity traces for existing mobile user behavior to facilitate extraction of salient features
3. Facilities for measurement based studies should be developed as well as a methodology for extracting and cataloging a consistent archive of such information
4. The research community should emphasize the archival and management of data products
5. The development of an environment that allows code to be moved from simulation environments to actual deployment (at the PHY and MAC layer) should be encouraged and standardized.
6. Development of purely theoretical models for already well-understood areas should not be a high priority for researchers

3.7 Security and Privacy

3.7.1 Background and Prior Successes

The security and privacy stakes for wireless networks have been significantly heightened, not only by the proliferation of smartphones carrying much personal information, but also by the emergence of the Internet of Things including body-area networks and home-area networks, and by wireless edges to cyber-physical systems (e.g., smart-cars and critical infrastructure such as smart grids). A variety of attacks targeting wireless networks and devices have been observed, e.g., device impersonation and other types of authentication related attacks. In addition to "selective" jamming strategies, observed attacks have employed deliberate misinformation and non-conformance to standard protocols (including not obeying commands from base stations or access points) to deny access to resources including communication. Moreover, sophisticated eavesdropping attacks can infer important information from traffic analysis or mere monitoring of localized physical layer activity.

Securing network operations are more challenging in a wireless setting simply owing to the limited resources of wireless devices and the limited capacities of the channels over which they communicate, including resources for securing data integrity, privacy, authenticity and network availability. Achieved efficiencies due to cross-layer design and cognitive radios have often overlooked security issues. Cross-layer and cognitive radio frameworks may introduce attack vectors of more diverse attack surfaces.

Efficient cryptographic mechanisms have been developed in the past to achieve data integrity and privacy, in particular those leveraging tamper-resistant hardware for cryptographic computation with "burned-in" keys. Moreover, existing wireless communication protocols, e.g., 802.11 family and Bluetooth, do offer some measure of security. Indeed, academics have pointed out vulnerabilities to earlier versions of 802.11 standards, for example, which were corrected in subsequent security standards. Similarly, academics have contributed to increased security of wireless interfaces to cars and pacemakers.

Though there has been significant research to develop security systems (e.g., leveraging cryptographic primitives), and to secure different proposed wireless PHY/MAC frameworks, the need remains for comprehensive, feasible and adaptive security solutions to defend wireless networks handling increasingly sensitive and diverse types of information from increasingly sophisticated attacks.

3.7.2 Challenges and Opportunities

Overall Challenges: There is a need to clearly ascertain and model existing and emerging threats to data privacy, integrity, authenticity, and network/resource availability for emerging wireless systems and the diverse set of applications they are expected to support, in particular related to Internet of Things and Cyber-Physical Systems applications.

Defenses and attacks must factor practical resource constraints in their deployment contexts, such "work factors" need to be quantified by metrics. For example, the work factor of an attacker targeting delay performance of delay-sensitive application may be far less than one trying to achieve complete denial of communication.

Finally, defenses must be able to adapt, again within their means, to an evolving adversary, again which needs to be modeled. This is critical to, e.g., ensuring availability (e.g., from jamming or protocol misuse), providing location proofing, hiding the identity of communicator (anonymity) and the presence of communication (unobservability).

Securing Emerging Applications and Systems: There are heightened security concerns on wireless edge due to emerging "Internet of things," i.e., chatty devices that may carry a wide variety of sensitive, personal information, such as devices from body-area networks or home-area networks. Moreover, wireless edges of important cyber-physical systems, e.g., networked vehicles and critical infrastructures, may need to meet strict control (delay, loss) requirements. Thus, quality of service guarantees must be provided against various types of threats.

Moreover, wireless apps downloaded to cellular smart-phones may be equipped with special-purpose networking mechanisms that have not been fully vetted from a network security perspective. There is a clear need to be able to safely validate the network security of an increasing diversity of networking protocols before they are deployed.

Finally, with the increase use of wireless I/O interfaces, the security concerns are also heightened. Wireless I/O interfaces may be vulnerable in unexpected ways, e.g. near-field communication (NFC), visible light communication (QR code), audio/vibration side-channels (covert channel example). There is a pressing need to understand how to protect against vulnerabilities from these interfaces.

Privacy-Preserving Wireless Access: As our society becomes increasingly reliant on wireless access, an important goal is to achieve "authorization without authentication" regarding a high degree of desired privacy, e.g., how to connect to the wired AP without the wireless access provider knowing who you are. This may be achieved by a type of (multihop) wireless onion routing involving a community of proximal and mutually trusted wireless devices.

Cross-Layer Design for Security and Privacy: Cross-layer designs may better exploit opportunities to improve resource utilization. For example, if traffic is encrypted at a higher layer, then re-encryption at a lower layer is unnecessary, leading to energy savings in the wireless device. Moreover, an IDS operating in the network layer may detect evidence of a PHY layer attack or vice versa. Similarly, CPS control applications may use custom cross-layer designs, in particular to better manage stringent delay/reliability requirements of wireless access.

However, there is a trade-off in design for performance and security/robustness. Specifically, there is a need to understand implications to security of cross-layer designs and to jointly design with both performance and security considerations. Further study is needed to better understand how cross-layer designs complicate security analysis and increase attack surfaces and risks. Finally, often network operators demand operational transparency and modularity of their networking components. Thus, operators may object to vendor’s cross-layer designs despite their efficiencies.

Joint Design of Software-/Hardware-based security: Some embedded systems already employ hardware support for cryptography. A basic question is how best to leverage tamper-resistant hardware to create wireless security solutions. Further, another basic question is how to protect from threats (installed trojans/backdoors, side-channels) in the design and manufacturing process of putatively tamper-resistant hardware. Finally, because the software operating on compromised hardware is inherently unreliable, there is a need for remote (software) strategies for verification of wireless devices with putatively trusted platform modules (TPM).

Privacy Problems in an Age of Ubiquitous Geolocation: An increasing problem in the mobile era is to achieve self-localization without divulging private information to other parties including the wireless access provider. This issue brings up the important aspect of informed consent in privacy matters and the challenges in conveying to the users what information is being revealed and the associated risks.

MAC/PHY Layer Security: End-to-end security and security of composed systems needs to be robust to a wide variety of possible threats from the PHY layer and for different PHY-layer variations/operational states. A comprehensive theory of leveraging/composing proposed PHY/MAC security solutions is necessary to achieve end-to-end security. Furthermore, new and practical security evaluation mechanisms are needed for emerging PHY layer security primitives (e.g., for wireless key agreement). There is a basic need to understand the feasibility of MAC protocols that are resistant to traffic analysis. This may be achieved through a joint design with new PHY interfaces, e.g., directional antennas, but such new interfaces involve costs/work-factors that may outweigh the risks associated with divulging the information in question.

3.7.3 Recommendations

Below we summarize the recommended immediate and long-term research efforts in the wireless security and privacy area.

1. **Threats to cyber-physical systems with delay-critical applications.** Cyber Physical Systems (CPSs), e.g., individual vehicles or entire "smart" power grids, require applications whose communication is both delay- and loss-sensitive. Moreover, these systems rely upon wireless infrastructure for I/O including access, control and sensor networking. Simple intrusions by conventional means have been demonstrated on existing CPSs, by exploiting the vulnerabilities of different wireless devices (e.g., audio, camera, NFC). Thus, threats to availability and integrity of the wireless channel and associated resources (e.g., energy for transmission) need to be mitigated, and protocols need to be minimally functional to accommodate the needs of the supported applications and allow for rigorous specification-based defenses. CPSs may also employ custom SDR and cross-layer frameworks, which had been typically designed for delay and throughput performance, but without consideration of the vulnerabilities they introduce. These frameworks need to be assessed for vulnerabilities to realistic threats and their design reconsidered with built-in security.

2. **Light-weight privacy protections and authentication for the Internet of Things.** A plethora of devices, many with low energy resources and computational ability, will compose the future "Internet of Things". Some of these devices will participate in critical applications such as health maintenance. To secure such devices from threat through their wireless interfaces, new mechanisms need to be developed that are very resource efficient from the devices' perspective. The work-factors of exploits may be quite low for such devices, e.g., minor variations of networking protocols to deplete their energy resources.

3. **Privacy-preserving wireless access.** "Authorization without authentication" involves authorized access granted via a wireless interface without divulging the identity of the grantee of access, or other sensitive information about the grantee such as their physical location, to the authorizing entity. This type of privacy is important to prevent the potential accumulation of private information by third parties who have no need to know, let alone accumulate, such information.

4. **The use of tamper-resistant hardware.** Tamper-resistant hardware can be leveraged to create many useful security primitives at relatively low cost. When interacting with a device on a platform of putatively tamper-resistant hardware, there is a need for lightweight mechanisms to remotely verify that the hardware is tamper-resistant and has indeed not been compromised.

5. **Custom networking mechanisms of downloaded apps.** Smartphones are engaging an increasing number and variety of apps, many of which have not been vetted by the smartphone hardware or operating system vendors. These apps may engage in unorthodox networking practices that may be harmful to the hosting device and/or other proximal devices. In particular, apps may be trojans for malware and it has been demonstrated that a base station can be compromised by a small number of cooperating smartphones. From security and privacy perspectives, how can such activities be detected and mitigated, with or without the aid of the base-station?

4. Tools for Education and Outreach

4.1 Background and Prior Successes

Using broad community impact as the metric, the following can be identified as the best examples of prior success in education/outreach tools that have resulted from NSF support: Motes/TinyOS, USRP/GNURadio, WARP, and ns-2/ns-3.

The first three represent a combination of hardware/software platforms that are in significant use by the sensor and wireless networking communities, for both educational and research purposes. All these technologies started as NSF funded projects that have now been absorbed into commercial ventures. The open source network simulation engine ns-2/ns-3 is by far the most popular freeware among the networking education and research community. CISE should be appropriately recognized for all these achievements.

Besides the above, other notable achievements have been the establishment of research-oriented wireless networking testbeds. Of these, the ORBIT grid at WINLAB, Rutgers University is the best-known for its cope and scale. Another significant one is the Community Wireless Mesh Testbed in Houston using the WARP platform that is currently operated by a non-profit and serves as broadband access network for a low-income neighborhood. Other notable mentions for similar efforts go to the QuRiNet testbed at Napa Valley, by UC Davis, those developed under the aegis of the (now graduated) CENS ERC at UCLA, the hardware emulation platform at CMU, the MoteLab at Harvard etc. These have had significant impact on the education and training of graduate as well as undergraduate students in terms of introducing wireless networking technology at an operational level.

Despite the above, there is a clear lack of sustained and pervasive impact on the broader education/outreach front. For example, the research testbeds noted above have remained largely limited to a small user community. Development of courseware based on USRP/GNURadio or Motes/TinyOS for broader educational use has been spotty at best. The development of new and validated wireless layers of the protocol stack for ns-3 has been hampered by insufficient funding, and development of interfaces between ns and a hardware test-bed that might allow novel co-emulation-simulation approaches to flourish have not taken off. The initial success of Summer Wireless Schools (U. Illinois) that engaged a broad cross-section of undergraduate students from a variety of institutions in learning wireless fundamentals has also diminished.

4.2 Opportunities and Challenges

The commodification of wireless hardware and simple user interfaces has created a familiar problem – wireless networking has become a victim of its own success, at least to a general audience. The ease of use of an Android phone and the simple programming abstractions (‘apps’) have successfully obscured the underlying technological achievements that are wonderfully advanced and represent examples of significant engineering achievements. This dichotomy creates a great dilemma: fewer students are motivated to look beyond the surface features, at a time when deeper understanding and a greater synthesis of a variety of skill-sets is needed for the next generation of wireless innovation! A clear impact of this has been the well-documented lack of skilled embedded systems engineers relative to demand within various industry segments.

It is possible to turn this challenge into an opportunity by turning the problem on its head. Strategically, this needs to be pursued on at least two fronts – one, for broader public outreach and second, in terms of overdue substantive curricular reforms. Regarding the first, it is vitally important to convey to the broader public a compelling narrative that they can relate to – in terms of the critical contributions of wireless to various transforming technologies – cellular telephony (past), mobile Internet and social networking (current) and Internet of Things inclusive of Smart Home/Environments, Smart Grids, Connected Vehicles and Tele-medicine (future).

On the second, a fundamental re-definition of Computer Engineering education must occur within EE and CS departments jointly, in the direction of recognizing Network Engineering as a *fundamental* component. Within this perspective, introductory courses based on familiar programming environments for mobile computing platforms (smartphones and tablets) at an introductory stage (freshmen) with few programming pre-requisites, could be the point of departure of a new curriculum. This would then evolve into an embedded-systems centric curriculum with an appropriate mix of emphasis on hardware and software and experience. Such a curriculum would prepare a new generation of students with the right skill-sets and knowledge base for design and innovation of future wireless systems/sub-systems.

Producing new curricular material at various levels within the context of such significant changes calls for bold action – but one that CISE is uniquely capable of driving.

4.3 Recommendations

Education and training of the next generation is an important and continued need of this community. But to do so, it was unanimously felt that the practice of funding educational outcomes as part of research proposal should be discontinued, and replaced with a *concerted attempt to fund education oriented proposals on their own, by de-coupling it from research oriented funding*. Thus programmatically, the most emphatic recommendation is creation of a new *Community Education Infrastructure* Program (akin to Community Research Infrastructure) with the explicit charter of funding the following efforts:

1. Development, maintenance and distribution of free public domain courseware, lab manuals, software and video-based instruction tools for wireless network education aimed at undergraduate and graduate students.
2. Support of a new national program for standardized entry-level networking education at the high school with innovative components – such as design competitions that encourage development and integration of new applications or functionalities on existing platforms.
3. Development of popular content on wireless for usual media channels as well as the Internet (consider use of on-line learning portals as well outlets such as YouTube for dissemination) to deliver a clear and unified message to the broader society that wireless technology is fundamental to many aspects of our modern digital lives.
4. Offering summer schools for wireless networks (including the fundamentals) at multiple locations with experts in the field and involving some hands-on experiences.

5. Engagement with Industry and Broader Impact

Industry impact is considered one of the important measures of success in applied science and technology funded by NSF. The goal of this section is to understand the level of engagement between academia and industry, the extent of impact wireless research is having on industry, and whether modifications are necessary in any aspect of the research operation.

5.1 Background and Prior Successes

There is a clear benefit in engaging the industry in order to have an eventual impact. The final impact could be through direct adoption (licensing) of the technologies (developed by the universities) by the industry or through direct marketing of the developed technologies through university-incubated start-ups.

Industrial impact also takes a variety of other forms: producing high-quality students, research ideas that act as seeds, influencing bodies like the FCC, and technology transfer. Participants from the industry recognized these contributions unanimously, acknowledging that many of the research activities are indeed influencing the way they think about problems and execute projects. As an example, it was mentioned that theoretical results, such as the Gupta-Kumar capacity of wireless networks, was instrumental in one case to cancel a major industry project. This saved a wastage of $7 million dollars.

Engaging the industry can be done through consortiums established in the university. Often, the monetary contribution from the industry may not be significant and might need to be supplemented with funds from government funding agencies. However, it does pave the way for active involvement of the industry and create a channel to better understand the challenges facing the industry and hence contribute towards them. Further, from the perspective of NSF, its I-Corps and GOALI programs have also helped universities in having an industry impact. While I-Corps helps research personnel to directly take their ideas and make them viable for commercialization through start-ups, GOALI allows industry personnel to collaborate actively with research personnel from the university, thereby providing a channel for easier technology transfer. Both these programs have emerged to be useful and need to be further encouraged.

5.2 Challenges and Opportunities

Broad Challenges: Realizing industry impacts entails a series of challenges from multiple bodies. First, the whole process would require a conscious decision on the part of the university personnel to devote time and effort. Several components of this process could involve time-consuming "non-research" tasks. Second, it is important to involve the industry on a periodic basis to understand the key challenges facing them as well as those that competitors are focusing on. This will help the university personnel focus on the right research challenges that are of priority to the industry, thereby aiding the adoption of their solutions later on, when they become available. Moreover, it will also expose the university students to real challenges facing the industry, thereby giving them the opportunity to work on problems that can have a direct impact in the real world. Student-driven start-ups also get facilitated by such exposure --- when students realize that multiple industries are facing the same problem, they develop confidence in the success of a start-up in that direction.

Difficulties in quantifying impacts. Despite noticeable impacts of academic on industry, there was agreement that quantification of such impacts is lacking and is difficult. Licensed research does not directly enter into products -- much engineering effort, debugging, and tuning ensues before the technology is ready for commercialization. These efforts dilute the individual contributions from a university research group. Moreover, there is no incentive inside the industry to trace the roots of success from a project and offer credit to university researchers -- this task is difficult and lacks incentives. On similar lines, impact through students is also hard to quantify, since it is so much diffused through teamwork. Individuals who single-handedly lead a project to completion may also be extra-ordinary themselves, and it is often difficult to quantify the credit that should be attributed to the university.

The dilemmas in publication-driven research. There seemed to be wide agreement that "publications" has become the goal of research. The key factor underlying this arises from the observation that there is no other quantifiable metric of research quality. Moreover, with deep specialization and inter-disciplinary research, the number of individuals who are in a position to accurately judge the quality of work is shrinking. As a result, when it comes to job-hunting, students find themselves in the best position when their CVs are heavy in publications at top conferences and journals. In contrast, pursuing a research project all the way, taking it to the point of deployability, returns far fewer papers, while running the risk that such a project may not eventually "pan out". Worse, with rapidly changing technology and new ideas supplanting older ones, the risk of investing an entire Ph.D. career on one specific project appears to be too risky. Finally, while students from absolutely top-ranked universities may risk such decisions (given that the visibility for these students are high by design), this risky strategy is undesirable to students from other universities. This is a deep and systemic problem warranting careful treatment and action.

A side effect of publication-driven research is that researchers are inclined to pick short term problems that optimize paper counts. Pursuing long-term problems is again risky and perhaps difficult to motivate against the backdrop of current technology trends. With competitive conferences and journals, it is important to be able to make a convincing case for papers -- long-term problems and solutions often lack clarity on their applicability, and are candidates for rejection. The discussion highlighted that this again stems from the lack of adequate and precise recognition of attempting/solving difficult problems. As a result, it was pointed out that often industry proves to be ahead of academia, since they are likely to be in a better vantage point with respect to the technology and other resources in the near time horizon.

5.3 Recommendations

1. In response to quantifying industry impacts, one proposition is *to trace the roots of industry success, and highlight how university research was actually influential*. An example is to trace back the roots of say, Apple's iOS software, and decompose the several components, and trace which university projects underlie them from the time of conception. *Another proposed idea was to develop a public relations team that would actively highlight to the news media various research projects that are engaging with industry and contributing to real-world technology*. For instance, if a university team licenses an algorithm to a company, there is hardly any structured way to express this to the research community. It is advisable to develop a platform where such efforts are recognized or even rewarded. Perhaps PIs could call the funding program and inform them about impactful development, even if several years after the completion of the project. An alternative way is to create the equivalent of an MIT TechReport, and incentivize researchers to report on partial or diffused successes with industry. Perhaps such environments will prod students to actually "run" with their research, to the extent possible within the walls of the university.

2. Regarding the "unhealthy" aspects of publication-driven research topics, one key observation here is the inadequate exposure that students get toward problems faced by industry. Several participants pointed out that some of the above issues could be alleviated if students could truly realize where industry struggles, which may in turn inspire them to solve a real, often difficult, problem. *To this end, it was proposed that forums or workshops, in which academics pitch their projects to an industry panel, and the latter share their problems with academicians, would be of immense value.* Designed carefully, this may serve several purposes that were discussed earlier -- students may get motivated towards start-ups, attacking real problems, and not stopping until they believe that industry would express interest.

3. In relation to student-led start-ups, it was widely believed that almost all international students live in the US on an F1 visa status, which precludes them from joining start-ups. *Changing their visa status to take on a career in start-ups is indeed a deterrent* -- *all the*

participants unanimously agreed on this. Any form of relaxation in this direction would immensely help transfer of university research to technology in the real world.

It should be also noted that academicians must bear in mind that impact often has a long "turn-around time". Thus, they should be encouraged to be patient, and to continue working on important, high impact problems, and that impact may arise suddenly at a later and unexpected point in time.

6. Suggested Outcomes

This section combines all of the discussion in prior sections to present a short and succinct set of suggested next steps.

Cross-layer and integrative research: As the demand for spectrum continue to grow in significant ways, it is important to better promote cross-layer research. While the notion of cross-layer design has existed for more than a decade, true cross-layer research has only started to grow roots in the recent years with the advent of more capable and effective programmable platforms. It is believed that orders of magnitude improvement in wireless capacity is possible by considering problems across all layers --- from hardware design, the physical layer, the control algorithms, all the way up to the applications. In particular, future cross-layer research should translate the PHY-layer capacity gain from advanced communications technologies, e.g., full-duplex, distributed MIMO and interference alignment, into network-layer performance gain that is tangible for end-users. Such research should be validated rigorously by using experimental testbeds.

Exploration of “clean-slate” spectrum access technologies: As new types of spectrum become available (60 GHz and beyond) and new modalities of access becomes popular (wide-band dynamic spectrum access), traditional access mechanisms need to be re-evaluated and re-designed. Hence, it is important to promote ideas and platforms that allow researchers unconstrained explorations in such new spectrum access technologies. A lot of research efforts in this domain limit themselves to the capabilities of existing hardware and the regulatory constraints in place today. In other cases, experimentation is limited to specific unlicensed frequencies since researchers do not have access to other spectrum outside it. To achieve significant strides outside of the present day constraints, researchers should be encouraged to explore these problems beyond these existing sandboxes. New platforms that allow experimentation across different layers --- from physical to applications, e.g., from highly capable radio front-ends to sophisticated heads-up displays --- are quite critical to the research community. Making experimental spectrum available to researchers is also important. The next generation of wireless networking technologies can be enabled through such bold and forward-looking efforts.

Building bridges to application and technology domains: Wireless networking is a significant enabler for real-world applications. While studying techniques to improve wireless networking in isolation is important, it is equally important to connect this community to synergistic application (e.g., vehicular, sustainability, health) and technology (e.g., big data, core networking, security) domains. It is expected that various societal grand challenges (e.g., zero traffic accidents, robotic operation theaters, predicting epidemic outbreaks) can be addressed when research in wireless networking technologies is complemented with domain expertise in each of these areas and deep collaborations between such communities are fostered. While various such cross-cutting programs exist amongst federal funding agencies, it is important to nurture these linkages further and make them stronger in the future.

Connecting theoretical research and practice through abstractions: It is anticipated that the research community will continue to push the envelope on technical results that hold great potential for improving performance across different dimensions. However, in order to transition good theoretical results into practical applications, it is important to ensure that suitable abstractions are used to define the right problems, and that such solutions can be more easily translated into realistic use cases. Engagement between researchers and practitioners in creating the right abstractions is particularly important for leveraging good theoretical results.

Research at scale: As the demand for wireless services continues to grow rapidly, some of the key challenges manifest themselves only at the scales in which they operate in the real world. For instance, the challenges encountered by a large, city-wide wireless operator is observable only when thousands of users attempt to access spectrum resources in the same vicinity and at the same time. To evaluate and address such challenges, it is important to support and enable research that can evaluate design of systems, services, applications, and algorithms at scale. Investments in large-scale, nation-wide testbeds, potentially with significant and realistic user populations, are necessary to meet such goals.

Enabling early exposure to wireless technologies for students: Educating the next generation of researchers is a key goal for driving successes in this domain. Given the great social impact of wireless technologies in human lives today, this domain can serve as a great advertisement for STEM careers. Hence, the community should be encouraged to develop popular content for broad dissemination among current and future students that showcases the diverse successes in society. Further, students in high schools and early stages of college education could be inspired through more accessible educational content that is integrated in their early curriculum.

Appendix --- Schedule of the Workshop

Monday, November 4, 2013

8:00 am
Breakfast

9:00am – 9:05 am
Introduction
Keith Marzullo, NSF

9:05 am – 9:20 am
Opening Remarks
FarnamJahanian, NSF

9:20 am – 9:45 am
Workshop Overviewby Workshop Co-chairs
Suman Banerjee (UW-Madison), Dapeng Oliver Wu (U Florida)

9:45 am – 10:45 am
The Critics View
Victor Bahl, Director, Microsoft Research
Bob Brodersen, Professor Emeritus, UC Berkeley
PR Kumar, Professor, Texas A&M University
Craig Partridge, Chief Scientist, Raytheon-BBN
MahadevSatyanaryanan, Carnegie Group Professor of Computer Science, CMU

10:45 am – 11:15 am
Coffee Break

11:15 am – 12:30 pm
Breakout Session 1
-**Wireless Systems** (leads: Lakshmi Subramaniam, ArunVenkataramani)
-**Mobile Applications** (leads: David Kotz, Kamin Whitehouse)
-**Wireless Networking Architectures and Testbeds** (leads: Dipankar Ray Chaudhuri, Peter Steenkiste)
-**Controls and Algorithms** (leads: EytanModiano, Sanjay Shakkottai)
-**Emerging Spectrum Access Technologies** (leads: Junshan Zhang, Ty Znati)

12:30 pm – 1:30 pm
Lunch
Speaker: John Smee, Qualcomm

1:30 pm – 2:30 pm
Initial report from each breakout with Q&A

2:30 pm – 3:00 pm
Break

3:00 pm – 3:30 pm
Feedback from critics

3:30 pm – 5:00 pm
Breakout Session 2
- **Metrics, Measurements, and Management** (lead: Dirk Grunwald)
- **Broadening scope of research** (lead: Kang Shin)
- **Security and Privacy** (lead: George Kesidis)
- **Engagement with industry and broader impact strategies** (lead: Romit Roy Choudhury)
- **Tools for education and greater outreach** (lead: Sumit Roy)

6:30 pm
Dinner
Speaker: John Chapin, DARPA

Tuesday, November 5, 2013

8:00 am
Breakfast

9:00 am – 9:45 am
Talk by Henning Schulzrinne, FCC

9:45 am – 10:45 am
Updates from breakout session 2

10:45 am – 11:15 am
Break

11:15 am – 12:15 pm
Breakout session 3 (same as session 1 groups)

-**Wireless Systems** (leads: Lakshmi Subramaniam,ArunVenkataramani)
-**Mobile Applications** (leads: David Kotz, Kamin Whitehouse)
-**Wireless Networking Architectures and Testbeds** (leads: Dipankar Ray Chaudhuri, Peter Steenkiste)
-**Controls and Algorithms** (leads: EytanModiano, Sanjay Shakkottai)
-**Emerging Spectrum Access Technologies** (leads: Junshan Zhang, Ty Znati)

12:15 pm – 1:15 pm
Lunch

1:15 pm – 2:15 pm
Plenary discussions (breakout updates and feedback)

2:15 pm – 2:30 pm
Remarks
Dr. Keith Marzullo, Dr. ThyagaNandagopal and Dr. Min Song

2:30 pm – 3:00 pm
Break

3:00 pm – 4:00 pm
Wild Ideas Breakout Session

4:00 pm – 4:30 pm
Wild Ideas Presentations

4:30 pm
Adjourn the meeting

Questions posed to different breakout groups a priori

Day 1: Breakout session 1 (these questions are common to all groups and groups are encouraged to pick the most relevant question for their topics):

- Define key past successes in your domain, in terms of impact on research, industry, public awareness and education.
- What are the broad applications and needs that are likely to drive significant wireless throughput demands?
- What novel innovations have come up that are paving the way for future research in this area?
- Identify important future challenges or approaches that can get us 1000X improvement in wireless throughput in the next 10 years.
- How can wireless research in this domain contribute and affect the emerging areas, such as Internet of Things, Cyber Physical Systems, Machine to Machine Communication, and Smart Health? What are the research problems of interest in those areas?
- What are the advances in communication and coding theory that will impact the growth of wireless throughput? How do these shape wireless resource consumption?
- Is the SDN paradigm relevant for wireless networking? What are the challenges posed by software-defined wireless systems for this domain?
- What challenges do cloud RAN architectures pose to last-hop wireless delivery? Can they be transformational in non-cellular domains?
- How is this domain addressing challenges posed by emergent Future Internet Architectures?

Day 1: Breakout session 2 (Mixers):

Metrics, Measurements, and Management (lead: Dirk Grunwald)

- What is the right balance between experimentation, analysis, and simulations?
- Is our community missing appropriate benchmarks?
- Is our science sound? Are experiments 'reproducible' ?

- Are there improved strategies for wireless management and instrumentation?
- Could we define what a “typical” wireless environment looks like?
- What is the role for ‘simulation’ of wireless networks in emerging research with its diverse application spaces?
- Can we collect systematic measurements to characterize different types of environments of interest?

Broadening scope of research (lead: Kang Shin)

- How can we enable greater research in other emerging wireless technologies --- whitespaces, DSRC, WiGig, ultrasonic and visible light communication, free space optics, mm-wave communications, completely open spectrum policy, body area networking and communication? What about inter- or intra-chip communication over wireless? Finally, very lower power RF communication systems, including energy harvesting systems?
- Is wireless going to be restricted to last-hop for the most part? What is the future of multi-hop wireless communications?
- What are the interesting opportunities to support interdisciplinary research --- cyber physical systems, smart grids, machine to machine communication?
- What is the role of big data analytics in our community?

Security and Privacy (lead: George Kesidis)

- What can PHY/MAC security do that higher layer security cannot address, or is PHY/MAC security going to be less relevant in the future?
- What are the major research directions for privacy research in an age of ubiquitous geo-location?
- Should we focus more on software-level strategies for security relevant, or hardware-level strategies be adopted widely?
- Are there current research topics in wireless security (at the network, device, communication level) that are likely to be less important in the future?
- What is the role of cross-layer research in addressing security and privacy?

Engagement with industry and broader impact strategies (lead: Romit Roy Choudhury)

- Is the academic community suitably connected to industry?
- How does the research community stay ‘agile’ and ‘lead the curve’ on developments in wireless?
- Are there other sectors of industry that we are not engaging in? e.g., smart-grids, antennas, advanced materials, safety, 3D-printing?

Tools for education and greater outreach (lead: Sumit Roy)

- How should we train the next generation of students?
- What are the challenges posed to students seeking a graduate education in wireless networking and systems? What is their current and future employer profile?
- How do we ensure the curriculum stay ‘agile’ with current research methods and encourage creative thinking?
- Are there strategies to engage the general population and help them understand the scientific advances being made in our field?
- Are there specific ways to better engage women and minorities?

Day 2: Breakout session 3 (same as Session 1 groups):

Topics to discuss further (in addition to the topics in Sessions 1 and 2):

- Do we need a ‘clean-slate’ design for wireless networking, and if so, what could that be?
- Propose a “grand challenge” scenario that can be addressed by this community and explore above mentioned future challenges in that context (try to limit this to just 1 or at most 2)

- What would be a good example of research success 10 years from today?
- How can wireless research in this domain contribute and affect the emerging interdisciplinary areas of Internet of Things, CyberPhysical Systems, and Machine-to-Machine communications? What are the grand challenges of interest to the wireless community in those spaces?
- What is missing in terms of nationwide resources to conduct the necessary research identified already?
- What collaborations and connections are missing that we should address (e..g, connections to the CyberPhysical Systems, Internet of Things, SmartHealth communities)?
- Can we learn from international partners on unique research topics and aspects that are currently poorly addressed in the US?