WIRELESS SYSTEMS RESEARCH IN 2020

Wireless systems research is the subfield of wireless research that studies networks of computing devices that communicate over a wireless medium. The main characteristic of this subfield is its emphasis on practical design, implementation, and empirical evaluation. In the last few decades, wireless systems research has played a key role in advancing wireless communication. Its contributions include: *ALOHANet,* which is the world's first computer communication network; *packet radio networks*, which is the predecessor of today's WLANs; Carrier Sense Multiple Access (CSMA), which is the mechanism used by today's computers to access a wireless network; and Software-Defined Radios (SDRs), which are the key enabler of cognitive communication systems. Recent years have witnessed a great momentum that led to innovations along multiple axes: (i) disruptive network designs have emerged including sensor networks, delay tolerant networks (DTN), vehicular networks, and cognitive networks, (ii) system research has shown an unprecedented breadth; its innovations now encompass signal processing at the physical layer as well as novel environment-aware applications, (iii) and a long-awaited convergence between empirical system research and information theory has arisen delivering the first implementations of wireless network coding, physical-layer security, and interference alignment.

GREAT OPPORTUNITIES

Wireless systems research has traditionally proved its ability to cut across many disciplines and connect hardware and radio advances with useful applications and social services. Further, because of its emphasis on practical design and empirical evaluation it has proven the vehicle that delivers basic science to industrial success. As a result, it is today positioned to contribute to multiple critical services that can affect all aspects of human lives such as:

HealthIT: Body area networks (BAN) allow a continuous monitoring of a patient condition in her/his own natural environment. Small radios with sensing devices are likely to be implanted under the skin to measure various body functions and report them over a wireless network to wearable radios. The latter collect the measurements and deliver them wirelessly to a central unit, eventually to reach appropriate medical experts in a hospital. This process extends healthcare beyond the hospital, reducing its cost, and significantly increasing its efficiency.

Broadband for all: High bandwidth Internet access is becoming a necessity in today's world. Various studies have shown that people with rich Internet connectivity have access to diverse information and use such information to significantly improve their social and economic conditions. Today, barely 25% of the world's population has any Internet access, leading to a widening standard-of-living gap with the remaining 75%. This inequality has created a new form of socio-economic disparity, referred to as the digital divide. Given the prohibitive costs of deploying a wired Internet infrastructure to every corner of the inhabited world, it is anticipated that a diverse set of wireless technologies will be the only feasible means that can realize the `broadband for all' dream.

Smart Grid: Efficient and clean energy usage may be improved with a better understanding and tighter control of energy consumption. A smart grid enables energy sources, carriers, and consumers to better understand and control various processes in the energy grid. Hence, it can help reduce waste, improve efficiency, and divert energy usage during emergencies to where it is most needed. This vision requires entities in the energy grid, including each individual power outlet, to be remotely accessed and manipulated. Because of its tetherless nature, wireless networks have been proposed as the likely medium for such communication.

Smart transportation systems: People regularly move from place to place as never seen before in human history. Wireless communication systems are likely to play an important role in making various modes of transportation safer and more efficient. Communication between pilots and air traffic control, navigation satellite systems, safety mechanisms that detect driver fatigue, and systems that can provide automated safeguard between vehicles, all depend on improving the robustness and quality of wireless communication.

Understanding our environment: An increasing number of mobile devices are equipped with a rich array of sensors, including cameras, microphones, accelerometers, electronic compasses, and GPS receivers. While each such sensor can facilitate new applications in these devices, in aggregate they open up unprecedented opportunities of understanding our environments. If data from billions of such sensors can be effectively collected, they can serve as a rich information telescope for the world. For example, chemical sensors in future mobile phones can track spread of pollutants within a city, feeds from user cameras can provide immediate feedback on outbreak of forest fires and their nature, and GPS receivers in phones can pinpoint locations of trapped individuals after a devastating earthquake.

GREAT CHALLENGES: INFINITE BANDWIDTH, INFINITE MOBILITY, INFINITE DATA

With great opportunities come great challenges. The applications described above place greater demands on wireless networks that are hard to satisfy with today's technologies. We envision that research in the coming ten years should try to answer the following three big challenges:

- How do we obtain unbounded wireless bandwidth?

So far, wireless bandwidth has always lacked behind wired bandwidth. But the future is likely to make wireless the dominant medium for connectivity; not only will it be used for Internet access but also it will replace the piles of wires at home and in the office. The scarcity of the wireless spectrum is the main factor that limits the vision for wireless access everywhere anytime, and perhaps the toughest challenge that wireless research has to undertake. Overcoming this challenge requires innovations in various areas including novel ways for spectrum sensing and reuse, mechanisms for using higher frequencies such as 60 GHz and the visible light, novel ideas for dealing with interference, high-density MIMO, and generally more efficient protocols and systems.

- How do we deal with fast and continuous mobility?

While full mobility is the main premise of wireless communications, current wireless networks fall short of delivering that vision. Today, 802.11 and other computer networks provide a nomadic mode in which users can occasionally move from one place to another. They cannot however support fast continuous mobility. Cellular networks are better at supporting mobile users but they pay for this feature very highly in terms of efficiency with today's 3G systems having a low spectral efficiency about 1 bit/s/Hz. The problem is that fast movements cause unpredictable variations in channel SNR that could be as high as 15-20 dB. These SNR variations make it difficult to select the right modulation and FEC for transmission. The net result is that existing wireless networks either give up on fast mobility (802.11 WLAN) or resort to a very conservative design (cellular networks). Addressing this problem requires innovations at all levels of the network stack.

- How does a network deal with humongous amount of data?

The future will carry a plethora of sensors that monitor every event anywhere. Some will be carried by human users including cameras and microphones; other will be embedded in the environment including sensors that monitor our energy grid, our factories, and our transportation systems. The result is a flood of sensor readings, some of which have a very rich content including video. Understanding this information, and eliminating the huge correlation and redundancy constitute a major challenge. Furthermore, given the size and speed at which information is generated, the network cannot simply carry the information to a centralized destination where it is processed. The network itself has to actively participate in processing the information to eliminate redundancy locally, transform the information into more manageable formats, and route it to where it is most useful.

GREAT INNOVATIONS REQUIRED: GET PHYSICAL, GET VISIBLE, AND GET AWARE

In order to address the above challenges, wireless systems research needs to innovate along multiple axes.

GET PHYSICAL

Connection between the PHY and higher layers: Conventional wireless design has a strict contract between the layers of the network stack. Specifically, the PHY and lower layers deliver fully correct packets and the higher layers route these packets to interested destinations. Recent studies have shown that such strict separation between layers leads to lower throughout and reliability. Future applications are likely to emphasize this trend. For example, implanted and wearable radios proposed for HealthIT cannot afford inefficiencies that stem from a traditional layered network stack. Innovative designs of a `layerless' network stack that avoids most protocol overheads while maintaining robustness are highly desirable. Significant research is needed to design systems in which applications and other traditional higher layer components are made aware of the physical layer characteristics, i.e., a PHY-aware stack. Similarly, it is also important to seek innovative PHY layer designs that leverage application context and treat the entire network as the communication channel.

Connection to theory: The battle to push network efficiency to capacity limits needs to continue unabated. Breakthroughs that will impact practical systems are expected to arise from collaboration between information and communication theorists and the experimentalists represented by wireless systems researchers. This is likely to leverage ongoing advances and collaboration at the physical layer, including those in the areas of interference alignment, interference cancellation, and high density MIMO. The net result could be an information-theory driven practical PHY design, where the most advanced theories are realized under practical constraints of feedback overhead, limited synchronization, and unknown and quickly changing environment.

Connection to devices and technologies: The future also depends on innovative uses of advanced wireless technologies and devices in networked systems. Promising ideas include using 60 GHz links to eliminate wiring in datacenter, or the use of UWB radios for spectrum sensing and access. Other innovative uses of directional antennas, MIMO, and energy-scavenging radios are also important.

Connection to security: Physical-layer security is a new field that has both systems and theoretical foundations. It advocates new PHY-based security mechanisms that augment the existing repertoire of cryptography-based techniques. For example, the fading patterns on a wireless channel can create a source of randomness to hide private information, the characteristics of the wireless channel can be used to generate device identity, and frequency hopping can be used to combat jamming.

GET VISIBLE

Free space optics and high frequencies: Spectrum scarcity will continue to be the biggest challenge for wireless networks. As a result, there is a great interest in exploiting higher frequencies such as those around 60 GHz and in the visible light range. Exploiting this part of the spectrum, however, requires novel network solutions to overcome fast degradation of the signal over distance, and its inability to penetrate obstacles in the line of sight. Innovations could cast traditional network protocols in a completely new light such as having visible-light networks route data using controllable mirrors and lenses.

Spatial spectrum: MIMO technology promises a linear gain in capacity with the number of antennas. Unfortunately, for RF communication, this linear gain stops after a few antennas. This is because the linear gain is dependent on the antennas being separated by at least half a wavelength. For 2.4 and 5.5 GHz networks, the wavelengths are on the order of tens of centimeters, which prevents packing many antennas on a small device. In contrast, the wavelengths in the visible light range are much smaller allowing the packing of millions of antennas on a small device. This is apparent from today's cameras and LCD, which carry millions of pixels. These LCD and cameras can be used as high density MIMO transmitters and receivers, capable of delivering much higher throughput than existing RF MIMO channels. Leveraging such channels requires innovative research that combines networking with image processing techniques. Innovations in this direction could extend the existing frequency spectrum with a spatial spectrum that captures spatial frequencies. Further, since transmissions in the visible light range are highly directional and do not suffer major interference, the spatial spectrum expands with the physical space.

GET AWARE

Cognitive communication: Today, the available spectrum is fragmented across a number of incumbents governed through strict control mechanisms. Approaches to facilitate dynamic access to the spectrum in an unlicensed manner without interfering with existing incumbents can yield a significantly more efficient spectrum usage. Cognitive communication across all spectrum bands has the potential to enable broadband-for-all. There are many ongoing activities in the domain of spectrum sensing, spectrum access, and cognitive radios that are focused towards this goal. More research is needed however to address questions like: how do we detect and avoid occupied frequencies? How can a user communicate efficiently over a highly fragmented spectrum? How do we leverage frequency selectively for increased throughput?

Context awareness: Wireless networks can leverage context-awareness to decrease overhead and increase information utility. For example, nodes in a MANET may use their accelerometer readings to detect mobility, which gives a coarse indication of the speed at which the channel may change, and hence how fast the routes need to be recomputed. Also, in a participatory sensing application, roaming users and vehicles may participate in monitoring the occurrence of a forest fire. To decide whether this application should report data, the node has to realize the unusually high temperature around a specific region. This might cause the camera to be turned on to take a picture. Further, the node uses its GPS readings to geo-tag each piece of information. Finally, all other users in the area close to the forest fire get an automatic alert based on their geo-locations as well. Innovative ways for leveraging context-awareness can lead to novel and useful applications.

From data-networks to information-networks: Future sensing applications are likely to generate a humongous amount of data to be communicated over the wireless medium. Much of this data however is highly correlated and, in principle, could be compressed. The problem however is that compression needs to occur at the sensors themselves before the data consumes the wireless bandwidth. Said differently, there is a need to move from traditional network design, which focuses on packet delivery to a more encompassing design that addresses information delivery. Recent advances in compressive sensing have the potential of addressing this issue but they need to be integrated with network design. Future research should provide innovations that allow a network to actively participate in compressing the data, manipulating the information, understanding it, and routing it to where it is needed.

REQUIREMENTS: PLATFORMS, BENCHMARKS AND EDUCATIONAL SUPPORT

Various platforms and activities can enable the wireless community to address the above challenges and produce the necessary innovations.

Testbeds and platforms: Many past successes were enabled by the accessibility of testbeds and wireless platforms. Without implementation measurements and testbed evaluation, wireless research risks losing its ability to provide practical and realistic designs. However, over time, the complexity and hence the cost of studied wireless technologies has significantly increased. For example, while a traditional wireless node composed of a PC and an 802.11 card costs a few hundred dollars, a MIMO software radio can cost as much as ten thousand dollars. To continue providing cutting-edge experimental wireless research, the funding agencies should incorporate this cost in their model.

Additionally, the choice of platform will impact whether other researchers can reproduce the results or leverage the produced software. It is important that there is some form of standardization in choice of hardware, software codebase, and testbed topologies. Much like NSF NOSS program supported and encouraged the use of Berkeley Motes and FIND program will encourage use of GENI, wireless systems research should encourage use of common platforms. This will build a consistent and coherent body of research results, and ensure high-impact of the supported research activities.

Reproducibility: To build a consistent body of research results, which can be reproduced, verified and advanced by independent researchers, it is important to develop common metrics and benchmarks. While there is a common understanding of relevant metrics such as loss rate, throughput, and SNR, it is difficult to translate these metrics into concrete benchmarks because such benchmarks need to involve the whole network setup and the interaction between different nodes. Furthermore, no good characterization exists of the common case behavior, e.g., typical loss patterns or hidden terminal probability in operational WLANs. Characterizing the range of behavior in existing wireless networks is an important area that needs major research efforts. Such characterization however requires monitoring a large number of academic and enterprise networks of various sizes. Thus, it will require the involvement of a large fraction of the community and the support of the funding agencies.

Sustainability: A major problem that hampers efficient and fast research outreach is that every research group has to build its experimental effort from scratch. Ideally, codebase produced by a research project is made publically accessible to other researchers. For this to happen however research groups need engineering staff, who can package the software for public release and maintain the package over time to continue to work with new software releases. A similar argument applies to the sustainability of platforms. More generally, projects and platforms that have significant community impact and are crucial for the success of wireless systems research should be supported as community assets beyond project support years. This will ensure that there is an uninterrupted support of the hardware and codebase.

Educational support: Traditional courseware describes wireless systems by partitioning them into silos (e.g., network protocols and applications vs. information and communication theory) and do not cover the inter-disciplinary skills crucial for future cutting-edge wireless research. There is a need for inter-disciplinarily classes as well as hands-on courses and workshops to equip students with a complementary set of skills that allows them to understand and redesign the whole network stack from the PHY layer to the network and application layers.