Final Report from the NSF Visioning Workshop on Extreme Wireless Networking

Held in Salt Lake City, Utah on October 15, 2017

Workshop Chairs

Marco Gruteser (Rutgers University) and Kyle Jamieson (Princeton University)

Lead Workshop Scribes

Chenren Xu (Peking University), Pengyu Zhang (Stanford University),
Xinyu Zhang (UC San Diego)

Participants

Ashwin Ashok (Georgia State University)
Chenren Xu (Peking University)
Dinesh Bhardia (Stanford University)
Elahe Soltanaghaei (University of Virginia)
Eric Rozner (IBM Research)
Eugene Chai (NEC Laboratories America)
Giovanni Pau (UMPC)
Haitham Hassanieh (UIUC)
Heather Zheng (University of Chicago)
Jie Xiong (Singapore Management University)
Kannan Srinivasan (Ohio State University)
Kyle Jamieson (Princeton University; Chair)
Lili Qiu (UT Austin)
Marco Gruteser (Rutgers University; Chair)
Pengyu Zhang (Stanford University)
Robin Kravets (UIUC)
Shubham Jain (Old Dominion University)
Sneha Kumar Kasera (University of Utah)
Swarun Kumar (CMU)
Wenjun Hu (Yale University)
Josh Smith (University of Washington)
Xia Zhou (Dartmouth University)
Xinyu Zhang (UCSD)
Yan Wang (SUNY Binghampton)
Yingying Chen (Rutgers University)

Table of Contents

Introduction

The wireless technologies we discuss in this workshop are key to enabling myriad transformative applications of information and communication technology (ICT) that hold the promise of a profoundly positive impact on business, government, and citizens.

Smart Cities. Wireless technology, in particular long-range, low-power wireless is key to enabling the vision of the Smart City, where transportation infrastructure, communication infrastructure, and utilities infrastructure are transformed to be more resilient, efficient, and sustainable through the application of ICT.

Smart Manufacturing. Improving the efficiency of assembly-line manufacturing holds the promise of less costly-produced products, increasing US manufacturing competitiveness, and eventually improving the standard of living for people domestically and worldwide. Wireless technology holds the promise of transforming manufacturing with ultra low-power wireless communication techniques that can locate, image, and inventory items as they move through the manufacturing chain and even beyond, into warehousing, and shipping logistics chains, eventually reaching the consumer.

Virtual Reality and Augmented Reality. Applications such as remote surgery, online learning and job training, as well as immersive online retail all benefit from effective virtual reality (VR) and augmented reality (AR) technology. High capacity wireless networking under high mobility allows significant improvement in the size, weight, and battery life of wearable VR / AR devices by offloading rendering tasks to servers without constraining user mobility.

Automated Vehicles. Wireless networking can play an important role in developing robust driver assistance and automated vehicle technologies. While current automated vehicle prototypes and emerging products demonstrate impressive performance, it is widely recognized that fully realizing level 5 automated driving over a broad range of driving conditions remains a long-term challenge. Low-latency and high-bandwidth communication between vehicles and with the surrounding traffic infrastructure could enhance situational awareness of a vehicle beyond its immediate line of sight through sharing of rich sensor data. This can provide important information that allow addressing a number situations that could otherwise result in accidents. It will also allow coordination across vehicles and traffic signals to enhance the overall efficiency of traffic flows.

Goals of the Workshop

This workshop is part of a series of visioning workshops that strives to enable more focused and more regular research challenge discussions within the networking community in a one-day format that is co-located with major conferences. The overall goal is to discuss grand research challenges relevant to the community.

As this workshop was co-located with the ACM MobiCom Conference, we placed particular emphasis on the following technology trends:

- Extremely High Mobility: Technology trends such as drones or automated and connected vehicles increasingly generate network load while actively in motion. On the flip side, they also offer new opportunities for exploiting moving network infrastructure.

- Extremely High Frequency: Providing ever higher capacity is expected to require extremely high frequency communications such as mm-Wave or visible light.

- Extremely Low Power: The desire to provide coverage for an increasing array of Internet of Things devices is leading to demand for lower data rate coverage at extremely low power.

We asked participants to consider the above from a practical, systems and networking research-oriented perspective that is well aligned with SIGMOBILE.

For each of the above areas, we asked the workshop participants first to identify and discuss the most exciting research opportunities in these areas. Second, we asked them to formulate major, 10-year horizon research initiatives that the community can work on together in the mid-term future. Third, we asked participants to design “Grand Challenge” competitions relating to these research initiatives that could be held at a future MobiCom conference to spark further interest and research progress in the broader SIGMOBILE community on the research initiatives identified here.

To summarize the outcomes of the workshop, this report documenting our discussion of these questions provides feedback to the NSF as well as the greater research community.

Structure of the Workshop

To enable focused discussions in smaller groups, we invited 20 participants to this workshop, considering factors such as expertise in the key topic areas and diversity of the participants.

Workshop participants received a detailed briefing about the goals of this workshop. In particular, that the goal of the discussions is to take a longer-term view and to identify grand challenges for the community that will remain important in a 10-20 year timeframe.

Workshop participants discussed each of the three topics in sequence for about 2.5 hours. The discussions of each topic started with 2-3 panel-style presentations covering topic introductions and initial grand challenge proposals, which were intended to stimulate further discussion. After an initial plenary question and answering period, participants joined three breakout groups, to discuss key questions as well as to develop a vision for a grand challenge and suitable competition under this topic. As part of this breakout discussion, participants were asked to consider the state-of-the-art in each area, to identify low hanging fruit, meaning problems that can be expected to be solved over the coming years, and to define fundamental grand challenges that will remain significant ten years on. Each topic discussion concluded with reporting to all workshop participants and a final plenary discussion.

The following sections synthesize the discussions of the breakout groups for each of the three workshop topics: Extremely High Mobility, Extremely High Frequency, Extremely Low Power.

Topic 1: Extreme Mobility

As wireless technology increasingly pervades the physical world, it is enabling the emergence of WiFi connectivity in devices with high mobility including cars, trains, and airplanes. On the other hand, the usage of drones for surveillance, object transportation, or coordinated constructions as well as autonomous vehicles have necessitated a reliable V2V/V2I wireless communication. Although cellular networks provides a reliable internet connection, or wireless communication protocols such as DSRC are designed for mobile vehicles, they only support limited bit rate and are not scalable to the potential high user demands in near future. On the other hand, existing wireless protocols (e.g. TCP) are not suitable for mobile platforms due significant variations in the physical arrangements. In terms of networking and topology design, adhoc solutions are considered as the state of the art methods, but they are designed for individual classes of mobile platforms (e.g drones, autonomous vehicles) and manufacturers (e.g. DJI, Tesla).

Current State of the Art

Applications

High mobility support is necessary in multiple platforms, such as realtime VR, autonomous driving, and drone management.

Realtime VR: The goal of VR is to provide an immersive experience, to enable users to be fully integrated into the virtual environment. In order to maximize the reach of VR platforms, it is thus necessary to ensure that VR can be used in different and diverse environments. For example, a VR experience of a First-Person Shooter game will involve sudden, high speed movements and gestures. However, existing VR platforms fall short of this goal since existing wireless networks are unable to provide the high bandwidth, low latency links needed to track rapid, user movements.

Autonomous Driving: Autonomous driving platforms today are managed through cloud services (e.g. fleet tracking, mapping) that do not offer the critical features (e.g. cooperative obstacle detection, collision avoidance, high speed maneuvering) needed for full, level 5 autonomous driving.

Drone Operations: Drones are an important, upcoming platform for providing multiple services, particular when fixed infrastructure is degraded or unavailable (e.g. in disaster scenarios). Current approaches are focused on single-pilot, single-drone applications, with limited, support for multiple, high mobility drone solutions.

Protocols

The most popular high mobility communication protocol is the global system for mobile rail (GSM-R), which supports maximum data rate of 200 kbps, while it could only support train operation control. The fourth generation (4G) long term evolution (LTE) system is designed to provide different quality of services for low to high mobility. However, experiments have shown that the existing 4G systems can provide the maximum data rate of 2-4Mbps. Dedicated Short-Range Communication (DSRC) is a wireless technology that has been designed to support vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications. However, the small available frequency bandwidth limits DSRC to a small number of users and is not scalable to dense network environments.

Low Hanging Fruit

Drones and high-speed train will likely to play important roles in addressing the massive highly mobile Internet access in either planned or unplanned mobility in the near future.

Drone: Being mobile in 3D space by nature, low-altitude (< 300 ft.), low-speed (< 50 mph) drones can provide many new applications or provide solutions for existing problems with much lower cost. For example, in the foreseeable future, drones will be used to perform video surveillance for border protection, search-and-rescue in disaster, high-tech agriculture, and wireless connectivity. Specifically, for wireless connectivity, drones may serve as a mobile AP and distribution system to dynamically provision the access network to provide fast handover, plan the coverage and capacity for vehicular hosts and thus optimize the QoE of Internet applications such as video-on-demand and skype calls. In planned mobility scenario where the vehicles' moving trajectories are predetermined or predictable, the drones can even move together with the vehicles to maintain a stable "backhaul" link for them. Drones are also expected to form a mesh network to extend communication range and improve the routing problem.

High-speed Train: The high mobility nature presents communication and networking challenges rooting from large Doppler shift, coverage gap, frequent handoff, and the tidal effect -- hundreds of clients migrating from one base station to another. One of the unique characteristics is that the mobility pattern of high-speed trains is highly predictable. Leveraged properly, this has great potential to improve the overall throughput for on-board passengers because the throughput of the cellular link would be location-dependent. As most of today's high-speed trains already provide Wi-Fi access through a cellular/LTE gateway, it is natural to upgrade this mobile relay infrastructure. From the perspective of link technology, one can leverage massive MIMO, millimeter wave radio, and free-space optical techniques to provide a multi-Gbps link solution. Caching as a service should also be used to facilitate the multicast and prefetching.

Big Challenges

While drones are deployed to accomplish different tasks, but two communication modalities are common: drone-to-ground and drone-to-drone. Our big challenges will thus revolve around these two main themes.

Drone-to-Ground: Indy Car Race Tracking

The big question for this challenge is: how well can the drone maintain connectivity to a high speed indy car?

The typical backhaul of existing wireless networks is assumed to be wired (or using fixed wireless), reliable and high bandwidth. Due to the high mobility of the drone platform, wireless links will necessarily form an integral part of the drone management plane. It is thus important that we design an efficient, reliable backhaul network over unreliable wireless links.

The challenge will have two components: drone backhaul provided by fixed wireless installations, and drone backhaul provided by high speed, mobile base stations with potentially unpredictable.

The solution will be evaluated according to their ability to:

(a) Maintain a "connected" string of drones from one side of the airspace to the other side. Provide maximum end-to-end throughput across the string.
(b) Move the string of drones, maintaining connectivity and throughput.
(c) Drive a car under the string of drones, maintaining connectivity with at least one drone at all times.

The quality of the solution will be judged according to the following metrics:

(a) Data rate: what rate between the drone and the base station (either fixed or mobile) can be sustained?
(a) Reliability: how resilient is the wireless management link to mobility of both the drone, and possibly the mobile base station?
(b) Latency of the connection: how well can the wireless management link support low latency services such as VoIP calls, video streaming, augmented reality, and drone fleet management applications?

Drone-to-Drone: High Bandwidth, Low Latency

In this challenge, participants design a high bandwidth, sub-millisecond latency communications platform for V2X (e.g. drone-to-drone) communications.

Sub-millisecond latency is necessary for real-time interactions between mobile platforms, for real-time mission critical applications such as mobility coordination, obstacle avoidance and overall network management. Such networks must also have high bandwidth to support management and monitoring services (e.g. vision-based navigation).

The solution will need to include innovative approaches to manage data streams of different priorities between drones. For example, data exchange for collision sensing and avoidance must operate at a lower latency than simple location updates.

The solution will be evaluated in the following environments:

(a) Multiple drones intentionally flying in collision paths. The solution must be able to detect this collision, communicate with each other and stop the forward motion. Drones can be encased in lightweight protective frames to avoid damage in the event of a crash.

(b) A fleet of drones maneuvering through a dense obstacle course (e.g. jungle gym, forest). Drone-to-drone communications is used to coordinate movement through the course. The goal is to clear the obstacle course in the smallest amount of time.

The outcome of the challenge will be evaluated according to the following two metrics:

(a) Latency: what is the minimum latency that can be supported by the drone-to-drone channel?

(b) Throughput: What is the maximum and variance of the bandwidth that can be supported between any pair of drones?

Topic 2: Extremely High Frequency

Current State of the Art

Extremely high frequency spans a range of bandwidth starting from the mmWave frequencies (30 GHz - 300 GHz) and moving up all the way to visible light (100s of THz).

mm-Wave: There has been a lot of exciting activity in industry and academia in the area of extremely high-frequency wireless networking. At the mmWave band, standards like 802.11ad have already been ratified, and commercial products have been available on the market for years. In addition, new standards like 802.11ay and 5GPP are being rolled out. Commodity 802.11ad products already demonstrate 2+ Gbps of throughput, with point-to-point static links.

However, these products are not designed to handle high mobility, they usually take 300 ms to re-steer the beams. The current use cases are limited to cable replacement and wireless backhauls. There remains a gap for a technology in the context of 5G and beyond, where there will be more demanding applications such as instant file sync, v2x/v2v context sharing, uhd virtual reality, 3d video surveillance (high bandwidth, dynamic traffic, low latency).

Visible Light: Recent years have recognized the opportunity of repurposing LED lights not just for Illumination but also for communication and sensing. Modulation of light is very different from that of RF signals, which uses strategies such as OFDM or phase modulation. In visible light, modulation mostly relies on changing the light intensity. State-of-the-art solutions have demonstrated kbps of bandwidth with off-the-shelf LED lights. Laser-based communication and free-space optics are much better options for high-speed communication (tens of Gbps) but the cost of the systems is significantly higher.

While solutions to increase the bandwidth of VLC require innovation both in hardware and associated software, presently hardware remains the most significant bottleneck. Indeed, despite the wide bandwidth of visible light (10,000x higher than that of the radio frequency band), achieved data rates tap into a small fraction of available link capacity. This motivates pushing the envelope on two present realities, based on state-of-the-art literature: (1) Under cost constraints, akin to a standard LED light bulb and coupled with the requirement of providing illumination, current bandwidths are limited to few kbps; (2) Removing cost constraints, and designing custom laser-based free-space optics like communication, achieve state-of-the-art bandwidth of at best 10 Gbps, several orders of magnitude below the full potential of visible light (hundreds of Tbps, in theory).

Low Hanging Fruit

(i) mm-Wave: Low-hanging open research questions in the space of mm-wave networking start from the very basic unit of the point-to-point link which needs better optimization under mobility, disruption and blockages.

Beam-steering in Point-to-point Links: mmWave networks use an order of magnitude narrower beams than most prior directional microwave systems, which leads to new challenges in beam-steering. The point-to-point link rate needs to be maximized under two important constraints: (1) high mobility, where the two end nodes move or rotate relatively at least at pedestrian speeds or higher, and (2) disruption, where moving obstacles, such as humans and pets, appear in between the transmitter and receiver. More efficient beam steering algorithms are needed, with much lower latency less overhead than the naive exhaustive search approach, presently deployed in 802.11ad. New radio architectures like mmWave MIMO (e.g., in 802.11ay) may help alleviate the problem. In addition, out-of-band information may help the decision making. For example, with an on-board antenna array at the node, the angle/bearing of the signal source can be obtained to guide more efficient directional transmission. Inertial measurement units can also be useful to track node orientation and steer beams accordingly.

Support from the Higher Layers: The higher-layer protocols need to be jointly designed to improve system level performance, because any outage or bandwidth variation at the low layer tend to be amplified at TCP or application level. Developing a holistic protocol stack that is mm-wave aware is an open problem space.

Bridging the Gap between WiFi and WiGig: Assuming that commercial devices of the future will be both WiFi and WiGig enabled, what would be the correct balance between the two and how do they interplay on user devices? Would the role of WiGig be that of a backhaul or would it function well under mobility as well? How can WiFi serve to complement 802.11ad in terms of ensuring reliability and a minimum guarantee on performance? Further, there is a great potential for several the core protocol ideas (rate adaptation, frequency selection, etc.) already well developed in the WiFi to be re-designed in the mm-wave context.

(ii) Visible light:

Smart Lighting: Lighting has reportedly consumed about one fifth of the world's electricity. An effective way to reduce this high energy footprint is to adopt solid state (LED) and smart lighting. However, smart lighting affects the throughput of the VLC system greatly. Thus, careful design of new modulation schemes is required to achieve fine-grained dimming for smart lighting and at the same time maximize the throughput for VLC.

VLC sensing applications: One can envision a variety of sensing applications, such as body tracking, highly-accurate localization, occupancy sensing and health/fitness tracking using VLC.

Many such applications have tight cost constraints, especially when deployed at large scale. So a big challenge lies in how to leverage incumbent light infrastructure, or designing smart lighting/sensing systems with the minimum retrofitting cost.

Big Challenges

(i) mm-Wave: The major challenges in the mm-wave space stem from looking at mm-wave as a network, as opposed to a series of point-to-point links.

Scalability: Running dense mmWave networks for crowded large spaces (stadiums, conventions, concerts) is challenging. In fact, imagine "painting" the mmWave antennas on the ceilings or walls in large buildings. Due to hundreds of beam directions per radio, the problem escalates the interference management problem to a scale that goes far beyond traditional omni-directional wireless systems. In addition, the coordination between densely deployed access points or base stations can be harnessed to improve the reliability for mobile users. To accelerate such mmWave network research, it is urgent to develop large scale, indoor and outdoor mm-wave testbeds and an experimental platform that can be reused for controlled, reproducible experiments.

High Mobility: All the mmWave challenges will be compounded under mobility scenarios that significantly surpass pedestrian speeds, e.g., beaming mmWave signals to drones, cars, trains, etc. Example research questions involve: How to realign the Tx/Rx beams efficiently without exhausting the precious channel time? How to map the interference and schedule the directional transmissions to maximize spatial reuse? The predictability of routes for such mobile devices may be a niche property that can be leveraged to improve the beam steering efficiency.

mm-Wave Sensing: On the other hand, the sensitivity to blockage, mobility and environment may be leveraged to enable interesting wireless sensing capabilities. For example, mmWave signals may be reused for location sensing, physiological status monitoring, human-mobile interaction, and robotic control.

Coordination to Avoid Blockage: Blockage is another critical issue. The current 802.11ad mmWave links tend to be broken if the link is blocked by body (due to hand-holding position changes, device being put in pocket, or blockage of people passing by). Coordination of densely deployed APs may help resolving the issue, but the deployment cost, and backhauling may become an immediate constraint. Self-backhauling can potentially reduce the cost substantially, by using dual mmWave interfaces on each base station, for backhauling and access link respectively. The challenge lies in dynamically adjusting the backhaul topology, and coordinating the backhaul and access links (through channel allocation or MAC) to avoid interference.

Extremely Low Energy-per-bit mm-Wave: The mmWave radios tend to have much higher power consumption compared with low-frequency counterparts, because of the wide-band data converters and power-amplifiers being used. Although owing to higher bit-rate, the average bits per Joule can be even higher than low-frequency radios, this is true only if the mmWave MAC and higher layer protocols can maximize the link utilization, i.e., minimizing link idling time, and timing spent in coordination, signaling, channel sensing, along with beam steering. Innovative solutions need to be designed to fit mmWave for extremely low power devices such as IoT sensors or wearables. Existing work has proposed single-bit data-converters ad PA-free mmWave radios, but the hardware feasibility and compatibility with other mmWave radios need further investigation.

(ii) Visible Light:

Maximize VLC rate per unit cost: Today's VLC light bulbs are restricted to a few kbps, while free space optics (FSO) achieve 10 Gbps at significantly higher cost. Designing a system that achieves the bandwidth of FSO and beyond at the cost of an LED light bulb is a challenge.

Between mm-Wave and VLC: The present discussion on mm-wave is largely centered around the 60 GHz frequency band. There is a need to investigate mmWave bands other than 60 GHz, and supporting multiple mmWave bands simultaneously (e.g., spanning 28 GHz and 60 GHz); Indeed, Terahertz frequency bands lie in between the GHz bands and visible light and are yet to have mature wireless frontends that can function reliability even over small ranges (few meters). Once the radio technology is mature, this will open up new opportunities to rethink wireless communication and sensing systems in the THz context.

Grand Challenges

The discussion groups proposed following two ideas for the Grand Challenges:

Idea 1: Seamless coverage and mobility support for mmWave networks. This challenge aims to encourage competing groups to develop algorithms that push the limit of the mmWave networks to the maximum throughput, maximum link availability, and minimum latency.

Two evaluation scenarios will be involved in the challenge, one outdoors and one indoors. A limited number of access points will be setup on the ceiling, walls and building infrastructure. The mmWave mobile devices could be carried by robots or drones moving on designated trajectories with predefined orientation changes.

We will use human-sized robots (filled with water to emulate human body) to randomly block the LOS. The groups will design their system competing with each other to achieve the best performance in terms of the link outage ratio, throughput, recovery latency from blockage, application related metrics (video stream quality over time). An example metric can be: "Can 100 users each get a reliable X Gbps of data rate 99% (or 99.9999, etc.) of the time, as they move around". Each team should aim to maximize the number X. Alternatively, this X metric can be broken out into multiple sub-metrics, like link outage ratio, median throughput, link recovery latency (time from blockage/mobility occurs to the time when the link is reestablished with maximum quality).

Different constraints can be imposed for a fair comparison under realistic settings. For example, we can have 1) Free-style design, with customized hardware and without considering the hardware constraints; 2) Cost-limited design, where teams try to maximize the metric X per unit cost. For this, the teams need to minimize the number of access points used to cover the test space. 3) platform constrained design, where teams work on standard-compatible hardware, but with reasonable reconfigurability at MAC layer and above, to maximize the X metric.

The competition can be carried out across multiple phases with increasing complexity: 1) All nodes are static; 2) Nodes are mobile, at different speed levels, to emulate different usage scenarios, e.g., human users holding mobile devices, or mmWave radios mounted on vehicles on freeways or in urban environment; 3) With robots or people roaming around in the environment intentionally blocking the static nodes; 4) Both the obstacles and radios are mobile. For outdoor tests, drones can achieve high moving speed to test the extreme mobility scenarios.

Idea 2: Designing a system supporting symmetric high data rate using high-frequency bands, either mmWave or visible light, entailing low latency, in the presence of possible blockages from other surrounding objects (i.e., users).

More specifically, the grand challenge is enabling synchronized dancing with dancers wearing VR/AR headsets. The idea is to have groups of dancer pairs, where each pair of dancers need to synchronize their dance moves based on what each sees (on the VR screen). Each dancer wears a VR/AR headset, streaming video to the other dancers in the pair, so that they can synchronize their dance moves. The VR/AR headsets also receive a large amount of data from the infrastructure. Thus, we need symmetric high-rate, low-latency links for the VR/AR headsets.

The participants of the challenge need to come up with solutions to provide these links, using either mmWave or VLC, while considering that other dancers are moving around and can potentially block the links. The participants are provided with a fixed budget, so cost is a constraint. The systems challenges include how to enable high rate in both directions, how to effectively track other dancers to steer the beams and overcome the blockage.

The evaluation metrics include the maximal number of dancer pairs the system supports, the latency a dancer perceives when receiving streamed data from the other dancer in the pair, and the aggregated throughput, all within the limit of the provided budget.

This grand challenge has broader implications beyond the synchronized VR dancing itself. Symmetric high bit-rates may be required in demanding mobile applications in 5G and beyond. For example, the future autonomous vehicles need to share the 3D scenes with each other and merge them to form a bird's-eye-view of the ambient environment. Or, they may need to upload each of their own scenes to an infrastructure (e.g., a base station with edge computing capability), who will then merge the scenes and stream them through the mmWave downlink to other users.

Topic 3: Extremely Low Power

Current State of the Art

Today, users are deploying more and more complex applications on Internet-of-Things (IoT) devices. As a result, the complexity of executing computational tasks and their associated data transfers increases. However, the energy efficiency of computing systems and wireless networks has not improved much over the last decade. As a result, there is a big gap between the energy needed for running various applications versus the energy available to the IoT devices. Therefore, we need to investigate novel solutions to design and implement power efficient IoT systems.

The current state of the art is summarized in the following paragraphs. On the wireless radio side, there are several low-power radios developed by industry and academia. Examples include ZigBee, WiFi HaLow, NB-LTE, SigFox, LoRa, and backscatter communication techniques. A common characteristic shared by these radios is that they are optimized for power. In other words, they can operate with microwatts of power. However, such low power consumption does not come for free. Some radios sacrifice communication distance and others data rate. It is currently very difficult to design a wireless radio that consumes low-power, supports high data rate, communicates at long distances, and scales well.

On the IoT system side, we identify several limitations of existing solutions. The first limitation is that many schemes contain redundant hardware or software components. Such redundancy prevents the power consumption reduction of IoT systems. To minimize power consumption, we can eliminate redundant components and processing in an IoT computing system.

The second limitation we identified is that the process of connecting and configuring low-power IoT devices is cumbersome. Users need a significant technical background in order to set up an individual IoT device. This calls for a system that allows IoT devices to become plug-and-play. Such a system could act similar to an operating system, but designed for managing IoT devices instead of server resources, while conserving energy.

The last limitation identified is that most of existing IoT devices operate independently and in isolation from one another. They do not communicate and collaborate with other devices. Individual decisions might cause inefficient energy usage when there are multiple IoT devices deployed. Therefore, we would like to have a centralized controller that can connect many IoT devices and control them for actuation.

Low Hanging Fruit

Several approaches were identified for short term research focus. The first technique centers around Software-defined Networking (SDN). SDN is a relatively new paradigm to separate the control plane from the data plane on network devices. As industrial vendors begin to adopt SDN technology on routers and switches, it has become easier for an administrator to program and configure her wired network. Inspired by software-defined networking, we believe that many benefits can arise from centralizing the control plane for low power wireless devices. A centralized scheme can assign tasks to appropriate devices based on additional internal and external information. For example, a smart, battery-operated doorbell that detects a person approaching the door may be able to better eliminate false positives if additional sensors can be utilized. In this case, the doorbell can query a centralized controller before uploading a video, and in the event of a false positive, this allows for more power-efficient performance. A centralized scheme also more easily enables a plug-and-play architecture by motivating a technical standard that allows devices to inter-operate more seamlessly. Today, many deployed energy-constrained smart home devices have proprietary standards and interfaces. A centralized, software-defined scheme will motivate vendors to adhere to a common standard.

Outside home environments, researchers can examine how to provide low power efficiency on a city-wide scale. Networks of the future should provide city-scale sensing, such as facial recognition, pollution and traffic monitoring and safety analysis. In general, city-scale sensing will require communication ranges much larger than those needed in the home. Therefore, a central issue city-scale networks must face is how to increase transmission ranges of the sensors in a low-powered fashion. There is much to be gained by building on ideas from the sensor network community. Rather than understanding how to modify the low-power devices, researchers can examine how to modify the infrastructure to support the low-powered devices. Researchers can investigate the usage of mobile infrastructure. For example, entities in an airborne platform could act as a low-power base station that serves nodes in the field. The base station can be used to not only limit transmission ranges, but also provide power to devices. The base station can provide power through wireless charging techniques or energy harvesting can be utilized on the sensing devices with the base station providing mechanisms to make harvesting efficient.

Last, all areas of the protocol stack and software stack can be analyzed for power efficiency. It is likely that a careful instrumentation and measurement of the energy consumption of sensors and similar devices will indicate where unnecessary energy usage is occurring. Simplifying network protocols, such as TCP, can enable more efficient wireless communication and minimize network-based processing. By analyzing the needs of a specific application, network designers may be able to eliminate or simplify network functionality from the physical layer to the application layer. For example, application-specific objectives can be analyzed under different rates and service levels, and designers can ensure that the wireless technology serving a specific application uses the lowest-energy network settings that do not affect application performance. Additionally, administrators may benefit from schemes that allow them to sacrificate application performance in a deterministic way if certain power gains can be realized. Therefore, techniques to accurately model application performance and corresponding device energy usage will be useful. The sensor networks research community has already developed a foundation for such approaches and one can expect that near-term research will successfully build on these ideas to deliver solutions for city-scale networks of more heterogeneous sensor and actuator devices.

Big Challenges

Minimizing Energy Per Bit and Peak Power
Generally, the key challenge is to achieve order of magnitude gains in network energy efficiency, as measured in energy per bit but also being able to adapt to a peak power constraint. While communication at higher bitrates can often reduce the energy consumption per bit these gains do not necessarily imply a reduction in peak power consumption, which is often an important system design constraint.

Design for Energy Asymmetry
It is increasingly common for wireless networks to consist of energy constrained mobile devices that communicate to wired infrastructure that has no energy constraints. Backscatter communication is a primitive that dramatically improves the energy efficiency of uplink communication in asymmetric networks. But additional research is needed to fully explore the design space of asymmetric networks. Example topics: (1) high performance, energy-efficient passive downlink [fully passive downlink solutions exist, but their performance in terms of data rate and range is limited; there is a need for better fully passive schemes, as well as for partially passive, high-performance & low-energy schemes] (2) hybrid active-passive network design: use passive uplink and downlink primitives in conjunction with active radio primitives; for example, passive radio can be used as flow control (e.g. wake up or ack) for active networking links to improve energy efficient performance (3) improving network energy efficient performance by joint delivery of energy and information; for example, power optimized waveforms can be thought of as simple codes that improve power delivery; design of downlink codes to deliver both energy and information, as well as provide for backscatter uplink, is an open problem.

The discussion group also developed three proposals for a grand challenge competition for *Extremely Low Power Wireless.* These proposals are summarized as follows:

A. GRAND CHALLENGE 1
The competition asks for participants to build *low-power primitives* for a high density wireless communication scenario. The testing environment will include 10,000 - 100,000 devices in a

specified square footage of indoor space (i.e. warehouse). The participants must demonstrate their solution and can innovate across software, hardware or both. The solution can include techniques, protocols, and mechanisms in any layer or can span across multiple layers.

The effectiveness of the solutions will be evaluated across static and dynamic (controlled and random movements) use-cases. The solution will be evaluated based on the success of the demonstration in achieving the following goals:

1. Centralized Device discovery -- A wall-powered device must be able to discover the presence and/or location of every deployed device.

2. Distributed Device Discovery -- Each deployed device must be able to discover every other deployed device.

3. Information dissemination/collection -- each deployed device must receive the same piece of data from the wall-powered device; The wall powered device must collect a fixed amount of data from each deployed device.

4. Reconfigurability -- the deployed devices must be able to reconfigure and/or reset to a specific state either independently or from specifications from a fixed device.

The specific evaluation metrics will be:

1. Latency: time taken to achieve the goal with the same amount of energy on each node
2. Energy efficiency: Joules/bit and power constraints per node

Current state-of-the-art for each of the metrics will be defined, and the challenge will iteratively require teams to improve upon these metrics by 2x, 10x, 100x, and so on.

Teams will be given 2 days to deploy and test their solutions in the test environment. The evaluation will be conducted over the following two days.

B. GRAND CHALLENGE 2

This grand challenge competition asks for participants to demonstrate an extremely low-power, *application-specific,* high-density deployment. The participants will be required to deploy low-power nodes with high density in a realistic deployment environment such as a hospital, traffic safety testbed, etc. The participants are free to choose any off-the-shelf hardware or build their own hardware for the nodes. This challenge expects the participants to consider one application in a dense environment and the challenge also clearly characterizes the application-centric metrics, the enabling software and hardware designs based on those application metrics, and shows how disparate technologies (for example, different wireless technologies or different hardware or software layers) can be combined in a low-power fashion.

Overall the effectiveness of the solution will be evaluated over the following metrics:

a. What is the maximum number of nodes and bit rate for a given fixed spectrum, coverage area and deployment scenario?
b. What is the maximum bits per joule but at the same time ensuring reliability and privacy? The calculation should include the computation, communication for non-wall powered devices.
c. How to maximize the total data rate given a specific deployment?
d. How to minimize the power consumption to meet the specific application-dependent data rate?

C. GRAND CHALLENGE 3

This grand challenge, titled “Watt Can You Do”, asks for participants to reduce energy consumption to enable the network ecosystems of the future.

Homes of the future will be outfitted with many connected smart devices to sense and then modify our environments in real-time. While many of today’s smart devices are hardwired for electricity, devices of the future cannot be constrained with wires to reach the full potential of their deployment. In order to facilitate the adoption of future energy-efficient smart devices, this challenge asks to develop solutions to power a smart home with X Watts of total power budget (with an eye on an 1 Watt goal).

To ensure the solution is not too application-specific, each year, a series of objectives or themes will be set for the challenge. For example, high-level themes of “Sensing” and “Virtual Reality” can be defined for a competition in which teams aim to provide an infrastructure that can solve arbitrary problems within those given areas. Teams are encouraged to deploy their own hardware and software solutions to satisfy challenge-defined, application-level objectives.

With the proposed sample themes, many tests will be defined at the competition, and the success of tests will serve as the evaluation metric for this challenge. In particular, the tests include:

a. Sensing: How many people are in a given room? What are the people doing in the room?
b. Virtual Reality: Allow for 5 simultaneous HD streams to 5 different users in the room.
c. A mix of both sensing and virtual reality: Equip each user with a virtual reality headset and require them to stream a HD dance video and then mimic the dance. Find the best dancer.

Teams will have 2 days to deploy, configure and test their ecosystems at a centralized location before the competition. Benchmarks should accurately track the amount of useful information transfer in an application-centric manner. For example, given the same energy budget, a system that determines a user is typing on a keyboard should score lower than a system that determines what the user is actually typing. A side competition can award additional points based on other metrics such as: number of devices deployed, lifetime of network, densities supported, and lowest overall power consumed. Overall winners are determined by normalizing the benchmark scores to power consumed.