Looking Beyond the Internet

Steering Group Final Report - March 28, 2016

This report documents the observations and recommendations of the "Looking Beyond the Internet" Steering Group; please see Appendix A for a member list. The Steering Group effort, including its three workshops, is based upon work supported by the National Science Foundation under Grant No. 1546769[1]. The report is structured as:

1. Introduction & Recommendations
2. Background for this Report
3. Highlights of the "Looking Beyond" Workshops

1. Introduction & Recommendations

The world is now in the early stages of simultaneous revolutionary deep changes in the foundations of computer and communications technology. These changes range from innovations in large-scale storage, software defined networks and infrastructure, to next-generation wireless access technologies and the emergence of ubiquitous sensors, on our bodies, and in our homes and automobiles, which are rapidly becoming sensor-packed "clouds on wheels."

Two clear trends emerge: 1) most forms of infrastructure going forward will be software-defined, and thus questions of how to abstract, architect, and program such systems, whose scale eclipses today's infrastructure by orders of magnitude, will be fundamentally important research topics; 2) there will be tremendous ferment and diversity at the edges of the network with increasing high-bandwidth sensors and actuators, computing power and local storage, and flexibility, with a consequent huge increase in "data torrents." These trends will turn today's global Internet upside down, as they will enable revolutionary personalized low-latency, real-time services at unprecedented scales.

The "stars are aligned" for research with extremely high benefit for our society and our economy. There are tremendous opportunities for research and innovation in three deeply interwoven themes:

- pervasive software defined infrastructure,
- highly diverse, heterogeneous "edge clouds" ranging from billions of untethered devices (small sensors, automobiles, drones) through premises, campuses and cities, and

[1] Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

- wireless city research with the simultaneous availability of spectrum, active city participation, and highly synergistic industry work

Realizing seamless programmability across tomorrow's city-scale, multi-domain systems is a major research challenge. Heterogeneous hardware, owned and operated by many organizations, will need to support abstractions necessary to support large-scale software systems such as a "programmable city." Key issues will include architectures and abstractions for large-scale software systems, verifiability and robustness, managing and synthesizing the "data deluge," and the ever-growing roles of data analytics and machine learning as we create and operate these new systems.

The challenges and opportunities for tomorrow's national-scale, end-to-end network that will connect geographically distributed edges and support new generations of national-scale applications are no less exciting. Today's gigabit networks will be superseded by terabit and petabit networks, all components of which will be software defined. Virtualization, network slicing, software defined infrastructure (SDI) and software defined exchanges (SDX) will enable the dynamic allocation of network and compute resources leading to potentially dramatic changes in today's Internet paradigms.

Many of the most important success stories in the United States economy during the past generation have arisen from forward-looking research in computer science and engineering, ranging from the Internet to Google, which were sustained initially by DARPA and later by NSF research investments from the 1960s through the 1990s. The past 15 years' experience with medium and large-scale systems projects – projects larger than a single Principal Investigator – have continued to show their importance. A substantial portion of today's systems research almost by definition needs collaborations and interactions across groups. Significant benefits are likely to come from large-scale projects that attempt to think big, focus the community, engage multiple strong systems groups, and deploy and try out research at scale.

Now is the time for the Computer Science systems research community to draw up its next-generation research agenda, and as important, to set in place plans to grow and strengthen the nation's systems research capabilities. CISE should look for ways to continue to help the community grow stronger, e.g., by enhancing opportunities to build and try out large-scale systems, and share experiences, methods, practices, and curated ecosystems of tools. Sharing of data, experiments, insights and infrastructure should be first class goals. Research should extend beyond the specific technical disciplines to include advances in managing and reasoning about the socio-technical aspects of systems (security, privacy, ownership, economics).

- **Overarching recommendation:** CISE should launch an intensive, long-term new program on "Beyond the Internet," sponsoring teams/consortia to carry out projects that combine research with the creation and operation of broadly-shared, interoperable infrastructure, to focus the community on key technical challenges and maintain US leadership in the Internet field.

This program should engage a research agenda with an emphasis on understanding how to deploy and try out novel applications within large-scale suites of software defined infrastructure designed to provide seamless services across clouds, core networks and wireless edge networks. The experimental network deployment should build in key features such as security, resilience, and managing safe cross-domain sharing of data, resources, and services.

A key part of this agenda will include creating and running a heterogeneous but sliced, deeply programmable, federated and interoperable suite of software defined infrastructure that spans the United States for both city-scale and global scale experiments. This infrastructure will need a national core plus multiple edge, campus and community infrastructures that will support cross-domain technical and application research experiments related to the broad theme of "Beyond the Internet." Architecting, building, and using this system will provide a tremendous opportunity for the community to work together in building and running large-scale experimental systems.

Research observations and recommendations

The global Internet will be almost unrecognizable ten years from now – driven by a deluge of low-latency sensing and actuation data, increasingly diverse "edge cloudlets" (comprising Internet of Things, Cyber-Physical Systems, etc.), within a planetary-scale network largely built from Software Defined Infrastructure (SDI) and a decentralized, diverse cloud. This report report recommends research programs that will build upon the opportunities sparked by these expected paradigm shifting changes.

At the edge we envision billions of clients ranging from tiny embedded sensors, to wearable devices, all the way up to high-end systems such as citywide or regional electrical and water systems. These clients will run a wide range of applications that will be personalized using large-scale data analytics. Based on their communication and computational needs, clients may choose from cloud services offered on a variety of platforms, such as traditional centralized clouds, smaller geographically distributed clouds, or cloudlets (or "fog") at the edge. The system must optimize these low latency, high availability and high bandwidth needs for critical applications such as personal health management, augmented reality, autonomous driving, and drone management.

These devices will communicate through a network based on software defined infrastructure with significant embedded computing and storage, building on, and extending, current technologies such as software defined networking, network-function virtualization, and software defined radios. Software defined infrastructure "cracks open" the calcified network core/stack, opening it to innovation, enhanced agility (perhaps incorporating machine learning), and an expanded diversity of networking approaches. Since DevOps cycle times are measured in days rather than years, we foresee an enormous acceleration in innovation as large-scale systems become software defined, both in the multi-domain core network and at the edge.

The software defined infrastructure model changes the nature of programming in a fundamental way, emphasizing distributed services for real-time networked control. In particular, the factoring of functions between programmable devices and control servers raises new challenges to building networks that are agile but also reliable, scalable, and secure. New frameworks and abstractions are needed to realize this goal. Evaluating and refining software for software defined infrastructure requires new research infrastructure to enable evaluation on realistic networks and with realistic usage conditions, involving a diversity of traffic patterns.

Edge networks will be increasingly wireless, putting pressure on spectrum availability. They will be highly diverse, to meet a wide range of constraints in bandwidth, range, dynamics, mobility, and power. Basic research is needed to develop new paradigms for efficient spectrum sharing suitable for devices with very diverse characteristics in the presence of the concurrent use of diverse wireless protocols.

In the new generation of application services that will leverage these wireless edge networks, there will be a shift toward systems that localize processing (where possible) in edge clouds under local control, e.g., for reasons of efficiency, privacy, or latency. Such new infrastructure systems open the door to a range of exciting new applications research opportunities, such as personal health management, real-time cognitive aids, augmented reality, and coordinated control of robotic systems. At the same time, some services will demand dynamic elastic provisioning and/or interaction beyond the edge to aggregate data from multiple locations. New models must emerge for multi-cloud software systems that manage the locality and trust issues inherent in multi-domain distributed cloud infrastructures.

Core network hardware infrastructure will need to accommodate new capabilities such as dynamic resource provisioning and higher performance while also enabling software defined infrastructure programmability for customization and innovation. Research is needed to understand how to architect and build the new, global scale, multi-domain, software defined, end-to-end network that will connect its edges flexibly and securely, bring computation to bear on data, incorporate new networking and cloud paradigms, and create new planetary-scale applications. Developing, verifying, deploying and testing new software versions will be central to the robustness, security and performance of this "Beyond the Internet." Data analytics and machine learning may become a dominant means for creating and driving such systems, enabling continuous adaptation to personalize applications and to optimize and manage the infrastructure. Security and privacy concerns will be acute and pervasive.

- **Recommendation R-1.** CISE should initiate a new program which provides substantial funding to teams or consortia that address key "Beyond the Internet" research challenges while simultaneously building and running the suites of open, broadly-shared experimental infrastructure necessary to address these challenges.

  These research challenges include but are not limited to: real-time edge cloud services, software defined infrastructure, software defined exchanges, multi-domain federation architectures, scaling (e.g., geographic or numbers of devices in diverse wireless edges), security, privacy, resilience, mobility, new network paradigms (e.g., protocols, architecture), in-network computing, and interfacing with the physical world. Each team should both perform research and build and operate open, broadly-shared infrastructure. To achieve the goals of this program, annual funding per team or collaboration of $7M for each of five years is recommended.

- **Recommendation R-2.** As this new world is raising critically important new economic, social, policy, security, and privacy challenges, a multi-disciplinary research approach is needed and is just as urgent as the technical components. Such challenges can be best addressed with in-situ experimentation involving significant new applications with real-world end users.

- **Recommendation R-3.** CISE should initiate a research program that uses applications in a living lab context to drive wireless cities research.

  This approach can take advantage of today's remarkable current opportunity in the Wireless Cities space, and sponsor significant research experimentation with next-generation wireless and cloudlet technologies that span diverse campuses and communities and cities, thereby providing an opportunity for their residents to "live in the future" five to ten years before such technologies become pervasive. This program should engage campus IT organizations and, where feasible, be undertaken in partnership with industry in order to leverage emerging new technologies.

Infrastructure observations and recommendations

Facilitating experimentation at scale, across multiple domains, is critical for the next wave in research outcomes. Central to the recommendations of this report, is the inclusion of significant activities around the development, operation and ongoing extensions of a large-scale suite of interoperable, reusable research infrastructures for experimentation. While a nationwide fabric of software defined infrastructure and software defined exchanges and will be needed for such experimentation, a substantial fraction of this infrastructure should be located in "edges," more specifically academic campuses, cities, and communities, as we foresee tremendous ferment and innovation in the edges.

Research infrastructure is an ecosystem comprising a variety of elements from experimental platforms to experiment management tools and artifacts. Shared infrastructure with useful, reliable tools promoting the development of good experimental methodology in computer science will play an important role in supporting the proposed "Beyond the Internet" research initiatives.

The majority of research infrastructure work to date has been in support of "realizing" experiments in various forms of infrastructure. Far less effort has been expended on experimental workflows and tools supporting automation, repeatability, reusability and fault diagnostics. The potential of large-scale, distributed suites of research infrastructure lies largely unlocked today due to the lack of these essential experimental mechanisms. Perhaps even worse, systems and networking progress has long been hindered by the lack of repeatability and reusability of distributed experiments.

While the word "infrastructure" has traditionally implied "hardware assets," the Planning Group strongly believes that over the next 10 years key innovations will arise from *software* running in easily-programmable hardware. This trend has extremely important implications for infrastructure funding models - in essence, the infrastructure must be designed to permit *rapid and fluid innovation in software*, along with excellent debugging tools and easy reuse of existing software artifacts. As a very important consequence, NSF funding models should also evolve to explicitly fund the support environments required for software innovations (such as tool development, archives, etc.) as well as hardware.

Research infrastructure is closely related to the research that will use it. The research topics being investigated (core network, software defined infrastructure, clouds, edges, wireless) are in fact the very components of advanced research infrastructure for experimentation. Just as the SDI trend enables increasingly more diverse forms of slicing and programmatically customized network computing environments, which may be interconnected on an on-demand basis, there is a growing shift from fixed, standalone, and often specialized testbeds to a new era of highly dynamic, composable, distributed, experimental research infrastructure with a strong emphasis on edge systems and new forms of multi-domain interconnect.

Creating such research infrastructure is a first class research activity in its own right, as there are many open research questions in developing this sort of experimental capabilities in support of communities of research interests. The infrastructures and their software support tools should be built by researchers as part of their exploration process, and as a community working together.

- **Recommendation I-1.** CISE should sponsor multiple, large-scale projects for the architecture, design, development, operation, and continual advancement of interoperable suites of experimental research infrastructure in support of the research programs described in R1 - R3.

  The federated suite of research infrastructure built by these teams should be freely available for experimentation by all US computer science researchers. It should provide services for rapid and fluid innovation in software, along with processes, tools and easy reuse of existing software artifacts. Programs should be developed to create both larger experimental research centers and satellite teams that are highly exploratory and or deeply specialized. The infrastructure itself will include components of the research topics being investigated, with an emphasis on network core, edge technology, clouds and wireless. The infrastructure may be driven by real-world application challenges, and should be designed to capture important characteristics such as real end users, real usage patterns (e.g., time of access, mobility, application traffic), and physical infrastructure (e.g., radio cells, vehicle platforms, infrastructure-assisted wireless).

- **Recommendation I-2.** CISE should view such infrastructure as an ecosystem consisting of shared focus, tools, methodology, and community.

  As an ecosystem, these will also provide long-term support for community-building activities, the creation and use of shared tools and archives, conferences and demonstrations, and other mechanisms to grow and strengthen the research communities building and using this infrastructure.

- **Recommendation I-3.** CISE should support a near term planning process for creating and operating suites of research infrastructure to support research into Wireless Cities.

  Considerable planning will be needed for such efforts, and this planning should begin as soon as possible. These new suites of infrastructure should be interoperable where feasible with existing research infrastructure suites, and support long-term application-driven experiments.

Process observations and recommendations

We believe that it is imperative to grow and strengthen the nation's current capabilities in academic systems research. At present, only a small number of academic groups have the ability to create major experimental systems. For historical comparison, systems such as MULTICS, created by MIT, GE and Bell Telephone Laboratories, 1964-1970, BSD Unix, created by the University of California, Berkeley, from 1977-1995, and the Mach kernel, created by Carnegie Mellon University from 1985-1994, pioneered exciting new research concepts and directions and, in addition, spawned an explosion of experimental systems research and education at universities beyond these institutions. The ideas embodied in each also played an influential role in the design of commercial systems offerings.

"Beyond the Internet" systems will be orders of magnitude more complex, encompassing national, regional and institution scale networks; distributed edge and core clouds; billions of connected devices and sensors; technologies ranging from wireless/mobile to optical; and paradigm shifts to software defined infrastructure, extensive use machine learning and peta-scale data set production, movement and analysis. U.S. leadership in this "Beyond the Internet" world is critical to our future national economic well being. To get there will require significant investments to create a critical mass of multi-institution centers of excellence in education, research and experimental infrastructure related to the future "Beyond the Internet."

To this end, CISE should encourage academic systems research teams to create and operate suites of shared, interoperable research infrastructure, with an ultimate goal of providing sustained and significant funding for a modest number of highly-capable teams (perhaps 10). This will probably initially require some community workforce capability-building. Hence CISE might each year fund a few (2-3) highly-capable teams at the full level proposed in the recommendations, plus a portfolio of teams with short term funding at a lower planning level. Those receiving planning level grants will be primed to compete in subsequent solicitations for highly capable teams. Staging the awarding of the major awards over several years will provide an opportunity for continuing evaluation and tuning.

Building and operating such research infrastructure requires professional research staff. They provide greater expertise, long-term consistency, and are less subject to the academic calendar than students. Equally important will be sufficient long term funding for researchers and graduate students. With staff as mentors, students can meaningfully contribute to research infrastructure projects. In many cases, students are running experiments for their research projects, and are thus well-placed to identify infrastructure needs, and to do early prototyping of new infrastructure and tools. This will encourage collaboration between researchers, software tool developers and system (software and hardware) maintenance staff enabling potentially significant synergies.

Last but not least, an important component of US leadership in networking over the past 45 years has been the multi-faceted, synergistic collaborations between academic researchers and the computer / communications industry. These interactions have ranged from scientific collaborations between university and industry researchers, to industry providing pre-product hardware/software for evaluation and use by academic researchers, to discussions of topics of mutual interest, to working with industry personnel to specify experimental research and education infrastructure, to industry participation on project advisory committees, to industry grants, contracts, donations and discounts in support of university research and infrastructure, and to a continuous two-way personnel flow between university and industry. Continuation of this industry/university relationship will be important to the future success of the "Beyond the Internet" research/infrastructure agenda.

- **Recommendation P-1.** CISE should in its funding for "Beyond the Internet" cultivate, nurture, and sustain a modest number (5-10) of American integrated academic multi-institutional systems research teams to create large, high quality, open, innovative, interoperable suites of infrastructure for investigating and learning about tomorrow's Internet architectures and to engage in research on new network paradigms that is supported by this infrastructure. In contrast to earlier infrastructure projects funded by NSF CISE, we believe it is important to fully integrate ambitious research goals with infrastructure development and operations.

- **Recommendation P-2.** To be effective, each team needs a "critical mass" that combines sufficient funding (~7M$ per team per year for five years) for hardware, researchers, software developers, and operations staff, typically in close proximity. CISE should support a balanced approach between infrastructure and research that provides adequate team funding for experimental research, software development and professional staff for deployment and operations, with lengthened award periods to provide more stable funding for staff. Our sense is that a 50-25-25 split of funding for research, hardware and for software development/operations would be appropriate.

- **Recommendation P-3.** In its solicitations for these programs, CISE should encourage collaboration/cooperation between university submitters and industry partners with the precise forms of collaboration/cooperation to be determined by the scientific needs of proposed projects.

2. Background for this report

This report forms one link in a chain of related NSF reports on the roles of testbeds and infrastructure within Computer Science research, dating back at least to the 2002 *Report of the NSF Workshop on Network Research Testbeds* [NRT 2002]. The Steering Group has been particularly struck by the direct relevance of observations within the most recent report in this series, *The Report of the NSF CISE AC Midscale Infrastructure Committee* [MIC 2014], and our report can be read as a companion piece and sequel to this earlier report.

We agree with the conclusions and recommendations drawn in the MIC 2014 report. Two years later, with two highly successful NSF Future Cloud projects underway and already experiencing very heavy experimental research use, we strongly agree with that report's focus on "cloud / network / grid systems as a first priority." It has proved a well-chosen first step. We also strongly concur with the report's other key conclusions, including the important role of experimental applications, the pathways to practice, the major benefits to human capital, and the right kinds of approaches to building and sustaining research infrastructure. The past two years have provided convincing evidence to back up these observations.

As the MIC report notes, a number of notable research testbeds have greatly advanced experimental computer science: well-known examples include the Gigabit Testbeds, PlanetLab, Emulab, DeterLab, ORBIT, PROBE, and WAIL, and more recently GENI, US Ignite, Chameleon, and CloudLab. It observed that rapidly growing interest in virtualization and converging interests in the cloud were likely to drive mid-term interests, which has proved true, and highlighted the following vision for experimental mid-scale infrastructure for Computer Science:

> "A nationwide, multi-tiered system (national / regional R&E backbones, data centers, campuses) that is sliced, deeply programmable, virtualized, and federated so that research experiments can run 'end to end' across the full suite of infrastructure."

To a first approximation, the GENI project provided an early step towards this vision, with the NSF Future Cloud data centers adding a sliced, deeply programmable "cloud" tier within the end-to-end suite. This report can be viewed as recommending that a distributed "edge-cloud" tier of campuses and communities be federated as the next tier.

What is new?

To the extent this this report serves as a companion piece to the earlier MIC report, we have now gained sufficient experience to understand better the back and forth between research, experimental infrastructure, and education and curriculum development, and can incorporate the "lessons learned" from real use. GENI now has over 6,500 users, mostly for research but also now for dozens of classes, and CloudLab and Chameleon are ramping up even more rapidly in their use -- in fact, both are close to "full" at the time of writing. (Of course, PlanetLab, Emulab, DeterLab, ORBIT, etc., have been "full" for years!) Thus we have increasing evidence that the MIC report was correct in its observations.

However, it is equally important to note that technology is rapidly evolving – indeed, we believe that the Internet itself is now beginning to undergo a profound transformation. This ongoing, rapid evolution brings major consequences for both research need and infrastructure needs in experimental computer science. New factors include:

- The success of software defined networking (SDN)
- The dawn of "software defined infrastructure," (SDI) and rising dominance of software running within cheap, commodity "white box" hardware
- The inevitable merging of cloud and network technologies
- Exponential growth in mobile data and the emergence of "5G" cellular systems
- Network Function Virtualization (NFV), with service chaining now starting to replace dedicated strings of hardware devices
- Rapid introduction of inexpensive, large-scale storage technologies
- The rise of the Internet of Things (IOT) and Cyber-Physical Systems (CPS)
- A huge "data explosion" from connected sensors
- And the prospect of entirely new "data torrents" from automobiles, drones, etc.

Two clear patterns emerge: 1) software systems will play the central role in most forms of infrastructure going forward, and thus questions of how to abstract, architect, and program such systems will be extremely important research topics; 2) there will be tremendous ferment and diversity at the "edges" of the network, with a consequent huge increase in "data torrents."

The Steering Group, workshops, and this report

The Steering Group for this report was organized in autumn 2015. Steering Group members were selected to form a distinguished, accomplished, and diverse panel of researchers, creators and operators of research infrastructure, and university Chief Information Officers (CIOs). Appendix A provides a list of the Steering Group members.

The Steering Group held a series of four meetings as follows: October 2015 (phone), December 2015 (at NSF), February 2015 (phone), and March 2015 (at NSF).

During the December meeting, CISE AD Jim Kurose gave the group his perspectives on the subject, and pointed out areas in which our observations would be particularly helpful. In that meeting, the Steering

Group also invited Nick McKeown (Stanford), Scott Shenker (Berkeley), and Henning Schulzrinne (Columbia) to provide their perspectives on what might or might not be useful in research infrastructure. McKeown and Shenker presented their views jointly; while the Steering Group did not agree with every point they made, it was extremely helpful to receive their inputs and recommendations, and the Steering Group did find itself in full agreement with the speakers as to the central role of software in innovation, and the importance of funding software development and maintenance efforts in CISE infrastructure projects. Schulzrinne participated for the full day, and provided unique and valuable insights into “hot” research areas going forward and the interactions of wireless research with federal policy makers.

As a major part of the Steering Group effort, Steering Group members co-chaired three workshops in early 2016 as follows, which engaged a total of roughly 160+ researchers, infrastructure operators, industry personnel, and CIOs in discussions:

- Applications and Services in the Year 2021 (Calyam, Ricart)
- Future Wireless Cities (Banerjee, Raychaudhuri)
- Software Defined Infrastructure / Software Defined Exchanges (Nick Feamster, Ricci)

These workshops took place in Washington, DC, during late January and early February 2016. The workshop reports form detailed companion pieces to this report, and provide much deeper insights into the specific research areas that focused each workshop. Most members of the Steering Group participated in at least one workshop, and some participated in all three. Thus the workshop outputs, and the workshop participation itself, provided important inputs to the Steering Group’s deliberations and drove many of the group’s key insights. Highlights of each workshop are included in the next section.

3. Highlights of the "Looking Beyond" workshops

This section provides highlights of the three workshops organized by the Steering Group. Taken as an ensemble, these workshops engaged 160+ computer science researchers, infrastructure operators, and participants from industry, cities, and the US government.

These workshops provided inputs and insights to the Steering Group. Each workshop is documented in its own workshop report, which provides background and further amplification to the major points summarized here. Please see Appendix A for details.

3a. Applications and Services in the Year 2021

The coming data deluge will change everything. An explosion of data, especially from massive streams of local Internet of Things data, will upset architectures and demand new approaches. In addition to the small IoT devices such as thermostats, we will soon see automobiles and drones become very high capacity sensors contributing tremendous volumes of 3D real-time measurement (e.g. lidar). Data will be at the center of much that drives future research and development, and the economy. It will be the glue that unites the physical and cyber worlds, the key to understanding human environments, the intelligence behind personal cyber coaches, the comparative weight that allows us to balance privacy against security in each case, and the grounding for artificially intelligent agents acting on our behalf.

An interesting part of the applications space is its growing need for "smart services" for continuously operating ecosystems of data, interactive analysis, artificial intelligence interpretation, and human-guided visualization and intervention.

The privacy and security implications of this new world will be profound and beyond today's comprehension. Privacy and security is very ill-understood in this new world -- technically, socio-technically, and also socio-economically. At-scale research experimentation is needed across campuses and communities, with real users with diverse preferences in order to create a privacy- protected repository of preference data that can be used/re-used by researchers. Academic research on data privacy should recruit subjects with life experience or broader perspective to reflect the diversity and complexity of the population at large.

We foresee tremendous innovation at the "edge," and campuses and communities are perfect places for edge experimentation. These suites of research infrastructure should be designed to support broad categories of applications, not simply with particular "killer apps" in mind, as such predictions of the future tend to be incomplete. With broadly available infrastructure that can be operated as software-defined infrastructure, new solutions for national priority application areas can be realized.

These suites of infrastructure (campuses, communities) should be open, federated, deeply programmable, sliceable, and sharable for experimentation, similar to current GENI, CloudLab, Chameleon and beyond. Locavore (cloudlet, fog) infrastructure is especially needed at the edge for applications with low- latency requirements.

We can gain huge benefits from "living lab" research throughout campuses, cities, and communities. Cities, enterprises and university communities can be leveraged as interesting living labs to understand and tackle the application and service challenges in a IoT world of the year 2021. Students and citizens should be encouraged to "live in the future" as part of these living labs, and push the limits of what we can imagine in terms of futuristic possibilities, and unexpected new frontiers for innovation.

Such living lab infrastructure can foster interdisciplinary research, marrying the IT fields of computer science and engineering with the human-oriented fields of sociology and psychology and economics. Further, these campus and community experimental infrastructures, when jointly shared by industry and academic researcher teams, could lead to new innovations that can be rapidly translated to benefit enterprises and even citizen applications. We note that applications in the national priority areas (healthcare, education, public safety, citizen innovation, etc.) are good drivers for interdisciplinary research, which is essential for addressing Global Challenges.

3b. Software Defined Infrastructure / Software Defined Exchanges

We are at the dawn of a new era: Software Defined Infrastructure (SDI). Today's relatively static cyber-infrastructures, implemented via hardware with predetermined control systems, are now beginning to morph into fluid, planetary-scale software systems – highly interconnected, deeply programmable, and virtualized within end-to-end slices across many administrative domains. SDI's forerunners include multi-tenant clouds, software defined networking, network functions virtualization, and software defined radios. Individually, each presents major research challenges. But viewed within the broader SDI context, they are simply starting points of a very deep revolution that will reshape our global computing infrastructure.

Today's Internet is already beginning to undergo a deep transition from multi-tenant clouds and relatively static infrastructure to a rapidly flexible, deeply programmable new infrastructure based on Software Defined Networks (SDN), Network Function Virtualization (NFV), Software Defined Radios (SDRs), and forward-looking 5G cellular system concepts such as Virtualized Radio Access Networks. We expect this trend to intensify and proliferate in the near future, leading to a rapidly thickening fabric of multi-domain, heterogeneous, edge clouds and interclouds at a scale far beyond today's Internet, and will accelerate as it incorporates mass-market Internet-of-Things devices and large-scale cyber-physical systems including cities, autonomous automobiles and ubiquitous drones providing massive "data torrents" from their mobile, high-bandwidth sensors.

The transition of today's R&E cyber-infrastructure to SDI has already begun. As we move from today's Software Defined Networking to the emerging vision of SDI, we will need to describe, program, trouble-shoot, and reason about systems that span many different kinds of component devices and subsystems. Chief among them are those that support computation, storage, connectivity, and many types of sensors and actuators. In addition, services will likely be incorporated into these over-arching systems. As is the case today, each of these components (resources) will be owned by someone. Therefore the basic issue is of creating end-to-end systems from components provided from multiple administrative domains.

Looking forward, we envision a world in which all aspects of the planet's cyber-infrastructure form an interconnected, multi-tenant (sliced), and deeply programmable planetary-scale ensemble composed of trillions of devices owned and operated by millions of partially-cooperating, partially-competing organizations. It is conceivable that today's Internet will run in just one "slice" across this infrastructure, with many other novel services populating other slices.

Many fundamental research challenges arise with this vision, and a vigorous community debate is now underway regarding the defining features and capabilities of this new SDI-based cyber-infrastructure. Given this very lively debate, it is intriguing that the workshop participants came to express a high level of agreement on three fundamental research challenges:

1) What are the right abstractions for representing, programming, troubleshooting, and reasoning about planetary-scale, sliced, multi-domain SDI systems that incorporate an enormous variety of devices and services?

2) How can we understand, reason about, troubleshoot, and control the dynamics of large-scale SDI systems, and how can we ensure the robustness and resilience of the services they host in the face of unexpected events?

3) How can we understand, reason about, and manage the socio-technical aspects of such systems, including the security, privacy, and data-ownership issues that arise in multi-domain systems that weave together many layers of software?

3c. Future Wireless Cities

Recent emphasis on experimental wireless research has led to critical breakthroughs. The last decade has seen the community move from mostly simulation and modeling based research to the creation of experimental systems in real-world deployments. In this transition, it has become abundantly clear that the effects of large scale, of complex interactions among large numbers of applications are critical to the design of wireless systems, but increasingly elusive in simulation. Experimental testbeds are needed to capture important city characteristics such as real end users, real usage patterns (e.g., time of access, mobility, application traffic), and physical infrastructure (e.g., radio cells, vehicle platforms, infrastructure-assisted wireless). Looking forward, such testbeds must span multiple domains and city geographies as opposed to single campuses.

Tomorrow's wireless edge is something completely new: "cloud systems with radios." They will consist of sliced, virtualized software programmable clouds with software-defined network infrastructure, radios that will be to a large extent software-defined, and of course real-world end users. These software-defined radios, along with the software-defined infrastructure "behind" the radios, create enormous opportunities for research and rapid innovation. The success of infrastructure at such scale requires the alignment with ongoing industry/open source 5G, SDN, NFV R&D efforts, exploration of new opportunities in spectrum, from existing commercial bands to 3.5G to others, and finally the need to study range of client devices and infrastructure.

Research in the next-generation "wireless edge" currently shows significant potential for breakthroughs of high economic importance. Specific areas of high potential impact include: innovative next-generation commercial wireless edge systems ("5G") both in terms of advanced functionality and the introduction of programmability/virtualization; emerging city-scale approaches to the Internet-of-Things (IoT) and Cyber Physical System (CPS); integration of computing and networking for mobile edge cloud scenarios; architectural research on future wireless systems addressing fundamental challenges of scale, latency, security, robustness and usability; and shared spectrum access in dense urban environments.

The "stars are aligned" for city-scale wireless research. Recent FCC rules for experimental spectrum and the availability of new frequency bands such as 3.5 Ghz and mmWave now make it possible to deploy experimental broadband wireless technologies outdoors on a realistic scale. We have strong evidence that a number of cities are eager to serve as host for future deployments, thanks to US Ignite. (Representatives from the following cities participated in the workshop and are eager to host such experimental infrastructure with their university partners: Ammon ID, Burlington VT, Chattanooga TN, Flint MI, Kansas City KS-MO, Philadelphia PA.) And equally important, today's industry focus on "5G" technology and services has significantly improved prospects for incorporating advanced technology into city scale experimental infrastructure. Emerging programmability and virtualization technologies for wireless such as NFV, SDN, open LTE and cloudRAN are expected to enable a great deal of experimental flexibility previously considered impossible.

Now is the time to act. NSF and the research community should work towards development of open, programmable, large-scale wireless city testbeds in collaboration with industry and ongoing city projects as well as other government agencies. Multiple testbeds in parallel may increase the probability of success. Sufficient critical mass investment will be required to ensure success of deployment. Infrastructure of such scale cannot likely be managed by individual academic entities alone. It requires academic institutions to work with local government (from local communities), and with industry that have suitable experience in managing such infrastructure. Creative solutions in collaboration with industry will be required for development of low-cost mobile and IoT devices with the right form factor and power consumption for widespread deployment. it is clear that a significant planning phase is necessary to move this process forward. The stars are aligned; now is the time to act.

Appendix A

The Steering Group effort, including its three workshops, was sponsored by the National Science Foundation (Award #1546769). We thank the NSF for this support. All opinions in this report are those of the Steering Group, and not of the National Science Foundation.

Steering Group members

Members represent a mix of researchers, creators and operators of experimental infrastructure, and university Chief Information Officers (CIOs). An asterisk (*) indicates the co-chair of a workshop.

Suman Banerjee*, University of Wisconsin - Madison
Terry Benzel, University of Southern California / Information Sciences Institute
Jim Bottum, Clemson University
Prasad Calyam*, University of Missouri
Jeff Chase, Duke University
Tracy Futhey, Duke University
Kate Keahey, University of Chicago / Argonne National Laboratory
Dipankar Raychaudhuri*, Rutgers WINLAB
Glenn Ricart*, US Ignite
Rob Ricci*, University of Utah
Peter Steenkiste, CMU

Steering Group organizers

Chip Elliott, BBN
Dave Farber, University of Pennsylvania
Larry Landweber, University of Wisconsin - Madison

Related Reports

[NRT 2002] *Report of the NSF Workshop on Network Research Testbeds*
[MIC 2014] *The Report of the NSF CISE AC Midscale Infrastructure Committee*

[ASW] *Applications and Services in the Year 2021* (Calyam, Ricart)

https://asw2016.wordpress.com

[FWC] *Future Wireless Cities* (Banerjee, Raychaudhuri)

http://www.winlab.rutgers.edu/events/wicities/

[SDI-SDX] *Software Defined Infrastructure / Software Defined Exchanges* (Nick Feamster, Ricci)

https://www.flux.utah.edu/beyond-internet-workshops/sdi

Appendix B

Observations and Commentary

This section provides observations and commentary made by members of the Steering Group, which in turn have driven our discussions and ultimately our recommendations. These are observations and commentary made by individual Steering Group members, and are not necessarily endorsed by all the other members. They do not represent Steering Group recommendations or consensus, but rather give a flavor for the group's discussions. Although the Steering Group freely intermingled observations on research opportunities, infrastructure concepts, and process issues in our discussions, for expository clarity we have organized this section as follows: research, infrastructure, process.

Research observations

We must prepare for profound and pervasive changes in Internet and Cloud architectures over the next five years due to three simultaneous but related explosions.

(1) <u>An explosion of communicating devices.</u> In five years, communicating devices will outnumber communicating people by perhaps by a factor of 100. It's been called the Internet of Things, the Industrial Internet, smart homes, smart cities, etc.
(2) <u>A data explosion</u>. We are beginning to collect, keep, analyze, and act on mountains of data, and the trend is still very young. New cognitive algorithms will be needed to integrate and make sense of all this data, but the impact also will require profound changes in computer systems, networks, and methods of federated systems.
(3) <u>An explosion in software-defined innovation</u>. Bespoke hardware devices are quickly giving way to software-defined devices. Software cycle times are fractions of hardware cycle times and will be a significant contributor to rapid cyber-innovation in the next five years. Software also permits virtualization or slicing, and slicing is expected to provide the first widely-deployed new tool for computer security and privacy in many years.

At the edge we envision client devices ranging from tiny embedded sensors, to wearable devices, all the way up to high-end systems such as sophisticated entertainment systems, autonomous automobiles, and fleets of drones. These devices will run a wide range of applications that will be customized using large-scale data analytics. Based on their communication and computational needs, client devices can choose from cloud services offered on a variety of platforms, such as traditional centralized clouds, smaller geographically distributed clouds, or cloudlets (or "fog") at the edge.

These devices will communicate through a software-defined infrastructure with significant embedded computing and storage, building on, and extending, technologies such as software-defined networking, network-function virtualization, and software defined radios. The edge network will be increasingly wireless, putting pressure on spectrum availability, and diverse, to meet a wide range of constraints with respect to bandwidth, range, dynamics and mobility, and power. In contrast, the core network infrastructure might be relatively homogeneous, relying on programmability for customization and innovation.

Important future research areas include appropriate technical, social, and economic architectures for these new systems and their interactions with extremely large numbers of devices. Specific research areas will include abstractions, infrastructure, and applications for secure distributed clouds and their data (locavore computing) and at least three edges (city, wireless, and personal). There will also be significant research challenges in managing end-to-end reliable application performance and economics when interconnecting distributed data and clouds across multiple administrative domains (e.g. via Software Defined Exchanges) while allowing for software-defined, end-to-end deep programmability.

All three of the explosions mentioned above have an impact on personal privacy. As so much information is being gathered, how can its owners manage the uses to which it is put? What does privacy now mean? Is there so much data it's impossible to keep secure? What policy issues should be addressed by government and by industry? Will there be open economic models or proprietary solutions?

All three explosions apply to the wireless domain. Devices and their data are expected to be predominantly wireless and there will be profound changes in the wireless world to accommodate such extensive growth. Software innovation will be critical, and next-generation wireless infrastructure will be dominated by a software-defined wireless edge approach.

Infrastructure observations

The overarching architecture for the suites of experimental infrastructure identified in this report will likely be distributed, largely software based, and with a particular emphasis on the edge. Research directions will probably dictate a collection of research infrastructures that can be specialized for domain specific explorations while also providing a heterogeneous set of loosely coupled and widely distributed resources. One approach to increasing realism and access is to leverage "edge-clouds" through a tier of federated campuses and communities. Procedurally, there are many benefits to be gained through university-city alliances for infrastructure and industry involvement. It should be noted that advanced research in cybersecurity and privacy may impose additional requirements on infrastructure architectures and use models.

Given such kinds of research infrastructure, many useful research "testbeds" will actually have no need to own or operate any hardware whatsoever. Given the rise of slicing and virtualization technologies, many (though certainly not all) testbeds can be layered on top of existing hardware platforms such as GENI, the NSF Future Cloud projects, other NSF investments, or commercial clouds. Such testbeds will be defined by *software toolsets* that enable specific communities of experimentation, rather than by dedicated collections of hardware.

One particularly interesting research area will be in architectures that provide generic experimental research frameworks that allow for specialized, domain-specific instantiations. This architectural principle may ultimately be more important than the actual connection fabric as it will move the discussion of research infrastructure from one of "realization" mechanisms to the higher-level design and representation requirements for experimental research.

In addition to the core capabilities discussed here, we see several meta-properties of importance. Research infrastructure needs to be easily usable by both a wide range of experimenters and by the owner/operators of the infrastructure. Then, in support of experiment validity, the community needs mechanisms and processes to provide confidentiality, availability and integrity of the experiment ecosystem. Finally, there are a number of cultural and social changes, along with community building, that are needed to facilitate future capabilities.

Experimentation with community-based (wireless) edge infrastructure has a strong potential to improve the lives of Americans living in those communities, and thus help show the way towards improving lives of everyone across the United States. Research into novel classes of applications and services, enabled by these future technologies, can both address national priorities, such as improved healthcare or transportation, and add value to fundamental and applied research in two very different ways:

(1) They can translate the fundamental advances of curiosity-driven research into valuable advantages and benefits for Americans and the communities in which they live.
(2) They can pose unsolved challenges that may provide grounding for fundamental and applied research and suggest new and untraveled approaches.

Both modalities will help to add innovation and growth to the national economy and improve the quality of life for all Americans.

One of many inputs for evaluating research ideas may be the expected opportunities to apply that research in practice, such as through today's smart and connected communities or other programs addressing national priorities. A good example is the use of campuses and their surrounding cities as experimental infrastructure for wireless research. Given the flexibility now available for spectrum licensing on campus, extending wireless research to a nearby understanding and welcoming community provides substantial benefits for all involved.

Process observations

To help strengthen academic systems research in the United States, NSF may wish to shift its infrastructure funding focus from an emphasis on "funding hardware" to a more balanced approach that provides adequate funding for software development and professional staff for deployment and operations, with lengthened award periods to provide more stable funding for staff. Research infrastructure is an ecosystem comprising a variety of elements from experimental platforms to experiment management tools and artifacts. Shared infrastructure with useful, reliable tools will play an important role in promoting the development of good experimental methodology in computer science.

Building and operating production-quality infrastructure requires professional research staff. They provide greater expertise, long-term consistency, and less subject to the academic calendar than students. With staff as mentors, students can meaningfully contribute to research infrastructure projects. In many cases, students are running experiments for their research projects, and are thus well-placed to identify infrastructure needs, and to do early prototyping of new infrastructure and tools.

Effective research teams are both an important outcome of cyber-infrastructure initiatives and as a crucial ingredient for their success. It is important to improve continuity for the teams themselves through funding programs that are stable over longer time horizons, provide adequate funding to retain skilled development staff over longer periods, and allow accomplishment-based renewals. These teams can also act as a "force multiplier" for researchers to advance a more ambitious cyber-infrastructure agenda that views innovation in this space as itself a research activity. For example, programmatic options to facilitate "embedding" of graduate students and perhaps postdoctoral students with these teams may be useful, e.g., through internships or local interaction. This approach can establish a role for these teams in the education and training mission, and can also help to engage students with research in cyber-infrastructure development.

University research in this area can significantly benefit by interactions with related industries. Such interactions may take many forms ranging from discussions of topics of mutual interest to industry researcher participation on advisory committees to funding support for university-based projects to collaborative projects. Key to success of such university-industry interactions is that they involve researchers or advanced developers/planners who are in a company's research or CTO organization. This helps insure that considerations of current product cycles are avoided and that the interactions remain focused on the overlap between the NSF funded research and similar research or advanced development interests of the company.

The same university research in these areas may also benefit from engagement with the campus IT infrastructure operators. For example, interactions may take the form of campus-based components of these research infrastructures that extend the scope of participation beyond researchers while still maintaining a controlled and limited population. To be successful, such researcher - IT infrastructure collaborations need to engage the central IT infrastructure operators at the planning stage of grants so that a fuller understanding of the needed approaches for at-scale testing can be incorporated into proposed projects. This will also help to ensure that as advances emerge from the research program, they are fully understood and integrated into the planning process for the evolution of the central campus IT infrastructure. Likewise, as experience at scale, as observed by the IT infrastructure operators, can inform the research, there will be a natural conduit for that exchange. Such engagements have the potential to contribute to a "virtuous cycle," furthering the impact of the research programs and funding.

A substantial portion of today's systems research almost by definition needs collaborations and interactions across many groups, both within computer science fields, as well as domain sciences (e.g., physics, bioinformatics, geosciences). The first driver is that most traditional science fields are increasingly becoming data-intensive due to the growing data resolutions from scientific instruments that are engineered to become inexpensive, and more pervasive across university campuses. The second driver is that domain scientists direly need systems experts to help them in managing, processing and sharing their data between geographically distributed collaborators for research and education purposes. In turn, such needs of domain scientists may create exciting laboratories for systems researchers to integrate and experiment with their latest software-defined infrastructure innovations.

NSF investments in recent years have successfully engaged computer science, and domain science researchers in collaboration with university IT staff. They have stimulated advanced integration and accelerated deployment of experimental software for managing networks, computing, storage, security and end-to-end performance measurement. This momentum has been very positive.

CISE might also encourage long-term (5-10 year) collaborations between the Computer Science systems research community and researchers in specific non-CS scientific domains / applications, e.g., via a series of targeted programs that leverage the technologies discussed in this report, and continue to encourage deep, ongoing engagement between researchers and their students, campus CIOs and IT staff, and industry. Examples to date of fruitful collaborations between these groups are the various CC* and the GENI and NSF Cloud projects. Collaboration with industry – particularly industrial researchers – is also to be commended.

A continuation of this process might include: (a) maturing the infrastructure and data management innovations for inter-disciplinary collaborations, (b) investing in research-support personnel across campus groups with expertise needed to effectively use advanced technologies to suit domain scientist needs across multiple institutions, (c) learning and workforce development in areas of advanced cyber-infrastructure that encourages sharing best practices and building knowledge bases for re-use of infrastructure and policy templates, (d) creating models for sustainability that will evolve with pertinent social, behavioral and economic considerations between researchers and university IT staff.

To this end, we contemplate the idea of a “testbed incubator.” Such a program might provide funding for relatively small, short-term, innovative, and high-risk testbed projects (either new testbeds, or new features for existing ones). Testbeds in this program would be expected to develop enough of a proof of concept to attract early, users, and there would be a path for testbeds that are successful in attracting users to “graduate” to larger, longer-term infrastructure projects.

Finally, to ensure long-term productive relationships between researchers and the available and wished-for forms of research infrastructure, NSF might wish to create an effective Coordination Committee. The primary task of such a committee would be to ensure that capabilities provided by existing and new experimental infrastructure are driven by current and emerging research needs of broad CISE community. This can be done by orchestrating an ongoing conversation in which the research community defines the infrastructure needs and infrastructure operators explain the capabilities and opportunities that they can provide. Given the emergence of multiple suites of infrastructure, such Committee would also be in a position to assess and advise on their lifecycle/maturity, potential and options for federation, as well as need for new capabilities. Finally, the existence of the Coordination Committee would help assess the progress of experimental infrastructures and provide a useful contribution towards a practical plan allowing for longer-term funding of experimental Computer Science infrastructure.

Appendix C

Grand Challenge Application Examples

One good way to motivate relevant and grounded areas of research is to pose categories of Grand Challenge applications that require research of the kinds described in this report. This appendix suggests some areas that may be especially relevant in the timeframe of this report.

Grand Challenge #1: Very secure, fault-tolerant critical applications infrastructure

The banking and medical industries would greatly value secure, fault-tolerant critical applications infrastructure. How can and should such trustworthy systems be designed and constructed? Can software-defined slices provide additional security and privacy for sensitive information, perhaps in addition to encryption?

As the amount of cyberinfrastructure grows, could slicing and software-defined infrastructure provide a better building block than simply attaching more devices to a common Internet? Could failures be handled as more routine items, becoming annoying retries instead of disastrous outages? Could the Netflix Chaos Monkey teach us more about how to create more fault-tolerant applications infrastructure?

There would seem to be an opportunity for the financial industry to save billions lost to cyber-fraud if we could better protect sensitive streams, perhaps by putting them in dynamically-created and software-partitioned slices.

Grand Challenge #2: Disaster Preparedness and Response

The group noted that “covering” challenges can motivate requirements in many areas and could be viewed as an “engage everyone” type of challenge. The main suggestion in this category was disaster preparedness and response. It can engage multiple disciplines and multiple areas of computer science in a single grand challenge. There was also sentiment that we should not only look at disastrous scenarios, but also look at more mild or “everyday” disasters as well since they occur far more frequently.

Grand Challenge #3: Health Maintenance and Management

Improving the quality of life and decreasing healthcare costs are both possible with better health maintenance and management. CISE research can make a substantial difference by providing automated monitoring and measurement tools that can feed algorithms giving people advice on improving their health and calling for assistance when needed.

Having computers watch and understand your behaviors seems to be an accepted practice. In addition to the growing use of fitness bands, cameras and sensors could aid with aging in place, providing more comfortable home environments for millions of seniors while making sure they have assistance when needed. Robots and other monitoring agents would feed algorithms that provide feedback to the senior and notify caregivers when outside assistance is needed. Seniors would want the comfort of knowing that people were not watching the cameras; only computers looking out for their welfare.

Chronic lifestyle diseases can be prevented in many cases. These include atherosclerosis, heart disease, and stroke; obesity and type 2 diabetes; and diseases associated with smoking and alcohol and drug abuse. Research could provide patient-specific real-time and longitudinal information needed to provide real-time behavioral feedback as well as longer-term cumulative reports for the person and their doctors and care-givers. Early indications are that using changes in easily-observable long-term continuously-recorded data functions well as indicators for chronic/acute disease propensity and onset. Because only computers are watching the data, there can be automated aggregation / summarization of data while maintaining patient privacy/security by using slices and private clouds. Cognitive processing of continuously delivered wireless data via edge processing appears appropriate.

Grand Challenge #4: Improving nonstop autonomous vehicle flow by scheduling traffic flows.

With adequate vehicle-to-vehicle and vehicle-to-infrastructure communication and smart intersection monitoring, we should be able to reduce pollution and autonomous vehicle travel time by scheduling the passage of each vehicle through the city and its upcoming intersections. Solving this challenge will involve advanced wireless techniques, interconnecting software-defined systems at the edge and via software-defined exchanges. Campus and city infrastructure is helpful here.

Grand Challenge #5: Cognitive systems for intelligent assistants and coaching

Better intelligent assistants and coaches could be significant productivity boosters for the economy and for helping people accomplish more than they could otherwise accomplish by themselves. Artificial intelligence now appears to be coming of age, and current results from Siri, Cortana, and Google Now shows that there are tremendous opportunities to do such things as:

(a) create records of creative work
(b) capture information that facilitates collaboration and sharing
(c) assist with maintaining security and privacy
(d) help fight negative addictive behaviors
(e) provide assistance with just-in-time learning
(f) assist with giving individual advice such as how to escape during disasters or other emergencies

Accurate configuration and programming of the multiple systems involved might be accomplished through cognitive neural networks from information provided by wireless sensors and aggregated and processed at the edge. Campus and city infrastructure is helpful here.

Grand Challenge #6: Realistic Virtual Reality (including augmented reality and mixed reality)

Virtual, augmented, and mixed reality holds the hope of providing high impact educational and health-maintenance experiences. Realistic, Holodeck-quality virtual and augmented reality triggers areas of the brain that reinforce learning. A small experiment with a learning game for 8th graders to teach them Cartesian Coordinates and slopes was far more effective and had longer-lasting impact than the same material taught in a traditional classroom and was far more engaging for he students. The possibility of using high-quality virtual reality for STEM education for higher-risk children and helping them into higher-paying higher-impact careers is compelling. It is easy to think of other uses:

- Accurate surgery visualization would be helpful for both surgeons practicing for surgery and for patients to understand the scope and risks of the surgery for which they are about to give informed consent.
- An elderly population may appreciate the virtual mobility that realistic virtual reality could provide.
- Cultures could be preserved across generations via accurate and detailed virtual reality.

Currently these kinds of applications can be demonstrated using expensive virtual reality headsets and high-powered graphics cards in computers tethered to the headsets. All of this creates an expensive solution. By using gigabit wireless networking, VR/AR could be streamed where needed for only the cost of the headset. This application has been suggested as a “killer application” for software-defined edge computing and wireless.

Grand Challenge #7: Urban Sciences

Urban sciences can also engage multiple disciplines and multiple areas of computer science and other disciplines in a single but encompassing subject. Urban sciences could include cyberphysical systems, people-oriented systems, the interaction of people with their physical and digital environments, and systems to help improve quality of life. Our communities would serve well as living laboratories to understand some of society’s most pressing issues and suggest both preventative and remedial intervention.