Report of NSF Wireless Mobile Planning Group (WMPG) Workshop

New Architectures and Disruptive Technologies for the Future Internet:
The Wireless, Mobile and Sensor Network Perspective

August 2-3 2005

New Brunswick, NJ

The Wireless Mobile Planning Group Workshop was supported by the National Science Foundation under grant CNS-053645

Table of Contents

Table of Figures

1 Executive Summary

Over the next 10-15 years, it is anticipated that significant qualitative changes to the Internet will be driven by the rapid proliferation of mobile and wireless devices, which may be expected to outnumber wired PC's as early as 2010. Wireless devices on the Internet will include laptop computers, personal digital assistants, cell phones (over 2 billion in use as of 2005 and growing!), portable media players, etc. along with embedded sensors used to sense and control physical world objects and events. As mobile computing devices and wireless sensors are deployed in large numbers, the Internet will increasingly serve as the interface between people moving around and the physical world that surrounds them. Emerging capabilities for finding, querying and controlling physical world objects and events via the Internet will result in a broad new class of applications tightly integrated with the way people live and work. The potential impact of the future wireless Internet is very significant because the network combines the power of computation, search engines and databases in the background with the immediacy of information from mobile users and sensors in the foreground.

The data flows and interactions between mobile users, sensors and their supporting computing infrastructure are clearly very different from those of today's popular Internet applications such as email, instant messaging or the world-wide web. Supporting these requirements efficiently in an integrated wired-wireless global network will require major changes. Thus, the broad architectural challenge facing the wireless and network research communities is that of either evolving or redesigning the Internet architecture to efficiently incorporate emerging wireless network elements such as mobile terminals, ad-hoc routers and embedded sensors and to provide end-to-end service abstractions that facilitate application development. A top-down approach to the problem starts by identifying a set of canonical wireless scenarios that cover a broad range of environments such as cellular data services, WiFi hot-spots, Infostations, mobile peer-to-peer, ad-hoc mesh networks for broadband access, vehicular networks, sensor networks and pervasive systems. These wireless application scenarios lead to a diverse set of service requirements for the future Internet that need to be analyzed and considered for protocol design and implementation. Some examples of notable requirements identified at the WMPG workshop are:

1. *Naming and addressing* flexibility
2. *Mobility support* for dynamic migration of end-users and network devices
3. *Location services* that provide information on geographic position
4. *Self-organization and discovery* for distributed control of network topology
5. *Security and privacy* considerations for mobile nodes and open wireless channels
6. *Decentralized management* for remote monitoring and control
7. *Cross-layer support* for optimization of protocol performance
8. *Sensor network features* such as aggregation, content routing and in-network processing
9. *Cognitive radio support* for networks with physical layer adaptation
10. *Economic incentives* to encourage efficient sharing of resources

Taken together, the above wireless/mobile/sensor requirements represent a significant new network design challenge. Although steady evolutionary improvements to the Internet Protocol (IP) cannot be ruled out, there is a broad consensus among panelists that a "clean-slate" architectural framework for the Internet will be required to adequately meet future needs. Competing strategies for deploying new protocol architectures for wireless include top-down transformation of the network using protocol overlays built on IP vs. bottom-up replacement of the entire IP network and protocol layers above it.

The wireless, mobile and sensor network research communities have produced many new results on future protocols during the past decade, but much of it in an isolated context. The next challenge for this community is to collaborate with end-to-end Internet protocol, overlay network, distributed systems and domain-specific application communities to find viable architectures and protocols for the next-generation Internet with billions of wireless devices. In addition to foundational theory and design, any such effort will require a dramatic change in experimental network research capabilities, not just for wireless edge networks but also for large-scale end-to-end system evaluations. Wireless systems involve real-world mobility and channel impairments which are difficult to model via simulation, so that there is an important need for real-world testbeds and experimental networks for validation, comparison and selection of new protocol ideas.

Existing wireless research testbeds do not support evaluation of large-scale networks in real-world settings, nor do they provide for end-to-end experimentation over a wide-area network with programmable protocols – this motivates the construction of a new generation of flexible experimental wireless networks. Recent work on wireless testbeds and platforms has led to the availability of several key technologies for building flexible systems. These include open API radio modules; cognitive radios with software-defined physical layer; MAC-layer virtualization techniques; general purpose platforms for sensors, ad hoc nodes and radio routers; measurement techniques; and network emulators and simulators. These building blocks now make it possible to implement a new generation of flexible experimental wireless networks that can support a variety of protocol concepts and evolve to accommodate rapid changes in core radio technology. Specific recommendations for building experimental wireless networks in support of future network protocol and applications research are:

- *Wireless emulation testbeds* integrated with wide-area research networks
- *Urban ad hoc mesh network* with ~1000 radio nodes in a dense city deployment
- *Suburban wide area network* with open access 3G/WiMax and 802.11 hot spots
- *Cognitive radio technology demonstrator* with shared spectrum in dense deployment
- *Application specific sensor network* deployments
- *Wireless infrastructure* for the Internet backbone

To facilitate end-to-end protocol experimentation, each of the above wireless testbeds needs to be connected to a programmable wide-area Internet backbone capable of supporting a range of architectures and protocols. Concepts for a large-scale experimental network with innovative programmability and virtualization features as described in a related NSF study entitled “Overcoming Barriers to Disruptive Innovation in Networking” [Pet05] are considered appropriate for this purpose. Such flexible wireless networking infrastructure connected to a programmable wide-area backbone may be expected to foster critical research and experimentation on innovative protocols, middleware and applications for the future Internet.

Overall recommendations of the WMPG panel to NSF are summarized as follows:

Recommendation 1: Recognize that wireless will drive a fundamental transformation of the Internet during the next 10-15 years, and invest in research aimed at creating the necessary technical foundations.

Recommendation 2: Increase research focus on central network architecture questions related to future mobile, wireless and sensor scenarios.

Recommendation 3: Invest in the development of flexible wireless technologies and platforms necessary to implement programmable and evolvable experimental networks.

Recommendation 4: Fund development of large-scale experimental wireless networks for validation and competitive selection of new architecture and protocol concepts in presence of real users and applications.

Recommendation 5: Encourage collaborative research that would result in end-to-end deployment and evaluation of future wireless/mobile and sensor networks and applications over the global Internet.

2 Introduction

2.1 Project Background

In July 2005, the National Science Foundation (NSF) funded a planning project aimed at involving the wireless and mobile network research communities in a discussion of novel architectures and disruptive technologies for the future Internet. The objective of the planning effort is to solicit service requirements and innovative network architecture concepts from wireless, mobile and sensor net researchers at universities and industrial/government research labs, and to consolidate this input into a coherent vision/agenda for future research and experimental infrastructure needs.

The NSF planning project is being carried out by a broad cross-section of academic and industrial researchers referred to as the NSF Wireless Mobile Planning Group (WMPG). After a period of discussion on mailing lists, the WMPG group held its first workshop on Aug 2-3 at Rutgers University, New Brunswick, NJ. This workshop, which had about 25 panelists, also included participants from the Internet architecture research community and representatives from NSF. A second workshop is planned for late 2005/early 2006 to finalize recommendations and to harmonize conclusions with other planning groups working on the future Internet.

This interim report is intended to record the findings and recommendations emerging from the Aug 2-3 WMPG workshop. A more complete final report will be issued by the WMPG group in 2006. This draft report is based on material contributed by workshop participants at an all-day plenary meeting held on Aug 2, followed by a second ½ day on Aug 3 with the following breakout sessions:

- The Future Wireless Internet – scenarios and requirements
- Network Architecture Considerations
- Experimental Infrastructure for Wireless Internet Research

2.2 Wireless/Mobile Rationale for the Future Internet

Wireless, mobile and sensor networks represent an increasingly important segment of networking research as a whole, driven by the explosive growth of portable computing, communication and embedded devices connected to the Internet. For example, laptop sales exceeded those of desktop PC's in 2003 and this trend towards compact and portable computing devices continues unabated. As of 2005, it is estimated that there are over 2 billion cell phones in use worldwide as compared with 500 million wired Internet terminals, and a significant fraction (~20%) of these phones now have data capabilities as 2.5G and 3G cellular services are deployed. In another 5 years, all cell phones will be full-fledged Internet devices, implying inevitable changes both in applications and network infrastructure to support mobility, location-awareness and processing/bandwidth limitations associated with this class of end-user terminals.

Another major category for growth in wireless devices is that of embedded wireless devices or sensors that help to monitor and control objects and events in the physical world via the Internet. Embedded wireless technologies are still at an early stage, but various market predictions estimate that the price of a wireless sensor chip could drop to ~$1 by 2015 resulting in potentially tens of billions of these devices on the Internet in the long-term. Wireless sensors are fundamentally different from other Internet terminals due to severe power and processing constraints and the greater importance of function, content or location over network address. Data flows and interactions between mobile users, sensors and their supporting computing infrastructure are clearly very different from that of today's popular applications such as email, instant messaging or the world-wide web. Thus, embedded wireless devices and the new classes of applications associated with them may be expected to drive further changes to the Internet architecture and protocols in order to account for very different service needs.

Overall, it is clear that mobile, wireless and sensor devices will certainly outnumber wired end-user terminals on the Internet in the near future, strongly motivating the consideration of fundamentally new network architectures and services to meet changing needs. The currently deployed TCP-IP model of the Internet was originally designed for communication between wired PC's, mainframes and servers, and is clearly not optimized for wireless devices. Several attempts at incorporating mobility-related extensions to IP have been made within the IETF (in particular, mobile IPv4 [Per02a, Per02b], and later IPv6 [Jo04]), but adoption of these evolutionary protocol standards has lagged behind expectations. Also, the mobile IP initiative represents only one part of the technical requirements posed by wireless computing devices. For example, changes to the TCP transport layer service [Ba95, Ba97, Ch04b, and Ca95] are required to deal with wireless channel impairments, while end-to-end security for mobile devices on the Internet remains an open problem. Some of these issues are being considered by IEEE 802.11 committees [IEE04], but these solutions remain specific to a single radio standard and are not necessarily harmonized with IP. Support of ad-hoc networks and emerging sensor network applications poses even more complex technical issues, which can to some extent be avoided by creating customized overlay networks [Pet02, Pet04] with IP tunnels between the network nodes at the expense of creating specialized systems that can be used by only a subset of Internet users.

The broad architectural challenge facing the wireless and network research communities is that of evolving or redesigning the Internet architecture to be able to efficiently incorporate emerging wireless network elements such as mobile terminals, ad-hoc routers and embedded sensors and to provide end-to-end service abstractions that facilitate application development.

2.3 Wireless Technology Roadmap

Wireless and mobile networks represent an active research and new technology development area. The rapid evolution of core radio technologies, wireless networks/protocols and application scenarios is summarized for reference in the technology roadmap given in Fig. 1 below.

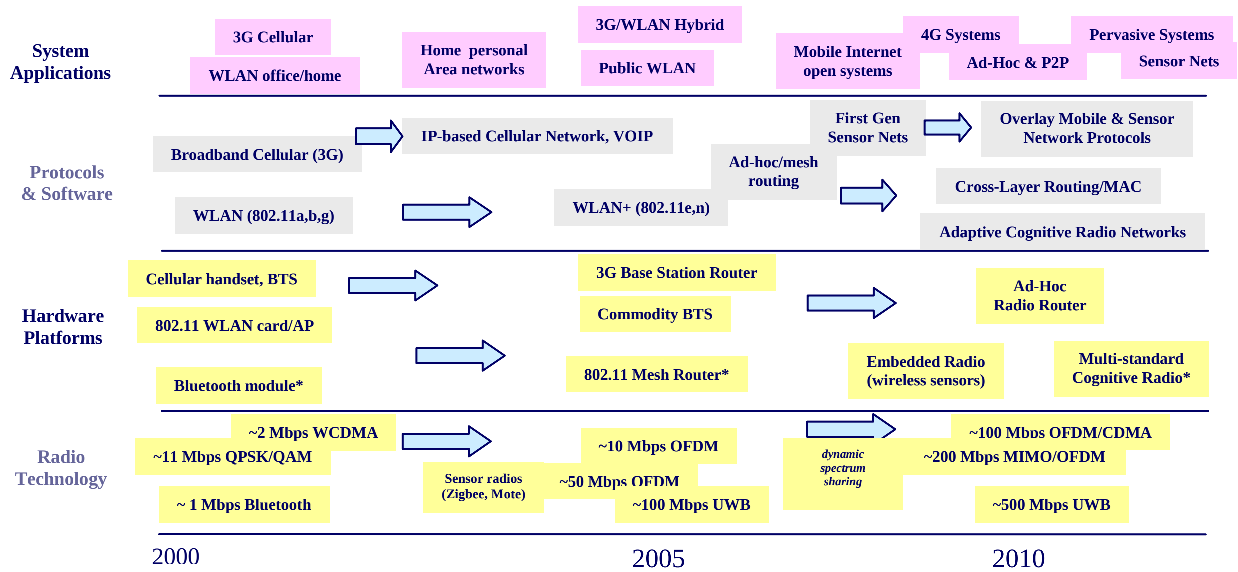

Figure 1 Wireless Technology Roadmap for the Period 2000-2010

It can be seen from the chart that in addition to 2.5G/3G cellular data and WLAN systems developed during the 1990's, emerging wireless scenarios include personal-area networks [Jo02, Ge01], wireless peer-to-peer (P2P) [St01], ad-hoc mesh networks [MIT, LOC], cognitive radio networks [Mi

99,Yat02,Ch04a, Bo04], sensor networks [Yar02, Cu04, Es99], RFID systems [Ra99, Hu98, TEX] and pervasive computing [He03, Al03].

Each of the above mentioned wireless scenarios is associated with unique network architecture and service requirements that affect both the access and infrastructure portions. The default approach adopted by most of the research community is to treat the wireless access portion as a “layer 2” local area network connecting to the Internet (i.e. layer 3 IP) through a gateway. This approach is pragmatic, but it precludes uniform dissemination of control and routing information through the entire network and creates a potential processing bottleneck at the gateway. A more integrated end-to-end control and routing architecture is important for optimizing mobile/wireless service features such as location management, dynamic handoff, quality-of-service (QoS) or cross-layer transport. Also, a local-area wireless network may contain one or more routing elements which can create inconsistencies in protocol layering and addressing. If compatibility with the current IP network is not viewed as an essential constraint, it may be possible to develop a “clean slate” network architecture that can accommodate emerging wireless networks in a single unified protocol structure. The next section examines emerging wireless, mobile and sensor scenarios in further detail with the objective of identifying specific networking requirements that should be supported by the future Internet architecture under consideration.

3 Wireless Scenarios in the Future Internet

The first task of the panel was to analyze the networking requirements associated with prominent future wireless scenarios. In order to introduce some structure in the requirements definition process, the panel followed a top-down approach in which representative scenarios and the applications driving them were identified first, and the requirements that these scenarios placed upon the network were then mapped out. The group converged on three canonical scenarios that together capture the key features of most emerging wireless systems:

3.1 Scenario A: Individual Wireless Devices Interfacing with the Internet (“Mobile Computing”)

The simplest scenario involves a single wireless device that interfaces with the broader Internet. The mobile device may be a cellular phone, a PDA, a media player, a digital camera or some type of combination consumer device. Mobile computing devices may connect through a wireless local area network, a mesh-style wireless network, or a wide-area wireless technology (such as cellular 3G or WiMax). Service models to be considered include mobile services, hot-spot services with limited mobility, as well as cached content delivery via opportunistic wireless links. High mobility, the potential for intermittent connectivity, and heterogeneity of radio access are key characteristics of this scenario.

A typical example of this mode of operation is that of a mobile customer downloading a real-time video stream (e.g., a live sporting event) to a portable media player from the Internet. Seamless connectivity should be maintained as the customer moves from a shopping mall (WiFi coverage) to outdoors (2.5G or 3G cellular connectivity) and then to the car (Bluetooth within the car, WiMax radio to the Internet). At each step, the wireless media player needs to be aware of available connectivity options and then select the best service. The multimedia server must also be aware of current connectivity constraints so that it can deliver a stream with parameters (data rate, format, etc) consistent with the configuration. The same mobile customer should be efficiently tracked by the network and reachable by VoIP calls, if he/she so chooses. Location- or context-aware queries (such as “where is the nearest pharmacy?”) and delay-tolerant services (e.g., seamless suspension and resumption of a large file transfer when the user walks or drives through areas without coverage) should be supported. Caching of files for rapid downloading within a hot-spot may also be useful in this scenario.

3.2 Scenario B: Constellations of Wireless Devices ("Ad-hoc Nets")

The second type of wireless scenario is motivated by a variety of settings in which multiple radio devices may be in close physical proximity and can collaborate by forming an ad-hoc network. For example, wireless devices in an office or home environment can set up an ad-hoc network between themselves to improve coverage and communications quality. Another popular application involving constellations is that of community mesh networks formed by rooftop radios for the purpose of shared broadband access. In the important emerging application of automotive telematics, clusters of cars on the highway may participate in an ad-hoc network for the purpose of collision avoidance and traffic flow management. Constellations may include heterogeneous radio and computing devices with different capabilities and resource levels. Emerging cognitive radio technologies also offer the capability of highly adaptive wireless ad-hoc networks with physical layer negotiation between nodes, practically scavenging unused spectrum at low cost to support a private ad hoc network. Opportunistic association, changing network topologies, varying link quality and potentially large scale (in terms of number of nodes) are some of the characteristics of this scenario.

A simple example of opportunistic constellations is the formation of an ad hoc network between several user laptops in a meeting room with limited Internet access coverage. The ad hoc network enables high bandwidth communication between participants at the meeting and allows them to use a favorably positioned (e.g., with good cellular network throughput) node as a forwarding relay to the Internet. Another example is the cooperative downloading of popular files from the Internet by drivers on a highway. Suppose hot-spot "infostations" with WiFi service are spaced by several miles on the highway and a car traveling at 60 miles/hr can download only fractions of a MB during each contact with an infostation. If several drivers are interested in the same file, it is possible for the cars to exchange segments in a P2P opportunistic networking arrangement similar to that used in Bit Torrent. This way, the download can be completed without requiring a car to stop at a hot spot, saving time for the end-user and avoiding traffic congestion problems. The same ad hoc networking capability can also be used by cars to exchange control information necessary for traffic flow management or collision avoidance.

Ad hoc radio constellations also apply to civilian disaster recovery and in tactical defense environments. These applications usually involve communications between a number of first responders or soldiers who work within close proximity of each other. The response team may need to exchange text messages, streaming media (e.g. voice or video), and use collaborative computing to address a shared task such as target recognition or identification of a spectral jammer. Individual nodes may also need to access the Internet for command and control purposes or for information retrieval. This application has similarities with the ad hoc mesh network for suburban or rural broadband access mentioned earlier.

3.3 Scenario C: Pervasive Systems and Sensor Networks ("Sensor Nets")

Sensor nets refer to a broad class of systems involving embedded wireless devices connected to the Internet. The first generation of sensor networks involves collecting and aggregating measured data from large numbers of sensors in a specified geographic area. In the near future, sensor net applications will also include closed-loop sensor/actuator systems for real-time control of physical world objects. Current sensor net applications are in science (ecology, seismology, ocean and atmospheric studies, etc.), and engineering (water quality monitoring, precision agriculture, livestock tracking, structural monitoring), as well as consumer-oriented applications (home security and energy management, hobbyist and sports enthusiast applications of distributed imaging, eldercare, pet monitoring, etc). Sensor networks share several characteristics of ad-hoc scenarios but are differentiated by the fact that tiny sensor devices have more stringent processing power, memory and energy constraints. These constraints generally imply the need for a hierarchical ad-hoc network structure in which low-tier sensor nodes connect to the Internet via one or more levels of repeating wireless gateways. Other important characteristics of this scenario are the data-centric nature of applications, potential for large scale (in terms of numbers of sensors) and geographic locality.

Traditionally, large "sensor fabrics" such as those installed to monitor the environment have been designed as vertically optimized systems, with an ad hoc network designed to meet specific energy and processing constraints and optimized to support specialized queries dictated by the application at hand. The interface to the Internet has been via edge nodes that isolate the Internet stack from the sensor fabric architecture. However, more recent trends indicate an increased need for sensor networks that provide open access via the Internet, in a more extensive and capillary way that can be supported via edge nodes. For instance, scientists interested in the correlation between data found in different data bases, (e.g. soil characteristics, pollutants carried in the local water supplies, productivity of local vineyards, production and sale of local wines), may be permitted to access specific regions within a sensor fabric directly from the Internet, to extract the required data rather than overburdening the access gateways . Moreover, new types of sensor networks based on "mobile" sensor platforms are becoming available—for example vehicles in the urban grid or firefighters in a disaster recovery operation equipped with a variety of sensors (video, chemical, radiation, acoustic, etc). These sensor platforms have practically unlimited storage, energy and processing resources. The vehicle grid then becomes a sensor network that can be accessed from the Internet to monitor vehicle traffic congestion and to help investigate accidents, chemical spills and possible terrorist attacks. Likewise, firefighters carry cameras and several other sensors, allowing the commander to be aware of the conditions in the field and to direct the operations to maximize the use of the forces, while preserving the life of his responders. These latter examples also show that the gap between sensor networks and ad hoc networks tends to diminish in mobile sensor systems at least in terms of communications capabilities and Internet access. In the longer term, pervasive systems involving personal mobile devices, smart offices/homes and densely deployed multimodal sensors/actuators will serve as a platform for development of various new applications ranging from tracking and inventory control to personal productivity, public safety and resource management.

These scenarios are not all-inclusive, but serve as canonical examples of future wireless networks. They do capture many of the requirements that wireless will place upon the future network architecture. Based on the scenarios and keeping the main objectives in view, the panel proceeded to a more detailed discussion of architectural considerations for the future Internet.

4 Network Architecture and Protocol Design Considerations

The wireless scenarios outlined above are associated with new network service requirements that motivate a rethinking of several Internet architecture issues. Several mobile/wireless features may require mechanisms that cannot be implemented through the conventional IP framework for the Internet, or if they can, may suffer from performance degradation due to the additional overhead associated with network protocols that were originally designed for static infrastructure computing. In this section, we discuss a set of architectural and protocol design issues related to the networking requirements of the representative wireless scenarios identified earlier.

4.1 Distinguishing Identity and Address for Mobility Management (Scenarios A, B)

The Internet network layer of today equates an address with a network attachment point. This facilitates reasonably simple routing and address assignment because it can take advantage of logical locality, but it poses challenges when nodes are mobile. Today, limited node mobility is supported through techniques at different levels in the protocol stack. When mobile users roam inside the same IP subnet, they can continue to use the same IP address and mobility can be implemented at the link layer. Examples include roaming in 802.11 and cell phones. When mobile users cross IP subnets, they will typically have to change the IP address they use, which can be visible to applications and consequently affect their operation. It is possible to deal with cross-subnet mobility at the application level but this can complicate application development and degrade performance. The level of naming indirection DNS introduces can allow very coarse grained mobility, but fails when nodes require uninterrupted communication across address changes (e.g., a telephone call). An alternative is to use Mobile IP, an IETF standard that supports mobility by allowing applications to continue to use the same IP address even as the node moves between IP subnets [Per02b]. Adapting these partial solutions for an Internet where most nodes, including routers and switches, are mobile, will likely lead to solutions that are expensive and inefficient.

To fully support mobility, a next generation Internet must provide ways to name and route to a much richer set of network elements than just attachment points. It should support routing in terms of *names*, which identify the actual desired end points, rather than some particular characteristic, such as an address. These names may well resolve to addresses at some layer in the architecture, for routing purposes; what is important is that the architecture supports seamless changes in name-address binding. Names may or may not have global consistency: they can represent anycast elements, supporting local resource discovery. Examples of names include "CNN's web server," "Vint Cerf's cell phone," and "a printer nearby." Another option is to directly support naming and addressing of clusters of nodes to improve support for communication with moving “constellations” such as car and body-worn networks [Na05]. Making “location” an explicit concept in the Internet architecture might also help support mobility across IP subnets.

4.2 Delay tolerant, disconnected operations (Scenarios A, B, C)

The next generation Internet needs to support communication with end points that are not always connected. These disconnections can be short or long-lived. The current approach handles disconnections above the network layer, as in the email system. This is adequate when most devices are perpetually connected to the Internet. With mobility, however, disconnections will be common. This requires a new solution to avoid every network service implementing mechanisms or protocols to support disconnection, which would result in a complex mix of layers above.

One way to support delay-tolerant networking is through new network services, such as the “Infostation” model [Fr96, Go97] or the “data-bus” service in which data is cached and delivered when an appropriate receiver is within range. In order to support “Infostation” services, it is necessary to have mechanisms that exploit mobility information to pre-fetch data and facilitate quick authentication as mobiles pass through coverage zones [Ku01]. Similarly, for data-bus services, the network must manage data caching by identifying appropriate points within the network to cache data, route data to these cache points, and support the delivery of data when suitable receiver devices approach.

Two mobile devices should, by themselves, be able to communicate within a next generation Internet architecture, as should larger collections of nodes that have local, but not global infrastructure. Fully decentralized infrastructure will expand the architecture from a monolithic entity to a set of principles, protocols, and services that operate in all of the edge cases that mobility and wireless introduce. Moreover, if the mobile has multiple interfaces and thus different path options, and is engaged in a delay tolerant application, it must be given the opportunity to transmit now or later depending on path cost, quality, security, etc [Fa03].

4.3 Location Awareness (Scenarios A, B, C)

The current Internet was designed for stationary terminals, thus it does not track node location. In mobile wireless systems with localized interaction, however, the identities of terminals participating in the local network change so frequently that location becomes a more natural primitive for addressing and scoping the distribution of computing resources relevant to a local area. This trend is further strengthened by improvements to 802.11 localization and global positioning system receivers that increase the availability of precise location information at mobile terminals. Beyond location as an addressing mechanism the architecture must enable the use of location information for internal network optimizations at all layers and support the development of location-aware applications. Specific point solutions such as geographic routing, mobicast, localization, or location fusion algorithms have been described in the literature [Ka00, Na02, Li00]. These need to accommodate additional requirements, such as high node mobility, which may be expected in automotive networking scenarios and foremost these solutions need to be integrated into a coherent architecture [Yu04].

An architectural proposition is to construct a location service that enables potential communication partners to determine each other's position and to identify the terminals in a specified geographic region [Pr01]. This service can be extended across the Internet backbone as an overlay [Da04]. The architecture should also define interfaces for network protocols to acquire location information of sufficient accuracy using localization techniques. One aspect of these interfaces is how location information can be represented in the system. Many applications prefer symbolic representation such as postal addresses or room and building numbers, which users can more readily understand. In the lower network layers, however, geographic coordinate systems based on an ellipsoid model of the earth are more amenable to processing requirements such as distance comparisons. Specifically, a Cartesian coordinate system, such as the Universal Transverse Mercator, enables fast distance computation compared to a latitude-longitude system (except when crossing zone boundaries). The architecture must likely include provisions for translation between different representations across different network layers.

4.4 Security and Privacy (Scenarios A, B, C)

Wireless networks are increasingly the platform of choice for launching a variety of attacks on the Internet [Bo03, Bell03, Bor01, Ko02, Per02c]. This is due to the lack of a physical connection, potentially weak authentication, and lack of accountability. New network architectures will need to elevate security and trustworthiness issues to an equal level of importance as other networking concerns. At the most basic level, wireless devices will likely have evolving naming and addressing schemes and the addresses used must be verifiable and authenticated. Wireless devices will connect to each other and to a backbone network in order to relay information. Authentication must also cover link layer associations with access points or peer devices and network control information (e.g. routing tables) necessary to form mesh networks or connect to the backbone network. Since the wireless medium allows for adversaries to easily sniff network traffic, it will be essential to provide mechanisms that prevent traffic analysis by adversaries.

Location information that nodes advertise must be trustworthy and reliable, but a location service that tracks the position of network nodes also raises significant location privacy concerns [Gr03]. Location conveys a rich set of potentially private contextual information. Even if devices are only identified by pseudonymous addresses, an adversary could learn the user's identity if the location is linkable to publicly available, identified location records. For example, if a query originates from a car on a suburban driveway, the likely household could be identified by looking at a public address database. The location architecture should incorporate privacy mechanisms across all layers to address location privacy concerns. At the link layer this requires changing interface identifiers and scheduling sufficient communication pauses to prevent the linking of new and old identifiers. At the network and higher layers anonymity can be improved by automatically adjusting the resolution of reported location information based on the density of users in a region. In addition, access control mechanisms based on user-defined policies can control the disclosure of location records when user identity is required. Research is needed in understanding the tradeoffs between privacy and accountability of different architectural propositions.

The wireless network should also be robust and flexible enough to ensure the availability of the network in the presence of adversarial (intentional or otherwise) threats. In particular, wireless networks should be able to adapt to the presence of radio interference (jamming) attacks. Further, defense mechanisms should be in place in order to defend against denial-of-service (DoS) attacks involving message flooding, such as those targeting various network buffer resources. In addition, there must be mechanisms in place to diagnose and prevent individual devices from improperly following resource management protocols (perhaps for the purpose of achieving their own improved QoS at the expense of other device's degraded performance). All security mechanisms should carefully consider the usage of device resources, the associated performance tradeoffs, and the latency involved in achieving security.

4.5 Management and Diagnostics (Scenarios A, B)

Deployment and use of new wireless edge networks (such as ad-hoc nets or sensor nets) will not conform to the traditional organization/enterprise boundaries corresponding to today's administrative domains. Data exchange is likely to occur not only within each subnetwork but across networks with multiple owners. In such scenarios, the fundamental problem is the one of manageability. How are failures diagnosed and communicated? How should end-to-end applications be structured for exploiting multiple interfaces and work around? In these networks, non-traditional metrics, such as manageability, diagnosability, configurability, and trustability will take precedence over traditional metrics of performance. This will require mechanisms for distributing functionality and addressing, instrumentation of network components across all layers, decomposition of information for automated processing, reasoning and response.

Specifically, network architectures should be developed that are self-configuring and self-healing, while not requiring expensive (and complex) infrastructure or protocol overhead. As wireless networking scales in size, loading, and mobility, auto-configuration services that have been provided in small networks, must now be extended to the Internet at large. Auto-configuration services will likely be needed at all levels of the protocol stack [Va02]. Potential research directions might involve forwarding based on semantic names (e.g. "AV Controller"), eliminating the need for (DNS-like) name-to-address translations, self-protection against nodal security and network intrusion attacks (auto-immunization), and establishing trust without a central authority. Perhaps most importantly, many new research challenges arise when networks with different technologies, and/or networks belonging to different organizations with different policies must be auto-configured.

Furthermore, the ability to manage a network (whether automatically or manually) is inextricably tied with the ability to observe and measure the network and its components. Most currently available approaches towards measurement provide coarse-grained counters such as SNMP variables, or long time-scale flow-based summarizations. From an architectural standpoint, research is needed to determine what finer-grained currently-hidden information (e.g., wireless channel characteristics, MAC protocol and routing information) can be exposed to the user, to network management functions, and to others (e.g., carried in protocol headers) to help manage (e.g., diagnose and (auto-)configure) a network. Finally, realizing that enhanced manageability may result in additional overhead, tradeoffs between increased manageability and consequent decreases in data plane performance must be quantified.

4.6 Exposing network information to higher layers (Scenarios A, B)

The end-to-end principle, on which the current Internet architecture is based, simplifies the internals of a network. While this has been a very useful principle at the network layer, experience has shown instances where higher layers can work around it to improve performance or introduce new functionality. Examples include HTTP proxy servers, peer-to-peer networks, and CDNs such as Akamai. The diversity and growth of these services means that trying to include them into the architecture would be problematic at best. However, the next generation architecture should acknowledge they represent an important usage of the network and, if possible, should introduce functionality that could simplify or improve their deployment.

The complexity of this functionality introduces a tradeoff between its benefit to layers above the network and its cost to the network layer itself. Recent work on DHTs and overlays has demonstrated how important route quality estimation (in terms of latency, bandwidth, etc.) is, while CDNs spend a significant amount of effort to determine what CDN distribution point is closest to a client [Ng03]. One of the complexities these problems deal with is the opaqueness of routing information. While these services can use heuristics such as address closeness or DNS server selection, the architecture itself provides little information on the paths and topologies in the network. Providing such information could be a very inexpensive way to provide tremendous benefit to these sorts of services. Topology information can also be used by edge systems that do not always obey the end-to-end principle. For example, embedded sensor networks have resource constraints that make in-network processing and aggregation important primitives. Rather than build customized solutions, these networks could take advantage of the topology information that the infrastructure provides.

Applications will also benefit from a network architecture that exposes topology and path quality information. If a server discovers that the mobile client is now connected via a "thin" 50Kbps connection with high loss rate, it may reconfigure the format of the delivery (say still frames and audio, instead of motion video). End users can learn about path quality in several ways: (a) through end-to-end probing techniques (e.g., active techniques like packet pair derivatives such as CapProbe [Ka04] and PathRate [Ch05]; or through passive techniques embedded in the transport and streaming protocols, e.g. TCP , TFRC etc); (b) via feed back in the end-to-end data stream from intermediate routers (e.g. ECN, Loss Notification, etc); or (c) with the assistance of network proxies at the edges (e.g., "distance" service, congestion information, etc) [Do01]. In the future Internet one will expect a combination of techniques (end to end and network supported) to provide timely assistance to mobile users. Moreover, some of the path monitoring techniques will apply across multiple domains; others will be confined to specific subnetworks.

4.7 Ad-hoc constellations and groupings - identity hierarchy (Scenario B)

Constellations illustrate how the flexibility of logical naming can enable a wide range of new communication services and applications. In addition to each element in the constellation having its own name (or set of names), the constellation itself can have names. Exactly how this name resolves is up to the constellation to decide: it might refer to a body network server that dispatches various communication streams or packets to other devices. While names can have hierarchies, the Internet architecture should not impose particular hierarchies, as they can and will depend on semantic information.

Nodes outside the constellation can communicate not only with each element-- which has its own name--- but also with the constellation as a whole. The services such a name provides are greatly dependent on its semantic meaning: body networks might provide traffic aggregation to egress points, while infrastructure networks (e.g., a LAN) might have management services that deal with the network as a whole [Ge02].

4.8 "Infrastructure free" communications and self-organization (Scenario B)

Wireless devices should be able to operate independently of the broader Internet. In particular, there may be times during which the connection of a wireless device or network to the Internet is not available. During these times, wireless devices should be able to operate stably in modes disconnected from the rest of the infrastructure, as well as be able to opportunistically establish "local" ad-hoc networks using their own native protocols. This will require that devices are capable of discovering each other without the assistance of the broader network, as well as require that devices use local addressing and routing schemes. Further, when the Internet infrastructure is made available at later times, there should be simple and transparent methods for reinsertion of the disconnected device or network into the broader network. In particular, this means that issues such as authorization and updating the device state should be seamless, with minimal latency.

4.9 Efficient use of multiple outgoing connections (Scenarios A, B)

Future wireless devices will not have a single interface to a single network, but will be capable of supporting multiple, simultaneous communication technologies. A new class of transport layer protocol will need to be developed to support multiple outgoing connections. In particular, issues such as rate adaptation and buffering will be critical at the interfaces between different communication standards. Further, issues of connectivity management and maintenance are made more challenging by the fact that many of the potential connections will be opportunistic connections to ad-hoc networks with dynamically changing topologies. Establishing and tearing down these cross-network connections must be efficient in terms of protocol overhead and latency. Further, the "opportunistic" exploitation of multiple resources must be supported by techniques that provide efficient path quality monitoring. For example, it will be desirable to know the level of congestion on a specific subnet, or the pricing per packet on a particular subnet. Physical layer characteristics, such as the type of radio, whether a particular interface feature is supported, and which portions of the spectrum are available will play a role here.

4.10 Decentralized trust models (Scenario B)

In applications involving multiple constellations of wireless nodes, and it will be desirable to allow network devices to be exchanged between them. Therefore, the future wireless Internet should support secure dissociation and re-association of devices across constellations. Such handoff procedures will be part of the requirements for automated management of constellation membership, and consequently should ensure low-latency and seamless continuity of service. These handoffs should be authenticated, and access privileges should be verified prior to successfully completing network level associations. Further, it will be desirable to support the ability for devices to maintain connections to more than type of network, and even allow for information to be delivered to multiple devices simultaneously.

Further, as devices will move across constellations, mechanisms should be in place to allow for networks to assess and convey measures of entity of trust during handoffs between constellation domains. Devices that move between networks should have a consistent basis of quantifying their security privileges in each domain. The protocol specifications to establish and convey trust must be decentralized as it is unreasonable to expect a static, centralized authority in the future Internet.

4.11 Incentives and economic issues in spectrum usage (Scenario A, B, C)

Wireless technologies will play an increasingly important role in the next generation of the Internet. In particular, the role of wireless technologies will not be limited to last-hop networks, but will also serve to replace traditionally wired infrastructure. As an example of how wireless technologies have had an economic impact, there are many municipalities that are reevaluating their communications needs and looking to wireless mesh networks as an alternative to providing city-wide coverage. Whether considering wireless technologies at the microscopic device level, or at the macroscopic infrastructure level, the new network should have methods in place to take advantage of the economic advantages that wireless technologies promise, as well as encourage the deployment of new forms of wireless networks that have yet to be developed. At the macroscopic level, there must be support for AAA services that will take care of accounting and billing as wireless users move across different networks. Further, appropriate tools must be integrated into future networking protocols to encourage service providers to adopt the new wireless Internet, as well as provide a means for vendors to add value to wireless services. At the microscopic level, however, methods must be in place to support new forms of commerce that are possible by negotiating wireless resources, such as spectrum. Protocols that coordinate spectrum and pricing mechanisms that guarantee the efficient and fair sharing of spectrum will provide new forms of commerce and encourage the development of new software and hardware industries.

Another type of incentive is required to make opportunistic networking work using third party nodes as store and forwarders. In a world of “constellations” and dynamically changing connectivity in urban and “mall” areas, it will be extremely beneficial to exploit neighbors’ resources for a number of services (data forwarding, data ferrying, information stations, etc). For example, in an urban vehicle network, one may wish to download location significant advertisements from neighboring cars rather than from the cellular network or from the next "infostation" down the road. This sharing of resources can be achieved only through incentives and with strong security guarantees, as discussed above. The incentive mechanism will have an impact on the "network layer" as credits should be transferable across (and therefore recognized by) different network domains [Sa03].

4.12 Cross-Layer Protocol Support (Scenario A, B)

There are a number of scenarios in wireless networking in which cooperation between physical, data link, network and transport layers [To03, Pa03, Ba02, and De03] can provide significant performance improvement. As a very simple example, packet loss on a noisy wireless link can be overcome by local channel coding as well as end to end application coding. Often, the optimal solution requires a combination of both. This optimum can be achieved only if the radio layer and the application cooperate and exchange information in some way. As another example, in an urban environment the cognitive radio installed in a car can establish "links" of different radio range at different data rates. The network layer may need this information to determine the best route, say, between two distant cars – cellular network or direct link. Similarly, in 802.11 based ad-hoc networks, rate adaptation in radios can complicate routing – for example, a short path with a 1Mbps 802.11b link would be less desirable than a longer path with all 11 Mbps links. The issue in most of these cases is, how much of this "radio layer optimization" must be visible to upper layers through "cross-layer" interaction; and, correspondingly, what mechanisms must be provided to make it happen. In some cases, PHY layer link quality feedback can be sufficient. In other cases, like in the joint channel and source coding example, explicit mechanisms must be put in place.

It may be desirable to go beyond cross-layer alone and re-examine what functionality belongs to which layer. In the past, protocol stacks have largely treated the wireless links as equivalent to wired links possibly with a higher error rate. While such an approach has served the purpose so far, it appears that this approach may not suffice in the future. In "cooperative wireless networks" now at an early research stage, the notion of a "link" may have to be revised. Some examples:

- *Ad-hoc networks*: In such networks, hosts route packets on each other's behalf, requiring the use of multi-hop routing among the hosts to deliver packets. Potentially, the "routing" functionality can be implemented at the link layer, such that the entire ad-hoc network appears as a single "link" from the perspective of the rest of the network.
- *Cooperative diversity*: With cooperative diversity, hosts other than the end-points of a "link" cooperate to help deliver packets on the link. For instance, when A transmits a packet to B, a third host C may relay a function of the "packet" received from A to C; host C then can combine the receptions from A and C both to determine the received packet.
- *Network coding*: Network coding allows hosts to relay not just a received packet, but instead a function of packets received from different sources.

Such approaches allow the use of multiple "hops" to deliver packets over what may be construed to be a single "link". Present link layer abstractions may not be sufficient for these more general notions of a link.

Adaptation mechanisms: Capacity limitations of wireless networks have motivated a large range of adaptation mechanisms including power or bit rate control, channel adaptation, and directional beamforming. While such adaptations have the potential for improving performance, they also complicate the notion of a network "link". For instance, given that a host can transmit at two different power levels, should a "broadcast" be performed at the lower power level, or the higher power level? Should the broadcast be performed on any one available channel, or all available channels? In general, the answer to such questions will have to depend on the higher layer's requirements, making the adaptations necessarily cross-layer. Suitable interfaces need to be developed to allow for such cross-layer interactions. Similarly, given that the physical layer behavior is defined by many parameters (e.g., rate, power, channel, beam angle, etc.), the network layer may wish to specify the parameters (or parameter range) to be used for transmitting each packet, based on network level information (as opposed to local information available to each host).

A brute force approach to deal with the above complexities would be make the network transparent, i.e. not supportive of such optimizations or cooperative mechanisms. Confining these advanced techniques to local scopes, however, would result in performance degradation. When the available capacity is limited, such degradation is not likely to be acceptable. On the other hand, if other approaches dramatically increase the capacity of wireless networks (e.g., using larger amounts of spectrum), then it may become desirable to trade-off some performance to simplify the protocol design. In the meantime, however, we suggest that approaches be developed to facilitate cross-layer cooperation, particularly at the lower layers of the protocol stack, in a manner that is "future-proof". Towards this end, the protocol stack needs to be "flexible" to be able to accommodate future developments in wireless networking and communication.

4.13 Support for cognitive radio technology (Scenarios A, B)

The combination of waveform-agile radios and diversity of emerging waveforms will likely result in a wireless devices being able to use a multitude of radio waveforms as well as allow for agility between different radio communication protocols. As a result, wireless devices will be able to create any kind of communication link they want, with whatever combination of capacity, error-rate, etc. they need (subject to what is physically achievable, of course). These links will become an integral part of the next generation Internet, especially as many wireless network architectures, such as mesh networks, are being proposed as alternative backbone infrastructures.

Architecturally, waveform agility enables the concept of a *definable link*, as opposed to the conventional notion of simple, fixed links. In other words, a link is no longer just an *input* to the topology and the algorithms that operate on the topology (e.g. OSPF). It is a *variable*, a parameter that can be controlled as necessary by the protocols. The next generation architecture should not only tolerate the highly dynamic nature of links, but should also be able to tailor the link characteristics in a cross-layer fashion to create the kind of topology that an application requires. *Topology control*, i.e. control of *which* nodes should be able to communicate as well as *how* they should communicate, should be an integral part of the architecture. This will allow control mechanisms at several layers of the stack to marshal resources as they desire.

An architectural proposition for the accommodation and exploitation of dynamic and definable link characteristics is as follows. A cross-layer module called *Dynamic Topology Control* (DTC) is present in an infrastructure node. The DTC module functions as a "broker" between link advertisements (what is possible) and link/topology requirements (what is needed). Link advertisements may be expressed in declarative form to capture the possibilities as well as constraints of resource availability[1]. Research is

[1] For example, consider a wireless router with two software radios, each capable of switching between three waveforms, for instance, 802.11, UWB and MIMO. Then, a link advertisement could be (802.11 AND UWB) OR (802.11 AND MIMO) OR (UWB AND MIMO). Other equivalent representations are clearly possible.

needed in languages for the expression and logical reasoning of resource availability, usage and constraints.

Research is also needed in how conventional control mechanisms, such as routing and transport, can utilize DTC appropriately. A wide range of combinations of link characteristics are possible, and a routing protocol needs to decide which to track/disseminate and which to not (otherwise it may not be scalable). Similarly, a question is whether a transport protocol, looking at the needs of a session can dynamically determine the best topology and get the links created as desired. For example, if capacity is important but latency is not, a topology with a larger number of short-range high-data-rate hops may be used whereas if latency is important, a waveform that has a single wireless hop may be best (e.g. a spectrum agile radio may be made to switch to a lower frequency and use a higher processing gain).

4.14 Integrating lightweight sensor protocols (Scenario C)

The fundamental building block of the Internet's data plane has been the globally addressed IP packet, and the TCP connection. As sensor networks are often comprised of numerous small devices to facilitate physical coupling, the one-to-one communication model will be surpassed by group communication primitives. Sensor devices are also small and battery powered, which makes supporting TCP/IP challenging and inefficient. As a result of these constraints as well as the driving applications, these nodes are usually deployed as clusters whose actions and processing is coordinated through a micro-server. The simplicity of the IP narrow waist leads to ease of interoperability, but edge sensor networks, which may mostly communicate locally, need the flexibility to use lightweight network protocols suited to the sensor domain.

One approach is to connect local sensor networks to the Internet through gateways. The gateway can be a central point of coordination for the sensor patch, combining externally requested communication and data acquisition into power-efficient and optimized requests in terms of sensor network protocols. Querying or collecting data from these networks typically involves naming well above the level of a single node. Therefore, the entire network could be presented under a single Internet name or address, which refers to the gateway. In this case, the new internet architecture would not govern the specific sensor network protocols used, but rather it would merely acknowledge that disparity in protocols exists. However, network management may also require naming particular nodes within the network. Being able to name the sensor nodes in the IP namespace would allow systems to take advantage of the existing infrastructure of monitoring tools. More research is needed to evaluate such a gateway-based architecture compared to defining new end-to-end narrow-waist protocols for the network.

4.15 Programming model for in-network services (Scenario C)

The fundamental application model of the Internet's data plane has been based on the end-to-end principle of data processing only at the edge of the network. This has provided important benefits for both application and network design in the current Internet. However, in-network processing may become necessary in sensor networks and pervasive systems because of scalability and response time considerations. Also, efficient data aggregation in sensor networks requires storage and processing at network nodes. A new programming model is needed to define the data processing behavior of the network.

The architecture and programming model must allow pushing filters and functions into different network elements. Network elements that must be programmable range from tiny sensor motes at the edge, micro-servers/gateways, and powerful sensorweb servers deep in the network. Complex applications will be based on statistical computing with multi-scale inputs, combining fine and coarse spatial and temporal scale data, and will involve relatively complex triggers and functions that make the collection and processing of data adaptive to the dynamic physical world.

4.16 Dynamic discovery of sensor devices (Scenario C)

The current Internet is text-dominated with relatively efficient search engines for discovering textual resources, while connecting to the few available sensors and actuators (e.g., cameras and printers) requires manual configuration. Traditional Internet applications conceptually support data served from a smaller number of servers to a larger number of data consumers. As we embed the Internet in the physical world, we turn this inside out and have larger numbers of data sources sending information to a smaller number of data consumers. An Internet dominated by unstructured information supplied from large numbers of sensor devices must support efficient mechanisms for discovering available sensor resources.

The new architecture must support methods for the registering of a new sensor system in the broader network. This is especially important for multi-purpose sensor networks that will likely be deployed. Since every sensor network is restricted to producing specific types of data, the network needs provide methods for applications to query the available data types, and the geographic coverage area of the network.

4.17 Data Integrity (Scenario C)

The current Internet merely provides a conduit for sensor data streams. Since future network traffic will likely be dominated by information from billions of sensors, more attention to data integrity challenges is warranted. In particular, physically embedded sensors are difficult to keep calibrated. They are inherently subject to noise and dynamics from the physical world, that while the transducer itself might maintain its factory calibration, the integrity of the data is still at risk due to changing physical coupling, or interference, for example.

A new architecture should provide data cleansing mechanisms that prevent corrupted data from propagating through the sensor network. In particular, services that maintain device calibration and monitor/detect adversarial manipulation of sensor devices should be integrated into sensor networks. This could be realized through obtaining context information, metadata, and statistical techniques to locally detect faulty inputs. Although most data cleansing might occur within the sensor network, the interface between the sensor network and the Internet should be able to convey measures of the reliability of sensor data. Since messages originating from the network might instruct the sensor system to activate more devices, or increase resource consumption, it is necessary that messages from the network can be authenticated and verified by the sensor network. When sensor networks are connected through a gateway, it is essential to ensure the trustworthiness of the interface between the sensor system and the broader network. Sensor network gateways should be authenticated in order to prevent impostors from reporting fallacious data.

4.18 Low latency interaction (Scenario C)

Sensor networks make sensor-actuator systems that react to physical inputs in real time the norm rather than the exception. The only way to richly observe most physical phenomena is with multiple modalities of observation. This implies multiple sensor types, multiple perspectives, and a tier-ing of devices from the smallest, most widely distributable to the mobile or manual intermittently available. This means that low latency communication among physically local devices is particularly important. As applications begin to look for correlations across large distances, low latency triggering across wider areas might also grow in importance.

The architecture could support low-latency requirements through a multi-tier architecture, which allows for direct communication between local sensors and actuators. If protocol translations are necessary these could be handled by a micro-server in the local network, rather than a distant gateway.

5 Summary of Network Architecture Challenges for Wireless

From the discussion of various scenarios and their requirements in Sec 4, it is clear that emerging wireless technologies drive a variety of new requirements for the Internet. Taken together, these new architecture and protocol design considerations represent a major research, design and implementation challenge for the networking research communities, both wired and wireless. An important design issue is that of distilling critical new features from the broad set of requirements discussed and then mapping them to practical network implementations. Since many of the wireless/mobile/sensor network protocol capabilities under consideration apply only to edge networks, it may be beneficial to adopt hierarchical structures for the future Internet in which edge networks with high functionality levels can be aggregated into more streamlined core networks optimized for high speed wide-area transport. However, it will be important to design the aggregation interfaces to pass necessary control information from the wireless edge to servers or other wireless edge networks so that end-to-end communication can be suitably optimized.

The following is a summary of major architectural and protocol design considerations for wireless, mobile and sensor networks:

5.1 Naming and Addressing Flexibility

Today's Internet addressing scheme is rather rigid; it is well suited to a static, hierarchical topology structure. It provides a very efficient way to label (and find) each device interface in this hierarchy. To support mobility, the next generation Internet must provide ways to name and route to a much richer set of network elements than just attachment points. It must support routing in terms of names, which identify the actual desired end point, rather than some particular characteristic, such as a physical address. These names may well resolve to addresses at some layer in the architecture, for routing purposes, what is important is that the architecture supports seamless changes in name-address binding. A clean architectural separation between name and routable address is a critical requirement.

5.2 Mobility support

Mobility is the most fundamental characteristic of a wireless network with mobile users, and it is therefore imperative that future networks provide mobility support as an integrated, first-class service.

Mobility scenarios anticipated in future networks include simple end-user migration from one subnetwork to another (as in cellular or WLAN hot-spot services), as well as more complex mobility patterns involving movement of radio routers and ad-hoc network clusters. Protocol features are also needed to address a new class of network services that take advantage user or platform mobility – these include location-based services, the "Infostation" model and the "data-bus" service in which data is cached and delivered when an appropriate receiver is within range. Mechanisms are needed for classical location management and low-latency handoff within or between wireless networks, as might occur as a device changes access points or moves between ad-hoc networks. Mobile nodes, constellations of ad-hoc routers and more general entities must be reachable as they move across different wireless domains. Scalable indirection schemes (more efficient and general than existing mobile IP) need to be designed to allow for this functionality. Infostations and data-bus services involve mobility prediction, pre-fetching and data caching, and it is of interest to evaluate what level of support should be provided by the network layer. Mobility can also induce temporary network disconnections. The network layer should support disconnected, delay-tolerant operations, a capability that is current not provided within IP.

5.3 Location Services

Another key feature needed for wireless networks is that of a location service, which would provide information about the location of a packet's source or destination. Beyond location as an addressing mechanism the architecture must enable the use of location information for internal network optimizations at all layers and support the development of location-aware applications. Specific point solutions such as geographic routing, mobicast, localization, or location fusion algorithms have been developed and these ideas may be expected to mature further during the next decade. A general solution for integrating location into the next-generation Internet architecture is an important research direction. Protocols for location service should provide the means for communicating entities to indicate their current position and to discover other devices in a specified geographic area. Design considerations include the granularity and type of representation for location information, which can vary from conventional descriptions like postal address or room number to geographic coordinate systems.

5.4 Self-Organization and Discovery

Emerging wireless system architectures involve ad-hoc network formation based on opportunistic identification of resources in the environment. As a result, protocols used should support discovery of neighboring radios and the topology of existing networks. Self-organizing ad-hoc networks have been built for homogeneous radio environments, but future networks will require more general capabilities for organizing a mix of wired and wireless components [Ga04]. For example, a local network may consist of several ad-hoc wireless clusters interconnected by a set of wired routers. The protocols developed need to incorporate distributed algorithms and measurements for efficient formation of desirable network topologies in an environment with dynamic mobility. Self organization should work in a distributed environment without requiring continuous connectivity to the global network.

5.5 Security and Privacy

Wireless networks can be expected to be the platform of choice for launching a variety of attacks targeting the new Internet. At the most basic level, wireless devices will likely have evolving naming and addressing schemes and it will be necessary to ensure that the names and addresses that are used are verifiable and authenticated. There are also a variety of complex security issues related to ad-hoc networks of peers, including management of authentication and trust in purely distributed environments. In addition, wireless networks are subject to a variety of denial-of-service or man-in-the-middle attacks and future approaches to Internet security will need to take this into consideration.

Further, one parameter uniquely associated with wireless networks is the notion of location. Location information provided by the network should be trustworthy. The wireless component of the network should provide both privacy and forensic capabilities, while allowing for the means to manage the tradeoffs between these two complementary considerations. Since the wireless medium allows for adversaries to easily sniff network traffic, it will be essential to provide mechanisms that prevent traffic analysis by adversaries. Additionally the architecture should provision hooks for future extensions to accommodate legal regulations.

5.6 Decentralized management

From an edge network dominated by the wired Ethernet, the edge is evolving into disparate wireless ecosystems such as cellular, 802.11, Bluetooth, RFID, and sensor networks. Managing such disparate dense networks requires augmenting the existing protocol stack for providing efficient remote manageability, diagnosis, configurability and trust. Wireless devices are not always connected to the Internet core. In a next generation Internet, these devices need to be able to form ad-hoc networks that operate with the same services and expectations as the larger Internet, but with limited connectivity. Ad-hoc constellation management involves not only knowing the status of device membership in each constellation but also the nature of interaction with others.

5.7 Cross-layer protocol support

As a mobile client crosses different wireless domains, its route characteristics continuously change. Awareness of these dynamic changes (at client, server or both) can be critically useful when supporting applications that require some quality of service, such as multimedia delivery. Important path characteristics include available capacity, loss rate, delay, hop distance, energy, stability, and security. Exposing this information would allow connection end points to make more intelligent routing, association, and protocol decisions without requiring complex probing or estimation techniques. There an obvious trade-offs between layering and complexity, so that future protocol designs need to find the right balance after evaluation of actual performance gains in realistic environments.

5.8 Sensor Network Integration

Efficient integration of sensor networks with the global Internet involves several additional requirements. These include the ability to interface with a lightweight sensor network protocol, possibly through a unified hierarchical protocol framework. In addition, sensor networks require the self-organization discovery and location services discussed in 5.3, 5.4. Other capabilities needed for sensor scenarios are in-network programming, content awareness, data aggregation and data integrity mechanisms. An attribute or location based dynamic binding service is also important for development of real-time sensor applications involving opportunistic associations. In general, sensor networks require identification of suitable new socket layer abstractions and programming models that differ significantly from today's TCP/IP paradigm.

5.9 Cognitive Radio Networks

Cognitive or software-defined radios are expected to be an important technology driver for wireless networks of the future. Cognitive radios will enable wireless devices to flexibly create many different kinds of communication links depending on required performance and spectrum/interference constraints. These links will become an integral part of the next generation Internet, especially as many wireless network architectures, such as mesh networks, are being proposed as alternative backbone infrastructures. Architecturally, waveform agility enables the concept of a *definable link*, as opposed to the conventional notion of simple, fixed links. In other words, a link is no longer just an *input* to the topology and the algorithms that operate on the topology (e.g. OSPF). It is a *variable*, a parameter that can be controlled as necessary by the protocols. *Topology control*, i.e. control of *which* nodes should be able to communicate as well as *how* they should communicate, should also be an integral part of the architecture.

5.10 Economic Incentives

Incentive policies are required to make opportunistic networking work using third party nodes as store and forwarders. In a world of “constellations” and dynamically changing connectivity in urban and “mall” areas, it will be beneficial to exploit neighbors’ resources for a number of services (data forwarding, data ferrying, information stations, etc). However, by helping the neighbor incurs a cost in terms of energy consumption, increased load, throughput reduction and security risk. There should be some form of incentive ranging from credits to micro payments to facilitate this type of cooperation. Similar incentives may be used to encourage cooperative sharing of radio spectrum.

6 Deployment Strategies for Future Wireless Networks

A number of new mobile/wireless/sensor service requirements and protocol features have been discussed in Sec 5. This list of potential enhancements to the network protocol is just a first step – each of these proposals and their specific implementation needs to be analyzed in terms of anticipated cost/benefit and protocol performance. Those items which survive careful scrutiny may be appropriate for further implementation in prototype wireless/sensor edge networks that connect to the existing or future Internet for end-to-end experiments. Several deployment options may be considered, taking into account the fact that the global Internet is inherently hierarchical in terms of access (edge), regional and national/global networks with protocols to support aggregation of routing information and traffic at the boundaries. The first option (shown in Fig. 2 below) for deployment with the least impact on the regional and national Internet is to develop an “IP+w” or “IP+s” specification customized for wireless or sensor access networks, with a suitable “border router” at the interface to the legacy Internet. This approach (which has been adopted in the past for mobile IP and for ad hoc MANET specifications in the IETF) has certain advantages with respect to legacy networks, but suffers from potential gateway processing bottlenecks and difficulties with end-to-end services involving existing servers, etc.

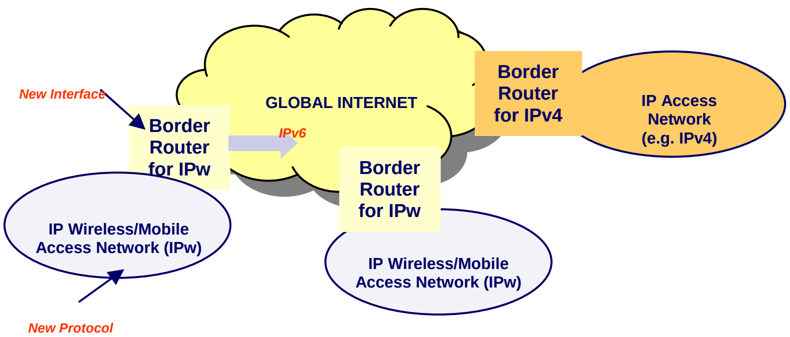

Fig. 2. Deployment of future wireless-specific IP access network with legacy backbone

The second deployment approach under consideration is the overlay network [PLA, Pet04], shown in Fig. 3 below. In this approach, new wireless/mobile access networks are interconnected by a global overlay network with wireless-specific features such as location services and attribute-based address resolution. The overlay network uses IP tunnels for basic transport and can thus reuse the existing Internet as its underlying infrastructure. The overlay approach has many advantages including the ability to rapidly roll out fundamentally new network services with modest investment and the fact that each network can be customized to meet a specific application need. The main drawbacks are the increased packet overhead, processing complexity and latency from the addition of a new networking layer, and the fact that multiple overlays may lead down the path of fragmenting networks by application category, thus reducing "network effect" benefits associated with a single Internet API. The overlay approach is quite feasible for many of the new mobile or sensor related services discussed in Sec 5, and may be the most pragmatic way to experiment with significantly new functionality in regional and core networks. It is noted that implementation of some end-to-end service features (e.g. cross-layer transport) may require the existence of a separate global control and management plane for low-level link and node status information.

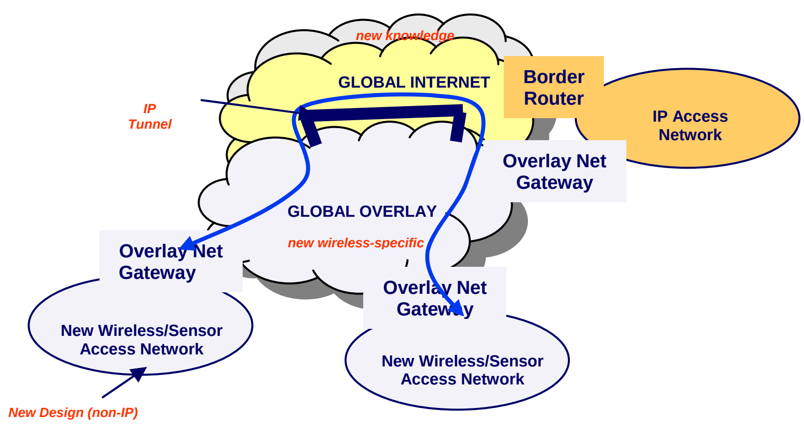

Fig. 3. Deployment of new wireless/sensor access network with overlay global backbone

The third approach to deployment (shown in Fig. 4) is based on a complete replacement of both the wireless/sensor access network and the global Internet with a single new end-to-end architecture. This approach has the highest cost and risk, but also has the greatest potential for an optimized, lightweight networking solution that is responsive to emerging needs. As with earlier attempts to replace large systems with new ones, this approach has the risk of being marginalized by Internet evolution and legacy staying power. A key challenge in building experimental systems of this sort is the design of programmable network elements (both wired and wireless) that are not constrained to any particular packet format or addressing structure.

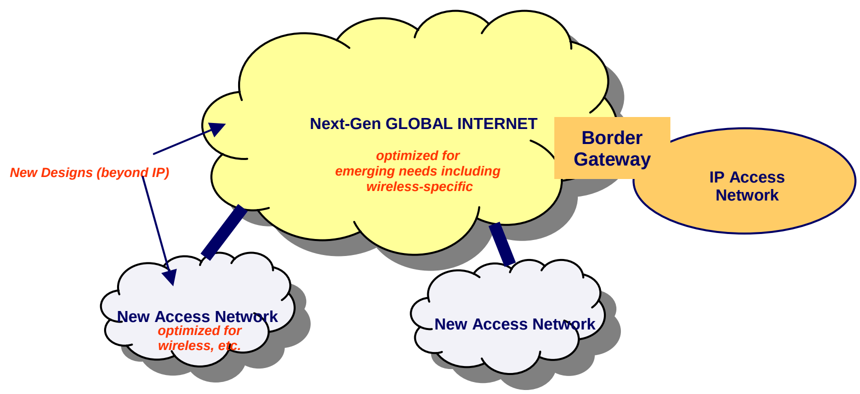

Fig. 4. Revolutionary deployment of new wireless/sensor access network and global Internet backbone

7 Experimental Infrastructure for Wireless Network Research:

One of the breakout panels at the WMPG workshop focused on identifying experimental infrastructure necessary to support research on the future Internet. The proposed experimental infrastructure should be designed to incorporate a wide range of wireless networking capabilities in order to provide experimenters with access to emerging radio technologies and sensor applications, which are becoming increasingly important at the Internet edge. In addition, such wireless/sensor testbeds should be integrated with a flexible wide-area wired network infrastructure that can be used to study architecture and protocol ideas in an end-to-end sense. The design objective for wireless networks mirrors that of the wired core network, i.e. the ability to support experimental flexibility with new protocol concepts in a variety of controlled and real-world scenarios. The concept of a large-scale experimental network with innovative programmability and virtualization features described in a related NSF study entitled "Overcoming Barriers to Disruptive Innovation in Networking" [Pet05] is considered appropriate for this purpose.

Network programmability and virtualization are two important features to be considered in the design of experimental wireless networks under consideration. Programmability refers to the experimenter's ability to program various radio nodes and network elements to run protocols and software supplied by the user. The user should also be able to select run-time, measurement parameters and network topology in a flexible manner. Virtualization refers to the ability of the system to experimental network to simultaneously support more than one "virtual network" or protocol implementation. Virtualization in wired networks has been demonstrated in PlanetLab [PLA], generally involving the concept of partitioning of node and link resources ("slices") in a non-overlapping manner between multiple virtual networks. Similar concepts can be applied to wireless networks taking care to account for shared-media interactions between wireless links and nodes. The degree of virtualization that can be practicably achieved depends on the hardware platforms used and the specific wireless networking scenario under consideration.

7.1 Techniques for Flexible Experimental Wireless Networks

Wireless technology has progressed significantly during the last decade, both in terms of quantitative Moore's law improvements (from ~10 Kbps to 100 Mbps and higher) and qualitative capabilities (compact sensors, ad-hoc radios, etc.). Flexible and programmable components are essential for building future networks with yet-to-be-determined protocols and changing wireless technology. In this section, we describe some examples of emerging wireless technology components that can be used to build flexible experimental networks.

7.1.1 Open API Radios

A general-purpose experimental wireless network should be based on the concept of a single radio API that provides uniform services across different radio technologies. The API supports both data transfer (packets) and a control interface for transfer of radio control and measurement (see Fig. 5, which also shows an example implementation of an open API radio card that provides this type of capability), or may be multiplexed over a single base protocol in other cases.

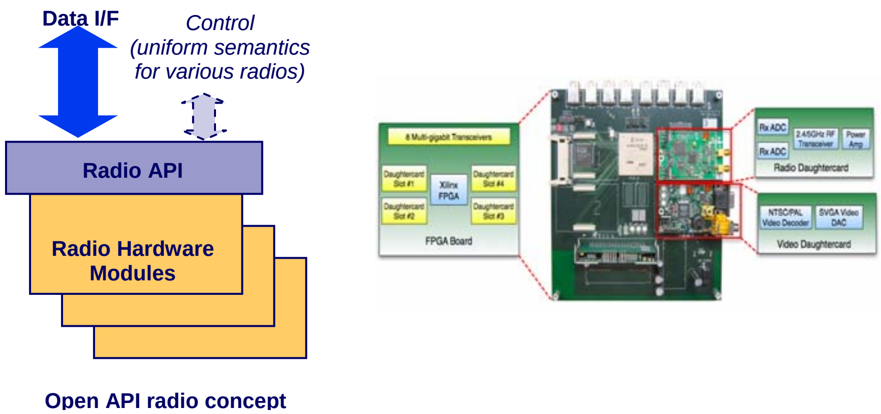

Figure 5 Open API Radio and Rice University Prototype

A general-purpose experimental wireless network may be expected to include a range of short range radios (such as Bluetooth, Zigbee, Mote), medium range radios (such as 802.11a,b,g or 802.11n) and long-range radios (such as WCDMA, CDMA2000, WiMax and new OFDM/MIMO alternatives). The panel believes that a critical task for the experimental wireless research community is the development of open API implementations for a representative set of these radios including 802.11x, UMTS and WiMax. Cooperative agreements with chip vendors along with new firmware/driver development work will be required in order to achieve this objective.

7.1.2 Cognitive Radio

Software radios are quickly emerging as the platform of choice for future wireless communications systems, from commercial [VAN] to open source [GNU] to military [JOI]. In a software radio, many signal processing functions such as modulation, coding, and spreading are done in software. This enables agile radios - a software radio that uses its flexibility to dynamically change waveform characteristics and behavior in response to instruction and even cognitive radios - radios that know about their capabilities, internal state and the spectrum environment to opportunistically communicate (without explicit instruction) using the best combination of transceiver parameters [Mi99]. The trend is clearly toward an increasing flexibility, agility and autonomy in the selection of wireless communication waveforms.

The idea of a cognitive radio extends the concepts of a hardware radio and a software defined radio (SDR) from a simple, single function device to a radio that senses and reacts to its operating environment. A cognitive radio incorporates multiple sources of information, determines its current operating settings, and collaborates with other cognitive radios in a wireless network. Some features of cognitive radio networks include: (1) spectrum sensing and scanning for interference analysis and to find spectrum gaps; (2) policy and configuration databases to determine radio capabilities and regional rules of use; (3) adaptive algorithms capable of reacting to measured radio signals, adhering to policy constraints and negotiating with peers to best utilize spectrum and meet user demands; and (4) distributed control by which radios can communicate with each other to exchange control information, and where appropriate, form collaborative peer networks to avoid destructive interference [Ray05].

Cognitive Radio Platforms: The cognitive radio outlined above requires new hardware platforms and related software capable of implementing flexible operation at radio physical, data link and MAC layers. Cognitive radio platforms for research are currently being developed at several organizations including Utah [GNU], University of Kansas [Mi05] and Rutgers University [Ac04], and these boards are expected to become available for early experimental work in 2006-07 (see Fig 6). The panel believes that cognitive radios represent an important component for the future Internet in view of their flexibility and performance, and should therefore be included in a large-scale experimental network implementation plan.

GNU Software Radio

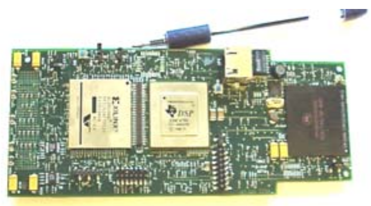

Lucent/WINLAB Cognitive Radio

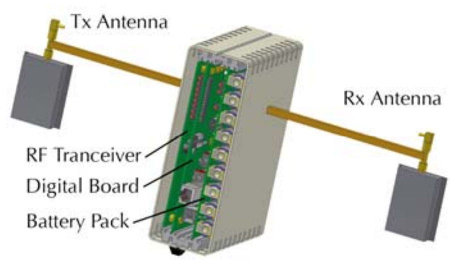

KU Agile Radio

Figure 6 Examples of cognitive/SDR platforms under development with NSF research funding

7.1.3 Virtualization of Wireless MAC

A virtual network, such as PlanetLab [PLA], uses single hardware instances and virtualizes the single hardware instance to allow multiple applications or simulated networks to use the hardware. Unlike traditional virtual networks, the problems are more significant in wireless networks because the performance and characteristics of wireless networks are greatly influenced by their MAC and PHY layers. Recent research indicates that it is feasible to build inexpensive, virtualized wireless interface that can also be used for efficient overlay networks for experiments at the MAC layer; this could be extended to virtualized physical layers but at considerable additional cost and effort. The structure of the proposed *MultiMAC* approach is shown in the following diagram.

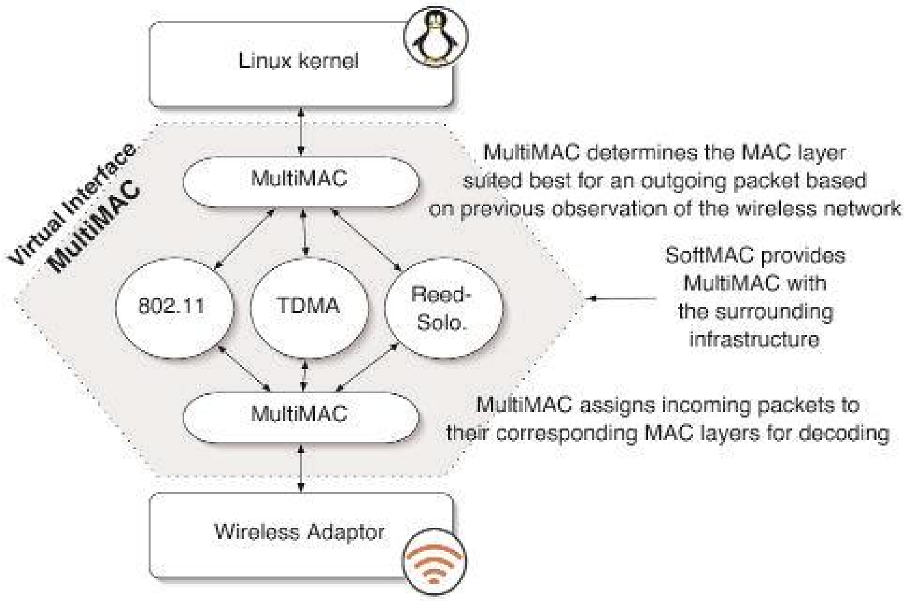

Fig 7. Architecture for virtualization of wireless MAC layer

The software is based on a software layer, SoftMAC, which provides software control over a specific wireless network hardware interface made by the Atheros Corp. The basic SoftMAC layer can be extended to create MultiMAC. As illustrated in the Fig 7, MultiMAC allows the concurrent use of multiple SoftMAC implementations. Packets are received using a common PHY layer implemented by the underlying radio and are then presented to the different MAC layers. The resulting MultiMAC implementation can be used to build both virtualized and overlay wireless networks at the MAC layer, and to do so using commodity hardware, although on a limited set of computing platforms (the source for the wireless driver is needed). This means that virtualization can be used both in testbeds and for operational network deployments. Likewise, rather than resorting to ad-hoc wireless overlays to effect new network functionality, such as in the paper [Ra05], it is possible to directly implement the desired functionality.

7.1.4 Wireless Network Monitoring and Measurement

Wireless networks introduce numerous challenges that make measurement significantly more complex than that of wired networks. Some of these factors include the time- and space-variability of the wireless channel, the need to simultaneously measure and correlate physical-layer, link-layer and network-layer performance, the difficulty of establishing universal metrics, and the fact that available wireless equipment or networks used may not provide access to necessary parameters. Typical radio measurements of interest include parameters such as RSSI (radio signal strength indicator), link quality (packet error rate), MAC congestion indicator, link throughputs and overall node throughput. From an implementation standpoint, wireless network measurements create significant overhead which can affect the outcome of the experiment being performed if the data is carried over the wireless medium. This implies the need for either a wired backhaul from each experimental radio node or an out-of-band wireless measurement infrastructure based on independent radio sensors. For example, in the ORBIT radio grid the measurement library gathers data for each radio node using a Gigabit Ethernet wired connection [Si05]. In another experimental system DIMSUMnet [Bu05], field data is collected by an independent array of spectrum sensors operating in a different frequency band. A second approach to wireless measurements in dense deployments is to leverage self-reporting by mobile devices in the field. As the number of mobile devices grows, this technique is expected to provide reasonably good measurements across space and time. An approach which has been used with some success in 802.11 networks with Prism radios is that of piggybacking measurement information along with data packets by modifying drivers associated with commercial wireless cards [Ye04].

Wireless measurements are useful not only from a network management perspective, but may also be made available to end-users to facilitate service selection in a heterogeneous environment. For example, the concept of a "spectrum server" has been proposed to facilitate shared usage of unlicensed band spectrum [Ma04]. In this approach, an Internet-based spectrum server gathers radio signal strength and location measurements and makes this available for possible coordination among users. Another method under investigation is the "wireless network repository" which stores service availability and quality information at various levels of spatial and temporal granularity. Users of the system may use this information to select among service providers or to optimize application performance.

7.1.5 Emulation and Simulation Testbeds for Wireless

Research on wireless networks faces a unique challenge of conducting efficient, repeatable experiments which are reasonably close to real world environments. Extensive simulation methods have been developed by the research community over the past 10-15 years, including network simulators, such as ns-2, and these can be used for early stage research. For more realistic experimentation, emulators have been proposed as a means for evaluating complete network protocol stacks with real platforms, while avoiding the complexities of physical world deployment. Several wireless network emulation testbeds are currently under development under the NSF Network Research Testbeds (NRT) project initiated in 2003. These include the WhyNet hybrid simulator/emulator at UCLA [WHY], the Emulab emulation testbed at Utah [EMU] and the ORBIT radio grid testbed at Rutgers [ORB, Ray05]. Each of these systems provides useful capabilities for experimental research on future wireless networks. In particular, WhyNet provides the ability to integrate network simulation with a variety of hardware PHY's, while Emulab and ORBIT support complex network emulation in real-time. Emulab has extensive wired + wireless facilities, while ORBIT is oriented towards large-scale ad-hoc, pervasive system and sensor net scenarios with up to 400 radio nodes.

7.2 Wireless Networking Platforms

A number of wireless platforms ranging from sensors to radio routers are required to build a wireless network. Design objectives for wireless platforms include:

- Designs should support multiple radios technologies (as "plug in"), provide open API's to layers 1-3 (radio, MAC and network), and have a common open-source software framework.
- Network nodes should support software download, multiple protocols and interoperate with virtualized network nodes used in core infrastructure.
- Network nodes should support remote monitoring and rebooting features – important for network operation and management of devices to be used in the field.

To simplify development, experimental systems for wireless network research should use a modular building block approach with hardware components that include commercial mobile platforms, embedded CPU boards, switching hardware, plug-in radio modules and remote monitoring/control modules. All these components (with the possible exception of certain embedded sensor devices) share a common software model based on open-source Linux with unified radio API's.

The following wireless platforms may be considered for experimental wireless networks:

1. Embedded wireless sensors (small, low-power devices with integrated computing, communication and sensing) with built-in radios such as Mote [CRO], Zigbee (see Fig 8.). These devices will generally be used with a lightweight operating system such as TinyOS, and as a result may have limited flexibility for multiple applications. Virtualization of multiple protocols on a single sensor is not considered practicable due to compute resource limitations; however, different sets of sensors can be used with spatial redundancy for multiple virtual networks.

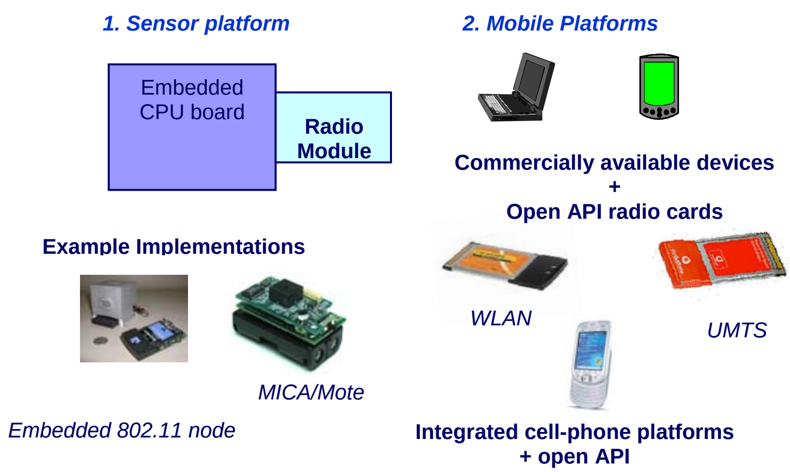

Figure 8 Sensor and Mobile Computing Platforms

2. Mobile computing devices (commercially available laptops, PDA's and cell phones) with external plug-in radios such as 802.11 or UMTS [UNI]/WiMax. A fully integrated cell phone platform will also be used for end-user applications. Note that the "open API" Linux driver for 802.11, 3G, WiMax or other radios represents a critical component for flexible wireless testbeds and larger scale experimental networks (see Fig 8.). Virtualization of these devices is possible (for example, using dual-radio cards and orthogonal frequency assignment), but the resulting resource assignment complexity may be avoided in favor of spatial redundancy, i.e. different devices for each network protocol being studied.

3. Generic radio access point or forwarding nodes with dual network interfaces – see Fig. 9. This is an internal network element with the capability of bridging or routing between multiple radios (for example, IEEE802.11b to 3G/UMTS) or between a radio link and wired network (for example, between 802.11b and Ethernet), and can thus be used as a general-purpose ad-hoc routing node, sensor gateway or wireless access point. The dual radio node used in the ORBIT testbed or the Intel Stargate [STA] board are examples of such platforms. Network virtualization across multiple radio technologies or a single radio technology and multiple frequencies is readily achieved with this platform.

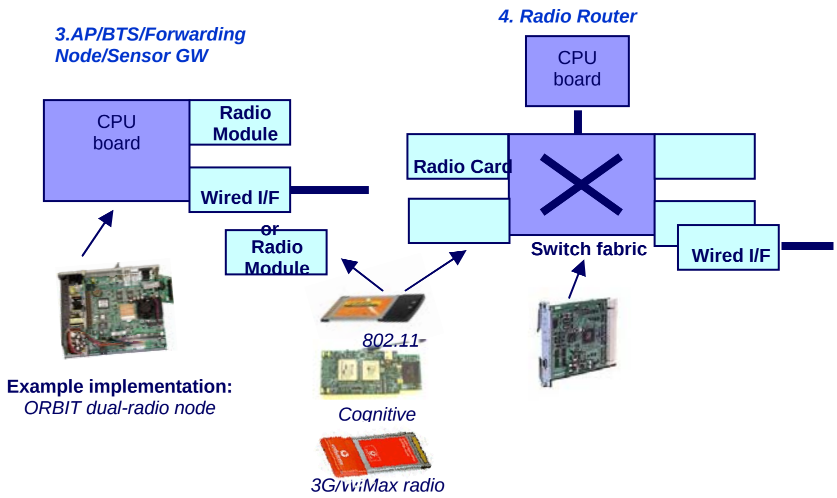

Figure 9 Radio Access Point/Forwarding Node and Radio Router Platforms

4. Radio router with switching fabric and multiple radio and wired ports – see Fig. 8. This is a more general realization of the radio forwarding node described in (3) above, with the capability of supporting multiple radio and wired interfaces using a switching fabric for interconnect and separate CPU board for protocol processing. As for the AP/forwarding node, this platform can support virtual networks using multiple radio technologies or via non-interfering frequency assignments.

7.3 End-to-End Software Architecture

The end-to-end software architecture of an open-access wireless network integrated with a flexible wired network testbed is outlined in Fig. 10. As shown in the figure, there are control, data and management/monitoring functions which are implemented on logically separate "planes" in the network. These planes may use separate virtual networks where available

Each radio node runs a single control program along with one or more experimental protocols for each virtualized network to be studied. Experiment control including code downloading is processed through a "node handler" service in the network. Once it is loaded onto the specified platform, the experimental protocol runs through the same radio API and interfaces with corresponding modules on radio routers, access points or edge routers. Management and data collection is handled by a separate service in the network, and is preferably associated with an independent control interface (via an Ethernet or DSL cable or general-purpose cellular data networks such as GPRS) that can be used for remote diagnostics, rebooting, etc. Further work will be needed on harmonizing the wireless software architecture with experimental wired network infrastructures to ensure end-to-end operation with a single user experiment interface.

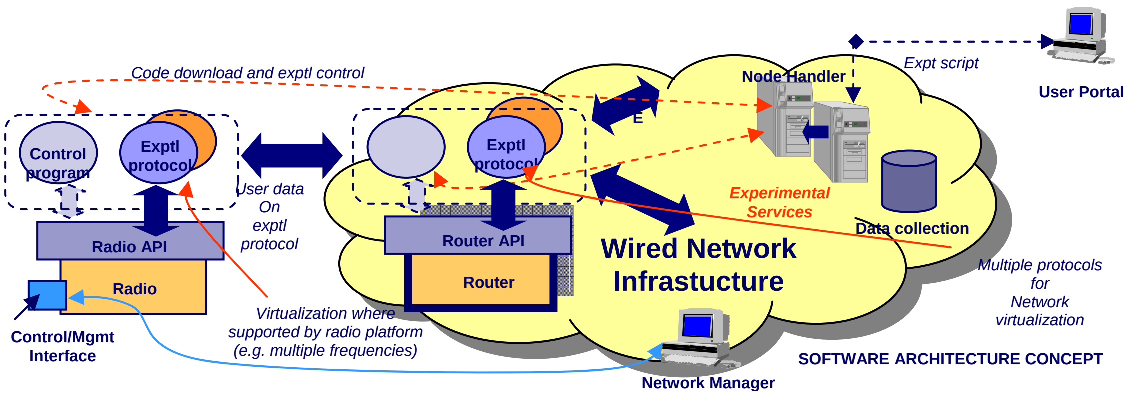

Figure 10 End-to-end software architecture

7.4 Experimental Wireless Networks

Recognizing that many new Internet application concepts involve wireless-end user devices and sensors, the panel recommends research infrastructure that will be equipped to support flexible experimentation in a variety of realistic wireless networking scenarios – see Fig. 11. Specific wireless network scenarios that may be considered for implementation in conjunction with a research infrastructure initiative for the future Internet include:

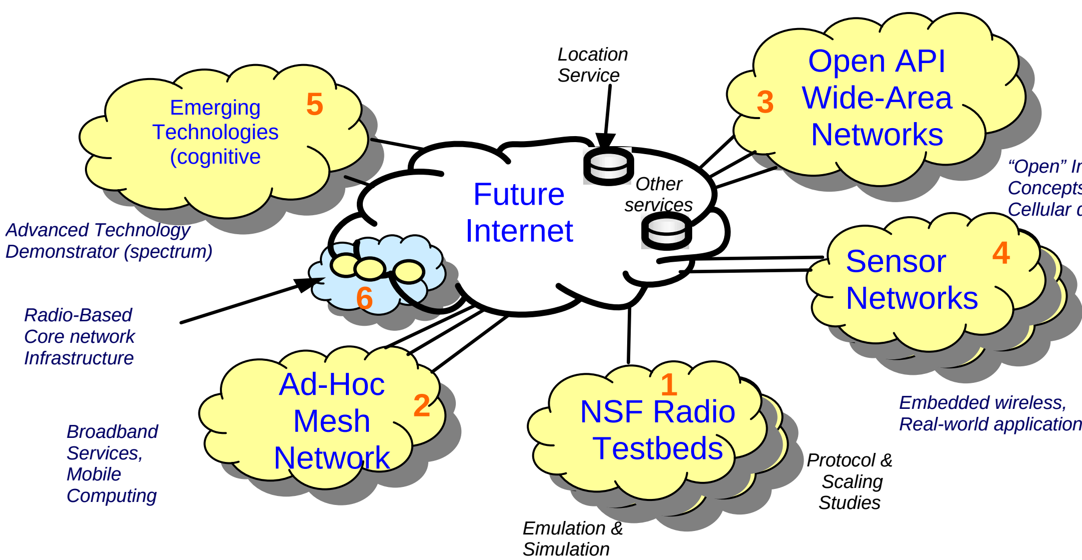

Figure 11 Experimental Wireless Networks for Future Internet Research - Overview

7.4.1 Wireless emulation and simulation

Wireless emulation and simulation environments that provide facilities for repeatable protocol validation studies of a quantitative nature – see Fig. 12. Several large-scale emulation and simulation testbeds have previously been developed with NSF funding and are currently in operation – these include the ORBIT, Emulab and WhyNet testbeds.

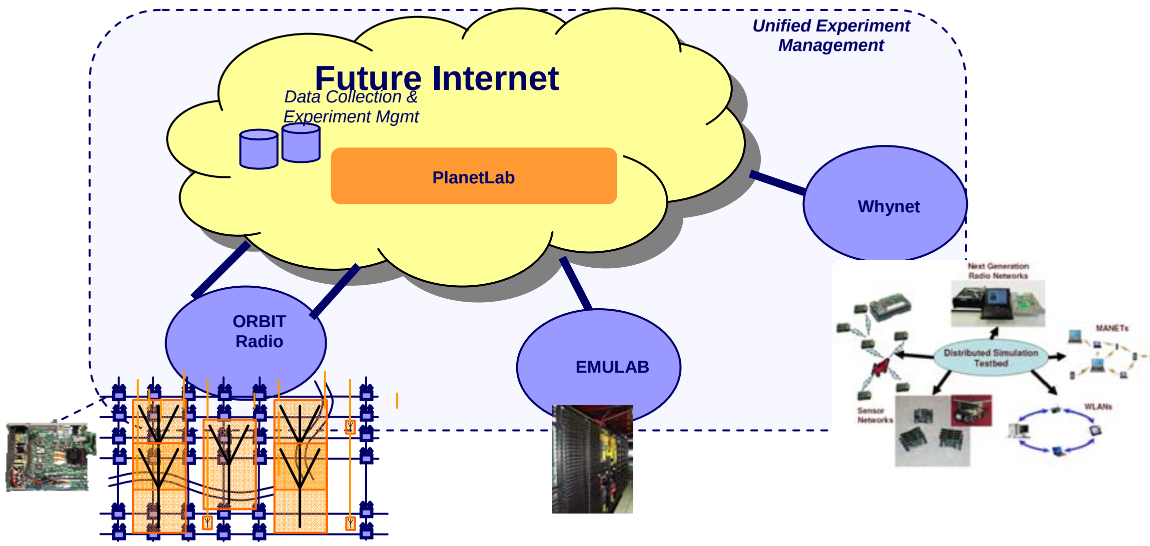

Figure 12 Integration of Existing NSF Wireless Testbeds into Future Internet Research Network

Research topics being studied using these emulators and simulators include ad-hoc network protocols, networks with heterogeneous radio PHY, large-scale sensor networks, wireless security scenarios and mobile multimedia applications. Existing NSF wireless testbeds should be integrated with the proposed wide-area future Internet infrastructure in order to facilitate end-to-end experimentation with new protocols. Integration will involve modification of existing control and user support software in each of these testbeds to a common experimental network usage model. Each of the existing testbeds provides support for experimental flexibility and some level of network virtualization, and these will need to be harmonized to work with the proposed future Internet system.

7.4.2 Urban 802.11-based mesh/ad-hoc network

Urban 802.11-based mesh/ad-hoc network designed to support real-world protocol experience with emerging short-range radios – see Fig. 13. Ad-hoc mesh networks represent an important area of current research and technology development activity, and have the promise of providing lower-cost solutions for broadband access, particularly in medium- and high-density urban areas. While protocols for ad-hoc mesh networks have been maturing, the research community has limited field experience with large-scale systems and application development. Research topics to be addressed using experimental wireless networks of this type include ad-hoc network discovery and self-organization, integration of ad-hoc routing with core network routing, cross-layer protocol implementations, MAC layer enhancements for ad-hoc, supporting broadband media QoS, impact of mobility on ad-hoc network performance and real-world, location-aware application studies. The proposed ad-hoc mesh network would consist of ~1000 open API radio routers/forwarding nodes densely deployed in one or more urban areas or campus settings with coverage area ~10 Sq-Km. A typical node would be installed at ~8m height using available lighting poles or other utilities, and will have electrical power and a remote management interface. Nodes in the network would be programmable in terms of layer 2 and layer 3 protocols and will support applicable forms of virtualization corresponding to specific platforms used.

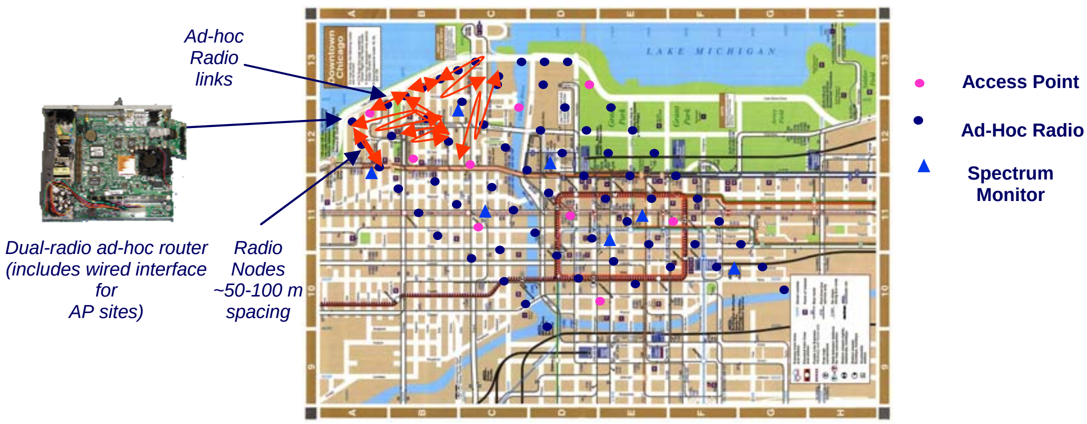

Figure 13 Deployment Concept for Ad-Hoc Mesh Network

7.4.3 Wide-area suburban wireless network

Wide-area suburban wireless network with open-access 3G/WiMax radios for wide-area coverage along with short-range 802.11 radios for hotspot and hybrid service models – see Fig 14. This wireless scenario is of particular importance for the future Internet as cellular phone and data devices are expected to migrate from vertical protocol stacks such as GSM, CDMA and 3G towards an open Internet protocol model. Experimental research on future cellular networks and their integration into the Internet is currently restricted by the lack of open systems that can support new types of protocols and applications. Research topics to be studied using the proposed experimental network include transport-layer protocols for cellular, mobility support in the future Internet, 3G/WLAN handover, multicasting and broadcasting, security in "4G" networks, information caching/media delivery, and location-aware services. A typical experimental network would include ~10 open API 3G or WiMax base station routers along with ~100 802.11 forwarding nodes and access points covering a suburban area of about 50 Sq-km. 3G nodes will need to be mounted at heights of ~30m or higher on buildings or towers, while 802.11 nodes will be at ~8m on utility poles, etc. Both 3G/WiMax and 802.11 nodes in the network will support flexible programming of layer 2 and layer 3 protocols, and will also support applicable forms of virtualization.

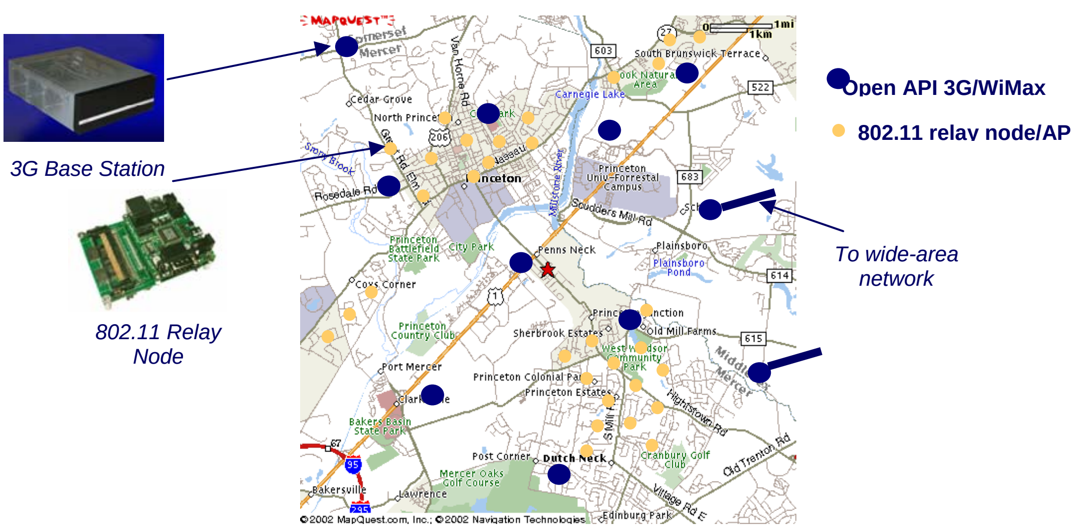

Figure 14 Deployment Concept for Open API Wide-Area Wireless Network

7.4.4 Sensor networks

Sensor networks capable of supporting research on both protocols and applications – see Fig. 15. Since the design of a sensor network tends to be somewhat application specific, we recommend providing necessary wireless infrastructure, leveraging either urban 802.11 mesh networks or wide-area suburban wireless networks listed above, along with a "sensor deployment kit" consisting of network gateways (from sensor radios to 802.11 or cellular), sensor modules and related platform software. Research topics to be studied using experimental sensor net systems include general-purpose sensor network protocol stacks, data aggregation, power efficiency, scaling and hierarchies, information processing, platform hardware/software optimization, real-time, closed-loop sensor control applications, vehicular, smart space and other applications.

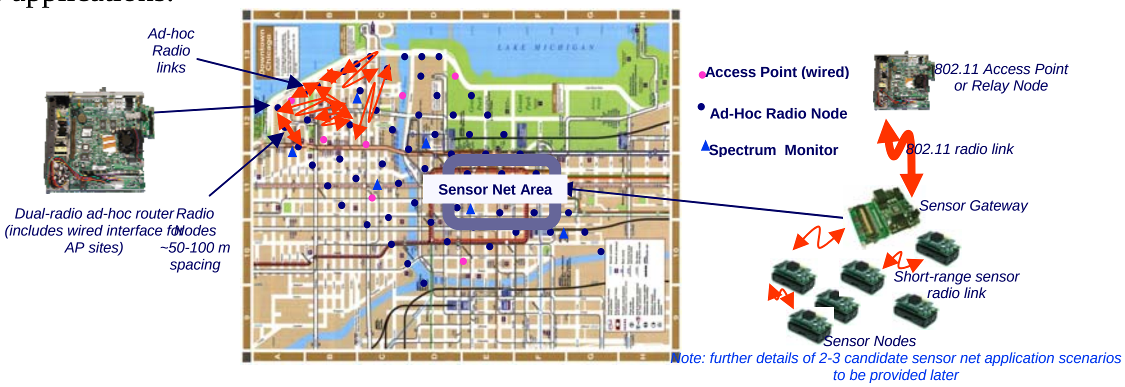

Figure 15 Deployment Concept for Sensor Networks

Specific sensor deployments in areas such as environmental monitoring, security, traffic control, vehicular safety or smart spaces should be solicited through a proposal process leading to selection of 2-3 large-scale sensor net projects. Each sensor net application is expected to involve up to 1000 sensors and 100 network gateways. The low-tier sensor nodes interface with a gateway typically at distances ~10m or less, while gateways would in turn connect to either 802.11 or cellular nodes with commensurate coverage areas. Both sensors and gateways will support flexible programming of protocols used, but the smaller sensor devices may be limited to a single networking protocol at one time due to CPU and other constraints. Sensor gateways will support network virtualization using multiple radio frequencies or via spatial segregation, as discussed earlier.

7.4.5 Cognitive radio network

Cognitive radio networks intended as an advanced technology demonstrator with focus on building adaptive, spectrum-efficient systems with emerging programmable radios – see Fig. 16. The emerging cognitive radio scenario is of current interest to both policy makers and technologists because of the potential for order-of-magnitude gains in spectral efficiency and network performance. NSF and industry funded R&D projects aimed at developing cognitive radio platforms are currently in progress and are expected to lead to equipment that can be used for field-deployable experimental networks in the 2007-08 time-frame. Protocol research to be supported with the planned experimental system includes discovery and self-organization, cross-layer protocols for PHY adaptation, cooperation and competition mechanisms, spectrum etiquette procedures, and cognitive radio hardware/software performance optimization.

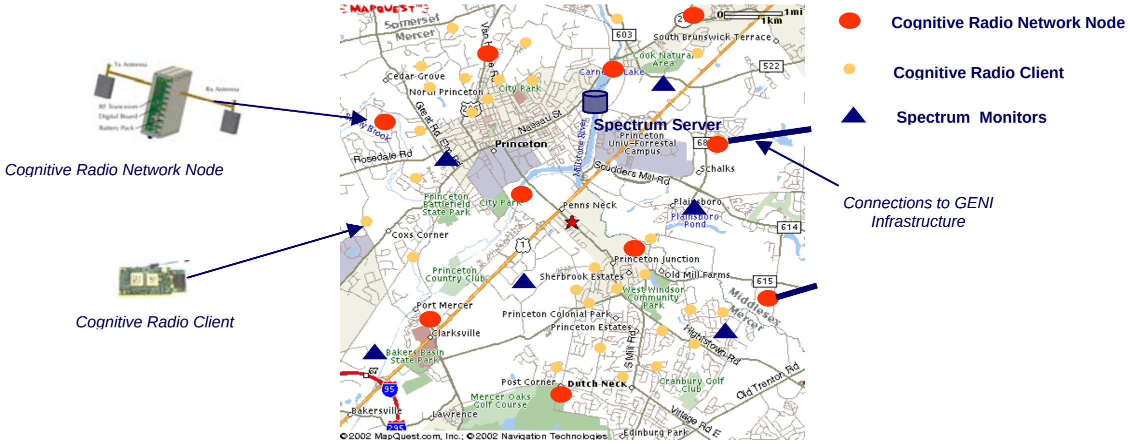

Figure 16 Deployment Concept for Cognitive Radio Demonstrator Network

A typical trial cognitive radio deployment would be in a suburban/medium-density coverage area ~50 Sq-Km with the objective of demonstrating and evaluating this technology as an alternative to available cellular and hybrid cellular/WLAN solutions. Deployment of this system also involves construction of a distributed spectrum measurement infrastructure along with centralized spectrum coordination resources (such as spectrum broker, spectrum server). A new wideband experimental spectrum allocation would also be required to support this trial network.

7.4.6 Wireless networks as core Internet infrastructure

Wireless networks as core Internet infrastructure using radio routers and emerging WiMax or MIMO 802.11 extended range wireless physical layers – see Fig 17. The focus of this experimental system would be on the use of wireless technologies for rapidly deployable and robust Internet infrastructure. Wireless core networks of this type are of particular interest as a building block for emergency response systems. Research topics to be studied using this experimental network include use of MIMO/directional antennas and WiMax, super high-speed wireless data link protocols, self-organization and network topology adaptation, integration with wired network protocols, spectrum efficiency studies, comparison with wired network solutions for robust/ emergency infrastructure application. In this case, a typical experimental wireless network would consist of a total of ~20 wireless routers including ~5 with interfaces to Tier 1 Internet points-of-presence. Each radio router would have software flexibility and virtualization features.

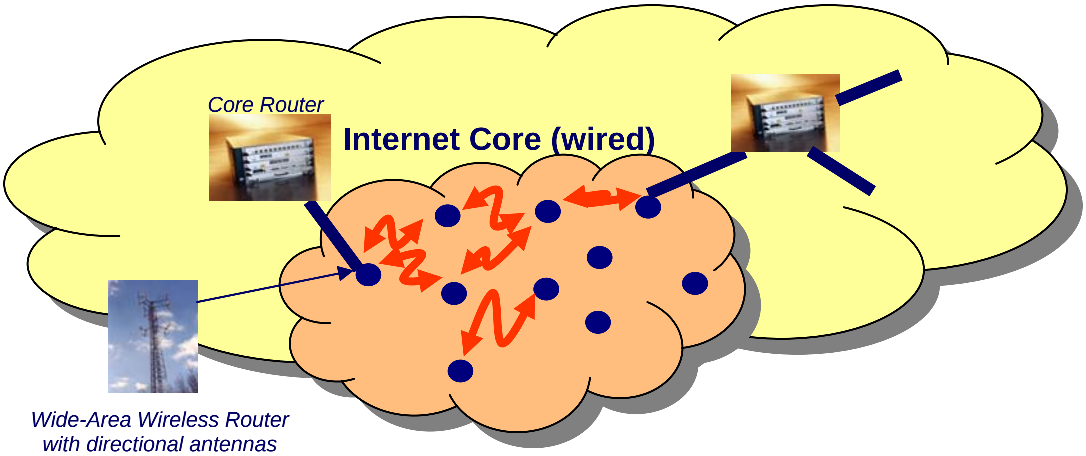

Figure 17 Experimental Wireless Core Network for Internet Infrastructure

In addition to the above six experimental networks, several Internet based network services will be required to support experimentation – these include location, spectrum coordination and sensor address resolution services.

8 Roadmap for Future Wireless Network Research

This document has presented a rationale for research on future wireless networks and the Internet as a whole, along with general discussions on requirements, architecture and experimental infrastructure. Panelists at the WMPG workshop believe that design, analysis, implementation and validation of future wireless networks and their integration with the global Internet represents a major challenge for our research community.

Technical challenges related to network protocol design are summarized in Sec 4, and include topics such as naming and addressing, integrating location, self-organization, cross-layer protocols, cognitive radio support, security/privacy, and sensor net features (such as content-awareness, lightweight protocols and data integrity). Each of these topics is being studied by the research community under NSF CNS programs such as NeTS, but most results are presented in isolation for self-contained wireless or sensor networks. Looking ahead, it would be important to incorporate end-to-end networking considerations in terms of the interface to the wide-area network (both data and control) and clearer specification of the protocols used. As discussed in Sec 6, there are several options for deploying new wireless network protocols including IP evolution, overlays or clean-slate architectures. It would be of interest to evaluate cost/performance trade-offs for each of these scenarios and to propose specific protocols that can be used as input by those working on end-to-end network protocols, whether IP extensions, overlays or new concepts.

The single most popular area for wireless networking research during the past 10 years has been ad hoc network routing. Considerable progress has been made on this topic, and focus needs to shift to broader problems associated with ad hoc networking, such as discovery and topology optimization, or the interaction between medium access control and routing. Since most wireless medium access protocols were originally designed for a star topology, it may be appropriate to examine fundamental changes to medium access control protocols for the multi-hop ad hoc network environment. An open research issue in ad hoc networking that needs further resolution is that of cross-layer protocol design. This topic has been studied for a few years, and early results on cooperative physical layers, PHY-aware routing, joint MAC/routing algorithms, etc. have demonstrated the potential for significant performance improvements. However, this is at the expense of protocol complexity and difficulties with carrying lower-layer protocol information across the network. There is a real need for more specific cross-layer protocol designs along with validation of performance and overhead on an end-to-end basis.

The topic of location-aware networking is an important theme for the future Internet. Several panelists agree that introducing location capabilities is an important goal not just for wireless edge networks but for the Internet as a whole. Research is needed to determine the protocol framework for providing location awareness – for example, are overlay services sufficient for this purpose, or should the networking protocol support a geographic routing mode? Also needed are underlying methods for location determination using available radio signal measurements, GPS and other technologies.

Adaptive networks of cognitive radios represent a major new challenge for the wireless research community. Emerging cognitive radio technologies provide a great deal of flexibility right down to the physical layer, making it possible to build highly adaptive networks that share spectrum efficiently and achieve significant improvements networking performance when compared with today's technology. New protocols and algorithms are needed for spectrum coordination, resource discovery, physical layer negotiation, network bootstrap and topology management. As with the cross-layer protocols mentioned above, specific cognitive network protocols need to be evaluated in context of end-to-end networking scenarios. Progress on cognitive radio will require extensive experimental validation in addition to design and simulation work.

In the area of sensor networking, the past 5 years have produced a significant body of research results that serve as a foundation for future Internet research. In particular, sensor network research has shifted from abstract consideration of ad hoc routing and energy efficiency issues to more complete hierarchical systems integrated with the Internet. At the same time, application-specific implementations have revealed a number of important design issues such as in-network processing and content-aware services. There is a need for generalized sensor network protocol designs along with end-to-end service/API models similar in spirit to the widely used TCP/IP socket. Closed loop, real-time sensor/actuator systems represent an important new thrust for sensor network research – these networks will require low latency and tight integration with Internet services that result in new protocol design considerations. As with the other scenarios discussed above, more consideration should be given to end-to-end protocol design issues at both network and higher layers.

Security and privacy issues for mobile/wireless and sensor network scenarios remains an important research focus for the community. It is recognized that as wireless networks proliferate, they will provide the easiest point of entry for security attacks ranging from eavesdropping to denial-of-service, and hence the wireless security problem takes on an added importance for the Internet as a whole. Research on security in distributed wireless networks (such as ad hoc or sensor networks) is still at an early stage, and significant effort is needed to build authentication, encryption and other security features into the protocols adopted rather than fix problems when they arise later. An important area for future work is the problem of decentralized trust models required to address the increasingly heterogeneous and distributed nature of wireless networks. Solutions for denial of service attacks in wireless networks that take into account the open nature of the radio channel (for both attacker and defender) are worth pursuing further. Sensor network security issues include problems of verifying data integrity in a distributed and physically accessible network environment. Privacy in wireless and sensor networks is another critical new area as it relates to fundamental social impact and acceptability of emerging location-aware network technologies.

The WMPG panel’s discussions also focused on experimental research aspects considered critical to sustainable advances in this field. It is noted that the wireless research community has made significant progress with experimental research in the past 3-5 years due in part to support from NSF programs such as Network Research Testbeds (NRT) and NeTS. The NSF community has recently produced a number of notable results both with core radio network technologies as well as with network emulators, testbeds and field systems. Taken together, these experimental initiatives have had a remarkable impact on research culture with numerous papers and thesis projects based on prototyping and testbed/field evaluation rather than simulation. The community is encouraged to share experimental resources (platform designs, driver source code, measurement software, etc.) so as to create a critical mass capability for wireless networking research. Taking stock of current capabilities, it is clear that 802.11x based experimental capabilities are now relatively widespread, and researchers also have access to large scale shared testbeds such as ORBIT. Similarly, experimentation with sensor nets has been enabled by wide distribution of the Mote platform, and a number of shared sensor net testbeds (such as Motelab or ENL) have become available for general use. At the same time, significant progress has been made with application-specific field trials, such as the CENS/UCLA projects on environmental monitoring.

Future experimental work on wireless networking should leverage existing testbed and platform resources to validate protocol designs and evaluate end-to-end networking issues. Since much of the existing experimental base is with 802.11 or Mote radios, there is a need to develop open API platforms with other radio technologies such as Zigbee, WiMax or 3G. In particular, it is important that open API platforms be developed for wide-area radios, such as WiMax or 3G, in order to close the gap between research and practice. Significant opportunities for new networking concepts and applications exist for hybrid networks with both wide-area and short-range radios directly connected to the Internet, free from vertical protocol stacks associated with today's cellular systems. Another important core technology target is the cognitive radio platform which may be expected to emerge from a number of ongoing research projects funded by NSF and DARPA. Even after initial hardware alternatives are evaluated, considerable research effort will be needed to develop software frameworks and the necessary protocols for cognitive radio networks. Collaborative efforts will be required for critical mass projects in this new technology area. A related experimental goal is that of dynamic spectrum measurement and management. While some projects have been started in this area, much remains to be done to establish the feasibility of agile radio and adaptive networking concepts in dense deployments.

Clearly, the next major step for experimental wireless research is the establishment of large-scale field trial systems similar to those proposed in Sec 7. The WMPG panel feels that the research community should make an attempt to go beyond controlled testbeds to develop large-scale field trial systems for in-depth validation of emerging protocols and algorithms for mobile, wireless and sensor networks. These wireless field systems should connect to a flexible experimental infrastructure for end-to-end networking studies and application development. Some topics related to large-scale experimental research include development of flexible radio platforms, implementation of virtualization and overlays on wireless networks, programmability and end-to-end software models, network management and remote monitoring/control, spectrum measurement, and protocol performance measurement across various layers. Since large-scale experimental networks require major investments in infrastructure, the community is encouraged to collaborate and consolidate resources with other research communities including those working on end-to-end network protocols, distributed computing and optical networks.

9 Summary of Major Findings and Recommendations

9.1 Major Findings

Wireless, mobile and sensor devices will drive fundamental transformation of the Internet: Continuing proliferation of wireless laptops, cell phones, PDA's, sensors, media players and other new consumer devices will have a fundamental impact on the Internet architecture. In 2015, there will be billions of wireless devices on the Internet and these will certainly outnumber wired PC's used for most of today's applications. Future applications based on mobile devices and sensors will involve the integration of physical-world people, objects and events with the data and computing on the Internet and will lead to a very different set of networking requirements.

Emerging wireless scenarios can be classified into three general categories: mobile data, ad-hoc nets and sensor nets: These three scenarios cover a broad range of application environments such as cellular data services, WiFi hot-spots, Infostations, mobile peer-to-peer, ad-hoc mesh networks for broadband access, vehicular networks, sensor networks and pervasive systems. These wireless application scenarios lead to a rich diversity of networking requirements for the future Internet that need to be analyzed and validated experimentally.

Future wireless network architecture requirements distilled from the above scenarios are:

1. *Naming and addressing flexibility* with a clean separation of name and routable address.

2. *Mobility support* which provides for dynamic migration of end-users, network elements such as routers, and ad hoc clusters.
3. *Location services* that provide communicating entities with information on geographic position.
4. *Self-organization and discovery* for rapid formation of networks in a distributed environment
5. *Security and privacy* considerations that take into account node mobility and the open nature of the wireless channel.
6. *Decentralized management* for remote monitoring, configuration and control of distributed and heterogeneous wireless networks.
7. *Cross-layer protocol support* for optimization of communications performance in complex radio environments with link layer adaptation and with integrated design of physical, MAC and routing protocols.
8. *Sensor network integration* that supports lightweight and power efficient sensor operation, including data aggregation, content awareness and in-network programming models.
9. *Cognitive radio support* with dynamic negotiation of link parameters and opportunistic ad hoc formation of adaptive wireless networks.
10. *Economic incentives* to encourage efficient sharing of spectrum resources and for ad hoc forwarding between peers.

Several alternative approaches for deployment of new network services can be identified: Options range from evolution of IP, top-down transformation of the network via protocol overlay, to bottom-up replacement of the IP layer and protocols above. Overlays are considered to be a promising approach to transforming the network without requiring complete replacement of the network infrastructure. Large-scale experimental studies with real users and applications will be needed to evaluate the effectiveness of these deployment strategies.

Future wireless network protocols will require extensive experimental evaluation under realistic conditions: Wireless systems involve real-world mobility and channel impairments which are difficult to model via simulation, so that there is an important need for real-world testbeds for validation, comparison and competitive selection of new protocol ideas. Existing wireless research testbeds do not support evaluation of large-scale networks in real-world settings, motivating construction of a new generation of flexible experimental wireless systems connected to a suitable wide-area network research infrastructure.

Several enabling technologies and methods for building flexible wireless systems are becoming available for building flexible experimental wireless networks: Flexible experimental networks require programmability of radio physical layer, medium access control and higher layer protocols along with tools for debugging and measurement. Recent work on wireless testbeds and platforms has led to the availability of several key technologies for building flexible systems. These include

- *soft MAC* for virtualization of wireless networks
- *cognitive radio* technology for flexible PHY layers
- *platform hardware and software* for general purpose sensor node, ad-hoc node and radio router/access point
- *measurement techniques* for wireless networks
- wireless network *data repository*
- wireless network *emulators and simulators*

A number of specific experimental wireless networks are recommended in support of future network protocol and applications research: Research on future wireless network protocols and their integration with the global Internet requires coverage of several emerging mobile computing, ad-hoc and sensor network scenarios. Realistic, large-scale deployments in urban or suburban environments are required to adequately evaluate emerging wireless networking protocols, as well as to facilitate applications research. Access to state-of-the-art simulation and emulation facilities is also considered to be a useful supplement to real-world experimental systems. Specific concepts for experimental wireless networks include:

- *Wireless emulation testbeds* integrated with future Internet research network
- *Urban ad hoc mesh network* with ~1000 radio nodes in a dense city deployment
- *Suburban wide area network with open access 3G/WiMax* and 802.11 hot spots
- *Cognitive radio technology demonstrator*
- *Application specific sensor network deployments* (in conjunction with urban mesh)
- *Wireless infrastructure* for the Internet backbone

9.2 Recommendations

Recommendation 1: Recognize that the Internet will undergo a fundamental transformation over the next 10-15 years, and invest in research programs aimed at creating technical foundations for the future network.

Panelists at the WMPG workshop feel that the Internet will be fundamentally transformed by emerging wireless, mobile and sensor network applications that will enable people to view and interact with the physical world around them. TCP/IP protocols used in today's Internet do not provide the necessary tools for efficient application development and need to be redesigned and possibly replaced with a new generation of Internet protocols both at the edge and on an end-to-end basis. Disruptive innovation in future Internet architectures and protocols will require sustained NSF and other research investment aimed at creating the necessary technical foundations including network analysis, architecture, protocol design, prototype implementation and operational experience.

Recommendation 2: Increase research focus on central network architecture questions related to future mobile, wireless and sensor scenarios.

The research community is encouraged to focus on critical network architecture issues related to efficient support of wireless devices in the future Internet. Panelists identified several major architectural directions in wireless/mobile/sensor networking including: naming and addressing flexibility, dynamic mobility support, integrated location services, ad-hoc self-organization and discovery, security and privacy in distributed wireless networks, decentralized management, cross-layer protocol support, sensor network integration (including data aggregation, content awareness and in-network programming), adaptive wireless networks based on cognitive radios, and economic incentives for spectrum sharing and peer networking.

Recommendation 3: Invest in development of flexible wireless technologies and platforms necessary to implement programmable and evolvable experimental networks.

Wireless technology has progressed significantly during the last decade, both in terms of quantitative Moore's law improvements (from ~10 Kbps to 100 Mbps and higher) and qualitative capabilities (compact sensors, ad-hoc radios, etc.). Flexible and programmable components are essential for building future networks with yet-to-be-determined protocols and changing wireless technology. Panelists agreed that it is important for the NSF to invest in the development of flexible wireless platforms for experimental use, a need that is less likely to be addressed by industry. Specific technologies and platforms that were identified include software/cognitive radios; open API 802.11, 802.14, 3G and WiMax PHY; 802.11 with virtualizable MAC layer; general purpose ad hoc radios, forwarding nodes and radio routers; wireless network measurements and data repositories; and wireless network simulators and emulators.

Recommendation 4: Fund development of large-scale experimental wireless networks for effective validation and competitive selection of new architecture and protocol concepts.

Since wireless systems involve real-world effects that are difficult to model (such as signal impairments and mobility), complete validation and competitive selection of future network protocols will require large-experimental networks accessible to the research community. The panel recommends construction of the following experimental wireless networks to cover a reasonable cross-section of anticipated usage scenarios and technologies: (1) urban mesh network with 802.11 radios (this includes the emerging class of ad hoc vehicular applications); (2) suburban wide-area cellular network with open API 3G/WiMax radios; (3) cognitive radio technology demonstrator in urban/suburban setting; (4) several application specific sensor networks; and (5) wireless Internet infrastructure network. These experimental wireless systems should be connected to a wide-area networking research network capable of supporting flexible experimentation on new architectures, protocols and application concepts.

Recommendation 5: Encourage collaborative research resulting in end-to-end deployment and evaluation of future wireless/mobile and sensor applications over the global Internet.

WMPG panelists felt that progress towards the future Internet would require improved collaboration between research groups with expertise in wireless access networks, end-to-end Internet protocols, overlay protocols, distributed systems and domain-specific applications. NSF should encourage systems projects with end-to-end networking solutions that support novel mobile/wireless applications with real end-users. Partnerships with industry and software development communities are considered necessary for effective deployment of new networking concepts and applications.

10 References

[Ac04] B. Ackland, et al, "NeTS-ProWiN: High performance cognitive radio platform with integrated physical and network layer capabilities", NSF grant CNS-0435370, September 2004.

[Al03] Alastair R. Beresford and Frank Stajano, "Location privacy in pervasive computing", *IEEE Pervasive Computing*, 2(1), 2003.

[Ba95] A. Bakre and B.R. Badrinath, "Indirect TCP for mobile hosts", *Proceedings of the 15th International Conference on Distributed Computing Systems*, 1995.

[Ba97] H. Balakishnan, V. Padmanabhan, S. Seshan and R. Katz, "A comparison of mechanisms for improving TCP performance over wireless links", *ACM Transactions on Networking*, 5(6), 1997

[Ba02] C. Barrett et al., "Characterizing the interaction between routing and MAC protocols in Ad-hoc Networks," *Proceedings of MobiHoc*, 2002 .

[Bell03] J. Bellardo and S. Savage, "802.11 Denial-of-Service attacks: real vulnerabilities and practical solutions." *In Proceedings of the USENIX Security Symposium*, pages 15--28, August, 2003.

[Bo03] M. Bohge and W. Trappe, "An authentication framework for hierarchical ad-hoc sensor networks," *Proceedings of the ACM Workshop on Wireless Security*, 2003.

[Bo04] V. Bose, "A software driven approach to SDR design", *COTS Journal*, January 2004.

[Bor01] N. Borisov and I. Goldberg and D. Wagner, "Intercepting mobile communications: the insecurity of 802.11", *In MobiCom '01: Proceedings of the 7th annual ACM International Conference on Mobile computing and networking*, pg. 180—189, 2001.

[Br02] R. Braden, T. Faber and M. Handley, "From protocol stack to protocol heap -- role-based architecture". *Proceedings of ACM HotNets-I Workshop*, October 2002.

[BRO] Broadband ISDN standard. http://www.itu.int

[Bu05] Milind M. Buddhikot Paul Kolodzy Scott Miller Kevin Ryan, Jason Evans, "DIMSUMNet: New Directions in Wireless Networking Using Coordinated Dynamic Spectrum Access", IEEE WoWMoM 2005, www1.bell-labs.com/user/mbuddhikot/ psdocs/**dimsum**-wowmom05-revised.pdf

[Ca95] R. Caceres and L. Iftode, "Improving the performance of reliable transport protocols in mobile computing environments", *IEEE Journal Of Selected Areas in Communications, Special Issue on Mobile Computing*, June 1995.

[Ca00] A. Carzaniga, D. S. Rosenblum and A. L. Wolf, "Content-based addressing and routing: a general model and its application", *Technical Report CU-CS-902-00, Department of Computer Science, University of Colorado*, January 2000.

[Ch04a] K Challapali, Stefan M. and Z. Zhong, "Spectrum agile radio: Detecting spectrum opportunities*", 6th Annual International Symposium On Advanced Radio Technologies*, March 2004.

[Ch04b] M. Chan and R. Ramjee, "Improving TCP/IP performance over third generation wireless networks", *Proceedings of IEEE Infocom*, 2004.

[Ch05] L. Chen et al, "Adhoc probe: Path capacity probing in wireless ad-hoc networks", *Proceeding of WICON*, July 2005

[Cl03a] D. Clark, K. Sollins, J. Wroclawski and T. Faber, "Addressing reality: an architectural response to real-world demands on the evolving Internet", *Proceedings of ACM SIGCOMM FDNA Workshop*, August 2003

[Cl03b] D. Clark, R. Braden, A. Falk and V. Pingali, "FARA: Reorganizing the addressing architecture", *Proceedings of ACM SIGCOMM FDNA Workshop*, August 2003

[Cl05] D. D. Clark, C. Partridge, R. T. Braden, B. Davie, S. Floyd, V. Jacobson, D. Katabi, G. Minshall, K.K. Ramakrishnan, T. Roscoe, I. Stoica, J. Wroclawski and L. Zhang, "Making the world (of communications) a different place", *Report of a working session of the End-to-End Research Group*, January, 2005, http://www.ir.bbn.com/~craig/e2e-vision.pdf

[CRO] Crossbow. http://www.xbow.com/Products/Wireless_Sensor_Networks.htm

[Cu04] D. Culler, D. Estrin, and M. Srivastava. "Overview of sensor networks". *IEEE Computer, Special Issue in Sensor Networks*, August 2004

[Da04] F. Dabek, R. Cox, F. Kaashoek and R. Morris. "Vivaldi: a decentralized network coordinate system", *Proceedings of ACM Sigcomm* , 2004.

[De03] D. S. J. De Couto, D. Aguayo, J. Bicket and R. Morris, "A high-throughput path metric for multi-hop wireless routing", *Proceedings Mobicom*, September 2003.

[Do01] C. Dovrolis, P. Ramanathan and D. Moore, "What do packet dispersion techniques measure?", *Proceedings of IEEE Infocom*, 2001.

[Ef01] F.J. Effenberger, H. Ichibangase and H. Yamashita H, "Advances in broadband passive optical networking technologies", *IEEE Communications Magazine*, (39)12, 2001.

[EMU] EMULAB project web site. http://www.emulab.net

[Es99] D. Estrin, R. Govindan, J. Heidemann and S. Kumar, "Next century challenges: Scalable coordination in sensor networks", *Proceedings of the fifth annual ACM/IEEE international conference on Mobile computing and networking,* August 1999.

[Fa03] K. Fall, "A delay-tolerant network architecture for challenged internets", *Proceedings of the Conference on Applications, technologies, architectures, and protocols for computer communications,* 2003.

[Fr96] R.H. Frenkiel and T. Imielinski, "Infostations: The joy of many–time, many–where communications", *WINLAB TR–119,* April 1996

[Ga04] S. Ganu, L. Raju, B. Anepu, S. Zhao, I.Seskar and D. Raychaudhuri. "Architecture and prototyping of an 802.11-based self-organizing hierarchical ad-hoc wireless network (SOHAN)", *Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 2004.

[Ge01] M. Gerla, R. Kapoor, M. Kazantzidis and P. Johansson, "Ad-hoc networking with Bluetooth", *Proceedings of WMI at Mobicom*, July 2001.

[Ge02] M. Gerla et al., "UMTS-TDD: A solution for internetworking Bluetooth piconets in indoor environments", *Seventh International Symposium on Computers and Communications*, July 2002.

[GNU] GNURadio. http://www.gnu.org/software/gnuradio/

[Go97] D. Goodman, J. Borras, N. Mandayam, and R. Yates, "Infostations: a new system model for data and messaging services", *Proceedings of IEEE VTC*, 1997.

[Gr03] M. Gruteser, G. Schelle, A. Jain, R. Han, and D. Grunwald, "Privacy-aware location sensor networks", *Proceedings of Workshop on Hot Topics in Operating Systems (HotOS)*, 2003

[Gr04] P.E. Green, "Fiber to the home: the next big broadband thing", *IEEE Communications Magazine*, (42)9, Sept 2004.

[He03] S. Helal, B. Winkler, C. Lee, Y. Kaddourah, L. Ran, C. Giraldo and W. Mann, "Enabling location-aware pervasive computing applications for the elderly", *Proceedings of 1st IEEE Conference on Pervasive Computing and Communications (Per-Com)*, March 2003

[Hu98] D.R. Hush and W. Cliff, "Analysis of tree algorithms for RFID arbitration", *Proceedings of IEEE International Symposium on Information Theory*, 1998

[Hu03] Y.C. Hu, A. Perrig, and D. B. Johnson. "Packet Leashes: A Defense against Wormhole Attacks in Wireless Networks." *In Proceedings of IEEE INFOCOM*, pg. 1976-1986, 2003

[IEE04] IEEE 802.11i Standard Extension for Medium Access Control (MAC) Security Enhancements, http://standards.ieee.org/getieee802/download/802.11i-2004.pdf

[Jo02] P. Johansson, R. Kapoor, M. Kazantzidis and M. Gerla, "Personal area networks: Bluetooth or IEEE 802.11?", *Proceedings of International Journal of Wireless Information Networks, Special Issue On Mobile Ad-hoc Networks (MANETs) Standards, Research Applications*, April 2002

[Jo04] D. Johnson, C. Perkins and J. Arkko, "Mobility support in IPv6", RFC 3375, June 2004

[JOI] Joint Tactical Radio System. http://jtrs.army.mil

[Ka00] B. Karp and H. T. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks", In *Proc. ACM/IEEE MobiCom*, August 2000, pp 243-254.

[Ka02] D. Katabi, M. Handley and C. Rohrs, "Congestion control for high bandwidth-delay product networks", *Proceeding of ACM SIGCOMM*, August 2002.

[Ka04] R. Kapoor, L.-J. Chen, L. Lao, M. Gerla and M. Y. Sanadidi, "CapProbe: A simple and accurate capacity estimation technique," *Proceedings. of ACM SIGCOMM*, 2004.

[Ko02] J. Kong, H. Luo, K. Xu, D. Gu, M. Gerla, and S. Lu, “Adaptive security for multi-layer ad-hoc networks,” *Special Issue of Wireless Communications and Mobile Computing*, 2002.

[Ku01] U. Kubach, and K. Rothermel. “Exploiting location information for infostation-based hoarding”, *Proceedings of ACM Mobicom*, 2001.

[Li00] J. Li, J. Jannotti, D. De Couto, D. Karger, and R. Morris, “A scalable location service for geographic ad-hoc routing”, *Proceedings of ACM MobiCom*, August 2000.

[LOC] LocustWorld Project. http://www.locustworld.org

[Ma04] N. Mandayam, C. Rose, R. Yates, P. Spasojevic, "Cognitive radios for open access to spectrum,", WINLAB Research Review, Oct 2004, Rutgers University, http://www.winlab.rutgers.edu/~crose/papers/ spectrumserver.ppt

[Mi99] J. Mitola III and G.Q. Maguire, "Cognitive radio: making software radios more personal," *IEEE Personal Communications*, August 1999.

[Mi05] Gary Minden, “KU Agile Radio Technology”, http://www.ittc.ku.edu/techreview2005/presentations/ Minden_Agile%20Radios.ppt

[MIT] MIT RoofNet Project. http://www.pdos.lcs.mit.edu

[Na02] B. Nath and D. Niculescu, "Routing on a curve", *Proceedings of HotNets-I*, October, 2002.

[Na05] A. Nandan, S. Das, G. Pau, M.Y. Sanadidi and M. Gerla, “Cooperative downloading in vehicular wireless ad-hoc networks”, *Proceedings of Wireless On-Demand Networks and Services*, January 2005.

[NAP] The New Arch Project: Future Generation Internet Architecture, http://www.isi.edu/ newarch/iDOCS/final.finalreport.pdf

[Ng03] T.S.E Ng, Y. Chu, S.G. Rao, K. Sripanidkulchai, and H. Zhang, ``Measurement-based optimization techniques for bandwidth-demanding peer-to-peer systems", *Proceedings of IEEE INFOCOM*, 2003.

[ORB] ORBIT Project web site. http://www.orbit-lab.org

[Ot04] M Ott, L. French, R. Mago and D. Makwana. “XML-based semantic multicast routing: An overlay network architecture for future information services”, *Proceedings of IEEE Global Telecommunication Conference)*, 2004.

[Pa03] G. Pau, D. Maniezzo, S. Das, Y. Lim, J. Pyon, H. Yu, M. Gerla, “A cross-layer framework for wireless LAN QoS support”, *Proceedings of IEEE ITRE*, August, 2003.

[PAS] Passive Optical Networking Forum. http://www.ponforum.org

[Per02a] C Perkins, “IP mobility support for IPv4”, RFC 3344, August 2002.

[Per02b] C. Perkins, “Mobile IP”, *IEEE Communications Magazine 50th Anniversary Ed*, May 2002.

[Per02c] A. Perrig, R. Szewczyk, D. Tygar, V. Wen, and D. Culler, “SPINS: Security protocols for sensor networks”, *Wireless Networks*, (8)5, 2002.

[Pet02] L. Peterson, T. Anderson, D. Culler, and T. Roscoe, “A blueprint for introducing disruptive technology into the Internet”, *Proceedings of First Workshop on Hot Topics in Networking (HotNets-I)* , October 2002.

[Pet04] L Peterson, S. Shenker and J. Turner, “Overcoming the Internet impasse through virtualization”, *Proceedings of the 3rd ACM Work on Hot Topics in Networks (HotNets-III)*, November 2004

[Pet05] L. Peterson at al, “Overcoming Barriers to Disruptive Innovation in Networking”, Report of NSF Workshop, Jan 2005.

[PLA] PlanetLab Project Website. http://www.planet-lab.org/.

[Pr01] N. B. Priyantha, A. K.L. Miu, H. Balakrishnan and S. Teller, “The cricket compass for context-aware mobile applications”, *Proceedings of the 7th annual international conference on Mobile Computing and Networking*, June 2001.

[Ra99] N. Raza, V. Bradshaw and M. Hague, “Applications of RFID technology”, *IEE Colloquium RFID Technology (Ref. No. 1999/123)* , 1999.

[Ra05] A. Rao and I. Stoica, "An overlay MAC layer for 802.11 networks", *Proceedings of ACM Mobisys*, June 2005.

[Ray04] D. Raychaudhuri, Adaptive Wireless Networks Using Cognitive Radios as a Building Block, *MobiCom 2004 Keynote Speech*, Philadelphia, PA, September, 2004. http://www.winlab.rutgers.edu/pub/Index.html

[Ray05] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu and M. Singh "Overview of the ORBIT Radio Grid Testbed for Evaluation of Next-Generation Wireless Network Protocols", *Proc. IEEE WCNC*, March 2005.

[Sa03] N. B. Salem, L. Buttyan, J. P. Hubaux, M. Jakobsson. "A charging and rewarding scheme for packet forwarding." *Proceedings of MobiHoc*, June 2003.

[Si05] M. Singh, M. Ott and I. Seskar, "ORBIT Measurements Framework and Library (OML): Motivations, Design, Implementation, and Features," *Proc. IEEE Tridentcom '05*, Feb 2005.

[St01] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: a scalable peer-to-peer lookup service for internet applications", *Proceedings of ACM SIGCOMM*, August 2001.

[STA] Stargate. http://www.xbow.com/Products/XScale.htm

[TEX] Texas Instruments Radio Frequency Identification (RFID) Solutions Home Page. http://www.ti.com/tiris

[To03] S. Toumpis and A. J. Goldsmith, "Performance, Optimization, and Cross-Layer Design of media access protocols for wireless ad-hoc networks", *Proceedings of the IEEE ICC 2003*, 2003.

[UNI] Universal Mobile Telecommunications System (UMTS), http://www.umts-forum.org.

[Va02] N. H. Vaidya, "Weak duplicate address detection in mobile ad-hoc networks". *Proceedings of MobiHoc*, 2002.

[VAN] VANU. http://www.vanu.com

[WHY] WHYNET.project web site. http://chenyen.cs.ucla.edu/projects/whynet/

[Xu05] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks," *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (Mobihoc05)*, pg. 46-57, 2005.

[Yat02] R. Yates, N. Mandayam, D. Raychaudhuri, C. Rose and P. Spasojevic. "ITR: collaborative research: Achieving innovative and reliable services in unlicensed spectrum", NSF grant CCR-0205362, in collaboration with Michigan State (CCR-0205362) and Cornell (CCR-0205431).

[Yar02] M. D. Yarvis, W. S. Conner, L. Krishnamurthy, J. Chhabra, B. Elliott and A. Mainwaring, "Real-world experiences with an interactive ad-hoc sensor network", *Proceedings of the International Conference on Parallel Processing Workshops*, 2002.

[Yeo04] Jihwang Yeo, Moustafa Youssef, and Ashok Agrawala, "A Framework for Wireless LAN Monitoring and Its Applications" *ACM Workshop on Wireless Security (WiSe 2004)* in conjunction with ACM MobiCom 2004, Philadelphia, PA, US, October 2004.

[Yu04] Y. Yu, G. Lu, and Z. Zhang. "Enhancing location service scalability with HIGH-GRADE", *IEEE MASS Conference Proceedings*, October 2004.