Technical Document on Overview – Wireless, Mobile and Sensor Networks

GDD-06-14

GENI: Global Environment for Network Innovations

September 15, 2006

Status: Draft (Version 2.0)

Note to the reader: this document is a work in progress and continues to evolve rapidly. Certain aspects of the GENI architecture are not yet addressed at all, and, for those aspects that are addressed here, a number of unresolved issues are identified in the text. Further, due to the active development and editing process, some portions of the document may be logically inconsistent with others.

This document is prepared by the Wireless Working Group.

Editors:

S. Paul, *Rutgers*

J. Evans, U Kansas

D. Raychaudhuri, *Rutgers*

Contributing workgroup members:

V. Bahl, Microsoft

C. Bostian, Virginia Tech

C. Elliott, BBN

D. Estrin, UCLA

E. Belding-Royer, UCSB

M. Gerla, UCLA

G. Minden, U Kansas

J. Heidemann, ISI / USC

E. Knightly, Rice

M. Welsh, Harvard

R. Govindan, USC

S. Seshan, CMU

W. Trappe, Rutgers

1 Executive Summary

This is an interim report on wireless, mobile and sensor network aspects of GENI, and is intended to be a work-in-progress snapshot from the GENI Wireless Working Group (WWG) as of July 2006. The GENI Wireless Working Group activities were initiated in March 2006 under the framework of the GENI Project Development Plan (PDP). The working group is led by D. Raychaudhuri and J. Evans as co-chairs, and currently has 13 other members with expertise across various wireless networking specializations including areas such as ad hoc networking, mobile and vehicular networks, sensor networks, cognitive radio and security.

The GENI wireless working group's scope-of-work and primary deliverables (during the one-year period from 3/06-3/07) as defined in the GENI PDP document are:

- Baseline design specifications for GENI wireless subnets
  - Draft reports on wireless subnet implementation plans
  - Final report, including design specs and system engineering details
- Design studies and laboratory evaluations of candidate experimental wireless technologies and platforms
  - Wireless platform specification document
  - Critical platform and radio technology studies
  - Selected platform evaluation results and recommendations
- GENI integration plans including control, management and experimental measurement frameworks
  - Report on GENI integration including control/management interface specification for wireless
- Guidance for GENI development projects and wireless networking research
  - Proposals for proof-of-concept projects on key hardware and software technologies
  - Experimental wireless network research agenda and roadmap
  - Planning for community events and demos in support of the GENI readiness process

The WWG's plenary discussions during the period 3/06-6/06 have focused on updating the baseline spec for each of the wireless GENI subnetworks under consideration, i.e. wireless network emulators, urban mesh, suburban wide-area, sensor networks and cognitive radio. The WG has also discussed experimental research in the wireless/mobile/sensor areas in order to better identify infrastructure requirements, and to provide wireless community input to the GENI Research WG. This working document is intended to summarize the results of these discussions and to identify future work and inter WG coordination topics. Sec 2 below provides an overview of the wireless/mobile/sensor network research agenda in the form of scenario descriptions, some Internet architecture considerations, and a "top 10" experiment list contributed by various members of the WG. This is followed by Sec 3 which provides updated descriptions of each of the wireless subnetworks planned for GENI, with technical details being filled in as they become available. Note that the high-level system descriptions in Sec. 3 may be expected to evolve during the preparedness process as technical studies on applicable platforms, key technologies, system software and deployment planning become available from the WG.

The WWG also focused on a number of critical technical issues that need to be resolved to finalize the GENI wireless development plan. In particular, several sub-working groups were formed to address the following open technical issues:

- Virtualization in wireless networks
- Sensor network deployment details
- Vehicular mobility requirements and experimental scenarios
- Integration of control and management across wired and wireless networks in GENI
- Programmable and cognitive radio technology issues and deployment plan

Preliminary reports on each of the above sub-working group study topics are given in Sec. 4. Sec 5 provides a work-in-progress view of the system engineering methodology by which a specific wireless subnet in GENI (the urban mesh) is being designed in accordance with WBS guidelines. The draft report concludes in Sec. 6 with a list of future work items and inter-WG coordination items.

2 Wireless research vision and sample GENI experiments -

The Wireless Mobile Planning Group (WMPG) report[1] issued in Aug 2005 identified several usage scenarios associated with emerging wireless and sensor network scenarios. Summaries of these scenarios are reproduced from the WMPG report as follows.

2.1 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

[1] "New Architectures and Disruptive Technologies for the Future Internet – the Wireless, Mobile and Sensor Net Perspective", WMPG workshop report, Aug 2-3, 2005.

(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.

2.2 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.

2.3 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 above scenarios, it is possible to identify some architecture and protocol design areas for the future Internet that will be driven by wireless requirements.

2.4 Impact on Future Internet Architecture

The Aug 2005 WMPG workshop also identified several architectural drivers for the future Internet based on the scenarios discussed in Sec 2.3 above. Summaries of key wireless/mobile/sensor related features are reproduced from the report as follows:

2.4.1 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.

2.4.2 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.

2.4.3 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.

2.4.4 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.

2.4.5 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.

2.4.6 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.

2.4.7 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.

2.4.8 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. 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.

2.4.9 Vehicular Network Integration

Vehicular communications are becoming a reality, driven by navigation safety requirements and by the investments of car manufacturers and Public Transport Authorities. Safe navigation support through wireless car to car and car to curb communications has become an important priority for Car Manufacturers as well as Municipal Transportation Authorities and

Communications Standards Organizations. New standards are emerging for car to car communications (DSRC and more recently IEEE 802.11p). The availability of powerful car radios, and of abundant spectrum (when not used for emergencies) will pave the way to a host of new applications for the Vehicle Grid (V-Grid). These emerging applications span many fields, from office-on-wheels to entertainment, mobile Internet games, mobile shopping, crime investigation, civic defense, etc. Some of these applications are conventional "mobile internet access" applications, say, downloading files, reading e-mail while on the move, etc. Others involve the discovery of local services in the neighborhood (ex, restaurants, movie theaters, etc) using the vehicle grid as Peer to Peer network. Others yet imply the close interaction among vehicles such as interactive car games. More generally, integration of vehicular networks with GENI will help investigate the Internet network and transport functions that will support these new applications.

2.4.10 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.

2.4.11 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.

2.5 Sample GENI Experiments

The GENI infrastructure needed to enable wireless networking architectural research is illustrated through the following example experiments contributed by members of the wireless working group.

2.5.1 Participatory urban sensing experiments

Participatory urban sensing will leverage mobile devices as a platform for data gathering in cities. This will drive development of network architecture through applications that have civic and cultural significance, such as participatory urban planning, community documentation, or tracking of environmental concerns across wide urban areas. This capability might eventually expand to applications that take advantage of network privacy and integrity to support distributed medical and personal health monitoring (e.g., common, repeated tests such as glucose levels and blood pressure) and finally into first responder and other critical urban applications.

This experiment will use the GENI facility to explore new network architecture to encourage data-sharing applications by providing novel privacy controls (e.g., in-network resolution control of data and context), data integrity information (e.g., in-network confirmation of location and time of data gathering nodes), and naming/dissemination services. It will expose new network features through simple APIs that can support top-down implementation of sophisticated urban sensing applications but also encourage 'bottom-up' participatory sensing.

The participatory urban sensing experiment will:

- Explore new features at three layers: low-level network services (close to IP layer), fabric level (above transport protocols), web services.
- Design a data sharing specification framework that allows end-users to easily specify sharing/privacy rules for their own data.
- Implement and evaluate an architecture of mediating access points and routers that support *selective sharing* based on those policies: resolution control of location and time context, provisions for operating on physical context information based on the sensor readings (adding jitter/noise to data on-the-fly to decrease resolution)
- At the fabric level, implement and evaluate naming and dissemination services that respect sharing policies provided by the mobile users. Additionally, consider fabric-level services that manage coverage estimation and sampling of these autonomous mobile nodes.
- Focus on imager, acoustic and GPS sensors that will be common in the handheld devices. Explore in-node signal processing to implement selective sharing controls, and even perform basic assays (e.g., pH) that use visual indicators with a calibration chart.

- Design simple web services APIs to promote platform usage in urban planning and participation tools created by collaborators in other disciplines.

This experiment will require campus, council district or other small-scale city deployments. Sample applications will need to be created in collaboration with community covered by the network. Each application will have components that run on the mobile nodes and as web applications. In order to experiment at reasonable scale, approximately 100 mobile devices will be needed for development (per deployment); with expansion targets to over 1000 existing mobile phone nodes. Access points and routers will need to be placed into experimental cellular and wireless networks. Blade-style servers will be needed to implement fabric services and basic web service APIs.

2.5.2 Internet-scale sensor network experiments

Sensor networking has demonstrated great potential in many areas of scientific exploration, including environmental, geophysical, medical,and structural monitoring. However, sensor networks have largely beenfocused on dense, small-scale deployments to monitor a specific physical area. GENI offers the opportunity to bridge across multiple discrete sensor networks to provide monitoring of physical phenomena at a global scale. In addition, GENI can provide the infrastructure for querying and fusing data across multiple (possibly overlapping) sensor networks in different scientific and administrative domains.

As an experiment to demonstrate these ideas in the context of GENI, we propose building an infrastructure to interface real-time sensor data into large-scale distributed setting, focusing on collecting and processing large volumes of sensor data within the GENI network. This would involve developing an appropriate overlay network infrastructure for querying individual sensor data sources; constructing flows of sensor data through multiple stages of filtering, processing, correlation, and aggregation; and delivering the resulting data to the end user. Pushing computation into the overlay network can distribute server load and reduce network bandwidth requirements, thereby increasing scalability in terms of the number of data sources and simultaneous queries that can be supported.

A rich application domain for this approach is geophysical monitoring. Examples include monitoring seismic activity along fault lines and at volcanoes, and GPS-based geodesic measurements of plate movements. The NSF *EarthScope* initiative is building the sensor infrastructure, and tying this source of data into GENI would enable a 'continental scale sensor network' supporting a wide range of real-time geophysical monitoring applications.

This experiment would provide the opportunity to bridge between sensor networks and the rest of the GENI infrastructure in the context of a realistic application. Doing raises a number of architectural issues, such as:

- Discovery and addressing of sensor data sources;
- Design of overlay networks to support queries over high-volume streaming data, including optimal placement of query operators within the overlay;
- Reliability and data fidelity in long-running queries subject to intermittent failures of data sources and overlay hosts;
- Load balancing and migration of computation within the overlay network, as well as from the overlay to the underlying sensor networks;
- Optimizations across multiple queries within the overlay, e.g., combining computation for similar subqueries.

Studying these problems will undoubtedly impact the overall GENI design as it pertains to support for sensor networks as well as large-scale stream data processing.

This experiment will require two primary components: (1) seismic and geodesic sensors deployed over a large geographic area; and (2) a number of overlay hosts in the GENI infrastructure for collecting and processing sensor data streams. In terms of sensor requirements, we expect that several (two to five) arrays of 10-100 nodes each can provide the data sources for this experiment. It may be possible to tap into existing or planned EarthScope sensors for this purpose, or else deploy specially-tasked sensor arrays, e.g., on a fault line or active volcano, to collect interesting data. Each sensor array will require connectivity to GENI through a wired, long-distance wireless, or satellite uplink, depending on the deployment location.

Collecting and processing data from these sensors will require on the order of 20-100 overlay network hosts within GENI, each capable of performing data filtering, signal processing, and aggregation. In addition, services are required for archiving sensor data, supporting queries on archived data, and visualization/presentation of results to end users.

Of course, these resources do not need to be (and in some sense should not be) unique to this experiment. Much of the planned GENI infrastructure can be leveraged to develop this system, as well as existing sensor emplacements (notably EarthScope).

The scope of this experiment is fairly broad and the evaluation goals encompass the following metrics:

- Scalability - in terms of the number of users and number of sensors;

- Latency - the ability of the infrastructure to deliver results to end users in a timely fashion.
- Robustness - ability of the infrastructure to respond to failures of network hosts and links;
- Fidelity - the accuracy and precision of results returned from queries spanning many (possibly faulty) data sources.

Evaluations should be based on fairly realistic queries defined by end-user needs. Working closely with experts in the relevant application domains is essential to this goal. For example, a scientist studying teleseismic signals may have greater interest in the timing precision of signals rather than the actual number of sensors that respond to a query. Understanding these requirements will undoubtedly shape the goals of the GENI architecture..

Scientific study based on large-scale monitoring of geographically diverse data sources is an emergent area. Only recently, with the advent of high-speed networks, has it been possible to perform such studies in near-real-time. Our current Internet infrastructure provides the means to connect sensor data sources with users and software services for archiving and processing that data. However, no framework exists for deploying, managing, optimizing, and sharing a computational infrastructure across experiments and different scientific domains. As a result, the ability of an individual researcher, or even an entire scientific community, to build the appropriate infrastructure to undertake this kind of scientific work is severely limited. GENI offers the potential to enable entirely new kinds of distributed services, as well as to greatly simplify their deployment. As a result, we expect GENI to foster a revolution in scientific progress similar to that enabled by the current Internet.

2.5.3 Emergency response communications with adaptive ad-hoc networks

This experiment will enable and test deployment of adaptive network layers for cognitive radios in an emergency response environment. The use case is a hurricane relief effort, with four distinct environments: a Command Center with extensive wired connectivity and wireless for rapid growth, an Impacted Area with all connectivity provided by wireless units deployed after the incident, an Evacuee Center with rapidly growing center with changing requirements met by cognitive radio deployments, and a Relief Relay Center with a quickly deployed center with mix of wired and wireless to catalog and direct relief. Each environment has its own messaging requirements that will drive the specific experiment, and the experiment will explore deploying protocols for adaptive communications depending on messaging needs.

This experiment will analyze and develop applications suited to emergency communications that take advantage of flexible radio and network protocols to allow for tailored communication environments. Multiple network layer protocol implementations will be developed and mechanisms to allow researchers to select and utilize a particular network layer on a per application basis, over the cognitive wireless radio network, will be created. The experiment will allow networking researchers and emergency responders to evaluate models including multiple preemptible one-to-one communications and group messaging of key information, public announcements, etc. It will improve understanding of how to configure cognitive radio stacks based on communication model (one-to- one, one-to-many, bandwidth, density) as well as physical environment. It will also improve understanding of how to adapt cross-layer (PHY/MAC/network) cognitive radios to changing application needs in a rapidly changing wireless environment. It will lead to software applications that can choose the most appropriate network stack based on the experimental results, given a particular application, layer 1 and layer 2 environment, and constraints/guidelines such as: best overall coverage, lowest power (preserve batteries), etc. It will lead to creation of the tools necessary to allow security and policy constraints to guide cognitive radio deployments that interact with end-to-end applications.

An expected outcome from this experiment is the development of protocol stacks that are aware of layer 1, layer 2, network, and application constraints and requirements, and self configure to deliver the application the most appropriate communication network. A societal objective is that it will increase the ability of early responders and longer-term support to reach victims.

This experiment will require 3-4 GENI access points and 5-10 wireless subnets of cognitive radios. GENI uniquely provides geographically separate but linked test environments which can be provide a realistic test system. In the experiment scenario, the Impacted Area will have multiple wireless subnets and be the target for point to multipoint announcements, and point to point coordination. The Evacuee Center will utilize cognitive radios' to provide coverage and access to a growing population, expected communication flows of broadcasts (finding relatives) and point to point communications (requesting supplies etc), will transition to a more traditional mix as the center works to provide longer term communication needs (telecommuting, etc.). The Relief Relay Center will utilize wired and wireless capabilities with mobile scanning and tracking devices (RFID, etc) to catalogue shipments of supplies, and provide directions to get the supplies to the best locale. The Command Center Setup will utilize wired and wireless systems to receive updates from the impacted area on needs, available routes, local conditions etc, and then coordinate the delivery of supplies via the relief tracking center, and the initiation of rescue missions based on information received from the Impacted Area (broadcasts for help) or specific leads provided by people at the Evacuee Center.

For each environment and application the infrastructure must measure at each node parameters such as receiver signal and noise, power utilization, connectivity stability, delay to receive communication, and available application bandwidth. Measurements such as these and surveys of emergency responders can be used to determine effectiveness of particular architectures.

2.5.4 End-to-end integration of cognitive radio subnetworks with global Internet

This experiment is designed to investigate and validate concepts for interfacing cognitive radio subnetworks with the global Internet. Cognitive radios represent a new kind of Internet device (router, end-ser terminal) with a high-degree of adaptation in terms of PHY/MAC and network topology, potentially requiring new protocols and algorithms for efficient addressing and routing.

Specific areas of research focus include:

- Identification of addressing structures that support ad-hoc organization
- Design of hierarchical aggregation methods for addressing and routing
- Specification of control interfaces for routing metrics and QoS control
- Evaluation of end-to-end cross-layer support requirements.

The experiment will require 2 or 3 cognitive radio subnetworks (each with ~5-10 mobile end users) connected via different GENI access nodes, a cognitive radio protocol stack running on GENI cognitive radio platform, and software for the aggregated cognitive radio / Internet routers, and protocol modifications at corresponding hosts, and spectrum measurement capabilities deployed within the cognitive radio subnetworks.

Experimental measurements that will be needed include cognitive radio subnet parameters (throughput, spectrum efficiency, intranet routing overhead, etc.), network layer parameters such as efficiency of routing and addressing schemes, cross-layer control and processing overheads vs. performance gains, end-to-end performance at the network, transport, and application layer, and long-term service stability versus degree of adaptation.

2.5.5 Securing the interface between wireless networks and the future global Internet

The objective of this experiment is to understand how to secure the interface between wireless networks and the future global Internet. Wireless networks represent a significant point of entry for launching security attacks against the global Internet. This experiment will help analyze and classify threats,specifically:

- Identify what types of security attacks can be launched from wireless devices on each other
- Identify what types of security attacks can be launched from wireless devices on other edge infrastructure
- Identify what types of security attacks can be launched from wireless devices on the core network
- Differentiate between wireless threats and conventional threats
- Develop suitable metrics to quantify the relative severity of different attack modalities on network functionality as well as on various types of applications running on users across the network
- Develop a suite of security solutions to address different classes of threats. For client-to-client threats, solutions must operate locally and might have to consider efficiency or resource issues (e.g. sensor networks). For client-to-edge infrastructure threats, asymmetric assignment of security responsibilities may be delegated to clients and infrastructure. For client-to-core threats, solutions can operate within the core itself.
- Develop a suite of security solutions for different wireless threats
- Identify a prioritization of security threats based on severity level and the practicality of their resolution

This experiment will require wireless experimental networks: a city-wide mesh/ad hoc network, and a sensor network. Both networks will interface with the broader GENI facility to facilitate wireless-to-wired security evaluation. Application services will run from wireless-to-wired-to-wireless (e.g. client accessing content delivery service located on the core network), and wireless-to-wired (e.g. sensor network supplying information for a monitoring application). Adversarial nodes will have to be developed that can monitor wireless traffic and inject arbitrary format-compliant data into the medium. Denial-of-service attack modules will need to be developed , as will spoofing and device masquerading tools. Tools for monitoring internal buffer states at wireless nodes and wired nodes on the network will needed to quantify the severity of denial-of-service attacks and the efficiency of packet filtering mechanisms. Cryptographic libraries will need to be implemented (starting from current publicly available libraries) in order to test various solutions. Measurements will also need to be conducted to evaluate threat severity.

2.5.6 Evaluating global location services for the future Internet

The objective of this experiment is to evaluate the efficiency, accuracy and overhead associated with alternative methods for offering global location services in the future Internet. Options that might be evaluated include overlay architectures as well as more integrated layer 3 type solutions associated with the routing protocol.

The infrastructure needed for this experiment includes existing wireless technology, and the deployment community should be a relatively sophisticated in order to have existing demand and usage patterns. The platforms need to be management-capable devices as well as emerging technologies that allow us to "time travel" and evaluate ideas for the future. Detailed instrumentation is needed, as many experiments need to understand the behavior of the system at a low level (physical signals).

2.5.7 Transformational health-care applications in urban mesh networks

The objective of this experiment is to enable and test deployment of pervasive high performance wireless and understand operational issues. It would study and develop applications with potential for high societal impact that are enabled by pervasive high performance wireless, study the impact of pervasive high performance wireless on healthcare and quality-of-life issues, and evaluate societal benefits such as reduced emergency room visits in low income neighborhoods, increased access to health-care educational resources and quantifiable cost and health benefits to community.

The infrastructure needed for this experiment is coverage over a reasonable area, perhaps greater than 2 sq. km. with broadband wireless access, and deployed protocols with predictable, reliable, resilient performance. The team for this experiment needs to include networking researchers, medical practitioners, and pervasive computing experts. Participants need to be equipped appropriately; for example, nurse practitioners with PDAs and video uplinks. Evaluation would employ surveys, usage studies, and measurements (over a multi-year period).

2.5.8 At-scale studies of emerging and disruptive wireless technologies

The objective of this experiment is to study emerging wireless standards in deployed and operational network as well as co-existence and interoperability with popular technologies. The experiment is based on deploying the community's advances in wireless data rates (approaching 1 Gb/s) years before commodity hardware arrives, potentially speeding up standards and architectural advances. This will allow fundamental performance studies of disruptive technologies, for example providing network traces at multiple temporal granularities.

This experiment will need programmable wireless kits in deployed networks in conjunction with existing hardware (say in mesh networks). The infrastructure must support transparent switching between the experimental and existing network in order to operate the emerging networks with real traffic users with safeguards against failure during testing.

2.5.9 Cognitive radio system evaluation

The objective of this experiment is to understand how to minimize use of constrained radio resources while maintaining effective and robust communication. It will attempt to understand how to predict the behavior of many individual radios operating in collaborative and non-collaborative environments, and how to design individual radios so that when placed into a collaborative radio network a desired behavior emerges (e.g. robust communications, or power efficient communications).

This experiment will also seek to understand and explain the different experimental methodologies we use: analysis, simulation, emulation, and field trials. Each has advantages and disadvantages and we should understand the appropriate application of each methodology and how the methodologies are inter-related. For example, field trials are good for collecting observations of physical phenomena that are useful in emulation and simulation, but are difficult to exactly repeat and are complex and expensive to set up and run.

This experiment will require a cognitive radio subnetwork, and facilities to measure relevant parameters. For example, measurement capabilities need to be able to understand whether Algorithm A or Algorithm B adapts faster, uses fewer radio resources, uses less power, etc., whether radio network architecture A or B (e.g. cellular versus mesh) use fewer radio resources, or whether approach A or approach B provides more accurate/robust location information. Infrastructure capabilities will also be needed to allow for collection of network traces for off-line analysis.

3. System-level outlines of GENI wireless subnets

The proposed GENI infrastructure will be designed to incorporate a wide range of wireless networking capabilities in order to provide experimenters with access to emerging radio technologies which are becoming increasingly important at the Internet edge. The design objective for GENI 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. Specific wireless capabilities to be included in GENI include integration of large-scale emulators, a real-world large-scale 802.11-based mesh network, an open API “4G” cellular/WiMax system, a cognitive radio technology demonstrator, and 2-3 large scale sensor net deployments (see Fig. 3.1 below).

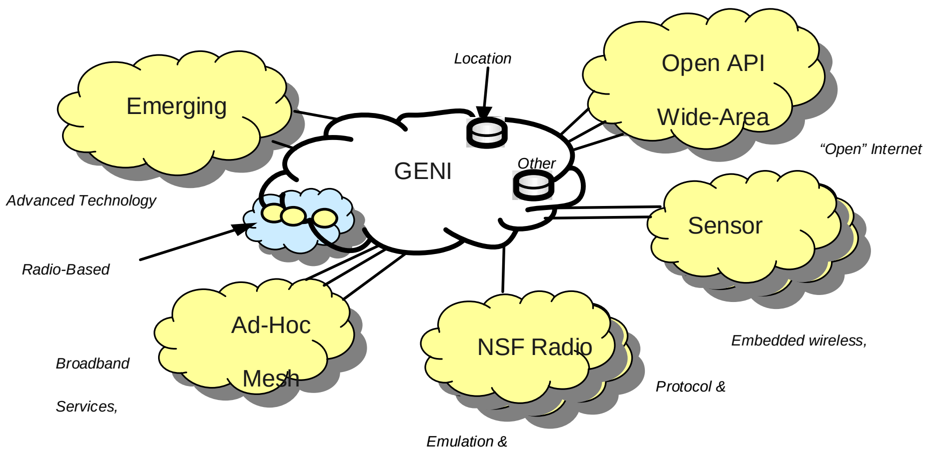

Fig. 3.1. Experimental Wireless Networks in GENI – Overview

High-level descriptions of each of the planned wireless GENI subnets are given in Secs. 3.1-3.5 below.

3.1 Wireless Emulation Subnetworks:

Wireless network emulation environments that provide facilities for repeatable protocol validation studies of a quantitative nature are viewed as an important feature of GENI (see Fig. 3.2 below). Because the cost of conducting a large-scale real-world wireless experiment involving sensors, vehicles, mobile terminals, etc. is quite high, most experimenters will benefit from a phased approach in which early protocol concepts are validated at scale in a controlled emulation environment before being migrated to a long-running field trial with real users. Availability of emulators will also help reduce experimental traffic load on GENI core and access networks, an important consideration given practical bandwidth, processing and slice constraints. A number of wireless and sensor network emulators have been developed over the past 5 years with NSF funding, including Emulab at Utah, ORBIT at Rutgers, and the Kansei (Ohio) and Harvard sensor net labs. It is recommended that at least one large scale emulator of each type (wireless, mobile, sensor net) be incorporated into GENI, with sufficient scale and realism to cover anticipated future experimental research needs. Integration of these testbeds will require support for GENI's standard virtualization model, as well as support for GENI's control and management interfaces details of which are yet to be determined.

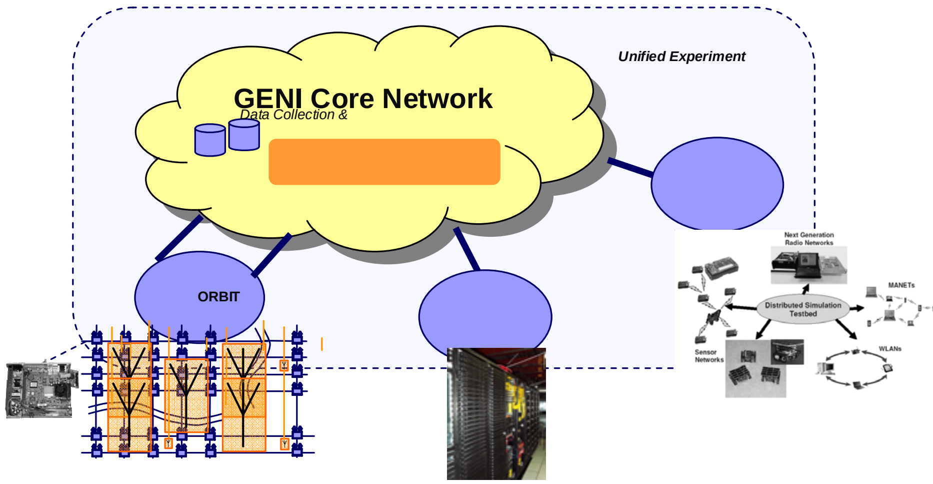

Fig. 3.2. Integration of Existing NSF Wireless Testbeds into GENI

3.1.1 Support for Virtualization:

Wireless network emulators in GENI are expected to support virtualization that permits more than one experiment to run simultaneously on the same physical infrastructure. Virtualization capabilities can be added using appropriate combinations of techniques such as "virtual MAC" (VMAC), space division multiple access (SDMA), frequency division multiple access (FDM) and time division multiple accesss (TDMA). Specific techniques to be used will depend on the granularity with which slicing is to be carried out, and the type of experiment (short-term repeatable vs. long-term service experiment) to be supported. A global radio resource measurement and management framework is required for implementation of virtualization schemes across large numbers of radio nodes. More details on virtualization are given in Sec 4.3.

3.1.2 Control and Management:

The emulation and simulation testbeds should be integrated in a seamless manner with the GENI infra-structure, and the key to the integration is a common control and management architecture. Existing emulators such as Emulab or ORBIT will require a retrofit of their central experimental control software to interface with the GENI control and management standard. The GENI control and management protocol specification will need to incorporate features for specifying wireless subnet attributes such as network topology, node location, frequency or power. Further discussion on the control and management requirements for wireless can be found in Sec. 4.4.

3.2 Urban Mesh/Ad-hoc Network:

The proposed urban mesh/ad-hoc network is designed to support real-world protocol experience with emerging short-range radios – see Fig. 3.3 below. 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 the GENI system 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 in GENI 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 will be installed at ~8m height using available lighting poles or other utilities, and will have electrical power and a wired or wireless (i.e. cellular) remote management interface. Approximately 25% of the nodes will be designated as access points with wired interfaces (typically VDSL or fiber) to GENI access routers. Preferably all radio nodes (not only AP's) should have a wired control and management interface to support code downloading, remote monitoring and experimental data collection functions. Nodes in the network will be programmable in terms of layer 2 and layer 3 protocols and will support applicable forms of virtualization corresponding to specific platforms used. The urban deployment will also include support for ~100 vehicular mobile nodes associated with a suitable bus or taxi fleet operating in that environment. End-user mobile nodes can also opt into the GENI urban mesh if they adhere to the specified control and management protocol. Note that wireless control and management interfaces such as GPRS or EVD0 will be required at vehicular or mobile nodes, limiting the degree of programmability and remote data collection.

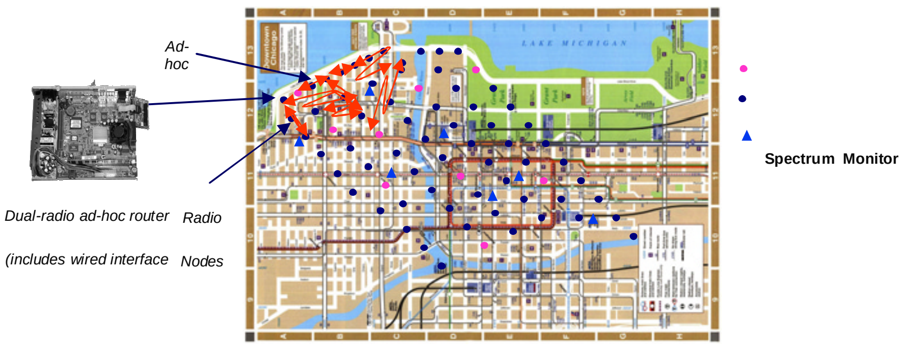

Fig. 3.3. Deployment Concept for Ad-Hoc Mesh Network in GENI

3.2.1 Support for Virtualization:

Virtualization can be provided in the urban mesh network using methods described in Sec. 4.2. VMAC may be used for certain classes of long-running applications that do not require strict partitioning of resources, while combinations of SDMA, FDMA and TDMA are applicable to shorter term experiments involving tighter timing constraints. It is noted that the system will require an independent radio resource monitoring infrastructure in order to support the global algorithms for allocation of resources to different slices.

3.2.2 Open Platform:

The nodes in 802.11-based Mesh/Ad-hoc networks should be built on open platforms so that it is possible not only to manipulate the parameters of wireless node, but also to modify the lower-level protocols and algorithms, such as, altering the MAC protocols (for example: SoftMAC). Thus, an "open" 802.11x chipset and related firmware and drivers are an important technology component required to build the wireless platform. This will require appropriate open source licensing arrangements with one or more 802.11 chip vendors.

3.2.3 Support for Mobile Nodes:

As mentioned above, the urban deployment will include support for ~100 vehicular mobile nodes, implemented by leveraging either a bus or taxi fleet that operates in that area. The mobile node will also follow the standard GENI control interface, and will allow for remote control (rebooting, etc.) via a cellular data interface (such as GPRS or EVD0).

3.2.4 Control and Management:

The urban mesh deployment will support to-be-determined GENI control interfaces for programming, resource allocation, slice management, monitoring, etc. Wireless nodes in the field will be accessed via an out-of-band control channel, either high-speed wired access where available, or cellular data access such as GPRS. In view of variations in wireless node connectivity and control bandwidth, a wireless network control gateway approach may be used to provide proxy experiment management services and global radio resource control for the entire deployment. The radio resource manager will check for compatible allocation of frequencies, power, etc. across nodes and slices in the network and will be supported by an independent monitoring infrastructure. An appropriate number of RF monitoring devices (~25-50) will be deployed collocated with some of the wired access point sites where control bandwidth is considered plentiful.

3.2.5 Wireless Kits:

The urban grid deployment described is limited to one major metro area due to cost constraints. We believe that proliferation of wireless access networks in GENI can be encouraged by providing a kit consisting of an open API GENI compliant AP, radio forwarding node and mobile user node. These can be provided by making a suitable agreement with the equipment vendor for wireless platforms used in the deployment.

3.2 Wide-area Suburban Wireless Network:

The GENI system will include a wide-area suburban wireless network with open-access cellular 3G/WiMax radios for wide-area coverage along with short-range 802.11 radios for hotspot and hybrid service models – see Fig. 3.4. 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. GENI will include one or more wide-area wireless experimental networks with ~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 are installed at ~8m on utility poles, etc. This wireless subnet will also include provision for ~25-50 vehicular mobile nodes as part of the deployment, and will also support user-supplied mobile devices which conform to the GENI control interface. 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. Control and management of all radio nodes will use an out-of-band channel, either wired backhaul where available at base stations and AP's, or cellular GPRS/EVD0 at mobile nodes or forwarding nodes without wired connections.

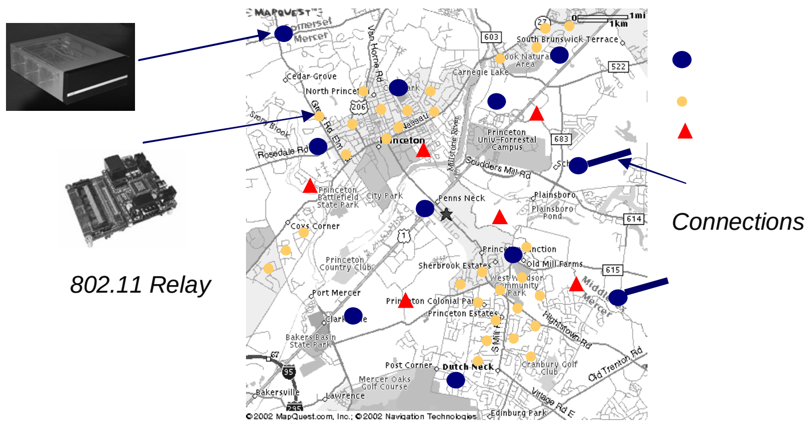

Fig. 3.4. Deployment Concept for Open API Wide-Area Wireless Network in GENI

3.3.1 Support for Virtualization:

Virtualization of cellular or WiMax base stations can be achieved using circuit-mode assignment supported by these standards, in addition to FDMA and SDMA methods applied across multiple nodes. As in Sec 3.2, 802.11 forwarding nodes and AP's can be virtualized using VMAC, or combinations of SDMA, TDMA and FDMA. An independent radio resource monitoring infrastructure consisting of ~10-20 nodes will be used to support global radio resource management in this network.

3.3.2 Open Platform:

The design goal for open cellular or WiMax base stations is to provide fully programmable layer 3 (routing) functionality with as much layer 2 (DLC, MAC) flexibility as possible. Open layer 3 can be achieved with existing cellular or WiMax technologies, but layer 2 programmability will require an open-source agreement with one or more PHY/MAC chip vendors. This is an open issue requiring further study on what is feasible and practical during the GENI deployment time-frame.

3.3.3 Support for Mobility:

Just as there will be static (non-mobile) OFDM/CDMA/WiMax nodes mounted on fixed infra-structure, there will also be ~25-50 mobile nodes with OFDM/CDMA/WiMax mounted on mobile nodes (such as, buses, taxis, other vehicles) that move around, in an area covered with 802.11 and/or OFDM/CDMA/802.16 radios. Mobile nodes will support the standard GENI control interface accessed through a cellular data (GPRS, EVD0) interface.

3.3.4 Control and Management:

The suburban wireless deployment will support to-be-determined GENI control interfaces for programming, resource allocation, slice management, monitoring, etc. Wireless nodes in the field will be accessed via an out-of-band control channel, either high-speed wired access where available, or cellular data access such as GPRS. In view of variations in wireless node connectivity and control bandwidth, a wireless network control gateway approach may be used to provide proxy experiment management services and global radio resource control for the entire deployment. The radio resource manager will check for compatible allocation of frequencies, power, etc. across nodes and slices in the network and will be supported by an independent monitoring infrastructure. An appropriate number of RF monitoring devices (~10-20) will be deployed collocated with some of the base station sites where wired backhaul of control data is available.

3.3.5 Wireless Kits:

As mentioned in Sec 3.2.5, suburban wireless network “kits” will be made available to the community to enable additional deployments based on user-supplied equipment. In this case, open cellular/WiMax base stations, open 802.11 AP’s and vehicular nodes may be provided by the equipment vendor(s) chosen for this subnetwork.

3.4 Sensor Networks:

The GENI deployment plan includes sensor networks capable of supporting research on both protocols and applications – see Fig. 3.5. Since the design of a sensor network tends to be somewhat application specific, GENI will provide experimental access to a small number of selected sensor net deployments which leverage the urban or suburban wireless infrastructure described in secs 3.2,3.3. There will also be a "sensor deployment kit" consisting of network gateways (from sensor radios to 802.11 or cellular), sensor modules and related platform software to enable users to build additional sensor networks for specific research or application objectives. 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.

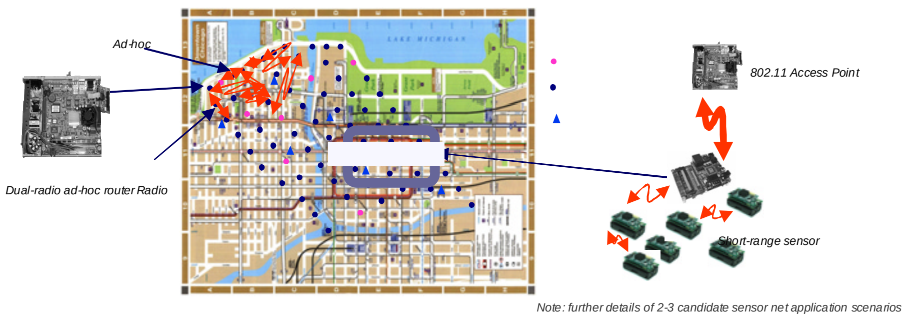

Fig. 3.5. Deployment Concept for Sensor Networks in GENI

Specific sensor deployments in areas such as environmental monitoring, security, traffic control, vehicular safety or smart spaces will 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 10's of 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 and virtualization used in GENI, 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.

We envision two broad classes of wireless sensor testbeds in GENI:

3.4.1: Indoor Testbed

These testbeds each operate on a scale of a few hundred nodes deployed in either a very dense laboratory setting or more widely distributed throughout a building. Each sensor is typically attached to a serial or network interface for programming and debugging, and is permanently powered. A front-end interface provides scheduling, programming, debugging, and logging capabilities. These testbeds will be connected more directly to applications through coordination with domain experts in their design and deployment.

A typical indoor testbed will consist of 500-1000 sensor nodes such as the current MicaZ or TMote Sky, equipped with an embedded microcontroller and 802.15.4 radio. Each unit will be connected to an Ethernet or serial/USB backchannel for programming, debugging, and data logging. Each sensor will be interfaced to a suite of sensors, depending on the application tie-in focus of the testbed; one testbed can also support nodes with varying sensor types. Each testbed will include the order of 20-40 more powerful gateway devices, such as embedded PCs, permitting studies of multi-tier architectures, and which act as local aggregation points. Finally, each testbed will have a front-end server providing a user interface, scheduler, and multi-terabyte database for logging traces collected from the network.

3.4.2 Urban-scale testbeds

An urban sensor testbed will consist of 200 or more nodes mounted in an outdoor, urban setting, such as on the sides of buildings or on utility poles. Each node will be permanently powered and will consist of an embedded PC-class device with a moderately powerful radio interface (e.g., 802.11a or g). Multiple nodes will have a wired Internet connection, acting as gateways between the testbed and the rest of the GENI infrastructure. The remaining nodes will rely on wireless communication.

We also anticipate that these urban testbeds will also serve as a "backbone" providing connectivity and aggregation points for colocated fixed or mobile sensor deployments. For example, a vehicular sensor network experiment could use the urban testbed for delivering data from mobile sensors to a central server, as well as tracking locations of vehicles without on-board GPS.

3.4.3 Support for Virtualization

An important aspect of building open testbeds is to allow sharing of resources where sharing can be done in both space- and time-dimesions depending on the requirements of the experiments.

Different projects have proposed varying approaches to resource management in a shared testbed environment. MoteLab makes use of a simple time-slot scheduler with a per-user quota on outstanding jobs. The Intel Mirage testbed requires that users "bid" on needed resources in a virtual market environment to support differentiated service during periods of contention. Each of these approaches is tailored for somewhat different usage scenarios and degrees of user sophistication striking the right balance between ease-of-use and flexibility for advanced experimentation.

3.4.4 Sensor kits

Realistic applications and application traffic are one of the most important things sensornets bring to GENI. Augmenting traditional testbeds, we see sensornet "kits" as essential to encourage a diverse set of realistic applications. Sensornet kits have two components:

1) A common set of designs for a sensor platform, a standard set of sensor packages, and the systems software necessary to make this hardware reasonably easy to use.

a) The sensor platform is a simple processing and communications module, possibly in a two or three flavors with different capabilities. An example might be today's MicaZ and Intel Stargates. There would be a standard sensor interface board with analog and digital I/O capabilities that will enable easy addition of other sensors as required for unanticipated or specialized applications.

b) There will be a few default sensor packages for applications of known interest (examples might include seismic monitoring, habitat monitoring, and simple localization and tracking).

c) The software platform should be open source with minimal licensing restrictions (BSD-style license or similar), with a group that maintains the point and coordinates bug reports, fixes, and releases.

2) Approximately 50 "sensornet patches" at diverse sites that have a moderate number (20-100) of these devices, connected to the GENI backbone either through a regional sensornet testbed (described below) or connected directly via a gateway host.

The gateway host that connects a sensornet patch to GENI will be responsible for allowing sensornet traffic into the broader GENI network, and for mediating use of the patch by external researchers.

3.4.5 Control and Management:

Consistent with the urban and suburban wireless deployments mentioned earlier, the sensor net gateways are expected to support the to-be-determined GENI control protocol and serve as proxy nodes for individual sensors in an experiment. The sensor gateways will interface with global radio resource control managers in each deployed network in order to coordinate frequency selection, etc. The global manager will identify unused sensors for "SDMA" assignment to a new slice when an experiment is being set up. As for the other wireless devices discussed earlier, each sensor gateway will have support for an out-of-band wired or wireless control channel for control and measurement data.

3.5 Cognitive Radio Networks:

The cognitive radio network deployment in GENI planned as an advanced technology demonstrator with focus on building adaptive, spectrum-efficient systems with emerging programmable radios – see Fig. 3.6. 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 GENI in the 2008-09 time-frame. Research teams are interested in investigating radio spectrum architectures, hardware platform (SDR) integration, cognitive networking adaptation algorithms, mobile/wireless network control and management functions, developing new protocols supporting new network services, and are further interested in testing and evaluating their work in both simulated/emulated and realistic environments. In order to perform experimentation on architectures that inherently build upon larger network services, the CR infrastructure will be integrated into and interoperate with the larger GENI infrastructure. GENI will include a cognitive radio deployment 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 from FCC will be required to support this trial network.

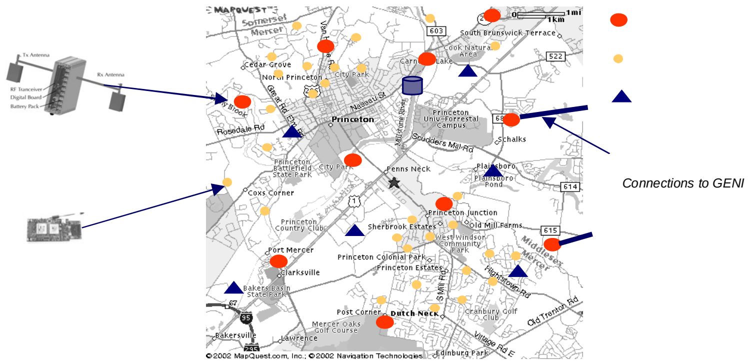

Fig. 3.6. Deployment Concept for Cognitive Radio Demonstrator Network in GENI

3.5.1 Cognitive Network Components:

The Cognitive Radio Infrastructure consists of three major components. These components are selected to support a wide range of research teams interested in radio networks, including those interested in the physical, data link, media access, resource allocation, network scaling, resource allocation, inter-operability, and content delivery aspects, among many other possible topics. The three components are:

1. *A CR Test Environment* – One or more facilities to extensively test and evaluate innovative radio networks in a repeatable manner.
2. *CR deployment in a typical urban/surban environment* – One or more facilities to test and evaluate Cognitive Radio networks in real-world situations.
3. *Experimental Platform Kits* – Easy to construct CR kits to expand the number of research teams working with the technology and developing network architectures and designs. (see Sec xx for further details)

3.5.2 Control and Management:

Cognitive radio provides programmability of the radio PHY and MAC functions, thus increasing the complexity of the control and management interface relative to ASIC-based industry standard radios discussed earlier. The GENI control and management interface needs to have extensibility features to support future cognitive radio control and measurement requirements as they become fully understood. An out-of-band control channel (preferably wired, but potentially using a broadband virtual channel in the cognitive radio band itself), will be needed to support control and measurement data without interfering with experiments themselves. In addition, a new class of agile, wideband spectrum monitors will be needed as part of the measurement infrastructure.