

An Interdisciplinary Collaboration on Performance Aspects of a High Performance Wireless Research and Education Network

Project Description

Introduction

A most important and daunting task for the evolving Internet, including its HPC component, remains the attainment of ubiquity. This issue is often discussed in terms of global ubiquity, with a common focus on third world countries. While global ubiquity is important, the same problem persists much closer to home: in rural America. About 20 miles or so from cities, and some distance from central offices of telephone companies, any notion of High Performance evaporates on dial-up modems. For researchers in the field, that kind of real-time data communication remains unattainable. The coverage of cellular phones and pagers in rural areas, even those frequented by people (e.g., between Ramona, via Julian, to the Anza-Borrego Desert) is in an embarrassingly marginal state, where currently deployed commercial infrastructure often resembles an illusion of reachability. Of course, there is no economic incentive for commercial service providers to change that, let alone provide high performance Internet access.

Available solutions for this situation are either too costly, or too far into the future. Geostationary satellite networks require a significant startup cost, and the services are expensive to maintain. Low and medium earth broadband satellite systems are still years away from implementation, much less wide-spread use. At the same time, the research community has immediate connection needs for researchers living in rural areas, researchers in the field, autonomous telemetry sensors, as well as education programs, at reasonable performance levels to the Internet.

At a recent workshop of the national Ecological Observatory Network (NEON) a researcher commented on about 10 million species on earth visible to the human eye, and that there will never be enough researchers to investigate, let alone understand, them all. The only choice is to improve the efficiency and productivity of the researchers we have, and widen the possible areas and scope of their work. This includes real-time access to their data and sensors in the field, access to network resources (such as computing, data storage, visualization), and communication capabilities to exchange findings with colleagues in real time.

To allow for this, we have to understand the networking and performance needs of disciplinary researchers, as well as the parameter space for implementation opportunities. An obvious path to pursue is to work with disciplinary researchers to create environments that fulfill their networking needs. Examples of such researchers are those running remote telemetry machines and sensors (e.g., earthquake sensors), those working in remote observatories (e.g., creating space imagery), and those working in the field (e.g., biological field stations). An example of communications needs not being met is that the only data link from San Diego State University's observatory on Mt. Laguna is a dial-up modem. Another example is the need to instrument earthquake faults in the desert, which have not gone off in hundreds of years, and when (not if!) they do, will create major problems across wide areas, which may be heavily populated. Connecting the desert with sensors (see Figures 1 and 2) that require real-time data communication is less than straightforward, but with the resultant improvement

in public safety and the ability to assess evolving catastrophic emergency conditions, it is well worth the effort, and quite inexpensive when compared to the alternatives.



Figures 1 and 2

Remote sensors in desert regions constitute a challenge for the distribution of telemetry data. The image on the left shows an underground earthquake sensor, with the power generation and antenna being mounted on a pole above ground. The image below shows a central point of an earthquake sensor array, the vertical component being below the several steel boxes mounted in the desert sand.



The performance assessment activity will be leveraged with the measurement and network analysis activities of the National Laboratory for Applied Network Research (NLNR), with an immediate application for geophysics and astronomy disciplinary research, including an infrastructure for the Scripps Institution of Oceanography (SIO) earthquake sensors and astronomy observatory connectivity. Other short term objectives include the needs of the ecology and environmental science community. Hans-Werner Braun of NLNR (who has also been a PI for the NSFNET backbone previously) is the Principal Investigator for this proposal, with SIO geophysicist Frank Vernon being Co-PI.

Overview

Over the last several months, multiple departments at University of California, San Diego (UCSD) began to collaborate on creating a prototype implementation for wireless connectivity. This includes the Measurement and Network Analysis Group of the National Laboratory for Applied Network Research (NLNR) and the San Diego Supercomputing Center (SDSC), the Scripps Institution of Oceanography (SIO), as well as the School of Engineering and its Center for Wireless Communications (CWC). The early result of this collaboration is a working, approximately 20-mile connection, between UCSD and Ramona, via an intermediate hop on a commercial microwave tower on Mt. Woodson (Figure 3).

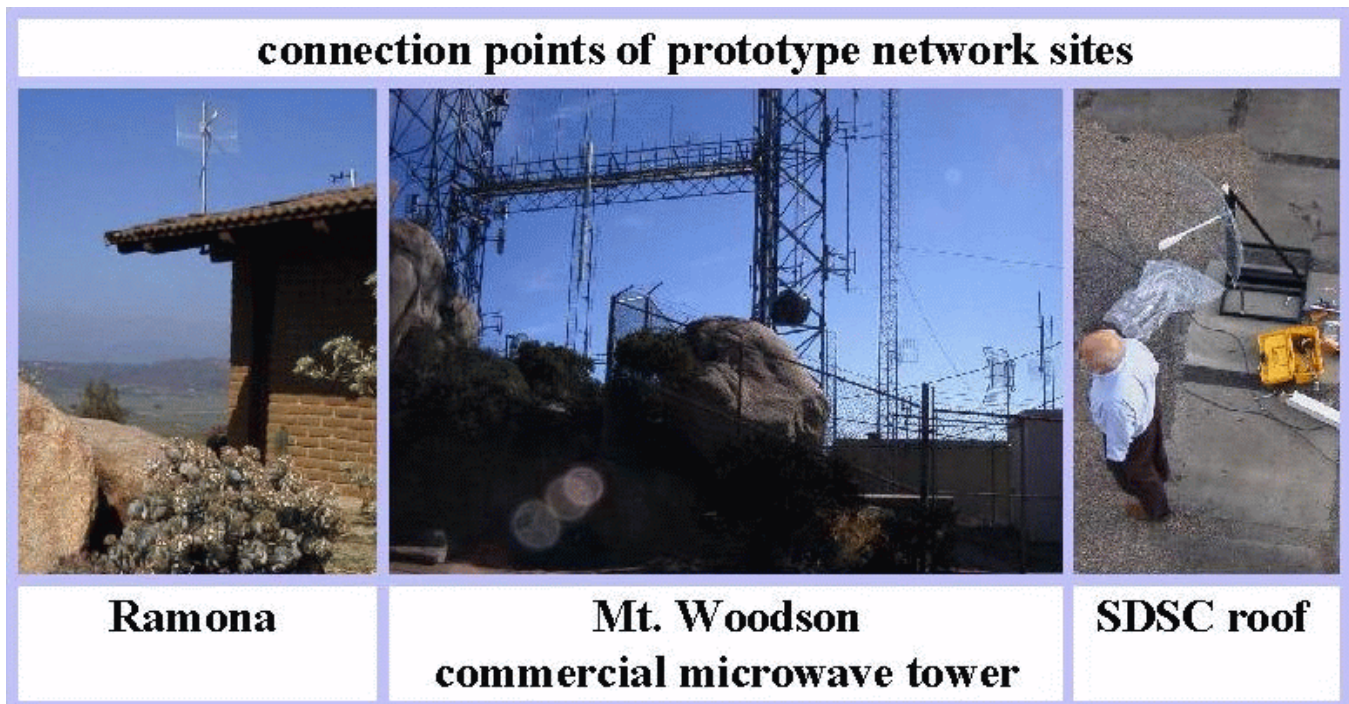


Figure 3

Outside collaborators are already emerging, they include the astronomy department at the San Diego State University (SDSU) for their observatory on Mt. Laguna, ecology and environmental researchers with remote site and sensor needs, interest in connecting educational facilities located on Native

American reservations, and possibly, Caltech astronomers for the observatory at Mt. Palomar. We intend to remain open to new collaborative activities throughout the lifetime of the project; at this time, it is unknown what other interesting opportunities may develop.

This three year proposal requests funding to expand upon the prototype connection, and create a substantial wireless backbone network to provide access for research and education applications in San Diego County, and to allow assessments of the technology and performance parameters to be determined for such an environment.

Hence, while it will provide services to researchers and telemetry stations in the field and a delivery mechanism for distance education in disadvantaged areas, it will function primarily as an applied test bed to address distance access issues over a relatively large rural area in general, and to assess performance characteristics of such a network.

The NLANR Measurement and Analysis Group has created and continues to build upon a network analysis infrastructure to support network research and engineering of high performance research networks (<http://moat.nlanr.net>). This includes a passive monitoring project and an active monitoring project; the experience and expertise gained from this will be used to build the wireless network mesh, as well as to examine issues to be addressed in an integrated system. The following are some of the issues in an integrated system that we want to address:

- system performance assessments and metrics
- how to centrally control access (MAC filters?)
- how to prioritize traffic (at RF level, based on MAC addresses?)
- aspects of utilizing unlicensed spectrum
- how to integrate new and interoperable technologies
- access to microwave towers to create the backbone
- RF radiation concerns
- integration of capillary nodes (e.g., laptops)
- integration of data generators (e.g., seismic sensors)
- integration of a campus infrastructure (e.g., UCSD's wireless trial)

Multiple Applications for a Variety of Disciplines: Plans and Opportunities

There are a number of scientific disciplines with remote area communication needs. The initial ones related to this proposal are seismology, ecology, and astronomy; in addition, there are educational components to the implementation of this proposal. Each of these applications is discussed in more detail, and brings a different workload onto the network, to allow for an understanding of various network performance needs.

Seismology

In order to provide Internet access to remote field sites where researchers are collecting and processing data, our proposal will build upon existing point-to-point wireless Internet links, (designed originally for transfer of seismic data collected at Toro Peak in southern California), by developing a reliable infrastructure that will deal with a variety of challenges and expand the usefulness of the existing links. In the current configuration, a dedicated digital microwave link connects Toro Peak in southern California to the Institute of Geophysics and Planetary Physics (IGPP) at Scripps Institution of Oceanography (SIO) via a repeater on Mt. Soledad (Figure 4). Real-time seismic data are sent over wireless asynchronous serial lines to a data concentrator which converts data into TCP/IP packets and retransmits these data to IGPP for processing.

This example of data collection for seismology is typical of earth sciences (geoscience, environmental science, and ecology), some areas of biological sciences, as well as other disciplines. The field sites are located in remote areas, which have difficult or limited access, with no available power or communications infrastructure. This remote access dilemma is known as the "Last Kilometer Problem."

The continuing explosive development in low power IP-based data acquisition and telemetry systems, solar panel technology, and wireless Internet capability leads to the possibility of a radically different system design. This new design starts with the assumption that all autonomous remote field stations are IP nodes. To connect each IP node to the Internet will require low power radios which operate in the unlicensed 902-928 MHz and 2.4-2.5 GHz spread spectrum bands. The limitation of these spread spectrum bands is that their frequency ranges require direct line-of-sight in order to work. To accommodate this "last kilometer" constraint we will install network nodes at several selected mountain peaks which have existing structures with excellent aerial coverage. The large yellow triangles in Figure 4 connected with black lines show the existing infrastructure currently used by the principal investigators to begin this project. The large red triangles show locations where we plan to install network nodes to extend the line-of-site coverage to include Riverside, San Diego, and Imperial counties in southernmost California.

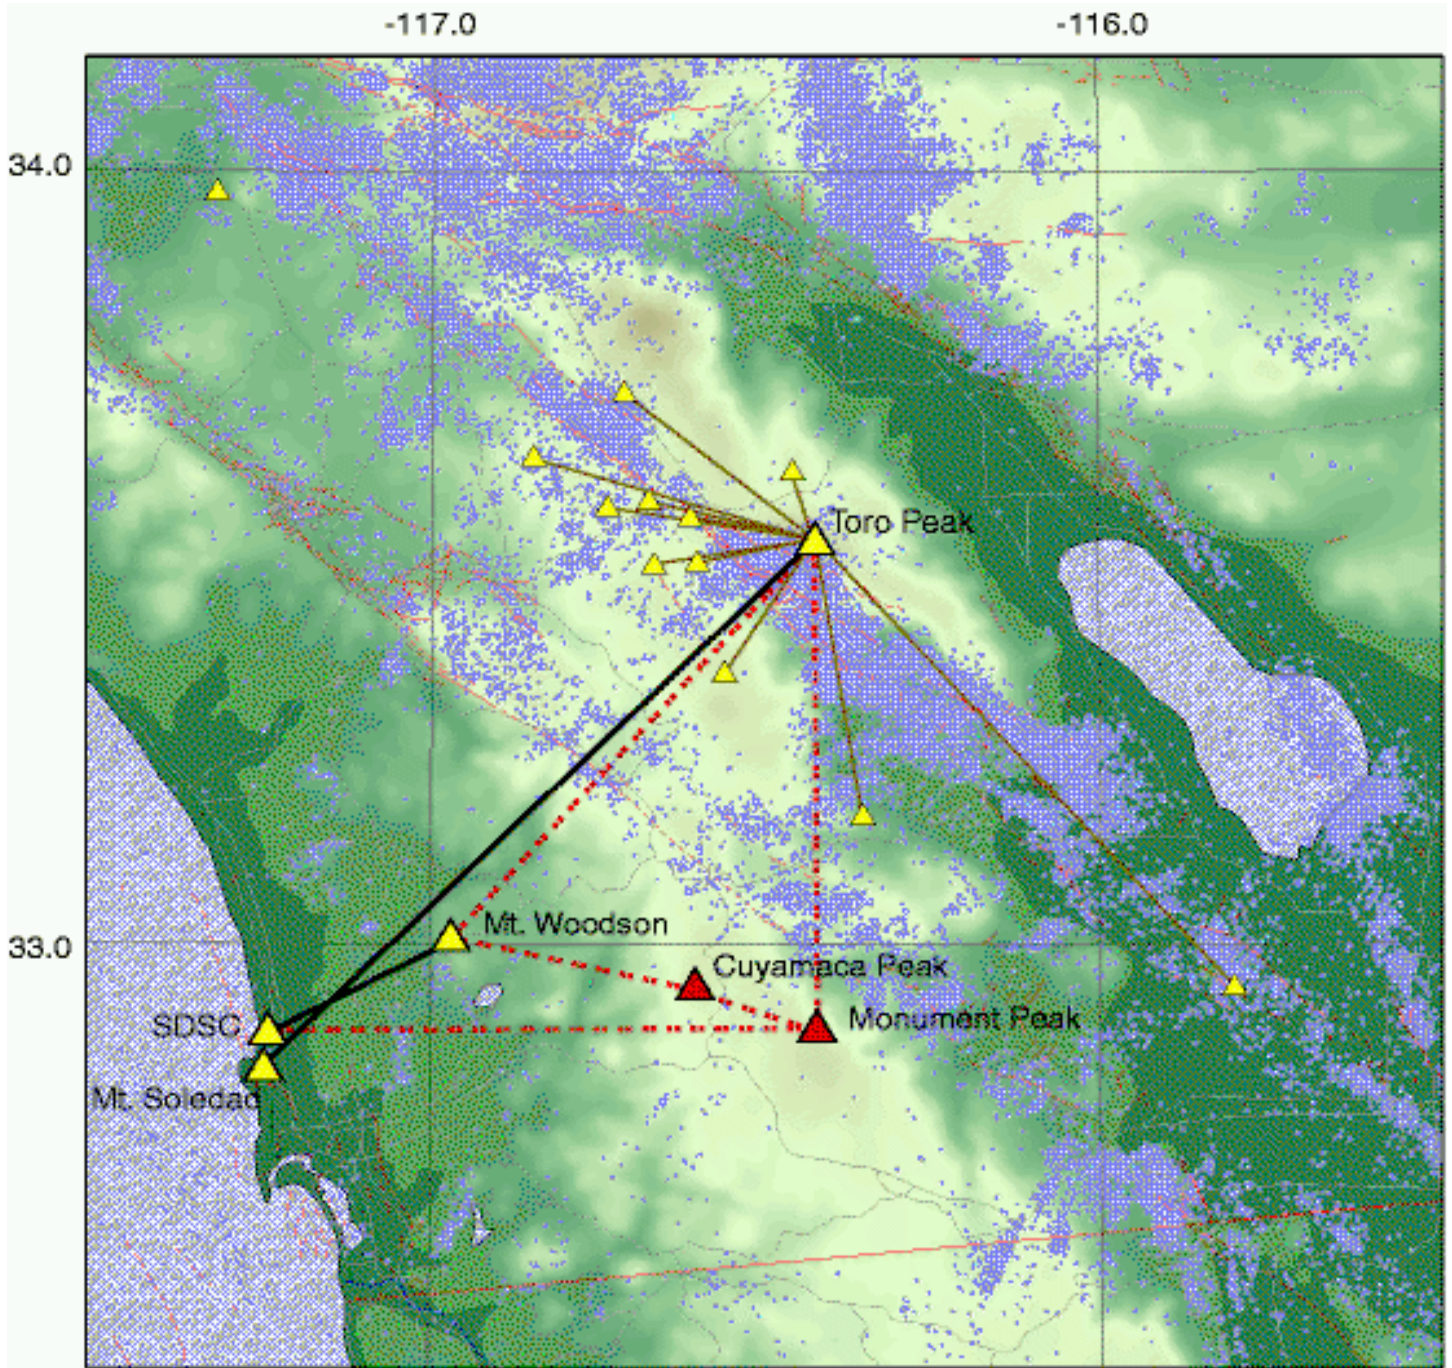


Figure 4

A color topographic map of a portion of southern California showing a possible wireless Internet implementation, and how it would interact with earthquake sensors. The large yellow triangles show sites where the PIs have installed point-to-point wireless IP access (black lines) at SDSC, Mt. Soledad, Mt. Woodson, and Toro Peak. New sample sites at Monument Peak and Cuyamaca Peak are located at the red triangles, with the backbone wireless links shown with dashed red lines. The seismicity between 1990 and 2000 is shown by purple dots. Existing wireless real-time telemetry seismic stations (small yellow triangles) are currently connected to Toro Peak.

Each field station is planned to have:

1. Low communications power, less than 10 watts for RF communications.
2. Multiple Mbps Internet access from sites with line-of-sight and within 30 km of a mountaintop node, using the 2.4 GHz spread spectrum band.
3. Up to 115 kbps Internet access from sites with line-of-sight and within 100 km of a mountaintop node using the 900 MHz spread spectrum band.
4. Low cost (\$200-\$1500) radios at field stations depending on bandwidth and distance requirements.

Initially, sensors in this network will be largely broadband seismographs and Global Positioning System (GPS) receivers capable of locating benchmarks to mm-level accuracy. However, the network can also include other environmental sensors (e.g., meteorological). We anticipate that the development of a successful solution to remote Internet access problems which have arisen with network growth, will allow a full range of geoscience, environmental, ecological and other data sensors to be deployed in a regional network. The network must be capable of delivering data continuously from remote locations. Furthermore, the data must be continuously and completely archived for future use while being prioritized, based on the needs of the actual user, who in some cases has extremely time critical needs, (e.g., during a catastrophic seismic event), while sharing the access network substrate with other users eventually, even perhaps including the public.

At the present time, subcomponents of the network are concentrated at several, individual data collection centers. While Internet communications between these data collection centers can provide access under most circumstances, situations (such as a large-scale earthquake), arise in which the concentrating center may become inaccessible, and immediate, direct access to the sensors would be essential. The network, as currently configured, is not capable of reconfiguring itself to such emergencies. While the seismic case is obvious, similar circumstances can be envisioned for other environmental data including meteorological and hydrological data. It is highly desirable to develop technologies that solve such needs via the evolving Internet, if the technology can conform to our real-time requirements.

Future opportunities - seismic researchers who may want to collaborate on the wireless network:

- 1) Pinyon Flat Observatory - A large geophysical crustal deformation observatory which has been in operation since the 1970's. It uses a dedicated telephone line for data communications. This is the premier crustal strain measurement facility in the world.
- 2) Bombay Beach Observatory - Located in the Coachella Valley to monitor crustal deformation at the southern termination of the San Andreas Fault. Currently has no Internet connectivity.
- 3) Infrasound Array. The field of infrasonic measurements is undergoing a renaissance particularly focused on Comprehensive Test Ban Treaty issues. New sensors and instrumentation will be deployed and tested in southern California starting later this year as the initial part of a larger field

program.

Ecology and Earth Sciences

Areas of collaboration include the earth sciences. There are applications in Ecology and Environmental Science for terrestrial remote sensing devices ranging across physical, chemical and biological domains as a means of observing the hydro, litho and biospheres. Such sensors are represented by meteorological packages, tide, current and streamflow gauges, chemical 'noses' and, eventually, field-deployable bioassay devices designed using the methods of molecular biology. Of course, there are also the conventional video cameras and microphones but the advent of the microbolometer array has made long-wave infrared sensors also available in addition to the conventional VNIR (visible, near-infrared) and multi-spectral cameras available and affordable to the individual investigator.

Ecologists at UCSD, in collaboration with the University of California, Santa Cruz (UCSC), as part of the NPACI Earth System Sciences thrust area, have developed an information architecture of heterogeneous sensors based on TCP/IP and wireless communications; typically using conventional low data-rate radio transceivers. This architecture is an ideal platform for the acquisition and dissemination of sensor-based measurements of ecological and environmental systems. The advent of the new high-bandwidth wireless modems, at relatively low cost, represents a tremendous opportunity to employ some of the high-bandwidth instruments such as the imaging systems mentioned above in networks used to measure and monitor ecological and environmental systems as well as to extend the number and range of conventional remote sensing devices in the terrestrial and aquatic domains.

This proposed high performance wireless network will have the technological capability to accommodate a great number of sensors. It is reasonable to expect a deployment of 20-30 sensors per year over the three year period, and depending on the connectivity and outside factors, this number could be larger by an order of magnitude. These sensors would range from tracking animals to sampling the physico-chemical environment, to visible and infrared imaging.

Astronomy (Mount Laguna Observatory)

San Diego State University's (SDSU's) Mount Laguna Observatory (MLO) is located at an elevation of about 6100 feet on the Eastern edge of the Cleveland National Forest. It is 35 air miles from the SDSU campus (on a direct line of sight), and 45 miles from the metropolitan center of San Diego. The observatory was dedicated in 1968, and was built with funding from the National Science Foundation, with matching funds provided to SDSU by the State of California. The original facilities included two research grade telescopes of 16- and 24-inch apertures, a 16-inch telescope for public viewing in conjunction with the Summer Visitor's program operated jointly with the United States Forest Service, and a four bedroom dormitory. The site has now expanded to include a full-service shop building built with State funds, a larger public-viewing telescope of 21-inch aperture donated by a private individual, a four building Visitor's Center built almost entirely with private funds, and their largest telescope of 40-inch aperture. SDSU is currently raising funds to construct a 100-inch

telescope. The 40-inch telescope was moved from Illinois in 1981 to MLO by the University of Illinois Urbana Champaign (UIUC). The telescope was originally purchased in 1966 under an NSF grant to UIUC and operated on a mound surrounded by corn fields. Since the telescope has been moved to the superior site at MLO, it has been upgraded with modern electronic detectors to produce much improved observational data on stars, star clusters, nebulae, and galaxies. Operation of the telescope is done jointly by SDSU and UIUC.

SDSU Principal Investigators on astronomy projects have been reasonably successful in obtaining Federal support for their research projects, and they expect grant activity to increase as retirements in the current faculty are filled. They typically have two or three NSF funded projects and two NASA funded Hubble Space Telescope projects in progress that make use of MLO. SDSU also provided about one-third of the ground based images to support the Astro-D mission's Ultraviolet Imaging Telescope, which was flown twice on the Space Shuttle.

Their current front-line detector is a 2048 X 2048 pixel (16-bit) Charge Coupled Device (CCD) for imaging. They also have a Near Infrared Imaging Camera, a spectroscopic CCD, and single channel aperture photometer. They have recently proposed to build and even larger 4096 X 4096 pixel CCD to cover more sky area. At present, all data transfer, whether to the SDSU or UIUC campuses, or to collaborators at other institutions, is done via Digital Audio Tapes (DATs) by "sneaker net" or "Fed-Ex". They do have a dedicated phone line direct to their campus computer, but because of the remoteness of MLO, a DSL connection is not possible. Cable modems will probably not be possible for some time to come since Cox Communication only provides their "@Home" service as far East as Alpine. Their present dedicated modem connection provides a maximum usable data transfer rate of about 30 kbps. Thus, the primary uses for the dedicated line are for communications with campus and collaborators, the transfer of smaller data files and digital schematic finding charts (not images), and system maintenance from the SDSU campus.

It is obvious that the current data transfer rates via modem are woefully inadequate to send data directly to the astronomer's home institutions or to collaborators. Their present front-line CCD provides 8.3 MBytes of data per image, which would require almost 40 minutes to transfer under best circumstances. An astronomer can typically generate 200 such images on a long winter night. The problem will only get worse as CCDs become larger, and DSL or Cable modem connections (if available then) would also become inadequate.

The issues created by the current slow data transfer rate are not just convenience, but timeliness and the need for increased flexibility. Many projects such as Nova or gamma-ray burst object searches require the rapid transfer of images to and from various observatories (including those that are space based). In the future, they would like to institute remote observing to carry video guiding images (at a relatively low scan rate) and data images simultaneously over a very high speed Internet connection. The 100-inch telescope will be a more modern instrument, and will in effect become an Internet appliance. They recently signed a Memorandum of Agreement with NASA's Goddard Space Flight Center Laboratory for Astrophysics and Solar Physics (LASP) to be a partner on the telescope project. LASP wants to test new instrumentation for spaceflight readiness, and also to carry out ground-based astrophysics to support their own research done from space. They are also soliciting other academic institutions to be partners on the project. Remote observing is attractive to astronomers because it not only improves efficiency, but it also lowers support and infrastructure

costs. Astronomy is increasingly a global enterprise.

There is also an educational component to the MLO portion of this proposal. File transfers would be of direct benefit to faculty directed student research and General Education. Observations could literally be planned in a GE class for the coming night, and then viewed and analyzed in class the next morning. Once on the Internet, a public database could be established for access by any K-12 or college student, or the casual arm-chair astronomer. SDSU is currently planning to pipe images live from MLO to the Reuben H. Fleet Science Center (located in the Balboa Park area of San Diego) when the 100-inch telescope comes on line. Astronomical Images from CCDs as well as Web-Cam images of the telescope and control room are envisioned.

In addition to the collaboration for astronomy purposes, MLO would also be willing to serve as a seismic station for the project.

Wireless Connectivity to Benefit Education

One aspect of this proposal is to include the capability in the wireless backbone to deliver Internet access to disadvantaged groups. In particular, we plan to provide wireless links (via local access equipment) to both the Pala Indian and the Rincon Indian Reservations through a cooperative effort with the San Diego State University (SDSU). These reservations are located in a rural area near Mt. Palomar, at the Cleveland National Forest. They are not only without high quality communication links, they are currently without any form of Internet access for their students. Representatives of these reservations have indicated that Internet access for their students is their highest priority educational need and we intend to collaborate with them to meet this need. These reservations will be accessed from a backbone station located in the Mt. Palomar area. If this installation works successfully on these two reservations, it will be determined if similar capability can be provided to other local reservations with a similar urgent need.

A second opportunity under investigation is the provision of educational programming to the reservations. San Diego County Public Schools, in cooperation with SDSC, is currently operating a distance learning program reaching about 10,000 elementary students in public schools via wireless links. There is a desire to further the reach of these programs and to make them available to Indian reservations on a pilot program to determine if they can be used effectively. We are evaluating the possibility of assisting the San Diego County on this expansion.

In addition to assisting researchers in their remote work and implementing the educational applications discussed, the planned wireless mesh will also be of use in other education and outreach activities, such as field station visitors' centers. Another potential use would be to develop a prototype emergency dispatch system over the network, in order to test the feasibility, reliability and usefulness of such a system.

In summary, the backbone network is designed to allow new users and usages, such as researchers from other disciplines, to participate at a relatively small incremental cost, not unlike the NSFNET backbone having allowed new sites to attach easily. The participation of researchers from a variety of fields increases the strength and usefulness of the network, both as a performance and analysis test bed and as a prototype to be emulated in a variety of areas for simultaneous, multiple purposes.

The technical overview describes the environment as we are currently envisioning it, although with the rapid introduction of new technology occurring now, we expect that by the time of implementation, some specifics may have been refined. The budget section also reflects only our current thinking.

Technical Overview

It is critical to implement our wide area wireless network as a hierarchy, to accommodate scaling and low incremental costs for new sites. We will create a backbone infrastructure as a geographic "corridor" through the county - from the ocean into the mountains and the desert. Around the backbone nodes, there will be fixed in-zone services to reach users, such as researchers in rural areas or in the field, and autonomous telemetry sites. We expect to experiment with the reach and performances of mobile services, which currently require omni-directional antennas and a reduced data rate to work at the lower signal levels. An additional aspect will be considerations of in-building communications, such as residential house, campus buildings, and other instrumented structures.

Each backbone node will include network measurement instrumentation to allow for passive and active measurements. This will allow us to assess workload profiles, as well as the performance between backbone nodes and peripheral sites. An initial test suite is already being evaluated as part of our overall measurement and analysis activities. This suite will be refined and adapted to the specific infrastructure that we will implement.

In addition, each network backbone node is planned to have:

1. Complete routing capabilities to eliminate single point of communication failures.
2. A data collection workstation, for seismic application data, to buffer and eliminate data loss even if receiver node is taken off-line.
3. Weather sensor to collect and store environmental information for comparison reasons with network performance parameters, such as signal strength information. Inexpensive weather stations provide approximated information for temperature, humidity, wind information, barometric pressure, and rain fall.

Implementation and Impact

We will implement a first phase of the wireless network deployment within a few months of the award. This will include critical backbone links and nodes, as well as the connections to collaborators. After that we will spend significant time researching the actual application needs and performance profiles, while continuing to evolve the network. (As new collaborators join, we will expand the capabilities and uses of the network to adapt to the new criteria.) Figure 5 shows the existing hardware setup of the initial prototype. Figure 6 shows line-of-sight maps from Mt. Laguna (Monument Peak), Cuyamaca Peak and Toro Peak. The line-of-sight from each site is indicated by the yellow/gold color. There is also a topological map of Southern California (upper left image).

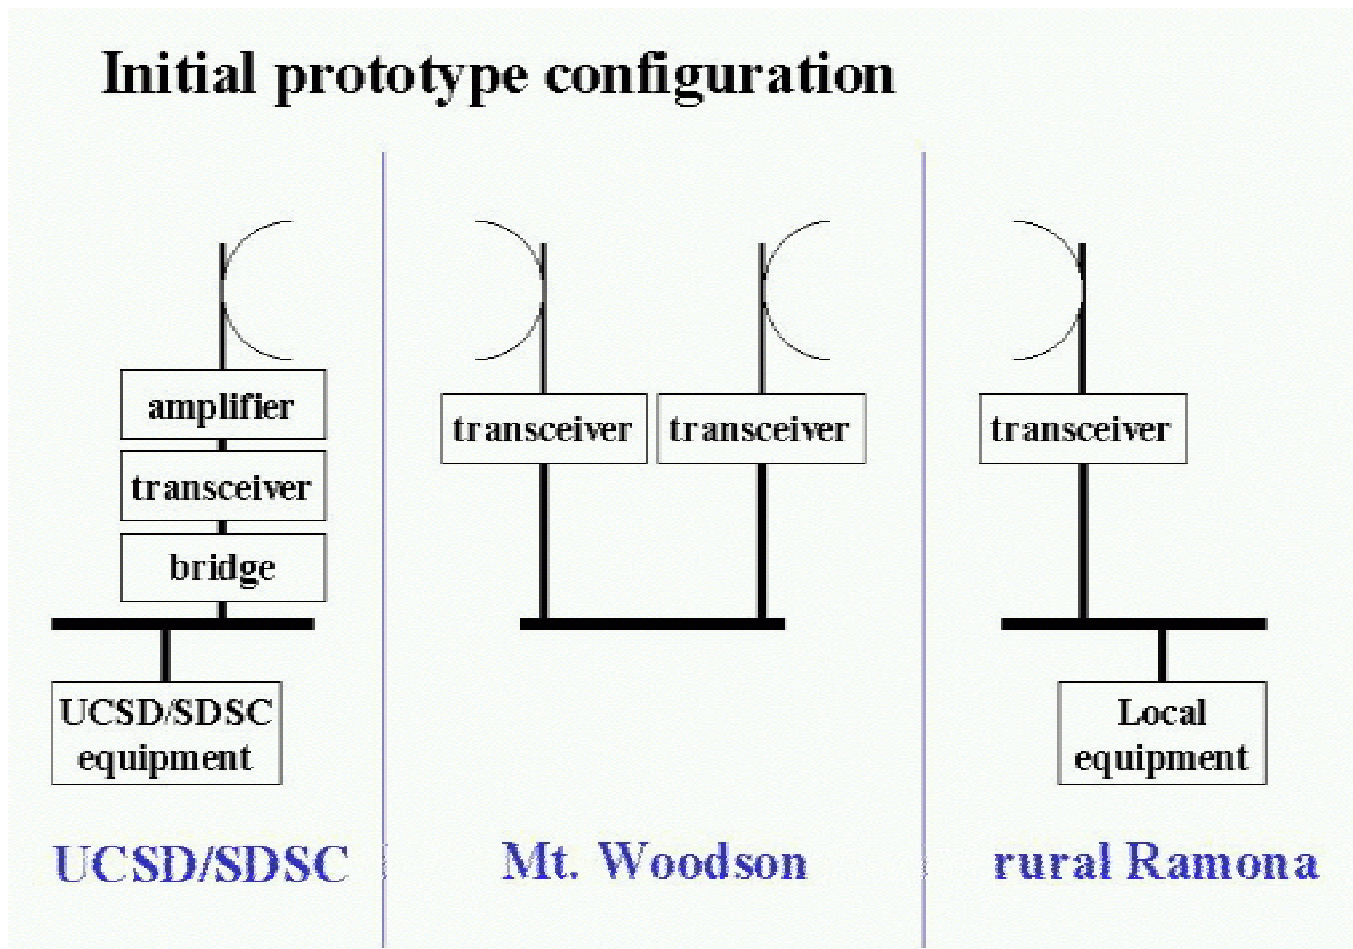


Figure 5 Hardware setup illustration of initial configuration

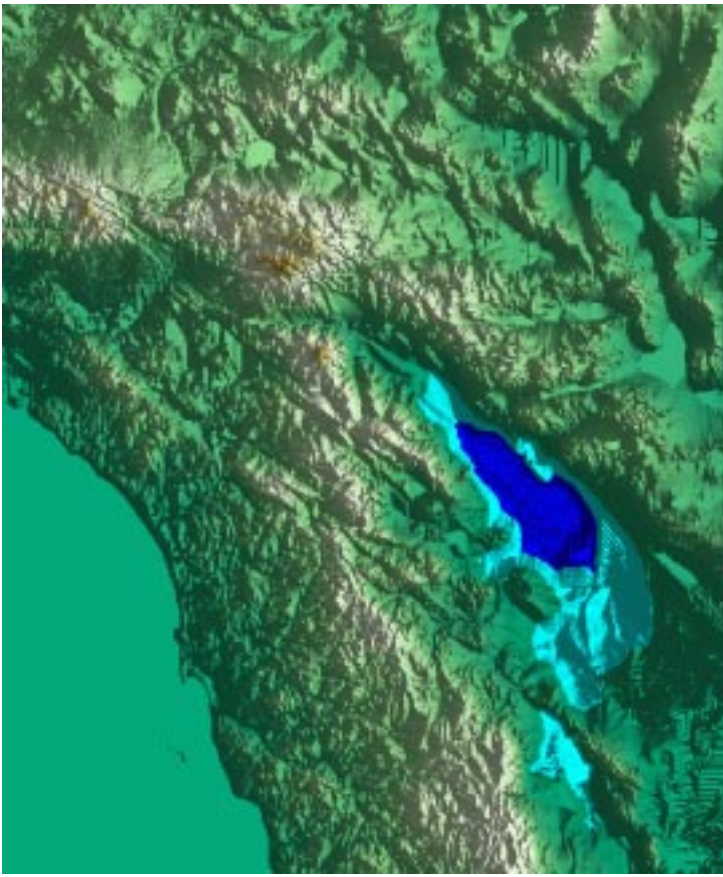
In the second and third years, the deployed network will be enriched by further nodes and connections to collaborators. We will continually investigate how to improve the wireless networking environment, to make it more applicable to actual uses. For example, an area of research will be to apply differentiated services to applications, e.g., to give real-time seismic data priority over background traffic.

The results of this activity will be disseminated by means of a Web site, as well as via published papers and articles.


Substantial benefits are likely to result from this activity, not only in advancing the field of applied network research and understanding of network performance needs, but also impacting disciplinary researchers, educational facilities in rural areas, and disadvantaged communities, including Native Americans living on remote tribal lands.

The high volume capabilities of this network will greatly impact researchers with remote telemetry and observation needs, enabling them to deploy many sensors without overloading the network, as well as extending the reach of currently used tools. The ecological applications alone expect to deploy 20-30 sensors per year over the three year period (depending on the connectivity and outside factors, this number could be larger by an order of magnitude). Another benefit of this proposed high performance wireless network is the possibility for researchers to utilize high-bandwidth tools that were previously unworkable in the field, e.g., imaging systems such as visible, near-infrared instruments. Facilitated communications also enables an increase (perhaps, significantly) in the number of researchers in the field; students could be sent into the field to act as observers (and still be connected to their campus advisors). Researchers could cover wide areas in varying locations, and still be working together in real-time. Collaborations between outside institutions will be facilitated (e.g., the Mt. Laguna Observatory being operated by San Diego State University's [SDSU] and University of Illinois Urbana Champaign [UIUC]).

As wireless networking becomes more and more important in the future, an early understanding of actual needs and performance parameters will be extremely valuable in helping others to implement and make available large wireless environments with good performance characteristics.



Southern California



Mt. Laguna, Monument Peak



Cuyamaca Peak



Toro Peak

Figure 6 Line-of-sight maps (indicated by the yellow/gold color)

Addendum

Results of previous NSF funding

Hans-Werner Braun is the Principal Investigator for NSF-funded ANI-9807479, "National Laboratory for Applied Network Research: Measurement and Operations Analysis Team (NLANR/MOAT)." As outlined in this proposal for a high performance wireless research and education network, the network analysis infrastructure and the measurement and analysis activities form a foundation which will be utilized and expanded upon with the wireless network. For extensive details of this current and previous work and related activities, please see <http://moat.nlanr.net/> and <http://moat.nlanr.net/Reports/>.

For details of the original NLANR cooperative agreement in which Hans-Werner Braun served as Project Director and Principal Investigator (Charlie Catlett, NCSA, Co-Principal Investigator), NCR-9415666 and ANI-9796124, "Creation of a National Laboratory for Applied Network Research," please consult: <http://moat.nlanr.net/Reports/>.