# Wireless Measurement and Analysis on HPWREN

Todd Hansen, Pavana Yalamanchili, Hans-Werner Braun

*Abstract*— **This paper gives a description of the High Performance Wireless Research and Education Network (HPWREN) Measurement and Analysis Infrastructure. We then discuss three case studies in which measurements provided a key component in determining specific problems related to radio frequency filters, weather conditions, and interference.**

*Keywords*— **wireless networks, measurement, analysis, performance, HPWREN**

## I. Introduction

Advances in network technology have made the world a smaller place. Messages can be delivered instantly across the global Internet, enhancing communication and potentially improving productivity. Connectivity has emerged as the single most important ingredient in the information technology revolution. We are pushing the envelope with regard to the speed at which information can be delivered via copper, fiber, and air. Major breakthroughs have been attained in the quest for bigger, better, and more efficient solutions for connectivity.

Measuring the parameters of various network devices and conducting performance tests allows us to study the characteristics of connectivity. These results are used to study and assess whether or not the quest for a better and more efficient connection is being met and to understand the reliability of the systems. Using these results, future installations can be planned and existing ones optimized.

In this paper, we first describe the High Performance Wireless Research and Education Network's (HPWREN) [1] infrastructure and the measurement and analysis activities which are currently being conducted [2]. We then describe a few case studies of some network problems which were resolved using our Measurement and Analysis Infrastructure.

The HPWREN network is a noncommercial, prototype, wide area network capable of speeds up to 45 Mbps. The backbone of the network runs at 45 Mbps using Western Multiplex [6] Tsunami radios. These radios operate in the 5.8 GHz & 5.3 GHz spectrums. The network has many users, some of which are connected at 45 Mbps: Mt. Laguna Observatory, Palomar Observatory and Santa Margarita Ecological Reserve. Other sites are connected using Lucent 802.11 radios: several Indian reservations and a fire station as well as other locations.

The HPWREN backbone network (see Figure 2) begins at the San Diego Supercomputer Center (SDSC), which provides it with connectivity to the major research high performance backbones. From there, two radios point in different directions. The first uses a Gabriel 8ft microwave dish to communicate with Mt. Woodson. Mt. Woodson has a number of 802.11 user links. Three of the 802.11 links have intermediate relays on other mountain tops, including two solar powered relays. The backbone continues on from Mt. Woodson to North Peak. At North Peak there is a 45 Mbps user link for Palomar Observatory; the backbone then continues to Stephenson Peak. At Stephenson peak there is a 45 Mbps user link to the Mt. Laguna Observatory (MLO). Additionally, in the near future we intend to provide connectivity to earthquake sensors in the desert from Stephenson Peak.

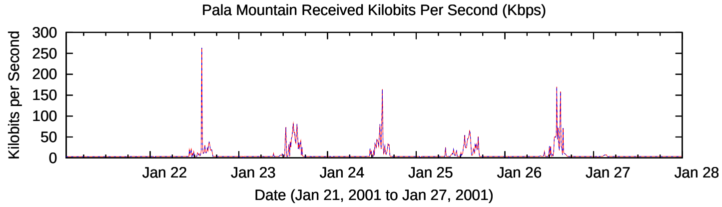

Fig. 1. Initial traffic from the Pala Indian reservation

The other direction from SDSC heads up to Mt. Soledad using a 5.3 GHz band radio at 45 Mbps. From Mt. Soledad, there is a link to Red Mountain (40 miles away). Red Mountain has a 45 Mbps user link to the Santa Margarita Ecological Reserve. Our goal is to eventually interconnect Red Mountain to Toro Peak, and to create a redundant backbone by connecting Toro Peak to the Stephenson Peak area. We have a very diverse user base, each with different network requirements, over a highly variant topology. For example, Palomar Observatory does bulk image transfers at night, while the reservations tend to do tutoring after school. This wide assortment of users and applications combined with wireless networking technology is what makes this network measurement and analysis so interesting and necessary.

## II. Measurement and Analysis Infrastructure

We have the advantage that the HPWREN network is completely under our control. We benefit from this in a number of ways. For example, we use equipment to monitor network performance, to diagnose problems, to measure network activities in a way that previously only network operators could, and finally, to study the impact of network measurements on the network itself. In short, it is an environment we can manipulate in order to conduct various measurement and analysis activities.

In order to utilize the network for these measurements and to provide the results to the network architects, the HPWREN Measurement and Analysis Infrastructure was developed. This infrastructure is still under development, measurements are added as necessary. Currently there are measurement machines at most of our nodes (see Figure 2) and a repository [2] on which we conduct analysis and make our results avail-

Todd Hansen is with the NLANR/MNA, University of California, San Diego, San Diego Supercomputer Center, 9500 Gilman Dr., La Jolla, CA 92093-0505. E-mail: tshansen@nlanr.net

Pavana Yalamanchili is with the HPWREN project, E-mail: pavana@hpwren.ucsd.edu

Hans-Werner Braun is the principal investigator of HPWREN and NLANR/MNA.

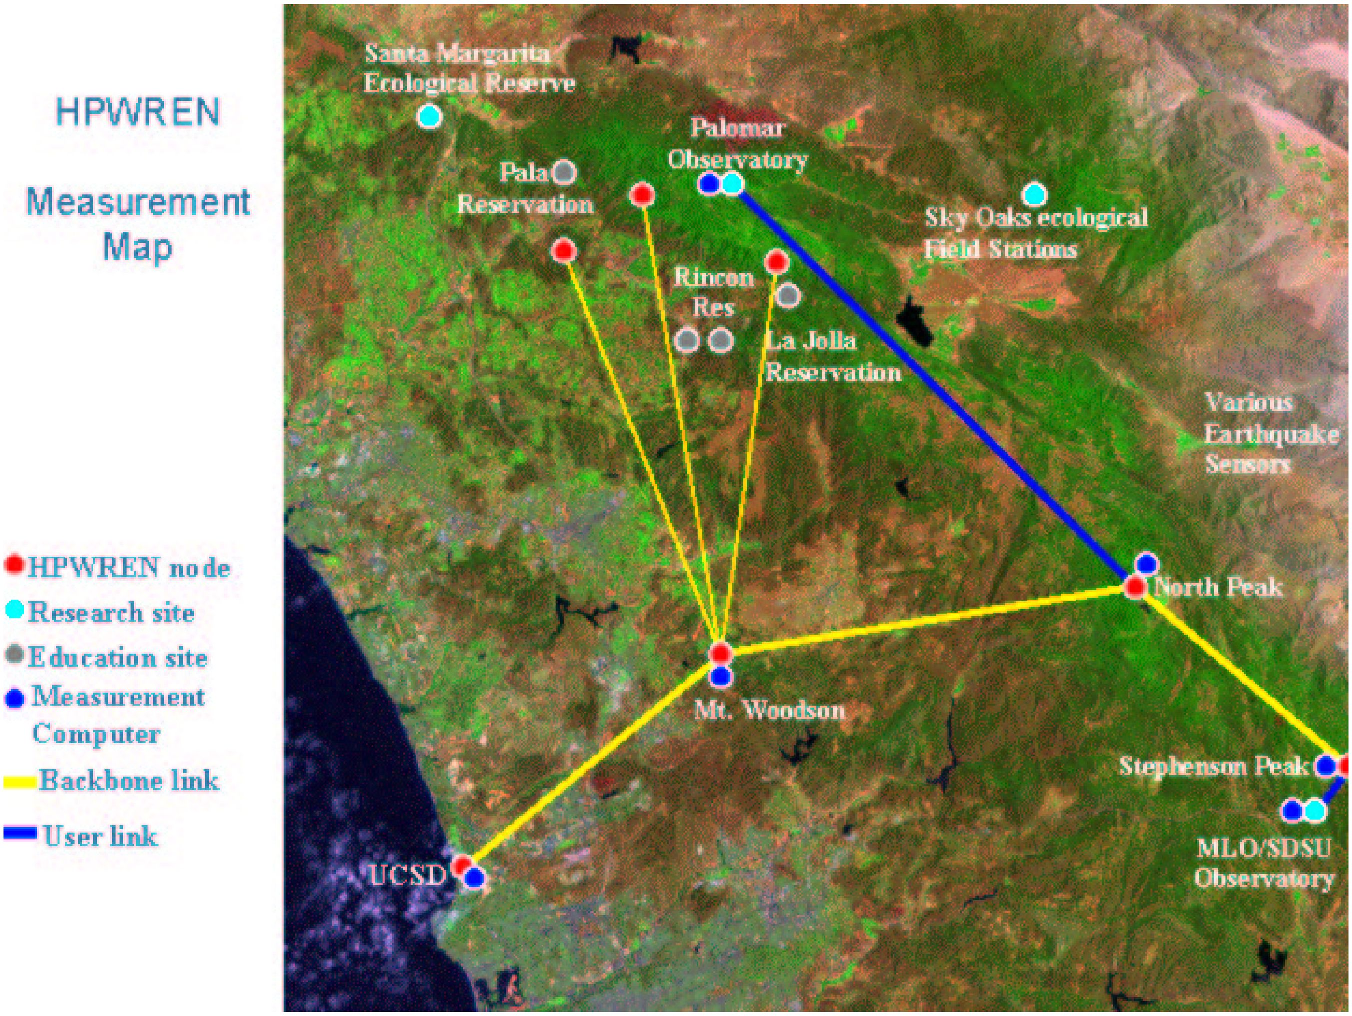

Fig. 2. Map of locations with dedicated measurement computers (San Diego, CA)

able to the public. Since we control the network, we are able to adjust the measurements quickly to meet various needs as they arise.

The measurements conducted on the network are the most important aspect of the HPWREN Measurement and Analysis Infrastructure. They provide the raw data we use to determine cause and effect on the network. They also allow us to gauge many other factors; for example, we were able to see a graph of some of the first network traffic originating from the Pala Indian Reservation (see Figure 1).

The measurements we conduct can be categorized into five basic types: Active Measurements, Passive Measurements, Management Information Base (MIB) Data, Network Status, and Other Measurements. We make all of this data available to the public via our Web site [2].

### *A. Active Measurements*

Active measurements can test the performance a network is capable of providing. As such, these measurements have a distinct impact on the network. They are important in determining the maximum performance an end-user can expect to get and they are even more important in determining when performance has been degraded by a change in network characteristics. For these reasons we run a limited number of active measurement tests on the network. We conduct three types of active measurements. The most significant of these measurements is the one conducted by the NLANR/MNA Active Measurement Project (AMP) [4], [5]. We run an HPWREN AMP mesh that measures round trip time (RTT) and route data between many points within our network. We have machines at most of our end nodes and at each of our major backbone nodes. in order to determine the network's connectivity relative to the NLANR AMP sites, we also measure RTT and route data from each of those sites to NLANR's National Center for Atmospheric Research and Harvard AMP monitors. By segregating the HPWREN AMP mesh from the NLANR AMP mesh we are able to reduce the amount of traffic and still get useful measurements about the performance of our network for comparison with others.

The second set of active measurements we run are throughput measurements (eight times a day on each link). In order to reduce the impact of throughput measurements, we do not run measurements over the entire network but rather over each individual link. This reduces the overall network load

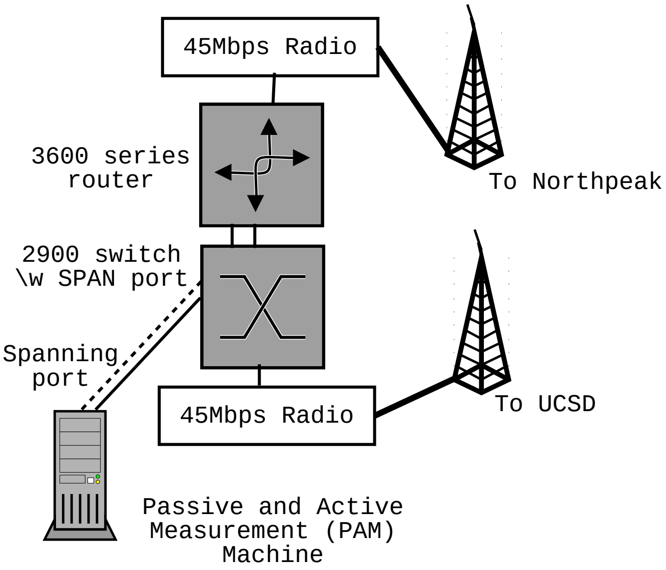

Fig. 3. Diagram of a node's measurement architecture

because we do not conduct multiple measurements over the same segments. However, it still allows us to notice performance changes on individual links as well as the network as a whole. It also allows us to notice consistently high throughput which would point to a well configured link.

The final active measurement we do is to conduct a set of *Windowed Ping (MPING)* [7] tests once per week. These measurements are conducted at night and allow us to better understand the network's capacity and buffering behavior. To reduce the administration of these tests all measurements are conducted from a central server at the University of California, San Diego / San Diego Supercomputer Center (UCSD/SDSC) to numerous destinations. This gives us a connection based data set which is useful for improving the end user experience. Because this test is very network intensive, we avoid performing it excessively.

### B. *Passive Measurements*

We have a distinctive architecture at each of our major backbone nodes that allows us to take encoded traffic traces of our network activity in a variety of ways (see Figure 3). This is accomplished through a spanning port on the switch. A spanning port allows us to view traffic on any subset of switch ports we desire. Currently, we conduct these measurements as needed.

In addition to these on demand passive measurements, we also conduct continuous *Netflow* [8] measurements on one of our Cisco Routers at a central backbone site on Mt. Woodson (see Figure 2). The results are sent back to our repository for analysis. This allows us to understand what type of data is flowing over our network, how much our network is being used, and when. We are currently beginning analysis of this data set.

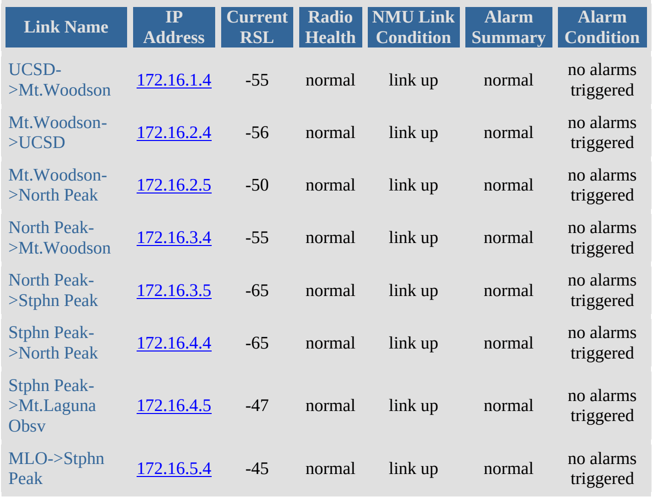

| Link Name | IP Address | Current RSL | Radio Health | NMU Link Condition | Alarm Summary | Alarm Condition |
|---|---|---|---|---|---|---|
| UCSD->Mt.Woodson | 172.16.1.4 | -55 | normal | link up | normal | no alarms triggered |
| Mt.Woodson->UCSD | 172.16.2.4 | -56 | normal | link up | normal | no alarms triggered |
| Mt.Woodson->North Peak | 172.16.2.5 | -50 | normal | link up | normal | no alarms triggered |
| North Peak->Mt.Woodson | 172.16.3.4 | -55 | normal | link up | normal | no alarms triggered |
| North Peak->Stphn Peak | 172.16.3.5 | -65 | normal | link up | normal | no alarms triggered |
| Stphn Peak->North Peak | 172.16.4.4 | -65 | normal | link up | normal | no alarms triggered |
| Stphn Peak->Mt.Laguna Obsv | 172.16.4.5 | -47 | normal | link up | normal | no alarms triggered |
| MLO->Stphn Peak | 172.16.5.4 | -45 | normal | link up | normal | no alarms triggered |

Fig. 4. Tsunami radio status page

### C. *Management Information Base Data*

Management Information Base (MIB) data refers to the data recorded in the memory of routers, uninterruptible power supplies (UPSs), and radios. We use this data to determine the status of the network as well as the individual links. As part of this, we look at the link integrity including signal level degradation and noise floor increases. Additionally, there are alarm sensors that tell us when something goes wrong. On the UPSs, we monitor parameters like power status, current draw, and battery temperature. These measurements allow us to determine the current state of the network and also to understand certain problems that have occurred. We make the raw data available for use by other researchers.

### D. *Network Status*

We maintain these pages so that users, network administrators, and the general public can view current information about critical components of HPWREN. We also have a mailing list that allows users to be notified of site outages. There are four status pages. Each page details the current condition of one aspect of the network. This involves a simple set of measurements, or for certain status pages, MIB based analysis. The first is the highest level and simply denotes the reachability of each site from UCSD/SDSC, the next two pages involve radio status and the last page covers UPS status conditions. Figure 4 shows an example of one of our radio status pages.

### E. *Other Measurements*

The last category is a category for all other measurements performed. We have weather sensors at various nodes that we can correlate with radio performance in order to better understand how weather effects radio performance at these frequencies. Another measurement tool under this category is our set of pan-tilt-zoom cameras. We have two cameras on the network; one is at Mt. Woodson and the other is at the Mt. Laguna Observatory (see Figure 2). There are

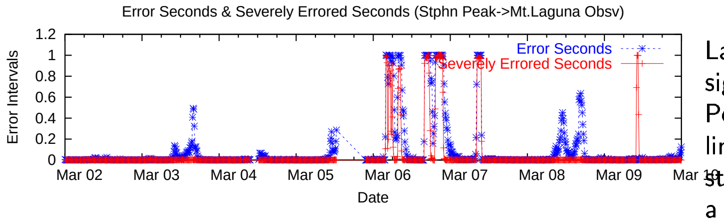

Fig. 5. High number of severely errored seconds

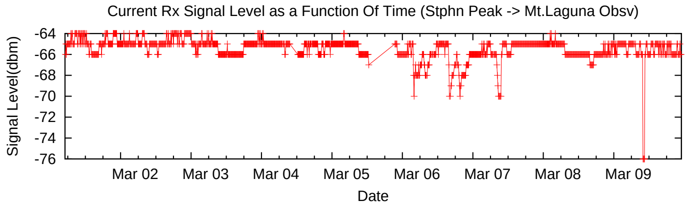

Fig. 6. RSL values from bad link

numerous uses for these cameras, but from a measurement perspective we are interested in correlating the performance of various links with visual weather conditions such as fog and cloud cover. For example, a fixed weather station may never give you enough information to detect an inversion layer, but often a person (or a camera) can easily spot one. An inversion layer is a region of air that quickly changes temperature, it can cause radio signals to bounce or bend upon impact which can severely effect the performance of a link. You can occasionally identify these inversion layers by the clouds caught below them [9, chapter 22, pages 16-18].

The preceding measurements are those we currently conduct on HPWREN. However, we are developing analysis tools on a daily basis and adding them to our Web site [2]. Some of the goals we hope to accomplish include: support of network development through the use of analysis, a better understanding of the HPWREN network and its evolution, and finally, an understanding of how best to instrument a wireless network for network measurement and analysis.

## III. CASE STUDIES

This section discusses a few critical network problems we have encountered and how the HPWREN Measurement and Analysis Infrastructure helped us investigate these problems. Measurement and analysis was a critical component in discovering the cause of each of these problems. In some situations we were unable to solve the problem easily, but in the end we did at least find out the cause. After reading this section the reader might get the impression that network measurements are the only thing needed to diagnose cause and effect. But in reality, network measurements must be coupled with a significant amount of field work. Only then, can we fully define the problem and determine the cause.

### *A. A Radio Frequency (RF) Filter Problem*

Shortly after installation of the HPWREN backbone to Mt. Laguna Observatory (MLO), we noticed high data loss and a significant bit error rate (BER) on the MLO to Stephenson Peak link (see Figure 2). This caused the data flow on the link to be slow and interactive sessions were very jumpy. We started collecting MIB data and soon had a graph showing a marked number of severely errored seconds (see Figure 5). BER is the rate at which bit errors occur in the data. Errored seconds are the number of seconds that have detected one or more bit errors. Severely errored seconds are the number of seconds with a BER greater than $10^{-6}$. The maximum value for errored seconds and severely errored seconds is 1.

We also noticed that the received signal level (RSL), a measure of the remote power being received, was the lowest of all of our backbone nodes (see Figure 6). This was unusual because this link was the shortest and had the best RSL when the network was being designed. (Note that on this link a good RSL level is approximately -48 dBm.)

Our initial reasoning was that the signal level was too low. Based on the original design calculations of -50 dBm, we decided that equipment must be faulty or not installed correctly. This is where the field work came into play. Both sites were visited numerous times, antenna alignment was checked, feed horns and feedlines were replaced, finally the radios were replaced (one at a time). All of this made no significant change in the RSL or the persistent errors. Finally, it was decided to temporarily turn up the power level to see if we could increase the RSL and reduce the errored seconds. At the same time the power was increased on one of our other links to see what would happen.

Within one week we noticed that the other link had fewer errored seconds. But errored seconds on the MLO link had increased drastically after we changed the power level (see Figure 7); we also saw asymmetry in the link. After increasing the power at both ends of the link, the errored seconds on the MLO to Stephenson Peak link remained unchanged; however the errored seconds on the Stephenson Peak to MLO link increased. At this point the problem was isolated to the radio at Stephenson Peak which pointed towards MLO. The fact that distortion increased with more power lead us to believe the problem might lie with the power amplifier or RF filter circuit.

We then recalled that we had not tried a replacement filter at Stephenson Peak (previously we did not know it was an issue). Had we not made the measurements and performed the field work we would never have realized that it was an issue. The next day, the radio and filter at Stephenson Peak were replaced.

Immediately, the signal level rose to over -50 dBm (see Figure 7). The errored seconds also stopped occurring and we had no detectable bit error rate. Later in the day we went to the Mt. Laguna Observatory (MLO) to turn back down the power. While there, we were able to achieve 40 Mbps from MLO to the coast (UCSD/SDSC). The top throughput you can achieve with these radios is approximately 40 Mbps, as protocol headers take up the remaining bandwidth. This was the first time we had been successful at achieving opti-

**Timeline of Changes:**
**Apr 1st - increased transmit power**
**Apr 13th - replaced SP radio**

**Results:**

**Prior to Apr 1st - RSL at -65 dB, medium error seconds and BER**

**Apr 1st - RSL at -62 dB, extermely high errored and severly errored seconds, BER is significantly higher as well**

**Apr 13th - RSL at -43 dB, no errored seconds BER=0, 40 Mbps from Coast to Mt. Laguna**

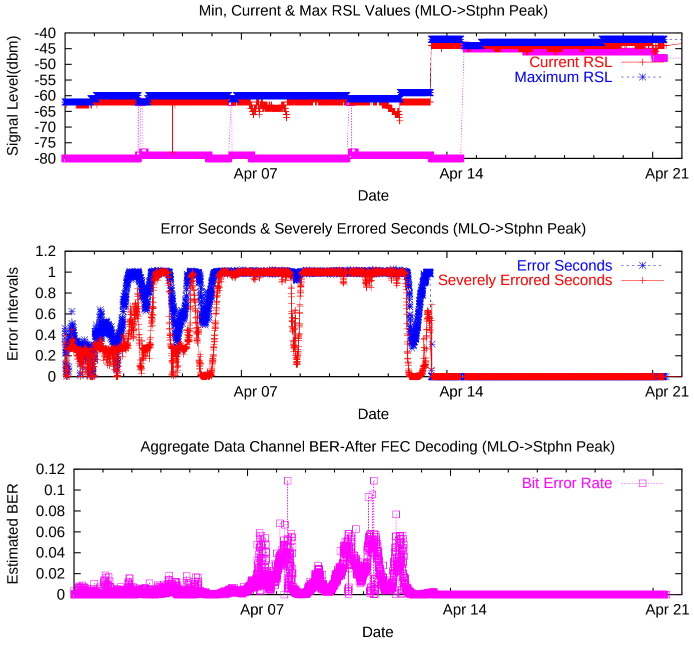

Fig. 7. Timeline of events including analysis data

mum throughput. By that evening researchers at Mt. Laguna Observatory were sending their images directly to San Diego State University for analysis. Within 18 hours, images from MLO were being used to teach a class; this was a great improvement. With the old dialup/tape drive network it would have taken significantly longer from mountain to classroom.

In the end, if we had not made these measurements and conducted this analysis we would not have found the problem. But on the other hand, if we had not done all of the field work we would have been equally unsuccessful. That is the most important lesson we learned from this exercise.

### B. A Weather Problem

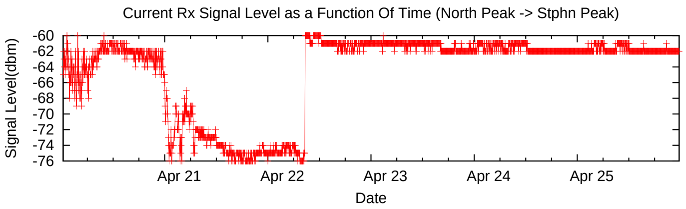

Fig. 8. RSL during the week of the storm; storm occurred on April 21, 2001

During a spring snow storm, the signal level at North Peak dropped drastically. It dropped so low that the North Peak to Stephenson Peak link did not have any connectivity for the entire duration of the storm. Even though some attenuation of the RSL due to rain had been observed in the past, it was never expected that RSL values would drop close to the threshold value (-79 dBm). Figure 8 shows the RSL values during the snow storm.

In Figure 8 it can be seen that before the onset of the storm, the RSL value fluctuated around -62 dBm. Within an hour of the onset of the storm the RSL dropped by 14 dBm to -76 dBm.

A fade margin of 22 dBm is generally enough for a connection to survive a storm in the San Diego County area. The fade margin is the expected maximum change in RSL due to weather conditions. It can be seen that the change in link RSL did not exceed the 22 dBm fade margin. However, this link was designed with a fade margin of approximately 15 dBm (because use of a bigger antenna was not feasible at this site) and hence, a 14 dBm drop was enough to kill the connectivity. Most of the other links in the network are designed with a fade margin of 22 dBm and to date these links have not exhibited such connectivity problems. If the North Peak to Stephenson Peak link had also been designed with a fade margin of 22 dBm, this problem might not have occurred.

An interesting factor to note is how quickly the link RSL dropped by 14 dBm to -76 dBm. Within an hour of the onset of the storm the connectivity was lost. This sharp drop occurred right before 1:00 AM on the April 21, 2001; and the RSL remained at this low level until after 8:00 AM on April 22, 2001.

Due to the storm, a field trip to these sites could not be arranged and the link characteristics could only be studied and monitored using the HPWREN Measurement and Analysis Infrastructure. By the time the storm ended, the RSL level was back to normal (around -62 dBm) and the connection was restored. This provided further evidence that the problem was entirely weather related.

This experience points out yet another important issue to consider when planning a network installation - how quickly the environmental conditions impact the RSL. While this problem was resolved using the measurement and analysis data alone, most problems require a site visit. During adverse weather conditions, a field trip to the problematic site is not feasible and measures such as the following should be taken beforehand to avoid similar problems.

One of the most obvious solutions is to increase the transmit power so that even with a drop of about 22 dBm, the RSL remains well above the threshold value and connectivity is maintained. Care should be taken not to violate the Federal Communications Commission (FCC) rules regarding the transmit power. Hence we could not increase the transmit

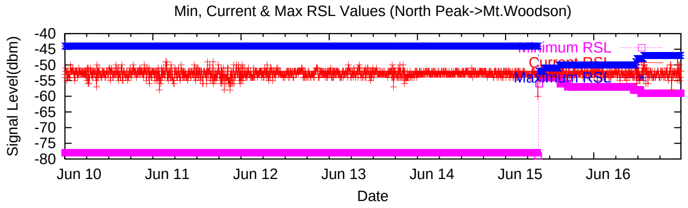

Fig. 9. RSL during week

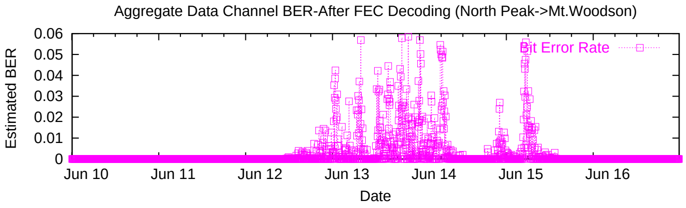

Fig. 10. Bit error rate during week

power for the problematic link.

Another solution to improve the link is to upgrade the feed line which is the cable that connects the antenna to the radio. Currently, we are using 5/8" heliax and lmr-400 feed line; we plan to upgrade to waveguide feed line. The waveguide feed line has much less attenuation as compared to 5/8" heliax and lmr-400 feed line.

We chose to replace the cable but another possible solution would have been to change the antenna on North Peak from a six foot antenna to an eight foot antenna. Using the eight foot antenna would have given us better gain and as a result, a better RSL.

In the end, this experience also highlights a critical problem for institutions and researchers who rely on connectivity to meet their daily needs, such as ecological field stations and seismic sensors. These groups need connectivity 24 hours a day, 7 days a week because they rely on real-time data. Loss of connectivity even momentarily can result in loss of data which could prove to be quite costly.

If it was not for the HPWREN Measurement and Analysis Infrastructure, we would not have discovered the problem during the snow storm itself. We would have had to wait until the storm had subsided and eventually a field trip would have been conducted which might not have discovered the problem. This approach would have been very time consuming; also it may not have led us to the cause.

### C. An Interference Problem

The last problem we would like to discuss is one that involved external events that in the end caused our network to behave poorly. The UNI-I/II bands are allocated to unlicensed operation on a secondary basis. The primary use is for military radar. They have a frequency range of 5.25-5.925 GHz [10].

As a result we observed an interesting situation where for approximately one work week [four days from June 12 (Tuesday) - June 15 (Friday)] the link was severely impaired by what we now suspect were radar tests.

Initially, we thought the problem was just occurring on one or two radios, specifically the radio at North Peak that pointed to Mt. Woodson. A site visit was made and everything tested, including alignment of the antenna. In the end, even replacing the radio left us with the same interference. During the field work it was noticed that more than just this link was being affected. We also noticed that none of the affected sites saw a decrease in the received signal level. We have included graphs of the RSL (Figure 9), bit error rate (Figure 10) and errored seconds (Figure 11) for North Peak. The data from other sites looked similar.

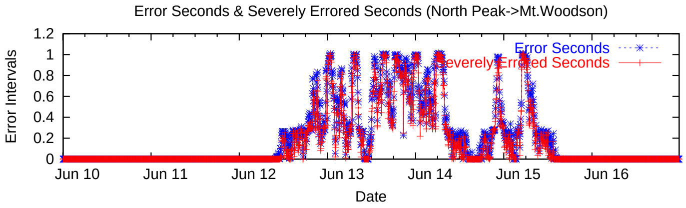

Fig. 11. Errored seconds during week

The only indication that this was occurring was that all of the A2 radios (A2 radios use a frequency of rx 5.750, tx 5.800) were getting very large bit error rates and the radios were almost continuously registering severely errored seconds. The A1 radios (A1 radios use a frequency of rx 5.800, tx 5.750) were having very little problem receiving the A2 radios. As a result we were able to determine the interference was on the 5.750 GHz frequency. In other words, the interfering signal was overpowering the receivers and therefore causing our data to be corrupted. Whatever was causing the interference was using much more power because they were affecting all of our A2 radios. The interference receded during the evening of June 15th. We have not seen the problem reoccur, so we suspect someone was trying some experimental hardware that they eventually turned off.

In this instance we were unable to find a way to avoid this problem. However, it serves as a reminder that with most unlicensed spectrum we are not the primary users of the allocated frequency range. And as such, we need to develop plans for this interference when further designing our wireless network. It also would have been beneficial to do a more thorough analysis of the data before we went into the field. It would also have been useful to have a measure of noise floor for the Tsunami radios. For the 2.4 GHz gear we selected, this information is measurable without taking the link offline. This would have allowed us to see the noise floor increase and would have led us to look into interference issues sooner.

## IV. CONCLUSION

In this paper we discussed HPWREN Measurement and Analysis Infrastructure as well as some case studies in which we used these measurements to determine the causes. Measurement and analysis activities while crucial for any network, are especially so for wireless networks. Additionally, by describing the HPWREN Measurement and Analysis Infrastructure we have defined a set of measurements we feel best represent the type of measurements needed in a wireless network.

## V. ACKNOWLEDGMENTS

This work was sponsored by the National Science Foundation under Cooperative Agreement No. ANI-9807479 and Grant No. ANI-0087344 to the University of California, San Diego. This work was done in close collaboration with the National Laboratory for Applied Network Research Measurement and Network Analysis Group (NLANR/MNA) [3]. We would like to thank Jose Otero and Dewayne Hendricks for their invaluable contributions to this paper. We would also like to thank Maureen C. Curran and Kim Bruch for their reviews of this paper.

## REFERENCES

[1] High Performance Wireless Research and Education Network (HPWREN), http://hpwren.ucsd.edu/
[2] HPWREN Measurement and Analysis Web Server, http://stat.hpwren.ucsd.edu/
[3] National Laboratory for Applied Network Research Measurement and Network Analysis Group (NLANR/MNA), http://moat.nlanr.net/
[4] NLANR/MNA Active Measurement Project (AMP), http://moat.nlanr.net/AMP/
[5] McGregor, A., Braun, H-W., and Brown, J., *The NLANR Network Analysis Infrastructure*, IEEE Communications Magazine, Vol. 38, No. 5, pp. 122-128, May 2000
[6] Western Multiplex: http://www.wmux.com/
[7] M. Mathis, "Windowed Ping: An IP Level Performance Diagnostic", Proceedings of INET'94, June 1994. http://www.isoc.org/inet94/papers/523.ps.gz
[8] CAIDA, *cflowd: Traffic Analysis Tool Web site*, http://www.caida.org/tools/measurement/cflowd/
[9] *The ARRL Handbook for Radio Amateurs 1993, 70th edition*, The American Radio Relay League Inc., Newington, CT, 1992
[10] Dewayne Hendrix, Member, FCC Technological Advisory Council, http://www.dandin.com/