

Connectivity maps: Measurements and applications

Theodoros Kamakaris, Jeffrey V. Nickerson
Stevens Institute of Technology
tkamakar@stevens.edu, jnickers@stevens.edu

Abstract

Network connectivity is becoming an important factor for many transportation problems, yet in general we do not have available to us sufficiently detailed maps of expected connectivity that might guide our decision making. Here we detail an approach to creating such maps. Connectivity contour maps are introduced, and the relationship between signal strength and throughput discussed. This work may contribute toward more widespread efforts to create such maps, and the presence of such maps may in turn have implications for emergency response, transportation planning, policy, and individual decision making.

1. Introduction

Improvements in mobile data communications are creating the potential for a new range of applications. If, for example, Wi-Fi roaming prevails, as many predict [1], then it will be possible to perform a variety of data-related transactions as one moves. These applications may be important both in applications related to emergency response, and in applications related to general commuting.

Yet the microcell / picocell architecture of Wi-Fi environments makes large scale RF planning extremely difficult and complex. Connectivity along a projected coverage area varies greatly due to both propagation and network loading characteristics. Real time travel guidance systems are now possible that take advantage of GPS location to optimize travel based on traffic conditions [2]. But it is true that many may want to consider available bandwidth while planning travel, and there is evidence to suggest that people will change patterns of behavior based on bandwidth [3-5]. It is conceivable that, if we understood the environment - if we had connectivity maps – we could change our movement patterns to optimize not just our travel time but our data transmission as well. If so, then algorithms which trade off travel speed for increases in bandwidth may be useful. These algorithms exist, and they are discussed in a concurrent publication [6].

As a prerequisite to the effective implementation of such algorithms, we need to be able to build a connectivity map. This paper discusses the theoretical and practical difficulties in doing so, and presents information

on an ongoing effort which keeps and up-to-date connectivity map of the Stevens Institute of Technology campus, as part of a larger research effort in wireless security [7].

Others have engaged in work to map out areas of Wi-Fi coverage [8, 9]. What distinguishes this work from other efforts may be our focus on not only determining coverage areas, but creating contour maps that can be used to determine the user's expected bandwidth at every location.

2. Figures of Merit

In typical radio frequency (RF) planning for wireless systems, the figures of merit used to determine the area and quality of coverage are the Signal to Noise Ratio (SNR) and the Signal to Interference Ratio (SIR). As packet based networks become prevalent for wireless data, metrics at higher layers, adapted from wired communications such as latency, throughput and quality of service (QoS) become more representative figures of merit for the end user.

Analyzing and troubleshooting a Wi-Fi wireless data network requires monitoring of the communication channel at the physical, data link and network layers by monitoring bit error rate (BER), collisions and availability seen at the MAC, packet error rate (PER), throughput and jitter. Yet, given an infrastructure, the characteristics of service available to the end user can be summarized by the mean and variance of the “goodput”, a metric measured at the transport layer containing information regarding throughput of end user data without associated overhead of delivery (retransmissions, overhead etc.)

3. Experimental Setup

The experiments presented were performed on the Wireless Network Security Center (WiNSEC) testbed developed in Stevens Institute of Technology, Hoboken NJ which includes among other infrastructure, a production level campus wide Wi-Fi network. In order to quantify the coverage characteristics of the Campus wireless network, two separate measurement campaigns were designed, providing two different approaches; RF planning and Network design.

During the first measurement campaign signal strength measurements were taken at 100 points distributed along

the Stevens Campus. The experimental setup consisted of a Windows based laptop with a PCMCIA network card running Netstumbler software and connected to a GPS (figure 1). The software was periodically scanning through all eleven Wi-Fi channels, monitoring for access point beaconing packets. When beaconing packets were received, the Received Signal Strength Indication (RSSI) along with the access point MAC address and the GPS location per beaconing packet were recorded. The data was post processed to acquire one representative signal strength measurement per location. This was achieved by averaging over all packets received at the same location for a particular access point. Whenever multiple access points were detected at a specific location, the one with the strongest signal was considered, assuming that this would also be the choice of a roaming laptop.

MAC	SSID	Name	Chan	Vendor	Signal	Noise	SNR
00078583C17E	PyneNET		1	Cisco	-85	-96	9
00022D2621382		McLean 4th floor lab		Proxim (Agero) ORINOCO			
00022D1A81897		McLean 4th Floor		Proxim (Agero) ORINOCO			
00022D1A45223		Play Lab Burchard 610		Proxim (Agero) ORINOCO			
00022D178939		Schact Lab Two		Proxim (Agero) ORINOCO			
00022D17A51F		Schact Lab One		Proxim (Agero) ORINOCO			
00022D1A45B4		Burchard CS Lab 1st Floor		Proxim (Agero) ORINOCO			
00022D6213DB		EAS201		Proxim (Agero) ORINOCO			
00022D17A673		Carnegie 1b		Proxim (Agero) ORINOCO			
00022D17A5F9		Carnegie 1a		Proxim (Agero) ORINOCO			
00022D15887C		agPond		Proxim (Agero) ORINOCO			
00022D1A8111		Machine Room		Proxim (Agero) ORINOCO			
00022D4D2C86		Pocco 2nd		Proxim (Agero) ORINOCO			
00022D1A45F5		Play Lab Burchard 619		Proxim (Agero) ORINOCO			
00022D17A5A2		Hayden Lounge		Proxim (Agero) ORINOCO			
00022D4D2C74		Library 1-st Fl		Proxim (Agero) ORINOCO			
00022D4D2C8A		McLean Basement		Proxim (Agero) ORINOCO			
00022D4D2C82		Griffith 1st		Proxim (Agero) ORINOCO			

Figure 1. Screenshot of NetStumbler

The second set of measurements was designed to monitor connectivity quality beyond availability. The experimental setup consisted of the same hardware, yet in this test a Linux based environment was used to achieve a more robust experimental platform. Measurements were taken using an automated bash script while students carried the laptop around the campus. The open source ttp software tool was used to measure throughput from the laptop to a server located at the backbone of the wireless infrastructure, isolating in this way network performance degradation to the wireless network infrastructure.

The results of these two initial experiments are illustrated in the following two figures. The topology of the measurements was ignored in this phase in order to concentrate on the implicated difference on network coverage from the two approaches comparing physical layer measurements to transport layer measurements. Figure 2 shows the interpolated map of network availability based on signal strength measurements. The color map scales used indicate the RSSI values from minimum to maximum as follows: *Blue – Cyan – Yellow – Orange – Red*.




Figure 2. Signal Strength Network Availability Map

The results of the throughput based measurements are presented in figure 3 with the same color map as before. Here the values range from 0 (blue) to maximum throughput measured (red).




Figure 3. Throughput Network connectivity Map

The discrepancies between the implicated results of the two approaches provide for an interesting discussion regarding testing methodologies and the importance of choosing the proper figures of merit. Based on the data analysis of the initial measurements we decided to utilize throughput as the figure of merit for assessing quality of wireless network coverage. The throughput measurement campaign was repeated for acquiring a larger set of data. The resulting coverage map is presented geo-referenced in figure 4, overlaid with a map of the Stevens campus – Stevens sits on top of a cliff overlooking the Hudson river, and the green area delineates the edge of the cliff.



Figure 4. WLAN coverage of Stevens campus network

4. Environment Dynamics

One of the first questions that arise regarding the measurement methodology is the ramifications of different measurement densities to the experimental results. Setting the priors for this experiment, it is important to notice the characteristics of the environment. The Stevens Campus is an urban environment with significant terrain variation despite its limited size. Although the measurements of this experiment focused on the outdoors coverage of the wireless network the initial coverage planning was geared towards creating individual hotspots for indoor locations. Therefore the placement of most access points is far from ideal for outside coverage which is mostly provided through signals “leaking” outside the buildings. It is characteristic that although more than 30 AP’s are installed throughout the campus, only 11 were detected to provide outdoors coverage during the measurements.

These complex propagation characteristics of the test environment necessitate high density measurement distributions in order to reconstruct an accurate RF propagation map. Furthermore, the urban environment is inherently a highly dynamic one. Mobility of vehicles as well as high population density allows large temporal variations of the signal propagation characteristics. Figure 5 shows the temporal variation of the signal strength of the beacons transmitted by two access points over a period of 10 seconds. The beacon transmission period of the access points was set to 100ms, allowing 100 measurements to be taken in 10 seconds. Although both the receiver and the transmitters were stationary, it is obvious that the propagation environment varied significantly within sub second time periods. For this reason, multiple packets were captured and averaged for

each location in order to acquire an indication of the mean channel condition.

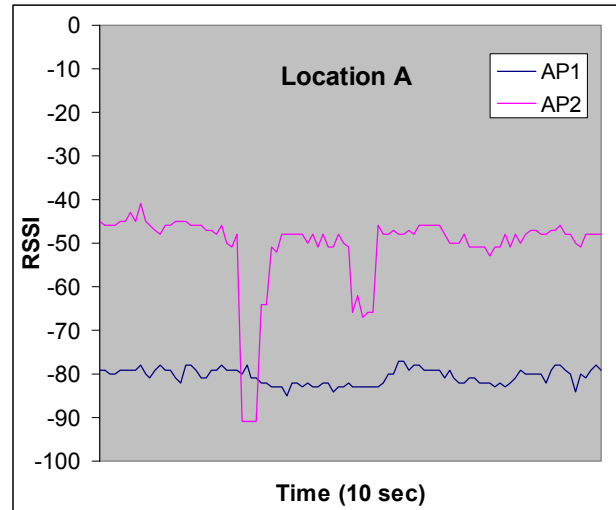


Figure 5. Temporal variance of RSSI

Similarly in the case of throughput measurements, the throughput achieved at each location, exhibited variations as can be seen in figure 6. Yet in this case the time scale associated with each measurement was different by an order of magnitude since during each measurement, 1MB was transferred allowing for a measurement period of 2.5 seconds or greater, depending on the connection quality.




Figure 6. Temporal variance of Throughput

Throughput variation can be attributed to a variety of factors other than propagation channel variations, including medium contentions by other users, increased access point traffic or other wired network latency issues. In order to isolate these factors a number of measures were taken. Medium contentions were minimized by enabling Request to Send / Clear to Send (RTS/CTS) to all access points. Furthermore, the experiments were conducted at low traffic periods to minimize potentially

interfering users. Finally the wired network latency is not considered to influence the measurements since the wired network at the time of experiments was underutilized.




Figure 7. A histogram of throughput of a file transfer over 2500 continuous trials.

Understanding the distribution of throughput is important in emergency response applications, where different situations may call for different attitudes toward risk. For example, the above distribution shows that the expected bandwidth is high – but there is a chance of much reduced bandwidth, and responders in a particular situation may not be willing to tolerate the risk.

5. Connectivity map applications

For movement through continuous space – as for a consumer walking in an open area – a contour throughput map is ideal. For road travel, we might want to first convert the map into a road grid. For each edge of the grid, we would want to know both the expected throughput while traveling over the edge and the expected time to drive along the corresponding road. So essentially, we would overlay a grid over the contour map and use the expected values of the contour map to label the route map.

To give one example, one might want to plan one's trips on the basis of the amount of bandwidth available. Consider the following situation:

An emergency responder wants to get to an incident as fast as possible, but needs to maintain connectivity at a particular throughput or greater along the way. What is the faster route?

We can represent this with the following more formal question: Given a spatial attributed graph G , in which a

function $bw(e)$ can provide us the expected bandwidth along an edge, what is the shortest path p between $s, t \in V$ such that $\forall e \in p, bw(e) > C$ where C represents a threshold bandwidth above which connectivity to maintain a particular application is certain.

This can be solved by mapping $G \rightarrow G'$, where $e \in G' \Leftrightarrow e \in G \wedge bw(e) > C$, and then finding the shortest path of G' . In other words, we drop out the low bandwidth edges. Figure 7 shows how this process might be implemented, starting with a contour map, and then transforming a graph into a modified graph with only high throughput edges showing.

The low throughput edges might represent routes through tunnels, or through areas of spotty coverage. One can imagine that this function above might prove a useful tool for emergency responders. It might also apply in more prosaic commuting situations, in which commuters wish to download data or engage in conference calls.

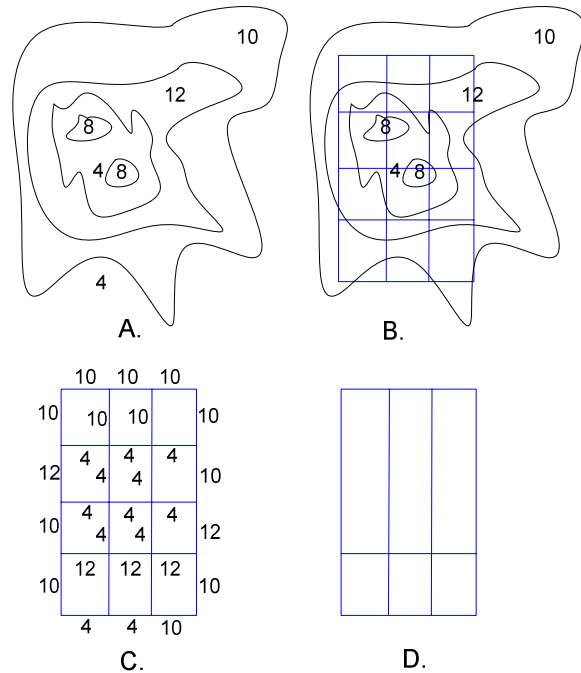


Figure 8. A roadmap (B.) is overlaid on a contour connectivity map (A.), to create C. By applying a threshold, a new map of high-bandwidth routes is created (D.)

As mentioned before, it is possible to enhance the algorithm to consider not just the expected throughput, but the entire distribution. One might construct a utility function that would adjust the choice of routes according to one's situation dependent degree of risk averseness; this issue is discussed more deeply in [6].

It is also possible to imagine applications in which the need for throughput may modify the speed with which

one is willing to travel. In emergency response, one might not want to arrive on certain scenes until a map has already been downloaded. The connectivity map we have discussed provides a platform on which to develop these new types of route planning.

6. Future Research

Given the dynamics of an urban environment, the generated map cannot be expected to be accurate over a long period of time. As indicated by the measurements in the previous sections (figures 5, 6), connectivity varies in a number of time scales. The propagation characteristics of the investigated environment are governed by Rayleigh and Ricean fading [10], allowing for large variations of the signal strength caused by minute changes of the environment such as cars moving, people passing by, or simply the motion of a sign caused by wind. In addition, environmental changes such as rain, fog and snow have an adverse effect on signal propagation [reference]. Furthermore, the network loading characteristics vary at a different time scale as the traffic depends on the time of day, day of week or other larger scale changes, such as an increased user density of the Campus network during a conference, exams etc.

In order to keep the connectivity map up-to-date, we are researching two approaches that can compliment each other. In order to provide with periodic updates of the map we have developed laptops that are installed inside the cars of the campus police that can be switched from regular operations (providing internet connectivity to campus police cars) to throughput measurement sensors. These laptops are equipped with GPS and can duplicate the measurement campaign described above, with the exception of more confined motion, since the cars have to remain on the campus streets. The second approach allows us to dynamically update the signal strength maps by retrieving information from the access points regarding the signal strength of each connected user and correlating that information with the location of the user. The location can be known by cooperative means (GPS) or non cooperative (triangulation).

As the number of mobile users utilizing GPS for location aware computing becomes larger, dynamic mapping of connectivity becomes more and more applicable. Furthermore, the WiNSEC team has been researching non-cooperative location finding through signal strength triangulation techniques with promising results. Although this method might never produce results as accurate as GPS based implementations, the density of measurements will greatly increase since each user will effectively become a sensor of the measurement network. The trade-offs between measurements density and position accuracy may provide a prolific research field. In the future, connectivity maps may provide the basis for

new forms of route planning that may contribute to our capacity to respond better to emergency situations.

7. Acknowledgements

The authors would like to acknowledge the ideas and insight of Dr. Paul Kolodzy, Director of the Wireless Network Security Center at Stevens Institute of Technology. Also, the members of the WiNSEC center at Stevens Institute of Technology, and in particular Jason Evans, Johnny Hsu and Ryan Gove who have helped in the implementation of the experiments.

8. References

- [1] S. J. Vaughan-Nichols, "The challenge of Wi-Fi roaming," *Computer*, vol. 36, no. 7, pp. 17-19, 2003.
- [2] L. Fu, "An Adaptive Routing Algorithm for In-vehicle Route Guidance Systems with Real-time Information," *Transportation Research, Part B: Methodological*, vol. 35B, no. 8, pp. 749-765, 2001.
- [3] T. G arling, T. Laitila, and K. Westin, *Theoretical foundations of travel choice modeling*, 1st ed. Amsterdam ; New York: Elsevier, 1998.
- [4] M. Kreitz, K. W. Axhausen, K. J. Beckmann, and M. M. Friedrich, "Mobiplan: an Internet-based personal mobility advisor," *Transport Policy*, vol. 9, no. 2, pp. 155-168, 2002.
- [5] E. Nash, "GNER to install WiFi on trains," <http://www.itweek.co.uk/News/1142986> 2003.
- [6] J. V. Nickerson, "A Concept of Communication Distance and its Application to Six Situations in Mobile Environments," *IEEE Transactions on Mobile Computing (in press)*, vol. 4, no. 1, 2005.
- [7] Stevens Institute of Technology, "WiNSEC," <http://www.stevens.edu/winsec> 2004.
- [8] S. Byers and D. Kormann, "802.11b access point mapping," *Communications of the ACM*, vol. 46, no. 5, pp. 41-46, 2003.
- [9] U. of Kansas, "Wireless Network Visualization Project," <http://www.itc.ku.edu/wlan/index.shtml> 2004.
- [10] Theodore S. Rappaport, "Wireless Communications: Principles and practice," Prentice Hall, 1996.
- [11] K.K. Leung et al., "Outdoor IEEE 802.11 cellular networks: MAC protocol design and performance." *IEEE International Conference on Communications*, vol.1, pp.595-599, 2002.