

Towards Mobile Computational Application Steering: Visualizing The Spatial Characteristics of Metropolitan Area Wireless Networks

J. S. Pascoe^{†‡}, V. S. Sunderam[‡], R. J. Loader^{†‡} and G. Sibley[‡]

[†]Department of Computer Science
The University of Reading
United Kingdom
RG6 6AY

J.S.Pascoe@reading.ac.uk
Roger.Loader@reading.ac.uk

[‡]Math & Computer Science
Emory University
Atlanta, Georgia
30322

vss@mathcs.emory.edu
gsibley@emory.edu

Abstract. Computational application steering from mobile devices is an attractive proposition and one that is being actively explored. In wireless networks, particularly metropolitan area infrastructural networks, a range of unique temporal and spatial performance characteristics may affect the ability to perform computational steering. The network coverage area contains zones, termed *trouble spots*, that are distinct pockets where garbled, poor, or no network coverage exist. They have been experimentally attributed to geographic features (e.g. tunnels and tall buildings). Point-to-point manifestations of trouble spots include stalled web pages, slow ftp connections and data corruption which, although detrimental, do not compromise system usability. However, computational applications such as steering are highly susceptible, and could suffer from serious operational and semantic problems in the presence of trouble spots. Previous experimental work has verified not only the existence of trouble spots in a range of wireless networks, but has also identified some of the issues surrounding their detectability. One of the difficulties encountered during the initial study was the collection of reliable data, primarily due to a lack of tool support. To alleviate this, a visualization package, termed *RadioTool*¹ has been developed; the underlying goal being to investigate the nature of trouble spots more reliably. It is envisaged that the tool will eventually serve as a detection and display mechanism for visualizing network quality as a user moves about a wireless network, and thereby complement and support computational steering applications. This paper describes the features incorporated in RadioTool and describes its use in relation to the exemplar Ricochet radio network.

Keywords – Computational steering, wireless networking, trouble spots

1 Introduction

The Ricochet radio network provides wireless Internet connections of up to 128 Kbps in several cities across the United States. While this paper predominantly focuses on RadioTool, it is noteworthy to explain the motivating factors for its development in the

¹ This project was supported in part by NSF grant ACI-9872167 and DoE grant DE-FG02-99ER25379.

overall project. Networks such as Ricochet and GUIDE [10] are facilitating many new directions in mobile computing and demand that legacy applications incorporate additional functionality [12]. In particular, programs that utilize distributed communication must be supplemented with schemes to adapt to connectivity problems or *trouble spots*. To facilitate trouble spot detection, a previous study suggested the following three metrics: *received signal strength indication (RSSI)*, *network response time* and *packet loss*. Furthermore, these metrics facilitate a user process to categorize specific trouble spot instances and in doing so, appropriate resolution strategies can be adopted. This functionality, coupled with a database of trouble spot locations and characteristics facilitates the basis of a preemptive scheme that minimizes the impact of trouble spots.

2 The Ricochet Network

In this section, a brief overview of the Ricochet network is presented. It must, however, be emphasized that the work described herein is architecturally independent of the actual wireless network; in fact, RadioTool and the other associated parts of the system may be used not only in MAN's, but also in wireless local area networks. The prototype implementation of RadioTool used the Ricochet network, but it has been designed to be easily portable to other network types. Ricochet is present in eight major US cities with further expansion into another thirteen being planned in the future. The network offers a fixed bandwidth of 128Kbps and is accessible to the general public through small portable wireless modems. The network itself consists of seven architectural components which are grouped hierarchically to provide the service.

- *Wireless Modems* – Ricochet modems are small devices that can be connected to any portable workstation.
- *Microcell Radios* – Microcell radios communicate with user workstations through wireless modems. In addition, they perform error checking and are responsible for sending acknowledgements.
- *Wired Access Points* – Wired Access Points or WAPs are installed within a 10-20 square mile area. Each WAP collects and converts wireless data into a form suitable for transmission on a wired IP network.
- *Network Interface Facility* – A Network Interface Facility connects a group of WAPs to the Network Operations Centers.
- *Ricochet Gateway* – The Ricochet gateway is part of each Network Operations Center and connects the Ricochet network to other systems.
- *Network Operations Centers* – The Network Operations Centers provide a means for monitoring and control within the network.
- *Name Server* – The name server maintains access control and routing information for every radio and service within the wireless network.

The overall topological structure of the networks wireless component is a mesh, since microcell radios are placed at regular intervals in a 'checkerboard' fashion. The authors argue in [6] that the additional redundancy in a mesh provides a more robust wireless network topology than that of a tree or a star.




Fig. 1. The Ricochet Communication Architecture

3 Trouble Spot Detection

The project and system described herein is based upon the premise that all wireless networks contain areas of degraded quality, some of which are static (due to topography or structures) while others are dynamic (caused by weather and related factors). While the existence of these “trouble spots” is an accepted fact, one outstanding issue is the pertinent question of how to detect and ascertain the semantics relating to a specific trouble spot. In this regard, it is useful to classify trouble spots into different categories ranging in intensity – from trouble spots that cause occasional garbling and packet loss to those that result in complete disconnection. Detection may be achieved in the first instance by the mobile device continuously monitoring network quality indications, and comparing these values against recent history as well as against a database of stored past values. At this stage, the system must determine (by executing a series of diagnostic tests) which category of trouble spot it has encountered so that it may invoke an appropriate failure handler. In order to achieve this, we propose the following metrics for evaluating the state of a service in a fixed topology wireless network.

- *Received Signal Strength Indication* – determining signal strength is usually accomplished by query the hardware of the network interface and is intuitively central to the detection of trouble spots.
- *Packet loss* – packet loss at the wireless network interface is a function of not only the power of the received signal, but also the path loss and the state of the network. Packet loss is inversely proportional to signal strength and as such, a high packet loss indicates a drop (or imminent drop) in service. Thus, packet loss is an important metric in determining trouble spots of the more severe categories.
- *Response time* – The response time of the network is the duration between a message being transmitted and the reply being received. It can often be useful in confirming trouble spot classification when the packet loss and path loss are both low. In this circumstance, a high response time will indicate that the network problem is related to the wired component of the infrastructure.

The above metrics were suggested by a number of sources, the most notable being [8], [14] and the IETF MANET² metrics for the evaluation of ad hoc routing protocol performance [5].

4 RadioTool Technical Operation

RadioTool is an event driven X-Windows (Gtk based) Graphical User Interface and in the prototype version for Ricochet, it utilizes the full range of services offered by

² Mobile Ad-hoc NETworks working group.

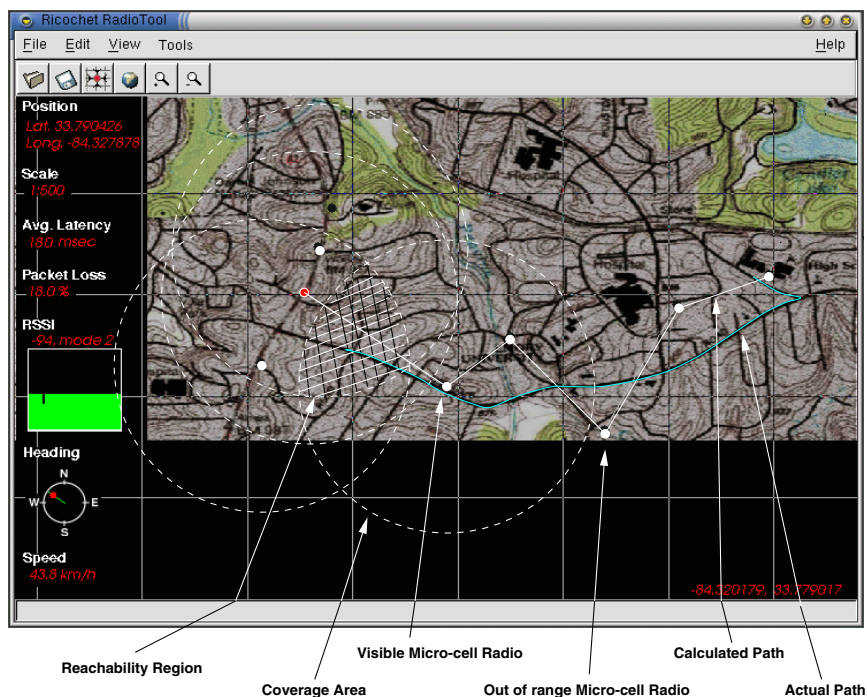


Fig. 2. RadioTool Graphical User Interface

the Ricochet Radio Network API (RRNAPI) [13], that is, an application programming interface developed to gather information specifically from the Ricochet network. By substituting alternate libraries, RadioTool can be utilized in conjunction with all other wireless networks.

Figure 2 depicts the layout of the RadioTool interface. When RadioTool starts, it calls `RRN_init`, a function within the RRNAPI that dictates how data should be collected. RRNAPI then connects to the wireless modem and launches a thread to regularly query the hardware and update RadioTool's data structures. After invoking `RRN_init`, RadioTool installs a callback function that is executed once every 250 milliseconds.

Location is measured in latitude / longitude format whereas heading is represented by two-dimensional polar coordinates. Received signal strength indication, packet loss and latency are all measured in standard units. Note that, latency is measured as the average round trip latency to the Ricochet wired access points (i.e. latency measurements do not traverse the Internet).

Furthermore, RadioTool calls `RRN_poletops`, a function which returns a list of the currently visible micro-cell radios. This list is used to visualize the users position relative to the poletops, as described in section 5.

4.1 Location Determination

Currently, location is determined through one of three methods. The first approximates user location to the position of the nearest microcell radio, that is, the transmitter with

the strongest signal. The second method utilizes readings from all reachable radios by taking an average of the contactable radios locations, hence arriving at a ‘center of mass’ approximation. Method three calculates an offset center of mass based on the relative signal strengths of the MCRs. Thus, in an unbiased system where only two radios are contactable, the hosts location will be approximated to a bisection.

Empirical results (obtained by comparing computed values to those measured via GPS) have shown that method one is the more accurate of the three, typically within 200 meters. Method two performs badly as it is often possible for distant transmitters to be contactable (albeit barely). The hypothesis of the third method is to address this skewing caused by distant radios through a weighting function for the average. In practice, a factor of attraction is empirically set at -125 dB, which is approximately the weakest signal typically observed. Increasing this factor causes method 3 to behave as accurately, and sometimes better, than method 1. Current experimentation is considering which weight results in the best estimate of location.

Additional Methods for Calculating Location Calculating location based on signal strength of known location radios is similar to the problem in mobile robotics of *spatial localization* based on beacon placement [2]. From these studies we can see that there are better ways to calculate location based on RSSI than the averaging scheme used above [2]. In short, there are essentially two general localization methods used in much of the research: proximity based localization [15, 4, 3] and multilateration³ based localization [1, 7, 9].

Both approaches use a standard method to approximate distance, such as radio frequency localization [15], RSSI [11], or a database of RSSI signatures [1]. Thus, RRNAPI’s location determination algorithms are based on the approach that uses signal strength to approximate distance. In either a proximity based approach or a multilateration based approach, the primary improvement in RRNAPI would be to add the concept of *reachability regions* to the `RRN_location` function. Reachability regions are areas where different radios coverage areas overlap in the manner of a Venn diagram. For example, if a host can detect two radios, then the reachability region is the intersection of their signal coverage areas. As the number of radios increase, the intersection shrinks thus increasing the accuracy of the approximation.

5 Radio Network Visualization

RadioTool iterates through a list of microcell radios, drawing each radio at its latitude / longitude position as well as indicating an approximate reachability region⁴. RadioTool also stores another list of inactive microcell radios which are shown on the interface but do not correspond to coverage circles. Latitude / longitude grid lines are drawn, along with an indication of scale. The cursor can be used to ‘grab’ the map and move the users view point; the cursor also displays its position relative to the map it is over.

With the *Map Tool*, geographic views can be loaded as background images. Maps scale and move as the user zooms or moves his view. This feature is useful in determining the position of microcell radios, WAPs and other users. The ‘status’ overlay (see fig. 2) displays the users location, current scale, packet loss, latency, recent RSSI

³ In multilateration, position is estimated from distances to three or more known points.

⁴ Coverage is idealized to a sphere of linearly decreasing signal strength.

measurements, heading, speed, and the modem connectivity statistics. User location is indicated by a small dot that is dynamically updated as the user moves.

6 Conclusions and Future Work

This paper has presented RadioTool, a simple yet robust package for visualizing and collecting data from wireless networks. The insight, information and data that RadioTool provides is being used to develop a solution to the problem of wirelessly steering computational applications. Although RadioTool has some caveats, the data it provides has proved to be largely more reliable than that gathered from ad-hoc experimentation. Future work will focus solely on these goals, the envisaged culmination being a freely available software system.

References

1. P. Bahl and V.N. Padmanabhan. Radar: An in-building user location and tracking system. In *Proc. of the IEEE Infocom 2000*, volume 2, pages 775–785, Tel-Aviv, Israel, March 2000.
2. N. Bulusu, J. Heidemann, , and D. Estrin. Adaptive beacon placement. In *21st International Conference on Distributed Computing Systems*, Phoenix, AZ, April 2001.
3. N. Bulusu, J. Heidemann, , and D. Estrin. Gps-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, pages 7(5):28–34, October 2000.
4. F.A. Chudak and Shmoys. Improved approximation algorithms for capacitated facility location problems. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages S875–S876, Baltimore, Maryland, USA, January 16-17 1999.
5. S. Corson and J. Macker. *Routing Protocol Performance Issues and Evaluation Considerations*. IETF Mobile Ad Hoc Networking (MANET) Working Group, 1999. RFC 2501.
6. S. Lee et al. Ad hoc Wireless Multicast with Mobility Prediction. In *Proc. IEEE International Conference on Computer Communications and Networks*, pages 4–9, 1999.
7. L. Girod. Development and characterization of an acoustic rangefinder. Technical Report USC-CS-TR-00-728, University of Southern California, April 2000.
8. F. Halsall. *Data Communications, Computer Networks and Open Systems*. Addison-Wesley, fourth edition, 1995.
9. B. Hofmann-Wellenhoff, H. Lichtenegger, and J. Collins. *Global Positioning System: Theory and Practice*, volume Fourth. Springer Verlag, 1997.
10. Keith Cheverst and Nigel Davies et al. The Role of Connectivity in Supporting Context-Sensitive Applications. In H. W. Gellersen, editor, *Handheld and Ubiquitous Computing*, volume 1707 of *Lecture Notes in Computer Science*, pages 193–207. Springer-Verlag, 1999.
11. F. Koushanfar, A. Savvides, G. Veltri, M. Potkonjak, and M. B. Srivastava. Iterative location discovery in ad hoc sensor networks. Slides from UCLA CS review, April 2000. <http://www.cs.ucla.edu/farinaz/location2-1.ppt>.
12. J. S. Pascoe, G. Sibley, V. S. Sunderam, and R. J. Loader. Mobile Wide Area Wireless Fault Tolerance. Technical report, University of Reading and Emory University, 2001.
13. G. Sibley and V. S. Sunderam. Tools for collaboration in metropolitan wireless networks. In V. Alexandrov, J. Dongarra, B. Juliano, R. Renner, and K. Tan, editors, *Proc. 2001 International Conference on Computational Science (ICCS 2001)*, volume 2073 of *Lecture Notes in Computer Science*, pages 395–403, May 2001.
14. Y. Tu, D. Estrin, and S. Gupta. Worst Case Performance Analysis of Wireless Ad Hoc Routing Protocols: Case Study. Technical report, University of Southern California, 2000.
15. R. Want, A. Hopper, V. Falcao, , and J. Gibbons. The active badge location system. In *ACM Transactions on Information Systems*, pages 10(1):91–102, January 1992.