

The Case for a Hybrid Passive/Active Network Monitoring Scheme in the Wireless Internet

Björn Landfeldt*, Pipat Sookavatana*,** and Aruna Seneviratne**

Dept. of Electrical Engineering and Telecommunications

The University of New South Wales, Sydney 2052 Australia

**Department of Computer Engineering

Mahanakorn University of Technology, Bangkok 10530, Thailand

[bjornl, pipats, aruna]@ee.unsw.edu.au

Abstract

This paper discusses monitoring within wired and wireless networks and the type of monitoring information needed to support different applications. We suggest a hybrid active/passive monitoring approach with a dynamic time window mechanism and interchangeable filters to extract requested information. The paper also shows our initial experimental results and presents our conclusions.

Keywords-- Network monitoring, wireless links, dynamic window generation, passive monitoring and active probing

1. Introduction

The past few years have seen an exponential growth of networked applications based on the Internet. Many of these applications are interactive and use data with real time characteristics. In parallel, there has been a rapid growth in cellular and other mobile communication systems such as wireless LANs, and infra red access nodes. The wireless Internet will comprise of a fixed IP based core network and multiple overlaid wireless access networks. This will provide users with simultaneous access to several different networks [14]. In the ideal case, this will provide the capability of matching the application and user requirements to network characteristics. However, mapping of user/application requirements to available networks is not trivial. To facilitate this process, we have designed a reactive quality of service manager called "User Services Assistant", USA [15]. A fundamental requirement for USA to be able to provide the assistance is having the ability to ascertain the system status, especially the network characteristics. Thus, it is of vital importance to USA to have an effective network-monitoring tool.

In this paper, we discuss the requirements for performing network monitoring in wireless environments and suggest a framework that provides flexibility in statistical processing and presentation of monitoring. The rest of this paper is organised as follows. Section 2 presents related work and section 3 discusses general aspects of network monitoring. Section 4 then examines the relevance of the proposed monitoring framework, section 5 describes our experimental set-up and results and finally in section 6 we present our conclusions and future work.

2. Related work

A good survey of network monitoring techniques and the available tools can be found in [2]. These monitoring schemes can be divided into two categories, namely active and passive monitoring.

2.1. Active Monitoring Schemes

Active monitoring schemes operate by inserting probing traffic into the network. These probes are used to extract information about network characteristics such as available and bottleneck bandwidth, round trip delay etc. Bottleneck bandwidth can be detected by a technique called back-to-back or packet-pair probing [6]. Several monitoring tools use this technique [5,6,7,11]. For example, Pathchar [7] uses this method in combination with traceroute [3], in order to derive per-hop network information.

A more common way of estimating the available bandwidth is to simulate a bulk data transfer. Ttcp [8] makes an estimation of the available bandwidth based on

data from a TCP bulk data transfer. Treno [9] uses ICMP echo packets with flow and congestion control in order to simulate a TCP bulk data transfer, and from that derive the available bandwidth.

Active monitoring provides an accurate method of determining network characteristics. However, the monitoring traffic competes with application data flows for network resources and is therefore not scalable, especially for resource scarce environments. Furthermore, the delays associated with obtaining necessary results may be large, thus making it unsuitable for use with reactive QoS managers such as USA[15].

2.2. Passive Monitoring

Passive schemes overcome the disadvantages of active schemes associated with overheads and delay by monitoring streams in progress. From these observations, they infer the current network status, rather than introducing monitoring traffic. SPAND [1] is a monitoring scheme that is based on the principles of passive monitoring. It extends the basic passive monitoring by providing facilities for sharing of measurement results among hosts in order to increase the accuracy.

Although the assumption is valid in the current fixed Internet environments as shown in [4], it is not directly valid in a wireless Internet environment. The performance of a wireless access network not only depends on the current load, but it also depends on other factors such as multipath fading, other forms of interference, and the transmitted power. These factors in turn depend on the geographic location of the mobile host. Consequently, sharing measurement results collected by mobile hosts are likely to lead to unreliable estimates. Therefore, shared passive monitoring cannot be directly used in the wireless Internet environment.

3. Hybrid Monitoring Frame Work

Neither active, nor passive monitoring can be directly used for monitoring wireless Internets. However, it is possible to use a hybrid scheme that uses passive monitoring where possible and resorts to active monitoring when passive information is unavailable. The scheme also need be coupled with a mechanism, which accounts for the rapid fluctuations of performance of wireless environments. The proposed scheme is such a hybrid scheme, which provides a universal monitoring tool for the emerging wireless Internet environment.

3.1. Hybrid Monitoring

The hosts attached to the wireless segments will only be able to use passive information if the host has an open connection to the wanted destination due to the volatility of the wireless communication environments. If no open connects are present, the mobile host need to resort to sending active probing.

This is realised by treating the data-path as consisting of two independent segments, namely the wired and wireless parts. And using an extended passive monitoring server. The proposed scheme is shown in Figure 1. The extended passive monitoring server - EXPAND provides support for passive or active monitoring of the wireless access networks as described below, in addition to the passive or active monitoring of the fixed network segment.

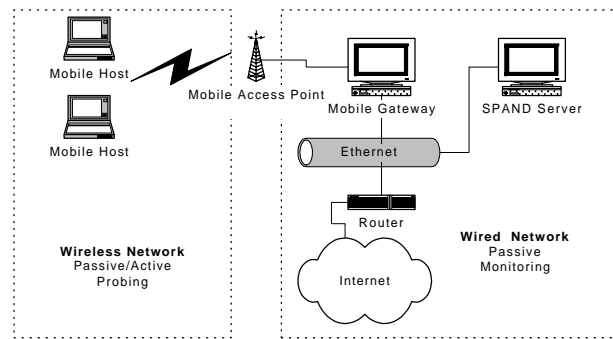


Figure 1. The hybrid passive and active monitoring approach

In the case there are no open connections, the mobile host transmits a request packet to the EXPAND server. This will requesting information from the EXPAND server about the connection to the destination. The EXPAND server will respond with the requested information about the fixed network segment. The request-response packets will be used by the requesting mobile host as active probes to determine the characteristics of the wireless access network. The requesting host will use the results from the EXPAND server about the fixed segment of the network, and the information from the request response packets about the wireless segment of the network, to estimate the end-to-end characteristics of the connection.

A wireless link exhibit highly varying characteristics. It is therefore not possible to make reliable estimates of the characteristics of the link using raw data from the active probes. Furthermore, the capture effects reported in [16], will lead to false estimates over CSMA links such as 802.11. This is overcome by searching for the boundaries

by making a snapshot of bottleneck bandwidth and packet loss, and comparing the results with the fixed link estimates. Since the active probing only takes place on demand, and only when passive information is unavailable, in the above scheme, the probing traffic overhead is minimised.

3.2. Dynamic Windowing

As mentioned, the raw measurement data provided by the split-monitoring scheme in itself, is insufficient to obtain a good estimate of the network characteristics. Firstly, in order to obtain a good estimate the raw data needs to be filtered to suit the application requesting the data. Secondly, as the characteristics of the wireless access segment will vary significantly, it will not be possible to use simple statistical techniques. This is overcome by using a dynamic windowing mechanism to filter the raw data. The dynamic windowing scheme exploits the fact that the Internet displays quasi-stable characteristic [4]. This will be true even for the wireless networks as these networks will be used to access information from stationary locations, rather than whilst moving, i.e. from a meeting room, a laboratory or a client's premises. This will be even more true in the emerging wireless LAN environments as they reduce the data rate when the signal quality drops, thus giving different link speeds at different locations within the network [12, 13]. Thus, within these networks, the data rate will stay stable within a region and the only difference between the wireless and wired segments will be the period of stability.




Figure 2. Dynamic window data regions

The dynamic windowing scheme thus attempts to determine a step change in the network characteristics, and then uses a simple statistical analysis of the data from the step in order to determine the characteristics. This is schematically shown in Figure 2. At time $t = t_x$, $0 < t_x < t_1$, the data set used for calculation is between t_x and t_0 . At time t_y ; $t_1 < t_y$, i.e. when the step change has been detected,

the data set used for the calculation is between t_y and t_1 and the data between t_0 and t_1 is discarded.

Thus, the windowing scheme consists of two parts. A mechanism for detecting step changes in instantaneous measured data and a statistical filtering mechanism. The step changes are determined by the variance of the data normalised by the mean value. The variance indicates the variability of the measurements and will therefore reflect the step change. The normalisation is done so that the variability can be quantified. In order to find the step, we define a threshold value δ and compare it against the cumulative variance to mean ratio, i.e. a step change is found if

$$\delta \leq \frac{\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2}{\frac{1}{n} \sum_{i=1}^n x_i}$$

is true where x is a sample value and n the total number of samples.

4. Experimental Results

To verify the viability of the above framework we conducted several tests within the experimental test-bed shown in Figure 3. RTT measurements were used for convenience in these experiments. The WaveLAN network covered one floor of the Electrical Engineering building at the University of New South Wales, and provided varying SNR at different locations. The wireless tests consisted of a person moving with a laptop within the wireless LAN, measuring the available bandwidth to the PC from fixed locations. The wired tests were carried out over the Internet between the PC in Australia and the servers in Sweden and Thailand respectively.

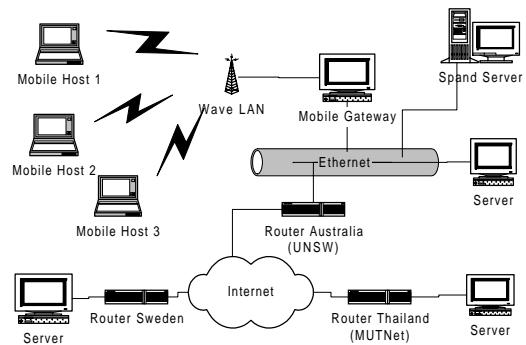


Figure 3. The Experimental Test-bed

4.1. Active versus Passive Monitoring

In order to demonstrate the advantages of passive monitoring we conducted a series of tests between Australia, Sweden and Thailand.

Passive monitoring introduces fewer overheads to the network traffic, and the overhead traffic is only introduced locally and can therefore be better controlled and accounted for when dimensioning networks. We illustrate this in Figure 4.

In the example, we compare some of the most widespread tools for active probing with the passive EXPAND. The figures for EXPAND are taken from our experimental implementation and estimated as; 20 reports to the server of 45 bytes each over one hop, plus active probes of a wireless link coupled with the information request, of 10 packets of 1500 byte each. Together this introduces 15900 byte over one hop. The other tools are measured with default values and the figure compares the introduced traffic over all links. The figure shows how the active tools inject significantly more traffic into the network than EXPAND does.




Figure 4. Comparison Active and Passive traffic

4.2. Wireless experiments

Figure 5 shows a typical plot of available bandwidth measured over the WaveLAN wireless LAN. Our wireless tests aimed at investigating the traffic behaviour over the wireless link. In figure, at the initial state, only one client accessed to the fixed network passing through the mobile gateway while the average available bandwidth was 1.4 Mbps. Consequently at the time 210-sec., another client came and joined to the wireless network. As shown in the figure, the available bandwidth dropped significantly to 0.6 Mbps. The client then moved out from the network at time 530-sec and the available bandwidth rapidly increased to the maximum to value. In our measurements,




Figure 5. Available Bandwidth Measurement

we could consistently see this behaviour. This behavior might occur periodically with in Mobile IP, where many mobiles are roaming form network to network.

Because of the rapid changes, the need for a dynamic windowing scheme is accentuated. The stable behaviour also means historical data is useful even within a wireless network.

4.3. Dynamic window experiments

In order to instigate the viability of the dynamic full window scheme, we used the experimental result form the pervious path and also conducted a series of RTT measurement between Australia and Thailand.

The following example demonstrates the performance improvement of the dynamic window scheme. As can be seen the measurements vary significantly with time. The Figure 6 shows available bandwidth estimations using three different methods, mean value, least square linear approximation and least square linear approximation using the dynamic window scheme. In order to find the edges of the dynamic windows we used the variance to mean formula described above and compared it to a δ value of 0.04. Determining the appropriate δ value is still an open research issue. For now, our experiences show that a value between 0.02 and 0.07 makes the scheme perform well. The figure clearly shows that the least square linear approximation with dynamic windowing estimates the available bandwidth better than the other two methods.

To show that dynamic windowing scheme also performs good estimations in slow changing behaviors, we measured a round trip time (RTT) between Australia (UNSW) and Thailand (MUT) during 10 hours. Again, although the RTT was not change rapidly, the least square linear approximation with dynamic windowing still performs better than the other two methods.




Figure 6. Available Bandwidth Estimation using three different methods

Our experimental results show that using the dynamic windowing scheme gives consistently better result than either using instantaneous measurements or fixed windows.




Figure 7. RTT Estimation using three different methods.

5. Conclusion and Future Work

In this paper, we discussed how network monitoring should be carried out in a mixture of wired and wireless environments. The purpose of this discussion is to form guidelines for an overall framework for network monitoring. We have proposed a split network monitoring approach with passive/active monitoring. Central to this framework is the method of dynamic windowing to find a region with a relevant set of data on which to conduct statistical computations. In the future, we intend to focus on statistical computations of monitoring results. Firstly, we will investigate possible ways of determining the optimal time window size for filtering transient events while still reacting on sustained bandwidth shifts.

Secondly, we will work on statistical methods for use in different filters in order to provide accurate and for applications appropriate estimates of the performance parameters. Thirdly, we will work on implementation issues such as a protocol between a terminal and an EXPAND server for specifying filter options etc. Finally, we intend to implement a full prototype of the framework in order to evaluate its usefulness.

6. References

- [1] S. Seshan, M. Stemm and R. Katz. "SPAND: Shared Passive Network Performance Discovery". In Proc 1st Usenix Symposium on Internet Technologies and Systems (USITS '97) Monterey, CA December 1997.
- [2] heplibw3.slac.stanford.edu/xorg/nmtf/nmtf-tools.html.
- [3] W. R. Stevens, TCP/IP Illustrated Volume 1 Addison-Wesley, Reading MA, Nov 1994.
- [4] V. Paxson. "Measurements and Analysis of End-to-End Internet Dynamics. PhD thesis, U. C. Berkeley, May 1996.
- [5] R. L. Carter and M. E. Crovella. "Measuring bottleneck-link speed in packet switched networks". Technical Report BU-CS-96-007, Computer Science Department, Boston University, March 1996.
- [6] S. Keshav. "Packet-Pair Flow Control". IEEE/ACM Transactions on Networking, Feb 1995.
- [7] Pathchar – A Tool to Infer Characteristics in Internet Paths. [ftp://ee.lbl.gov/pathchar](http://ee.lbl.gov/pathchar), 1997.
- [8] IEEE RFC 1470 June 19937
- [9] M. Mathis and J. Mahdavi "Diagnosing Internet Congestion with a Transport Layer Performance Tool". In Proc. INET '96, Montreal Canada, June 1996.
- [10] D. Clark and D. Tannenhouse "Architectural considerations for a new generation of protocols". In SIGCOM'90 Philadelphia Pennsylvania September 1990.
- [11] K. Lai and M. Baker, "Measuring Bandwidth." Proceedings of IEEE Infocom'99, March 1999.
- [12] www.3com.com/promotions/wireless
- [13] www.wavelan.com/products
- [14] R. Katz and E. Brewer, "The Case for Wireless Overlaid Networks", SPIE Multimedia and Networking Conference (MMNC'96) San Jose January 1996.
- [15] B. Landfeldt, A. Seneviratne, C. Diot, "User Services Assistant, an End-to-End Reactive QoS Architecture", IWQOS98 Napa California May 1998.
- [16] K. K. Ramakrishnan and H. Young, "The Ethernet Capture Effect: Analysis and Solution", in Proceedings of IEEE 19th Conference on Local Computer Networks, MN, USA, October 1994.