

Inferring Queue State by Measuring Delay in a WiFi Network

David Malone, Douglas J Leith, Ian Dangerfield

Hamilton Institute, NUI Maynooth.

Abstract. Packet round trip time is a quantity that is easy to measure for end hosts and applications. In many wired networks, the round trip has been exploited for purposes such as congestion control and bandwidth measurement because of relatively simple relationships between buffer occupancy and drain time. In 802.11 networks, the buffer drain times show considerable variability due to the random nature of the MAC service. We examine some of the problems faced when using round-trip-time-based queue estimates in these networks, particularly in relation to congestion control.

1 Introduction

Network round-trip time is a useful measurement that is easily estimated by end hosts. It is often used as a measure of network congestion either implicitly (e.g. a human looking at the output from traceroute or ping) or explicitly (e.g. TCP Vegas [3], FAST [13] or Compound TCP [12] use RTT as a proxy measure of buffer occupancy). The assumption is that queueing is the main source of variation in RTTs, and so RTTs can be used to estimate queueing. This has led to tools such as pathchar [7].

In wired networks, this is often a reasonable assumption: there is usually a linear relationship between queue length and queue drain time. However, this relationship is not universal. In WiFi networks there can be a significant random component associated with transmitting packets. A device usually has a back-off period before sending. The duration of this period is a combination of a randomly selected number and the duration of busy periods due to other traffic on the network [6]. Also, a packet may suffer a collision or corruption, requiring further back-off periods and retransmission.

Figure 1 shows observed queue drain times plotted against queue length from a device transmitting packets over a contended WiFi link. A striking feature of this graph is the overlap between RTTs associated with different queue lengths: RTTs observed for a queue length of one packet could have come from a queue length of 10 packets; RTTs from a queue of 10 packets could easily have come from a queue of 20 packets. Even before other sources of delay are considered, this is a challenging environment for making inferences about queue length from RTTs.

Previous work has touched on the impact of this variability. A comparison of bandwidth estimation tools over wireless was conducted in [11]. They suggest that some errors made by bandwidth estimation tools may be due to variable service, but they do not conduct an in-depth investigation of this. Other work, such as [5] looks at various TCP metrics over WiFi. They consider RTT averaged over connection lifetimes, but are not concerned with the relationship between measured RTT and buffer occupancy.

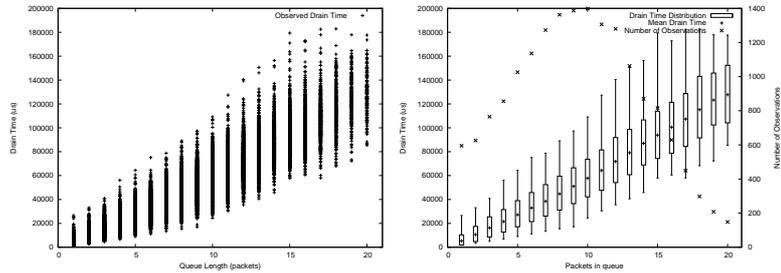


Fig. 1. Impact of Queue Length on Drain Time. (a) Scatter plot of observed values. (b) 10–90% box-and-whiskers plot and mean; number of samples is also shown on right hand axis.

In this paper we investigate the complications introduced by the random nature of the service in 802.11. We note that there have been complications in application or transport layer measurement of RTTs in wired networks (for example, some filtering is necessary to remove artifacts caused by TCP’s delayed ACKing or TSO [8]). However, in order to focus on the issues raised by the varying delay of a wireless link, in this paper we will assume that accurate RTT measurements are available.

We show that raw RTT measurements don’t allow sharp conclusions to be drawn about the queue length, but are well correlated with it. We also show that variability in measurements grows as \sqrt{n} . We then look at filters that might be applied to the RTT measurements and find that normal RTT filters *decrease* correlation. Linux’s Vegas implementation deals relatively well with these challenges and we consider why this is.

2 Testbed Setup

We consider network delay associated with winning access to transmission opportunities in an 802.11 WLAN. We measure both the queue drain time (the time from when a packet reaches the driver to when the transmission is fully complete) and the MAC service time (the time from reaching the head of the hardware interface queue to when transmission is fully complete), using techniques described in [4]. The MAC service time can vary by orders of magnitude, depending on network conditions.

The 802.11 testbed is configured in infrastructure mode. It consists of a desktop PC acting as an access point, 15 PC-based embedded Linux boxes based on the Soekris net4801 [2] and one desktop PC acting as client stations. The PC acting as a client records measurements for each of its packets, but otherwise behaves as an ordinary client station. All systems are equipped with an Atheros 802.11b/g cards.

All nodes, including the AP, use a Linux kernel and a version of the MADWiFi [1] wireless driver modified to record packet statistics at the driver layer with a fixed queue of 20 packets. While we focus on the queueing at the driver layer, Section 3 shows the statistics of drain time as the number of packets increases. All of the tests are performed using the 802.11b physical maximal transmission rate of 11Mbps with RTS/CTS disabled and the channel number explicitly set.

3 Raw RTT Signal

The data for Figure 1 is taken from a run from our testbed where 4 stations are uploading using TCP. Measurements are taken from one TCP sender, so all packets are 1500 bytes. The results are taken over about 110s where network conditions are essentially static.

Briefly consider the simple problem of determining if the queue in Figure 1 contains more than ten packets based on the observed drain time. For example, consider a simple threshold scheme: set a threshold and if the observed time is greater than the threshold, we infer it has more than ten packets, otherwise we infer it has less than ten packets. Even if we have prior knowledge of the drain time distribution in Figure 1, how effective can such a scheme be for WiFi?

Figure 2 shows how often this scheme makes a mistake for a range of different thresholds. The upper curve in Figure 2(a) is the chance that the delay threshold incorrectly indicated that the queue length was either above or below 10 packets. The upper curves in Figure 2(b) show how this breaks down into situations where the queue was small but the delay was big or the queue was big but the delay was small. The best choice of threshold, around $60,000\mu s$ (about 10 times the mean service time), makes a mistake just over 10% of the time. Thus a congestion control scheme based on such a threshold could make an incorrect decision about once in every ten packets.

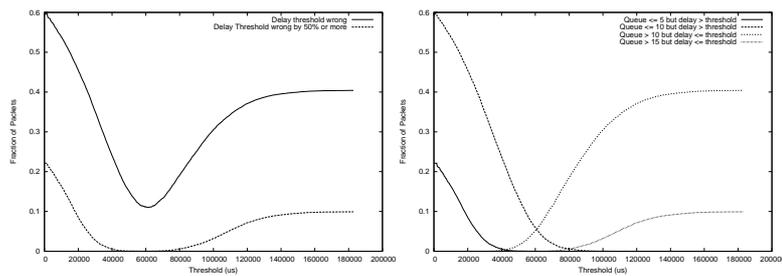


Fig. 2. Simple thresholding of delay measurements: (a) how often simple thresholding is wrong, (b) breakdown of errors into too big or too small.

Of course, it is possible that these mistakes occur mainly when the queue is close to 10 packets. To check this we also calculate the chance that while the queue had five or fewer packets that the delay is less than the threshold and the chance that the queue has more than fifteen packets while the delay is short. These represent large mistakes by the threshold technique. The results are the lower set of curves in Figure 2(a), with a large flat section for threshold values from 40,000 to $80,000\mu s$. While it is making mistakes regularly these are not gross mistakes. A range of thresholds produce reasonable results.

This suggests that though delay measurements are quite noisy, there is hope of learning information about queue length from them. Basic statistics for the drain times against queue lengths are shown in Figure 1(b). We show the mean drain time and a box-and-whiskers plot showing the range and the 10th and 90th percentiles.

Figure 3(a) shows the estimated autocorrelation for the MAC service times, queue drain times and queue lengths. We see that the MAC service times show no evidence of correlation structure. This is what we intuitively expect from an 802.11 network operating without interference. In contrast, the queue lengths show a complicated correlation structure. The queue lengths are sampled at the time a transmission completes; because the queue length will not change much between these times we expect strong correlation over lags comparable to the queue length. The longer term structure in the queue length will be a function of TCP’s congestion control behaviour in this network. Finally, the queue drain times show a *similar* structure to that observed for the queue lengths. This is encouraging: the drain time and the queue length are in a sense carrying similar information. We can confirm this by calculating the Pearson correlation value of 0.918.

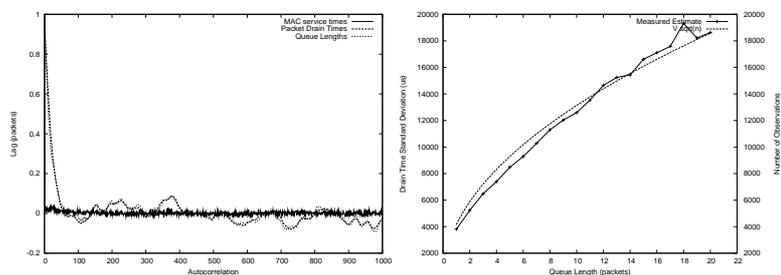


Fig. 3. Drain time statistics. (a) Autocorrelation for the sequence of MAC service times, queue drain times and queue lengths. (b) Estimate of standard deviation of drain times as a function of queue length.

Based on the low autocorrelation of the MAC service times, it may be reasonable to approximate the drain time of a queue of length n as the sum of n independent service times. The variance of the sum of random variables grows like the sum of the variances. Thus we expect the range of the 10–90% percentiles to scale like \sqrt{n} . This is confirmed in Figure 3(b), where we plot standard deviation of the drain times and compare them to \sqrt{n} . Larger buffers will make queue estimation even more challenging.

4 Smoothed RTT Signal

Most RTT measurements are smoothed before use, and based on the statistics we have seen in the previous section, there is a reasonable possibility that this may help in understanding queue behaviour. In this section we look at the impact of a number of commonly used filters on our ability to estimate the queue length.

A well-known example of the use of a smoothed RTT is the sRTT used in TCP to estimate round-trip timeouts. This estimator updates the smoothed estimate every time a new estimate arrives using the rule

$$\text{srtt} \leftarrow 7/8\text{srtt} + 1/8\text{rtt}. \quad (1)$$

We'll refer to this as 7/8 filter. It is also used in Compound TCP for delay based congestion control. We can do similar smoothing based on the time between packets:

$$\text{srtt} \leftarrow e^{-\Delta T/T_c} \text{srtt} + (1 - e^{-\Delta T/T_c}) \text{rtt}. \quad (2)$$

ΔT is the time since the last packet and T_c is a time constant for the filter. This filter approximately decreases the weight of RTT estimates exponentially in the time since the RTT was observed. We'll refer to this as the Exponential Time filter.

TCP Vegas and derivatives use a different smoothing. Ns2's implementation of Vegas uses the mean of the RTT samples seen over a window of time that is about the same as the current RTT. In order to avoid spurious spikes due to delayed acking, the Linux implementation of Vegas uses the minimum of the RTT's seen over a similar window. We'll refer to these as the Windowed Mean and Windowed Minimum filters.

We applied these filters to the drain time data to see if the resulting smoothed measurement was a better predictor of the queue length. We used a window size/time constant of 100ms, which is comparable to the actual RTT in our experiment. The results of our simple threshold test and calculation of autocorrelation are shown in Figure 4.

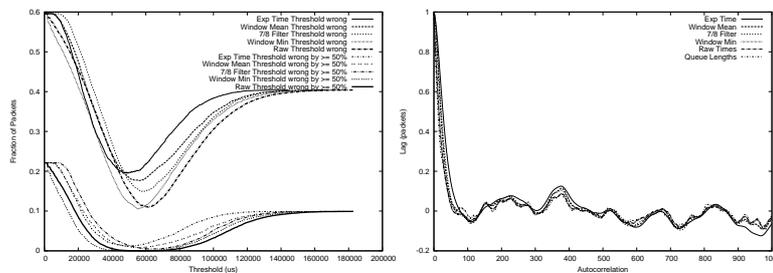


Fig. 4. Thresholding of filtered delay measurements: (a) errors while thresholding simple measurements, (b) autocorrelation of filtered measurements.

Interestingly, the filters *except for* Windowed Minimum have made things worse. Achievable error rate for thresholding has increased from 11% to 15, 18 and 20% for 7/8s, Windowed Mean and Exponential Time filters. The Windowed Minimum achieves an error rate of around 10.5%, which is comparable with the raw drain time error rate.

The autocorrelation graph tells a similar story: raw times and Windowed Minimum follow the queue length most closely. The Windowed Minimum measurements have the highest Pearson correlation with the queue length (0.922) closely followed by the raw measurements (0.918). There is then a gap before the 7/8th filter, the Windowed Mean and the Exponential Time results (0.836, 0.797 and 0.752 respectively).

5 Variable Network Conditions

As noted, the length of 802.11's random backoff periods are not just based on the selection of a random number, but also on the duration of busy periods due to the trans-

missions of other stations. In addition, the number of backoff periods is dependent on the chance of a collision, which is strongly dependent on the number of stations in the network and their traffic. Thus the RTTs observed by a station depend on cross traffic that may not even pass through the same network buffers.

For example, consider Figure 5. This shows the time history of queue lengths and drain times as we shut down the competing stations from the setup described in Section 3. By 242s there is little competing traffic in the system, and Figure 5(a) shows that the mean drain time and variability have been radically reduced. However, if we look at Figure 5(b) we see that this is not because the queue size has been reduced. In fact TCP Reno is keeping the queue closer to full because of reduced RTTs and contention.

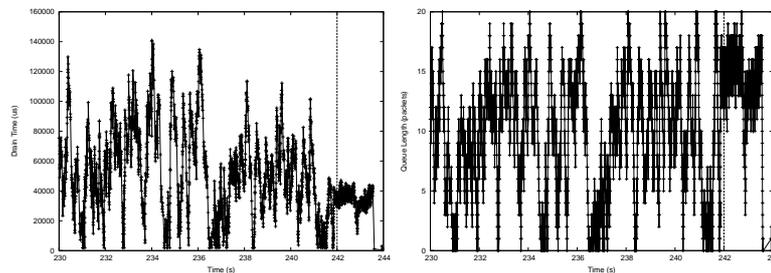


Fig. 5. The impact of other stations leaving the system: (a) drain times and (b) queue lengths.

When stations join the system the impact can also be dramatic, as shown in Figure 6. 4 TCP uploads are joined by another 4 TCP uploads just after 120s (note, to get 8 TCP uploads to coexist in a WLAN, we have used the ACK prioritisation scheme from [9], resulting in smoother queue histories). We see basically no change in queue length, but almost a doubling of round trip time.

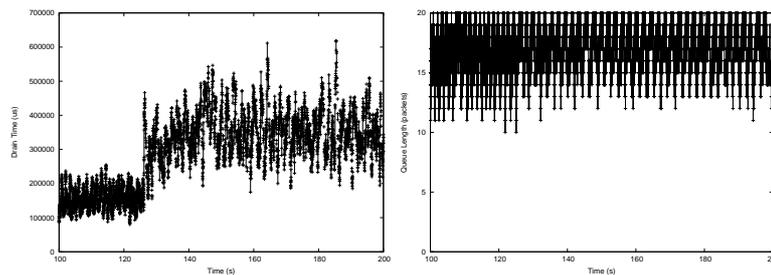


Fig. 6. The impact of other stations joining the system: (a) drain times and (b) queue lengths.

These changes in drain time are caused by a change in the mean service rate for the queue. Clearly, any scheme for detecting queue length based on round-trip time would have detect changes in network conditions and re-calibrate. This also creates a problem for systems that aim to measure the base RTT, i.e. the round-trip time in the absence of queueing. Because the mean service rate depends on traffic that is not in the queue, a change in other traffic can cause a shift in the base RTT. As queueing time is usually estimated as $RTT - \text{baseRTT}$, this could be an issue for many schemes.

6 Impact on TCP Vegas

We now look at the performance of Linux's TCP Vegas in light of our observations. We consider Vegas because it is one of the simplest delay-based congestion control schemes. We expect other delay based schemes, such as FAST and Compound, to face similar challenges. Linux's Vegas module alters congestion avoidance behaviour but reverts to Reno-like behaviour in other situations. Over each cwnd's worth of packets it collects a minimum RTT observed over that cwnd. It also maintains a base RTT, which is the smallest cwnd observed over the current period of congestion avoidance.

The target cwnd is then estimated as $\text{cwnd} \times \text{baseRTT} / \text{minRTT}$. The difference between this and the current cwnd is compared to the constants $\alpha = 2$ and $\beta = 4$. If the difference is less than α cwnd is increased, if it is greater than β it is decreased. Vegas aims to introduce a few packets more than the bandwidth-delay product into the network resulting in a small standing queue.

We anticipate two possible problems for Vegas. First, because Vegas is using RTT measurements, it is possible that the noise in these measurements will cause Vegas to incorrectly manage the queue, either resulting in an empty queue (reducing utilisation) or overfilling the queue (resulting in drops, which delay-based schemes aim to avoid). Second, after a change in network conditions, Vegas may use an incorrect baseRTT. If this change results in an increased baseRTT then Vegas might continually reduce cwnd in an attempt to reduce the observed minRTT, resulting in poor utilisation.

To investigate these potential issues, we run a TCP flow across our testbed with various round trip times introduced with Dummynet [10]. After 60s we change the network conditions by introducing additional 11 stations, one per second, sending UDP packets at a high rate. First, as a baseline, we run a set of experiments with very large buffers and TCP Reno. Reno keeps these buffers from emptying, and so gives an indication of the best achievable throughput. Results for Reno with a 5ms, 50ms and 200ms RTT are similar to the throughput for Vegas shown in Figure 7.

Figure 7 shows throughput and cwnd histories for Vegas with a 5ms RTT (results for Vegas with a 50ms RTT are broadly similar). We observe that in terms of throughput, it compares well with Reno, both before and after the introduction of additional flows. Can we understand why Vegas does not keep reducing cwnd? If we calculate the min-RTT that is the threshold for increasing cwnd, we get a value of $\text{baseRTT} / (1 - \frac{\alpha}{\text{cwnd}})$. The upper threshold for decreasing cwnd is the same, but with β instead of α . When cwnd is small, the band for maintaining or increasing cwnd becomes larger. Thus, as cwnd becomes smaller Vegas can accommodate increased variability, though it may decrease cwnd below the bandwidth-delay product before this comes into play.

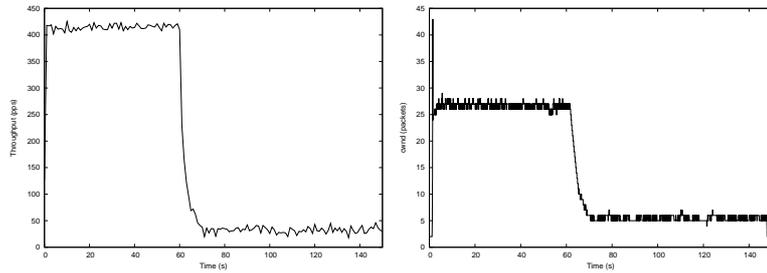


Fig. 7. Vegas with 5ms additional RTT in an initially uncontended WLAN with additional fws introduced around 60s: (a) throughput, (b) cwnd.

Figure 8 shows results for Vegas with a 200ms RTT. Vegas is behaving in a different way: it experiences losses even when not competing with other stations. This may be due to Vegas maintaining a longer queue, and consequently seeing larger fluctuations due to the random service. At 200ms the queue fluctuations are large enough that packets are lost, resulting in Vegas reverting to Reno until it re-enters congestion avoidance. This resets the baseRTT, allowing Vegas to recover when new flows are introduced.

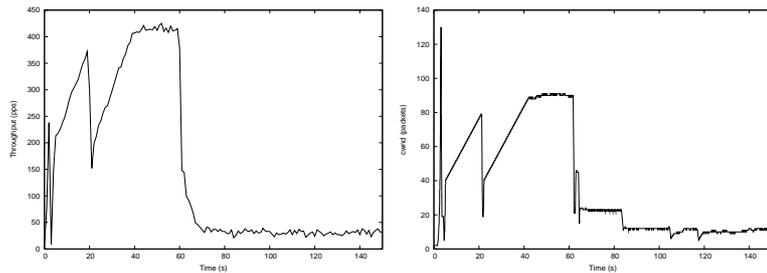


Fig. 8. Vegas with an additional 200ms RTT in an initially uncontended WLAN with additional fws introduced around 60s: (a) Throughput. (b) Cwnd.

7 Conclusion

In this paper we have studied a number of interesting problems faced when inferring buffer occupancy from RTT signals in a WiFi network. We have seen that the raw RTT signal is correlated with buffer occupancy, but there is significant noise that grows as buffer occupancy increases. Standard smoothing filters seem to reduce our prospects of estimating buffer size. We have also seen that traffic that does not share a buffer with our traffic may have a significant impact on the RTT measurements, possibly creating

problems for estimation of queue length under changing network conditions. We have briefly looked at the implications of these observations for Linux's Vegas implementation. While Vegas performs well in our simple tests, possibly due to its use of a Windowed Minimum filter. We believe these observations will prove useful in designing delay-based congestion-control for WiFi.

References

1. Multiband Atheros driver for WiFi (MADWiFi). <http://sourceforge.net/projects/madwifi/>. r1645 version.
2. Soekris engineering. <http://www.soekris.com/>.
3. L. Brakmo and L. Peterson. Tcp vegas: End to end congestion avoidance on a global internet. *IEEE Journal on Selected Areas in Communication*, 13(8):1465–1480, October 1995.
4. I Dangerfeld, D Malone, and DJ Leith. Experimental evaluation of 802.11e EDCA for enhanced voice over WLAN performance. In *Proc. WiNMee*, 2006.
5. M Franceschinis, M Mellia, M Meo, and M Munafo. Measuring TCP over WiFi: A real case. In *WiNMee*, April 2005.
6. IEEE. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE std 802.11-1997 edition, 1997.
7. V. Jacobson. pathchar - a tool to infer characteristics of internet paths. MSRI, April 1997.
8. G. McCullagh. Exploring delay-based tcp congestion control. Masters Thesis, 2008.
9. ACH Ng, D Malone, and DJ Leith. Experimental evaluation of TCP performance and fairness in an 802.11e test-bed. In *ACM SIGCOMM Workshops*, 2005.
10. L. Rizzo. Dummynet: a simple approach to the evaluation of network protocols. *ACM/SIGCOMM Computer Communication Review*, 27(1), 1997.
11. N Sundaram, WS Conner, and A Rangarajan. Estimation of bandwidth in bridged home networks. In *Proc. WiNMee*, 2007.
12. K. Tan, J. Song, Q. Zhang, and M. Sridharan. A compound tcp approach for high-speed and long distance networks. In *INFOCOM*, 2006.
13. D.X. Wei, C. Jin, S.H. Low, and S. Hegde. FAST TCP: motivation, architecture, algorithms, performance. *IEEE/ACM Transactions on Networking*, 14:1246–1259, December 2006.