

Statistical analysis of the IEEE 802.11 MAC service time

Hector Velayos and Gunnar Karlsson

KTH Royal Institute of Technology
Department of Signals, Sensors and Systems
SE-10044 Stockholm, Sweden
{hvelayos,gk}@kth.se

Abstract: We present a statistical analysis of the IEEE 802.11 MAC service time. Our analysis complements the results from mathematical models, which focus on the mean delay. The location analysis shows that the service time's distribution is skewed: the mean is always larger than the median, and the mode is always the smallest. A deeper analysis of the mean illustrates its dependence on the number of stations in the cell, the offered load, the packet size and the bit rate. The analysis of variability includes the coefficient of variation of the service time and its cumulative distribution function. The high variability found indicates that the 95-percentile of the service time may be a more meaningful measure than the mean. Finally, the analysis of correlation demonstrates that the service time of a packet is not predictable from the service time of previous packets.

Keywords: IEEE 802.11 MAC, service time, analysis.

1. INTRODUCTION

Information about the expected delay in communication networks is vital for most applications. Interactive multimedia require a limited end-to-end delay to reach acceptable quality levels. Streaming of multimedia contents uses delay measurements to compute the size of the play-out buffers for jitter compensation. Elastic flows such as web browsing rely on the ability of TCP to predict the end-to-end delay for triggering retransmissions. These examples explain the recent research interest in delay studies for different networking technologies. In particular, the delay in wireless LANs (WLANs) based on the IEEE 802.11 standards has received much attention lately, since it is the most popular option for high-speed, packet-based wireless access to the Internet.

The delay of a packet in a WLAN can be split into three parts: the delay at the medium access control (MAC), the transmission delay and the propagation delay. The propagation delay can be neglected due to the small size of the cells. The transmission delay can be easily calculated from the packet size and bit rate used. The MAC delay is the most difficult to determine because it depends on the traffic in the cell. We divide the MAC delay into two parts: service time, or time to gain access to the shared channel following the rules specified in the IEEE 802.11 standard, and the time spent in the queue waiting for earlier packets to be transmitted. We focus our study on service time, since queuing delay is the consequence of

packet inter-arrival times being shorter than the service times.

The growing interest in quality of service has fueled the publication of models for the delay in the IEEE 802.11 MAC protocol. The analysis of MAC service time is a key part of these models. A common assumption for early models is that stations are always ready to transmit (saturation). There are several examples of such models. Chatzimisios et al. studied the packet delay in presence of transmission errors [1]. Their model is an evolution of Bianchi's model based on Markov chains for throughput analysis in ideal channel conditions [2]. Tay and Chua suggested a different model based on stochastic analysis that provides throughput and packet delay [3]. Carvalho and Garcia presented another model for the MAC delay as a function of the channel state probabilities [4]. Unfortunately, these probabilities can only be calculated under the assumption of saturation. All these models provide mean service time in saturation; they cannot be used for non-saturating loads. Nevertheless, their output values can be used as upper bounds for the service time. Banchs suggested an approximated expression for the distribution of the backoff delay (equivalent to the service time) in saturation [5]. His work permits a better understanding of the service time in saturation compared to previous models that only provided the mean.

The saturation condition was relaxed in two models recently published. Tickoo and Sikdar presented a queuing model for the average service time valid for non-saturating loads and arbitrary arrival patterns [6]. Their model determines the mean service time from average inter-arrival times for traffic sources. Li and Battiti suggested another model for non-saturation in which the mean service time can be derived from the probability that a station's transmission queue is empty after the successful transmission of a packet [7]. None of these models provide information about service time for individual packets.

There are two limitations in existing models. First, they only provide the mean MAC service time; the variability of the packet delays is not modeled. Second, they assume that the number of stations in the cell is large enough so that the probability of packets colliding is constant and independent of the transmission time. However, measurements with real equipment are not a valid alternative to these models to obtain more information about the MAC delay. Packet timestamps in commercial WLAN cards are not accurate enough and include other delays not related to the MAC protocol operation.

In this paper, we present a statistical analysis of the IEEE 802.11 MAC service time based on simulations. Our analysis extends the mathematical models providing packet level information for small number of stations. Output from our analysis includes histograms, mean and variability, cumulative distribution functions, and autocorrelation plots. We show that the service time distribution is skewed and the variability increases with the load. Applications with strict requirements on delay per packet may consider the 95-percentile better than the mean as an indicator of the expected delay. We also show that there is no significant correlation between the service times experienced by the packets of a flow. Hence, the service time of future packets cannot be inferred from previously measured values.

The rest of the paper is organized as follows. Section 2 describes the channel access procedure of the IEEE 802.11 MAC protocol. Section 3 details our methodology for the analysis. Section 4 presents the analysis of location in which the distribution of service time is analyzed; Section 5 contains the analysis of variability, and Section 6 provides the analysis of autocorrelation. Finally, Section 7 closes the paper by summarizing the main findings.

or buffer overflow; the cell operates with an ideal radio channel. Stations are uniformly distributed along a circumference of radius 20 meters with the access point in its center. Physical and MAC layer parameters are set according to the IEEE 802.11b standard.

During the simulations, a warm-up period of one minute precedes the measurement of the service time. The service time is measured for a single station during 10 seconds. Since the seed of the simulator's random generator heavily influences this result, we repeat the measurement 20 times using different seeds. The results shown in the next sections are averages of the 20 runs.

The four parameters affecting the service time were varied during simulations to study their influence on the results. The number of stations was 4, 7 or 10. The packet size was 40 bytes, 500 bytes or 1500 bytes. According to CAIDA² reports, the packet size distribution on the Internet backbones peaks around these values. The bit rate varied between 2 Mbps, 11 Mbps and 54 Mbps, which is the highest bit rate standardized to date for IEEE 802.11 networks. The cell's offered load was normalized to the bit rate, and ranged from 5% up to 100%, increasing in steps of 5% (i.e. an offered load of 50% corresponds to 1 Mbps when the bit rate is 2 Mbps). Each station in the cell hosted one traffic source. The source's rate was 1/n of the cell's offered load, where n was the number of stations in the cell.

In addition to the rate, the behavior of the source affects the MAC service time. We have compared sources generating traffic according to three different patterns: constant bit rate (CBR), exponentially distributed on/off with 20 ms average on time and 35 ms average off time, and Poisson inter-arrivals. Fig. 1 shows the mean MAC service time for the three source types. The service time exhibits a state change for all sources. Numerical values are similar for the different sources for low load (below 50%) and in saturation (above 65% load). A small difference occurs as the load approaches saturation because different sources have a slightly different saturation value. Fig. 2 shows the coefficient of variation of the service time for the three source types. As expected, the variability of the CBR source in non-saturation is the lowest, while the variability of the other two sources is similar. All sources show the same variability in saturation. According to these results, we have selected exponential on/off

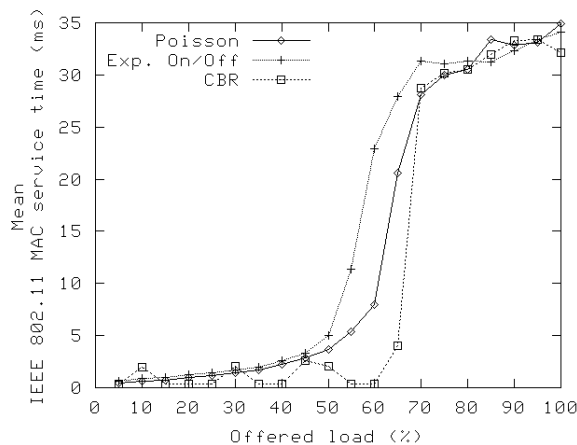


Fig. 1. Comparison of the mean MAC service time for different source behaviors.




Fig. 2. Comparison of the coefficient of variation for different source behaviors.

² <http://www.caida.org>




Fig. 5 Mean service time vs. offered load with number of stations as parameter.

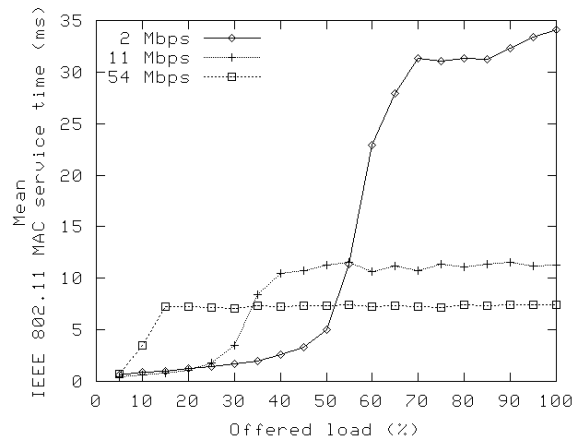


Fig. 6. Mean service time vs. offered load with bit rate as parameter.

analysis of the service time. We now analyze the dependence of the mean service time on its parameters. We first look at how competition for the channel affects the mean. Competition is a function of offered load and number of stations. Fig. 5 plots the mean as a function of the offered load with the number of stations as parameter. The other two factors influencing the mean are fixed to 500 bytes for packet size and 2 Mbps for bit rate.

Two states are visible in the figure: saturation for loads above 70%, and non-saturation for loads below 50%. We denote “near saturation” the loads in the range of the quick transition between these two states. Service times in saturation are much larger than in non-saturation and limited by the buffer size of the stations. The number of stations influences the mean in saturation and near saturation, but it does not affect it in non-saturation. Admission control to maintain low mean service time should control offered load rather than number of stations [8].

We now look at how the time during which the channel is captured for other stations’ transmission affects mean service time. Capture time is controlled by the bit rate and the packet size. Fig. 6 and Fig. 7 show the mean service time versus offered load using as parameter the bit rate and packet size respectively. Both figures show a state transition from non-saturation to saturation as Fig. 5 did.

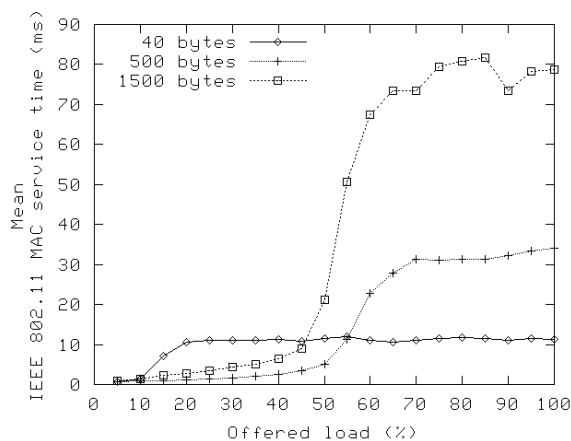




Fig. 7. Mean MAC service time vs. offered load with packet size as parameter.


Increasing the bit rate reduces the offered load to reach saturation. However, service time in saturation is smaller because competing stations capture the channel shorter periods. Changing the packet size affects the mean service time in a different way. Packets of 500 bytes, or larger, only experience significant differences in service time during saturation. The saturation load is similar for all packet sizes from 500 bytes and up. Reducing the packet size has a larger impact on the mean delay. It becomes smaller but the saturation load is severely reduced. Fig. 7 shows that the cell is saturated as




Low load




Low load




Low load



High load



High load



High load

Fig. 10. Service time's CDF for different no. of stations.

Fig. 11. Service time's CDF for different bit rates.

Fig. 12. Services time's CDF for different packet sizes.

precision above. If the bit rate is 11 Mbps, the packet size fix to 500 bytes and the sources behave according to an exponential on-off pattern, our simulations indicate that it would take 31.1 seconds. Probably, the situation will not be stable for such a long period if the measurements would be taken in a real cell, in which stations enter and leave, bit rate changes depending on radio conditions, and packet size varies depending on applications. Protocols, such as admission control or load balancing, relying on measurements should find a way to deal with this high variability, or to work with less accurate mean service times.

The high variability and skewed distribution of the service time makes it interesting to look at upper bounds for the service time of the packets rather than mean values. The cumulative distribution function (CDF) of the service time shows the probability that the service time is




Fig. 13. Comparison of the 95-percentile and mean MAC service time in non-saturation.

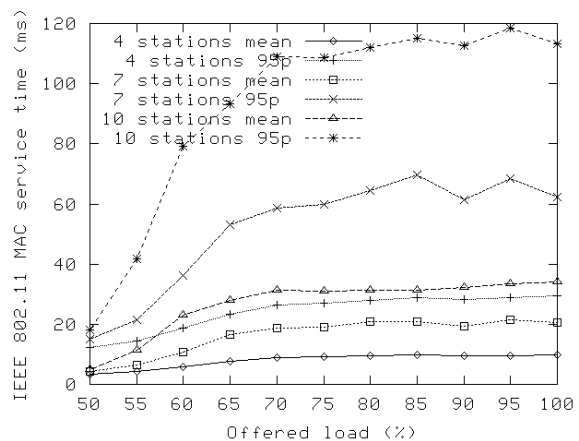


Fig. 14. Comparison of the 95-percentile and mean MAC service time in saturation.

loads, while the number of stations only affects the mean value in saturation. The analysis of variability showed a high spread of the service time around the mean for non-saturation. The variability increased with the load for the three parameters under study. Due to the high variability, we argued that applications with strict requirements on per packet delay would benefit from using the 95-percentile instead of the mean as indicator of the expected delay. We also indicated that measuring an accurate mean value in a real cell might be impossible due to the large number of packets required. Finally, the analysis of correlation showed no significant autocorrelation in the service time of the packets in a flow. Hence, the service time of a packet cannot be predicted from the service time of previous packets.

The starting point for this study was the modeling work that has been presented recently. We conclude that models, which are valid only in saturation and for a high number of stations, provide only worst case measures and give limited insight of the performance of wireless LANs. We also conclude that the arithmetic mean might not be a suitable metric for the performance of a WLAN cell due to the high variability. However, if admission control is provided in the cell to limit the load below saturation, then the performance is good and practically invariable with respect to the system parameters considered herein. The main issue in designing such a control mechanism is to determine the cutoff point for the load since this decision level is determined by the transmission bit rates of the stations as well as the characteristics of the offered load with respect to packet arrival patterns and packet size distributions. A promising technique is to use probing for distributed admission control [8].

REFERENCES

- [1] P. Chatzimisios, A.C. Boucouvalas, V. Vitsas, "Performance analysis of IEEE 802.11 DCF in presence of transmission errors", in *Proc. of IEEE ICC 2004*, pp. 3854 – 3858, June 2004.
- [2] Giuseppe Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function", *IEEE Journal on Selected Areas in Communication*, March 2000.
- [3] Y.C. Tay and K.C. Chua, "A capacity analysis for the IEEE 802.11 MAC protocol", *Wireless Networks*, vol 7, pp 159-171, Kluwer Academic Publisher, 2001.
- [4] Marcelo Carvalho and J.J. Garcia-Luna-Aceves, "Delay analysis of IEEE 802.11 in single-hop networks", in *Proc. of the 11th IEEE ICNP'03*, 2003.
- [5] A. Banchs, "Analysis of the distribution of the backoff delay in 802.11 DCF: a step towards end-to-end delay guarantees in WLANs", in *Proc of QoFIS 2004*, Barcelona, Spain, 2004.
- [6] Omesh Tickoo and Biplab Sikdar, "A queuing model for finite load IEEE 802.11 random access MAC", in *Proc. of IEEE ICC'04*, Paris, France, June 2004.
- [7] Bo Li and Roberto Battiti, "Analysis of the IEEE 802.11 DCF with service differentiation support in non-saturation conditions", in *Proc. of QoFIS 2004*, Barcelona, Spain, 2004.
- [8] Hector Velayos, Ignacio Mas and Gunnar Karlsson, "Distributed admission control for WLANs", KTH Technical report TRITA-S3-LCN-0502, Stockholm, Sweden, May 2005.
- [9] R. Jain, *The Art of Computer Systems Performance Analysis. Techniques for experimental design, measurement, simulation, and modeling*, John Wiley & Son Inc., 1991.
- [10] NIST/SEMATECH *e-Handbook of Statistical Methods*, Dec. 2004, <http://www.itl.nist.gov/div898/handbook/>.
- [11] Matthew Roughan and Oliver Spatscheck, "What Does the Mean Mean?", in *Proc. of International Teletraffic Congress (ITC) 18*, 2003.