Empirical Analysis of Local Round Trip Time for Wireless Traffic Differentiation

Guangzhi Qu and Michael M. Nefcy

Computer Science and Engineering Department, Oakland University,
Rochester, MI, 48309, USA
{gqu,mmnefcy}@oakland.edu

Abstract. This paper focuses on how to differentiate wireless traffic from wired peer by using the temporal TCP characteristics of SYN, FIN, and ACK local round trip times (LRTT) found in all TCP sessions. With these session-based temporal characteristics, traffic from wireless and wired nodes can be differentiated by exploiting the fundamental differences between Ethernet and 802.11b/g/n. The effort of this paper is then on analyzing the resulting empirical LRTT data extensively and designing several algorithms for effective wireless host discovery. Most algorithms are light-weight, with little memory overhead, and can be easily implemented on commodity hardware. Ultimately, SYN, FIN, and ACK LRTTs can be compared against each other to discover wireless hosts regardless of network speeds.

Keywords: Wireless Network, Temporal Analysis, Local Round Trip Time.

1 Introduction

Network security and resource management are vital components to the productivity of any modern day business network. Whether the network is a small home office or a large university, significant time and effort must be devoted towards protecting computers and services against threats. This all but demands that network administrators keep a close watch on their networks by planning hardware and software resources to carefully keep unauthorized users out. Such a picture becomes further complicated as more and more innovative technologies become available for business adoption, ultimately creating further network vulnerabilities. How successfully network administrators maintain security in the face of old and emerging threats depends largely on the deployment of firewalls, intrusion detection systems, and other network defense tools.

One particular piece of technology that network administrators know they need to keep a close eye on is 802.11 wireless networking. Wireless networking has seen tremendous growth and widespread business adoption in the past decade. In fact, commodity 802.11 devices have become nearly ubiquitous today, emerging as an easy, convenient solution that can already exceed megabit Ethernet speeds [1]. Challenge in differentiating the wireless network traffic from the traditional Ethernet traffic draws huge interests and attention to the system administration.

For example, take the simplified network illustrated in Fig. 1. The local area network has both wired (LAN) hosts and wireless (WLAN) hosts present. These hosts

X. Zhang and D. Qiao (Eds.): QShine 2010, LNICST 74, pp. 222–237, 2011.

are potentially very diverse, with some being servers and workstations, and still others as PDAs and phones. The hosts have different operating systems and hardware processing speeds, as well as different networking technologies. For instance, an internal host could use wired networking based on either 10/100/1000 BaseT gigabit IEEE 802.3ab Ethernet or 10/100 BaseT IEEE 802.3u megabit Ethernet (referred to as either gigabit or megabit, respectively). Alternatively, a host could use wireless networking based on IEEE 802.11b, 802.11g, or 802.11n wireless LAN (referred to as 802.11b, 802.11g, or 802.11n respectively). Further, the network infrastructure can have similar diversity as well. Typically, such infrastructures are wired 10/100/1000 BaseT IEEE 802.3ab gigabit Ethernet infrastructures or 10/100 BaseT IEEE 802.3u megabit Ethernet infrastructures (referred to as either gig-backbones or meg-backbones for simplicity's sake). Lastly, multiple wireless access points are also present, with similar diversities as the rest of the wireless hosts.

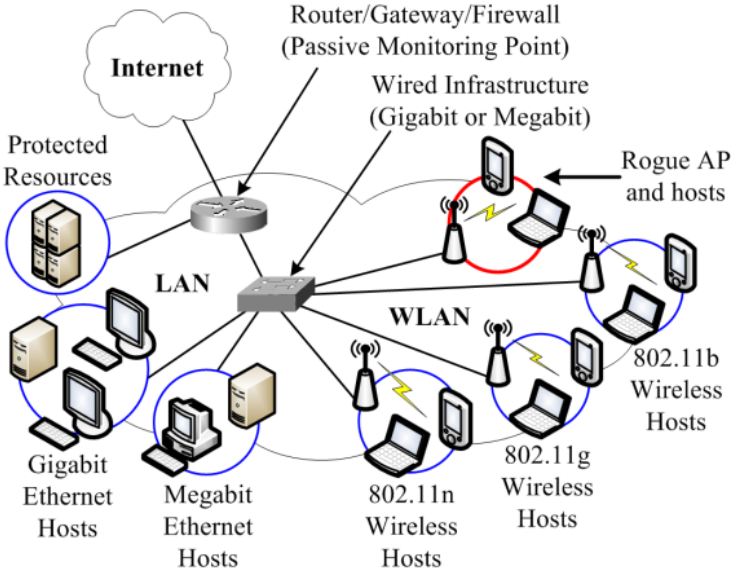

Fig. 1. A Typical Network Topology

In order to manage this heterogeneous network environment, the system administrators have to understand the network traffic. The question is how we can classify the network traffics in real-time? More specifically, how we can discover wireless hosts? What measurable metrics differentiate wireless from wired host traffic properties? And what algorithms can decide on such metrics in a scalable, real-time fashion with low overhead requirements?

To answer these questions robustly, this work takes the approach of using local round trip time (LRTT) metrics and some simple algorithms to discover wireless hosts in a heterogeneous, real-world, small office networking environment. Previous research has demonstrated that LRTT is an effective metric for wireless host differentiation from the rest of network traffic [3, 4]. The specific definition for LRTT

used here is the time any TCP/IP session packet pair takes to traverses the local side of a LAN between the gateway/router and the host itself. This study's novelty is similar to Watkins et al. [4]: it uses only packet pairs found within TCP/IP sessions.

Several simple, stateless algorithms are used to decide on the developed LRTT metrics. These just use empirically-derived thresholds to classify wireless traffic in a straight-forward fashion. More complicated learning algorithms will be evaluated in later research, but are presently excluded from analysis.

2 Analysis of LRTT Metrics

An in depth analysis of local round trip time metrics is provided in this section, demonstrating how they reveal wireless hosts on different host and network infrastructure technologies. This research also differs from previous work done by Watkins et al [4] and others in that a much more robust analysis of LRTTs collected within TCP/IP sessions is performed. Analysis of these separate metrics illustrates how Ethernet hosts (gigabit and 100-megabit) are separated from 802.11 wireless hosts (802.11b, 802.11g, and 802.11n) on different network infrastructure hardware generations (10/100/1000 BaseT gigabit and 10/100 BaseT megabit). The following analysis starts with the settings and assumptions for LRTTs. From them, different results for Ethernet and 802.11 wireless networking are derived using a high-level approach. Further, how these results are affected by network infrastructure changes is presented as well.

2.1 Settings for LRTT Analysis

Consider two hosts communicating with each other, as depicted in Fig. 2. An external host communicating with an internal network host first sends an incoming packet to the local gateway. The incoming packet is then propagated on the network infrastructure, represented by link L_1. If the internal host uses wired networking, it receives the incoming packet directly from L_1. Otherwise, if the host uses wireless networking, it must receive the incoming packet from a radio-frequency link between a wireless access point and the host, represented by link L_1. Once the packet is received, the internal host sends a responding outgoing packet back to the gateway via L_1 (for the wired host) or L_2+ L_1 (for the wireless host). Finally, the gateway relays the outgoing packet back out to the external host.

Due to the diversity of networking technologies, the two communicating hosts may have slightly different settings. The result is the three primary variants shown in Fig. 2 (based on the host's networking technology). In the top scenario, the internal host uses gigabit or megabit. In the middle scenario, the internal host uses 802.11b or 802.11g wireless LAN. And finally, in the bottom scenario, the internal host uses IEEE 802.11n wireless LAN. These different network technologies each translate to different properties. For instance, the link speeds can vary, allowing the internal host to communicate at maximum rates of 1000mbps, 100mbps, 11mbps, 54mbps, or 600mbps for gigabit, megabit, 802.11b, 802.11g, or 802.11n respectively. Further, Ethernet hosts can transmit and receive packets simultaneously (full-duplex), whereas 802.11b and 802.11g hosts can either only transmit or receive at once (half-duplex). Unlike 802.11b or 802.11g, the newer 802.11n also allows full-duplex just like Ethernet (although not as guaranteed).

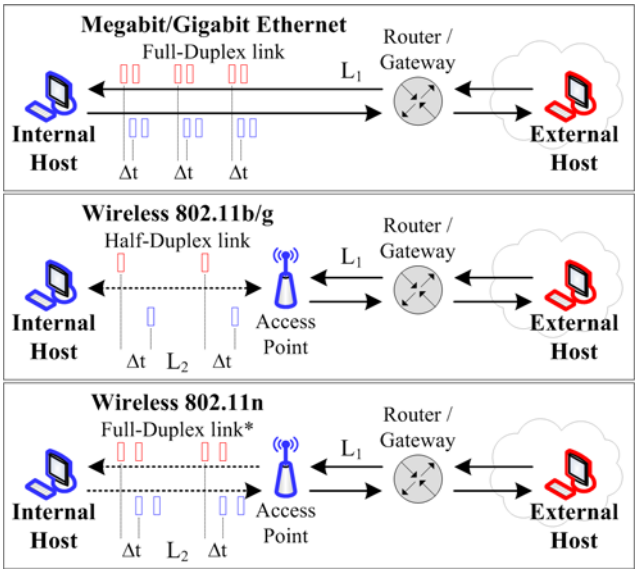

Fig. 2. The settings for analysis of LRTT metrics

Both the incoming and outgoing packets must pass over the network infrastructure, or backbone, represented by link L_1. This infrastructure can have similar diversity as well, and is typically either a gigabit-backbone or megabit-backbone. The network infrastructure limits the overall speed of the packets, as well as the ultimate link speed of the internal hosts. Additionally, the path of L_1 may have multiple chained links, where the packets must be propagated through a series of backbone switches.

As packets traverse the infrastructure, they pass through a central router or gateway. This gateway is the monitoring point, where all TCP/IP packet headers of network traffic are continuously, passively captured. From this traffic, unique local hosts are identified and their TCP/IP connection streams are analyzed. If a packet is incoming, its header is stored in a single, large, temporary queue for finding LRTTs (an M/D/1 queuing model). Conversely, if a packet is an outgoing response, its header is matched against its corresponding incoming packet from the queue, creating an incoming-outgoing-packet-pair. The arrival times of the packets in the packet-pair form a time difference, which is called the LRTT.

2.2 Assumptions of LRTT Analysis

There are two main types of assumptions made, the first involve time and the second involves TCP/IP itself. Though covered in more detail below, essentially this approach assumes LRTTs will only be affected by network factors and can readily be derived from TCP/IP session streams.

Time Assumptions. Pairs of corresponding incoming and outgoing packets are used to take time measurements for each unique internal network host. A core assumption of

these metrics is that packet-pair LRTTs will only vary with a significant difference based on the host's network technology and packet size.

Consider a packet-pair travelling on the network setting shown in Fig. 2. The LRTT of this packet-pair is influenced by various factors of the path the packets travel on. From the incoming packet (P_{in}) arriving to the outgoing packet (P_{out}) leaving, these factors are: (1) The processing (proc), queuing (queue), and transmission (trans) of P_{in} at the gateway; (2) The processing , queuing, transmission and propagation (prop) of P_{in} along the network backbone (L_1) to the internal host; (3) The processing, queuing, and transmission of P_{out} by the internal host; (4) Finally, the processing queuing, transmission, and propagation of P_{out} back along the infrastructure (L_1) to the gateway. Note that if the internal host is wireless, both P_{in} and P_{out} must also pass through the wireless access point and its time delays as well (L_2). Each factor adds a specific Δt_i to the total LRTT for P_{in} and P_{out}. Further, each Δt_i is specific for a given packet size. If P_{in} and P_{out} arrive at times t_{in} and t_{out}, then all Δt_i factors for the packet-pair can be stated as ΔT.

Ultimately, the dependence of each time delay factor on the size of each packet and each device is not as confounding as it seems. For a fixed packet size, some Δt_i factors will remain mostly constant between network technologies. Further, any Δt_i that remains constant is not scientifically significant and can thus be ignored (it may event drop out of the equation if both P_{in} and P_{out} are the same size). For instance, most hosts have CPU speeds fast enough to render host processing and queuing times insignificant. The same goes for the gateway itself, especially since all incoming packets experience the same gateway time delays for a given packet size. Therefore, the only factors that remain are infrastructure related delays and host transmission delays, as seen in the final equation for ΔT. And thus, for a given P_{in} and P_{out}, only the variability that remains significant are changes in the host's network technology and changes in the infrastructure network technology.

Packet size cannot be ignored, nor can it be dismissed as an easily calculation based on transmission speed alone. True, a faster network can transmit a same-size packet in less time than a slower one, but this is not always the case. In wireless networks, larger packet time delays are dominated by transmission rates, whereas smaller packets are dominated by processing, queuing, and propagation rates. The effects of large vs. small packets on time delay have been studied on wireless LANs [2, 3, 5], but also appear in wired networking too (similarly due to the timing of MAC protocols, but with opposite results; see Section 6). In effect, the overall rate of transmission and propagation depends on the per packet time spent in the link protocol vs. spent in the data itself. Thus, if the link-time to data-time ratio is low, then smaller packets transmit faster due to the high overall time efficiency (low time overhead per frame). But if the ratio of link-time to data-time is high, then larger packets transmit faster due to the low overall time efficiency (high time overhead per frame). In addition, larger packets are take longer to transmit, longer to propagate, and are more affected by congestion and interference. Also, larger packets need more network process and queuing time for a given host or infrastructure as well. In summation, a nonlinear time dependency exists between packet size and total time delay for each network technology variant which can't be explained by link speed alone. Thus packet-pairs of different sizes need to be kept separate in order to achieve the most accurate measuring of ΔT.

Stated another way, within each host network technology type, within each network infrastructure technology type, the total time delay will remain constant for a fixed packet-pair size. Any other variables simply don't have the right variance to be significant. Again, the ultimate goal is to show the host's network medium type. So if the host processing takes too long, the resulting ¢T won't reveal the underlying technology and would be discarded anyway. Furthermore, with larger packets showing transmission effects the most, and smaller packets showing the medium effects the most, this approach of taking both into account will scale with faster networking speeds.

TCP/IP Assumptions. Producing pairs of sequential, corresponding packets relies on several TCP/IP assumptions. First, both the external and internal hosts use TCP/IP for most, if not all, of their communications. This guarantees that packets from these hosts will appear in the monitoring queue of the central gateway. Second, the most basic TCP/IP standards for managing a connection need to be followed. This means that hosts need to follow the 3-way establishing handshake, the 2-way continuing transfer, and the 4-way termination handshake, as depicted in Fig. 3. Either host can initiate each of those three general pieces of the TCP / IP session, leading to the six possible permutations shown. Each permutation has specific packet sequences, which must follow standards. Each packet, besides having IP addresses and port numbers, also has identifiable TCP/IP flags, acknowledgement numbers, and sequence numbers. The packet is matched to which sequence it belongs to based on such information. Note that not all packets in a sequence are usable. In Fig. 3, the unusable packets are shown in black, whereas the usable packets are shown in color (with red denoting incoming packets, and blue denoting outgoing packets). Only the packet-pairs which form incoming and then outgoing communications can be used, and they are given specific packet case numbers in Fig. 3. Packet-pairs form case-pairs of either 0-1, 2-3, 4-5, 6-7, and 8-9. Further, these case-pairs fall into three different categories based on their TCP/IP roles: 3-way establishing SYN-pairs, 2-way continuing ACK-pairs, and 4-way terminating FIN-pairs. Each set of SYN-, ACK-, or FIN-packet pairs form their own local round trip time categories, which are referred to as $LRTT_{FIN}$, $LRTT_{SYN}$, and $LRTT_{ACK}$, respectively. Each LRTT category is assumed to have packet-pairs of distinct, relatively constant sizes (based on the total size of both packets in a packet-pair). Further, each packet-pair is assumed to be directly sequential in the overall packet stream – i.e. no other packets are transmitted between the incoming packet "stimulus" and the outgoing packet "response." In totality, these assumptions guarantee that matching packet pairs form effective, fully passive "pings" of a local host, estimating the local host's LRTT latencies for three different packet-pair class sizes.

All assumptions made here hold well by default, or can be enforced through filtering rules. In order to communicate, hosts need to follow basic TCP/IP standards, so these assumptions are well met most of the time. Any packets not meeting the packet cases expected are simply discarded. Any packet-pairs that are not directly sequential (with no other packets being transmitted from the host in between) are also discarded. Such rules further ensure that the assumptions hold. As discussed in Section 6, the three classes of LRTTs to, in fact, fall into mostly constant sizes. The median pair-pair sizes for FIN-, SYN-, and ACK-pairs are 114, 122, and 1574 bytes

respectively (see Table 1). Of final note is that the assumption on packet-pair sizes was not enforced in this specific approach, as only ACK-pairs had noteworthy size variation.

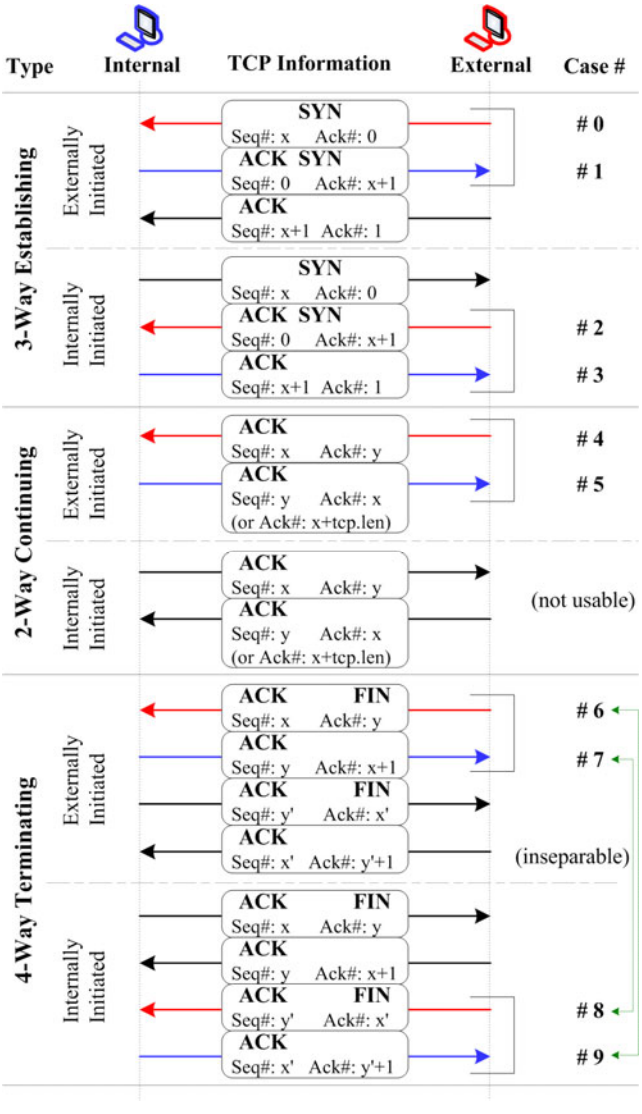

Fig. 3. Assumptions for LRTT metrics

Overall, this approach utilizes small amounts of overall network traffic. Less than 10% of the network's data stream is used. Further, of that utilized traffic, less than 10% is made up of FIN- and SYN-pairs. This is significant, since this approach places more importance is such the smaller packet pairs. Thus a large weight is placed on by far the smallest overall percentages of packets, which delays detection times to an

average of one second. Aside from this disadvantage, these packet-pair rules nevertheless create three powerful indices of LRTTs, which can be used to effectively differentiate wired from wireless host activity.

2.3 Ethernet Analysis

A closer inspection of local round trip time in the Ethernet host setting is provided in this subsection. As illustrated in the top scenario of Fig. 3, an Ethernet host typically uses full-duplex communications with link rates of 1000mbps or 100mbps for either gigabit or megabit network interfaces, respectively. Full-duplex communications allow for efficient, long distance, simultaneous transmitting and receiving of packets. Today, practically all Ethernet hosts use full-duplex networking on switch-based infrastructures, making the use of Ethernet's half-duplex CSMA/CD MAC protocol a rare enough event to ignore. Lastly, the noteworthy differences between megabit and gigabit hosts are that the later has a faster link rate, faster processing, and lower overall latency.

Network infrastructure complications aside, this all has several implications for Ethernet host LRTTs. First, full-duplex communications allow for very low noise and congestion. Second, the efficiency of the Ethernet MAC layer makes inter-packet spaces and time-overhead per packet very small (low time overhead per frame). Third, since hosts can send whenever they want, packets have very small queue waiting times before they're transmitted. All these factors mean Ethernet time delays are very short to begin with, even without considering the very fast link rates Ethernet provides. Thus Ethernet LRTTs will be the smallest in general, with gigabit hosts showing smaller LRTTs than megabit hosts (Equation (1)). Further, since the medium has high overall time efficiency, smaller packets will transmit faster within a given Ethernet type (Equation (2)).

Short packet queues on network interfaces and high time efficiency mean that the overall variance in packet time delays remain low. Stated another way, the faster an interface can process its packet queues, and the less link-layer negotiating that takes place, the more consistently a given packet size is transmitted. Thus gigabit hosts are expected to have smaller LRTT variability than megabit ones due to their faster processing and transmission rates (Equation (1)). Similarly, since smaller packets can be processed faster and spend less time in transmission per frame, their variability should be less than larger Ethernet packet transmissions (Equation (2)). Let $f^1(.)$ be a function that calculates the mean value of the input, $f^2(.)$ for median and $f^3(.)$ for variance of LRTT categories (FIN, SYN, or ACK). In summation, these LRTT derivations in Ethernet are as follows:

$$f^{1,2,3}(\mathrm{LRTT}_x(\mathrm{Gigabit})) < f^{1,2,3}(\mathrm{LRTT}_x(\mathrm{Megabit})) \,. \tag{1}$$

$$f^{1,2,3}(\mathrm{LRTT}_{FIN}(\mathrm{Eth})) < f^{1,2,3}(\mathrm{LRTT}_{SYN}(\mathrm{Eth})) < f^{1,2,3}(\mathrm{LRTT}_{ACK}(\mathrm{Eth})) \,. \tag{2}$$

2.4 Wireless Analysis

For Now a closer analysis of local round trip times associated with wireless hosts is given. Referencing the middle and bottom views in Fig. 2, wireless networking

technology encompasses more diversity than Ethernet does. Despite this, most wireless technology shares many of the same traits. For instance, 802.11b, 802.11g, and 802.11n all use the CSMA/CA algorithm in their MAC protocols, as well as random back-off and wait DCFs when collisions occur. Further, each technology often shares the same channels, causing interference and contention due to their inherent shared medium nature.

There are also many notable differences between wireless network technologies too. The link-rate speeds of 802.11b, 802.11g, and 802.11n are typically 11mbps, 54 mbps, and 300+ mbps. These correspond to actual data-rate speeds of approximately 5 mbps, 22 mbps, and 94 mbps respectively (see Table XX). The previous generations of wireless technology, 802.11b & 802.11g, use half-duplex communications, which often leads to significant network throughput reductions with only a handful of users. The latest generation of wireless technology, 802.11n, uses almost full-duplex with the help of MIMO and SDM. Other differences are that 802.11b/g uses the 2.4GHz spectrum, with 20Mhz channel bands, where as 802.11n has the option of using the 5GHz spectrum, with 40Mhz channel bands. By using an effective Greenfield spectrum and larger channels, 802.11n achieves significantly less interference and faster throughput. Other advances of 802.11n include MAC-layer packet aggregation with block level acknowledgments, and up to 4x4 spatial antenna streams. All of this allows 802.11n have link rates of over 600mbps, which may achieve effective gigabit Ethernet performance for the average user.

Again ignoring network infrastructure complications for now, all these factors elicit differences in wireless LRTT behavior. Half-duplex communications, random back-off and wait DCFs, 3-way handshakes, and other MAC-layer properties make wireless networking have high time-overhead per frame. This is even more pronounced in 802.11n, which, even though it uses MIMO/full-duplex, still uses the DIFS and SIFS spaces of wireless MAC frames. With high inefficiency per frame, this situation becomes the opposite of Ethernet. Here, the larger packet-pairs will be dominated by transmission delays. Thus the faster wireless medium is, the faster it will transmit larger ACK-pairs (Equation (3)). Furthermore, the greater time inefficiencies will manifest the most in the smallest packets, giving FIN- and SYN-pairs the largest delays (Equation (4)).

Similarly, wireless hosts will have to wait and synchronize carefully before sending, leading to larger queue time delays on top of larger processing and propagation delays. The faster the wireless link, the faster queues can be serviced, leading to smaller variability for the larger packets (which are dominated by transmission rate) (Equation (3)). This is due to larger frames needing to be carefully controlled to ensure fair throughput, which thus reduces $LRTT_{ACK}$ variability. At the same time, the smallest packets (already experiencing large medium delays) will show the largest variation in total time delay since they will often have to wait the longest relative to their size (Equation (4)).

Lastly, due to the high time efficiency of Ethernet and the low time efficiency of 802.11b/g/n, the smallest packets experience opposing effects on LRTTs. Wireless hosts see the most delays with the smallest packets, while Ethernet hosts see the least (Equation (5)). Moreover as wireless hosts see the most variance with the smallest packets, Ethernet hosts see the least (Equation) as well. These derivations are summarized below:

$$f^{1,2,3}(LRTT_{ACK}(w_n))<f^{1,2,3}(LRTT_{ACK}(w_g))<f^{1,2,3}(LRTT_{ACK}(w_b)) . \quad (3)$$

$$f^{1,2,3}(LRTT_{ACK}(w_x))<f^{1,2,3}(LRTT_{SYN}(w_x))<f^{1,2,3}(LRTT_{FIN}(w_x)) . \quad (4)$$

$$f^{1,2,3}(LRTT_{FIN,SYN}(Eth))<f^{1,2,3}(LRTT_{FIN,\,SYN}(w_x)) . \quad (5)$$

The differences between these equations and those for Ethernet are due to the ambiguity surrounding smaller packet sizes, since smaller packets are dominated by the network medium in wireless networking. However, using the link rates as approximations for the data rates of each host network technology type, equations (1) and (3) can still be combined to a general result. Since the largest packets will be dominated by transmission delays in both Ethernet and wireless, and since gigabit, megabit, 802.11n, 802.11g, and 802.11b are decreasing in link rates, then:

$$f^{1,2}(LRTT_{ACK}(w_{b,11m}))>f^{1,2}(LRTT_{ACK}(w_{g,54m}))>$$
$$f^{1,2}(LRTT_{ACK}(w_{n,300m}, Megabit))> f^{1,2}(LRTT_{ACK}(Gigabit)) . \quad (6)$$

$$f^{1,2}(LRTT_{ACK}(w_{b,11m}))>f^{1,2}(LRTT_{ACK}(w_{g,54m}))>$$
$$f^{1,2}(LRTT_{ACK}(Megabit))> f^{1,2}(LRTT_{ACK}(W_{n,600m}, Gigabit)) . \quad (7)$$

Note that megabit is comparable to 300mbps 802.11n due to the high wireless link overhead yielding nearly equivalent actual data rates. Thus assuming an 802.11n link rate of 300mbps, megabit and 802.11n cannot be intuitively ordered (Equation (6)). Similarly, assuming an 802.11n link rate of 600mbps, data rates for gigabit and 802.11n are ambiguous, so they cannot be intuitively sorted either (Equation (7)). As 802.11n becomes faster, its data rates may become indistinguishable from those of effective gigabit data rates. However, since no 600mbps 802.11n technology was available for evaluation, equation (7) will have to be verified at a later date.

2.5 Network Infrastructure Analysis

This subsection analyzes how network infrastructure influences the network time delays of internal hosts. Principally, this research only considers gigabit and megabit twisted pair Ethernet networking backbones. The differences between these two are an important consideration for maintaining RWAP detection accuracy on both current and legacy networks. Also, 802.11n wireless hosts currently achieve 300mbps link rates, making them effectively indistinguishable from megabit and gigabit Ethernet on megabit-backbones. Finally, since most previous studies were done using only megabit infrastructure, providing analysis of both is critical for comparing studies.

Gigabit and megabit network backbones alter LRTTs due to various factors. The link rate (i.e. transmission) differences have the largest effects, but disparities in processing and queuing times also exist. Since gigabit and megabit networking have link rates of 1000mbps and 100mbps respectively, packets on gigabit-backbones experience smaller transmission delays. In addition, gigabit-backbones often have faster, more efficient processing and queuing rates than megabit-backbones as well.

Thus on gigabit-backbones, hosts will experience smaller, less distorted LRTTs, making wireless host differentiation easier. Conversely, on megabit-backbones, hosts experience more latency and distortion of their medium's time characteristics, leading to larger LRTTs.

Moreover, gigabit Ethernet and 802.11n wireless hosts must accept lower data rates while on megabit-backbones. This means gigabit and 802.11n LRTTs will be no smaller than similar megabit hosts. In addition, since larger packets are dominated by transmission rates, the large ACK-pair LRTTs of gigabit, megabit, and 802.11n will be very similar on megabit-backbones. This may invalidate equations (6) and (7), with respect to such hosts. Finally, megabit-backbones are optimized for megabit hosts by default, so LRTTs may be smaller specifically for megabit hosts. In contrast, gigabit hosts will forced into a legacy mode, possibly leading to larger LRTTs than megabit hosts.

Overall, the ultimate goal is to create LRTT metrics and rules of host behavior that are agnostic with respect to network infrastructure speeds. Since only the faster gigabit, megabit, and 802.11n hosts are affected by infrastructure changes, a megabit-backbone essentially squeezes LRTT measurements together, distorting results, and making wireless host discovery less accurate. However, this loss of accuracy will be shown to be minor (but measurable).

So consider the worst case scenario: very fast, contention-free 802.11n hosts share a megabit-backbone with highly congested gigabit and megabit hosts. The high Ethernet congestion increases dropped and retransmitted packets, while contention free wireless hosts conversely see fewer delays. In this situation, the differentiation of wireless hosts will be more problematic, since both absolute and relative LRTT trends are affected. However, such extreme conditions would show only a slight decrease in the high MAC time efficiency of Ethernet. Further, retransmitted packets would experience delays too large to be useful and thus would be ignored anyway. And the already low MAC time efficiency of wireless would stay the same regardless. Therefore the overall increased delays in Ethernet LRTTs would be negligible. Lastly, the more normal scenario sees Ethernet as free from contention while wireless quickly degrades with just small numbers of hosts. So under normal conditions, wireless host behaviors increase their separation, aiding in wireless detection. Thus, in either extreme or normal conditions, wireless host are still differentiated.

3 Evaluation

This section presents the evaluation environment, the evaluation procedure, and the collected network traffic data sets.

To achieve real world results, a small, extremely diverse test environment was created following the setting shown in Fig. 1. The evaluation procedure is specified as follows. First, the network infrastructure was constructed as either gigabit or megabit Ethernet. Second, network traffic was captured on the router's LAN interface. Third, traffic captures were fed into a packet analysis program for convenient packet summary generation. Fourth, packet summary outputs were sent through a tracking program that continuously classified hosts using wireless host detection algorithms.

Due to the page limitation, we show only the cumulative distribution functions of how the behavior of 3 different packet-pair sizes from 5 different host types change over 2 different network backbone technologies. Figures 4-6 show the cumulative distribution functions (CDFs) of FIN-, SYN-, and ACK-LRTTs for Gig-Backbone data. Similarly, Figures 7-8 show the corresponding CDFs for Meg-Backbone data.

As an example, consider Fig. 4. This figure shows CDFs of FIN-pair local round trip times. Solid lines represent Gig-Backbone data, and each color represents a specific host network technology type. So the solid red line is the cumulative distribution of gigabit host LRTTs on the gigabit network infrastructure, etc. Since the graph is logarithmic on the horizontal axis, roughly 95% of all LRTTs are $1*10^{-3}$ seconds or less. In fact, most Ethernet traffic in general falls before $1*10^{-3}$s and is clearly distinct from wireless traffic. More importantly, despite the different properties of 802.11n, 802.11g, and 802.11b wireless types, all wireless $LRTT_{FIN}$ CDFs are virtually identical. The separate grouping of Ethernet and wireless technologies confirms results from other research showing that the smallest packets show affects from the transmission medium the best [34].

Compare the Gig-Backbone trends from Fig. 4 to the Meg-Backbone ones of Fig. 7. The Meg-Backbone data shows that $LRTT_{FIN}$ CDFs for wireless host types are similar to those from the gigabit infrastructure. However, gigabit host LRTTs increase (shift to the right) and become more variable (smaller slope). This is most likely due to gigabit hosts operating in legacy modes when on megabit infrastructure. Thus LRTTs from gigabit hosts become slightly more like those from 802.11n hosts when on megabit infrastructure. Interestingly enough, megabit hosts show the opposite: their LRTTs are slightly smaller and less variable in the Meg-Backbone dataset, increasing separation from wireless hosts. This behavior is seen repeatedly in later figures as well.

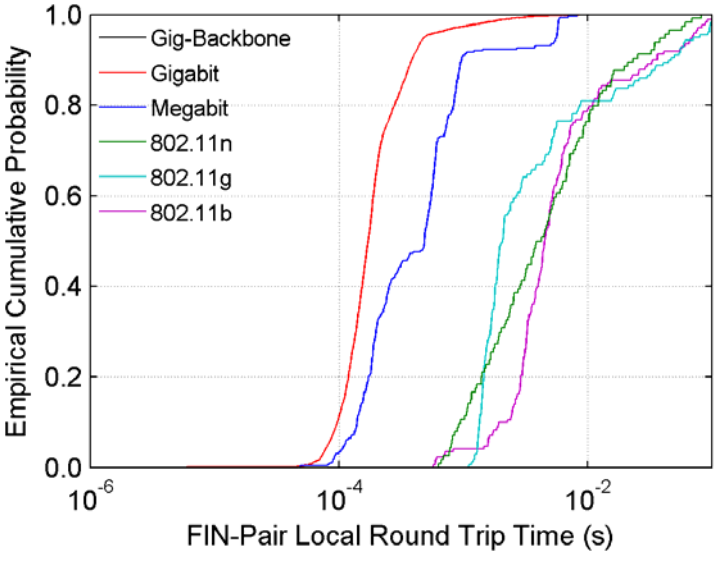

Fig. 4. CDF of Gigabit-Backbone FIN-LRTTs per host network technology type

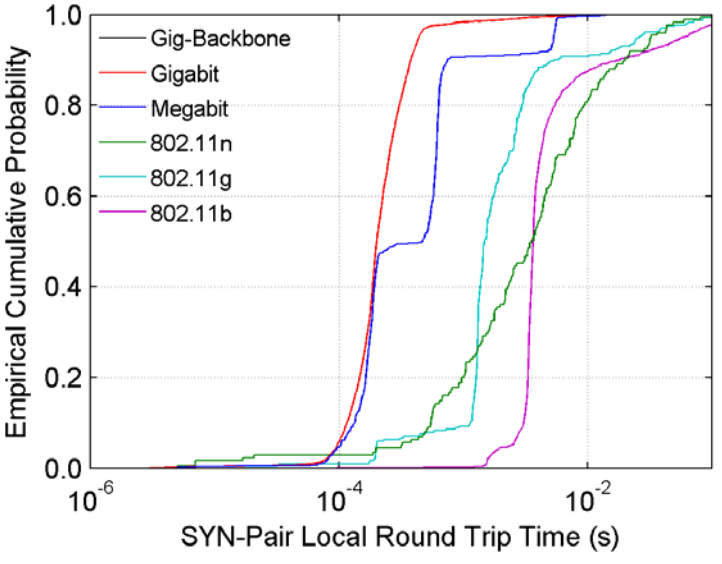

Fig. 5. CDF of Gigabit-Backbone SYN-LRTTs per host network technology type

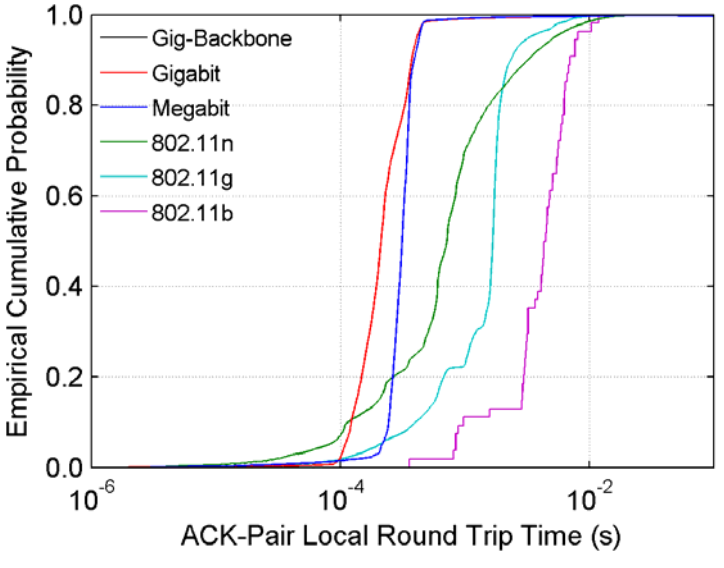

Fig. 6. CDF graphs of Gigabit-Backbone ACK-LRTTs per host network technology type

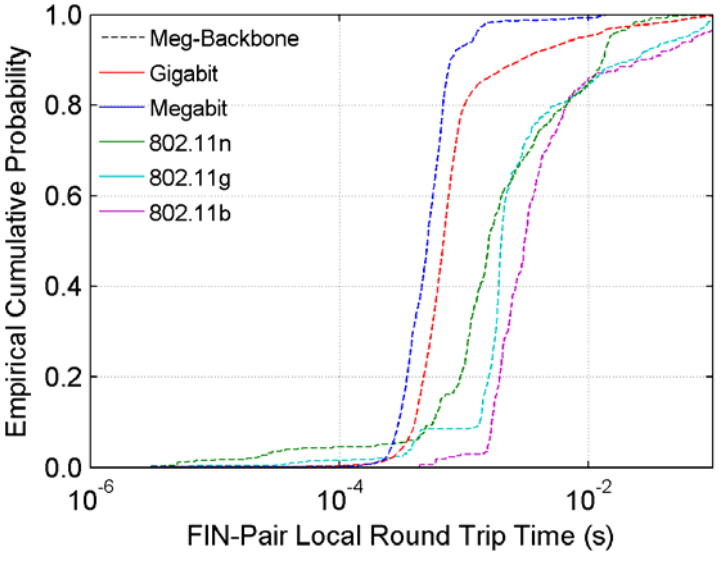

Fig. 7. CDF of Megabit-Backbone FIN-LRTTs per host network technology type

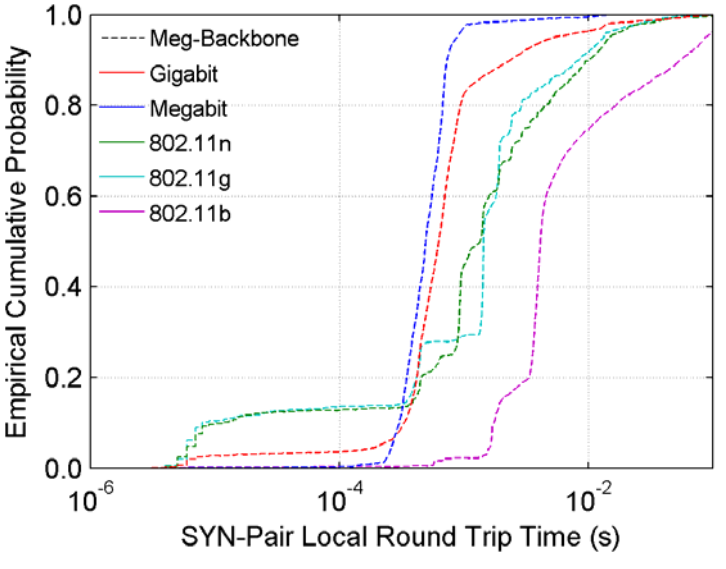

Fig. 8. CDF of Megabit-Backbone SYN-LRTTs per host network technology type

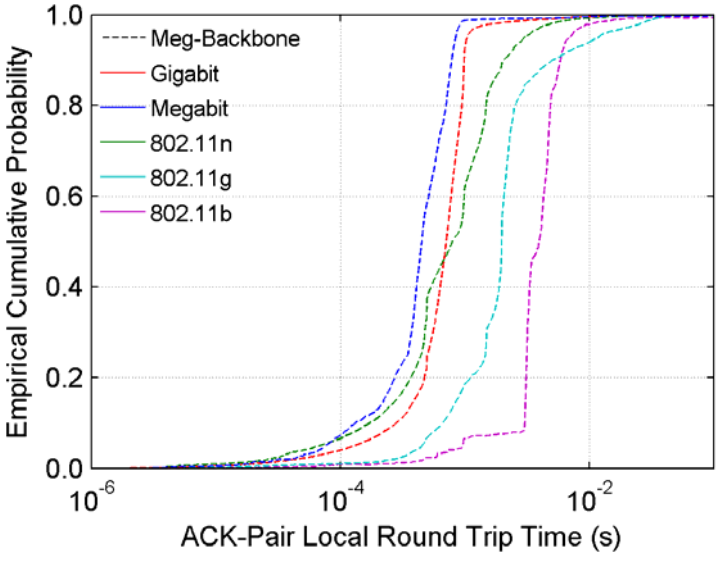

Fig. 9. CDF of Megabit-Backbone ACK-LRTTs per host network technology type

4 Conclusions

The analysis covered in this work reveals that TCP session based LRTTs can discover wireless hosts regardless of host network technology or network infrastructure type. Individual trends illustrate wireless $LRTT_{ACK}$ delays being higher in 802.11b/g wireless hosts than in Ethernet hosts. These ACK-pairs also become smaller as link rates increase. Further, the variance of wireless $LRTT_{ACK}$ delays will also become smaller as link rates increase. Yet no matter how fast a wireless host is, its smaller packet-pair $LRTT_{FIN,SYN}$ delays will remain large and highly variable. Relatively, this means that wireless $LRTT_{ACK}$ and $LRTT_{FIN,SYN}$ values will show greater separation as wireless technology advances. Thus, simple LRTT cut-offs classify slower wireless hosts, and increased separation of LRTT categories classify 802.11n hosts.

As for infrastructure, gigabit-backbones allow further separations of 802.11n metrics, increasing detection as wireless technology improves. Gigabit-backbones also show better detection overall by distorting LRTTs the least, while megabit-backbones confound wireless host separation. The worst case scenario sees the similar "squeezing" of host LRTTs together as in megabit-backbones, but wireless LRTTs are still distinct enough to stand out. So regardless of network backbone technologies and packet traffic conditions, wireless host differentiation is maintained.

Furthermore, these patterns in LRTT differences hold over different network gigabit- and megabit-backbone infrastructures, regardless of traffic conditions. Lastly, these derivations use only the characteristic link-layer timing differences of Ethernet and wireless, which are as unlikely to change as they are to be faked. Thus this approach to wireless host discovery remains robust against spoofing and wireless technology advances.

References

1. Cisco. 802.11n: The Next Generation of Wireless Performance. (2009), http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps6973/ps8382/prod_white_paper0900aecd806b8ce7_ns767_Networking_Solutions_White_Paper.html
2. Cheng, L., Marsic, I.: Fuzzy Reasoning for Wireless Awareness. International Journal of Wireless Information Networks 8(1), 15–26 (2001)
3. Mano, C.D., et al.: RIPPS: Rogue Identifying Packet Payload Slicer Detecting Unauthorized Wireless Hosts Through Network Traffic Conditioning. ACM Trans. Inf. Syst. Secur. 11(2), 1–23 (2008)
4. Watkins, L., Beyah, R., Corbett, C.: A Passive Approach to Rogue Access Point Detection. In: Global Telecommunications Conference, GLOBECOM 2007. IEEE, Los Alamitos (2007)
5. Mano, C.: Defending against malicious rogue system threats. University of Notre Dame (2006)