On the Impact of 802.11n Frame Aggregation on End-to-End Available Bandwidth Estimation

Arsham Farshad†, Myungjin Lee†, Mahesh K. Marina†, and Francisco Garcia‡
†The University of Edinburgh, ‡Agilent Technologies

***Abstract*—We consider for the first time available bandwidth estimation (ABE) in the context of 802.11n, which is fast replacing the legacy 802.11a/b/g networks. We experimentally show that the frame aggregation (FA) feature of 802.11n is the dominant one among 802.11n features affecting the ABE. Using an indoor 802.11n wireless testbed, we compare three ABE tools (WBest, DietTopp and pathChirp) in various cross-traffic scenarios. We find that FA significantly hurts the accuracy of all ABE tools; DietTopp and pathChirp are relatively more robust than WBest. Because faster available bandwidth estimation and less intrusiveness are desirable properties of any ABE tool and WBest satisfies them relatively better than the other two tools, we conduct an in-depth investigation into the harmful effect of FA on ABE using WBest. This in turn led us to come up with two key design principles to counter FA effects: (1) treating aggregated probes as one jumbo probe; and (2) generating a larger number of probes. We then develop an enhanced version of WBest termed WBest+ that incorporates these principles. Our evaluation shows that the new version is effective in achieving accurate ABE in the presence of FA.**

I. Introduction

End-to-end available bandwidth estimation (ABE) has a wide range of uses including adaptive application content delivery; transport-level transmission rate adaptation and admission control; traffic engineering and peer node selection in peer-to-peer/overlay networks [1], [2]. For instance, adaptive media streaming services (e.g., [3], [4], [5], [6]) keep multiple bitrate versions of each video with different encodings and stream the content to a client with a bitrate that closely matches the available bandwidth along the end-to-end path to the client. Media servers therefore typically employ bandwidth estimation techniques (e.g., packet-pair) within their streaming protocols (e.g., [7]) to accurately estimate available bandwidth for adaptive streaming and enhanced user experience.

As Internet access increasingly happens from wireless enabled devices (laptops, tablets and smartphones), the need for available bandwidth estimation over paths with wireless links – especially WiFi links – takes greater importance given that people spend most of their time indoors where WiFi is prevalent (e.g., homes). The WiFi technology based on the 802.11 suite of standards has evolved significantly in the past decade and a half in view of its widespread use and the growing demand for wireless speeds to match wired Ethernet. The current in the 802.11 series of standards is 802.11n [8], which can provide throughput above the MAC layer nearly reaching 400 Mbps. The follow-on standards in the making (e.g., 802.11ac [9]) promise gigabit wireless speeds. While the physical layer enhancements such as the use of multiple antennas and channel widths are responsible for increased physical layer bit-rates in 802.11n and 802.11ac, MAC layer enhancements especially *frame aggregation (FA)* are key to translating those bit-rates to higher throughputs above the MAC layer [10]. Frame aggregation, as the name suggests, aggregates several frames together and amortizes the protocol overhead (e.g., headers, inter-frame spaces, backoff) over the set of aggregated frames, thereby significantly improving MAC protocol efficiency.

978-1-4799-4657-0/14/$31.00 © 2014 IEEE

Our main goal in this paper is to study the impact of frame aggregation feature in 802.11n on end-to-end available bandwidth estimation. As available bandwidth estimation typically involves the use of active measurement with probing packets (packet pairs, packet trains, etc.), the flow of probing packets can be affected by the use of frame aggregation, which is the rationale underlying our goal. While there exists a substantial body of work (e.g., [11], [12], [13], [14], [15], [16], [17]) examining available bandwidth estimation with 802.11 wireless links, it only considers the case of legacy 802.11 networks (802.11a/b/g) which are fast becoming out-of-date. We begin to address this deficiency in this paper by shedding light on the effect of FA – a key feature of current 802.11n – on available bandwidth estimation.

Towards this end, we conduct a measurement based study using an indoor 802.11n wireless testbed. For our study, we consider WBest [13], DietTopp [12] and pathChirp [18] as three representative yet widely different ABE tools given that all three of them have been previously evaluated in the legacy 802.11a/b/g context. A more detailed justification behind the choice of these specific tools is given in §II. We study their available bandwidth estimation accuracy in the presence of 802.11n frame aggregation across a wide range of cross-traffic scenarios.

Our study leads to the following contributions and findings:

- We consider available bandwidth estimation in the 802.11n context for the first time, and show that frame aggregation is the most dominant feature among 802.11n features affecting ABE (§II-C).
- Our comparison of different ABE tools in various cross-traffic scenarios focusing on the impact of FA leads to the following observations (§IV):
  - The FA feature significantly hurts the accuracy of all ABE tools considered.
  - DietTopp and pathChirp, the tools that follow the

  - Probe Rate Model (PRM), are relatively more robust in the presence of FA compared to WBest that belongs to the Probe Gap Model (PGM) category.
  - Different cross-traffic scenarios affect the ABE tools differently. We also find that with FA underestimation is more predominant across all tools and scenarios.
- Keeping in mind that PGM based tools like WBest are better suited for fast ABE needed for adaptive multimedia streaming services, our motivating use case, we take a deeper look at the FA effect on the working of WBest and come up with the two general principles of jumbo probes and a larger number of probes that together make up our solution approach for improved ABE in the presence of FA (§V). We then develop an enhanced variant of WBest which incorporates our approach and show that it is indeed effective in achieving robust and accurate ABE over FA-enabled 802.11n networks. Although our approach is evaluated only in the context of WBest, we believe that it is more widely applicable.

II. Background and Motivation

A. Frame Aggregation (FA)

Frame aggregation is a MAC layer enhancement and a mandatory feature of 802.11n/ac to improve protocol efficiency, i.e., to translate the physical layer bit-rates to comparable throughputs above the MAC layer. Protocol efficiency is a major concern underlying the design of 802.11n because of two reasons: (1) physical layer bit-rates with 802.11n are up to an order of magnitude higher compared to the earlier legacy standards of 802.11a/b/g (54Mbps vs. 600Mbps); (2) the protocol overhead (medium access, header overhead, inter-frame spaces) has a more harmful effect on higher layer throughput at higher physical layer rates.

The idea behind FA is simple: spread the protocol overhead over several frames. For FA to be effective, it needs a companion feature called block acknowledgments (similar to the selective ACK feature of TCP).

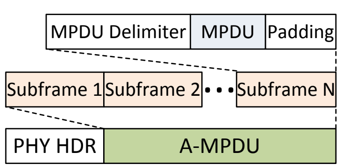

Fig. 1. A-MPDU type frame aggregation.

802.11n specifies two types of frame aggregation: the Aggregate MAC Protocol Data Unit (A-MPDU) and the Aggregate MAC Service Data Unit (A-MSDU). A-MPDU corresponds to aggregating multiple MPDUs (subframes) in the MAC layer, where MPDU (subframe) refers to a valid 802.11 MAC frame with MAC header, one IP packet as payload and a frame check sequence (FCS). A-MSDU, on the other hand, aggregates several IP packets above the MAC layer and puts them into one MAC frame with a common MAC header and FCS. A-MPDU is the most widely supported and popular option, so we only focus on A-MPDU frame aggregation and simply refer to it as frame aggregation or FA in the rest of the paper. The A-MPDU type FA is illustrated in Figure 1.

B. Available Bandwidth Estimation (ABE)

The available bandwidth on a link is the *unused capacity* on that link available for new flows. The available bandwidth on a path (or end-to-end available bandwidth) is the minimum available bandwidth across all links on the path. The link determining the path's available bandwidth is commonly referred to as the *tight link* [19]. Clearly, the end-to-end available bandwidth is upper bounded by the *path capacity*, which is the maximum data rate that can be supported by the path. A path's capacity is limited by the link with the smallest capacity (also called the *narrow link*).

On wireless links (including ones based on 802.11), link capacity is dependent on the maximum reliable physical layer bit-rate and that is in turn dependent on the channel condition at the time of measurement. Moreover, the *effective capacity* of a 802.11 link above the MAC layer is much lower than the actual link capacity due to overhead related to protocol headers, inter-frame spaces and backoffs. In this paper, we use the term 'capacity' to refer to effective capacity. Besides capacity, the available bandwidth is dependent on the utilization of the link/path by traffic from existing flows on any link (wireless or wired) on the end-to-end path; this traffic is commonly referred to as the *cross-traffic*. As cross-traffic can be time-varying, so is available bandwidth. Because of contention based access to the shared wireless medium in 802.11 networks, the type and amount of cross-traffic has a potentially big influence on the available bandwidth. Note that depending on the nature of cross-traffic the tight link needs not necessarily be the narrow link.

In this paper, we limit our focus to the *active measurement* approach. While passive measurement of available bandwidth or utilization (through measuring busy and idle times [20], [21], [22]) may be appropriate for the last hop path segment with 802.11 wireless link, it requires cooperation of intermediate access points (APs) for end-to-end ABE, making it less practical than the active measurement approach. Moreover, passive techniques usually rely on lower level information from the system or device drivers which may require superuser privileges (e.g., for packet captures) or device driver manipulation.

Most active ABE techniques can be classified as belonging to one of two models [23], [15], [17], [2]: 1) Probe Gap Model (PGM) and 2) Probe Rate Model (PRM). We briefly describe each of them in the following.

1) Probe Gap Model: This model relies on the observation that successive probe packets traversing a tight link undergo increased dispersion in time due to cross-traffic. The rate of cross traffic (alternatively, the utilization of the tight link) is estimated via increase in dispersion experienced by back-to-back probe packets. Available bandwidth is then computed using the estimated cross-traffic rate and capacity; the latter is either assumed to be known or separately estimated. Examples of PGM based tools include: Spruce [23] and WBest [13].

2) Probe Rate Model: This model is based on the notion of self-induced congestion. The essential idea is that when the probe traffic rate exceeds the available bandwidth of a path, the measured reception rate of probes starts lagging behind their sending rate (or equivalently, the inter-probe arrival times measured at receiver keep increasing). The maximum probe traffic rate at which this transition occurs is taken as the available bandwidth estimate. PRM based tools usually are iterative spanning several rounds to probe at different rates. Several tools such as Pathload [24], DietTopp [12] and pathChirp [18] belong to this category. These tools differ in their probing traffic patterns and receiver-side statistical analysis mechanisms.

We in our study choose WBest, DietTopp and pathChirp because they represent the two different ABE models (PGM and PRM) described above and also because they are known to work well in Wi-Fi environments. WBest has been specifically designed for (legacy) 802.11 wireless LANs and represents the PGM class of ABE tools. DietTopp, on the other hand, falls in the PRM class of tools and has been considered in various evaluation studies of ABE over 802.11 wireless LANs (e.g., [12], [14]). We include pathChirp, which also belongs to the PRM category, as it has been shown to yield good results in some ABE evaluation studies in wireless network settings (e.g., [15], [25]). Moreover, DietTopp and pathChirp, though both are from the same PRM class, differ in their probing traffic pattern, thereby allowing us to understand the relative effectiveness of different probing patterns. In addition, all three of them are publicly available. We next briefly describe these tools.

WBest. It consists of two steps. In the first step, a number of probe packet pairs are sent to estimate the effective capacity. The second step then transmits a probe packet train at the estimated effective capacity rate to estimate the available bandwidth. Finally, the packet loss rate experienced by the probe train is used to correct the estimated available bandwidth.

DietTopp. As a typical PRM based tool, it operates over multiple rounds, transmitting a probe packet train with an increased rate in each successive round. The highest sending rate with a matching receiving rate is reported as the available bandwidth estimate.

pathChirp. pathChirp employs a different probing traffic pattern called a chirp, consisting of multiple exponentially spaced probe packets of the same size, to improve efficiency and accuracy over the earlier TOPP [26] and Pathload tools. Wide range of rates can be probed within a single chirp, thereby improving efficiency. It performs statistical analysis at the receiver to estimate available bandwidth from multiple chirps.

C. Why Frame Aggregation?

Before we conduct our in-depth measurement study, we first elaborate on the rationale behind our choice to focus on frame aggregation among the various new features introduced as part of 802.11n that also include channel bonding (CB) and spatial division multiplexing (SDM).

Figure 2 shows absolute errors for the different ABE tools used when each of these three features is disabled relative to the case where all of them are enabled. Clearly, the absence of FA leads to the most reduction in ABE error for all tools, thus demonstrating that FA has the biggest impact on ABE in comparison with SDM and CB.

We therefore focus on FA and take an in-depth look at its impact on available bandwidth estimation with various ABE tools and in different cross-traffic scenarios. To isolate the effect of FA on ABE, we disable other 802.11n features including SDM and CB in the remainder of this paper and leave the characterization of the impact of those other features on ABE for future work. For the same reason, we factor out the link quality and rate diversity-related effects by ensuring all links operate reliably at the maximum bit-rate possible without SDM (i.e., MCS index 7 in 802.11n which corresponds to 65Mbps physical layer bit-rate).

D. Related Work

The available bandwidth estimation topic has received much attention from the research community with many ABE tools developed, mostly for wired networks (e.g., TOPP [26], Pathload [24], pathChirp [18], Spruce [23]) but some specifically for wireless networks (e.g., WBest [13]). There also exist several comparison studies of ABE tools for wired networks, some in the context of a new proposal (e.g., [18], [23]) while in other cases comparing a larger number of ABE tools in a common set of scenarios (e.g., [1], [27], [2]).

Lakshminarayanan et al. [11] were among the first who highlighted the unique challenges posed by 802.11 wireless networks to capacity and available bandwidth estimation due to use of multiple physical bit-rates, shared access and contention. More recently, a detailed analysis of the impact of multiple access contention related delays with 802.11 CSMA/CA on active bandwidth measurements is presented in [16].

A number of experimental performance evaluation studies have focused on available bandwidth estimation with wireless links, most of them in the context of 802.11 wireless LANs and mesh networks (e.g., [12], [13], [14], [15]) but a few also consider cellular networks (e.g., [25]). Koutsonikolas and Hu [17] consider both 802.11 and cellular links in their study. *However, none of the above examine available bandwidth estimation in now commonplace 802.11n wireless networks, the main focus of our work.*

III. METHODOLOGY

Frame aggregation can affect the available bandwidth estimation (e.g., by causing multiple probe packets to get packed inside a single frame). Our broad aim is to characterize the impact of FA on end-to-end ABE in various 802.11n wireless LAN (WLAN) scenarios. For this study, we choose three representative active measurement tools for ABE (as already noted in the previous section): WBest, DietTopp and pathChirp.

Testbed. We take an experimental approach using an indoor 802.11n wireless LAN testbed (illustrated in Figure 3) that emulates typical 802.11 WLAN deployments in home and hotspot environments in a simplified form. It consists of two co-located Wi-Fi networks—the available bandwidth estimation occurs

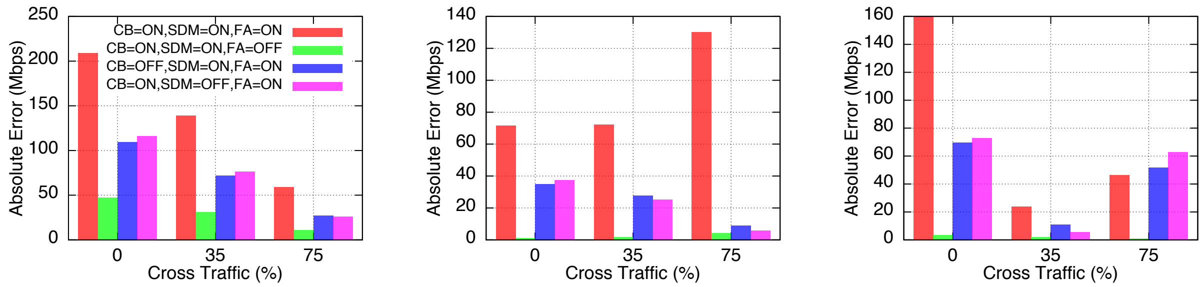

(a) WBest (b) DietTopp (c) pathChirp

Fig. 2. Relative impact of key 802.11n features (FA, channel bonding (CB) and spatial division multiplexing (SDM)) on the accuracy of different ABE tools in Scenario 1 (see Figure 3).

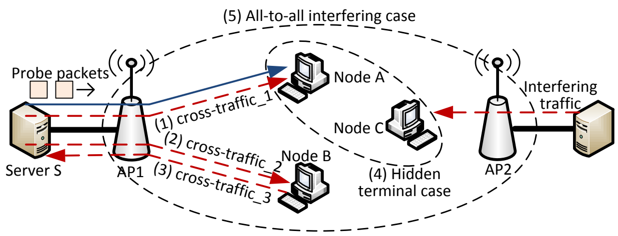

Fig. 3. Schematic of testbed and various cross-traffic scenarios.

from a server to a client (node A) in the left network while the right network causes interfering traffic. In the left network, two 802.11n clients (nodes A and B) are connected to 802.11n enabled access point $AP1$ over a real wireless channel and $AP1$ is connected to server S via a Gigabit Ethernet link. The Gigabit link is chosen to reflect the recent increased availability of high-speed broadband services [28], thus rendering even 802.11n wireless access links bandwidth-limited in our setting (wireless link capacity is 60Mbps with FA only enabled). In the right network, node C is connected to another server via access point $AP2$. All the 802.11n hardware in our testbed is based on Atheros chipsets and is used via the the ath9k wireless driver. To avoid external interference from other operational WiFi networks in the surrounding environment of our testbed, we set both $AP1$ and $AP2$ (and as a consequence for all client nodes A, B and C as well) to operate on channel 149 in 5GHz band which we identified to be unused by other external networks.

Cross-traffic scenarios. As the extent of FA influence on ABE potentially depends on the nature of cross traffic, we consider a wide range of cross-traffic scenarios, reflecting some of the key types of cross-traffic that would occur in practice. These scenarios are illustrated in Figure 3 and described below. The first three scenarios model various types of cross-traffic within a single AP WLAN setting, whereas the last two model cross-traffic due to interference from a co-located WLAN. We consider probing traffic in the downstream direction to WLAN client (node A) for all scenarios to reflect a case where a multimedia streaming server wants to determine the available bandwidth to a user. The level/amount of cross-traffic in each scenario is a variable parameter. In all scenarios, cross-traffic is generated as a UDP flow using the well-known Iperf [29] tool with default packet size of 1470 bytes and a specified generation rate to realize different levels of cross-traffic.

- *Scenario 1: Single node case.* In this scenario, there are two flows destined to node A from server S; one is a probing measurement flow using one of the three ABE tools considered (WBest, DietTopp or pathChirp) and the other is a cross traffic flow (cross-traffic_1 in Figure 3). This scenario models cross-traffic that reflects other downstream application traffic such as P2P file download.
- *Scenario 2: Cross-traffic to Node B.* Different from Scenario 1, this scenario models a situation where another user (node B) within the WLAN (e.g., a home WiFi network) competes with node A for network bandwidth, for instance, via file downloading or web browsing application traffic. It is the only source of cross-traffic in this scenario and shown as cross-traffic_2 in Figure 3.
- *Scenario 3: Cross traffic from Node B.* Cross-traffic (denoted as cross-traffic_3 in Figure 3) passes through $AP1$ from B to S. Thus, as opposed to Scenario 2 where $AP1$ arbitrates channel access to both measurement and cross-traffic clients, B and $AP1$ contend for channel access in this scenario.
- *Scenario 4: Hidden terminal case.* $AP1$ and $AP2$ are outside the communication range of each other but their associated client nodes (A and C) are close enough to hear each other and transmissions from both APs.
- *Scenario 5: All-to-all interference case.* All nodes (access points $AP1$ and $AP2$ and their associated clients A and C) hear from and talk to each other. This scenario models a commonly occurring situation with neighboring home WiFi networks with overlapping coverage areas. Note that in Scenarios 4 and 5, the only source of cross-traffic is from the other server to node C via $AP2$.

Obtaining ground truth. For each of the cross-traffic scenarios, we need *true* available bandwidth to assess the accuracy of different ABE tools under consideration. For this, we follow an approach similar to that taken in [17]. For Scenarios 1–3, we use backlogged Iperf UDP flow (for 100 seconds) on the probing measurement flow path to find out the true capacity and subtract the level of cross-traffic injected to obtain the true available bandwidth. In Scenarios 4 and 5 with interference traffic from another WLAN, the actual available bandwidth is computed as the throughput obtained for a backlogged Iperf UDP flow on the probing measurement flow path while cross-traffic on the interfering path is present at the specified level.

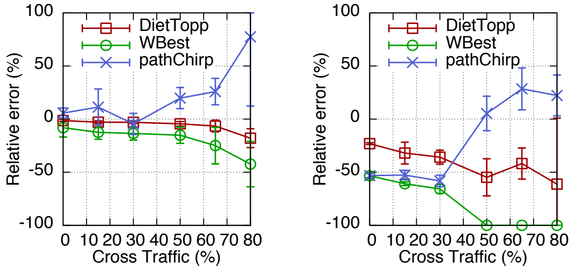

(a) FA-OFF (Capacity=30Mbps) (b) FA-ON (Capacity=60Mbps)

Fig. 4. Accuracy of DietTopp,WBest and pathChirp in Scenario 1 at varying levels of cross-traffic. The median relative error across all cross-traffic levels for DietTopp, WBest and pathChirp are 4%, 14% and 16% respectively when FA is OFF; and 39%, 83% and 24% when FA is ON.

Metrics. Our primary focus is on studying the accuracy of different ABE tools in the presence of FA in different cross-traffic scenarios and levels of cross-traffic. We consider the following two metrics to quantify accuracy of a tool: (1) *absolute error* computed as $|Estimated - True|$, i.e., the absolute difference between the estimated and true available bandwidth values; (2) *relative error* computed as $(Estimated - True)/True$. Each data point in our plots is an average of at least 20 runs and the error bars show standard deviations.

We also touch upon two other metrics commonly considered when evaluating an ABE tool, namely measurement duration/latency and intrusiveness (measurement overhead). Note that the method used for obtaining true available bandwidth (i.e., via a backlogged Iperf UDP flow) is not suitable for ABE in practice because it is highly intrusive in terms of overhead compared to ABE tools with carefully chosen probing traffic (packet pairs, trains or chirps).

IV. Performance of Current ABE Tools in FA-enabled 802.11n Scenarios

In this section, we compare WBest, DietTopp and pathChirp in various cross-traffic scenarios focusing on the impact of FA on their accuracy.

A. Scenario 1: Single node case

In Scenario 1, contention and interference among multiple links is absent, so FA can have a more pronounced effect. We evaluate the accuracy of WBest, DietTopp and pathChirp with(out) FA. All other 802.11n features are disabled. Results are shown in Figure 4. The amount of cross-traffic is varied and shown as percentage values on the x-axis with respect to the path capacity. For 100% cross-traffic rate, some of the tools under test did not converge so results are only reported till 80% cross traffic rate.

ABE accuracy when FA is OFF. Figure 4(a) illustrates that DietTopp, WBest and pathChirp achieve good ABE accuracy across most cross traffic rates; DietTopp performs better than WBest and pathChirp. DietTopp suffers at most 18% error at 80% cross traffic rate (*i.e.*, 80% of capacity) while WBest has 43% error at the same rate. DietTopp and WBest tend to underestimate ABE. pathChirp, on the other hand, exhibits overestimation, which coincides with the conclusion in [27].

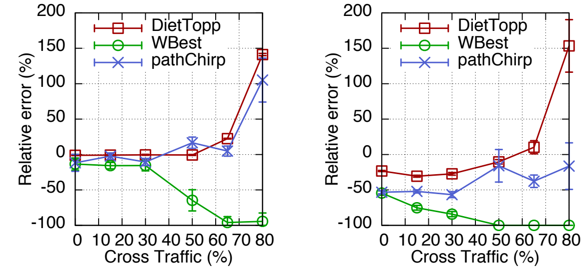

(a) FA-OFF, Scenario 2 (b) FA-ON, Scenario 2

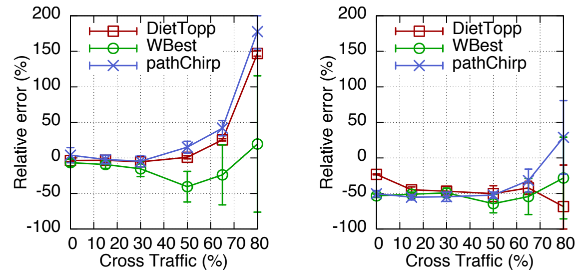

(c) FA-OFF, Scenario 3 (d) FA-ON, Scenario 3

Fig. 5. Accuracy of the three tools in Scenarios 2 and 3 at varying levels of cross-traffic. True capacity is 30Mbps for FA-OFF and 60Mbps for FA-ON. The median relative error across all cross-traffic levels for DietTopp, WBest and pathChirp are: 1%, 40% and 1% when FA is OFF in Scenario 2; 72%, 92% and 45% when FA is ON in Scenario 2; 1%, 12% and 10% when FA is OFF in Scenario 3; 46%, 52% and 51% when FA is ON in Scenario 3.

Also note that the relative error metric somewhat amplifies the error values with increasing cross-traffic rates as true value of available bandwidth (*i.e.*, denominator in relative error computation) correspondingly decreases.

ABE accuracy when FA is ON. Compared to the FA OFF case, enabling FA increases estimation error by at least 20% for DietTopp and 53% for both WBest and pathChirp even in an idle link with no cross-traffic (i.e., at 0% cross-traffic rate). The estimation error increases up to 60% for DietTopp and 100% for WBest as the amount of cross traffic increases (WBest reports 0Mbps available bandwidth when cross traffic rate gets larger than 50%). One notable observation is that pathChirp initially underestimates ABE at lower cross-traffic rates, works most accurately at 50% cross-traffic rate and then begins to overestimate ABE. On the other hand, other tools consistently output less ABE than the ground truth.

B. Scenarios 2 and 3: Cross-traffic to/from node B

We now study Scenarios 2 and 3 with wireless contention. Recall from §III and Figure 3 that the only difference between Scenario 2 and Scenario 3 is the direction of cross-traffic.

FA OFF case. From Figures 5(a) and 5(c), we see that DietTopp and pathChirp (both PRM based tools) work better than WBest (a PGM based tool) in most cross traffic rates in both scenarios. pathChirp is most stable in Scenario 2 while it experiences highest overestimation errors with higher cross-traffic rates in Scenario 3. Similar behavior is seen with DietTopp in both scenarios, which is somewhat different from the observation made in Figure 4(a). However, this is not too surprising because it is known that PRM based tools like DietTopp report the fair share bandwidth rather than available bandwidth in case of a fair wireless link [14]. Thus, DietTopp overestimates the available bandwidth by reporting the fair share when the cross traffic rate is more than 50% of the capacity. Another notable observation is that WBest in Scenario 2 exhibits stable estimation performance till 30% cross traffic rate but becomes very erroneous quickly, yielding -100% error from 65% cross traffic rate. On the other hand, in Scenario 3, as the amount of cross traffic increases (at the mark of 50% cross traffic rate), the further underestimation with WBest stops, the estimation error becomes smaller and the tool eventually produces overestimates at 80% cross traffic rate.

FA ON case. WBest behaves similarly as shown in FA OFF case (see Figure 5(a) and 5(c)) in both Scenarios 2 and 3 although its performance in FA ON case is worse than in FA OFF case. pathChirp exhibited overestimation trends in the FA OFF case whereas in the FA ON case it underestimates available bandwidth for most cross traffic rates in both scenarios. One unique phenomenon from Figures 5(b) and 5(d) is that DietTopp shows different behaviors between Scenario 2 and Scenario 3. More specifically, its behavior in Scenario 2 follows a similar trend to the corresponding FA OFF case, whereas underestimation becomes worse as the cross traffic rate increases in Scenario 3.

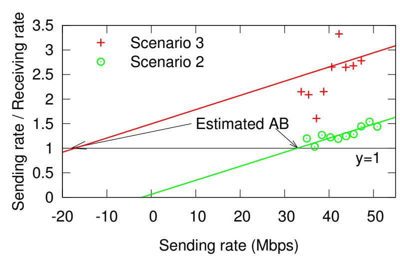

Fig. 6. Different estimation behaviors of DietTopp at 80% cross-traffic rate in Scenarios 2 and 3 when FA is ON.

To understand why, we analyze the behavior of DietTopp at 80% cross-traffic rate in Scenarios 2 and 3 (150% overestimation error in the former scenario while 70% underestimation error in the latter as shown in Figure 5(b) and Figure 5(d)). Figure 6 shows how DietTopp estimates available bandwidth based on measurement samples. In the figure, each point represents one measurement sample given the sending rate of a probe train. The intersection of $y = 1$ line with the trend line (obtained by linear regression of the measurements points) is the available bandwidth (denoted as 'Estimated AB' in Figure 6) estimated by DietTopp [12]. As mentioned before, in Scenario 2 where the AP coordinates the wireless medium access for both measurement and cross-traffic, the fair share nature of DietTopp is preserved. As a result, the receiving rates of 28–34Mbps are observed in the Scenario 2. These receiving rate samples form the regression line (presented in green color) that meets the $y = 1$ line at 32Mbps sending rate. Therefore, we have an over-estimate as the true available bandwidth is 12Mbps. On the other hand, the wireless channel is not equally shared in Scenario 3. The intersection of the $y = 1$ line with the regression line (presented in red color) constructed with the data points of the receiving rates of only 16–20Mbps is found at a negative sending rate, thus leading to an under-estimate.

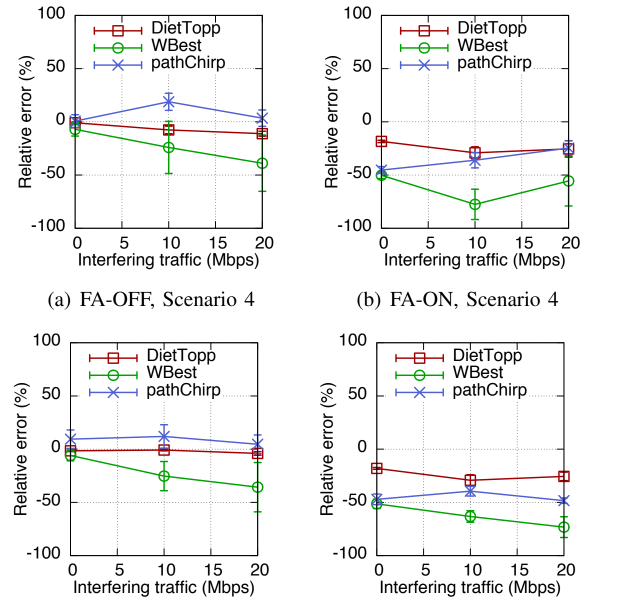

(a) FA-OFF, Scenario 4 (b) FA-ON, Scenario 4

(c) FA-OFF, Scenario 5 (d) FA-ON, Scenario 5

Fig. 7. Accuracy of the three tools in Scenarios 4 and 5. The median relative error across all cross-traffic levels for DietTopp, WBest and pathChirp are: 8%, 24% and 3% when FA is OFF in Scenario 4; 27%, 59% and 33% when FA is ON in Scenario 4; 2%, 25% and 9% when FA is OFF in Scenario 5; 26%, 63% and 47% when FA is ON in Scenario 5.

C. *Scenarios 4 and 5: Interference cases*

We now present results for Scenarios 4 and 5. Figure 7 shows results for interfering traffic only till 20Mbps (the maximum traffic rate supported in Scenario 4 when FA is OFF) for consistency across both scenarios and with(out) FA.

FA OFF case. Figures 7(a) and 7(c) show that both DietTopp and pathChirp achieve less than 20% error overall across both scenarios. On the other hand, WBest quickly loses its accuracy as the amount of cross traffic increases; the average error with interfering traffic of 20Mbps is almost 40%. This shows that WBest is more susceptible to interference than other tools. We also observe that DietTopp appears to be more robust in Scenario 5 than Scenario 4.

FA ON case. All three tools become more erroneous as compared to FA OFF case in both scenarios. DietTopp experiences at least $3\times$ higher error in Scenario 4 with FA ON than with FA OFF, and at least $5\times$ higher in Scenario 5. Similarly, pathChirp exhibits $3\times$ worse accuracy as compared to FA OFF case. Moreover, there is no noticeable difference in accuracy of the tools between the two scenarios.

D. *Summary*

We make three key observations from this section:

- FA has a detrimental impact on the accuracy of all the three tools (DietTopp, WBest and pathChirp). In many instances, the estimation errors increased by at least three times with FA compared to the case when FA is OFF.

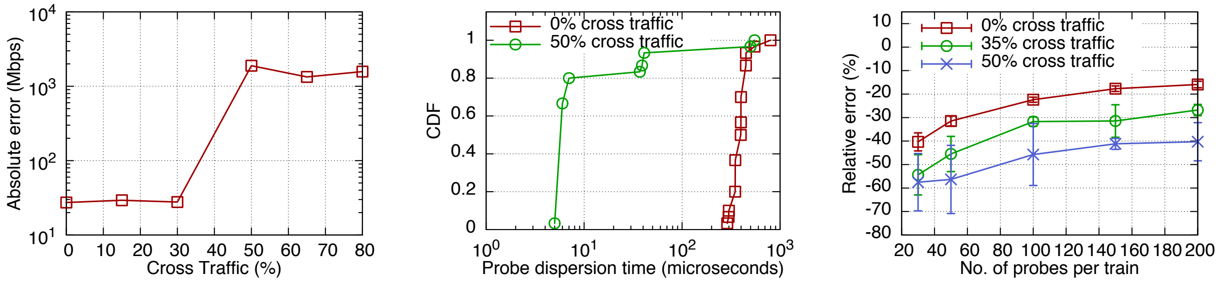

(a) (b) (c)

Fig. 8. (a) Accuracy of capacity estimation with WBest; (b) CDF of dispersion times from 30 probe pairs used for WBest's capacity estimation phase; (c) Accuracy of WBest with known capacity for varying probe train sizes and cross-traffic levels. All cases correspond to Scenario 1 when FA is ON.

- With FA ON, PRM based tools, DietTopp and pathChirp, generally outperform PGM based tools (WBest). As we show in the next section, frame aggregation distorts the dispersion times. PGM based tools which rely on packet dispersion times are therefore adversely affected.
- The tools behave differently in different scenarios. This is particularly evident between Scenarios 2 and 3. As the cross traffic amount increases, WBest completely fails to estimate available bandwidth in Scenario 2 (100% error because it only produces 0Mbps at some point onwards; we will look into this in the next section) but it achieves lesser error (about less than 70%) in Scenario 3. When FA is ON, DietTopp also shows very different behaviors between those two scenarios: overestimation vs. underestimation. Somewhat surprisingly, however, underestimation is more dominant than overestimation across all the tools.

V. Towards More Accurate ABE in Presence of Frame Aggregation

A. A Closer Look at the Problem using WBest

The previous section revealed that while the accuracy of all ABE tools is adversely affected by FA, the PRM based tools (DietTopp and pathChirp) fare relatively better. *However measurement latency is an issue for PRM tools. For instance, for DietTopp and pathChirp, we find that the measurement period can range from 5 to 11 seconds.* On the other hand, WBest representing PGM tools does estimation in less than a second. Faster ABE is crucial because applications such as multimedia services need available bandwidth estimate in a short time scale to effectively tune the streaming rate. So, even though WBest is seen to be more erroneous, given its faster measurement property, we choose to take a closer look at its behavior to better understand the effect of FA and identify the root causes of the problem. Also as noted at the outset, a packet pair technique similar to one used in WBest is already adopted in MS media server applications [7].

In the rest of this section, we first do an in-depth analysis on how FA impacts the performance of WBest. We then propose our approach for accurate ABE in presence of FA and apply it to WBest as a case study. Recall that WBest consists of two phases: capacity estimation and available bandwidth estimation. It sends out 30 packet pairs to estimate capacity and uses a packet train with 30 packets to estimate available bandwidth based on the capacity estimate. We start by examining the capacity estimate phase in presence of FA.

Analysis of capacity estimation phase. Our analysis shows that packet pair technique used for capacity estimation in WBest yields either underestimates (half of the true capacity) or extreme overestimates (21–31× higher than the true). This is mainly because probes in a pair either arrive together in an aggregated frame or separately in different frames. Both these patterns harm capacity estimation — separate arrival of a packet pair means that probes in the pair do not experience the benefit of FA (doubling the capacity as compared to the legacy 802.11); on the other hand, aggregation trips the capacity estimator with a small dispersion time which leads to too much overestimation. Figure 8(a) shows such trends. True capacity is 60Mbps, but estimates are only 30Mbps until 30% cross traffic rate. From that point onwards, WBest suddenly yields 1.3–1.9Gbps as its capacity estimate (note the log-scale for the y-axis). For deeper understanding, we choose two data points: 0% and 50% cross traffic rates and analyze the CDF of dispersion times from the 30 probe packet pairs (shown in Figure 8(b)). When there is no cross traffic, almost 90% of packet pairs have larger than 300μs dispersion time (because they do not get aggregated). With 50% cross traffic, on the other hand, 93% of dispersion times are less than 41μs (due to aggregation). We confirmed the segregation and aggregation phenomena by looking into the packet traces captured over the air using AirPcap [30].

Analysis of available bandwidth estimation phase. In this analysis, to shield WBest from the impact of wrong capacity estimate, we modify WBest so that it is configured with the true capacity (60Mbps) and let it only carry out the ABE stage. In addition, we also vary the length of packet train to see what impact different train lengths may have on the accuracy.

From Figure 8(c), we first learn that the modified version only obtains under-estimates regardless of the cross traffic rate and the length of probe train. Note that under-estimation with increasing cross-traffic rates has been seen before in Figure 4(b). Our second observation is that we obtain more accurate estimates as we increase the probe train length — we see an extra 22% improvement (e.g., from -40% error to -18% error in case of 0% cross traffic rate). This is related to average dispersion time of probes. Note that the available bandwidth estimate value in WBest is inversely proportional to the the average dispersion time [13]. Fewer number of probes

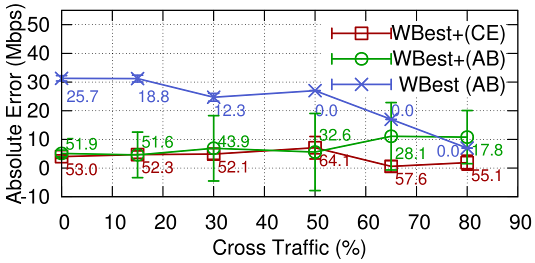

Fig. 9. Accuracy with WBest+ in Scenario 1 when FA=ON. Numbers next to data points denote actual estimation values.

in the train result in smaller number of aggregated probes, thereby making it more likely to observe large gaps between probes. On the other hand, as the probe count increases, the average becomes smaller due to higher probability of probe aggregation. We confirmed this trend by analyzing WBest dispersion time logs. In our experiments, 30 probes produce an average dispersion time of 258.5μs and 100 probes 228.2μs in the 0% cross traffic case. Similarly we find, in 35% cross traffic case, an average dispersion time of 290.28μs for 100 probes and 356.41μs for 30 probes. Thus, we see that the underestimation can be reduced by increasing the probe count.

B. Our Approach for Mitigating FA Effects

Our analysis on WBest in the previous subsection showed that FA can cause widely different dispersion times depending on whether probe packets can get aggregated or not. Directly using such dispersion times in the capacity or available bandwidth estimation can end up causing under-/over-estimations. While we have seen this happen in WBest's capacity estimation phase, the same applies even for the available bandwidth estimation. Another observation is that a larger number of probes are helpful in getting better estimates. Based on these observations, we identify two key principles for accurate ABE in the presence of FA: *i) treating aggregated probes as one jumbo probe* and *ii) increasing the number of probe packets.* These two principles make up our solution approach. We will describe each of them below before applying them to WBest.

1) Treating aggregated probes as one jumbo probe. As we already studied, FA creates minute gaps between probes, which makes WBest overestimate capacity too much. However, the small probe gaps are not a symptom that is unique to FA. Interrupt coalescing done in modern computer systems is another source with the same effect. Existing approaches including those used to mitigate interrupt coalescing related measurement noise (e.g., [18], [24], [31]) perceive small dispersion times as abnormal samples and discard them. However, given that FA actually plays a role of increasing capacity, unlike interrupt coalescing, we cannot simply apply the same approach to the FA problem. Doing so will mean considering only large dispersion times; this can lead to underestimations, which was seen in many cases of pathChirp (see Figures 4(b), 5(b), 5(d), 7(b) and 7(d)). Instead, we treat aggregated probes as one jumbo probe. Our rationale behind this principle is that if probes are aggregated, they are transmitted over the 802.11n link as part of the aggregated frame and not as individual probes. However, at an application level, probes are already decapsulated from a frame, which makes it difficult to identify which probes belonged to which frame. We reconstruct aggregated probes by using an observation that probes in the same frame tend to have a small interval between them. We find that this approach accurately clusters probes. Now that we have the notion of jumbo probes and have a way to identify such probes at the application layer in the receiver, the capacity (available bandwidth) is estimated by computing instantaneous samples of capacity (available bandwidth) with each received jumbo probe and applying a statistic across all samples (e.g., maximum for capacity and average or median for available bandwidth). Note that with our jumbo probe approach individual dispersion times have no bearing on the bandwidth estimations, thus contributing to robustness in the presence of FA.

2) Increasing probe packet counts. With a small number of probes, there is a possibility of probes getting aggregated in a single frame and leading to few or no measurement samples at the jumbo probe level. Thus, to overcome this issue, increasing the number of probes is necessary. In addition, it was also shown in §V-A that more probes help improving accuracy. Choosing the right number of probes depends on several factors including the probe generation model, 802.11n hardware etc. Given this we take an empirical approach to determine the number of probes required.

Case Study: application of our approach to WBest. We modified both capacity estimation and ABE steps of WBest to incorporate our proposed principles described above. We empirically found out that a minimum of 100 probes are required for capacity and ABE as we need to have enough samples at the jumbo probe level. This probe count increase roughly adds extra 35ms on average to the total ABE time according to our evaluation. For capacity estimation, a total of 100 probes are sent in bursts of 15 probes each (this burst size is empirically determined); 100ms gap is imposed between two successive bursts to avoid interference between them. This ensures at least two jumbo probes. An instantaneous capacity sample is obtained from each jumbo probe and maximum of all such samples is used as the capacity estimate. In the ABE phase, we send a single train of 100 probes (as opposed to 30 with WBest) at the rate of estimated capacity. Unlike vanilla WBest, we compute instantaneous available bandwidth estimates for each successive pair of jumbo probes; this is more similar to how Spruce [23] does ABE with packet pairs and is less sensitive to dispersion time variations. Altogether, we call this modified version WBest+.

Figure 9 shows capacity and available bandwidth estimation for WBest+. Comparing the capacity estimation of WBest+ (noted as WBest+ (CE) in Figure 9) with WBest (Figure 8(a)) shows that our principles can improve capacity estimation significantly even for high cross traffic rate. As for ABE, WBest for lower cross traffic has 25Mbps to 30Mbps error, and it only reports zero for cross traffic more than 50%. As WBest

(AB) and WBest+ (AB) curves in Figure 9 show, WBest+ still works better than WBest in terms of available bandwidth estimation for most of the cross traffic rates. While WBest seems to work better than WBest+ when cross traffic rate is at its highest at 80% rate, it is just an artifact of WBest reporting 0Mbps available bandwidth (indicating that the vanilla WBest completely fails to provide any estimate). We also found that WBest+ outperforms DietTopp and pathChirp (achieving about 20% higher accuracy than DietTopp and 5% than pathChirp) by looking into the median of relative errors across all cross traffic rates (not shown for brevity). Although our proposed mitigation approach (based on the principles of jumbo probes and a larger number of probes) is only applied to WBest in this paper, we believe it is more generally applicable (especially the jumbo probe principle). Validating this assertion is an issue for future work.

VI. Conclusions

The advent of 802.11n standard with its various features has made high-speed wireless Internet access possible. We have conducted the first investigation of the end-to-end available bandwidth estimation (ABE) on paths with 802.11n links. In particular, we have experimentally shown that frame aggregation (FA), one of the features of 802.11n, has a major impact on ABE accuracy in comparison with other 802.11n features. Given this, we have experimentally studied the impact of FA on ABE, considering three representative ABE tools (WBest, DietTopp and pathChirp) and comparing their accuracy in presence of FA in various cross-traffic scenarios.

Our results have shown that FA seriously harms the accuracy of all ABE tools as it distorts their probing traffic patterns. DietTopp and pathChirp belonging to the PRM class of tools are relatively robust to FA but the measurement latency with them is considerably longer to be suitable for some compelling ABE applications like adaptive multimedia streaming services. So we have conducted an in-depth analysis of the relatively faster WBest, a representative of PGM tools, to better understand the effect of FA on ABE tool behavior. This analysis has led us to our solution approach that is rooted in two key principles—treating aggregated probes as one jumbo probe and generating a larger number of probes—for robust ABE in presence of FA. We have developed an enhanced variant of WBest that incorporates our approach; the evaluation results of this variant confirm the efficacy of our proposed approach. As part of future work, we would like to apply our approach to other ABE tools (e.g., DietTopp and pathChirp) and also conduct a detailed investigation of the impact of other 802.11n features (e.g., channel bonding) on FA.

References

[1] C. D. Guerrero and M. A. Labrador. On the Applicability of Available Bandwidth Estimation Techniques and Tools. *Computer Communications*, 33(1), Jan 2010.

[2] A. Botta et al. *Active Techniques for Available Bandwidth Estimation: Comparison and Application*. Springer, 2013.

[3] Alex Zambelli. IIS Smooth Streaming Technical Overview. Microsoft Corporation, 2009.

[4] David Hassoun. Dynamic streaming in Flash Media Server 3.5 – Part 1: Overview of the new capabilities. http://www.adobe.com/devnet/adobe-media-server/articles/dynstream_advanced_pt1.html.

[5] R. Pantos and W. May. HTTP Live Streaming. IETF Internet-Draft, Sept 2011.

[6] Akamai HD Network. http://www.akamai.com/dl/brochures/Product_Brief_Sola_HDN.pdf.

[7] Packet-Pair Bandwidth Estimation. http://msdn.microsoft.com/en-us/library/dd763927.aspx.

[8] IEEE 802.11n. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Enhancements for Higher Throughput, 2009.

[9] E. Perahia and M. Gong. Gigabit Wireless LANs: An Overview of IEEE 802.11ac and 802.11ad. *ACM SIGMOBILE Mobile Computing and Communications Review*, 15(3), Jul 2011.

[10] E. Perahia. IEEE 802.11n Development: History, Process, and Technology. *IEEE Communications*, 46(7), Jul 2008.

[11] K. Lakshminarayanan et al. Bandwidth Estimation in Broadband Access Networks. In *Proc. ACM IMC*, 2004.

[12] A. Johnsson, B. Melander, and M. Bjorkman. *Bandwidth Measurement In Wireless Networks*. Springer, 2006.

[13] M. Li, M. Claypool, and R. Kinicki. WBest: A Bandwidth Estimation Tool for IEEE 802.11 Wireless Networks. In *Proc. IEEE Conference on Local Computer Networks (LCN)*, 2008.

[14] M. Bredel and M. Fidler. A Measurement Study of Bandwidth Estimation in IEEE 802.11g Wireless LANs Using the DCF. In *Proc. IFIP Networking Conference*, 2008.

[15] D. Gupta et al. Experimental Comparison of Bandwidth Estimation Tools for Wireless Mesh Networks. In *Proc. IEEE INFOCOM*, 2009.

[16] M. Portoles-Comeras et al. Impact of Transient CSMA/CA Access Delays on Active Bandwidth Measurements. In *Proc. ACM IMC*, 2009.

[17] D. Koutsonikolas and Y. C. Hu. On the Feasibility of Bandwidth Estimation in Wireless Access Networks. *Wireless Networks*, 17(6), Aug 2011.

[18] V. Ribeiro et al. pathChirp: Efficient Available Bandwidth Estimation for Network Paths. In *PAM*, 2003.

[19] R. Prasad et al. Bandwidth Estimation: Metrics, Measurement Techniques, and Tools. *IEEE Network*, 17(6), 2003.

[20] Q. Wang, A. Feng, and J. Cao. Available Bandwidth Estimation in IEEE 802.11 Ad Hoc Networks. In *Proc. 9th International Conference on Hybrid Intelligent Systems (HIS)*, 2009.

[21] H. Zhao, E. Garcia-Palacios, J. Wei, and Y. Xi. Accurate Available Bandwidth Estimation in IEEE 802.11-based Ad Hoc Networks. *Computer Communications*, 32(6), 2009.

[22] S. Tursunova, K. Inoyatov, and Y-T. Kim. Cognitive Passive Estimation of Available Bandwidth (cPEAB) in Overlapped IEEE 802.11 WiFi WLANs. In *Proc. IEEE Network Operations and Management Symposium (NOMS)*, 2010.

[23] J. Strauss, D. Katabi, and F. Kaashoek. A Measurement Study of Available Bandwidth Estimation Tools. In *Proc. ACM IMC*, 2003.

[24] M. Jain and C. Dovrolis. End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput. *IEEE/ACM Transactions on Networking*, 11(4), 2003.

[25] C. Castellanos et al. Comparison of Available Bandwidth Estimation Techniques in Packet-Switched Mobile Networks. In *IEEE PIMRC*, 2006.

[26] B. Melander, M. Bjorkman, and P. Gunningberg. A New End-to-End Probing and Analysis Method for Estimating Bandwidth Bottlenecks. In *Proc. IEEE GLOBECOM*, 2000.

[27] E. Goldoni and M. Schivi. End-to-End Available Bandwidth Estimation Tools, an Experimental Comparison. In *Proc. Traffic Monitoring and Analysis (TMA) Workshop*, 2010.

[28] C. Beaumont. Virgin to roll out 100mbps broadband service. http://www.telegraph.co.uk/technology/broadband/7314442/Virgin-to-roll-out-100mbps-broadband-service.html.

[29] Iperf. http://iperf.sourceforge.net.

[30] AirPcap Nx USB dongle. http://www.metageek.net/products/airpcap/.

[31] S. Kang and D. Loguinov. IMR-Pathload: Robust Available Bandwidth Estimation Under End-Host Interrupt Delay. In *PAM*, 2008.