



HAL
open science

Home Network or Access Link? Locating Last-Mile Downstream Throughput Bottlenecks

Srikanth Sundaresan, Nick Feamster, Renata Teixeira

► **To cite this version:**

Srikanth Sundaresan, Nick Feamster, Renata Teixeira. Home Network or Access Link? Locating Last-Mile Downstream Throughput Bottlenecks. PAM 2016 - Passive and Active Measurement Conference, Mar 2016, Heraklion, Greece. pp.111-123, 10.1007/978-3-319-30505-9_9 . hal-01294924

HAL Id: hal-01294924

<https://inria.hal.science/hal-01294924>

Submitted on 30 Mar 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Home Network or Access Link? Locating Last-mile Downstream Throughput Bottlenecks

Srikanth Sundaresan¹, Nick Feamster², and Renata Teixeira³

¹ ICSI

² Princeton

³ Inria

Abstract. As home networks see increasingly faster downstream throughput speeds, a natural question is whether users are benefiting from these faster speeds or simply facing performance bottlenecks in their own home networks. In this paper, we ask whether downstream throughput bottlenecks occur more frequently in their home networks or in their access ISPs. We identify lightweight metrics that can accurately identify whether a throughput bottleneck lies inside or outside a user’s home network and develop a detection algorithm that locates these bottlenecks. We validate this algorithm in controlled settings and report on two deployments, one of which included 2,652 homes across the United States. We find that wireless bottlenecks are more common than access-link bottlenecks—particularly for home networks with downstream throughput greater than 20 Mbps, where access-link bottlenecks are relatively rare.

Keywords: Bottleneck location, wireless bottlenecks, last-mile, passive measurements

1 Introduction

Many countries around the world are investing heavily to increase the speeds of access network infrastructure. As the downstream throughput of access links increases, a natural question is whether users are reaping the benefits of these faster speeds. The downstream throughput they are experiencing may be limited by other factors, such as their home wireless networks, which may face performance problems due to a variety of factors (*e.g.*, a poorly placed access link, interference from competing networks or even devices on the same network). In light of these trends, we study a simple question: *Do users tend to see downstream throughput bottlenecks more often in their access ISPs or in their home wireless networks?* To study this question, we design and implement an algorithm, *HoA (Home or Access)*, that can accurately locate these downstream bottlenecks on commodity home routers. We deploy HoA in 2,652 home networks in the United States and characterize the throughput bottlenecks that we observe across this deployment.

Despite the importance and widespread interest in answering this question, both data and conclusions have proved to be elusive. Although throughput analysis and wireless diagnosis tools exist, each existing tool has some limitation that makes it unsuitable for studying this question—typically, these tools require performing measurements

from multiple vantage points (which are hard to convince users to install in their home networks), performing active measurements (which can affect the performance of the wireless network), or custom hardware (which can hamper widespread deployment). (Section 4 explains how our work relates to previous throughput detection and analysis tools and why existing tools do not apply in our setting.) In contrast, we seek to develop a passive network measurement tool that can run from a low-cost, commodity home network router. This choice necessarily limits the extent of the data that we can collect (and, as a result, the conclusions that we can draw), but it also affords a relatively large-scale deployment. HoA’s simplicity allowed us to implement it on a commodity Netgear router for two in-home deployments: A deployment of BISmark routers across 64 homes and 15 countries; and another deployment that was sponsored by the US Federal Communications Commission (FCC) and included 2,652 homes across the United States. These deployments allowed us to conduct a first-of-its-kind large-scale study of last-mile bottlenecks. Section 2.6 describes the deployments in more detail.

Realizing HoA required tackling several challenges. First, we needed to properly isolate performance problems in the home network versus outside of the home; capturing measurements at the home router offers a convenient solution to this challenge, since it lies between these two parts of the network. Next, we had to identify and validate metrics that were lightweight enough to capture on a low-cost home router, yet sufficient to accurately locate downstream throughput bottlenecks. We also wanted to use performance metrics from passive network traffic capture, to avoid introducing conditions that might either alter the state of the wireless network or disrupt network performance for home network users. Ultimately, we identified two features—the coefficient of variation of packet inter-arrival time and the round-trip time on the wireless LAN—that can be measured passively, are lightweight enough to be deployed on a commodity home gateway, and can identify last-mile bottlenecks in many circumstances. Section 2 incorporates these metrics into a complete identification algorithm.

We offer two important contributions: (1) the design of HoA, a lightweight tool that both accurately detects home access link and wireless network bottlenecks; (2) a detailed characterization of the nature and extent of throughput bottlenecks that commonly arise in many home networks using data from a large-scale prototype deployment of HoA in home routers. We do not determine *why* a particular bottleneck exists (*e.g.*, it cannot determine whether a wireless problem results from poor device placement, non-WiFi interference, or other causes), but rather only *where* the problem exists, to the granularity of whether the problem is inside or outside the home. Our study yields the following important findings:

- Access link bottlenecks rarely occur in home networks where downstream access throughput exceeds 20 Mbps. Rather, in these cases, throughput bottlenecks are often introduced by the home wireless network.
- Access link bottlenecks only tend to be common for users whose downstream access throughput is less than 10 Mbps.
- In homes with multiple devices where we detect a wireless bottleneck, it is equally likely that only a single device experiences the wireless bottleneck as it is that all devices in the home experience the bottleneck simultaneously.

Our results suggest that it is worth spending effort to improve home wireless network performance, in addition to the extensive attempts to optimize performance in other parts of the network and end hosts.

2 HoA: Design, Implementation, and Deployments

We describe the design, implementation, and deployment of HoA.

2.1 Design Choices

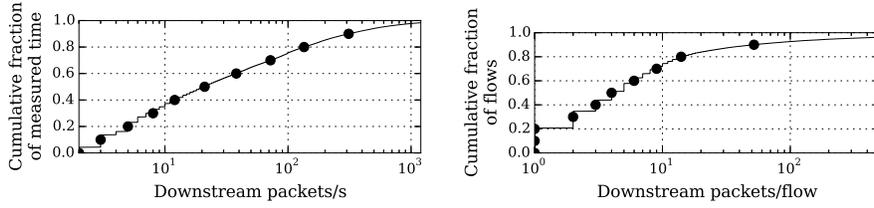
Our first design choice was to perform measurements from the home access point. Locating bottlenecks at the last mile becomes easier with a vantage point inside the home network. Although vantage points in the access ISP (such as in the DSLAM for a DSL ISP) can see all the home traffic, these locations outside the home obscure metrics that can provide important clues about whether the home wireless network is introducing a bottleneck. Inside the home, we can either instrument end-hosts or the access point itself. Client devices can observe wireless properties from their own traffic but may not be able to observe traffic properties of other clients. A device also cannot determine characteristics of the access link. End-host tools such as T-RAT [22] can monitor TCP properties such as congestion window or duplicate ACKs to identify the causes of throughput bottlenecks but cannot isolate the location of congestion.

Our second design choice was to use passive traffic measurements. Although active probing may yield useful information about the state of the network, it also carries potential drawbacks. It risks introducing extra load on the network, thereby affecting the conditions that we are trying to measure; it may also disrupt the users who are hosting our measurement devices. Thus, we rely on passive measurements of *in situ* user traffic as the main source of information for detecting performance bottlenecks. We aim to do so without custom wireless drivers or anything that could adversely affect the performance of the networks we are measuring, so we look for features at the IP layer that can indicate performance problems. Possible metrics thus include flow timings and sizes, packet timings and sizes, and information that we can retrieve from TCP headers. We briefly discuss our choices.

2.2 Network Metrics

Packet arrival timings and TCP RTT are promising metrics particularly because our vantage point at the access point allows us to separately compute these metrics for the WAN and LAN portions of the end-to-end path, potentially allowing us to disambiguate problems that occur on either side of the access point.

Packet Interarrival Time. We exploit an observation that is common to many bottleneck links: *packets traversing a bottlenecked link experience buffering immediately upstream of the link; as a result, they experience smoothed arrival patterns downstream of the bottleneck link.* To capture this effect, we use the coefficient of variation of packet interarrival times, c_v , which is the standard deviation of packet interarrival time divided



(a) Cumulative distribution of the number of packets per second in FCC deployment. (b) Cumulative distribution of the number of packets per flow in FCC deployment.

Fig. 1: Properties of test samples from the FCC deployment

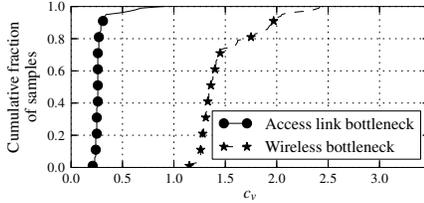


Fig. 2: Coefficient of variation of packet interarrival times. When the access link is a downstream throughput bottleneck, packet arrivals are smooth (i.e., the variance on packet interarrival time is lower).

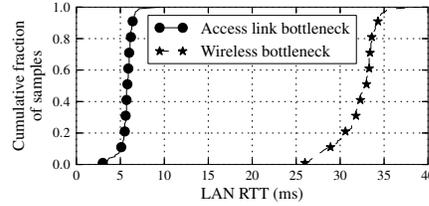


Fig. 3: TCP RTT between client and access point. When the wireless link is the throughput bottleneck, the TCP RTT between the device and the access point is significantly higher.

by the mean packet interarrival time. In our example, when the access link is the bottleneck $c_v = 0.05$; whereas when the wireless is the bottleneck c_v is 0.88. In Figure 2, the “access link bottleneck” curve presents the distribution of c_v for 100 experiments where we introduced a bottleneck at the access link; and the “wireless bottleneck” curve for 100 experiments where the bottleneck was on the wireless. There is no overlap between the two curves: c_v is lower when the access link is the bottleneck versus when it is not.

Wireless Round-Trip Time. The second effect is that devices in home networks are only one hop away from the access point, so the baseline latency between the access point and the device should be a few milliseconds (as we measured in our controlled experiments). We observe that the delays caused by buffering in the wireless network (i.e., those caused by throughput bottlenecks) are significantly higher. We measure this effect by capturing the TCP RTT, τ , between the device and the access point. Figure 3 presents the LAN TCP RTT (the RTT over TCP between the access point and a device in the home network) for two downstream throughput bottleneck scenarios: an access network bottleneck and a wireless bottleneck. In both experiments, we established (through repeated experiments) the wireless network capacity to be about 40 Mbps. In the first case, the access link is 30 Mbps, so it is always the bottleneck. In the second case, the access link is 70 Mbps so that the wireless network becomes the bottleneck. When the

access link is the bottleneck, the RTT is about 5 ms. In contrast, when the wireless is the bottleneck, packet buffering at the head of the wireless link (*i.e.*, the access point) increases RTTs to about 25–35 ms.

2.3 Detection Algorithm

For each device, d , we use two independent detectors. One detector uses a decision rule that determines whether an access-link bottleneck event, B , occurs, given a particular observed value of c_v . The other detector uses a decision rule that determines whether a wireless bottleneck event, W , occurs given a particular observed value of τ_d . We first compute likelihood functions $f(c_v|B)$ and $f(c_v|\overline{B})$ in a controlled setting, where we use our ability to control the throughput of the upstream link to introduce a bottleneck on the access link. We then define our decision rule in terms of the likelihood ratio:

$$A(c_v = v) = \frac{f(c_v = v|B)}{f(c_v = v|\overline{B})}$$

where v is the measured coefficient of variation of packet interarrival time for packets over the observation window. When A is greater than some threshold γ , the detector says that the access link is the bottleneck (*i.e.*, it is more likely than not, given the observation of $c_v = v$, that the prior is the event B). We can tune the detector by varying the value of γ ; higher values will result in higher detection rates, but also higher false positive rates. We use a similar approach for W . The next section presents our choices of threshold.

We can only perform bottleneck detection if the network is sending enough traffic. We set a minimum number of packets per second, T_{pps} , and a minimum number of packets per flow, T_{pff} , for running HoA. Figure 1 shows the distribution of the number of packets per second and packets per flow observed across homes in the FCC deployment. In approximately 40% of measured one-second intervals, we observe packet rates of less than 10 packets per second. We also tested T_{pps} values of 50, 100, and 150 packets per second, and T_{pff} values of 25, 50, and 75 packets per flow on real-world deployment data; none of these settings changed our conclusions.

2.4 Calibration

We built a testbed to run controlled experiments to calibrate detection thresholds. The testbed has an access point, its LAN, a network traffic shaper upstream of the access point, a well-provisioned university network, and servers in the university network. The access point is a Netgear WNDR3800 router running OpenWrt. To change the downstream throughput of the emulated access link, we use `tc` and `netem` on a second WNDR3800 router running OpenWrt. We run our tests against servers in the same well-provisioned university network to avoid potential wide-area bottlenecks. We run two sets of experiments using the testbed.

We use a traffic shaper to shape the link to different throughput levels while keeping the wireless link constant. In this case, identifying the ground truth is straightforward, as we know the capacities of both the wireless link and the shaped access link. We

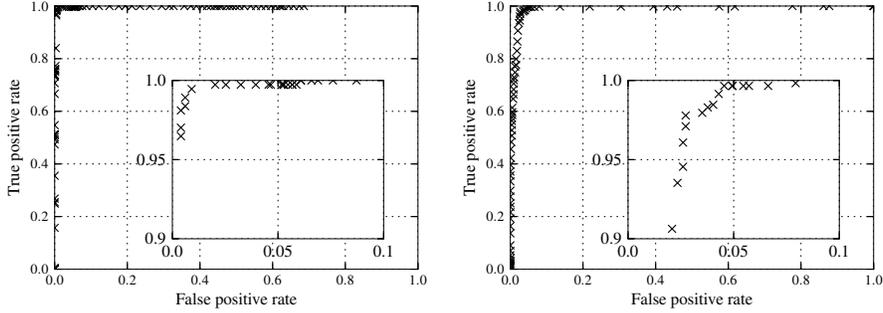


Fig. 4: Receiver operating characteristic for access link bottleneck detection using the co-wireless bottleneck detection using the TCP efficient of variation of packet interarrival time. **Fig. 5:** Receiver operating characteristic for access link bottleneck detection using the co-wireless bottleneck detection using the RTT between the access point and the client.

use 802.11a and 802.11n for the wireless link with respective capacities of 21 Mbps and 80 Mbps over TCP. We generate 1,356 experiments with 11 different emulated access links, with capacities varying from 3 Mbps to more than 100 Mbps. To introduce wireless bottlenecks, we conduct two sets of experiments. (1) Reduce capacity by degrading channel quality: we do this by positioning the host at different distances from the access point, and with multiple obstructions, and also transient problems by human activity. (2) Reduce the available capacity of the channel by creating contention with an interfering host that sends constant UDP traffic, with the interfering host close to the access point. For each setting, we run a TCP throughput test using `iperf`. To minimize interference that we do not introduce ourselves, we use the 5 GHz spectrum, which is less congested than the 2.4 GHz range in our testbed. In our repeated controlled experiments, we found that the wireless channel in our testbed delivers a TCP throughput of about 80 Mbps on 802.11n. We performed 1,356 experiments over many operating conditions.

Because there can only be one throughput bottleneck on an end-to-end path, by definition, the detectors should never detect bottlenecks simultaneously. Using the thresholds that we computed for each detector—as we describe for each case below—simultaneous detection occurs only 2% of all time intervals, typically in cases where the throughput values for the home wireless network and the access link were similar.

Packet Interarrival Time (T_{cv}) We use the results from the controlled experiments described above to compute the likelihood functions $f(c_v|B)$ and $f(c_v|\bar{B})$ to determine the detection threshold T_{cv} . We first evaluate the detection accuracy of the algorithm for different values of T_{cv} . Figure 4 shows the receiver operating characteristic for this detector. When T_{cv} is low (close to zero), the detector will always determine that the access link is not the bottleneck; when T_{cv} is high (close to one), the detector will always identify the access link as the bottleneck. Our results indicate that detection accuracy remains high for a wide range of threshold settings for T_{cv} , particularly between 0.7 and 0.9. Detection accuracy is very high in this range, with a true positive rate more than 95% and a false positive rate less than 5%. The range of good thresholds reinforces our

confidence of its robustness as a detection metric. We use a threshold $T_{cv} = 0.8$, which offers the best tradeoff between the true positive and false positive rates, to declare the access link the bottleneck.

Wireless Round-Trip Time (T_τ) We calibrate the thresholds for the likelihood functions $f(\tau_d|W)$ and $f(\tau_d|\overline{W})$ using a similar method. We choose a threshold $T_\tau = 15$ ms, which yields a detection rate of 95% and a low false positive rate of less than 5%. Similar to the T_{cv} parameter, T_τ is also robust; we get similarly high true positive rates and low false positive rates for values ranging from 12–17 ms. Higher LAN latencies in the wireless network can result from other wireless problems that may manifest as retransmissions or backoffs. We observe empirically that these wireless issues introduce up to 8–12 ms of delay, whereas delays caused by wireless throughput bottlenecks introduce more than 15 ms of extra delay, thresholds which yield a high detection and low false positive rate in our experiments.

2.5 Limitations

HoA has several limitations. First, because it relies on passive traffic analysis, the link must carry enough traffic to enable analysis. Section 2.3 how we determine minimum thresholds for detection, which are heuristics. Second, constant bit rate traffic could in some cases yield a low cv , thus causing HoA to mistakenly detect a throughput bottleneck on the access link; such cases may need to rely on other detection methods. With respect to bottlenecks, HoA cannot identify the root cause of bottlenecks, and it cannot identify bottlenecks far from the last mile, such as peering or server-side bottlenecks. HoA can only locate throughput bottlenecks where the link is work-conserving; because wireless links violate this assumption, HoA cannot detect upstream throughput bottlenecks. Additionally, detection thresholds may be sensitive to certain settings and configurations: T_τ may depend on the wireless driver and hardware; in cable access networks, T_{cv} may depend on the channel bonding configuration of the DOCSIS modem. The calibration methods from Section 2.4 may help determine the appropriate thresholds in various settings. Finally, to reduce CPU load, HoA collects data periodically, which does not allow us to capture aspects of the network that vary over small timescales.

2.6 Deployments

Table 1 summarizes our two deployments, which we briefly describe below.

BISmark Deployment. We deployed HoA on Netgear’s WNDR3700/3800, which has an Atheros chipset with a 450 MHz processor, one 802.11bgn radio, and one 802.11an radio. The 3800 has 128 Mbytes of RAM, and the 3700 has 64 Mbytes of RAM. The devices run OpenWrt, with the ath9k wireless driver. The driver uses the Minstrel rate adaptation algorithm, with the default setting to a maximum bitrate of 130 Mbps. Every 5 minutes, HoA collects packet traces from the WAN port for 15 seconds and extracts timestamps and per-flow RTTs on either side of the access point, as well as the number

	BISmark	FCC
Homes	64	2,652
Location	15 Countries	United States
Duration	March 6–April 6, 2013	November 4–5, 2014
Tests	52,252	73,193

Table 1: Deployments of HoA, including locations and study durations. In addition to the larger FCC deployment, we also performed a pilot deployment of HoA on 100 homes in the FCC deployment from August 24–30, 2014.

of packets for each connection using `tcptrace` [21]. `tcptrace` tracks packets and the corresponding ACKs to compute the RTTs.

FCC Deployment. We use the FCC’s deployment of Netgear WNR3500L, which has a Broadcom chipset and a 480 MHz processor, one 802.11bgn radio, and 64 Mbytes of RAM. The devices run a custom Netgear firmware based on OpenWRT. The resource constraints of the WNR3500L required two changes to our implementation. First, we imposed a packet limit and a time limit for every trace collection iteration. The collection runs for 10 seconds or until it has collected 10,000 packets, whichever comes first. We discard any trace for which the packet filters dropped at least 5% of packets from our analysis. Additionally, due to resource constraints, we do not perform any processing on the device, except for anonymization. Instead, we offload the packet header traces for offline analysis. To avoid conflicts with FCC’s Measuring Broadband America program, we could only perform our measurements three times per hour.

3 Results

This section explores our findings: (1) In home networks where downstream throughput exceeds 20 Mbps, the home wireless network is the primary cause of throughput bottlenecks. (2) Access link bottlenecks are prevalent in home networks where the downstream throughput is less than 10 Mbps. (3) In homes where HoA detects a wireless throughput bottleneck, it is about equally likely that the wireless throughput bottleneck is isolated to a single device or observed across all devices.

3.1 Prevalence of Last-mile Bottlenecks

In this section, we explore the prevalence of downstream throughput bottlenecks in access links versus home wireless networks using HoA. Specifically, we study the fraction of tests for which HoA identifies downstream throughput bottlenecks, and to what extent these bottlenecks are caused by the access link versus the home wireless network.

We perform more than 50,000 tests over a wide range operating conditions in the field. HoA identifies downstream throughput bottlenecks in 55% of tests in the BISmark deployment and 47% of tests in the FCC deployment. When HoA does not detect a bottleneck, the underlying cause may be low demand or bottlenecks being elsewhere in the network (*e.g.*, at a peering point). As expected, homes with access-link throughput

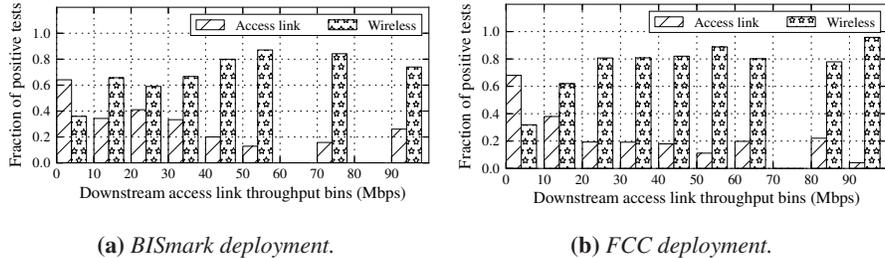


Fig. 6: Prevalence of access link and wireless bottlenecks home networks the two deployments deployment. When downstream access-link throughput exceeds about 20 Mbps, only about 20% of last-mile bottlenecks occur on the access link.

less than 10 Mbps experience the largest fraction of throughput bottlenecks; 55% of tests detect a bottleneck. The fraction of tests where HoA detects a bottleneck, however, remains close to 40% even for homes with access-link throughput above 90 Mbps. In the rest of this section, we further characterize the tests where HoA detects a downstream throughput bottleneck.

Figure 6a plots the fraction of downstream throughput bottlenecks in the BISmark deployment that are located either in the access link or in the home wireless network. We group home networks into bins of 10 Mbps according to the measured downstream throughput of their access links. The results show that many throughput bottlenecks in the BISmark deployment are due to the wireless network. Our analysis of the bottlenecks per home in the BISmark deployment shows that the fraction of wireless bottlenecks varies significantly across homes even for homes with similar access-link throughput. For example, homes with access-link throughput less than 20 Mbps had wireless bottlenecks in between 3–58% of downstream throughput bottlenecks, and 11–83% of downstream throughput bottlenecks. By default, we configured these home routers to use 802.11n, which can support significantly higher rates. The default 802.11n configuration supports frame rates of up to 130 Mbps (we observed about 85–90 Mbps over TCP), while 802.11g supports only framerates up to 54 Mbps. The fact that these networks are experiencing throughput bottlenecks suggests persistent problems with home wireless network deployments in practice.

Figure 6b shows the same results for the FCC deployment. First, access-link bottlenecks only occur frequently for home networks with downstream access throughput less than 20 Mbps. Homes with access throughput less than 10 Mbps experience access-link bottlenecks in about 66% of cases; however this fraction drops rapidly as access throughput increases: for homes with access throughput between 10 and 20 Mbps about 40% of downstream throughput bottlenecks are due to access-link bottlenecks, whereas for homes with access links exceeding 20 Mbps access-link bottlenecks explain only about 20% of downstream throughput bottlenecks. Conversely, wireless throughput bottlenecks become more prevalent in homes with higher access throughput: 33% of downstream throughput bottlenecks for homes with throughput less than 10 Mbps are due to wireless bottlenecks; 40% of the bottlenecks are in the wireless network for homes with

10–20 Mbps access throughput; and, nearly 80% of the bottlenecks are in the wireless network when access throughput exceeds 20 Mbps. That wireless throughput bottlenecks occur even for access links with such low speeds is surprising: the FCC access points support 802.11n, with default frame bitrates of up to 130 Mbps and a maximum frame bitrate of 300 Mbps. Some users had configured their routers to 802.11g, and those users did experience lower throughput. Yet, 802.11g comprised only 10% of all tests, so most of the problems that we observed occurred even with 802.11n.

In about 8% of downstream throughput bottlenecks in homes with access-link throughput less than 10 Mbps, HoA indicates that both the wireless network and the access link are introducing throughput bottlenecks. In principle, this should not occur as, by definition, there can be only one bottleneck. The prevalence of this result for primarily low-throughput access links suggests that in these cases, at least one device in the home network may be experiencing poor wireless conditions in conjunction with an access-link bottleneck.

3.2 Wireless Bottlenecks Within a Home

The previous section demonstrated that wireless bottlenecks are common; in cases where wireless bottlenecks exist, at least one device in the home experiences a wireless throughput bottleneck during the tests. For about 75% of tests when HoA detects a wireless bottleneck, we only observe traffic for one device in the home. For the remaining 25% of tests with a wireless bottleneck, we investigate whether the active devices experience a downstream throughput bottleneck in the wireless network simultaneously. Simultaneous throughput bottlenecks in the wireless network to independent devices might indicate a more systemic problem (*e.g.*, pervasive interference, poor signal from the access point, contention), whereas isolated throughput bottlenecks are more likely to indicate a problem with a particular device. About half of the cases we observed involve throughput bottlenecks that are isolated to a single device; in another 45% of cases, all of the devices in the home simultaneously experience a throughput bottleneck.

4 Related Work

HoA draws inspiration from several previous diagnosis techniques. Zhang *et al.* developed T-RAT [22] to analyze TCP performance. T-RAT estimates TCP parameters such as maximum segment size, round-trip time, and loss to understand flow behavior. Katabi *et al.* [11], used entropy in packet interarrival time to estimate shared bottlenecks. Biaz *et al.* [3] used packet interarrival times for distinguishing between different kinds of losses. HoA is similar to some of the approaches used in these papers (*e.g.*, it uses packet interarrival time as input to a detector for access link bottlenecks), but we tailor our approach so that it only relies on data that can be easily collected from a home router. Previous work has studied broadband access performance [4, 8, 9, 20]. In particular, Sundaresan *et al.* [20] study residential access performance from home routers (also using the FCC Broadband America dataset). There have also been many previous approaches to diagnosing wireless networks. One approach is to deploy passive traffic monitors throughout the network to diagnose wireless pathologies [1, 2, 6, 15, 16]

or to study wireless performance [14]. Kanuparth *et al.* [10] developed a tool to detect common wireless pathologies (such as low signal-to-noise ratio, congestion, and hidden terminals) by using both active probes and an additional passive monitor deployed within the network. Kim *et al.* [12] analyze wireless metrics such as frame bitrates, frame ACKs and retransmission rates to identify root causes of wireless performance problems. Other approaches have monitored wireless networks with custom hardware [5, 13, 16–18]. Unfortunately, it is difficult to deploy multiple monitoring points or custom hardware in many home networks, since it requires deploying equipment beyond what a normal user is typically willing to install or have installed in their home. Other efforts have characterized home networks in terms of connected devices and usage [7, 19]. None of these studies, however, have studied how often the home network constraint downstream throughput.

5 Conclusion

To identify performance bottlenecks in home networks, we developed an algorithm and tool, HoA, that passively observes traffic flows between the home network and the access network to determine the location of last-mile downstream throughput bottlenecks. Our prototype deployment of HoA in 2,652 home networks shed new light on the prevalence of downstream throughput bottlenecks in both home networks and access networks. We find that when the downstream throughput of a user’s access link exceeds about 20 Mbps, a high fraction of throughput bottlenecks are caused by the user’s home wireless network. This finding is significant in light of recent proposed regulations to change the definition of broadband Internet access to increasingly higher speeds. Our study opens several avenues for future work. First, we need methods to identify root causes that explain *why* various wireless performance problems exist in addition to where they are. Second, a follow-up to HoA could attribute problems that home network users experience to a more complete and more specific set of causes.

Acknowledgments We thank the FCC and SamKnows for helping us develop and deploy HoA in the Measuring Broadband America (MBA) platform. We also acknowledge the participants of the MBA platform. We would like to thank our shepherd, Mahesh K. Marina, and the reviewers for their helpful comments. This work was supported by NSF awards CNS-1535796, CNS-1539906, and CNS-1213157, and the European Communitys Seventh Framework Programme (FP7/2007-2013) no. 611001 (User-Centric Networking).

References

1. Adya, A., Bahl, P., Chandra, R., Qiu, L.: Architecture and Techniques For Diagnosing Faults in IEEE 802.11 Infrastructure Networks. In: *MobiCom*. pp. 30–44. Philadelphia, PA (2004)
2. Ahmed, N., Ismail, U., Keshav, S., Papagiannaki, K.: Online Estimation of RF Interference. In: *ACM CoNEXT*. Madrid, Spain (Dec 2008)
3. Biaz, S., Vaidya, N.H.: Discriminating congestion losses from wireless losses using inter-arrival times at the receiver. In: *IEEE Symposium on Application - Specific Systems and Software Engineering and Technology (ASSET)*. Washington, DC, USA (1999)
4. Canadi, I., Barford, P., Sommers, J.: Revisiting Broadband Performance. In: *ACM SIGCOMM Internet Measurement Conference (IMC)* (Oct 2012)

5. Cheng, Y., Bellardo, J., Benko, P., Snoeren, A.C., Voelker, G.M., Savage, S.: Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis. In: Proc. ACM SIGCOMM. Pisa, Italy (Aug 2006)
6. Cheng, Y.C., Afanasyev, M., Verkaik, P., Benkö, P., Chiang, J., Snoeren, A.C., Savage, S., Voelker, G.M.: Automating cross-layer diagnosis of enterprise wireless networks. SIGCOMM Comput. Commun. Rev. 37(4), 25–36 (Aug 2007)
7. Cioccio, L.D., Teixeira, R., Rosenberg, C.: Measuring home networks with homenet profiler. In: Proc. of Passive and Active Measurement Conference (2013)
8. Croce, D., En-Najjary, T., Urvoy-Keller, G., Biersack, E.: Capacity Estimation of ADSL links. In: Proc. CoNEXT (Dec 2008)
9. Dischinger, M., Haerberlen, A., Gummadi, K.P., Saroiu, S.: Characterizing residential broadband networks. In: Proc. ACM SIGCOMM Internet Measurement Conference. San Diego, CA, USA (Oct 2007)
10. Kanuparth, P., Dovrolis, C., Papagiannaki, K., Seshan, S., Steenkiste, P.: Can user-level probing detect and diagnose common home-wlan pathologies. SIGCOMM Comput. Commun. Rev. 42(1), 7–15 (Jan 2012)
11. Katabi, D., Blake, C.: Inferring congestion sharing and path characteristics from packet inter-arrival times. Tech. Rep. MIT-LCS-TR-828, Massachusetts Institute of Technology (2002)
12. Kim, K.H., Nam, H., Schulzrinne, H.: WiSlow: A Wi-Fi network performance troubleshooting tool for end users. In: IEEE INFOCOM. pp. 862–870 (2014)
13. Lakshminarayanan, K., Sapra, S., Seshan, S., Steenkiste, P.: Rfdump: an architecture for monitoring the wireless ether. In: Proceedings of the 5th international conference on Emerging networking experiments and technologies. pp. 253–264. CoNEXT '09 (2009)
14. Mahajan, R., Rodrig, M., Wetherall, D., Zahorjan, J.: Analyzing the mac-level behavior of wireless networks in the wild. pp. 75–86. SIGCOMM '06 (2006)
15. Niculescu, D.: Interference map for 802.11 networks. In: ACM SIGCOMM Internet Measurement Conference. pp. 339–350. San Diego, California, USA (Oct 2007)
16. Rayanchu, S., Mishra, A., Agrawal, D., Saha, S., Banerjee, S.: Diagnosing wireless packet losses in 802.11: Separating collision from weak signal. In: INFOCOM 2008. The 27th Conference on Computer Communications. IEEE. pp. 735–743 (april 2008)
17. Rayanchu, S., Patro, A., Banerjee, S.: Catching whales and minnows using wifinet: deconstructing non-wifi interference using wifi hardware. In: USENIX NSDI. San Jose, CA
18. Rayanchu, S., Patro, A., Banerjee, S.: Airshark: detecting non-wifi rf devices using commodity wifi hardware. In: ACM SIGCOMM Internet Measurement Conference. pp. 137–154. Berlin, Germany (2011)
19. Sanchez, M.A., Otto, J.S., Bischof, Z.S., Bustamante, F.E.: Trying broadband characterization at home. In: Proc. of Passive and Active Measurement Conference (2013)
20. Sundaresan, S., de Donato, W., Feamster, N., Teixeira, R., Crawford, S., Pescapè, A.: Broadband Internet Performance: A View From the Gateway. In: ACM SIGCOMM. Toronto, Ontario, Canada (Aug 2011)
21. tcptrace: A TCP Connection Analysis Tool, <http://irg.cs.ohiou.edu/software/tcptrace/>
22. Zhang, Y., Breslau, L., Paxson, V., Shenker, S.: On the characteristics and origins of internet flow rates. In: Proc. ACM SIGCOMM. Pittsburgh, PA (Aug 2002)