

Crowdsourcing ISP Characterization to The Network Edge

Zachary S. Bischof
Northwestern University
zbischof@eecs.northwestern.edu

John P. Rula
Northwestern University
jrula@eecs.northwestern.edu

John S. Otto
Northwestern University
jotto@eecs.northwestern.edu

David R. Choffnes
University of Washington
choffnes@cs.washington.edu

Mario A. Sánchez
Northwestern University
msanchez@eecs.northwestern.edu

Fabián E. Bustamante
Northwestern University
fabianb@eecs.northwestern.edu

ABSTRACT

Evaluating and characterizing Internet Service Providers (ISPs) is critical to subscribers shopping for alternative ISPs, companies providing reliable Internet services, and governments surveying the coverage of broadband services to its citizens. Ideally, ISP characterization should be done at scale, continuously, and from end users. While there has been significant progress toward this end, current approaches exhibit apparently unavoidable tradeoffs between coverage, continuous monitoring and capturing user-perceived performance.

In this paper, we argue that network-intensive applications running on end systems avoid these tradeoffs, thereby offering an ideal platform for ISP characterization. Based on data collected from 500,000 peer-to-peer BitTorrent users across 3,150 networks, together with the reported results from the U.K. Ofcom/SamKnows studies, we show the feasibility of this approach to characterize the service that subscribers can expect from a particular ISP. We discuss remaining research challenges and design requirements for a solution that enables efficient and accurate ISP characterization at an Internet scale.

Categories and Subject Descriptors

C.2.3 [Communication Systems Organization]: Computer Communication Networks—*Network Operations*; C.2.5 [Communication Networks]: Local and Wide-Area Networks—*Internet*; C.4 [Performance of Systems]: Measurement techniques

General Terms

Experimentation, Performance, Measurement

Keywords

Broadband access networks, ISP, Characterization

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

W-MUST'11, August 19, 2011, Toronto, Ontario, Canada.
Copyright 2011 ACM 978-1-4503-0800-7/11/08 ...\$10.00.

1. INTRODUCTION

Evaluating and characterizing ISPs is critical to subscribers shopping for alternative ISPs, companies providing reliable Internet services (e.g., customers of Keynote [10] and IneoQuest¹), and governments surveying the availability of high-speed Internet services to its citizens [18, 19]. While ISPs have increased broadband bandwidth and expanded coverage in the past decade, our ability to measure the performance of these networks has not kept pace [2].

Ideally, ISP characterization should be done (i) at scale, to capture the diversity of available providers and their services; (ii) continuously, to witness dynamic changes due to network management policies (e.g. traffic shaping or oversubscribed networks) and unscheduled events (e.g. service interruptions); and (iii) from end users, to guarantee accurate characterization of service.

Given the important business and policy [9] implications of ISP characterization, recent years have brought a variety of approaches to profiling edge network services ranging from Web-based testing platforms [11, 16] to dedicated active measurement devices located inside PoPs or home networks [17]. While these efforts have started to shed some much needed clarity on network conditions at the edge, they all exhibit apparently unavoidable trade-offs between the diversity of vantage points, the time granularity of monitoring, and the fidelity of user-perceived performance metrics.

We argue that popular network-intensive applications running on end systems avoid these tradeoffs, offering an ideal platform for ISP characterization. In this paper, we present a new approach to characterization that leverages the detailed views of Internet-wide ISP performance offered by these applications – we call this approach *C2E* for *Crowdsourced ISP Characterization at the Network Edge*. By passively monitoring user-generated traffic within these applications, C2E is able to capture the end user's view in a scalable manner. By extending existing applications, C2E has the flexibility and low-barrier to adoption of other software-based models. By combining passive monitoring with dynamically extensible active measurements, C2E can achieve the effectiveness of hardware-based solutions without their associated costs.

This paper presents preliminary results demonstrating the feasibility of C2E through a large-scale study of data gathered from BitTorrent users. We demonstrate that it is possible to identify and isolate offered levels of service for an ISP using passive bandwidth monitoring techniques. Our results indicate that by combining the views of multiple subscribers in the same ISP, one can obtain sufficiently continuous data to evaluate time-of-day effects on performance. Finally, we show that, at scale, such a monitoring approach enables performance comparisons between ISPs, regions, and countries.

¹<http://www.ineoquest.com/>

Approach	Scale	Continuous	End-user
Netalyzr [11]	✓	X	✓
Speedtest [20]	✓	X	✓
YouTube (myspeed) [21]	✓	X	✓
SamKnows [17]	X	✓	✓
Keynote [10]	X	✓	X
Dischinger et al. [8]	✓	✓	X
Croce et al. [7]	✓	✓	X
C2E	✓	✓	✓

Table 1: Existing approaches to ISP characterization.

While our results are encouraging, there remain a large number of research challenges and design requirements for a solution that enables ISP characterization at an Internet scale; we conclude with a brief discussion how we address some of these issues in Sec. 5.

2. ISP CHARACTERIZATION

We argue that ISP characterization should be done at scale, continuously and by end users. It should be done at scale to capture the diversity of available providers and their services. It should be done continuously to witness dynamic changes due to network management policies (e.g. traffic shaping or oversubscribed networks) and unscheduled events (e.g. service interruptions). Last but not least, it should be done by and from end users to guarantee accurate characterization of service.

This paper presents *C2E*, a new approach to ISP characterization that leverages the detailed views of Internet-wide ISP performance offered by network-intensive applications. Popular network-intensive applications such as peer-to-peer (P2P) file sharing systems, IPTV and VoIP offer a nearly ideal platform for the task. The popularity of these applications ensures the necessary coverage and provides a platform that can grow organically with the Internet itself and potentially enabling comparisons between ISPs, regions, and countries [1]. Their typical usage patterns and comparatively long session times mean one could achieve nearly continuous monitoring by combining the views of multiple subscribers in the same ISP. In addition, the user generated traffic of these network-intensive services could be sufficient to capture the end user’s view in an scalable manner avoiding costly active measurements [13]. Equally important, software platforms are generally more flexible and simpler to update than hardware based services. This is key for broadband service characterization as these services continue to evolve at rapid pace, partially driven consumer demand, and the research community continue to develop better ways to characterize them and present results to end users.

There is a growing body of work related to profiling ISP performance, both to improve transparency and to enable evaluations of SLA compliance. Existing approaches can be classified according to their deployment model (software or hardware), monitoring location (data centers, research environments, PoPs, home networks, or edge systems), targeted environment (single or multiple applications, core or edge networks) and measurement approach (active/passive, on-demand/continuous). Table 1 compares some of these key approaches in terms of the diversity of their vantage points, the granularity of their measurements and their ability to capture end-user perceived performance.

Some previous efforts have focused on measuring core networks [12,22] and thus do not address end user performance, which is typically limited by the quality of their last-mile links. To capture this view, Dischinger et al. [8] conducted measurements of cable and DSL last-mile links in North America and the UK, covering 1,894 hosts from 11 different residential ISPs. However, using

active, unsolicited probing of links (measuring TCP and ICMP packet responses) from dedicated servers limits both the coverage and accuracy (e.g., in high-speed networks) of their approach.

To capture the end-user’s perspective, a common approach relies on Web-based user-initiated tests against centralized servers [11, 16, 20]. While the model is appealing for its low barrier to user adoption and ability to capture end-to-end performance seen by users, these systems are limited by infrastructure and bandwidth costs proportional to the number of adopters. Further, they are only able to capture performance at the time when the test is run and are susceptible to white-listing by ISPs.

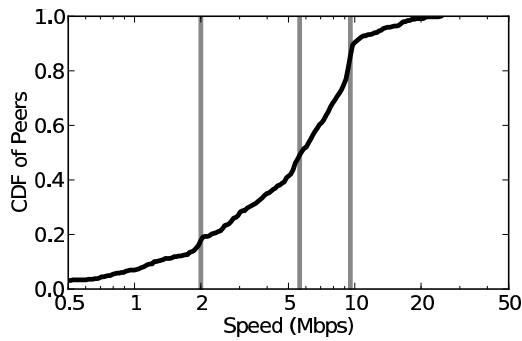
A number of companies provide ISP profiling through active measurements from dedicated devices in points of presence (PoPs) or home networks [17]. For example, companies like Keynote simulate flows for popular applications from inside PoPs, however, devices in PoPs do not cover last-mile links. SamKnows (currently working in collaboration with the U.S. and U.K. governments) conducts large volumes of active measurements from “white boxes” deployed inside subscribers’ home networks. This approach can easily account for home network traffic and enable detailed performance evaluations. However, while there are no fundamental technical limitations to this approach, coverage is limited by the cost to deploy infrastructure in today’s large and rapidly evolving network environments. For instance, SamKnows relies only on a few thousand “white boxes” to capture the large-scale view of the highly connected countries including the U.S. and the U.K.

We still face a number of challenges in realizing C2E. For instance, while many of the network-intensive applications that could host C2E tend to exhibit sufficiently long session times for our purposes, they are nevertheless eventually shutdown. Our approach, however, requires the existence of multiple instrumented clients in each ISP to ensure a sufficient measurement sample size, and to collectively provide a continuous view of performance. Closely related, since we lack control on when (and where) instrumented hosts will become available or what application maybe in use, our monitoring approach must identify and address issues of bias and the impact of uncorrelated externalities (e.g., home network cross traffic [6]) on the measurement results. Finally, to correctly interpret passively gathered performance data, we must understand the context in which the data was acquired. At the level of an application, e.g. in the case of BitTorrent, it is important to distinguish whether download performance is limited by the application being monitored (e.g., slow swarm speed in BitTorrent, or a user-defined bandwidth cap) or by the speed of the access link. At the ISP service level, knowledge of the user’s service level (e.g., maximum download speed) and employed traffic management techniques (e.g., interfering with certain applications’ traffic) impacts the interpretation of measurements.

In the following sections, we discuss some of these challenges and demonstrate the potential for C2E using traces gathered from BitTorrent users. In particular, we explore the feasibility of the proposed approach for ISP characterization using popular metrics – advertised transfer rates and broadband availability. Identifying other useful metrics for characterizing network services as well as alternative hosting applications for C2E is part of our future work.

3. A CASE STUDY: BITTORRENT

In this section, we use data reported by users of a popular BitTorrent extension for Vuze to evaluate the feasibility of the proposed approach and highlight the importance of performing ISP characterization at scale (Sec. 3.1), continuously (Sec. 3.2), and from the edge (Sec. 3.3).



(a) Per-User Max Download Speed

Derived	Advertised	Package
2.0	2.0	“Small” Cable
5.6	8.0	“National” ADSL
9.6	10.0	“Medium”/“Large” Cable
—	20.0	“Extra-Large” Cable

(b) Virgin Media Service Level Comparison

Figure 1: We derive service levels (vertical lines) from the distribution of Virgin Median users’ maximum download rates, and find that they appear to align with the company’s advertised service tiers. The observed tier at 5.6Mbps likely corresponds to the 8.0Mbps tier because ADSL performance is reduced when a user is too far from the exchange.

Dataset. The dataset includes 30 sec. samples of transfer rates, network and geographic location information of the monitoring hosts, as well as information about user-configured caps on BitTorrent transfer-rates, allowing us to address this potential source of bias. The data is contributed by users of Ono [4], an extension for the Vuze BitTorrent client, and includes no personally identifiable information. Ono implements a biased peer selection service for reducing P2P-related cross-ISP traffic without sacrificing performance.

For comparison, we use data from the Office of Communication (Ofcom) UK broadband speed reports [14, 15]. Performance data in the reports was collected by SamKnows Ltd., Ofcom’s technical partner in the study, which has recently partnered with the US Federal Communication Commission [19]. The measurements were collected over two time periods: between November 2008 and April 2009, and between November and mid-December of 2010.

3.1 Scalability

Performing ISP characterization at scale is necessary to capture not only the diversity of ISPs and their levels of service, but also to identify variations in performance within an ISP’s level of service. In the following paragraph, we show that passively collected bandwidth measurements can be used to empirically derive an ISP’s levels of service (i.e., speed tiers) and to reveal diversity in performance within and between geographic regions. These results, in turn, serve to demonstrate that extensive coverage is necessary to accurately characterize an ISP’s performance.

Extracting levels of service. Broadband network services typically advertise alternative plans in terms of their maximum download speed. A key challenge for a general approach to ISP characterization at the network edge is to automatically identify the, *a priori* unknown, level of service of a particular connection.

Network-intensive applications should offer a good understanding of the maximum transfer rates that subscribers can achieve. For

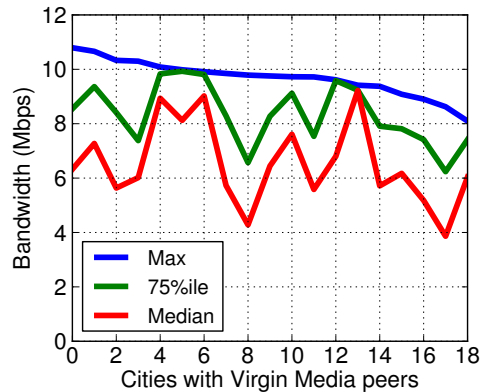


Figure 2: The maximum, 75th-percentile, and median users’ maximum download speeds achieved in Virgin Media’s network across 18 different cities.

instance, it has been observed that BitTorrent clients often saturate their access link [5]. Our methodology leverages this fact.

Using Nov. 2009 data from Virgin Media U.K. customers as an example, we show that we can empirically derive an ISP’s service levels using passive BitTorrent bandwidth measurements across users in an ISP. Figure 1a plots the cumulative distribution function (CDF) of the per-user maximum download speed. We detect these features by finding a steep region (maximum in the first derivative) followed by a flattened region (minimum in the second derivative). We identify three download tiers at 2Mbps, 5.6Mbps and 9.6Mbps, shown in the figure with dark vertical bands, which correspond to several of Virgin Media’s advertised broadband download speeds (see Table 1b).²

Clearly, less popular and higher speed tiers are not as easily observable in a CDF. By tuning the sensitivity of the detection algorithm, we can identify less prominent features in the CDF, albeit at the cost of a higher false positive rate. On faster tiers (e.g. 20Mbps), we find that users are less likely to obtain speeds near their advertised maximum, limited by the upload speed of other peers downloading the same content. In swarm-based systems like BitTorrent, the speed of each participant in a cooperative download group, the swarm, is in part determined by the combined upload speed of the group. The scarcity of peers with high-end connections, at this point, means that those connected would not be able to maximize their download capacity and reduces the “clustering” effect of users gathering near a single threshold.

Regional diversity in performance. We examine the same data to capture the diversity of an ISP’s performance between and within regions. For this analysis, we selected users with maximum speeds faster than Virgin’s 2Mbps service but no faster than their 10Mbps caps, including a 10% tolerance to include users on the threshold. We assume that users that are able to achieve speeds significantly faster than 2Mbps are subscribing to at least the 8 or 10Mbps services.

Figure 2 shows the maximum, 75th-percentile, and median users’ maximum download rates seen by Ono users in Virgin’s 8 and 10Mbps services across 18 different cities (sorted by the fastest user’s speed).³

²This approach could also potentially be applied to derive upload speed tiers.

³While these are two different levels of service, we have not been able to distinguish between them in practice.

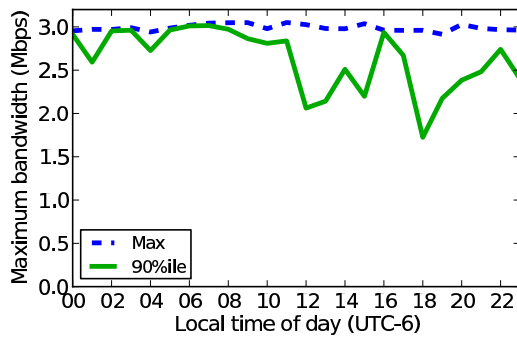


Figure 3: The maximum and 90th-percentile download speed of peers in Rogers’ network for each hour-long period during Nov. 2009.

The range of the “Maximum” curve across the cities shows that the maximum achievable speed is determined by the city where they reside. Virgin customers living in Belfast (pop. ~483,000) report a maximum download speed of 10.8 Mbps while residents of Crawley (pop. ~100,000) see a maximum of just 8 Mbps. The 75th-percentile and median lines show, in addition, that the maximum download speed a user obtained can range significantly within a given city. As an example, consider the case of Nottingham, the eighth city in the graph, where the median user’s maximum download rate is about 4.2 Mbps, less than half the fastest user’s speed of 10 Mbps.

These results point to the conclusion that any solution for ISP characterization must be able to not only scale across ISPs, but within them as well, across both regions and service levels.

3.2 Continuous Monitoring

The service that users receive from their ISPs at any time of day partially depends on the particular access network technology (e.g. DSL or Cable) and the traffic management techniques their ISPs implement. Dischinger et al. [8] show that while DSL ISPs have bandwidth rates that roughly correspond to those advertised, the performance of some Cable ISPs can vary significantly during the day.

Accurately characterizing ISP performance, thus, requires nearly continuous monitoring. In the following paragraphs, we argue that by passively monitoring the performance of network-intensive applications, it is possible to efficiently capture such time-of-day effects.

We use the Rogers (cable-based) ISP as an example as it has been reported to show clear diurnal patterns. Figure 3 shows the maximum and 90th-percentile download rate seen, over a 24-hour period, by Rogers customers subscribed to a level of service of 3 Mbps. Each data point is the average of all such rates from subscribers during a 1-hour interval (i.e., all statistics gathered between 2:00:00 PM and 2:59:59 PM on each day of the month are considered together) during November 2009. To ensure that all customers included in the analysis are subscribed to the same level of service, we only consider peers that had reported a maximum download rate within that level of service.

The graph shows that the reported maximum achieved transfer rate for each hour reaches the 3 Mbps limit fairly consistently. However, the 90th-percentile download rate drops from an average of 2.9 Mbps during the peak residential hours of 5-10PM. This corresponds to approximately 96% of the advertised service during low activity and about 60% during peak residential usage hours. This

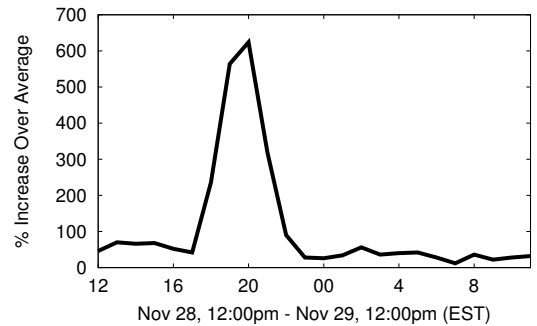


Figure 4: Percent increase over average of DNS timeout reports for Comcast customers during the ISP’s service outage.

drop is consistent with that reported by Dischinger et al. [8] for a different dataset.

3.3 Performed From the Edge

We argue that ISP characterization should be done by the end users to reduce potential biases. From a technical standpoint, conducting tests from the end user has several advantages. Systems that attempt to probe end users from the core of the network cannot diagnose several ambiguous failure modes – such as distinguishing between a client being down or simply not responding to probes.

Similarly, systems with hardware in an ISP’s PoP are unable to capture many aspects of the end user’s performance since the users’ “last mile” access links cannot be observed. Other potential issues may include modem configurations, routing, or access network failures.

One particular example that we focus on is the end user’s perspective on a 5-hour outage of Comcast’s DNS servers in the eastern United States during the evening of Nov. 28, 2010 between 7PM and midnight EST. This event, clearly invisible to any measurement technique *from the core*, resulted in widespread issues for Comcast subscribers in the region.

Our analysis focuses on DNS timeouts as reported by Ono users in Comcast networks. We define a DNS timeout to start when a request to a DNS server fails. We observe a 6x increase in the hours leading up to the outage (see Fig. 4). Before the outage, these failures were generally transient and a second request would be successful. However, the DNS servers eventually stopped responding altogether, at which point the number of reported timeouts fell back to normal levels, as Ono clients waited for DNS servers to return.

In total, we identified 8 unresponsive DNS servers within the Comcast network during this event. Seven of the servers were located in New Jersey and one was located in Virginia (based on a GeoIP lookup). Comcast customers in both these states were affected by the outage.

Events such as this one, which affect users’ ISP service, cannot be detected by ISP characterization systems that are not operated continuously and at the edge.

4. BEYOND CHARACTERIZATION

In this section, we explore some of the applications of our work beyond single-ISP characterization, including comparative studies of ISPs and broadband availability surveys. We show that Internet subscribers could use this information to compare alternative services (Sec. 4.1). Next, we use our preliminary results

Advertised Service	2009			2010		
	ISP (Speed)	Ofcom Speed	%	Ono Speed	Ofcom Speed	%
BT (8)	4.0	31%	6.7	4.4	33%	6.9
BT (20)	—	—	14.8	7.8	52%	19.3
O2 (8)	4.2*	44%	7.4	—	—	7.0
O2 (20)	—	—	10.9	10.8	18%	14.4
Virgin (10)	8.4	24%	9.9	9.6	20%	9.9
Virgin (20)	—	—	16.6	18	2%	15.3
Virgin (50)	—	—	—	45.6	3%	25.9

Table 2: We show download performance trends between 2009 and 2010. “Ofcom Speed” is the average download rate in the Ofcom report for that service level. “%” is the percent of Ono users whose maximum bandwidth was at least the Ofcom Speed. “Ono Speed” shows the 95th-percentile user’s maximum download speed. Empty cells indicate that Ofcom data was not available, or (as in the case of “Virgin (50)”) the tier did not exist in 2009. All speeds are given in Mbps.

for a study of broadband availability in two regions of the US, western Europe, and Japan (Sec. 4.2).

4.1 Comparing ISP Performance

We now investigate whether ISP performance data collected continuously from end users could be used to allow subscribers to compare between advertised and actual performance for different ISPs in their region.

For this analysis, we focus on three particular ISPs in the U.K. – BT,⁴ O2,⁵ and Virgin Media⁶ – and compare our findings with those reported by Ofcom. We compare our BitTorrent traces from two periods, the month of November 2009 and from November 1, 2010 to December 15, 2010, against the 2009 and 2010 Ofcom reports.

Table 2 summarizes Ofcom’s published results and those from our traces. Paralleling the Ofcom report’s results, our Ono data shows increases in achieved speeds for BT’s 8 Mbps service. Although the 95th-percentile for Virgin Media’s 10 Mbps service does not increase (the Ofcom reported values increase by 1.2 Mbps), more Ono users report higher download rates.

Considering the SamKnows/Ofcom report, our approach provides a more complete, virtually live view of performance trends during the studied period. Our analysis of BitTorrent performance as reported by Ono users could be used as an additional data point for side-by-side ISP comparison.

4.2 Broadband Availability

ISP characterization has attracted the attention of governments surveying the availability of broadband access to its citizens. In the next paragraphs, we show the range of broadband services in several key regions as reported by Ono users.

For this analysis, we use data collected by users of our Vuze extension during November 2009. We use GeoLite City⁷ to map each peer IP address to a region. To characterize the service an ISP provides in a region, we find the fastest download rate reported by a user in that region. We use this to classify the ISP into one of five categories: less than 80 Kbps, between 80 and 800 Kbps,

⁴<http://www.bt.com>

⁵<http://www.o2.co.uk/>

⁶<http://www.virginmedia.com>

⁷<http://www.maxmind.com/app/geolitecity>

Label	Description	Regions
US NE	USA Northeast	New York, Penn, New Jersey
US SE	USA Southeast	Kentucky, Tenn, Miss, Alabama
EU	Europe	France, Germany, Italy, United Kingdom
JP	Japan	All regions

Table 3: Groupings of regions used for Fig. 5.

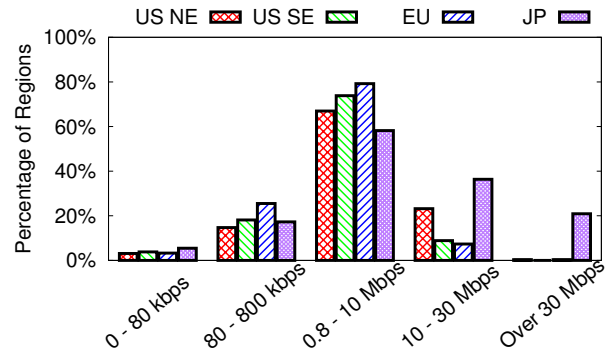


Figure 5: Percentage of subregions containing at least one ISP providing each level of service.

between 800 Kbps and 10 Mbps, between 10 and 30 Mbps, and over 30 Mbps.

We then group these regions to facilitate comparison. In the US, we cluster regions following the Census Bureau’s designated divisions. For brevity, we include only the divisions with the highest and lower rates of high-speed broadband availability: parts of the northeast and parts of the southeast, respectively. In Europe, we group together Germany, France, Italy, and the United Kingdom – four G20 countries representing major economies in that continent. Last, we include Japan, given its uniquely high penetration rate of fiber-based broadband access [3]. Table 3 summarizes the states and countries in each group.

Figure 5 shows, for each grouping, the distribution of regions by the fastest level of service available in a region. Within the US, the Southeast division has a higher percentage of ISPs in the 80 Kbps to 10 Mbps ranges than the Northeast division. The Northeast division, however, shows a higher number of customers with 10-30 Mbps services, suggesting that the Northeast has wider access to faster broadband.

This figure also shows that the European countries appear to have similar broadband access to the US. In contrast, subscribers in Japan clearly have access to significantly faster networks, indicated by a greater percentage of regions with 10-30 Mbps services and a significant number with services over 30 Mbps (20%). By monitoring from a ubiquitous application like BitTorrent, it is thus feasible not only to account for regional differences in a single country, but also to compare the state of broadband availability across multiple nations.

5. DISCUSSION

We have argued that effective ISP characterization must be performed at scale, continuously and by end users. Our preliminary results demonstrate the advantages of using popular network-

intensive applications as hosting platforms for crowdsourcing ISP characterization.

There are a number of interesting open issues that we leave as future work. For instance, C2E is hosted within an application at end hosts and thus lacks visibility into home network configurations and overall traffic. We are currently exploring the use of UPnP to communicate with routers to potentially detect cross traffic on the access link (monitoring TotalBytesSent and TotalBytesReceived on the WAN interface).

Ensuring privacy for participating subscribers is also essential to the success of any crowdsourcing solution. The examples in this paper do not require access to any personally identifiable information. We believe that many other applications can provide such anonymized data to conduct similar analyses. While addressing privacy concerns is essential to the feasibility of our approach, anonymization poses new challenges for ensuring the integrity of comparative analyses.

Finally, while we have discussed C2E as a stand-alone solution, we envision it as a complementary rather than an alternative approach for ISP characterization. For instance, instrumented gateways can help calibrate C2E measurements. On the other hand, data collected from C2E could help inform the deployment of hardware-based measurement points by identifying ISPs, networks and geographic regions that need further study.

6. REFERENCES

- [1] A. Carr. The real ISP report card: Why BitTorrent's network performance data is "miles better" than what Netflix found. *FastCompany*, February 2011.
- [2] M. Casado, T. Garfinkel, W. Cui, V. Paxson, and S. Savage. Opportunistic measurement: Extracting insight from spurious traffic. In *Proc. of HotNets*, 2005.
- [3] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *Proc. of ACM SIGCOMM*, 2006.
- [4] D. R. Choffnes and F. E. Bustamante. Taming the torrent: A practical approach to reducing cross-ISP traffic in peer-to-peer systems. In *Proc. of ACM SIGCOMM*, 2008.
- [5] D. R. Choffnes, F. E. Bustamante, and Z. Ge. Crowdsourcing service-level network event detection. In *Proc. of ACM SIGCOMM*, 2010.
- [6] L. D. Cioccio, R. Tiexeira, and C. Rosenberg. Impact of home networks on end-to-end performance: Controlled experiments. In *Proc. of HomeNets*, 2010.
- [7] D. Croce, T. En-Najjary, G. Urvoy-Keller, and E. W. Biersak. Fast available bandwidth sampling for ADSL links: Rethinking the estimation for large-scale measurements. In *Proc. of PAM*, April 2009.
- [8] M. Dischinger, K. P. Gummadi, A. Haeberlen, and S. Saroiu. Characterizing residential broadband networks. In *Proc. of IMC*, 2007.
- [9] G. Gross. New broadband projects score \$795 million in funding. *InfoWorld*, 2010.
- [10] Keynote. Internet health report. <http://internetpulse.net/>.
- [11] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Natelyzr: Illuminating the edge network. In *Proc. of IMC*, 2010.
- [12] R. Mahajan, M. Zhang, L. Poole, and V. Pai. Uncovering performance differences among backbone ISPs with Netdiff. In *Proc. of USENIX NSDI*, 2008.
- [13] Netflix, Inc. Netflix performance on top ISP networks. <http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html>, January 2011.
- [14] Office of Communication (Ofcom). UK broadband speeds 2009: Consumers' experience of fixed-line broadband performance. Technical report, Ofcom, London, UK, July 2009.
- [15] Office of Communication (Ofcom). UK fixed broadband speeds, November/December 2010. Technical report, London, UK, March 2011.
- [16] A. Ritacco, C. Wills, and M. Claypool. How's my network? A Java approach to home network measurement. In *Proc. of IEEE ICCCN*, 2009.
- [17] SamKnows. Accurate broadband information for consumers, governments and ISPs. <http://www.samknows.com/>.
- [18] SamKnows. Samknows & Ofcom UK broadband performance testing. http://www.samknows.com/broadband/ofcom_and_samknows, June 2009.
- [19] SamKnows. Samknows & the FCC American broadband performance measurement. http://www.samknows.com/broadband/fcc_and_samknows, June 2009.
- [20] Speedtest.net. The global broadband speed test. <http://www.speedtest.net/>.
- [21] YouTube. YouTube video speed history. http://www.youtube.com/my_speed.
- [22] Y. Zhang, Z. M. Mao, and M. Zhang. Detecting traffic differentiation in backbone ISPs with NetPolice. In *Proc. of IMC*, 2009.