A Time for Reliability – The Growing Importance of Being Always On

Zachary S. Bischof and Fabián E. Bustamante
Northwestern University
{zbischof, fabianb}@eecs.northwestern.edu

Categories and Subject Descriptors

C.2.3 [**Communication Systems Organization**]: Computer Communication Networks—*Network Operations*; C.4 [**Performance of Systems**]: [Reliability, availability, and serviceability]

Keywords

Access link reliability, Broadband access networks

1. INTRODUCTION

When a new technology reaches the market, we often focus on the want or need that it can fulfill. However, as the market for a technology matures, reliability often becomes a key differentiating factor between competing products. Examples of this abound, from cars to passenger flights, phones and televisions.

Broadband bandwidth capacities remain a challenge around the world. In the developed world, however, their rapid growth seems to have already passed current needs. Figure 1 shows how peak network usage (95th percentile hour) changes with growing bandwidth capacity (measured from over 7000 residential gateways in the US). We observe a law of diminishing returns in their relationship: the relative increase in demand is greater for lower capacity than for the larger capacity connections. At about 12 Mbps, most users 95th percentile usage remains below 15% of their capacity.

In this poster, we posit that in the developed world broadband reliability will soon become the dominant feature for service comparison. With increased capacity, a growing number of consumers switch their cable TV and landline phone subscription to video streaming or VoIP services over the Internet. For such users, service uptime becomes a serious concern. While previous outages resulted in mild inconveniences while browsing the web, today's network availability can determine the outcome of an emergency situation.

Previous works on broadband networks have focused on characterizing performance [1, 4], and works on reliability typically study wireless networks (cellular, WiFi, WiMax,

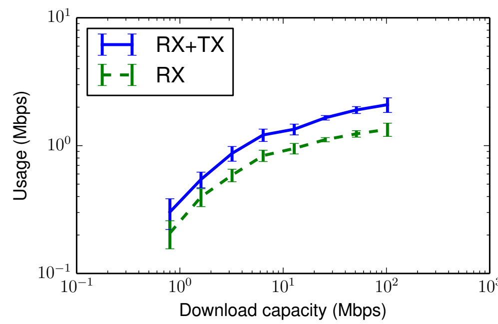

Figure 1: Peak (95th percentile) network usage by download link capacity. The lines represent total traffic ($RX + TX$) and download traffic (RX).

etc). In this poster, we use data collected from residential gateways (via FCC/SamKnows) and end-hosts (via Namehelp) to study the availability and reliability of fixed-line broadband networks. We use traditional metrics (e.g. failure rate, MTBF, MTTR) to quantify each broadband service. We study the reliability of DNS, a service typically provided by ISPs. Using natural experiments, we look at the impact that increased network downtime has on user demand. Also, we note the severity of an outage largely depends on when it occurs (e.g. time of the day, day of the week). As a result, we use typical network usage patterns to design a new metric for service reliability, the average expected traffic lost (AETL), to quantify the impact of each service's outages.

2. APPROACH

In this work, we leverage two datasets. First, we use data provided by the FCC [2], collected from gateways distributed to fixed-line broadband customers in the US. In our study, we use measurements of packet loss rates to M-Lab hosted measurement servers, DNS lookup failures, and byte counters of traffic over the WAN interface. Each of these are recorded at hour intervals. We consider an access link to be down when packet loss rates are above 10% for the hour.

Our end-host collected dataset comes from *namehelp* [3], a tool for improving DNS performance. We designed and deployed an extension for monitoring users' access links (at 30-second intervals) and ISP DNS servers, measuring service availability. In addition to providing more fine-grained information, this also gives us access to a diverse set of users (over 20,000 clients in 115 countries).

SIGCOMM'14, August 17–22, 2014, Chicago, IL, USA.
ACM 978-1-4503-2836-4/14/08.
http://dx.doi.org/10.1145/2619239.2631460 .

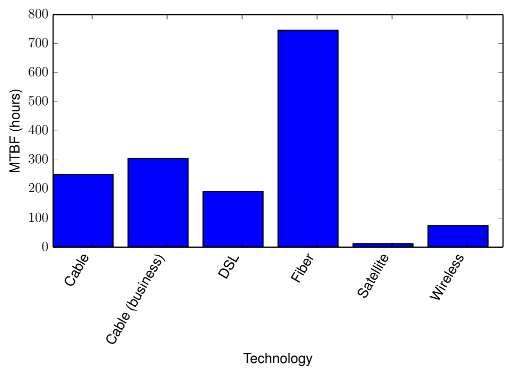

Figure 2: Mean time between failure by technology.

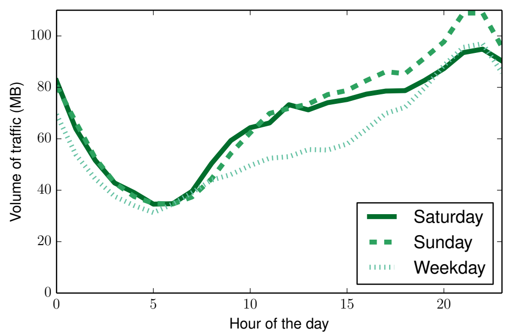

Figure 3: Hourly usage by time of day for weekdays, Saturday, and Sunday.

3. PRELIMINARY RESULTS

Our preliminary analysis has produced a number of interesting results. Below we provide a brief summary of our current findings.

First, we studied how service reliability was affected by the technology employed. Figure 2 shows the mean time between failure (MTBF), averaged for each technology. Overall, fiber services were by far the most reliable, with about 1600 hours (or about 67 days) between service outages. Next were DSL and cable services, with similar averages. Wireless services (such as Clearwire) and satellite services were the least reliable services, with less than 100 hours between failures.

Although DSL services had a higher MTBF than cable, we found that network outages in cable networks were typically more likely to occur during off-peak hours (i.e. in the middle of the night, between 2 and 4 am local time or in the early afternoon). In contrast, outages in DSL services were either uniformly distributed across the day or concentrated during peak usage hours.

With this in mind, we also want to compare services by the severity of their outages. To do this, we calculate the "typical" (median) demand at each hour of the day, for each day of the week, for all users, shown in Fig. 3 (Monday through Friday are grouped together as they showed similar usage patterns). This gives us an idea of how the impact of service outages differ on user traffic. For each service outage that we find, we use the "typical" number of bytes during that hour to estimate the volume of affected traffic.

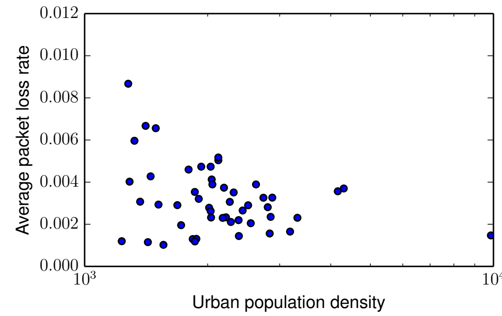

Figure 4: Average packet loss rate in each state versus the average urban-area population density.

We average this across all users to calculate the expected number of bytes of traffic lost, giving us an idea of the impact of each outage.

In this analysis, we found that outages on DSL services were typically more severe than in cable services. Also, we found that both cable and DSL services had similar mean times to recovery. The difference was largely due to differences in *when* the outages occurred. Therefore, despite the lower MTBF, we believe that cable networks in the US will tend to be perceived as more reliable.

We also investigated how service availability related to geographic settings. Since we only have course-grained (state) geographic information for users in the FCC/SamKnows dataset, we used census data from 2010 to find the fraction of the population in each state living in urban areas. Figure 4 shows the average packet loss rate in each state by the population density in urban regions of that state. For the most part, states with a lower urban population density tended to have higher packet loss rates rates. We expect that this is either due to the wider use of less reliable technologies, or that the implementations themselves are less reliable, as the have to cover greater distances. We are studying this, and other aspect of service reliability and availability in greater detail in our ongoing work.

4. ACKNOWLEDGEMENTS

The work was supported in part through NSF Awards CNS 1211375 and CNS 1218287, and a generous Google Faculty Award.

5. REFERENCES

[1] Z. S. Bischof, J. S. Otto, M. A. Sánchez, J. P. Rula, D. R. Choffnes, and F. E. Bustamante. Crowdsourcing ISP characterization to the network edge. In *Proc. of W-MUST*, 2011.

[2] FCC. Measuring broadband America. http://www.fcc.gov/measuring-broadband-america.

[3] J. S. Otto, M. A. Sánchez, J. P. Rula, and F. E. Bustamante. Content delivery and the natural evolution of DNS: Remote DNS trends, performance issues and alternative solutions. In *Proc. of IMC*, November 2012.

[4] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband internet performance: a view from the gateway. In *Proc. of ACM SIGCOMM*, 2011.