A First Look at 1Gbps Fiber-To-The-Home Traffic

Matthew Sargent*, Brian Stack*, Tom Dooner*, Mark Allman*+

*Case Western Reserve University, EECS

+International Computer Science Institute

PROJECT GOALS

Internet Service Providers (ISPs) have been offering service on the order of ten megabits/second for several years. Recently, various research projects, including the Case Connection Zone (CCZ), have connected residential neighborhoods with significantly higher capacity than that offered by ISPs. As these networks of the future come to fruition, we are left with two basic questions that we seek to answer:

1.What will users do with significantly higher capacity?

- Which applications are the most popular?
- How much data will users consume and generate?

2.Are our protocols up to the task of utilizing significantly higher bandwidth in edge networks?

- Are current protocols capable of utilizing a high-bandwidth network effectively?
- What factors contribute to the throughput of various connections?

EXPERIMENTAL SETUP

The CCZ consists of 90 residences each with its own bi-directional 1 Gbps fiber connection. We record all of the network traffic from January 25, 2011 through March 31, 2012 using the Bro Intrusion Detection System running on a monitor between the CCZ homes and the general internet.

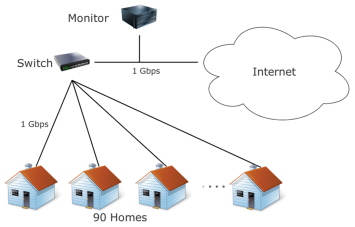

Figure 1—Network Setup: *The network configuration for the Case Connection Zone. Each link in the diagram represents a 1 Gbps fiber cable.*

DATA COLLECTION

Using our monitor, we are able to collect two types of data: (1) Bro logs and (2) packet traces. Each type of log offers us a view at a different level of granularity and helps us to answer very different questions about the network.

Bro Logs:

1. *Connection Logs*: Provide summaries about each transport-layer connection (including UDP "connections").
2. *HTTP Logs*: Contain records for each HTTP transaction including message type such as GET and POST messages.
3. *BitTorrent Logs*: Splits each BitTorrent connection into individual message summaries.

```
1299145806.755200 0.617250 IP_A IP_B other 54487 23546 tcp 68 566 SF X ShADdfFa
1299145805.235290 2.139151 IP_C IP_D other 54446 6901 tcp 451 669 SF X ShADadFf
1299145773.378212 34.085032 IP_E IP_F http 55361 80 tcp 6937 44919 SF X ShADadFf %3292
1299145774.166412 33.298208 IP_E IP_F http 55362 80 tcp 4180 10963 SF X ShADadFf %3293
```

Figure 2—Connection Logs

```
1299144735.587047    %274 start IP_A:62287 > IP_B:80
1299144735.596366    %275 start IP_A:62288 > IP_B:80
1299144735.682427    %274 GET EXAMPLE_URL_A.jpg (404 "Not Found" [162] X.com)
1299144736.793458    %274 GET EXAMPLE_URL_B (200 "OK" [43] X.com)
1299144737.475743    %275 GET EXAMPLE_URL_C.jpg (200 "OK" [7839] X.com)
1299144744.753481    %275 GET EXAMPLE_URL_D.gif (200 "OK" [1789] X.com)
```

Figure 3—HTTP Logs

Packet Traces: Packet-level traces offer the finest level of detail when analyzing network traffic. We collect two types of traces:

1. Header traces: Packet traces where each packet has had its payload stripped off before storage.
2. Full traces: Packet traces containing each packet seen in full.

NETWORK USAGE OVERVIEW

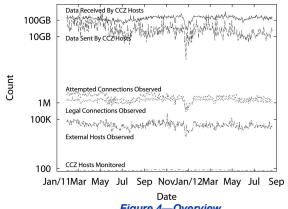

Figure 4—Overview

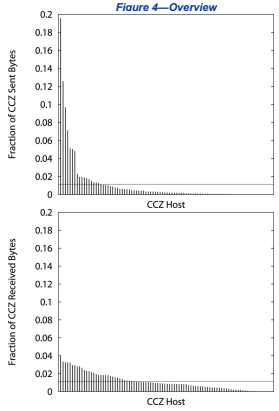

Figure 5—Traffic Volume Per Host

Service	Hosts	Conns.	Sent	Rcvd.
HTTP	90	242 M	789 GB	41 TB
Flash	89	343 K	4.4 GB	2.9 TB
BitTorrent	70	23.6 M	7.8 TB	2.3 TB
HTTPS	90	33 M	437 GB	1.1 TB
Steam	58	36 K	142 MB	584 GB
DNS	90	187 M	7.8 GB	43 GB
Other-1111	25	1.4 M	724 GB	37.4 GB
Other-8332	20	6.5 M	7.1 GB	8.2 GB
Minecraft	22	6.2 M	329 GB	7.2 GB
Unclassified	88	68 M	7.0 TB	3.7 TB
	98%	12%	41%	7%

Table 1—Popular Applications

PER-HOST SPEED

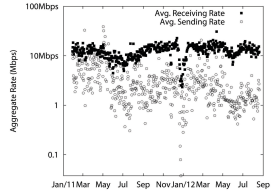

Figure 6—Aggregate Transmission Speed

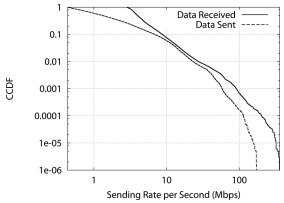

Figure 7—Top Sending and Receiving Bins: *This CCDF shows the sending and receiving rate for the top 1% of 1-second bins across the entire dataset for all users.*

TRANSMISSION SPEED CAUSES

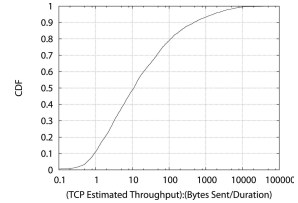

Figure 8—Comparison of TCP's observed and theoretical throughput

ADDITIONAL INFORMATION

Additional information can be found at the project website: http://www.icir.org/dash/