Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in P2P Systems

Choffnes and Bustamante

SIGCOMM 2008

Intro

- Some reports indicate that P2P applications account for 70% of Internet traffic worldwide.
- Current P2P applications do not consider underlying network topology and generate much cross-ISP traffic leading to higher operating costs.

ISPs and P2P Traffic

- ISPs often attempt to limit P2P traffic by using traffic shaping and blocking.
- Early efforts could shape or block traffic based on port number, but many applications now use random port selection.
- Deep packet inspection by ISPs have led to applications encrypting data.
- Some ISPs try to reduce traffic by using caches. This is questionable since the ISPs could be accused of illegally distributing copyrighted data.

Overview

- Summary: there are no good solutions that allow ISPs to reduce traffic generated by P2P applications.
- Proposed solution: use information provided by CDNs for biased peer selection.
- Features: no cooperation between users and ISPs, no infrastructure,and no network topology information required.

Overview

- Hypothesis: clients who are redirected to the same replica server by a CDN are likely to be good peers.
- Methodology: developed a plugin to Azureus BitTorrent client.
- Results overview: reduces cross-ISP traffic, identifies lower latency paths, and increases download rates.

CDNs

- Cache content close to the end user.
- Use various information such as geographic location and network measurements to determine which replica a client should use.
- When a client makes a DNS request for the CDN (e.g., a245.akamai.net), the IP returned is the IP of a "close" replica.

Approach

- Each peer maintains a ratio map that contains ratios of how often the peer is redirected to a particular replica.
- <r_1 =>.75, r_2=>.25> means that the peer has been directed to r_1 75% of the time and r_2 25% of the time.
- Peers with the same ratio map should be "close".
- Use cosine similarity of ratio maps to figure out how close peers are.

Example

- Recall

$$\text{cos\_sim}(a, b) = \frac{\Sigma_{i \in Ia}(\mu_{a,i} \cdot \mu_{b,I})}{\sqrt{(\Sigma_{i \in Ia}(\mu_{a,i}^2) \cdot \Sigma_{i \in Ib}(\mu_{b,i}^2))}}$$

P1 < r_1 = .25, r_2 = .5, r_3 = .25>

P2 < r_1 = 0, r_2 = .5, r_3 = . 5>

P3 < r_1 = .25, r_2 = .5, r_3 = .25>

P4 < r_1 = .3, r_2 = .5, r_3 = .2>

Example

- P1 < r_1 = .25, r_2 = .5, r_3 = .25>
  P2 < r_1 = 0, r_2 = .5, r_3 = . 5>
  P3 < r_1 = .25, r_2 = .5, r_3 = .25>
  P4 < r_1 = .3, r_2 = .5, r_3 = .2>

P1, P2 = .43

P1, P3 = 1.0

P1, P4 = .993

Implementation

- Created a BitTorrent (Azureus) plugin called Ono.
- Periodically, each peer does a DNS lookup for popular CDNs.
  - Air Asia
  - cnn.com
  - lemonde.com
  - Fox News
  - ABC Streaming Video
- Ratio maps are built using /24 addresses since CDNs return IPs of nodes in the same cluster.

Implementation

- Use exponential decay such that nodes not see for 24 hours are removed from map.
- If a peer has no redirection info, it does a DNS lookup once per 30 seconds for 2 minutes.
- After bootstrapping, lookup interval increases by 1 minute each time redirection information is the same as the previous lookup.
- If redirection info is different than the previous lookup, interval is halved.

Implementation

- When peers connect, they exchange their ratio maps.
- Each peer uses cosine similarity to compare its ratio map to that of its peers.
- Peer selection is biased toward peers with similar redirection behavior.
- Only a fraction of total connections are biased to maintain benefit of diversity provided by BitTorrent.

Experimental Setup

- During download, sample transfer rate every 5 seconds.
- Use average of 3 ICMP pings (issued continuously) to determine latency.
- Everytime a connection is established and torn down, a traceroute is executed.
- Use traceroute info to determine AS-level path information between peers.
- Data gathered Dec 1 - 16, 2007.

Reducing Cross-ISP Traffic

- Comparison of path characteristics between Ono-recommended peers and randomly selected peers.
- Used traceroute measurements.
- Results use lemonde.com CDN.
- Figure 2: Ono-recommended peers are closer in terms of IP hops.

Reducing Cross-ISP Traffic

- Doesn’t say anything about cross-ISP traffic.
- Translate traceroute data into number of AS hops
- Figure 3: Few AS hops between Ono client and peers.

Path Characteristics

- Does reduced cross-ISP traffic impact performance?
- Use ping data to determine RTT to peers.
- Figure 4: Ono-recommended peers are two orders of magnitude closer.
- Figure 5: Ono-recommended peers have average 31% lower loss rate.

Transfer Performance

- Do lower latency and lower loss rate translate into faster download/upload speeds?
- Figure 6(a), 6(b): Average download rates are higher by 31%, but median rates are only 2KB/s higher.
- This is an unexpected result, but suggests that peers are overloaded and the bottleneck is the access link to the ISP, not cross-ISP link.

Transfer Performance

- Figure 6 results are not true for all ISPs.
- RDSNET in Romania provides 50 Mb/s unrestricted transfer over fiber for in-network traffic. and 4Mb/s outside of the ISP.
- Figure 7: Average download is 207% higher and media is 883% higher.

Multiple CDN Names

- Figure 10: Some CDN names are better than others.
  - Depends on how many different data centers/servers the CDN has.

Discussion

- Good or bad approach?
- Limitations of this work?
- Extensions of this work?