8th International Workshop on Peer-to-Peer Systems (IPTPS '09)

Boston, MA
April 21, 2009

ROBUSTNESS

Summarized by Ghulam Memon (gmemon@cs.uoregon.edu)

- ***Bringing P2P to the Web: Security and Privacy in the Firecoral Network***
  Jeff Terrace, Harold Laidlaw, Hao Eric Liu, Sean Stern, and Michael J. Freedman, Princeton University

Jeff presented Firecoral, a P2P content distribution network, and addressed the security and privacy concerns in such a network. It runs a tracker to which the content provider delegates the responsibility of content distribution. To ensure that the tracker does not change the content, Firecoral uses a trusted Signing Service (SS). The SS has the responsibility to compute content hash and encrypt it with its own private key. The tracker can only distribute these encrypted hashes. Each client possesses the public key of the SS. This approach prevents the content from being modified.

Firecoral has three components: the tracker, 1000 lines of PHP running on Apache; SS, 700 lines of Python code; and the client (Firefox extension), 7000 lines of Javascript, XUL, and CSS. The Firefox extension uses a whitelist and a blacklist. The whitelist is for those Web sites for which Firecoral must be used, e.g., popular news aggregators and under-provisioned Web sites. The blacklist contains well-provisioned Web sites.

Terrace was asked if Firecoral is simply moving the problem from the content provider to the tracker. He was referred to a different paper that uses a similar approach. More information about Firecoral can be found at http://firecoral.net/.

- ***Deconstructing Internet Paths: An Overlay for AS-Level Detour Route Discovery***
  Sing Wang Ho and Thom Haddow, Imperial College London; Jonathan Ledlie, Nokia Research; Moez Draief and Peter Pietzuch, Imperial College London

This paper focuses on discovering detour paths through the Internet at the AS level. The idea is to exploit the property that detours at the AS-level exist because of BGP anomalies. They use traceroute from 176 PlanetLab nodes to obtain detour paths. They map the IP addresses to respective AS numbers. From this information they construct a graph in which each AS is represented by a node and a path between different AS nodes is represented by a link. Using this graph, they group different Internet paths based on the same detour nodes used.

The authors propose a hierarchical clustering algorithm that is used to group Internet paths with or without detours. Using the 176 PlanetLab nodes, they found that the algorithm can classify the 94.3% of paths with detours and the 83.1% of paths without detours. Out of the paths classified as detour paths, latency can be reduced for 85.3% paths when the suggested detour node is used. The authors also propose a decentralized mechanism for constructing the desired clusters. They construct an overlay network and use gossip to disseminate the acquired information.

The presenter was questioned about the appropriateness of using only 176 nodes for data collection. He was also questioned about the feasibility of mapping IP addresses to AS, given the complicated structure of autonomous systems.

- ***EigenSpeed: Secure Peer-to-peer Bandwidth Evaluation***
  Robin Snader and Nikita Borisov, University of Illinois at Urbana-Champaign

Robin Snader presented a technique for accurate bandwidth estimation in a peer-to-peer system. This is clearly useful, because of the heterogeneous nature of P2P networks. The key idea is to use a modified form of principal component analysis (PCA). The authors had to modify PCA because in its original form, PCA may allow some malicious activities. The focus of the paper is to prevent malicious users from disrupting the bandwidth estimation process.

The paper introduces the idea of consensus bandwidth estimation. Each node maintains the measured bandwidth information about the nodes it communicates with. Different nodes can then share this information to develop a consensus about the system. The bandwidth information obtained from other nodes is weighted based on that node's bandwidth information. For example, a high bandwidth node can have a better estimate than a low bandwidth node.

EigenSpeed solves the problem of node churn by marking newly arriving nodes unevaluated and leaving them out of PCA computation. EigenSpeed avoids the problem of near-sink by using symmetric values for bandwidth estimation by two nodes. If the values are not symmetric, then the lowest value is considered.

Snader was asked where this technique will be most useful. The Tor network is the primary customer for this approach. In general this work was greatly appreciated.

MEASUREMENT

Summarized by Jeff Terrace (jterrace@cs.princeton.edu)

- ***Dynamic Swarm Management for Improved BitTorrent Performance***
  György Dán, KTH, Royal Institute of Technology; Niklas Carlsson, University of Calgary

BitTorrent is widely used on the Internet today, measurements indicating that 54–70% of all Internet traffic is due to peer-to-peer technologies, of which 20–57% is BitTorrent traffic. Mininova, one of the most popular torrent Web sites, was the subject of a study including data from 800,000 torrents and 1700 trackers and covering seeds, leechers, downloads, and file hashes.

They found that performance on small swarms is low and that large swarms can get overloaded because they don't take advantage of multiple trackers. György said their goal was to increase the performance for small swarms and distribute load across multiple trackers for large swarms.

The solution is to use a new protocol, called Distributed Swarm Management (DISM), which allows trackers to work together. DISM uses an approximation algorithm for pairwise peer balancing. This allows for fine-grained swarm adjustment. The resulting analysis shows that a set of trackers implementing DISM is much more balanced, fewer torrents have a low number of peers or a low amount of bandwidth, and a 20–30% increase is gained in performance.

- ***Large-Scale Monitoring of DHT Traffic***
  Ghulam Memon and Reza Rejaie, University of Oregon; Yang Guo, Thomson; Daniel Stutzbach, Stutzbach Enterprises

Dynamic Hash Tables (DHTs) are a widely studied area of research. When deploying a real DHT, Ghulam Memon pointed out, it is often desirable to monitor traffic within the system for measurement studies or system monitoring. Since a DHT is inherently distributed, a central point of monitoring is not available as in traditional systems. Instead, monitors have to be deployed within the system itself, but to monitor all traffic, a large number of monitors must be deployed. This changes the properties of the system you are measuring, while deploying too few monitors might not accurately model the system.

The authors introduced a new model for monitoring DHTs called Minimally Visible Monitors (MVMs). The key idea of an MVM is to insert itself in the DHT but only become visible to the single node it is monitoring. The MVM doesn't respond to any other requests, making it invisible to the rest of the DHT (and treated like a stale, departed node). This allows it to still receive routing requests from its monitoree without affecting the behavior of the DHT. To distinguish between destination and routing traffic, multiple MVMs are inserted for every node within a "zone," defining an N-bit prefix in the DHT identifier space. For all traffic captured within the zone, the destination can be determined with post-processing.

To validate their method, experiments were run with the Kad DHT, and it was determined that Montra captures 90% of all DHT traffic within the zone and correctly determines the destination for 90% of traffic captured for prefixes up to six bits in length.

- ***On the Locality of BitTorrent-based Video File Swarming***
  Haiyang Wang and Jiangchuan Liu, Simon Fraser University;
  Ke Xu, Tsinghua University, Beijing

Haiyang Wang repeated the claim that peer-to-peer (P2P), specifically BitTorrent, traffic has become widely popular on the Internet. One of the problems with P2P traffic is that it is agnostic to the topology of the Internet, so peer selection is not optimized for locality. Locality-based peer selection attempts to minimize inter-ISP traffic, but it also negatively affects the performance of BitTorrent.

The authors did a large-scale measurement study of BitTorrent traffic from btmon.com which consisted of 30,000 video torrents and 44,000 non-video torrents, and they used PlanetLab to collect information on the BitTorrent swarms. The largest portion, 51%, was AVI files. The top AS measured had 16,000 thousand peers, and the top ten ASes had 97 to 165 thousand ASes.

Their measurement showed that large swarms do have poor locality and generate a lot of inter-AS traffic, but small swarms don't have enough diversity within each AS to apply locality-based algorithms. For large enough clusters a peer prediction method can be used, and the authors provide a conditional probability-based peer prediction method, used only when AS clusters become large enough.