An Analysis of Internet Content Delivery Systems

Himani Apte
CS 739: Distributed Systems
University of Wisconsin, Madison

January 20, 2006
Spring 2006

1 Overview

The paper analyzes four different content delivery systems from the arena of client-server oriented world-wide web, content delivery networks and peer-to-peer file sharing systems.

The important features of peer-to-peer systems are symmetry among the peers (since they behave as both servers and clients) and scalability up to as many as millions of machines. Dynamic membership, wide-area network and heterogeneity of the participating systems in terms of their bandwidth, connectivity and performance are some of the distinguishing characteristics of peer-to-peer systems. Such systems tend to be application specific. Some peer-to-peer systems may have a hierarchy among the members, for instance some peers in Kazaa are supernodes and maintain indexes for the content available at peers in the nearby neighborhood.

2 Problem Statement

The paper examined the traffic flow of content delivery systems, focusing largely on web versus peer-to-peer traffic flows, and specifically on HTTP web traffic, Akamai, Kazaa, and Gnutella delivery systems.

3 Methodology

The methodology employed was passive network monitoring of all traffic coming in and out of the border routers between the University of Washington (UW) and the rest of the Internet. The TCP flows were reconstructed at the monitoring hosts to extract information to categorize them into HTTP and non-HTTP traffic. The HTTP traffic is further distinguished into WWW, Kazaa and Gnutella based on the destination ports. Akamai traffic is identified based on whether it is served by an Akamai server.

This methodology misses the internal traffic within the local network of the University, for instance file sharing traffic among Kazaa users within the University network. The analysis also does not take into account the non-HTTP TCP traffic (that is 43% of the total TCP traffic) and non-TCP traffic (that is 3% of the total network traffic). Other peer-to-peer systems such as BitTorrent and Napster have also been excluded from the study.

4 Observations

This section summarizes the observations made in class about the analysis presented in the paper.

4.1 Data characteristics

The HTTP trace summary statistics (presented in Table 1) are unavailable for outbound Akamai traffic as there are no Akamai servers hosted within UW. Kazaa has the highest outbound traffic in terms of net bytes transferred, in spite of a much smaller server and client population within UW.

The total TCP bandwidth consumed by HTTP transfers for different content delivery systems (presented in Figure 1) show a typical diurnal pattern. WWW traffic peaks in

true daylight hour as opposed to Kazaa traffic which peaks late at night.

The HTTP trace was collected in May-June over a nine day period. *The trace could significantly vary based on the timing when it was collected as the network behavior of university students may be widely different during summer break than in the final exams week. Also a longer data sample would be desirable.*

Analysis of the UW client and server TCP bandwidth (presented in Figure 2) indicates that Kazaa peers within UW act as servers much more than the web servers at the university. *A possible reason could be high connectivity of the UW Kazaa peers.*

4.2 Content delivery characteristics

Most bytes are transferred in video objects, although most requests are for GIF and JPEG images. The median object size for WWW is 2 KB, while that of peer-to-peer systems is 4 MB.

The top bandwidth consuming UW clients (Figure 7) and UW servers (Figure 10) are the Kazaa peers. *Hence caching will be most beneficial for Kazaa file sharing system. The cause for large bandwidth consumption are large size objects and very popular objects (that result in large number of requests). As a result small number of clients consume a large amount of bandwidth in peer-to-peer systems.*

The Kazaa and Gnutella servers are not perfectly load-balanced in spite of the scalability of peer-to-peer systems. *Possible cause for this may be the existence of highly popular content on a single server or availability of large-size objects on only a few peers. However, it is incorrect to make conclusions about whether or not the Kazaa and Gnutella servers are load-balanced based on the data available as the trace does not include internal file sharing traffic within UW network.*

4.3 Role of caching

For studying the role of caching in CDNs and P2P systems, the authors have simulated infinite-capacity caches.

The three causes of cache misses (popularly known as the three C's) are cold, capacity and conflict misses. In case of infinitely large cache, the only kind of misses that can occur are cold misses. As a result the cache miss rate is actually a measure of the proportion of unique bytes accessed to the net bytes accessed.

The ideal byte hit rate (Figure 14) for outbound Kazaa traffic was found to stabilize at 85%, while that for inbound traffic did not stabilize by the end of the trace. The cache byte hit rate as a function of population size is presented in Figure 15. *On the one hand, increasing the number of Kazaa clients may increase the number of requests, thereby lowering the cache hit rate. On the other hand, this leads to the complementary effect of caching in the numerous clients, thereby improving the cache hit rate.*

Although the preliminary investigation presented in the paper suggests that caching would have a large effect on a wide-scale P2P system, potentially reducing wide-area bandwidth demands dramatically, it may not actually be feasible to employ caches in P2P systems due to legal issues pertaining to the content distributed in such a system. Also the paper does not give an insight into a realistic size of cache that would be sufficient to obtain improvements in bandwidth usage.

5 Conclusions

The paper presents a quantification of the domination of P2P systems in the modern day Internet traffic.

Although the global characteristics are not easily seen looking at a small part of the network (UW, in this case), it still makes interesting revelations about the network traffic flows. In the future, we would expect WWW traffic to show similar characteristics and an even larger P2P traffic.