Networking for Multimedia
Tech Topic #01: Content Delivery Networks

S. Saroiu, et al., “An Analysis of Content Delivery Systems,” ACM SIGOPS Operating Systems Review Special Issue on Network Behavior, vol. 36, Winter 2002.

Although this paper is clearly more appropriate for the USENIX Internet Technologies and Systems symposium, it was instead submitted and accepted to the USENIX Symposium on Operating Systems Design and Implementation. The assumption here being that the issue is simply one of timing, as the paper was published in 2002 and USITS is held only during odd-numbered years. Napster had already come, hit its peak, and gone, but P2P file-sharing networks were still relatively hot, at least enough so that a paper consisting mostly of taking traces garnered significant interest. As much as the authors at the time may have said to themselves, “ah, I remember Napster”, today the comparable sentiment is “ah, I remember Kazaa”. But it is doubtful that a similar paper that replaced Kazaa with Bit Torrent would be accepted today.

Related work is mostly a description of the 3 major systems the authors will be analyzing, as the paper is primarily a “look how things behave in the real world” expose.

In the methodology it would have been interesting to describe the monitoring hardware topology in greater detail. The section mentions 4 switches, each with a mirror port, but only describes one monitoring host with a single NIC to process the traffic from these interfaces.

The biggest fundamental problem the reviewer had with the methodology paper was the lack of consideration for non-HTTP or TCP traffic. Realistically, it is may be intractable to analyze *all* inbound and outbound network traffic from an entity as large as a university for over a week. Still, a better explanation of why these 5 particular classifications were chosen could have been provided. Eight years ago it may have been more acceptable to assume HTTP-based delivery systems were all that needed to be examined, but they should have at least mentioned why they were ignoring all other traffic such as UDP or even ARP. Given that the paper later turns into a commentary of how network use is changing, a rudimentary comparison of the classified traffic with the aggregate bytes collected by the network monitor would have been useful. It could have been the case that all the flows discussed were dominated in physical link utilization overall by some other protocol never mentioned.

Despite this overlook, the methodology section does a good job discussing the issues with tracing P2P traffic, including Kazaa's habit of downloading fragments. The reviewer wonders if the authors realized that such a methodology would dominate P2P systems in the second half of the decade.

The high-level data characteristics are largely unsurprising and even boring (at least from a 2010 perspective). It is debatable if the HTTP trace summary statistics (table 1) needs to be in a table; the large values make it so busy as to be marginally unreadable. Given the disparity in values across the selected categories, it may have been interesting to present this data on a logarithmic scale.

The references to the authors' 1999 traffic study was a neat data point to compare to, both for consistency and because it is nice to see year over year coherence in research.

The conclusions regarding traffic characterization and data import/export arrived at in the end of section 4 are undermined by the weakness discussed above.

Section 5 is largely spent making very intuitive statements such as “P2P traffic consists of large objects like music and movies while regular web traffic is mostly relatively smaller objects like web pages.” Anyone who had ever heard of Kazaa and used a browser before probably could have arrived at the same conclusion, but certainly it is nice to Do Science and have empirical data to back up these claims. Table 2 has the same readability issue as table one, although in this case it is not obvious how to fix it. In this case the details of the table probably could have been omitted.

Even the number of concurrent connections result the authors label as “surprising” does not seem like a conclusion that could not have been arrived at by considering the main objective of P2P networks and the object size this implies. Longer data transfers equates to more concurrent connections. The “important implications for network infrastructure” claimed at the end of section 5.2 because of all these connections are never explained. Not only is it not clear what these implications are, but such a statement makes certain non-trivial assumptions: if it is known that the number of concurrent connections is high due to P2P usage, what is the responsibility (if any) of the network infrastructure to change to support this? In fact, this statement is in direct violation of the end-to-end argument—perhaps a point the authors wished to avoid debating.

The P2P scalability discussion in section 5.4 is again adversely affected by the aforementioned weakness in methodology. It is inadequate to characterize network-wide behavior by only considering a sub-set of network traffic. The argument is presented that another 450 Kazaa peers would be the equivalent of doubling the entire campus web client population. This statement ignores, for example, the possibility that P2P traffic simply “expands” to fill all available bandwidth. If this is the case, another 450 P2P users may result in no additional bandwidth usage at all.

The caching contribution of the paper is, in a word, minimal. The previous related work (section 2) ignored this aspect. The simulation environment and parameters are not explained, other than the fact that the simulated caches are either “practical” or “ideal”, but even these are not well-qualified. The analysis is cursory and the section mostly seems to exist in a (meek) attempt to leverage some of the trends identified after sifting through a mountain of data.

The P2P cache portion, in addition to suffering from the general weaknesses above, is additionally ill-motivated. It is never explained why network administrators would be inspired to provide a reverse P2P cache when it is well-known assumption that a majority of the traffic on these networks is illegally-distributed content.

The need to append a traditional “contribution” to the work in order to avoid simply organizing collected data in graphs is noted, however in this case the contribution served only as a unnecessary sidebar and detracted from the overall theme of the paper.

In general the paper performs well in providing a large dataset to reinforce several intuitive statements about content delivery networks, specifically peer-to-peer systems. The scalability limits of these systems, however, are over-emphasized and under-analyzed in order to motivate an uninteresting foray into CDN cache design. Not knowing the specific circumstances of the authors, it is the opinion of the current reviewer that, given the state of P2P research at the time, the work could have stood on its own as a in-depth analysis of collected network traces.