

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

This paper is very well written from the abstract to the conclusion. The authors make this paper very easy to follow. Within their introduction, they layout the structure of the rest of the paper by outlining each section. This allows the reader to know what to expect when reading each section. Someone with little knowledge of content delivery networks would be able to read this paper and understand its contribution.

This paper has a significant contribution because it is the first study to compare the workloads of different content delivery architectures with content delivery networks. To make their results more clear, the authors describe the different kinds of traffic that they will analyze (WWW, CDNs, P2P). They follow by describing how they were able to distinguish one type of traffic from another within their traces. Using well explained charts and graphs, they show statistics such as the bandwidth consumed and the number of bytes transferred.

This research contributed to confirmation of many beliefs about CDNs and P2P systems. They were able to prove that P2P traffic transfers the most bytes of data (three times as much as WWW traffic). Intuitively, they show that the object size of P2P traffic is much larger than normal web content. This large object size directly causes a longer transfer rate. So even though there are more requests for web content, there are still more data going back and forth caused by P2P traffic. The paper also shows that, in P2P traffic, there is a small number of those large objects. One thing that the paper was surprised to find was that the P2P clients were getting a lot of their content from a small number of servers. However, this result is not surprising to me because most of the Kazaa users on campus are likely students who probably share a descent amount of similar interests (in movies and music). These students search the content and somebody who has the content they are searching for is likely to have other songs and content that interests the user. The students then probably view the rest of the other users content and begin to download more from that same user.

Another significant contribution that this paper has is that it lays the foundation for more research to be done on caching. The authors found that caching can be very effective in lowering the bandwidth consumption of a P2P network. They found that caching is effective in a small network and even gets more effective as the network grows.