

Zachary Bischof
EECS 443
2008.01.21

Title: The Feasibility of Supporting Large-Scale Live Streaming Applications with Dynamic Application End-Points

Authors: Sripanidkulchai, K., Ganjam, A., Maggs, B., and Zhang, H.

Summary: Application end-point architectures are not only feasible in distributed computing and file-sharing, but also in applications, such as live streaming of audio or video.

Central Ideas: Up until now, supporting applications with dynamic application end-points has been proven to work in cases for distributed computing and file-sharing, however, these cases generally include nodes in the range of 10 to 100. The main argument of this paper is that such architectures are also feasible on larger scales of nodes on the order of 1,000 to 100,000. The basis of this paper is that performance between a higher number of nodes can be maintained by creating trees that nodes together to form smaller clusters and share resources with each other throughout the tree. As mentioned in the paper, a few algorithms stuck out above the rest. In particular, when choosing a parent, results showed that choosing based on minimum depth was a good algorithm when compared to the optimal, impractical, solution. In addition, when choosing a clustering policy, grouping by network delay seemed to work better than geographic distance, since closeness in physical distance does not always equate to closeness on the network.

Flaws: This idea focus more on the feasibility of something rather than a particular implementation of something. As such, the paper is very strong in the sense that makes a valid case for its main argument, that supporting applications with high bandwidth requirements, such as live streaming, could be feasibly supported by dynamic end-points. However, there were some things that were not entirely clear to me. When discussing load balancing for multiple-tree protocols, did not seem as clear to me as it could have. The only major flaw I saw with this proposal was the same flaw that is inherent in these architectures. The main one being that it may be difficult to motivate application end-points to contribute their resources. I also believe that at least some testing with higher bandwidth requirements could have been useful for demonstrating the feasibility as bandwidth increases and services providing high-definition video become more popular, however, this paper is a few years old and this may not have been as doable at the time.

Relevance: As I mentioned previously, there has been an increase in the number of high-definition video streaming services available to consumers. As such, the use of application end-points may become more important as the number of users continues to increase. Then again, since many of these services may not have a significant number simultaneous broadcasts of the same video, it may not be that important to companies providing these streaming services.