
An Alternative Paradigm to Scalable On-Demand Applications...

This paper presents Jukebox a near Video on Demand system that utilizes multicast to deliver content to end users. It was published 1999 in IEEE Transactions on Knowledge and Data Engineering, volume 11 at a time when it seems video content delivery was primarily done through broadcast means such as broadcast TV and pay-per-view along with an increasing interest utilizing multicast technology on to deliver content simultaneously and efficiently to many end users.

The primary idea behind Jukebox is to use a combination of a voting system to schedule video streams to later be delivered to the users who desire to watch the specific video. The authors designed such a system as described in the paper in addition to performed a measurement study on the behavior of its users as various parameters were changed such as the number of votes required to schedule a video feed. Additionally the authors considered how to provide the much desired VCR functions (pause, fast forward, rewind) in such a system.

The authors did a fantastic job of presenting Jukebox and the relevant prior work and as such this was a terrific paper. There was, however, one problem I had which was with figure 8 and it's use of a 3d graph. I still have no clue how to interpret the graph and the figure commentary was no help.

·---
A Digital Found Approach to Reliable Distribution of Bulk Data
---
This paper, published in SIGCOMM 1998, presents a protocol for achieving a digital fountain. A digital fountain is analogous to a water fountain for data in the way that many people can concurrently drink from a fountain to quench their thirst, however no one drinks the same water molecules. Likewise with a digital fountain many people can receive the data they need, however they may each use different pieces to make up the whole. The primary method for accomplishing this is through erasure coding in which the authors came up with a faster method called Tornado code.

Tornado coding is a huge step forward compared to the Reed-Solomon erasure coding both in encoding and decoding speed, however with a slight inefficiency that is more pieces than the original file size are required to recover a file. The authors ensure to point out that this inefficiency is moot when considering the time to acquire additional missing pieces as opposed to the time Reed-Solomon takes to recover missing pieces.

In addition to tornado encoding this paper discusses performing congestion control in multicast via sending over multiple multicast streams where users add additional streams to receive increased throughput. This works in tandem with the tornado coding in such a way that the different streams send different data at different times to maximize efficiency.

The paper presented an a very interesting method for reliably and efficiently multicasting data to end users. The paper was quite well written and as such I have no complaints.