

CS Masters' Thesis Defense

Title: *PARE: A Partial Replication Strategy Adaptive to Movie Popularity Changes for Video-on-Demand Servers*

Speaker: Lantian Gai

Date: Monday, March 12, 2012

Time: 10:00 a.m.

Location: GMCS 418

Thesis advisor: Dr Tao Xie

Abstract:

The performance of a Video-on-Demand (VOD) system is significantly impacted by the throughput of its storage subsystem. Current VOD storage subsystem normally consists of an array of hard disks organized in some RAID organizations. Although it can achieve a large storage capacity, its actual throughput could be lower than its ideal bandwidth. This is mainly because some popular videos receive much more accesses than others, which leads to an imbalanced workload distribution among disks. As a result, the overall throughput of the disk array is substantially degraded. One approach to addressing this problem is to use data replication technology. Unfortunately, most existing data replication algorithms require prior knowledge of file access pattern, which may not be realistic. Further, they normally do not consider video popularity changes. A recently study on a large-scale VOD system discovered that about 50% of user requests require only the first 10 minutes of the target videos. Inspired by the discovery, in this research we propose a dynamic data replication strategy called PARE (partial replication), which makes partial replication for currently most popular videos and evenly redistributes them among disks to adapt to the video popularity changes. Experimental results from synthetic benchmarks and one real-world trace demonstrate that PARE achieves a good performance while saving energy.