NATIONAL UNIVERSITY OF SINGAPORE

School of Computing

PH.D DEFENCE - PUBLIC SEMINAR

Title:	On Flash Crowd Performance of Peer-assisted File Distribution
Speaker:	Ms Cristina Carbunaru
Date/Time:	22 May 2014, Thursday, 02:00 PM to 03:30 PM
Venue:	Executive Classroom, COM2-04-02
Supervisor :	Dr Teo Yong Meng, Associate Professor, School of Computing

Abstract:

Given the growing popularity of peer-assisted file distribution in commercial applications, it is increasingly important to understand the performance of file distribution. File distribution systems often have to cope with extreme conditions, called flash crowds, when the number of users suddenly surges. Flash crowds affect the efficiency of the file distribution and impact user download performance. Thus, content distributors have to ensure that the system has sufficient capacity to cope with flash crowds while maintaining the agreed quality of service with minimum costs. To cope with flash crowds, protocol designers need to understand how p2p protocols match both users and content distributors expectations.

The objective of this thesis is to develop methods for understanding and predicting the performance of file distribution systems during flash crowds. Contrary to current assumption that peer bandwidth utilization is constant throughout the download process, our measurement study on PlanetLab shows three distinct phases in the utilization of peer bandwidth over the download time, namely start-up, maximum utilization, and end-game. Furthermore, a key observation is that the last phase has a step-like function that corresponds with the number of classes of peers with different upload bandwidth. Based on these measurement observations, we propose a general analytical approach to predict the download performance of a file under flash crowd conditions and demonstrate the robustness of our approach for a number of applications.

Our analytical approach models flash crowds using two distinct scenarios. Firstly, a closed model is proposed where a large number of peers join the system in a short period of time and no peers arrive after the flash crowd. This corresponds with content being pushed to users in a staggered manner such as automatic software updates. Secondly, a more complex open model is proposed for multimedia content where arrival rate is not constant over time but decreases as the file popularity drops. Validation of the estimated average download time against PlanetLab measurements shows 14% and 6% error in closed and open systems with up to 150 peers, respectively. Furthermore, validation against simulation for up to 5,000

peers shows that our model maintains an average error of 9%. Our parameter sensitivity study shows that accurately estimating the duration of the maximum utilization phase decreases the model errors by up to 20%. Lastly, to demonstrate the robustness of our modeling approach, we used the simpler homogeneous closed model for heterogeneous closed systems and homogeneous open systems with a smaller than expected increase in model error of 4% and 14%, respectively.

A number of insights for users, service providers and protocol designers of peer-assisted file distribution systems are drawn from applying our model. As peers contribute their upload bandwidth to overall system capacity, both closed and open systems cope well with an increase in the number of peers downloading the file, independent of the peer upload bandwidth. In server provisioning, the closed model shows that the provisioned server capacity, and thus the cost, can be reduced by 40% by relaxing the download time by 10%. Uncoordinated allocation of server bandwidth disproportionately favors fast peers. Thus, increasing the server bandwidth allocated for slow peers and decreasing that for fast peers can be effective in reducing the download time and server provisioning costs without affecting the fairness of the system. In protocol design, coupling our model results with measurements, we discovered that improving fairness can sometimes lead to transient starvation with significant performance degradation. This thesis concludes that achieving high peer bandwidth utilization is essential for scaling peer-assisted file distribution. The handle to reduce both peer download time and provisioning costs is to manage server bandwidth.