In Favour of Posting

THE CHINESE UNIVERSITY OF HONG KONG
Department of Information Engineering

Seminar

Optimal Content Placement for a Large-Scale VoD System by Exploiting Viewing Behavior

by
Dr. K. K. Ramakrishnan
AT&T Labs Research
USA

Date : 16 October, 2012 (Tuesday)
Time : 2:30-3:30pm
Venue : Room 1009, William M.W. Mong Engineering Building
The Chinese University of Hong Kong

Abstract
IPTV service providers offering Video-on-Demand currently use servers at each metropolitan office to store all the videos in their library. With the rapid increase in library sizes, it will soon become infeasible to replicate the entire library at each office. We present an approach for intelligent content placement that scales to large library sizes (e.g., 100Ks of videos). We formulate the problem as a mixed integer program (MIP) that takes into account constraints such as disk space, link bandwidth, and content popularity. To overcome the challenges of scale, we employ a Lagrangian relaxation-based decomposition technique combined with integer rounding. Our technique finds a near optimal solution (e.g., within 1-2%) with orders of magnitude speedup relative to solving even the LP relaxation via standard software. We also present simple strategies to address practical issues such as popularity estimation, content updates, short-term popularity fluctuation, and frequency of placement updates. Using traces from an operational system, we show that our approach significantly outperforms simpler placement strategies. We further focus on how to take advantage of user viewing patterns, such as abandoning the video without viewing it completely, to place content in provider networks to further reduce their storage and network utilization. Our MIP-based solution can serve all requests using only half the link bandwidth used by LRU or LFU cache replacement policies. We also investigate the trade-off between disk space and network bandwidth. (joint work with David Applegate, Aaron Archer, Vijay Gopalakrishnan, Seungjoon Lee, Kyung-Wook Hwang, Vishal Misra, and Deborah Swayne)

Biography
Dr. K. K. Ramakrishnan is a Distinguished Member of Technical Staff at AT&T Labs-Research. He joined AT&T Bell Labs in 1994 and has been with AT&T Labs-Research since its inception in 1996. Prior to 1994, he was a Technical Director and Consulting Engineer in Networking at Digital Equipment Corporation. Between 2000 and 2002, he was at TeraOptic Networks, Inc., as Founder and Vice President.

Dr. Ramakrishnan is an AT&T Fellow, recognized for his fundamental contributions on communication networks and lasting impact on AT&T and the industry, including his work on congestion control, traffic management and VPN services. He is an IEEE Fellow, and has received other awards. His work on the "DECbit" congestion avoidance protocol was recognized in the 1995 retrospective issue of ACM Sigcomm Computer Communication Review as one of the 16 most important papers published over the previous 25 years in ACM Sigcomm publications. The work once again received the ACM Sigcomm Test of Time Paper Award in 2006. He has published nearly 200 papers and has more than 100 patents issued in his name. K.K. has been on the editorial board of several journals and has served as the TPC Chair and General Chair for several networking conferences and has been a member of the National Research Council Panel on Information Technology for NIST.

K. K. received his MS from the Indian Institute of Science (1978), MS (1981) and Ph.D. (1983) in Computer Science from the University of Maryland, College Park, USA. His web page can be found at: (http://www.research.att.com/people/Ramakrishnan_Kadangode_K/)

**** ALL ARE WELCOME ****

Host: Professor Dah-Ming Chiu (Tel: 3943-8357, Email: dmchiu@ie.cuhk.edu.hk)
Enquiries: Information Engineering Dept., CUHK (Tel.: 3943-8385)