CCIT Report #432 June 2003

Optimizing Hybrid Multicast and Unicast Overlay Networks

Oren Unger
Department of Electrical Engineering, Technion
and Zoran Microelectronics
Email: unger@tx.technion.ac.il

Israel Cidon
Department of Electrical Engineering
Technion - Israel Institute of Technology
Email: cidon@ee.technion.ac.il

Abstract

Overlay networks architecture should support high-performance and high-scalability at low costs. This becomes more crucial when communication, storage costs as well as service latencies grow with the exploding amounts of data exchanged and with the size and span of the overlay network. For that end, multicast methodologies can be used to deliver content from regional servers to end users, as well as for the timely and economical synchronization of content among the distributed servers. Another important architectural problem is the efficient allocation of objects to servers to minimize storage, delivery and update costs.

In this work, we suggest a multicast based architecture and address the optimal allocation and replication of dynamic objects that are both consumed and updated. Our model network includes application servers which are potential storage points connected in the overlay network, consumers which are served using multicast and/or unicast traffic and media sources which update the objects using multicast communication. General costs are associated with distribution (download) and update traffic as well as the storage of objects in the servers.

Optimal object allocation algorithms for tree networks are presented with complexities of $O(N)$ in case of multicast distribution and $O(N^2)$ in case of hybrid unicast/multicast distribution. A special case of the hybrid distribution problem automatically selects, for each user, between multicast and unicast distribution.

Using the techniques of the optimal tree algorithm we also present an efficient approximation algorithm for general networks in case of multicast only distribution.

Index Terms

Content Distribution, Location Problems, Multicast, Overlay Networks, Tree Networks

I. Introduction

Recent years have witnessed tremendous activity and development in the area of content and services distribution. Geographically dispersed consumers and organizations demand higher throughput and lower response time for accessing distributed content, outsourced applications and managed services. In order to enable high quality and reliable end-user services despite unpredictable Internet and Intranet conditions, organization and applications service providers (ASPs) employ content distribution networks (CDN) and overlay networks. These networks bring content and applications closer to their consumers, overcoming slow backbone paths, network congestions and physical latencies. Multiple vendors such as Cisco [1], Akamai [2] and Digital Fountain [3] offer CDN services and overlay technologies. Recently, more collaborative models such as distributed storage and peer-to-peer computational models require both consumption and modification of the content by multiple, geographically distributed users [4, 5].

An overlay network is a set of application servers that are connected through the general Internet Infrastructure. Naturally, organizations and ASPs try to optimize the overall cost of the overlay network mainly in terms of storage and communication costs. Efficient allocation of information objects to the overlay network servers reduces the operational cost and improves the overall performance. This becomes more crucial as the scale of services extend to a large number of users over international operation where communication and storage costs as well as network latencies are high. The optimization problem becomes more difficult as the service becomes dynamic and needs to be changed, updated and synchronized frequently.

The popularity of multicast for distribution of the content is increasing with the introduction of real-time and multimedia applications that consume high bandwidth and are delivered to a large number of consumers. Although multicast