

Multicast Service Overlay Design *

Li Lao¹, Jun-Hong Cui², Mario Gerla¹

llao@cs.ucla.edu, jcui@cse.uconn.edu, gerla@cs.ucla.edu

¹ Computer Science Department, University of California, Los Angeles, CA 90095

² Computer Science & Engineering Department, University of Connecticut, Storrs, CT 06269

Keywords

Overlay Design, Multicast Service Overlay Networks, Integer Linear Programming, Greedy

Abstract

Recently, the notion of Multicast Service Overlay Network (MSON) has been introduced to address some of the deployment and maintenance problems of IP multicast and application-level multicast. Basically, an MSON is a backbone service overlay provisioned by an MSON provider and designed to “multiplex” several user overlays. The MSON provider designs the backbone overlay for a large user population and can therefore enjoy large scale economies and benefits. Since the deployment of an MSON is a capital-intensive investment, it is very imperative to carefully design the MSON so that the ISP can make the best revenue out of this investment. In this paper, we formulate the backbone design problem by taking into account ISP operational costs. We divide it into three sub-problems: overlay proxy placement, overlay link selection, and bandwidth dimensioning. We propose several algorithms for each of these sub-problems. Through simulations, we study the effectiveness of different overlay design algorithms, analyze their impact on multicast performance, and suggest some practical solutions for MSON design.

1. INTRODUCTION

Multicast is a useful data delivery method for group communication (from one or many sender(s) to many receivers). Over the years, a lot of research, development, and testing efforts have been devoted to multicast support. However, in today’s Internet, multicast service is far from being widely deployed.

IP multicast is resource efficient by utilizing a tree delivery structure. However, there are many issues which make Internet Service Providers (ISPs) reluctant to deploy and provide IP multicast service ([11, 2]), such as the lack of a scalable inter-domain routing protocol, the requirement of global deployment of multicast-capable IP routers and the lack of appropriate pricing models.

To overcome the scaling limitations of IP multicast, there has been increasing interest in application level multicast, in which multicast functionalities are pushed to end systems. Several successful experiments have been reported in video

and real time multicast, and the state of the art of designing stable, efficient and scalable overlays has been steadily improving ([14, 9, 20, 4]). A challenge in this design and also a cause of substantial overhead has been the “provisioning” of the overlay topology, in part due to available bandwidth estimation on various candidate links. Compared with IP multicast, application-layer multicast protocols are much more flexible, with no requirement of global deployment of multicast-capable routers. However, they usually suffer from the inefficiency of bandwidth usage and the difficulty of managing group membership and maintaining multicast trees especially when multicast groups are large. Moreover, this approach excludes ISPs out of the picture, which, however, is the most motivated party to deploy multicast. Note that end users usually do not care about how group communication is implemented (using unicast or multicast), as long as they could get desired services at good prices.

To address these deployment and maintenance issues of multicast overlays, recently, the notion of Multicast Service Overlay Network (MSON) has been introduced. Basically, an MSON is a backbone service overlay provisioned by an MSON provider and designed to “multiplex” several user overlays. The MSON provider designs the backbone overlay for a large user population and can therefore enjoy large scale economies and benefits.

An MSON consists of service nodes or proxies deployed by an ISP. The design of an MSON relies on a well-defined business relationship between the MSON provider, network service providers (i.e., underlying network domains), and end users: the MSON provider provisions its overlay network according to end user requests, purchases bandwidth from network service providers based on their service level agreements (SLAs), and sells its services to end users via contracts. This characteristic allows the ISP to better design networks in order to improve application performance, and it is complementary to existing end system based approaches. In access domains, end users can communicate with each other by subscribing to MSON proxies. Indeed, communication among end users outside the MSON can employ application layer multicast schemes. For example, Two-tier Overlay Multicast Architecture (TOMA) has been proposed to provide large-scale multicast service [18]. It advocates MSON as the backbone service domain, and end users outside the MSON (i.e., in access domains) form a number of clusters, where an application-layer multicast protocol is used between clustered end users. With carefully designed MSONs, TOMA is claimed to be a scalable, efficient, and practical solution to multicast support.

*This material is based upon work supported by the National Science Foundation under Grant No. 0435515 and Grant No. 0435230.

Obviously, the deployment of an MSON is a capital-intensive investment. Thus, it is very imperative to carefully design the MSON so that the ISP can make the best revenue under given investment. In this paper, we focus our efforts on MSON design by taking into account the operational costs of MSON providers. We formulate the problem and divide it into three sub-problems: overlay proxy placement, overlay link selection, and bandwidth dimensioning. We propose several algorithms for each one of them: ILP and greedy algorithms for overlay proxy placement; Complete Graph, K-MST (Minimum Spanning Tree), and Adjacent Connection for overlay link selection; and a simulation-based bandwidth dimensioning approach. Through simulation studies, we investigate the effectiveness of these algorithms, and analyze their impact on the performance of TOMA. It is worth pointing out that our algorithms are not applicable to TOMA only. They can be adopted or easily adapted for the design of general multicast service overlay networks.

The remainder of this paper is organized as follows. In Section 2, we review some background information as well as related work. In Section 3, we formulate the overlay provisioning problem, and divide it into three sub-problems. We then present some practical algorithms for each sub-problem in Section 4. Finally, we evaluate the performance of the proposed algorithms through simulation study in Section 5, and draw conclusions in Section 6.

2. BACKGROUND AND RELATED WORK

In this section, we give a brief review of TOMA. We also go over some closely related work.

2.1 Review of TOMA

TOMA is a two-tier overlay multicast architecture [18], in which MSON is advocated as the service backbone domain. In the MSON, overlay proxies are strategically placed to form an overlay network, on top of which multicast distribution trees are built for data delivery. Outside the MSON, end users subscribe by transparently connecting to some proxies advertised by the MSON provider. Each proxy organizes some end users into a “cluster”, where an application-layer multicast tree (also denoted as peer-to-peer or P2P multicast tree in this paper) is formed for data delivery.

TOMA uses URL-like unique names to identify groups (e.g., *TOMA://groupname.xyzmson.com/*). An end user explicitly subscribes to a group g by sending out a join request containing the URL-like group name. Through DNS, this request will reach the DNS server of the MSON. After receiving the request, the DNS server will send back to the subscriber a list of IP addresses of the advertised proxies, from which an appropriate proxy will be selected by the user.

After finding a proxy, the end user sends its join request to the proxy, which will then “subscribe” to the multicast group inside the MSON on behalf of this member. The proxy will add the new user to the peer-to-peer multicast tree already set up for this group (or create a new tree) in the local cluster. It will relay the join request towards the source proxy (the proxy to which the source subscribes) in order to extend (or establish) an overlay multicast tree in the backbone domain.

A high level picture of TOMA is illustrated in Fig. 1. There is one multicast group with a source proxy A and destination proxies D and E in this figure. An overlay multicast tree is built in the MSON for this group. Three P2P trees (one for each proxy) are established outside the MSON.




Figure 1: A big picture of TOMA, where F is the group registry server/DNS server, and A, B, C, D and E are overlay proxies.

It is claimed that this architecture can provide scalable and efficient multicast support if the MSON is carefully designed [18]. In this paper, we will study different provisioning algorithms and examine their effects on the performance of TOMA.

2.2 Related Work

Besides TOMA, several other schemes have employed multicast overlay using proxies, such as Overcast [15], Scattercast [7], RMX [8], AMcast [22] and OMNI [5]. However, these architectures are based on different service models from TOMA. Namely, TOMA is based on a well-defined business relationship between MSON providers, network service providers, and end users (or group coordinators). In contrast, Overcast and RMX are deployed exclusively in the user domain; they do not involve the ISPs. Consequently, they mainly address user performance issues such as reliability. In [22], [21] and [5], although the authors mention the concept of overlay tree construction, they focus on optimizing end-to-end delay and access bandwidth usage at the proxies since they assume the bandwidth at proxies is scarce resource. The actual overlay proxy placement, overlay link selection, and bandwidth dimensioning within the service overlay network are usually not addressed. Instead, the authors in [22], [21] and [5] all assume that the proxies are pre-deployed, overlay links between any two proxies are possible (i.e., a full mesh among proxies is used), and the link capacity is simplified as the proxy out-degree bound. Moreover, these proposals only consider overlay tree construction for a single group. In contrast, our work addresses all the problems of proxy placement, link selection, and bandwidth dimensioning by taking into consideration large numbers of groups.

The overlay architecture closely related to our work is SON [12], which incorporates bandwidth dimensioning algorithms for a unicast service overlay network. Due to the many extra features required by multicast, bandwidth dimensioning for unicast is not sufficient for multicast services. Issues of end user distribution, physical network awareness, and multicast routing etc, which are not considered in SONs, would be major concerns for MSONs. To the best of our knowledge, this paper is the first work focusing on the design of multicast service overlay networks (MSONs).

3. PROBLEM FORMULATION

In this section, we formulate the MSON provisioning prob-

lem. To design a good MSON, we need to take into account traffic patterns (e.g., group and member distribution), group bandwidth requirements, end-to-end delay concerns, as well as multicast routing algorithms.

3.1 Defining the Problem

Graph Model We model the physical network topology as an undirected graph $G = (V, E)$, where V and E denote the sets of network nodes (or routers) and physical links respectively. The total number of routers in the network is denoted as n . Each link $e \in E$ has a bandwidth capacity $c(e)$. We also denote the set of all possible paths (between any two nodes) in G as $path(E)$.

Group Model From long-term measurements, we can obtain a set of groups $\{g_i\}$ with group member distribution and bandwidth requirements [3, 6, 10], and then use this information for MSON design. To quantify the group member population clustered around a router, we use a random-weighted model [13], in which we assign to each router i a weight w_i . This weight represents the number of directly attached end users participating in multicast groups.

Problem Formulation The overlay provisioning problem can be formulated as follows: Given the set of groups $\{g_i\}$, and a physical network topology $G = (V, E)$, find a virtual topology $G' = (V', E')$ on top of G (where $V' \subset V$ and $E' \subset path(E)$), in which each $e' \in E'$ is assigned bandwidth $b(e')$, such that G' can accommodate all the groups, and the cost of G' and the average end-to-end delay of group members are minimized.

The determination of the cost of G' depends on the MSON management policy. In most cases, we can use the sum of the assigned bandwidth (purchased from underlying network service providers) to estimate the cost of bandwidth, and the number of overlay links (i.e., $|E'|$) to measure the overlay maintenance overhead.

Obviously, the MSON provisioning problem defined above is a multi-objective optimization problem, which is NP-hard due to the combinatorial explosion. To tackle this difficult problem, we use a divide-and-conquer approach: we first divide the whole problem into several sub-problems, and then provide practical solutions for each of them.

3.2 Dividing the Problem

From the above definition, clearly, to obtain G' , the overlay ISP needs to make three decisions: 1) determine the locations of the proxy nodes, i.e., select V' ; 2) select the overlay connections between these proxies, i.e., choose E' ; 3) compute the bandwidth to be reserved on each overlay link, i.e., assign a bandwidth $b(e')$ to each $e' \in E'$.

Following the principle of divide-and-conquer, we divide the overlay provisioning problem into three sub-problems: overlay proxy placement, overlay link selection and bandwidth dimensioning (note that this approach may incur “sub-optimality” of the solutions since its main concern is the problem manageability). We will present algorithms to solve them accordingly in the next section.

4. PROVISIONING ALGORITHMS

4.1 Overlay Proxy Placement

For the proxy location problem, we assume the total number of overlay proxies is bounded due to the deployment cost of proxies. Since end users tend to receive data packets from

the closest proxies, the locations of overlay proxies directly affect the latency of data transmission. Intuitively, if a router is connected to a lot of users, a proxy should be placed near this router to reduce the average delay. Hence, we can minimize the total delay between users and their proxies.

In this way, the overlay proxy placement becomes an optimization problem: given the number of group members w_i ($1 \leq i \leq n$) for all n routers and the shortest distance d_{ij} between any two routers i and j ($1 \leq i, j \leq n$), find no more than K (a predefined value, and $1 \leq K \leq n$) routers as proxies, such that the weighted sum of distances from each router to its nearest proxy is minimized. We propose an ILP (Integer Linear Programming) formulation and a greedy algorithm to solve this problem.

4.1.1 ILP Formulation

Before presenting the ILP formulation, we define the following variables:

$$x_i = \begin{cases} 1, & \text{if router } i \text{ is selected as an overlay proxy} \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in [1, n]$$

and

$$y_{ij} = \begin{cases} 1, & \text{if router } j \text{ subscribes to proxy } i \\ 0, & \text{otherwise} \end{cases} \quad \forall i, j \in [1, n]$$

Here we use indicator variables y_{ij} in addition to x_i to overcome the difficulty of representing the closest proxy to a router, and our objective function guarantees that a router will subscribe to its closest proxy.

The objective is to minimize the total distance from any router to its closest proxy weighted by the number of members attached to the router:

$$\min \sum_{j=1}^n \left(\sum_{i=1}^n d_{ij} \times y_{ij} \right) \times w_j,$$

subject to the following constraints:

- 1) Every router subscribes to exactly one proxy:

$$\sum_{i=1}^n y_{ij} = 1 \quad \forall j \in [1, n];$$

- 2) A router cannot subscribe to a non-proxy node:

$$y_{ij} \leq x_i \quad \forall i, j \in [1, n];$$

- 3) Every proxy should have at least one subscriber (this constraint is not enforced in the objective function since the goal is to minimize the total distance):

$$x_i \leq \sum_{j=1}^n y_{ij} \quad \forall i \in [1, n];$$

- 4) No more than K nodes can be selected as proxies:

$$\sum_{i=1}^n x_i \leq K.$$

The complexity of this ILP formulation is $O(n^2)$ in terms of the number of variables and the number of constraints. Existing ILP solvers, such as lp_solve (an open source software) and CPLEX (a commercial mathematical optimization tool), can be used to solve for x_i and y_{ij} .

4.1.2 Greedy Algorithm

By solving the above ILP formulation, an optimal solution can be obtained for the overlay proxy placement problem. However, due to its high computation complexity, when the number of variables or constraints is very large, the problem becomes intractable. Hence, we develop a greedy algorithm (Algorithm 1) to efficiently solve this problem.

Algorithm 1 Greedy-Proxy(R, K)

```

1:  $P \leftarrow null$  //initialize proxy set
2: for  $i = 1$  to  $K$  do
3:   //  $K$  is the number of proxies to be selected
4:   for all  $r \in R$  do
5:     //  $R$  is the set of remaining routers
6:     compute the weighted sum of distance  $d_r$  if  $r$  is selected as proxy in addition to  $P$ 
7:   end for
8:    $r_{best} \leftarrow r$  with the minimum  $d_r$ 
9:    $P \leftarrow P \cup \{r_{best}\}$ 
10:   $R \leftarrow R - \{r_{best}\}$ 
11: end for

```

The greedy algorithm works as follows: in each step, a router r is selected as a proxy if its selection can reduce the weighted sum of distances (i.e., the objective function) by the largest amount. This procedure repeats until the maximum number of overlay proxies K is reached. In each iteration, computing the weighted sum of distance for each router requires $O(n)$ time, and the whole iteration takes $O(n^2)$ time. Therefore, the time complexity of the algorithm is $O(Kn^2)$.

This algorithm tries to minimize the objective function at each step, but it may be stuck in local minima and not reach the global minimum. In Section 5, we show that this algorithm is able to achieve competitive performance in comparison with the optimal ILP solution.

Considering Overlay Proxy Deployment Cost In the above algorithms, we assumed that the maximum number of proxies (K) is pre-determined. Alternatively, the overlay ISP may want to minimize the sum of data delivery cost and proxy deployment cost. In this case, this problem is equivalent to the classical *Warehouse Location Problem*, which tries to determine the appropriate locations of warehouses to minimize shipping cost and warehouse maintenance cost. Efficient branch-and-bound algorithms [16] and heuristic algorithms [17] can be adopted to solve this version of the overlay proxy location problem.

4.2 Overlay Link Selection

Once the locations of the overlay proxies have been determined, the next step is to connect these proxies into a mesh, on top of which overlay multicast trees will be constructed inside the MSON. Usually, there are two optimization goals when constructing an overlay mesh, namely, minimizing end-to-end latency and minimizing the cost of multicast trees constructed on this mesh. Since these two objectives generally can not be satisfied simultaneously, we propose different approaches for them separately.

4.2.1 Optimizing End-to-end Latency

A naive approach to construct an overlay mesh with optimized end-to-end latency is to use *Complete Graphs*, that is, an overlay link is established for every pair of proxies. In




Figure 2: An example of overlay link connection approaches. (a) Underlying physical network topology. Nodes B, C, F, G and H are proxy nodes. (b) Complete Graph. (c) 1-Minimum Spanning Tree. (d) Adjacent Connection.

a network with n proxy nodes, this approach results in a total of $\frac{n(n-1)}{2}$ overlay links. Consequently, it induces a large amount of overlay maintenance overhead since directly connected overlay nodes usually need to periodically exchange refresh messages. An example of overlay link connection approaches is illustrated in Fig. 2. As shown in Fig. 2(a), among the 9 routers, B, C, F, G, and H are selected as overlay proxies. When the Complete Graph approach is adopted to connect the proxies (Fig. 2(b)), every pair of proxies use the shortest overlay path as the overlay link. As a result, any shortest path tree constructed on top of this overlay is essentially a star topology, in which a source proxy is connected to destination proxies directly. In other words, multicast in this case is equivalent to multiple unicast. Thus, multicast trees on Complete Graphs will have the shortest delay but the largest total cost and the highest overlay maintenance overhead.

To reduce tree cost and overlay overhead while achieving minimum delay, we can adopt an approach called *Adjacent Connection* [19] to “reduce” Complete Graphs: an overlay link is established between two overlay nodes only if the shortest IP-layer path between them does not go through any other overlay nodes. In other words, we eliminate the overlay links that are implied by (or composed of) other overlay links. For example, in Fig. 2(d), the overlay link (B, H) is removed because it is composed of the overlay links (B, G) and (G, H). In this way, the shortest overlay paths between any two proxies are exactly the same as the network-level shortest paths, and it can create efficient multicast trees instead of multiple unicast paths. Therefore, the constructed overlay networks resemble the underlying network topologies and have high routing efficiency.

4.2.2 Optimizing Tree Cost

Since minimal spanning trees (MSTs) aim at minimizing the total cost of links on the tree, only overlay links with low costs will be included in this type of graphs. Thus we can use an MST as the overlay topology when trying to optimize

the cost of multicast trees constructed on top of it. Though the MST approach minimizes multicast tree cost, the limited set of overlay links may increase the end-to-end delay between some proxies. For example, in Fig. 2(c), an MST is constructed as the overlay network. The network-level shortest path from proxy B to F is 2 hops; however, the shortest path on the overlay network has to go through proxy G and thus has a length of 3 hops. Additionally, this approach has obviously the fewest number of overlay links (i.e., $n - 1$) and the lowest overlay maintenance overhead, but it tends to concentrate multicast traffic on a few overlay links, and it is also more susceptible to node/route failures. To improve the fault tolerance and load balancing, a more general approach is to construct k -MST, in which k least-overlapping minimum spanning trees are used to connect overlay proxies [24]. We will examine the trade-offs of the proposed overlay link selection approaches through simulations in Section 5.

4.3 Bandwidth Dimensioning

Having determined the overlay topology (overlay proxies and links), we then need to decide the bandwidth required on overlay links to accommodate the groups obtained from long-term measurements.

Bandwidth dimensioning is a complex procedure since it is tightly related to multicast routing algorithms. Thus, we suggest a simulation-based bandwidth dimensioning approach, taking multicast routing algorithms into account. The basic idea is as follows: we compute multicast trees based on the given routing algorithm for all the groups, then by summing up the traffic volume on each tree, the amount of bandwidth to be leased on every overlay link can be determined.

Based on long-term measurement results, the computed bandwidth is sufficient to satisfy the bandwidth requirement of these groups on average. In reality, the traffic peak rate is likely to exceed the average rate frequently. Hence, the MSON provider should consider over-dimensioning the network by reserving a certain percentage of extra bandwidth. If the bandwidth is still insufficient during peak hours, the MSON can either reject some groups or lease bandwidth from higher-tier ISPs for short-term usage (possibly at a higher price).

5. SIMULATION STUDY

In this section, we evaluate the effectiveness of our overlay provisioning algorithms. We investigate the impact of the overlay proxy placement, overlay link selection, and bandwidth dimensioning algorithms on the performance of TOMA.

5.1 Simulation Settings

Considering that different topologies may affect the effectiveness of the proposed algorithms, we use two different network topologies in the simulation experiments. The first topology is abstracted from a real network, AT&T IP backbone [1], which has a total of 123 nodes: 9 gateway routers, 9 backbone routers, 9 remote GSR (Gigabit Switch Routers) access routers, and 96 remote access routers. The abstract topology is constructed as follows: the attached remote access routers of a gateway or backbone router are “contracted” into one **contracted node**. In addition, we create a neighbor node called **exchange node** for each gateway router in the backbone, since gateway routers represent connectivity to other peering networks and/or Internet public exchange




Figure 3: Multicast tree cost for different proxy placement approaches.

points. This abstraction procedure results in a simplified network of 54 nodes.

For this topology, we use a weight-based model mentioned in Section 3 to generate group members: for each router i , we randomly attach end hosts to this router with a probability proportional to w_i . Thus, for a group with a fixed size, the average number of group members attached to a router is proportional to this router’s weight. Among the 54 routers, gateway nodes and backbone nodes are assumed to be core routers and are assigned weight 0. Each access router is assigned a weight of 0.01, and a contracted node’s weight is the summation of the weights of all the access routers from which it is contracted. Exchange nodes are assigned a weight of 0.9 since they usually connect to peering networks and tend to have a large number of group members.

The second topology is a router-level topology (which is referred to as 1755 topology in this paper) with approximately 300 routers (measured by Rocketfuel Project [23]). In the 1755 topology, we use a random group member distribution model to generate group members, since we have no information to differentiate the routers. On the other hand, by using the random distribution model, we also tend to explore how member distribution models affect the overlay provisioning algorithms.

Multicast group members join multicast groups during an interval of 400 seconds. Inside the overlay domain, source-based overlay trees are constructed for data delivery between proxies; outside the overlay network, application-layer minimum spanning trees rooted at the proxy nodes are established for communication among clustered end users. Data are collected after multicast trees stabilize.

For the overlay topology construction, we set the number of proxy nodes to 9 for the AT&T topology, which is equal to the number of exchange nodes (to which the majority of members are attached), and 20 for the 1755 topology. Unless otherwise specified, proxy nodes are placed in the locations determined by the greedy algorithm and connected using the Adjacent Connection approach as explained in Section 4.2.

To measure the quality of multicast trees, we use three performance metrics, i.e., multicast tree cost, link stress, and path length. *Multicast tree cost* is defined as the total number of physical links in a multicast tree, and it measures the resource usage of a multicast tree. *Link stress* is measured by the number of redundant packets delivered on a single network link, and it quantifies the efficiency of a multicast




Figure 4: Average link stress for different proxy placement approaches.

tree. *Path length* is defined as the number of hops on the end-to-end path from a source to a receiver, and it measures the quality of data delivery paths.

5.2 Overlay Proxy Placement

To investigate the effectiveness of our proxy placement algorithms, we compute proxy locations based on the weights assigned to the routers using the ILP and greedy algorithms, and compare the performance of multicast groups with different group sizes on the resulting overlay networks. In this set of simulations, we only use the AT&T topology due to the long computation time of ILP when the number of nodes is high. As a reference, we also include the multicast performance when proxies are randomly selected.

The results of multicast tree cost vs. group size are plotted in Fig. 3. The ILP solution yields the lowest tree cost, since it optimally places proxies as close to exchange routers (which have a bulk of members) as possible, and thus prevents the packets from traversing the links between proxies and access routers multiple times. In contrast, the random approach chooses proxies randomly from all the routers in the network. As a result, end users may use some fairly long paths to connect to the proxies, and some physical links may be used multiple times by different users. As for the greedy algorithm, it selects some gateway routers and exchange routers, so its performance lies in between. For the same reason, a similar trend can be observed for average link stress, which is shown in Fig. 4. Note that link stress is determined by the multicast tree cost and the number of distinct physical links in the tree, so even though tree cost increases with group size, link stress remains rather static.

The average end-to-end path length is depicted in Fig. 5. It is clear that by intelligently placing proxy nodes, both the ILP and greedy solutions reduce the average path length greatly. For example, the average path length is approximately 10.4 to 10.7 when proxies are randomly placed, and it is reduced to below 7.0 when the ILP or greedy approach is adopted. This remarkable improvement is again due to the fact that the two proposed approaches place proxies closer to end users to avoid using unnecessarily long paths for data delivery.

5.3 Overlay Link Selection

In this set of experiments, we connect proxies using the three approaches described in Section 4.2, namely, Adjacent

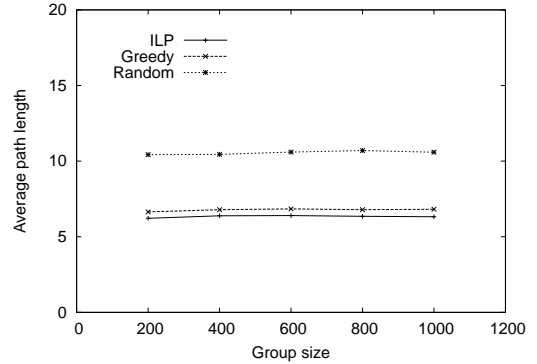


Figure 5: Average path length for different proxy placement approaches.

Table 1: Comparison of Overlay Link Selection Approaches in the AT&T topology.

Metrics	CG	AC	1-MST	2-MST
<i>Number of Overlay Links</i>	36	27	8	16
<i>Multicast Tree Cost</i>	27.9	23.7	18.0	19.7
<i>Link Stress</i>	1.78	1.52	1.16	1.25
<i>Path Length</i>	3.10	3.10	4.32	3.31

Connection (AC), Complete Graph (CG) and k -MST (Minimum Spanning Tree), and compare the performance of resulting overlays. For k -MST, we set k to be 1 and 2, separately. Since the quality of the application-level multicast trees outside the overlay network only depends on the locations of the proxy nodes and the end users rather than the connections between the proxies, we only consider the performance of multicast trees between proxy nodes inside the overlay domain. We present the results of overlay multicast performance in Table 1 for the AT&T topology and Table 2 for the 1755 topology.

As shown in these two tables, Minimum Spanning Tree and Adjacent Connection produce lower tree cost and link stress, whereas Complete Graph gives the highest values for both metrics. This is reasonable, since the Minimum Spanning Tree approach includes only overlay links with low cost, and the Adjacent Connection method removes overlay links that go through some intermediate proxies. On the other hand, as explained in Section 4.2, in the Complete Graph scheme, a shortest path tree has the same cost as the sum of the multiple unicast paths, and thus is the highest among all the algorithms.

From Table 1 and 2, it is clear that, in terms of path length, Complete Graph and Adjacent Connection perform the best among all the solutions. Since multicasting in a

Table 2: Comparison of Overlay Link Selection Approaches in the 1755 topology.

Metrics	CG	AC	1-MST	2-MST
<i>Number of Overlay Links</i>	190	76	19	38
<i>Multicast Tree Cost</i>	81.9	48.3	31.2	35.0
<i>Link Stress</i>	2.23	1.32	1.01	1.06
<i>Path Length</i>	4.24	4.24	6.05	4.98




Figure 6: Request rejection ratio for bandwidth dimensioning in the AT&T topology.

complete graph is equivalent to multi-unicasting, Complete Graph will have the shortest end-to-end delay. Adjacent Connection yields the same delay as Complete Graph, because it attempts to remove the overlay links without affecting the routing efficiency: if the shortest path between two proxies does not go through other proxies, they will be connected directly using the shortest paths; otherwise, they will connect to the intermediate proxy nodes via overlay links. Unlike these two approaches, MST eliminates overlay links with high cost and trades path length for higher tree efficiency (See the example given in Section 4.2).

Another important observation is that compared with 1-MST, 2-MST increases tree cost and link stress, but significantly reduces path length. This is consistent with our intuition: when we build two minimum spanning trees, more overlay links are included, so shorter paths can be usually found, while tree cost is not optimal any more. This is another evidence of the trade-off between end user delay and tree efficiency.

Furthermore, we find out that the numbers of overlay links in the generated overlay topologies using different approaches are different. For example, in the AT&T topology, the number of overlay links is 8 for 1-MST, 16 for 2-MST, 27 for Adjacent Connection and 36 for Complete Graph. This metric measures the overlay maintenance overhead. Obviously, the observed results are consistent with our analysis in Section 4.2.

5.4 Bandwidth Dimensioning

To determine the bandwidth for each overlay link, we model the traffic pattern as follows: 1) for every group, a proxy has attached end users participating in the group with a given probability as described in Section 5.1; 2) groups arrive according to a Poisson process, and the average life time is exponentially distributed; 3) each group requires the same bandwidth throughout its lifetime, and the bandwidth requirements for different groups follow a certain distribution. In the simulations, we define three types of multicast groups: 50% of the groups are low bandwidth (10K), 30% are medium bandwidth (100K), and 20% are high bandwidth (1M)¹.

In our study, we generate sample traces of multiple groups

¹These bandwidth requirements are typical for several common applications on the Internet, e.g., text message applications (10K), voice/audio or low-quality video applications (100K), and high-quality video applications (1M).




Figure 7: Request rejection ratio for bandwidth dimensioning in the 1755 topology.

and use the simulation-based approach proposed in Section 4.3 to determine the bandwidth assignment. For each trace of multiple groups, we compute their source-based multicast trees, determine the bandwidth required on each link at every sampling point, and take the average of the samples. Then we over-dimension the bandwidth by a certain amount and validate the dimensioning results by measuring *group join request rejection ratio* for another group trace with the same traffic pattern. A join request of a multicast group is rejected if the residual bandwidth of an overlay link is not enough to accommodate the multicast tree computed for that group.

Fig. 6 and 7 show the average request rejection ratio in the AT&T and 1755 topologies, respectively. In these two figures, the percentage of over-dimensioning (denoted as *od*) and the average number of active groups are varied. Due to the high computation overhead of the larger topology 1755, we use fewer groups for this topology in the simulation experiments.

We observe three trends in both figures. First, the rejection ratio remains very low for different numbers of co-existing groups, even without over-dimensioning. It is reasonable that the rejection ratio is not 0 because the bandwidth is reserved for the average traffic rate instead of the peak rate. In addition, the rejection ratio heavily depends on the specific topology and group membership model. For example, when there are 1000 active groups, the rejection ratio is approximately 3% for AT&T and 8% for 1755. One major reason is that the AT&T topology has a smaller number of overlay links and group members are very unevenly distributed; thus, multicast traffic is more “predictable” on each overlay link. Second, as *od* increases, the rejection ratio decreases, because more bandwidth is reserved. Third, the rejection ratio decreases when there are more groups in the network, since the sum of bandwidth requirements for a larger number of groups tend to vary less significantly. These results indicate that our bandwidth dimensioning scheme is indeed very effective.

5.5 Summary

Based on the simulation and analysis results, we can draw the following conclusions: 1) the performance of TOMA is significantly affected by the overlay provisioning algorithms; 2) for overlay proxy placement, the greedy algorithm is a practical and efficient solution; 3) for overlay link selection,

we need to consider the trade-off between end user delay and tree cost. MST and Adjacent Connection produce better results than Complete Graph; 4) slight over-dimensioning helps to improve the performance of the proposed simulation-based bandwidth dimensioning approach.

6. CONCLUSIONS

In this paper, we have studied the MSON design problem. Our contributions can be summarized as follows. We first formulate the difficult MSON provisioning problem, and divide it into three sub-problems: overlay proxy placement, overlay link selection, and bandwidth dimensioning. For each of these sub-problems, we propose several algorithms: an ILP formulation and a greedy algorithm for overlay proxy placement; Complete Graph, K-MST, and Adjacent Connection for overlay link selection; and a simulation-based approach for bandwidth dimensioning. Through extensive simulations, we find that the performance of TOMA is significantly affected by different overlay provisioning algorithms. Based on our simulation and analysis results, we make the following suggestions for MSON design: the greedy algorithm is a very effective solution for overlay proxy placement; MST and Adjacent Connection produce better results than Complete Graph; over-dimensioning is needed for bandwidth dimensioning in order to reduce the join request rejection ratio.

Future Work To make MSONs a reality, there are a large course of work to be done. In the future, we plan to study the impact of various network topologies and multicast group models (such as different group lifetime distributions) on the effectiveness of the proposed algorithms. In addition, we would consider the design of peer-to-peer trees in local clusters when MSONs are used as the backbone service overlays, and investigate the feasibility of using existing end system based approaches such as Narada [9] and NICE [4] in this case.

In short, we hope that our basic work in this paper could shed some light on the challenging MSON design problem and inspire some future interesting work in this direction.

7. REFERENCES

- [1] *AT&T IP Backbone*, 2001.
<http://www.ipservices.att.com/backbone/>.
- [2] K. Almeroth. The evolution of multicast: From the Mbone to inter-domain multicast to Internet2 deployment. *IEEE Network*, Jan./Feb. 2000.
- [3] K. Almeroth and M. Ammar. Multicast group behavior in the internet's multicast backbone (Mbone). *IEEE Communications*, June 1997.
- [4] S. Banerjee, C. Kommareddy, and B. Bhattacharjee. Scalable application layer multicast. In *Proceedings of ACM SIGCOMM*, Aug. 2002.
- [5] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, and S. Khuller. Construction of an efficient overlay multicast infrastructure for real-time applications. In *Proceedings of IEEE INFOCOM*, Apr. 2003.
- [6] R. Chalmers and K. Almeroth. Modeling the branching characteristics and efficiency gains of global multicast trees. *Proceedings of the IEEE INFOCOM' 2001*, Apr. 2001.
- [7] Y. Chawathe, S. McCanne, and E. A. Brewer. *An Architecture for Internet Content Distributions as an Infrastructure Service*, 2000. Unpublished, <http://www.cs.berkeley.edu/yatin/papers/>.
- [8] Y. Chawathe, S. McCanne, and E. A. Brewer. RMX: Reliable multicast for heterogeneous networks. In *Proceedings of IEEE INFOCOM*, Mar. 2000.
- [9] Y.-H. Chu, S. G. Rao, and H. Zhang. A case for end system multicast. In *Proceedings of ACM Sigmetrics*, June 2000.
- [10] J.-H. Cui, M. Faloutsos, D. Maggiorini, M. Gerla, and K. Boussetta. Measuring and modelling the group membership in the Internet. In *Proceedings of the ACM SIGCOMM/USENIX Internet Measurement Conference (IMC 2003)*, Oct. 2003.
- [11] C. Diot, B. Levine, J. Lyles, H. Kassem, and D. Balensiefen. Deployment issues for the IP multicast service and architecture. *IEEE Network*, Jan. 2000.
- [12] Z. Duan, Z.-L. Zhang, and Y. T. Hou. Service overlay networks: SLAs, QoS, and bandwidth provisioning. *IEEE/ACM Transactions on Networking*, 11(6):870–883, 2003.
- [13] A. Fei, J.-H. Cui, M. Gerla, and M. Faloutsos. Aggregated Multicast with inter-group tree sharing. *Proceedings of NGC2001*, Nov. 2001.
- [14] P. Francis. *Yoid: Extending the Multicast Internet Architecture*. White paper, <http://www.aciri.org/yoid/>.
- [15] J. Jannotti, D. K. Gifford, K. L. Johnson, M. F. Kaashoek, and J. W. O. Jr. Overcast: Reliable multicasting with an overlay network. In *Proceedings of USENIX Symposium on Operating Systems Design and Implementation*, Oct. 2000.
- [16] B. M. Khumawala. An efficient branch and bound algorithm for the warehouse location problem. *Management Science*, 18(12):B718–B731, Aug. 1972.
- [17] A. A. Kuehn and M. J. Hamburger. A heuristic program for locating warehouses. *Management Science*, 9(4):643–666, July 1963.
- [18] L. Lao, J.-H. Cui, and M. Gerla. Toma: A viable solution for large-scale multicast service support. In *Proceedings of IFIP Networking*, May 2005.
- [19] Z. Li and P. Mohapatra. The impact of topology on overlay routing service. In *Proceedings of IEEE INFOCOM*, Mar. 2004.
- [20] D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel. ALMI: An application level multicast infrastructure. In *Proceedings of the 3rd USNIX Symposium on Internet Technologies and Systems*, Mar. 2001.
- [21] S. Shi and J. S. Turner. Routing in overlay multicast networks. In *Proceedings of IEEE INFOCOM*, June 2002.
- [22] S. Shi, J. S. Turner, and M. Waldvogel. Dimensioning server access bandwidth and multicast routing in overlay networks. In *Proceedings of NOSSDAV'01*, June 2001.
- [23] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proceedings of ACM SIGCOMM*, Aug. 2002.
- [24] A. Young, C. Jiang, M. Zheng, A. Krishnamurthy, L. Peterson, and R. Y. Wang. Overlay mesh construction using interleaved spanning trees. In *Proceedings of IEEE INFOCOM*, Mar. 2004.