Scalable Application Layer Multicast

Suman Banerjee, Bobby Bhattacharjee, Christopher Kommareddy

Abstract— **We describe a new scalable application-layer multicast protocol, specifically designed for low-bandwidth, data streaming applications with large receiver sets. Our scheme is based upon a hierarchical clustering of the application-layer multicast peers and can support a number of different data delivery trees with specific desirable properties. We show that group members maintain state for a constant number of other members and the control overhead is also a constant.**

Next, we present extensive simulations of both our protocol and the Narada application-layer multicast protocol over Internet-like topologies. Our results show that for groups of size 32 or more, our protocol has lower link stress (by about 25%), improved or similar end-to-end latencies and similar failure recovery properties. More importantly, it is able to achieve these results by using orders of magnitude lower control traffic.

Finally, we present results from our wide-area testbed in which we experimented with 32-100 member groups distributed over 8 different sites. In our experiments, average group members established and maintained low-latency paths and incurred a maximum packet loss rate of less than 1% as members randomly joined and left the multicast group. The average control overhead during our experiments was less than 1 Kbps for groups of size 100.

I. Introduction

Multicasting is an useful primitive for scaling multi-party applications since it can potentially decouple the size of the receiver set from the amount of state kept at any single node in the network, including the data source. However, deployment of network-layer multicast [9] has not widely adopted by most commercial ISPs, and thus large parts of the Internet are still incapable of native multicast more than a decade after the protocols were developed. *Application-Layer Multicast* protocols [8], [10], [5], [12], [13] do not change the network infrastructure, instead they implement multicast forwarding functionality exclusively at end-hosts. Such application-layer multicast protocols and are increasingly being used to implement efficient commercial content-distribution networks.

In this paper, we present a new application-layer multicast protocol called SMOP (Scalable Multicast Overlay Protocol), which is specifically designed to support applications with very large receiver sets. Such applications include news and sports ticker services such as Infogate (See http://www.infogate.com) and ESPN Bottomline (See http://www.espn.com); real-time stock quotes and updates, e.g. the Yahoo! Market tracker, and popular Internet Radio sites. All of these applications are characterized by very large (potentially tens of thousands) receiver sets and relatively low bandwidth soft real-time data streams that can withstand some loss. We refer to this class of large receiver set, low bandwidth real-time data applications as *data stream* applications. Data stream applications present an unique challenge for application-layer multicast protocols: the large receiver sets usually increase the control overhead while the relatively low-bandwidth data makes amortizing this control overhead difficult. SMOP can be used to implement very large data stream applications since it has a provably small (constant) control overhead and produces low latency distribution trees. It is possible to implement high-bandwidth applications using SMOP as well; however, in this paper, we concentrate exclusively on low bandwidth data streams with large receiver sets.

The authors are at Dept. of Computer Science,University of Maryland, College Park, MD 20742, USA

A. Application-Layer Multicast Protocol Evaluation

The basic idea of application-layer multicast is shown in Figure 1. Unlike native multicast where data packets are replicated at routers inside the network, in application-layer multicast data packets are replicated at end hosts. Logically, the end-hosts form an overlay network, and the goal of application-layer multicast is to construct and maintain an efficient overlay for data transmission. Since application-layer multicast protocols must send the identical packets over the same link, they are less efficient than native multicast. Two intuitive measures of the "goodness" (*stress* and *stretch* were defined in [8]. The stress metric is defined per-link and counts the number of identical packets sent by a protocol over each underlying link in the network. The stretch metric is defined per-member and is the ratio of path-length from the source to a member versus the length of the unicast shortest path. Consider an application-layer multicast protocol in which the data source unicasts the data to each receiver. Clearly, this "multi-unicast" protocol minimizes stretch, but at a cost of $O(N)$ stress at links near the source (N is the number of group members) and would require $(O(N)$control overhead at some single point. However, this protocol is robust in the sense that any number of group member failures do not affect the other members in the group.

All the application-layer multicast protocols we discus, including SMOP, are purely end-to-end, i.e. they are oblivious to the underlying network topology. In general, application-layer multicast protocols can be evaluated along three dimensions:

- *Quality of the data delivery path.* The quality of the tree is measured using topological metrics such as stress, stretch, and node degrees.
- *Robustness of the overlay.* Since end-hosts are potentially less stable than routers, it is important for application-layer multicast protocols to mitigate the effect of receiver failures. The robustness of application-layer multicast protocols is measured by quantifying the extent of the disruption in data delivery when different members fail, and the time it takes for the overlay to restore delivery to the other members. We present the first comparison of this aspect of the application-layer multicast protocols.
- *Control overhead.* For efficient use of network resources, the control overhead at the members should be low. This is an important cost metric to study the scalability of the scheme to large member groups.

B. Existing Approaches

A number of different application-layer multicast schemes have been proposed in recent literature. They can be classified into two broad categories: tree-first approaches and mesh-first approaches. In the tree-first approach, members directly construct an overlay tree topology for data delivery, and additional

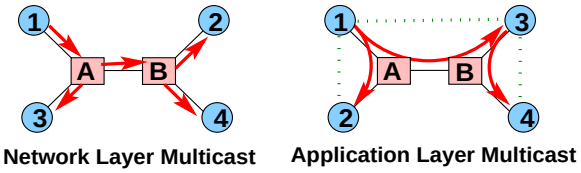

Fig. 1. Network-layer and application layer multicast. Square nodes are routers, and circular nodes are end-hosts. The dotted lines represent peers on the overlay.

control links are monitored and maintained to allow quick recovery from member failures. Yoid [10] and ALMI [13] are examples of the tree-first approach. As the name suggests, in the mesh-first approach, members distributedly construct a mesh (an overlay topology in which multiple paths exist between pairs of members). Each member participates in a routing protocol on the mesh topology, and generates a source-specific tree to all other members. Narada [8], and Gossamer [5] are examples of the mesh-first approach.

No end-to-end approach can produce an overlay topology with constant stretch and stress bounds simultaneously for arbitrary distribution of members in the network. It is possible to do better if (and only if) the underlying topology is known. In [11], a *centralized* topology-aware tree-building algorithm is described in which the stretch between any pair of end-points is bounded by a constant factor. The tree degree of a member, can however, be unbounded. However, with a slight modification to this algorithm, we can simultaneously guarantee a *constant* degree bound for the members, and a $O(\log N)$ bound for the stretch.

C. SMOP Trees

Our goals for SMOP were to develop an efficient, scalable, and distributed tree-building protocol which did not require any underlying topology information. Specifically, the SMOP protocol reduces the worst-case state and control overhead at any member to $O(\log N)$, maintain a constant degree bound for the group members and approach the $O(\log N)$ stretch bound possible with the topology-aware centralized algorithm. Additionally, we also show that an average member maintains state of a *constant number of other members*, and similarly incurs a *constant control overhead* in topology creation and maintenance.

Unlike previous approaches, SMOP is neither an exclusively tree- or mesh-first approach. We create a hierarchically-connected control topology which has higher connectivity than a tree. This control topology is equivalent to a mesh; however, the data delivery path is implicitly defined in the way the mesh is structured and no additional route computations are required. SMOP is, therefore, a hybrid of the tree- and mesh-first approaches. Unlike the centralized topology-aware algorithm described in [11], it is not possible to bound pair-wise end-to-end stretch using a distributed algorithm like SMOP. However, in our simulations and on wide-area experiments, *all* of the pair-wise end-to-end paths were, in fact, bounded by the $O(\log N)$ stretch bound.

Along with the analysis of the various bounds, we also present a simulation-based performance evaluation of SMOP. In our simulations, we compare SMOP to the Narada application-layer multicast protocol. Narada was first proposed [8] as an efficient application-layer multicast protocol for small group sizes. Extensions to it have subsequently been proposed [7] to tailor its applicability to high-bandwidth media-streaming applications for these groups, and have been studied using both simulations and implementation. Lastly, we present results from a wide-area implementation in which we quantify the SMOP run-time overheads and convergence properties for various group sizes.

D. Roadmap

The rest of the paper is structured as follows: In Section II, we describe our general approach, explain how different delivery trees are built over SMOP and present a few theoretical bounds about the SMOP protocol. In Section III, we present the operational details of the protocol. We present our performance evaluation methodology in Section IV, and present detailed analysis of the SMOP protocol through simulations in Section V and a wide-area implementation in Section VI. We elaborate on related work in Section VII, and present conclusions in Section VIII.

II. Solution Overview

The SMOP protocol arranges the set of end hosts into a hierarchy; the basic operation of the protocol is to create and maintain the hierarchy. The hierarchy implicitly defines the multicast overlay data paths, as described later in this section. The member hierarchy is crucial for scalability, since most members are in the bottom of the hierarchy and only maintain state about a constant number of other members. The members at the very top of the hierarchy maintain (soft) state about $O(\log N)$ other members. Logically, each member keeps detailed state about other members that are *near* in the hierarchy, and only has limited knowledge about other members in the group. The hierarchical structure is also important for localizing the effect of member failures.

In this paper, we use end-to-end latency as the distance metric between hosts. While constructing the SMOP hierarchy, members that are "close" with respect to the distance metric are mapped to the same part of the hierarchy: this allows us to produce trees with low stretch.

The SMOP hierarchy described in this paper, is similar to the member hierarchy used in [3] for scalable multicast group re-keying. However, the hierarchy in [3], is layered over a multicast-capable network and is constructed using network multicast services (e.g. scoped expanding ring searches). Our work builds the necessary hierarchy on a unicast infrastructure to provide a multicast-capable network. In fact, it follows that hierarchy construction and maintenance and the multicast services created by our application-layer multicast protocol can be efficiently leveraged by a scheme like [3] to scalably implement multicast group key distribution mechanisms.

In the rest of this section, we describe how the SMOP hierarchy is defined, what invariants it must maintain, and describe how it is used to establish scalable control and data paths.

A. A Hierarchical Arrangement of Group members

The SMOP hierarchy is created by assigning members to different levels (or layers) as illustrated in Figure 2. Layers are numbered sequentially with the lowest layer of the hierarchy being layer zero (denoted by L_0). Hosts in each layer are parti-

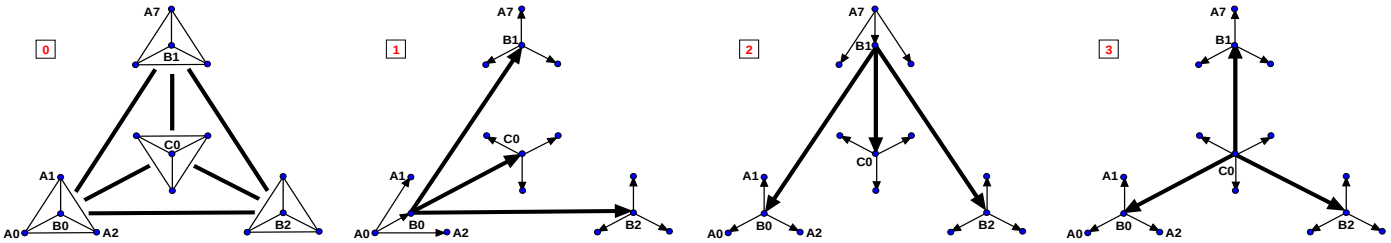

Fig. 3. Control and data delivery paths for a two-layer hierarchy. All A_i hosts are members of only L_0 clusters. All B_i hosts are members of both layers L_0 and L_1. The only C host is the leader of the L_1 cluster comprising of itself and all the B hosts.

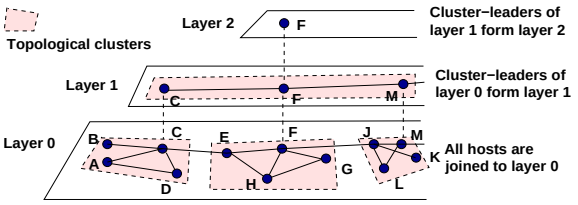

Fig. 2. Hierarchical arrangement of hosts in SMOP. The layers are logical entities overlaid on the same underlying physical network.

tioned into a set of clusters. Each cluster is of size between k and $2k-1$, where k is a constant, and consists of a set of hosts that are close to each other. Further, each cluster has a cluster leader. The protocol distributedly chooses the (graph-theoretic) center of the cluster to be its leader, i.e. given a set of hosts in a cluster, the cluster leader has the minimum maximum distance to all other hosts in the cluster. This choice of the cluster leader is important in guaranteeing that a new joining member is quickly able to find its appropriate position in the hierarchy, using a very small number of queries to other members.

Hosts are mapped to layers using the following scheme: All hosts are part of the lowest layer, L_0. The clustering protocol at L_0 partitions these hosts into a set of clusters. The cluster leaders of all the clusters in layer L_i join layer L_{i+1}. This is shown with an example in Figure 2, using $k = 3$. The layer L_0 clusters are [ABCD], [EFGH] and [JKLM][1]. In this example, we assume that C, F and M are the centers of their respective clusters of their L_0 clusters, and are chosen to be the leaders. They form layer L_1 and are clustered to create the single cluster, [CFM], in layer L_1. F is the center of this cluster, and hence its leader. Therefore F belongs to layer L_2 as well.

The SMOP clusters and layers are created using a distributed algorithm described in the next section. The following properties hold for the distribution of hosts in the different layers:

- A host belongs to only a single cluster at any layer.
- If a host is present in some cluster in layer L_i, it must occur in one cluster in each of the layers, $L_0, \ldots, L_{i-1}$. In fact, it is the cluster-leader in each of these lower layers.
- If a host is not present in layer, L_i, it cannot be present in any layer L_j, where $j > i$.
- Each cluster has its size bounded between k and $2k-1$. The leader is the graph-theoretic center of the cluster.
- There are at most $\log_k N$ layers, and the highest layer has only a single member.

[1] We denote a cluster comprising of hosts $X, Y, Z, \ldots$ by $[XYZ\ldots]$.

We also define the term *super-cluster* for any host, X. Assume that host, X, belongs to layers $L_0, \ldots, L_{i-1}$ and no other layer, and let [..XYZ..] be the cluster it belongs it in its highest layer (i.e. layer L_{i-1}) with Y its leader in that cluster. Then, the super-cluster of X is defined as the cluster, in the next higher layer (i.e. L_i), to which its leader Y belongs. It follows that there is only one super-cluster defined for every host (except the host that belongs to the top-most layer, which does not have a super-cluster), and the super-cluster is in the layer immediately above the highest layer that H belongs to. For example, in Figure 2, cluster [CFM] in Layer 1 is the super-cluster for hosts B, A, and D. In SMOP each host maintains state about all the clusters it belongs to (one in each layer to which it belongs) and about its super-cluster.

B. Control and Data Paths

The host hierarchy can be used to define different overlay structures for control messages and data delivery paths. The neighbors on the control topology exchange periodic soft state refreshes and do not generate high volumes of traffic. Clearly, it is useful to have a structure with higher connectivity for the control messages, since this will cause the protocol to converge quicker.

In Figure 3, we illustrate the choices of control and data paths using clusters of size 4. The edges in the figure indicate the peerings between group members on the overlay topology. Each set of four hosts arranged in a 4-clique in Panel 0 are the clusters in layer L_0. Hosts B_0, B_1, B_2 and C_0 are the cluster leaders of these four L_0 clusters. and form the single cluster in layer L_1. Host C_0 is the leader of this cluster in layer L_1.

In the rest of the paper, we use $Cl_j(X)$ to denote the cluster in layer L_j to which member X belongs. It is defined if and only if X belongs to layer L_j.

The control topology for the SMOP protocol is illustrated in Panel 0, Figure 3. Formally, on the control topology, the peers of a member, X, which belongs to layers $L_0, \ldots, L_i$, are exactly the other members of the corresponding clusters to which X belongs in each of these layers, i.e. $Cl_0(X), \ldots, Cl_i(X)$. Using the example (Panel 0, Figure 3), member A_0 belongs to only layer L_0, and therefore, its control path peers are the other members in its L_0 cluster, i.e. A_1, A_2 and B_0. In contrast, member B_0 belongs to layers L_0 and L_1 and therefore, its control path peers are all the other members of its L_0 cluster (i.e. A_0, A_1 and A_2) and L_1 cluster (i.e. B_1, B_2 and C_0). In this control topology, each member of a cluster, therefore, exchanges soft state refreshes with all the remaining members of the cluster. This al-

Procedure : *MulticastDataForward*(h, p)
{ $h \in$ *layers* $L_0, \ldots, L_i$ *in clusters* $\mathrm{Cl}_0(h), \ldots, \mathrm{Cl}_i(h)$ }
for j in $[0, \ldots, i]$
if $(p \notin Cl_j(h))$
ForwardDataToSet($Cl_j(h) - \{h, p\}$)
end if
end for

Fig. 4. Data forwarding operation at a host, h, that itself received the data from host p.

lows all cluster members to quickly identify changes in the cluster membership, and in turn, enables faster restoration of a set of desirable invariants (described in Section II-D), which might be violated by these changes.

The delivery path for multicast data distribution needs to be loop-free, otherwise, duplicate packet detection and suppression mechanisms need to be implemented. Therefore, we choose a tree to be the data delivery path in the SMOP protocol. More specifically, given a data source, the data delivery path is a source-specific tree, and is implicitly defined from the control topology. Each member executes an instance of the Procedure *MulticastDataForward* given in Figure 4, to decide the set of members to which it needs to forward the data. Panels 1, 2 and 3 of Figure 3 illustrate the consequent source-specific trees when the sources are at members A_0, A_7 and C_0 respectively. We call this the *basic data path*.

To summarize, in each cluster of each layer, the control topology is a clique, and the data topology is a star. It is possible to choose other structures, e.g. in each cluster, a ring for control path, and a balanced binary tree for data path.

C. Analysis

Each cluster in the hierarchy has between k and $2k-1$ members. Then for the control topology, a host that belongs only to layer L_0 peers with $O(k)$ other hosts for exchange of control messages. In general, a host that belongs to layer L_i and no other higher layer, peers with $O(k)$ other hosts in each of the layers $L_0, \ldots, L_i$. Therefore, the control overhead for this member is $O(k.i)$. Hence, the cluster-leader of the highest layer cluster (Host C_0 in Figure 3), peers with a total of $(O(k \log N)$ neighbors. This is the worst case control overhead at a member.

It follows using *amortized cost* analysis that the control overhead at an average member is a constant. The number of members that occur in layer L_i and no other higher layer is bounded by $O(N/k^i)$. Therefore, the amortized control overhead at an average member is

$$\leq \frac{1}{N} \sum_{i=0}^{\log N} \frac{N}{k^i} k.i = O(k) + O(\frac{\log N}{N}) + O(\frac{1}{N}) \rightarrow O(k)$$

with asymptotically increasing N. Thus, the control overhead is $O(k)$ for the average member, and $O(k \log N)$ in the worst case. The same holds analogously for the basic data path and the stress at the members.

While an $O(k \log N)$ upper-bound is acceptable for the control topology, for high data rates, hosts may be unable or unwilling to forward data to $O(k \log N)$ other hosts. Therefore, we define an enhanced data path (using local transformations of the basic data path) for applications generating high bandwidth streams (e.g. video), where the stress at each member on the data path is also reduced to a constant, at the cost of some additional (constant) control overhead at the members. We outline this enhancement mechanism in the Appendix.

D. Invariants

All the properties described in the analysis hold as long as the hierarchy is maintained. Thus, the objective of SMOP protocol is to scalably maintain the host hierarchy as new members join and existing members depart. Specifically the protocol described in the next section maintains the following set of invariants:

- At every layer, hosts are partitioned into clusters of size between k and $2k-1$.
- All hosts belong to an L_0 cluster, and each host belongs to only a single cluster at any layer
- The cluster leaders are the centers of their respective clusters and form the immediate higher layer.

III. Protocol Description

In this section we describe the SMOP protocol. Due to space constraints, we do not present packet formats etc.; instead, we concentrate on the high-level protocol details.

We assume the existence of a special host that all members know of a-priori. Using nomenclature developed in [8], we call this host the Rendezvous Point (RP). Each host that intends to join the application-layer multicast group contacts the RP to initiate the join process. For ease of exposition, we assume that the RP is always the leader of the single cluster in the highest layer of the hierarchy. It interacts with other cluster members in this layer on the control path, and is bypassed on the data path. (Clearly, it is possible for the RP to not be part of the hierarchy, and for the leader of the highest layer cluster to maintain a connection to the RP, but we do not belabor that complexity further). For a real application like streaming media transfer, the RP could be a distinguished host in the domain of the data source.

The SMOP protocol itself has three main components: initial cluster assignment as a new host joins, periodic cluster maintenance and refinement, and recovery from leader failures. We discuss these in turn.

A. New Host Joins

When a new host joins the multicast group, it must be mapped to some cluster in layer L_0. We illustrate the join procedure in Figure 5. Assume that host A_{12} wants to join the multicast group. First, it contacts the RP with its join query (Panel 0). The RP responds with the hosts that are present in the highest layer of the hierarchy. The joining host then contacts all members in the highest layer (Panel 1) to identify the member closest to itself. In the example, the highest layer L_2 has just one member, C_0, which by default is the closest member to A_{12} amongst layer L_2 members. Host C_0 informs A_{12} of the three other members (B_0, B_1 and B_2) in its L_1 cluster. A_{12} then contacts each of these members with the join query to identify the closest member among them (Panel 2), and iteratively uses this procedure to find its L_0 cluster.

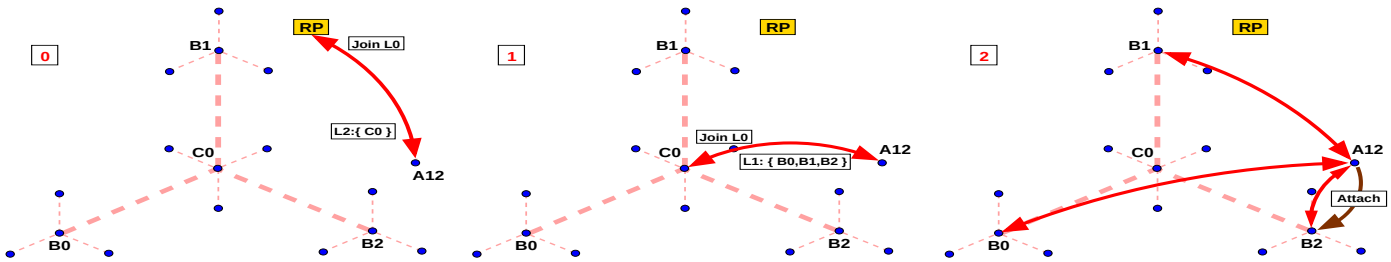

Fig. 5. Host A_{12} joins the multicast group.

It is important to note that any host, H, which belongs to any layer L_i is the center of its L_{i-1} cluster, and recursively, is an approximation of the center among all members in all L_0 clusters that are below this part of the layered hierarchy. Hence, querying each layer in succession from the top of the hierarchy to layer L_0 results in a progressive refinement by the joining host to find the most appropriate layer L_0 cluster to join that is close to the joining member. The outline of this operation are presented in pseudocode as Procedure *BasicJoinLayer* in Figure 6. This procedure, however, is not infallible. If d_h is the least distance from the joining host to a host in the highest layer, then it is possible for a host to join a cluster that is approximately d_h away from its nearest cluster leader. Therefore, we have built in refinement mechanisms, as described later in this section.

Procedure : *BasicJoinLayer*(h, i)
$Cl_j \leftarrow$ *Query*$(RP, -)$
$while\ (j > i)$
  Find y s.t. $dist(h, y) \leq dist(h, x), x, y \in Cl_j$
  $Cl_{j-1}(y) \leftarrow$ *Query*$(y, j-1)$
  Decrement j, $Cl_j \leftarrow Cl_{j-1}(y)$
$endwhile$
Join cluster Cl_j

Fig. 6. Basic join operation for member h, to join layer L_i. $i = 0$ for a new member. If $i > 0$, then h is already part of layer L_{i-1}. *Query*$(y, j-1)$ seeks the membership information of $Cl_{j-1}(y)$ from member y. *Query*$(RP, -)$ seeks the membership information of the topmost layer of the hierarchy, from the RP.

A.1 Join Latency

The joining process involves a message overhead of $O(k \log N)$ query-response pairs. The join-latency depends on the delays incurred in this exchanges, which is typically about $O(\log N)$ round-trip times. In our protocol, we aggressively locate possible "good" peers for a joining member, and the overhead for locating the appropriate attachments for any joining member is relatively large

To reduce the delay between a member joining the multicast group, and its receipt of the first data packet on the overlay, we allow joining members to temporarily peer, on the data path, with the leader of the cluster of the current layer it is querying. For example, in Figure 5, when A_{12} is querying the hosts B_0, B_1 and B_2 for the closest point of attachment, it temporarily peers with C_0 (leader of the layer L_1 cluster) on the data path. This allows the joining host to start receiving multicast data on the group within a single round-trip latency of its join.

A.2 Joining Higher Layers

An important invariant in the hierarchical arrangement of hosts is that the leader of a cluster be the center of the cluster. Therefore, as members join and leave clusters, the cluster-leader may occasionally change. When the leadership of a cluster, C, in layer L_j changes, the existing leader of C removes itself from all layers L_{j+1} and higher to which it is attached. The new leader of C first tries to join its current super-cluster (which by definition is a cluster in layer L_{j+1} to which the existing leader of C was joined to), i.e. the new leader replaces the existing leader in this cluster in layer L_{j+1}. However, if the super-cluster information is stale and currently invalid, then the new leader invokes the *BasicJoinLayer* procedure, which will terminate at layer L_{j+1} instead of L_0.

B. Cluster Maintenance and Refinement

Each member H of a cluster C, sends a *HeartBeat* message every h seconds to each of its cluster peers (neighbors on the control topology). The message contains the distance estimate of H to each other member of C. It is possible for H to have inaccurate or no estimate of the distance to some other members, e.g. immediately after it joins the cluster.

The cluster-leader includes the complete updated cluster membership in its *HeartBeat* messages to all other members. This allows existing members to set up appropriate peer relationships with new cluster members on the control path. For each cluster in level L_i, the cluster-leader also periodically sends the its immediate higher layer cluster membership (which is the super-cluster for all the other members of the cluster) to that L_i cluster.

All of the cluster member state is sent via unreliable messages and is kept by each cluster member as soft-state, refreshed by the periodic *HeartBeat* messages. A member H is declared no longer part of a cluster independently by all other members in the cluster if they do not receive a message from H for a configurable number of *HeartBeat* message intervals.

B.1 Cluster Split and Merge

A cluster-leader periodically checks the size of its cluster, and appropriately splits or merges the cluster when it detects a size bound violation. However, if a cluster that just exceeds the cluster size upper bound $2k-1$ is split, it creates two clusters of size k

each. Any single departure from these clusters will subsequently require a cluster merge operation to meet the size lower-bound. For this reason, we relax the size upper bound to be $3k$ and leave the lower bound unchanged. With this new upper bound, when a cluster is split into two equal parts, each of the parts is guaranteed to be at least $3k/2$, thus avoiding an immediate subsequent merge.

If the size of a cluster exceeds $3k$, the leader initiates a cluster split operation. Given a set of hosts and the pairwise distances between them, the cluster split operation partitions them into subsets that meet the size bounds, such that the maximum radius (in a graph-theoretic sense) of the new set of clusters is minimized. This is similar to the K-center problem (known to be NP-Hard) but with an additional size constraint[2]. We use an approximation strategy — the leader splits the current cluster into two equal-sized clusters, such that the maximum of the radii among the two clusters is minimized. It also chooses the centers of the two partitions to be the leaders of the new clusters and transfers leadership to the new leaders through *LeaderTransfer* messages. If these new clusters still violate the size upper bound, they are split by the new leaders using identical operations.

If the size of a cluster, $Cl_i(J)$ (in layer L_i) with leader J, falls below k, the leader J, initiates a cluster merge operation. Note, J itself belongs to a layer L_{i+1} cluster, $Cl_{i+1}(J)$. J chooses its closest cluster-peer, K, in $Cl_{i+1(J)}$. K is also the leader of a layer L_i cluster, $Cl_i(K)$. J initiates the merge operation of C_i with $Cl_i(K)$ by sending a *ClusterMergeRequest* message to K. J updates the members of $Cl_i(J)$ with this merge information. K similarly updates the members of $Cl_i(K)$. Following the merge, J removes itself from layer L_{i+1} (i.e. from cluster $Cl_{i+1}(J)$.

B.2 Refining Cluster Attachments

During a phase a period of rapid membership changes to the group, a joining member may not be able to locate its closest L_0 cluster, and therefore, attach to some other cluster. This may be true in some cases for higher layer members as well. Therefore, each member in any layer (say L_i) periodically probes all members in its super-cluster (they are the leaders of layer L_i clusters), to identify a closer cluster for itself in layer L_i. If a closer cluster is found, then the probing member leaves its current cluster and joins the closer cluster.

C. Host Departure and Leader Selection

When a host H leaves the multicast group, it sends a *Remove* message to all clusters to which it is joined. This is a graceful-leave. However, if H fails without being able to send out this message all cluster peers of H detects this departure through non-receipt of the periodic *HeartBeat* message from H. If H was a leader of a cluster, this triggers a new leader selection in the cluster. Each remaining member, J, of the cluster independently select a new leader of the cluster, depending on who J estimates to be the center among these members. Multiple leaders are re-conciled into a single leader of the cluster through exchange of regular *HeartBeat* messages using appropriate flags

[2]This is not always a 2-center problem, since we typically wait till a timeout to actually split a cluster, instead of splitting it exactly when it crosses the $3k$ for some efficiency reasons.

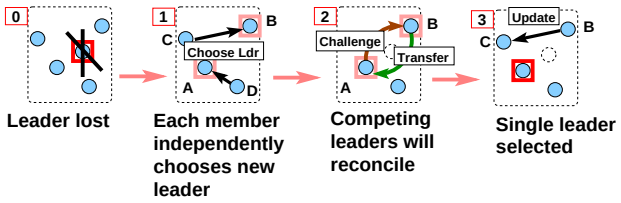

Fig. 7. Restructuring when a cluster-leader departs.

(*LeaderChallenge* and *LeaderTransfer*) each time two candidate leaders detect this multiplicity. This is shown in Figure 7.

It is possible for members to have an inconsistent view of the cluster membership, and for transient cycles to develop on the data path. These cycles are eliminated once the protocol reconciles the cluster view for all members, and restores the hierarchy invariants.

IV. Experimental Methodology

We have analyzed the performance of SMOP using detailed simulations and Internet-wide implementation. In the simulation environment, we compare the performance of SMOP to three other schemes: multi-unicast, native IP-multicast using the Core Based Tree protocol [2], and the Narada application-layer multicast protocol. In the Internet experiments, we benchmark the performance metrics against direct unicast paths to the member hosts.

Clearly, native IP multicast trees will have the least (unit) stress, since each link forwards only a single copy of each data packet. Unicast paths have the lowest latency [3] and so we consider them to be of unit stretch. Thy provide us a reference against which to compare the application-layer multicast protocols.

A. Data Model

In all these experiments, we model the scenario of a data stream source multicasting to the group. We chose a single end-host, uniformly at random, to be the data source generating a constant bit rate data. Each packet in the data sequence, effectively, samples the data path on the overlay topology at that time instant, and the entire data packet sequence captures the evolution of the data path over time.

B. Performance Metrics

We compare the performance of the different schemes along the following dimensions:

- *Quality of data path:* This is measured by three different metrics — tree degree distribution, stress on links and routers and stretch of data paths to the group members. In particular, apart from averages of these metrics, we also report their variances and distributions.
- *Recovery from host failure:* As hosts join and leave the multicast group, the underlying data delivery path adapts accordingly to reflect these changes. However, in transience, and particularly

[3]There are some recent studies [15], [1] to show that this may not always be the case; however, we use the native unicast latency as the reference to compare the performance of the other schemes.

after host failures, path to some hosts may be unavailable. It is also possible for multiple paths to exist to a single host and for cycles to develop temporarily. To observe these effects, we measured the fraction of hosts that correctly receive the data packets sent from the source. We also recorded the number of duplicates at each host. In all of our simulations, for both the application-layer multicast protocols, the number of duplicates was insignificant and zero in most cases.

• *Control traffic overhead:* We report the mean, variance and the distribution of the control bandwidth overheads at both routers and end hosts.

V. Simulation Experiments

We have implemented a packet-level simulator for the four different protocols. Our network topologies were generated using the Transit-Stub graph model, using the GT-ITM topology generator [4]. All topologies in these simulations had $10,000$ routers with an average node degree between 3 and 4. End-hosts were attached to a set of route, chosen at uniformly at random, from among the stub-domain nodes. The number of such hosts in the multicast group were varied between 8 and 2048 for different experiments. In our simulations, we only modeled loss-less links; thus, there is no data loss due to congestion, and no notion of background traffic or jitter. However, data is lost whenever the application-layer multicast protocol fails to provide a path from the source to a receiver, and duplicates are received whenever there is more than one path. Thus, our simulations study the dynamics of the multicast protocol and its effects on data distribution; in our implementation, the performance is also affected by other factors such as additional link latencies due to congestion and drops due to cross-traffic congestion.

As mentioned in Section I, the Narada protocol involves an aggregate control overhead of $O(N^2)$. In our simulation setup, we were unable to simulate Narada with groups of size 1024 or larger since the completion time for these simulations were on the order of a day for a single run of one experiment on a 550 MHz Pentium III machine with 4 GB of RAM.

A. Our implementation of Narada

We have implemented the entire Narada protocol from the description given in [8]. We did not implement the Narada high bandwidth extensions described in [7].As described before, Narada is a mesh-first application-layer multicast approach, designed primarily for small multicast groups. In Narada, the initial set of peer assignments to create the overlay topology is done randomly. While this initial data delivery path may be of "poor" quality, over time Narada adds "good" links and discards "bad" links from the overlay. Narada has $O(N^2)$ aggregate control overhead because of its mesh-first nature: it requires each host to periodically exchange updates and refreshes with all other hosts.

The protocol, as defined in [8], has a number of user-defined parameters that we needed to set. These include the link add/drop thresholds, link add/drop probe frequency, the periodic refresh rates, the mesh degree, etc. We experimented with a wide-range of values for these parameters to understand the behavior of Narada and observed some interesting trade-offs in choosing these parameters. Specifically, we found that:

• The mesh degree bound for hosts should not be strictly enforced to ensure connectivity. Instead additional mechanisms that limit the degree of the data path on the mesh should be used.

• There is a clear tradeoff between choosing a high versus low frequency for periodic probes to add or drop links on the mesh. A high frequency allows members to aggressively add and drop good and bad overlay links respectively. However, this leads to frequent changes to the data paths on the mesh, which can lead to a temporary loss of data path to other members. (This effect is different than when a route changes and state for the old route can be temporarily maintained to mitigate the effect of the route change). We observed this effect in our experiments where we use a high periodic probe frequency, especially if this parameter is set higher than the route packet exchange frequency. In contrast, using a low probe frequency leads to more stable paths; however, this implies that the mesh topology takes a long time to stabilize.

B. Simulation Results

In our experiments, we have simulated a wide-range of topologies, group sizes, member join-leave patterns, and protocol parameters. For SMOP, we set the cluster size parameter, k, to 3 in all of the experiments presented here. Broadly, our findings can be summarized as follows:

• SMOP trees have data paths that have stretch comparable to Narada

• The stress on links and routers are lower in SMOP, especially as the multicast group size increases

• The failure recovery of both the schemes are comparable.

• SMOP protocol demonstrates that it is possible to provide these performance with orders of magnitude lower control overhead for groups of size > 32.

We begin with results from a representative experiment that captures all the of different aspects comparing the various protocols.

B.1 Simulation Representative Scenario

This experiment has two different phases: a join phase and a leave phase. In the join phase a set of 128 members[4] join the multicast group uniformly at random between the simulated time 0 and 200 seconds. These hosts are allowed to stabilize into an appropriate overlay topology till simulation time 1000 seconds. The leave phase starts at time 1000 seconds: 16 hosts leave the multicast group over a short duration of 10 seconds. This is repeated four more times, at 100 second intervals. The remaining 48 members continue to be part of the multicast group till the end of simulation. All member departures are modeled as host failures since they have the most damaging effect on data paths. We experimented with different numbers of member departures, from a single member to 32 members leaving over the ten second window. Sixteen departures from a group of size 128 within a short time window is a drastic scenario, but it helps illustrate the failure recovery modes of the different protocols better. Member departures in smaller sizes cause correspondingly lower disruption on the data paths.

[4] We show results for the 128 member case because that is the group size used in the experiments reported in [8]; SMOP performs increasingly better with larger group sizes.

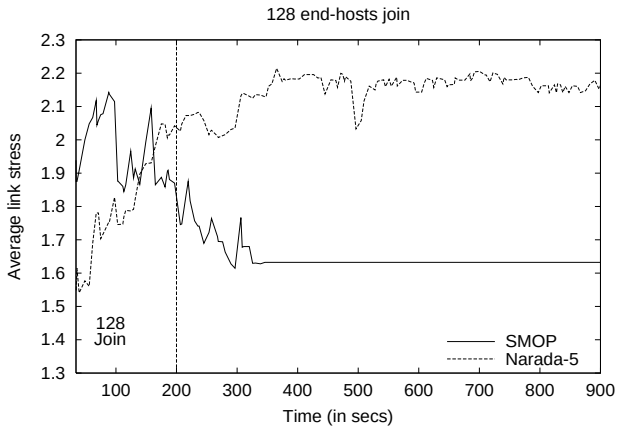

Fig. 8. Average link stress (simulation)

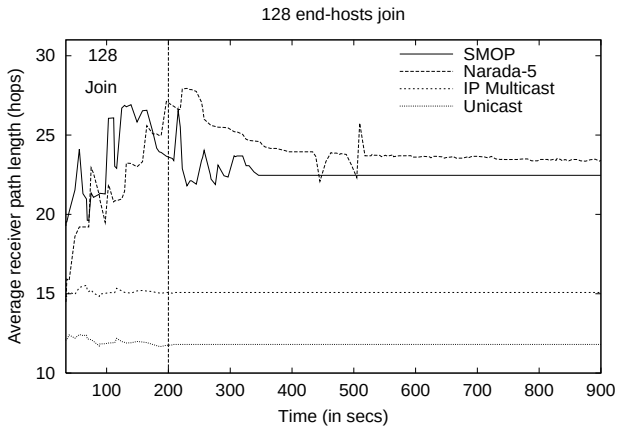

Fig. 9. Average path length (simulation)

We experimented with different periodic refresh rates for Narada. For a higher refresh rate the recovery from host failures is quicker, but at a cost of higher control traffic overhead. For Narada, we used different values for route update frequencies and periods for probing other mesh members to add or drop links on the overlay. In our results, we report results from using route update frequencies of once every 5 seconds (labeled Narada-5), and once every 30 seconds (labeled Narada-30). The 30 second update period corresponds to the what was used in [8]; we ran with the 5 second update period since the heartbeat period in SMOP was set to 5 seconds. Note that we could run with a much smaller heartbeat period in SMOP without significantly increasing control overhead since the control messages are limited within clusters and do not traverse the entire group. We also varied the mesh probe period in Narada and observed data path instability effect discussed above. In these results, we set the Narada mesh probe period to 20 seconds.

Data Path Quality

In Figures 8 and 9, we show the average link stress and the average path lengths for the different protocols as the data tree evolves during the member join phase. Note that the figure shows the actual path lengths to the end-hosts; the stretch is the ratio of average path length of the members of a protocol to the average path length of the members in the multi-unicast protocol.

As explained earlier, the join procedure in SMOP aggressively finds good points of attachment for the members in the overlay topology, and the SMOP tree converges quicker to a stable value (within 350 seconds of simulated time). In contrast, the Narada protocols gradually improve the mesh quality, and consequently so does the data path over a longer duration. Its average data path length converges to a stable value of about 23 hops between 500 and 600 seconds of the simulated time. The corresponding stretch is about 2.18. In Narada path lengths improve over time due to addition of "good" links on the mesh. At the same time, the stress on the tree gradually increases since the Narada decides to add or drop overlay links based purely on the stretch metric.

The cluster-based data dissemination in SMOP reduces average link stress, and in general, for large groups SMOP converges to trees with about 25% lower average stress. In this experiment, the SMOP tree had lower stretch than the Narada tree; however, in other experiments the Narada tree had a slightly lower stretch value. In general, comparing the results from multiple experiments over different group sizes, (See Section V-B.2), we concluded that the data path lengths to receivers were similar for both protocols.

In Figures 10 and 11, we plot a cumulative distribution of the stress and path length metrics for the entire member set (128 members) at a time after the data paths have converged to a stable operating point.

The distribution of stress on links for the multi-unicast scheme has a significantly large tail (e.g. links close to the source has a stress of 127). This should be contrasted with better stress distribution for both SMOP and Narada. Narada uses fewer number of links on the topology than SMOP, since it is comparably more aggressive in adding overlay links with shorter lengths to the mesh topology. However, due to this emphasis on shorter path lengths, the stress distribution of the links has a heavier-tail than SMOP. More than 25% of the links have a stress of four and higher in Narada, compared to $< 5\%$ in SMOP. The distribution of the path lengths for the two protocols are comparable. The multi-unicast scheme (presented for comparison) shows the shortest path length distribution if stress on links is ignored.

Failure Recovery and Control Overheads

To investigate the effect of host failures, we present results from the second part of our scenario: starting at simulated time 1000 seconds, a set of 16 members leave the group over a 10 second period. We repeat this procedure four more times and no members leave after simulated time 1400 seconds when the group is reduced to 48 members. When members leave, both protocols "heal" the data distribution tree and continue to send data on the partially connected topology. In Figure 12, we show the fraction of members that correctly receive the data packets over this duration. Both Narada-5 and SMOP have similar performance, and on average, both protocols restore the data path to all (remaining) receivers within 30 seconds, and on correctly serve over 90% of the members. We also ran the same experiment with the 30 second refresh period for Narada. The lower

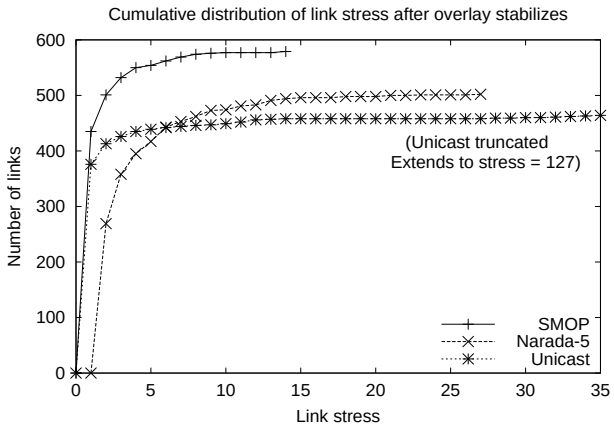

Fig. 10. Stress distribution (simulation)

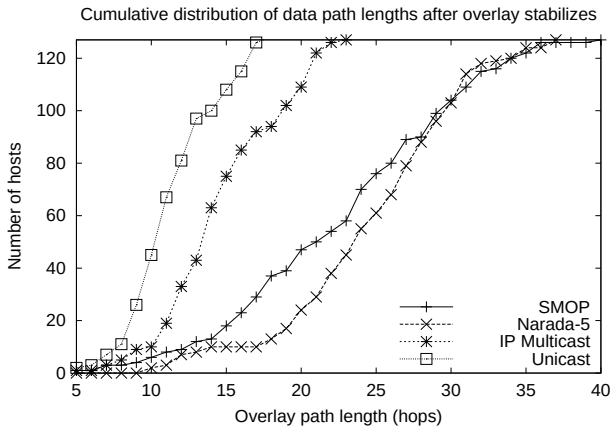

Fig. 11. Path length distribution (simulation)

refresh period caused significant disruptions on the tree with periods of over 100 seconds when more than 60% of the tree did not receive any data. Lastly, we note that the data distribution tree used for SMOP is the least connected topology possible; we expect failure recovery results to be much better if structures with alternate paths are built atop SMOP.

In Figure 13, we show the byte-overheads for control traffic at the access links of the end-hosts. Each dot in the plot represents the sum of the control traffic (in Kbps) sent or received by each member in the group, averaged over 10 second intervals. Thus for each 10 second time slot, there are two dots in the plot for each (remaining) host in the multicast group corresponding to the control overheads for Narada and SMOP. The curves in the plot are the average control overhead for each protocol. As can be expected, for groups of size 128, SMOP has an order of magnitude lower average overhead, e.g. at simulation time 1000 seconds, the average control overhead for SMOP is 0.97 Kbps versus 62.05 Kbps for Narada. At the same time instant, Narada-30 (not shown in the figure) had an average control overhead of 13.43 Kbps. Note that the SMOP control traffic includes all protocol messages, including messages for cluster formation, cluster splits, merges, layer promotions, and leader elections.

B.2 Aggregate Results

We present a set of aggregate results as the group size is varied. The purpose of this experiment is to understand the scalability of the different application-layer multicast protocols. The entire set of members join in the first 200 seconds, and then we run the simulation for 1800 seconds to allow the topologies to stabilize. In Table I, we compare the stress on network routers and links, the overlay path lengths to group members and the average control traffic overheads at the network routers. For each metric, we present the both mean and the standard deviation.

As we showed in our first experiment, Narada and SMOP tend to converge to trees with similar path lengths. The stress metric for both network links and routers, however, is consistently lower for SMOP when the group size is large (64 and greater). It is interesting to observe the standard deviation of stress as it changes with increasing group size for the two protocols. The standard deviation for stress increased for Narada for increasing group sizes. In contrast, the standard deviation of stress for SMOP remains relatively constant; the topolgy-based clustering in SMOP distributes the data path more evenly among the different links on the underlying links regardless of group size.

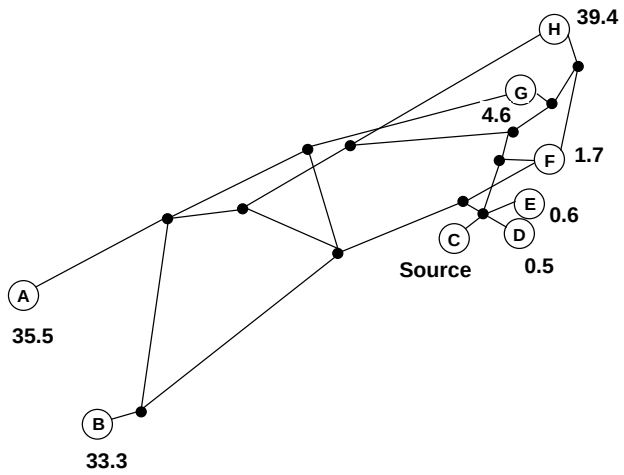

Fig. 14. Internet experiment sites and direct unicast latencies from C

The control overhead numbers in the Table are different than the ones in Figure 13; the column in the table is the average control traffic *per network router* as opposed to control traffic at an end-host. Since the control traffic gets aggregated inside the network, the overhead at routers is significantly higher than the overhead at an end-host. For these router overheads, we report the values of the Narada-30 version in which the route update frequency set to 30 seconds. Recall that this protocol, Narada-30 performs relatively poorly when members leave, but is much more efficient (specifically 5 *times* less overhead with groups of size 128) than the Narada-5 version. The refresh messages in SMOP were still sent at 5 second intervals.

VI. Internet-wide Implementation

We have implemented the complete SMOP protocol and experimented with our implementation a one-month period, with

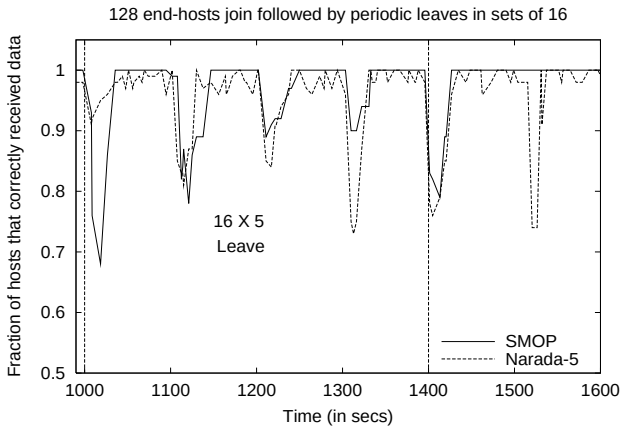

Fig. 12. Fraction of members that received data packets over the duration of member failures. (simulation)

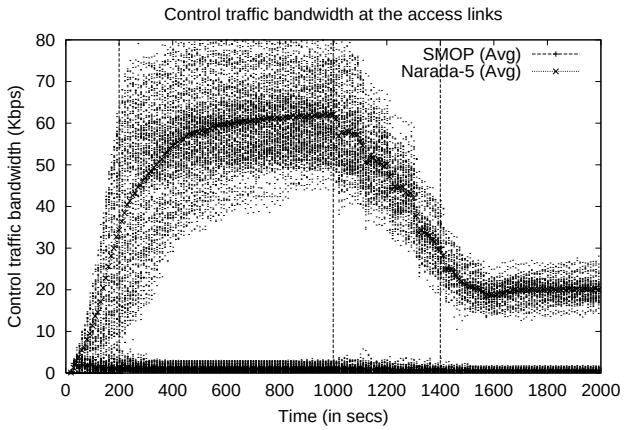

Fig. 13. Control bandwidth required at end-host access links (simulation)

Group	Router Stress		Link Stress		Path Length		Bandwidth Overheads (Kbps)	
Size	Narada-5	SMOP	Narada-5	SMOP	Narada-5	SMOP	Narada-30	SMOP
8	1.55 (1.30)	3.51 (3.30)	1.19 (0.39)	3.24 (2.90)	25.14 (9.49)	12.14 (2.29)	0.61 (0.55)	1.54 (1.34)
16	1.84 (1.28)	2.34 (2.16)	1.34 (0.76)	1.86 (1.39)	19.00 (7.01)	20.33 (6.75)	2.94 (2.81)	0.87 (0.81)
32	2.13 (2.17)	2.42 (2.60)	1.54 (1.03)	1.90 (1.82)	20.42 (6.00)	17.23 (5.25)	9.23 (8.95)	1.03 (0.95)
64	2.68 (3.09)	2.23 (2.25)	1.74 (1.53)	1.63 (1.39)	22.76 (5.71)	20.62 (7.40)	26.20 (28.86)	1.20 (1.15)
128	3.04 (4.03)	2.36 (2.73)	2.06 (2.64)	1.63 (1.56)	21.55 (6.03)	21.61 (7.75)	65.62 (92.08)	1.19 (1.29)
256	3.63 (7.52)	2.31 (3.18)	2.16 (3.02)	1.63 (1.63)	23.42 (6.17)	24.67 (7.45)	96.18 (194.00)	1.39 (1.76)
512	4.09 (10.74)	2.34 (3.49)	2.57 (5.02)	1.62 (1.54)	24.74 (6.00)	22.63 (6.78)	199.96 (55.06)	1.93 (3.35)
1024	-	2.59 (4.45)	-	1.77 (1.77)	-	25.83 (6.13)	-	2.81 (7.22)
1560	-	2.83 (5.11)	-	1.88 (1.90)	-	24.99 (6.96)	-	3.28 (9.58)
2048	-	2.92 (5.62)	-	1.93 (1.99)	-	24.08 (5.36)	-	5.18 (18.55)

TABLE I

DATA PATH QUALITY AND CONTROL OVERHEADS FOR VARYING MULTICAST GROUP SIZES (SIMULATION)

32 to 100 member groups distributed across 8 different sites. Our experimental topology is shown in Figure 14 (labeled A to H). Unfortunately, experiments with much larger groups were not feasible on our testbed. However, our implementation results for protocol overheads closely match our simulation experiments, and we believe our simulations do provide a reasonable indication of how the SMOP implementation would behave with larger group sizes.

A. Implementation Specifics

We have conducted experiments with data sources at different sites. In this paper, we present a representative set of the experiments where the data stream source is located at site C in Figure 14. In the figure, we also indicate the typical direct unicast latency (in milliseconds) from the site C, to all the other sites. These are estimated one-way latencies obtained using a sequence of application layer (UDP) probes. Data streams were sent from the source host at site C, to all other hosts, using the SMOP overlay topology. For our implementation, we experimented with different *HeartBeat* rates; in the results presented in this section, we set the *HeartBeat* message period to 10 seconds.

In our implementation, we had to estimate the end-to-end latency between hosts for various protocol operations, including member joins, leadership changes, etc. We estimated the latency between two end-hosts using a low-overhead estimator that sent a sequence of application-layer (UDP) probes. We controlled the number of probes adaptively using observed variance in the latency estimates. Further, instead of using the raw latency estimates as the distance metric, we used a simple binning scheme to map the raw latencies to a small set of equivalence classes. Specifically, two latency estimates were considered equivalent if they mapped to the same equivalence class, and this resulted in faster convergence of the overlay topology. The specific latency ranges for each class were 0-1 ms, 1-5 ms, 5-10 ms, 10-20 ms, 20-40 ms, 40-100 ms, 100-200 ms and greater than 200 ms.

To compute the stretch for end-hosts in the Internet experiments, we used the ratio of the latency from between the source and a host along the overlay to the direct unicast latency to that host. Specifically, during the experiments, when a host A receives a data packet forwarded by member B along the overlay tree, A immediately sends back a overlay-hop acknowledgment back to B. B logs the round-trip latency between its initial transmission of the data packet to A and the receipt of the acknowlegdment from A. After the entire experiment is done, we sum the overlay round-trip latencies for each data packet by referring back the logs at each host. We estimate the one-way overlay latency as half of this round trip latency. We obtain the unicast latencies immediately after the experiment terminates. This is clearly not ideal for long running experiments; however, unfortunately the concurrent computation of the unicast latencies perturbed the experimental data and we had to resort to computing the unicast latencies after the experiment completed.

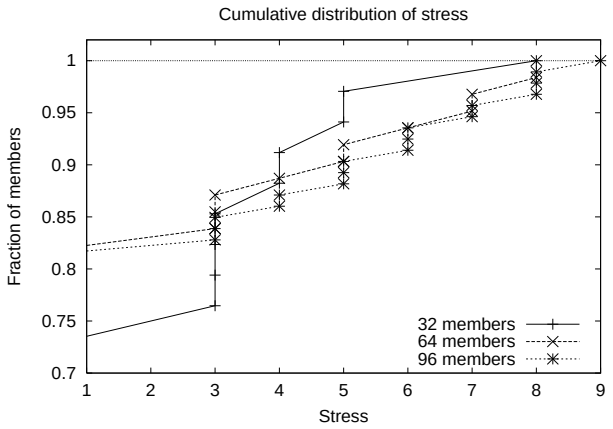

Fig. 15. Stress distribution (testbed)

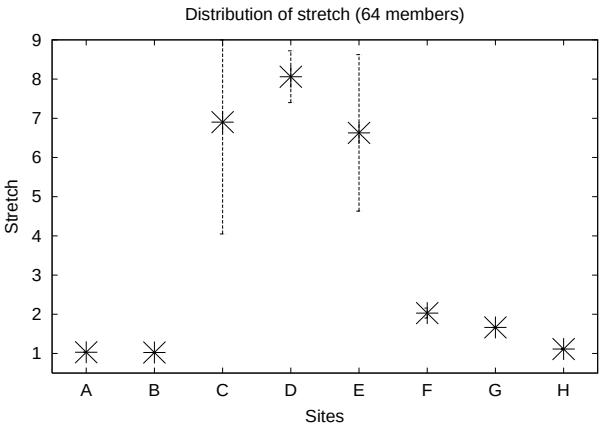

Fig. 16. Stretch distribution (testbed)

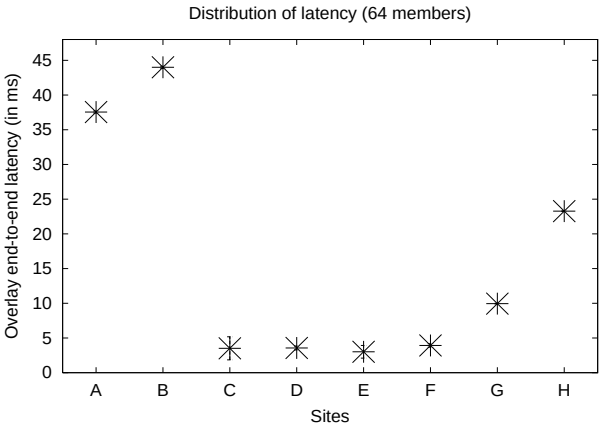

Fig. 17. Latency distribution (testbed)

B. Implementation Scenarios

The Internet experiment scenarios have two phases: a join phase and a rapid membership change phase. In the join phase, a set of member hosts randomly join the group from the different sites. The hosts are then allowed to stabilize into an appropriate overlay delivery tree. After this period, the rapid membership change phase starts, where host members randomly join and leave the group. The average member lifetime in the group, in this phase was set to 30 seconds. Like in the simulation studies, all member departures are ungraceful and allow us to study the worst case protocol behavior. Finally, we let the remaining set of members to organize into a stable data delivery tree. We present results for three different groups of size of 32, 64, and 96 members.

Data Path Quality

In Figure 15, we show the cumulative distribution of the stress metric at the group members after the overlay stabilizes at the end of the join phase. For all group sizes, typical members have unit stress (74% to 83% of the members in these experiments). The stress for the remaining members vary between 3 and 9. These members are precisely the cluster leaders in the different layers (recall that the cluster size lower and upper bounds for these experiments is 3 and 9, respectively). The stress for these members can be reduced further by using the high-bandwidth data path enhancements, described in the Appendix. For larger groups, the number of members with higher stress (i.e. between 3 and 9 in these experiments) is more, since the number of clusters (and hence, the number of cluster leaders) is more. However, as expected, this increase is only logarithmic in the group size.

In Figure 16, we plot the cumulative distribution of the stretch metric. Instead of plotting the stretch value for each single host, we group them by the sites at which there are located.For all the member hosts at a given site, we plot the mean and the 95% confidence intervals. Apart from the sites C, D, and E, all the sites have near unit stretch. However, note that the source of the data streams in these experiments were located in site C and hosts in the sites C, D, and E had very low latency paths from the source host. The actual end-to-end latencies *along the overlay paths* to all the sites are shown in Figure 17. For the sites C, D and E these latencies were 3.5 ms, 3.5 ms and 3.0 ms respectively. Therefore, the primary contribution to these latencies are packet processing and overlay forwarding on the end-hosts themselves.

In Table II, we present the mean and the maximum stretch for the different members, that had direct unicast latency of at least 2 ms from the source (i.e. sites A, B, G and H), for all the different sizes. The mean stretch for all these sites are low, in some cases we do see relatively large worst case stretches (e.g. 4.63 stretch for the worst case for a member in a group of size 96).

Failure Recovery

In these Internet experiments we observe the effects of group membership changes on the data delivery tree. To do this, we record how successful the overlay is in delivering data during changes to the overlay topology. This is the rapid membership change phase of the experiment, which begins after the initial member set has stabilized into the appropriate overlay topology. In this phase the average lifetime of a member in the group is 30 seconds, over a 15 minute period.

In Figure 18 we plot over time the fraction of members that successfully received the different data packets. A total of 30 group membership changes happened over the duration. In Figure 19 we plot the cumulative distribution of losses seen by the

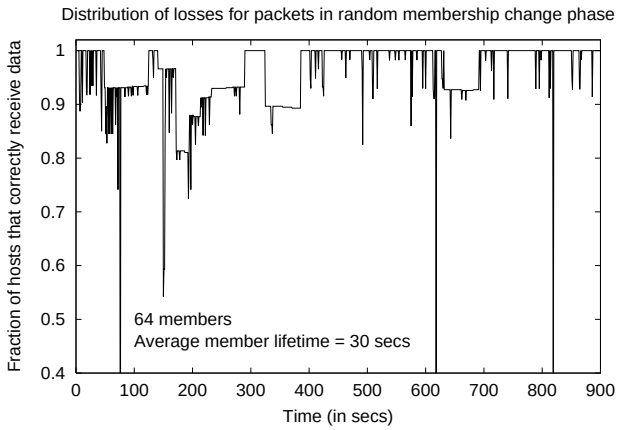

Fig. 18. Fraction of members that received data packets as group membership continuously changed (testbed)

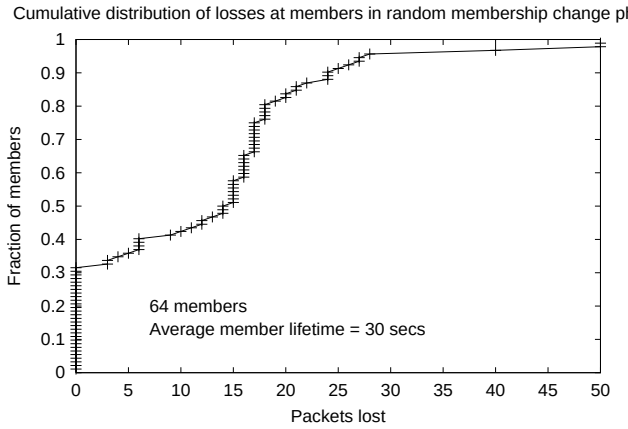

Fig. 19. Cumulative distribution of total number of lost packets for different members out of the entire sequence of 900 packets during the rapid membership change phase (testbed)

different members over the entire 15 minute duration. The maximum number of losses seen by a member was 50, and 30% of the members did not encounter any losses. Even under this rapid changes to the group membership, the largest continuous duration of packet losses for any single host was 34 seconds, while typical members experienced a maximum continuous data loss for only two seconds — this was true for all but 4 of the members. These failure recovery statistics are good enough for use in most data stream applications deployed over the Internet. Note that in this experiment, only three individual packets (out of 900) suffered heavy losses: data packets at times 76 s, 620 s, and 819 s were not received by 51, 36 and 31 members respectively.

Control Overheads

Finally, we present the control traffic overheads (in Kbps) in Table II for the different group sizes. The overheads include control packets that were sent as well as received. We show the average and maximum control overhead at any member. We observed that the control traffic at most members lies between 0.2 Kbps to 2.0 Kbps for the different group sizes. In fact, about 80% of the members require less than 0.9 Kbps of control traffic for topology management. More interestingly, the average control overheads and the distributions do not change significantly as the group size is varied. The worst case control overhead is also fairly low (less than 3 Kbps).

Group	Stress		Stretch		Control overheads (Kbps)	
Size	Mean	Max.	Mean	Max.	Mean	Max.
32	1.85	8.0	1.08	1.61	0.84	2.34
64	1.73	8.0	1.14	1.67	0.77	2.70
96	1.86	9.0	1.04	4.63	0.73	2.65

TABLE II

AVERAGE AND MAXIMUM VALUES OF OF THE DIFFERENT METRICS FOR DIFFERENT GROUP SIZES (TESTBED)

VII. RELATED WORK

There are a few closely related projects which explore implementing multicast at the application layer. They can be classified into two broad categories: mesh-first (Narada [8], Gossamer [5]) and tree-first protocols (Yoid [10], ALMI [13]). Yoid defines a distributed tree building protocol between the end-hosts, while ALMI uses a centralized algorithm to create a minimum spanning tree rooted at a designated single source of multicast data distribution. The JungleMonkey project[5] is a similar effort to create an application layer multicast overlay using a tree-first approach.

Bayeux [17] in another architecture for application layer multicast, where the end-hosts are organized into a hierarchy as defined by the Tapestry overlay location and routing system [16]. A level of the hierarchy is defined by a set of hosts that share a common suffix in their host IDs. Such a technique was proposed by Plaxton et.al. [14] for locating and routing to named objects in a network.

The Overcast [12] protocol organizes a set of similar proxies (called Overcast nodes) into a distribution tree rooted at a central source for single source multicast. A distributed tree-building protocol is used to create this source specific tree, in a manner similar to Yoid. RMX [6] provides support for reliable multicast data delivery to end-hosts using a set of such proxies, called Reliable Multicast proXies. Application end-hosts are configured to affiliate themselves with the nearest RMX. The architecture assumes the existence of an overlay construction protocol, using which these proxies organize themselves into an appropriate data delivery path. TCP is used to provide reliable communication between each pair of peer proxies on the overlay.

VIII. CONCLUSIONS

In this paper, we have presented a new protocol for application-layer multicast. Our main contribution is an extremely low overhead hierarchical control structure over which different data distribution paths can be built. Our results show that it is possible to build and maintain application-layer multicast trees with very little overhead. While the focus of this paper has been low-bandwidth data stream applications, our scheme is generalizable to different applications by appropriately choosing data paths and metrics used to construct the overlays. We believe that the results of this paper are a significant first step towards

[5] http://www.junglemonkey.net

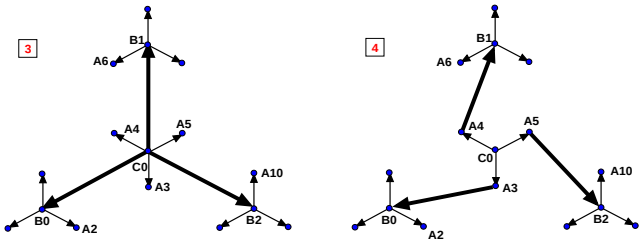

Fig. 20. Data path enhancements using delegation.

constructing large wide-area applications over application-layer multicast.

REFERENCES

[1] D.G. Andersen, H. Balakrishnan, M. Frans Kaashoek, and R. Morris. Resilient overlay networks. In *Proceedings of 18th ACM Symposium on Operating Systems Principles*, October 2001.

[2] T. Ballardie, P. Francis, and J. Crowcroft. Core Based Trees (CBT): An Architecture for Scalable Multicast Routing. In *Proceedings of ACM Sigcomm*, 1995.

[3] S. Banerjee and B. Bhattacharjee. Scalable Secure Group Communication over IP Mulitcast. In *Proceedings of Internation Conference on Network Protocols*, November 2001.

[4] K. Calvert, E. Zegura, and S. Bhattacharjee. How to Model an Internetwork. In *Proceedings of IEEE Infocom*, 1996.

[5] Y. Chawathe. Scattercast: An Architecture for Internet Broadcast Distribution as an Infrastructure Service. *Ph.D. Thesis, University of California, Berkeley*, December 2000.

[6] Y. Chawathe, S. McCanne, and E. A. Brewer. RMX: Reliable Multicast for Heterogeneous Networks. In *Proceedings of Infocom*, 2000.

[7] Y.-H. Chu, S. G. Rao, S. Seshan, and H. Zhang. Enabling Conferencing Applications on the Internet using an Overlay Multicast Architecture. In *Proceedings of ACM SIGCOMM*, August 2001.

[8] Y.-H. Chu, S. G. Rao, and H. Zhang. A Case for End System Multicast. In *Proceedings of ACM SIGMETRICS*, June 2000.

[9] S. Deering and D. Cheriton. Multicast Routing in Datagram Internetworks and Extended LANs. In *ACM Transactions on Computer Systems*, May 1990.

[10] P. Francis. Yoid: Extending the Multicast Internet Architecture, 1999. White paper http://www.aciri.org/yoid/.

[11] A. Gupta. Steiner points in tree metrics don't (really) help. In *Symposium of Discrete Algorithms*, January 2001.

[12] J. Jannotti, D. Gifford, K. Johnson, M. Kaashoek, and J. O'Toole. Overcast: Reliable Multicasting with an Overlay Network. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation*, 2000.

[13] D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel. ALMI: An Application Level Multicast Infrastructure. In *Proceedings of 3rd Usenix Symposium on Internet Technologies & Systems*, March 2001.

[14] C. G. Plaxton, R. Rajaraman, and A. W. Richa. Accessing nearby copies of replicated objects in a distributed environment. In *ACM Symposium on Parallel Algorithms and Architectures*, 1997.

[15] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collin s, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan. Detour: A Case for Informed Internet Routing and Transport. In *IEEE Micro*, 1999.

[16] B. Y. Zhao, J. Kubiatowicz, and A. Joseph. Tapestry: An Infrastructure for Fault-tolerant Wide-area Location and Routing. In *Tech. Report UCB/CSD-01-1141, University of California, Berkeley*, April 2001.

[17] S. Q. Zhuang, B. Y. Zhao, A. D. Joseph, R. Katz, and J. Kubiatowicz. Bayeux: An architecture for scalable and fault-tolerant wide-area dat a dissemination. In *Eleventh International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV 2001)*, 2001.

APPENDIX

I. DATA PATH ENHANCEMENTS FOR HIGH BANDWIDTH APPLICATIONS

The basic data path in SMOP imposes more data forwarding responsibility onto the cluster leaders. As a consequence, members that are joined to higher layers are cluster leaders in all the lower layers that they are joined to. Therefore, they are required to forward higher volume of data than those members that are joined to only the lower layers. This data forwarding path, is therefore, not suited for high bandwidth applications (e.g. video distribution). We define an enhancement to this basic data path by allowing the cluster leaders to *delegate* data forwarding responsibility to some of its cluster members in a deterministic manner. In this paper, we only explain data path delegation, assuming that data is originating from the leader of the highest cluster in the topology. However, the same delegation mechanism is equally applicable for data originating from any member (with minor modifications).

Consider a host h that belongs to layers, $L_0, \ldots, L_i$ and no other higher layer. The corresponding clusters in these layers are: $Cl_0(h), \ldots, Cl_i(h)$. In the basic data path (described in Section II-B), h receives the data from the leader, p, of cluster $Cl_i(h)$, i.e. its topmost layer. It is also responsible for forwarding data to all the members in the clusters $Cl_0(h), \ldots, Cl_{i-1}(h)$, i.e. the clusters in the remaining layers.

In the enhanced data path, h forwards data to only the other members of $Cl_0(h)$, i.e. its cluster in the lowest layer (L_0). Additionally, it delegates the responsibility of forwarding data to members in $Cl_j(h)$ to members in $Cl_{j-1}(h)$, for all $1 \leq j \leq i-1$. Since the cluster sizes are bounded between k and $2k-1$, each member of $Cl_{j-1}(h)$ gets a *delegated* forwarding responsibility to at most two members in $Cl_j(h)$. Only the cluster leader can delegate forwarding responsibility to another member of its cluster. A member that belongs to multiple layers belongs to only a single cluster in each layer and is the leader of the respective clusters in all but one layer. It is not a leader in its topmost cluster. Therefore, each member can be delegated forwarding responsibilities for at most 2 new peers. The member, h receives data from a member, q to which p (the leader of its topmost cluster, $Cl_i(h)$) has delegated the forwarding responsibility.

This data path transformation is illustrated with an example in Figure 20. Consider the basic data path with host C_0 as the source (Panel 3). Host C_0 is the leader of both its L_0 and L_1 clusters. Therefore, in the basic data path, it is required to forward data to the other members both its clusters (A_3, A_4 and A_5 in layer L_0, and B_0, B_1 and B_2 in layer L_1). In the enhanced data path (Panel 4), it delegates the other members of its L_0 cluster to forward data to the other members of its L_1 cluster. In particular, it sets up the new data path peers as: $\langle A_3 \rightarrow B_0 \rangle$, $\langle A_4 \rightarrow B_1 \rangle$, and $\langle A_5 \rightarrow B_2 \rangle$. Members which are not leaders of their L_1 clusters, i.e. B_0, B_1 and B_2 now receive data not from the cluster leader (i.e. C_0), and instead receive data from the members delegated by C_0 as described.

Any member, in the enhanced data path, forwards data to all members of only one of its clusters (i.e. its L_0 cluster), and additionally may be delegated to forward data to two other members. This the total number of data path peers for any member in this enhanced data path is bounded by $2k$, a constant that depends on the cluster size.

A member can further reduce its data path peers organizing its intra-cluster data path using structures other than a star, (e.g. using a ring path in the cluster implies that a member has at most 3 data path peers). This is, of course, at the cost of increased end-to-end latencies.