Encodings of Multicast Trees

Vijay Arya¹, Thierry Turletti¹, and Shivkumar Kalyanaraman²

¹ INRIA Sophia Antipolis, France

² Rensselaer Polytechnic Institute, NY USA

{vijay.arya, thierry.turletti}@sophia.inria.fr,
shivkuma@ecse.rpi.edu

Abstract. In this paper, we present efficient ways of encoding multicast trees. Multicast tree encodings provide a convenient way of performing stateless and explicit multicast routing in networks and overlays. We show the correspondence of multicast trees to theoretical tree data structures and give lower bounds on the number of bits needed to represent multicast trees. Our encodings can be used to represent multicast trees using both node identifiers and link indexes and are based on balanced parentheses representation of tree data structures. These encodings are almost space optimal and can be read and processed efficiently. We evaluate the length of these encodings on multicast trees in generated and real topologies.

1 Introduction

Current IP multicast routing protocols such as DVMRP [19] and PIM [8] install a per group state in the routers. To support these protocols, a router needs to maintain a multicast forwarding table entry corresponding to every multicast group whose distribution tree passes through the router. Since entries corresponding to different multicast groups cannot be aggregated, the forwarding table size grows with the number of active multicast groups, violating the stateless principle followed by routers. As the number of multicast groups increase, the cost and performance of routers are adversely affected. The state maintenance problem is further aggravated by the fact that even small multicast groups can introduce significant amount of state into the network. The amount of state introduced by a multicast group in the network depends on the size of its tree. The size of the multicast tree in turn is a function of both the relative locations of the source and various receivers, and the number of receivers itself. A small multicast group consisting of receivers which are spread uniformly in the Internet may introduce more state than a large multicast group concentrated in one region of the Internet. Thus, even a large number of small multicast groups may cause a state explosion. In order to reduce the state overhead at routers, one of the approaches followed is to move the state from the network to the packet, i.e., to use a representation of the multicast tree within every data packet. Each router can then function in a stateless manner and perform multicast forwarding

by reading and processing packet headers. To use this approach even in small to medium sized multicast groups, a compact encoding of multicast tree is required.

Due to deployment problems in Inter-domain multicast, several application level multicast protocols have been proposed [16, 21]. However, the problem of per group state maintenance remains even in application level multicast. Scribe [16] and Bayeux [21] are application level multicast protocols for Pastry and Tapestry overlays respectively. Pastry and Tapestry are large scale DHT overlays which allow interconnection and routing between millions of nodes in a scalable manner. To implement multicast in these overlays, both Scribe and Bayeux introduce a per group multicast state within overlay nodes which are part of the multicast distribution tree. Due to constraints of per group state maintenance, these overlays cannot support a large number of multicast groups (small or big). State maintenance is a serious issue in overlay nodes since they are constrained by memory and processing power (normally PCs/servers).

Reduction of per group state is just one of the several motivations for encoding Multicast trees. In Distributed Interactive Applications (DIAs) such as large scale virtual environments (e.g. multiplayer games) and distributed interactive simulations, participants act as senders and receivers in several multicast groups. Minimizing the control overhead due to group dynamics is a major design consideration here [10, 15]. In these applications, nodes are clustered according to communication needs and multicast distribution trees are constructed within the clusters. As nodes join and leave, they send control messages to each other to repair the multicast distribution trees. On average, the number of messages is linear to the number of nodes in the tree. Since there are a large number of multicast groups with frequent joins and leaves, significant control messages flow in the network. If the multicast trees are encoded within data packets, control overhead is reduced and on-the-fly tree repair is made possible at forwarding nodes [13].

Multicast tree encodings also provide a convenient way of choosing explicit multicast trees in networks and overlays. In overlay networks such as RON [1], OMNI [17], and Akamai, by embedding multicast trees within data packets, multicast traffic can be routed on specific links to perform load balancing or traffic engineering. Multicast tree encodings can also be used to communicate group keys needed for secure group communication [20] to specific receivers.

The above motivations notwithstanding, very little is known about optimal encodings of multicast trees. Indeed, there are two proposals which encode multicast trees non optimally [11, 13]. In this paper, we consider multicast trees which could occur at network or overlay level and show the correspondence of these trees to theoretical tree data structures. Using this correspondence, we give lower bounds on the number of bits needed to encode multicast trees and show efficient ways of encoding them.

Multicast tree encoding must satisfy two requirements. Firstly, the encoding must be of minimal length since the it is passed in *every* multicast data packet and processed by each forwarding node (router or overlay node). Secondly, the forwarding operation must consume minimal time. Each forwarding node pro-

Fig. 1. Link Index concept: A router with its link indexes

cesses the encoding to determine the list of downstream nodes to whom the multicast packet must be forwarded. To minimize both the end-to-end packet delay and the per packet processing load on a forwarding node, the encoding must support the execution of this task efficiently.

The networks in which we consider multicast trees can be divided into two categories: (i) Networks which can *potentially* support forwarding based on *link Indexes*. The link index of a node with value i refers to its i th link. For example, the link index of a router refers to one of its interfaces. A router can potentially order its link interfaces and refer to each interface by its index in this ordering (Fig 1). The link index of an overlay node refers to one of its neighbors. Using a link index, a node can be instructed to forward a packet to one of its neighbors. This operation can potentially be supported in IP networks, static overlays such as OMNI, RON, Akamai, partially upgraded IP networks such as Bananas [9], etc where neighbors of a node do not change frequently. (In Bananas, nodes do forward unicast packets based on link indexes to support multipath-routing). (ii) In dynamic peer to peer networks (e.g. DHTs, unstructured p2ps) both the degree and the neighbors of nodes change frequently as nodes join and leave and hence forwarding based on link indexes cannot be supported. In these networks, node identifiers (e.g. IP addresses) of nodes need to be used to encode the multicast tree.

We show encoding of multicast trees using both link indexes and node identifiers. The tree encodings we propose are independent of underlying tree construction protocols and are based on *balanced parentheses representation of ordinal trees*. The rest of the paper is organized as follows. Section 2 places related work in context, Section 3 shows correspondence of multicast trees to tree data structures and Section 4 presents various tree encodings. Section 5 evaluates the length of these encodings on real and generated multicast trees and Section 6 concludes.

2 Related Work

The encoding of a multicast tree depends on whether it represents the *default* multicast tree supplied by unicast routing framework of the underlying network or whether it represents a *specific or explicit* tree. To represent a specific tree, the multicast tree must be encoded using either the node identifiers or link indexes of the entire tree. To choose the default tree, identities of receivers are sufficient. Default trees can be chosen in any network which supports unicast

routing and their main application is multicast state reduction. On the other hand, explicitly representing trees provides the option of routing multicast traffic as desired which may be useful to create a multicast routing framework at application level, perform traffic engineering, or multicast state reduction.

In Xcast [7], a multicast tree is encoded by listing the IP addresses of receivers in the multicast tree. In IP networks, a packet can be routed to a destination node given its IP address. Due to this reason, explicitly listing the IP addresses of receivers of a multicast group suffices to represent the default multicast tree that exists between the source and the receivers. Forwarding routers perform a routing table lookup for each IP address in the packet and group them based on the common outgoing interface. The multicast packet is then forwarded on the interfaces, each packet containing the corresponding subset of IP addresses. Alternatively, if the last hop routers are allowed to store multicast state, the tree can be represented by listing the IP addresses of last hop routers instead of receivers (Xcast+). Xcast has two limitations. Since each IP address consumes 32 bits, Xcast encoding is suitable for multicast groups containing a small number of last hop routers (6-10 last hop routers need 24-40 bytes in IPv4). Secondly, each forwarding router needs to perform k IP lookups, where k is the number of last hop routers in its subtree. Large encodings require more IP lookups which are expensive.

The Xcast model is also suited to peer to peer DHTs such as Pastry and Tapestry, where packets are routed in the overlay using node identifiers. In Pastry, a node is represented using a 128 bit identifier. By representing multicast trees using receiver identifiers, the state overhead imposed by small multicast groups can be avoided in Scribe and Bayeux. The Xcast model is advantageous in DHTs since they guarantee routing in the presence of joins and leaves by intermediate peers.

In Linkcast [11], a multicast tree is encoded using the link indexes of links in the multicast tree. Each forwarding node reads *its* corresponding link indexes and forwards the multicast packet. The multicast tree which is encoded is a reverse path multicast tree constructed using join messages sent from last hop routers towards the source. Each router receiving the join message is expected to forward its identity along with the link index of the interface which received the join message. Thus the source can encode the multicast tree using link indexes. For forwarding, a pointer in the encoding points to the location of the current router in the encoding. Each router uses the pointer to read its outgoing link indexes and updates the pointer for the next hop routers. Although the number of links in the multicast tree is larger than the number of receivers, link indexes can be represented using fewer bits as compared to 32 bits for the IP address of a node. Also, expensive routing table lookups are avoided in Linkcast. We propose a new encoding for Linkcast which requires less space and two new encodings of multicast trees using link indexes.

In [13], authors propose two tree encodings to represent application level multicast trees occurring within clusters of DIAs. The multicast trees are encoded using IP addresses of participating nodes. Authors propose two encodings

- Per level header encoding and Preorder header encoding. In addition to IP addresses of nodes, these encodings spend significant number of additional bits. In a multicast tree of maximum degree d , the per-level header encoding spends $\log d$ bits per node and preorder header encoding spends approximately $\log n$ bits per node. We show an encoding using only 2 additional bits per node.

3 Multicast Trees and Tree Data Structures

Figure 2(i) shows a multicast tree occurring in an arbitrary network or overlay. The portion of the multicast tree which needs to be encoded is shown highlighted. The last hop routers, which serve one or more receivers are shown grayed. The link indexes corresponding to one of the nodes are shown. Figure 2(ii), (iii) and (iv) show the *Ordinal tree*, *Cardinal tree* and *Arbitrarily labeled tree* data structures corresponding to the multicast tree.

Fig. 2. (i) Multicast Tree in a Network (ii) Ordinal tree (iii) Cardinal tree (iv) Arbitrarily labeled tree (v) Balanced parentheses representation of ordinal tree

An ordinal tree is a rooted (a unique node is distinguished as the root), unlabeled tree in which each node has an arbitrary degree. A binary tree is an ordinal tree in which the maximum degree of a node is two. The number of ordinal trees on n nodes can be bounded by $O_n = \binom{2n+1}{n} / (2n+1)$. Thus, $\lg(O_n) \approx 2n$ is an information theoretic lower bound on the number of bits needed to represent ordinal trees ($\lg$ is $\log_2$). Every ordinal tree of n nodes can be represented using $2n$ balanced parentheses. Fig 2(v) shows the balanced parentheses representation of the ordinal tree in Fig 2(ii). This representation is obtained by a preorder (depth first) traversal of the tree, outputting a “(“ while visiting a node for the first time and a matching “)” while visiting the node after visiting its subtree (The correspondence of some parentheses to nodes is shown). By representing a “(“ by 1 and a “)” by 0, balanced parentheses can be represented using $2n$ bits.

Figure 2(iii) shows the cardinal tree which results when a multicast tree is represented using link indexes. A cardinal tree of degree d is an unlabeled tree in which each node has d positions for an edge to a child (if a child is at the i th position, the corresponding link gets the label i). There are $C_n^d = \binom{dn+1}{n} / (dn+1)$ cardinal trees of degree d on n nodes. $\lg(C_n^d) \approx (\lg d + \lg e)n$ is an information theoretic lower bound on the number of bits needed to encode these trees. A multicast tree in which the maximum link index is d is a cardinal tree of degree d . Figure 2(iii) is a cardinal tree of degree 5. Thus, $(\lg d + \lg e)n$ is a lower bound on the number of bits needed to represent a multicast tree using link indexes.

Figure 2(iv) shows the arbitrarily labeled tree which results when a multicast tree is represented using node identifiers. Node Identifiers a, b , etc are essentially IP addresses. A labeled tree of n nodes is an ordinal tree in which each node is labeled from 1 to n . An arbitrarily labeled tree is an ordinal tree in which each node has a *unique* arbitrary label. If each node is represented using a k bit label, then the number of arbitrarily labeled trees on n nodes can be bounded by $A_n^k = \binom{2^k}{n} n! \approx O_n$. Thus, $\lg(A_n^k) \approx kn + 2n$ is a lower bound on the number of bits needed to represent a multicast tree using node identifiers.

4 Multicast Tree Representations Using Link Indexes

For ease of explanation, we use the following notation. All nodes in the multicast tree have one incoming link (except the root) and zero or more outgoing links. A *terminal or leaf* node has zero outgoing links, a *relay* node has one outgoing link and a *branch* node has more than one outgoing links. In Fig 3(i), a, e, h and d are branch nodes, j, k, m, c, f and g are leaf nodes, and b, i and l are relay nodes. The links of a multicast tree can be divided into branch and relay links. All outgoing links of a branch node are branch links and all outgoing links of a relay node are relay links. For example, ac is a branch link and lm is a relay link. Let l denote the number of links and n denote the number of nodes, $n = l + 1$. Let $\bar{l}$ denote the number of branch links and $\bar{r}$ the number of relay links, $l = \bar{l} + \bar{r}$. Let b denote the number of branch nodes, r the number of relay nodes, and t the number of leaf nodes, $n = b + r + t$. For the tree in Fig 3(i), the preorder node traversal gives $a, b, e, h, j, k, i, l, m, c, d, f, g$ and the preorder link traversal gives $ab, be, eh, hj, hk, ei, il, lm, ac, ad, df, dg$.

4.1 Improvements to Linkcast Encoding

Linkcast consists of two encodings - SBM and DBM. These encodings consist of a series of elements. In SBM, each element is a link index or a pointer. A pointer is used to point to the representation of the child node in the encoding. In DBM, each element is either a link index, pointer or a node demarker. SBM requires $\bar{b}$ pointers. DBM reduces this to $\bar{b} - b$, but at the expense of introducing $l + 1$ node demarkers. Each node demarker requires at least two bits to be differentiated from both link indexes and pointers, consuming $2(l + 1)$ bits of space. Utilizing the concept of pointers, we now present a simplified encoding which is shorter than both SBM and DBM and refer to this encoding as Link+. Instead of using

node demarkers, our approach is to distinguish between links that terminate at routers which perform forwarding and routers which do not, using one bit per link. One bit is used to distinguish between pointers and link indexes.

To construct Link+ encoding, nodes are visited in preorder. When a node is visited, a string of pointers and outgoing link indexes is outputted. If a node has d outgoing links, then $d - 1$ pointers followed by d link indexes are outputted. The $d - 1$ pointers point to the output strings of 2nd to d th child nodes. The output string of 1st child node occurs immediately after the output string of current node. Fig 3(ii) shows the Link+ encoding for the tree in Fig 3(i). For example, root node a has three outgoing links, thus 2 pointers followed by three link indexes are outputted. The output string of the first child node b occurs after the output string of a . The 1st pointer points to the output string of 2nd child node c (null). The 2nd pointer points to the output string of 3rd child node d . The first bit in each element is used to distinguish a pointer (0) from a link index (1). The second bit in each link index is used to distinguish links which terminate on leaf nodes (0) from those which do not (1).

Routing. In Link+, each node receives the pointer to its position in the encoding (as Linkcast). The root node receives a pointer to the first element in the encoding. Each forwarding node determines the (i) outgoing links to forward the packet (ii) positions of child nodes in the encoding in turn to be forwarded to the children. If the position of the current node starts with a pointer, then it is a branch node, else it is a relay node. If a branch node finds d pointers, it reads the subsequent $d + 1$ elements to determine its outgoing links. A relay node receives the position of its outgoing link and its child is the next element. If the second bit of an outgoing link is 0, a null pointer is forwarded and thus a leaf node receives a null pointer. The forwarding operation at a node takes time proportional to its out degree. The detailed algorithm is given in [2].

In total, the encoding length of Link+ is $(S_l + 2)l + (S_p + 1)(\bar{b} - b)$ bits. S_l is the number of bits needed to represent the maximum link index in the tree. S_p is the number of bits needed to represent a pointer, $S_p = \lceil \lg(l + \bar{b} - b) \rceil$.

4.2 Encoding Based on Balanced Parentheses

We now present a link index encoding based on balanced parentheses representation. We refer to this encoding as Link*. Link* encoding consists of two parts (i) balanced parentheses representation of the tree (ii) preorder list of link indexes. The first encoding in Fig 3(iii) shows the encoding for the tree in Fig 3(i). The balanced parentheses representation is similar to the balanced parentheses representation of ordinal trees, except that each parenthesis now corresponds to the incoming link of a node, instead of the node itself. The encoding is constructed by a preorder traversal of the tree. When a link is visited for the first time, its link index and a "(" are outputted. When a link is visited after visiting all links in its subtree, a ")" is outputted. This results in $2l$ parentheses and l link indexes. In total, Link* requires $(S_l + 2)l$ bits. S_l as before denotes the number of bits needed to represent the maximum link index.

Fig. 3. (i) Multicast Tree (ii) Link+ encoding (iii) Link* encoding (iv) Link** encoding

Routing. In Link*, instead of passing a pointer to the position of a node in the encoding, the encoding itself is changed after each forwarding operation. Each node receives only the encoding of its subtree. Fig 3(iii) shows the encoding received by the root node a and the encodings it passes to child nodes b , c and d . A forwarding node reads the balanced parentheses once to determine the outgoing links and the encodings of subtrees to be passed to the child nodes. The i th open parenthesis "(" corresponds to the link index at the i th position in the list of link indexes within the encoding. For example, in Fig 3(iii), root node a reads its balanced parenthesis and determines that it has 3 children. The 1st "(" corresponds to child node b , the 9th "(" corresponds to c and the 10th "(" corresponds to d . Thus a 's output links are at positions 1, 9 and 10 in the list of link indexes, i.e., link indexes 1, 3 and 5 respectively. The detailed forwarding algorithm is given in [2]. The forwarding operation of a node takes time proportional to reading the balanced parenthesis representation of its subtree. (This encoding is similar to encoding of Munro, et al [3]. However, they consider efficient bit representations of balanced parentheses for very large trees, finding the parent of a node, size of subtrees, etc which are not needed for small and medium sized multicast trees. We use only the bare minimal representation of cardinal trees).

4.3 Improvements to Balanced Parentheses Representation

Several studies [5, 12, 6] have investigated the type of multicast trees which occur in the Internet. Trees occurring in the Internet are constrained by the underlying structure of the Internet. These trees have a large number of relay nodes compared to branch nodes. In Link*, other than the mandatory space of lS_l bits to represent the link indexes, 2 bits are spent for each link (for balanced parentheses). We now present a new encoding Link** which reduces the number of bits on those links which terminate on relay nodes and increases the number of bits on those links which terminate on branch nodes. In Link**, 3 bits are spent per incoming link of a branch node or leaf node, and only 1 bit is spent per incoming link of a relay node. In total, Link** requires $(S_l + 2)l + b + t - r$ bits. If $r > b + t$, the encoding spends less than 2 bits per link. In section 5, we shall see that for several trees in the Internet, the number of relay nodes is larger than the sum of branch and leaf nodes.

Link** encoding consists of three parts (i) relay bit (ii) balanced parentheses (iii) preorder list of link indexes. The encoding is similar to Link*. But it is obtained by encoding a virtual tree in which a path of consecutive relay nodes is treated as a single virtual link whose label is the concatenated list of its individual link indexes. Thus, the balanced parentheses represent only branch and leaf nodes. The first encoding in Fig 3(iv) shows the encoding for tree in Fig 3(i). For example, the path em is treated as a virtual link with label "415", and the path ae is treated as a virtual link with label "12". To determine the start and end of a virtual link, the first bit in each link index is used to distinguish links which terminate on intermediate relay nodes (1) from those which terminate on branch or leaf nodes (0). The relay bit is used for routing.

Routing. The forwarding operation is similar to Link*. Fig 3(iv) shows the encoding received by root node a and the encodings it passes to child nodes b , c and d . The i th open parenthesis "(" corresponds to the i th virtual link in the encoding. For example, in the encoding received by a , the 1st "(" corresponds to the 1st virtual link ae and the 5th "(" corresponds to the 5th virtual link em . During forwarding, a node reads the relay bit to check whether it is a relay node or not. The relay bit is 1 when the encoding is received by a relay node and 0 otherwise. The relay bit received by a node is simply the first bit taken from the link index of its incoming link. For example, when a performs forwarding, it sets the relay bit of child node b equal to the first bit of link index of link ab . In Link**, only branch nodes read the balanced parentheses and change it. The relay nodes read only the first link index to perform forwarding. The time complexity of forwarding is constant at relay nodes. The time complexity of forwarding at a branch node is proportional to reading the balanced parenthesis of its virtual subtree. The detailed forwarding algorithm is given in [2].

4.4 Representing Trees Using Node Identities

Both encodings Link* and Link** can be used *as is* by replacing link indexes of links by node identifiers of nodes which terminate on those links. Figure 4 shows the encoding Link* using node identifiers for tree Fig 3(i), as received by root node a . The forwarding algorithm remains the same, except that each forwarding node routes the packet to the corresponding next hop nodes instead of sending it on specific outgoing links [2].

Fig. 4. Link* encoding using nodes instead of links

5 Simulations

To evaluate these encodings, we conducted tests using various generated and real topologies of the Internet. Here, we present the results for three representative topologies - **ts1000**, **scan**, and **att**. **ts1000** is the transit-stub topology generated by GT-ITM [4] consisting of 1040 routers. **scan** is a partial map of the Internet collected by SCAN project [14] consisting of 284,805 routers. **att** is the router level ISP topology of AT&T collected by Rocketfuel [18] consisting of 731 routers. For each of these topologies, we chose random routers as source and last hop routers and constructed shortest path multicast trees from source to last hop routers. To compute shortest paths, for **scan** and **att**, hop count metric was used and for **ts1000**, the generated weights were used. Figure 5 (a),(c) and (e) show the properties of multicast trees for the three topologies (each point is the average of 1000 simulations).

(a) ts1000

(b) ts1000

(c) scan

(d) scan

(e) att

(f) att

Fig. 5. Properties of multicast trees and their encoding lengths in various topologies

As observed, in all topologies the number of relay nodes is significantly larger than the number of branch nodes. However, the total number of branch links is significant and grows linearly with the number of last hop routers. Figure 5(b),(d)

and (f) show the encoding lengths of Link+, Link*, Link**, and Xcast+. The Xcast+ encoding length is essentially 32 times the number of last hop routers. For link based encodings, the link index was represented using 5 bits (routers in the Internet rarely have degree larger than 32). The pointer in Link+ was represented using 8 bits. As seen, the Link+ encoding takes more space than both Link* and Link** to represent pointers corresponding to the branch links (Linkcast encodings take at least $l + 2$ more bits than Link+, for l links). For inter-domain topologies **ts1000** and **scan**, Link** encoding takes less space than Link* due to the presence of a large number of relay routers. For intra-domain topology **att**, as the number of last hop routers increase (> 25), Link* takes less space than Link** since the proportion of relay routers reduces compared to branch and last hop routers.

6 Conclusions

In this paper, we presented efficient ways of encoding specific or explicit multicast trees using link indexes and node identifiers. We presented an improved encoding Link+ and two new encodings Link* and Link** which consume space close to the minimum number of bits needed to represent multicast trees. These encodings can be used to represent multicast trees using either link indexes or node identifiers. We evaluated the space consumed by these encodings using shortest path multicast trees in real and generated topologies. These encodings can be used for various applications such as multicast state reduction, traffic engineering and application level multicast, and provide a feasible way of representing trees for small and medium sized multicast groups within data packets.

References

1. David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. The Case for Resilient Overlay Networks. In *8th Annual Workshop on Hot Topics in Operating Systems*, 2001.
2. Vijay Arya, Thierry Turetletti, and Shivkumar Kalyanaraman. Encodings of Multicast Trees. Technical Report RR-5442, INRIA, Dec 2004.
3. David Benoit, Erik D. Demaine, J. Ian Munro, and Venkatesh Raman. Representing Trees of Higher Degree. In *International Workshop on Algorithms and Data Structures*, 1999.
4. Ken Calvert, Matt Doar, and Ellen W. Zegura. Modelling Internet Topology. *IEEE Communications Magazine*, 1997.
5. Robert C. Chalmers and Kevin C. Almeroth. On the topology of multicast trees. *IEEE/ACM Transactions on Networking*, 11(1):153–165, 2003.
6. Danny Dolev, Osnat Mokryn, and Yuval Shavitt. On Multicast Trees: Structure and Size Estimation. In *IEEE INFOCOM*, 2003.
7. Rick Boivie et al. Explicit Multicast (Xcast) Basic Specification, Internet Draft draft-ooms-xcast-basic-spec-06.txt, 2004.
8. S. Deering et al. The pim architecture for wide-area multicast routing. *IEEE/ACM Transactions on Networking*, pages 153–162, 1996.

9. Hema Tahiramani Kaur, Shivkumar Kalyanaraman, Andreas Weiss, Shifalika Kanwar, and Ayesha Gandhi. BANANAS: An Evolutionary Framework for Explicit and Multipath Routing in the Internet. In *SIGCOMM FDNA Workshop*, 2003.
10. Michael R. Macedonia, Michael J. Zyda, David R. Pratt, Paul T. Barham, and Steven Zeswitz. NPSNET: A Network Software Architecture For Large Scale Virtual Environments. In *Presence: Teleoperators and virtual Environments*, 1994.
11. Mozafar Bag Mohammadi, Siavash Samadian-Barzoki, and Nasser Yazdani. Linkcast: Fast and Scalable Multicast Routing Protocol. In *NETWORKING*, pages 1282–1287, 2004.
12. Graham Phillips, Scott Shenker, and Hongsuda Tangmunarunkit. Scaling of multicast trees: comments on the Chuang-Sirbu scaling law. In *Conference on Applications, technologies, architectures, and protocols for computer communication*, pages 41–51, 1999.
13. George Popescu and Zhen Liu. Stateless Application-Level Multicast for Dynamic Group Communication. In *IEEE International Symposium on Distributed Simulation and Real-Time Applications (DS-RT'04)*, pages 20–28, 2004.
14. SCAN project. <http://www.isi.edu/scan/mercator/maps.html>.
15. J. Mark Pullen and David Wood. Network Technology for DIS. In *Proc. of IEEE*, pages 1156–1167, Aug 1995.
16. A. Rowstron, A-M. Kermarrec, M. Castro, and P. Druschel. SCRIBE: The design of a large-scale event notification infrastructure. In *Proc. of NGC*, 2001.
17. S.Y. Shi and J.S. Turner. Routing in overlay multicast networks. In *IEEE INFOCOM*, 2002.
18. Neil Spring, Ratul Mahajan, and David Wetherall. Measuring ISP Topologies with Rocketfuel. In *SIGCOMM*, 2002.
19. D. Waitzman, C. Partridge, and S.E. Deering. Distance Vector Multicast Routing Protocol. RFC 1075, November 1988.
20. Chung Kei Wong, Mohamed Gouda, and Simon S. Lam. Secure group communications using key graphs. *IEEE/ACM Transactions on Networking*, 8(1):16–30, 2000.
21. Shelley Q. Zhuang, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz, and John Kubiawicz. Bayeux: An Architecture for Scalable and Fault-tolerant Wide-Area Data Dissemination. In *NOSSDAV*, 2001.