



Institut Eurécom
Department of Corporate Communications
2229, route des Crêtes
B.P. 193
06904 Sophia-Antipolis
FRANCE

Research Report RR-04-107

MULTI+: Building Topology-Aware Overlay Multicast Trees

April 2004

L. Garcés-Erice, Ernst W. Biersack and Pascal A. Felber

Tel : (+33) 4 93 00 26 26
Fax : (+33) 4 93 00 26 27
Email : {Luis.Garces, Ernst.Biersack, Pascal.Felber}@eurecom.fr

¹ Institut Eurécom's research is partially supported by its industrial members: Bouygues Télécom, Fondation d'entreprise Groupe Cegetel, Fondation Hasler, France Télécom, Hitachi, ST Microelectronics, Swisscom, Texas Instruments, Thales

MULTI+: Building Topology-Aware Overlay Multicast Trees

Luis Garcés-Erice, Ernst W. Biersack and Pascal A. Felber

Institut EURECOM
06904 Sophia Antipolis, France
{garces|erbi|felber}@eurecom.fr

Abstract. TOPLUS is a lookup service for structured peer-to-peer networks that is based on the hierarchical grouping of peers according to network IP prefixes. In this paper we present MULTI+, an application-level multicast protocol for content distribution over a peer-to-peer (P2P) TOPLUS-based network. We use the characteristics of TOPLUS to design a protocol that allows for every peer to connect to an available peer that is close. MULTI+ trees also reduce the amount of redundant flows leaving and entering each network, making more efficient bandwidth usage. We use different procedures to measure or calculate the distances among peers, in order to validate the algorithms in MULTI+.

Keywords: Overlay networks, Topology-aware, Multicast

1 Introduction

IP Multicast seems (or, at least, was designed) to be the ideal solution for content distribution over the Internet: (1) it can serve content to an unlimited number of destinations, and (2) it is bandwidth-wise economic. These two characteristics are strongly correlated. IP Multicast saves bandwidth because a single data flow can feed many recipients. The data flow is only split at routers where destinations for the data are found in more than one outgoing port. Thus n clients do not need n independent data flows, which allows for IP Multicast's scalability. However, IP Multicast was never widely deployed in the Internet: security reasons, its open-loop nature, made IP multicast remain as a limited use tool for other protocols in LANs. The core Internet lacks of an infrastructure with the characteristics of IP Multicast.

Lately, with the advent of broadband links like ADSL and the generalization of LANs at the workplace, the *edges* of the Internet started to increase their bandwidth. Together with the ever-cheaper and yet more powerful equipment (computational power, storage capacity), they give the millions of hosts connected to the Internet the possibility of implementing themselves services that augment at the application level the capabilities of the network: the Peer-to-Peer (P2P) systems. Various application-level multicast implementations have been proposed [3, 12, 27, 20, 7], most of which are directly implemented on top of P2P infrastructures (Chord [23], CAN [19] or Pastry [21]). The good scalability of the underlying P2P networks give these application-level multicast one of the properties of the original IP Multicast service, that of serving content to a virtually unlimited number of clients (peers). However, these P2P networks are generally conceived as an application layer system completely isolated from the underlying IP network.

Thus, the P2P multicast systems that we know may fail at the second goal of IP Multicast: a LAN hosting a number of peers in a P2P multicast tree may find its outbound link saturated by *identical* data flowing to and from its local peers, unless those peers are somehow *aware* of the fact that they are sharing the same network. This situation remains valid for higher-tier networks (e.g., ISPs): even if the capacity of links to the Internet is designed using conservative parameters, normally the resulting system does not allow for all or many hosts to sustain a continuous throughput from and to the Internet. This is a critical issue for ISPs due to P2P file-sharing applications and flat-rate commercial offers that allow a home computer to be downloading content 24 hours a day. This problem also affects application-level multicast sessions if peers do not take care of the network topology.

We have based our P2P multicast protocol on TOPLUS because of its inherent topology-awareness. We aim at building a P2P multicast network that avoids network link stress while remaining scalable. In TOPLUS, nodes that are topologically close are organized into groups. In turn, groups that are topologically close are organized into supergroups, and close supergroups into hypergroups, etc. The groups within each level of the hierarchy can be heterogeneous in size and in fan-out (i.e., number of subgroups). TOPLUS derives the groups from the network prefixes contained in BGP tables or from other sources.

Before we elaborate on a Multicast deployment in a TOPLUS network, we must state some assumptions about the system: First of all, there is a large population of peers participating in the network, which justifies the utilization of multicast in the first place. Second, the TOPLUS Multicast infrastructure may be used to transmit to many different multicast groups, and thus many multicast trees must live together without interfering each other. Third, a peer is only willing to transmit data from a multicast tree if the peer is in the multicast group (that is, a peer will not use its bandwidth just to forward data for others). However, when participating in the TOPLUS Multicast infrastructure, a peer must accept to cooperate with others in low-bandwidth maintenance tasks.

Related Work. Some examples of overlay networks which introduce topology-awareness in their design are SkipNet [17], Coral[9], Pastry [5], CAN [22]. Application-level Multicast has given some interesting results like the NICE project [3] or End System Multicast [12]. Among those using overlay networks, there are application-level multicast implementations like Bayeux [27], or using CAN [20] and Pastry (Scribe) [7]. There is an interesting comparative in [8]. Content distribution overlay examples are SplitStream [6] and [13]. Recently, the problem of data dissemination on adaptive overlays has been treated in [26]. Our approach differs mainly in the achievement of efficient topology-aware multicast trees with no or very little active measurement.

In the next section we present the main aspects of TOPLUS. Then we sketch MULTI+. We test different measurement techniques, and discuss on their accuracy. Then we show and comment some results, before we conclude.

2 TOPLUS Overview

TOPLUS [10] is based on the DHT paradigm, in which a resource is uniquely identified by a key, and each key is associated with a single peer in the network. Keys and peers share a numeric identifier space, and the peer with the identifier closest to a key is responsible for that key.

Given a message containing key k , the P2P lookup service routes the message to the current active peer that is responsible for k . The message travels from source peer p_s , through a series of intermediate peers p_1, p_2, \dots, p_v , and finally to the destination peer, p_d . The principal goal of TOPLUS is simple: each routing step takes the message closer to the destination. We now formally describe TOPLUS in the context of IPv4.

Given a message with key k , source peer p_s sends the message (through IP-level routers) to a first-hop peer p_1 that is “topologically close” to p_d ; After arriving at p_1 , the message remains topologically close to p_d as it is routed closer and closer to p_d through the subsequent intermediate peers.

Let I be the set of all 32-bit IP addresses.¹ Let \mathcal{G} be a collection of sets such that $G \subseteq I$ for each $G \in \mathcal{G}$. Thus, each set $G \in \mathcal{G}$ is a set of IP addresses. We refer to each such set G as a *group*. Any group $G \in \mathcal{G}$ that does not contain another group in \mathcal{G} is said to be an *inner group*. We say that the collection \mathcal{G} is a *proper nesting* if it satisfies all the following properties:

1. $I \in \mathcal{G}$.
2. For any pair of groups in \mathcal{G} , the two groups are either disjoint, or one group is a proper subset of the other.
3. Each $G \in \mathcal{G}$ consists of a set of contiguous IP addresses that can be represented by an IP prefix of the form $w.x.y.z/n$ (for example, 123.13.78.0/23).

The collection of sets \mathcal{G} can be created by collecting the IP prefix networks from BGP tables and/or other sources [14]. In this case, many of the sets \mathcal{G} would correspond to ASes, other sets would be subnets in ASes, and yet other sets would be aggregations of ASes. This approach of defining \mathcal{G} from BGP tables require that a proper nesting is created. Note that the groups differ in size, and in number of subgroups (the fanout).

If \mathcal{G} is a proper nesting, then the relation $G \subset G'$ defines a partial ordering over the sets in \mathcal{G} , generating a partial-order tree with multiple tiers. The set I is at tier-0, the highest tier. A group G belongs to tier 1 if there does not exist a G' (other than I) such that $G \subset G'$. We define the remaining tiers recursively in the same manner (see Figure 1).

Peer State. Let L denote the number of tiers in the tree, let U be the set of all current active peers and consider a peer $p \in U$. Peer p is contained in a collection of telescoping sets in \mathcal{G} ; denote these sets by $H_N(p), H_{N-1}(p), \dots, H_0(p) = I$, where $H_N(p) \subset H_{N-1}(p) \subset \dots \subset H_0(p)$ and $N \leq L$ is the tier depth of p 's inner group. Except for $H_0(p)$,

¹ For simplicity, we assume that all IP addresses are permitted. Of course, some blocks of IP addressed are private and other blocks have not been defined. TOPLUS can be refined accordingly.




Fig. 1. A sample TOPLUS hierarchy (inner groups are represented by plain boxes)




Fig. 2. A simple multicast tree.

each of these telescoping sets has one or more siblings in the partial-order tree (see Figure 1). Let $\mathcal{S}_i(p)$ be the set of sibling groups of $H_i(p)$ at tier i . Finally, let $\mathcal{S}(p)$ be the union of the sibling sets $\mathcal{S}_1(p), \dots, \mathcal{S}_N(p)$.

Peer p should know the IP address of at least one peer in each group $G \in \mathcal{S}(p)$, as well as the IP addresses of all the other peers in p 's inner group. We refer to the collection of these two sets of IP addresses as peer p 's *routing table*, which constitutes peer p 's state. The total number of IP addresses in the peer's routing table in tier L is $|H_L(p)| + |\mathcal{S}(p)|$. In [10] we describe how a new peer can join an existing TOPLUS network.

XOR Metric. Each key k' is required to be an element of I' , where I' is the set of all s -bit binary strings ($s \geq 32$ is fixed). A key can be drawn uniformly randomly from I' , or it can be biased as we shall describe later. For a given key $k' \in I'$, let k be the 32-bit suffix of k' (thus $k \in I$ and $k = k_{31}k_{30} \dots k_1k_0$). Throughout the discussion below, we will refer to k rather than to the original k' .

The XOR metric defines the distance between two ids j and k as $d(j, k) = \sum_{\nu=0}^{31} |j_\nu - k_\nu| \cdot 2^\nu$. The metric $d(j, k)$ has the following properties:

- If $d(i, k) = d(j, k)$ for any k , then $i = j$.
- Let $p(j, k)$ be the number of bits in the common prefix of j and k . If $p(j, k) = m$, $d(j, k) \leq 2^{32-m} - 1$.
- If $d(i, k) \leq d(j, k)$, then $p(i, k) \geq p(j, k)$.

$d(j, k)$ is a refinement of longest-prefix matching. If j is the unique longest-prefix match with k , then j is the closest to k in terms of the metric. Further, if two peers share the longest matching prefix, the metric will break the tie. The peer p^* that minimizes $d(k, n)$, $p \in U$ is “responsible” for key k .

TOPLUS Network Initialization. In [10] we describe how a new peer can join an existing TOPLUS network, but how is this network created in the first place? Before explaining how our P2P multicast works in an operating, fully deployed TOPLUS network, we show a method to initialize the P2P network, which validates the TOPLUS algorithms and our multicast scheme even in the early stages of deployment. Our solution involves the existence of one first peer that knows the whole TOPLUS tree. This is not a special requirement: recall that the TOPLUS tree is calculated off-line, and its knowledge is necessary *before* the TOPLUS network is deployed. This peer is called the “primordial peer”, or p^0 , for short.

p^0 is deployed by some institution or company offering a service involving TOPLUS. p^0 should be a powerful well-connected host. If needed, p^0 can be implemented by a cluster of machines, provided they act as a coordinated entity.

When the system is first deployed, no peer can be found in any group. However, p^0 is the first “virtual” member of each tier-1 group. Thus, p^0 has a routing table where the delegate in all tier-1 groups is itself.

When a first peer p comes into the system, it must contact the only existing peer, p^0 . p asks p^0 to find p 's corresponding inner group G , and the contact peer in G is p^0 again. However, when p^0 in G passes the routing table to the new peer p , p^0 is not among the members of group G : p is alone in G . Then p , to maintain the “diversity property”, asks each peer in the routing table for an alternate delegate. If no new peer has arrived, p will remain with p^0 as the only delegate in each group. Instead, if one peer q came to any other group, p^0 must know about him, and q is returned as an alternate delegate.

Note that arriving peers cannot in the future use p^0 when they answer to a query for a peer in their group to satisfy the diversity property: when a peer receives the list of members of its group, the primordial peer is never among them. Thus, p^0 silently disappears from the system over time.

Of course, we are assuming here that once the system starts its deployment, there is a continuous flow of arriving peers. If the first peer arrives and afterwards leaves, p^0 should come back to its initial situation. We also assume that peers belong more or less uniformly to all groups. That should be the case at least in statically deployed TOPLUS trees. In other cases, the TOPLUS tree should contain only groups where peers indeed exist or may exist (e.g., personal dial-up and ADSL ISPs for file-sharing applications), eventually reaching a situation where all groups are more or less populated.

Recall that each peer must know another in each tier-1 group. The primordial peer avoids communication to all tier-1 groups each time one of them is populated for the first time. There's always somebody in all tier-1 groups: p^0 , at the very least. Peers just get to know new peers in other groups, through the "diversity property" conservation process, which is less costly than broadcasting to all groups the arrival of a peer to a previously empty group.

This approach can be introduced in other DHTs. Although the technical details may differ because of the different algorithms, the idea is to initialize the P2P network with a set of virtual peers uniformly distributed in the ID space. Those peers are implemented by a single physical host (or cluster), namely our p^0 . Our technique gives stability to the initial network.

3 MULTI+: Multicast on TOPLUS

A Multicast Tree. First we assume that all peers are connected through links providing enough bandwidth. A simple multicast tree is shown in Figure 2. Let S be the source of the multicast group m . Of course, the source should be located in some group of a TOPLUS tree covering the whole Internet. Peer p is receiving the flow from peer q . We say that q is the parent of p in the multicast tree. Conversely, we say that p is a child of q , its parent in the multicast tree. Peer p is in level-3 of the multicast tree and q in level-2. It is important to note that, in principle, the *level* where a peer is in the multicast tree has nothing to do with the *tier* the peer belongs to in the TOPLUS tree.

Peer p is in level-3 of the multicast tree and q in level-2. It is important to note that, in principle, the *level* where a peer is in the multicast tree has nothing to do with the *tier* the peer is in the TOPLUS tree.

In the kind of multicast trees we aim at building, each peer should be close to its parent in terms of network delay, while trying to join the multicast tree as high (close to the source) as possible. Each peer attempts at join time to minimize the number of hops from the source, and the length of the last hop. In the example of Figure 2, if p is a child of q and not of r , that is because p is closer to q than to r . Moreover, we want this condition to hold even if when p joins, q has not yet integrated the multicast tree: if q joins later and is closer to p than to r , then q should become a child of p . By trying to minimize the network delay for data transmission between peers, we also avoid rearranging peers inside the multicast tree, except when a peer fails or disconnects. For the kind of applications we are targeting the content distribution network should be as stable as possible. For a discussion on this topic, see *Membership Management* below.

Although P2P networks are often characterized by a frequent joining and leaving of its members, we expect that peers connecting to a given multicast group are interested in remaining connected at least for the duration of the multicast session. The TOPLUS network infrastructure offers a very interesting framework for the construction of such multicast trees.

Building Multicast Trees. We use the TOPLUS network and look-up algorithm in order to build the multicast trees. Consider a multicast IP address m , and the corresponding key that, abusing the notation, we also denote m . Each tier- i group G_i is defined by an IP network prefix a_i/b where a_i is an IP address and b is the length of the prefix in bits. Let m_i be the key resulting from the substitution of the first b bits of m by those of a_i . The inner group that contains the peer responsible for m_i (obtained with a TOPLUS look-up) is the *responsible inner group*, or RIG, for m in G_i (note that this RIG is contained in G_i .) Hereafter, we assume a single m , and for that m and a given peer p we denote the RIG in $H_i(p) \in \text{tier-}i$ simply as RIG- i of p . This RIG is a rendezvous point for all peers in $H_i(p)$. The deeper that a tier- i of a RIG- i is in the TOPLUS tree, the narrower the scope of the RIG as a rendezvous point (fewer peers can potentially use it).

In the simple 3-tier example of Figure 3, we have labeled the RIGs for a given multicast group (peers in grey are members of the multicast group), where all inner groups are at tier-3. The RIG- i of a peer can be found following the arrows. The arrows represent the process of asking the RIGs for a parent in the multicast tree. For example, p and q share the same RIG-1 because they are in the same tier-1 group. t 's inner group is its RIG-1, but t would first contact a peer x (white) in its RIG-2 to ask for a parent. Note that this last peer is not in the multicast tree (Figure 4.) Some inner groups can be RIGs for more than one higher-tier group. If a peer is alone in its tier-1 group, it asks the source for a parent.




Fig. 3. The RIGs in a sample TOPLUS network.




Fig. 4. Sample multicast tree.

Assume a peer p in tier- $(i + 1)$ (i.e., a peer whose inner group is at tier- $(i + 1)$ of the TOPLUS tree) wants to join a multicast tree with multicast IP address m , which we call group m .

1. The peer p broadcasts a query to join group m inside its inner group. If there is a peer p' already part of group m , p connects to p' to receive the data.
2. If there is not such peer p' , p must look for its RIG- i . A look-up of m_i inside p 's tier- i group (thus among p 's sibling groups at tier- $(i + 1)$) locates the RIG- i responsible for m . p contacts any peer p_i in RIG- i , and asks for a peer in multicast group m . If peer p_i knows about a peer p'' that is part of m , it sends the IP address of p'' to p , and p connects to p'' . Note that p'' is not necessarily a member of the RIG- i inner group. In any case p_i adds p to the peers listening to m , and shares this information with all peers in RIG- i . If p'' does not exist, p proceeds similarly for RIG- $(i - 1)$: p looks up m_{i-1} inside p 's tier- $i - 1$ group (i.e., among p 's sibling groups at tier i). This process is repeated until a peer receiving m is found, or RIG-1 is reached. In the latter case, if there is still no peer listening to m , peer p must connect directly to the source of the multicast group.

One can see that the search for a peer to connect to is done bottom up.

Property 1 When a peer p in tier $i + 1$ joins the multicast tree, by construction, from all the groups $H_{i+1}(p), H_i(p), \dots, H_1(p)$ that contain p , p connects to a peer $q \in H_k$ where $k = \max\{l = 1, \dots, i + 1\} : \exists r \in H_l \text{ and } r \text{ is a peer already connected to the multicast tree}$. That is, p connects to a peer in the deepest tier group which contains both p and a peer already connected to the multicast tree.

This assures that a new peer connects to the closest available peer in the network. Notice that even in the case of failure of a peer in a RIG- i , the information is replicated in all other peers in the RIG- i . If a whole RIG- i group fails, although MULTI+ is undeniably affected, the look up process can continue in RIG- $(i - 1)$. We believe this property makes MULTI+ a resilient system. For a proof of this and other properties below, we refer to the Appendix.

Multicast Tree Topology. Any peer in any group can become a multicast tree node; only the RIGs are fixed, for a given multicast group and a given TOPLUS network structure. The topology of the tree is thus completely random, depending only on the peers that want to join a given multicast tree. However, although not deterministic, this tree is constrained to follow the TOPLUS structure. As it has been shown before, when a peer joins a multicast tree it will always connect to another peer that is close (according to TOPLUS proximity).

In general, we can assume that for any networked group in the Internet it is better to keep as much traffic as possible *inside* the group and avoid outbound traffic: Between ASs, routing at peering interfaces is done based mainly on economical considerations, whereas inside an AS normally other policies stand, more closely related to networking. In most LANs, the outbound link to the Internet presents an order of magnitude less bandwidth than the network itself. We show here that Multicast on TOPLUS minimizes internetwork traffic.

Property 2 For each group defined by an IP network prefix containing at least one peer connected to the multicast tree, there is only one inbound data flow. **Proof:** We proof by induction for a peer p in tier- $i + 1$ contained in a set of groups (or networks defined by IP network prefixes) $H_{i+1}(p), H_i(p), \dots, H_1(p)$:

- If there is one peer q already connected to the multicast tree in p 's inner group H_{i+1} , p connects to q . Else, p must connect to another peer outside H_{i+1} . Thus only one flow may go into the network H_{i+1} .




Fig. 5. Inbound and outbound flows at each network.




Fig. 6. Inbound and outbound flows at each network, with $F = 4$ connections per peer.

- In the look-up of a peer already connected to the multicast tree, assume p gets to $H_k(p)$ in tier- k . Thus, by property 1 no peer exists connected to the multicast tree in $H_{i+1}(p), H_i(p), \dots, H_{k+1}(p)$. Only the data flow to p will go into $H_{i+1}(p), H_i(p), \dots, H_{k+1}(p)$. If there is one peer $q \in H_k(p)$ receiving the multicast flow, p connects to q , and p adds no data flow into $H_k(p)$. Else p proceeds the look-up in $H_{k-1}(p)$. In any case only one data flow can go into the network defined by $H_k(p)$. \square

Property 3 For each group defined by an IP network prefix containing at least one peer connected to the multicast tree, the number of outbound data flows in the worst possible case is bounded by a constant. **Proof:** Assume as before a peer p in tier- $i + 1$ contained in a set of groups (i.e., networks defined by IP network prefixes) $H_{i+1}(p), H_i(p), \dots, H_1(p)$. Assume peer p is looking for a peer to connect to. Without loss of generality, let $H_k(p)$ be the tier- k group where the RIG- k knows about a peer q connected to the multicast tree. p connects to q and adds an outbound data flow (from q to p) to $H_{k+1}(q)$. From the previous property, there can only be at most one inbound data flow in $H_{k+1}(p)$. The same is valid for every sibling group of $H_{k+1}(q)$. The number of siblings of $H_{k+1}(q)$ is a constant. Thus the number of outbound data flows from $H_{k+1}(q)$ towards tier- $(k + 1)$ groups is a constant. For a worst case scenario, assume now that q is also the peer known by RIG- $(k - 1)$ in $H_{k-1}(q)$. Applying the same reasoning to the siblings of $H_k(q)$, we have another constant number of outbound flows from $H_{k+1}(q)$ towards tier- k groups. A constant number of outbound flows at every tier, and the number of tiers being bounded by a constant, the number of outbound data flows is bounded by a constant for group $H_{k+1}(q)$. \square

Property 4 Using multicast over TOPLUS, the total number of flows in and out of a group defined by an IP network prefix is bounded by a constant. The proof is trivial from the two previous properties. \square

Note that this is however an ideal scenario where each peer can transmit the received data flow to as many peers as necessary. In the real world, there is a bandwidth limitation. We study later a modification of the protocol that allows each peer in the multicast tree to set a maximum number of children, depending on its bandwidth. We assume that most peers have good-quality links to the Internet (at least a medium-range ADSL).

Membership Management. Each peer p knows its parent q in the multicast tree, because there is a direct connection between them. Because p knows the RIG where it got its parent's address, if p 's parent q in level i of the multicast tree fails or disconnects, p directly goes to the same RIG and asks for a new parent. If there is none, p becomes the new tree node at level i , replacing q . Then p must find a parent in level $i - 1$ of the multicast tree, through a join process starting at said RIG. If p had any siblings under its former parent, those siblings will find p as the new parent when they proceed like p . If more than one peer concurrently tries to become the new node at level i , peers in the RIG must consensually decide on one of them. It is not critical if a set of peers sharing a parent q are divided in two subsets with different parents upon q 's depart.

Join and leave is a frequent process in a P2P network, but we expect the churn to be rather low due to the fact that in a multicast tree, all peers seek the same content concurrently, throughout the duration of the session. Thus we do not consider "freeriding", i.e., peers that connect just to download some data and disconnect immediately afterwards, which introduces high dynamics in the overlay network.

Multicast Flow Load Distribution. Depending on the structure of the TOPLUS network, a node in a multicast tree may have many descendants. For example, in a tier-1 group with 100 tier-2 subgroups, the level 1 parent may have up to 100 descendants. The bandwidth of a single peer may not be sufficient. We now introduce two load balancing techniques to deal with that issue:

1. In the protocol we have just presented, all peers arriving to a RIG- k looking for a parent in the multicast tree get assigned the same parent. This parent peer may become overloaded if a large number of children get the multicast flow from him. Assume that each peer can afford a maximum of F outgoing flows. For each set of F new peers arriving to a RIG- k , one of them should become the new parent for the next F peers to come later. Notice that this policy alone would increment the number of levels in the multicast tree, and consequently the delay from root to leaves. To see this, in Figure 6 only the first F groups can connect directly to q . The next F peers must connect to one of the first F . The multicast tree is increased by 1 level for them (compare the hops from p to q in Figures 5 and 6). The problem is that the RIG- k is only allowing F connections to q , but there may be other peers in $H_{k+1}(q)$ that can accept connections. After all connections to q have been allocated, q can suggest as a response a connection to another peer in $H_{k+1}(q)$ receiving the data flow. The height of the tree does not increase. Then, only if there are no available peers in $H_{k+1}(q)$, q can suggest a peer in one of the sibling groups of $H_{k+1}(q)$ in tier $k + 1$. There is another constrain though: independently of the peers willing to serve content in $H_{k+1}(q)$, the network outgoing link may not have enough bandwidth, and another sibling group may need to be used.
2. At RIG- k , we assign to a new peer a parent that depends on the RIG- $(k + 1)$ of origin of the new peer. Taking Figure 6 as an example, peer r whose RIG- $(k + 1)$ is not contained in $H_k(q)$ should have priority to connect directly to q . The reason of this is to allow peers from further groups to connect to a parent through fewer hops. Peers closer to q , that is, those contained in $H_k(q)$, can afford more hops when connecting to q through other peers in $H_k(q)$, because the delay is smaller. Thus delays become more balanced.

The properties presented previously are still valid with these policies. To understand why, consider that the number of inbound flows does not change and that the property for the outbound flows is based on the worst possible case: these policies just move a part of the outbound flows from the most loaded group to its siblings. Moreover, note that this only happens in the case where the outbound link of the network of the worst case group cannot sustain the required flow of data. Thus we still have a constant bound in the number of flows coming in and out of a group.

Limited Source Capabilities. We have considered so far a source capable of serving content to at least one peer in each tier-1 group in the TOPLUS tree. For compact TOPLUS trees, this sums up to 256 potential flows. Although 256 flows is perfectly affordable to commercial or institutional multicast (think of Radio stations on the web today), it does not allow end users in the Internet to serve content which may potentially be of interest to many.

The problem is that our method to find a nearby existing peer is not valid between tier-1 groups. In the absence of a peer to connect to, we connect by default to the source. We propose here another method: the source accepts a limited number of direct connections. Once the limit is reached, new incoming connections receive a list of peers currently receiving the flow. The peers can then connect to the closest non-overloaded in the list (according to some easily measurable metric like RTT, or calculating their distance in a pre-calculated coordinate space, as we shall see below.) This process can equally apply to normal peers, when the number of connections is limited: a peer refusing a new connection must give a list of peers already connected to him. To detect uncooperative or malicious behavior we can verify that those peers are indeed connected to the refusing peer. This process is repeated until a valid peer is found. Note that collusion is not possible without introducing loops in the tree.

We can further discriminate among peers in a similar fashion as described above: Peers far from the source should have priority at getting direct connections, or even a pool of reserved bandwidth (e.g., reserving a fraction of the connections for peers farther than 200 ms.) For peers close to the source, it is less critical to get the content directly.

Parent Selection Algorithms. From the ideas exposed before, we retain two main parent selection algorithms for testing the construction of multicast trees.

- FIFO, where a peer joins the multicast tree at the first parent found with a free connection. When a peer gets to a RIG to find a parent, the RIG answers with a list of already connected peers. This list is ordered by arrival time to the RIG. Obviously, the first to arrive connects closer (in hops) to the source. The arriving peer tests each possible peer in the list starting with the first one until it finds one that accepts a connection.
- Proximity-aware, where, *when the first parent in the list has all connections occupied*, a peer connects to the closest parent in the list still allowing one extra connection.

Note that we do not always verify if we are connecting to the closest parent in the list. The idea behind this is that, while we implicitly trust MULTI+ to find a close parent, we prefer to connect to a peer higher in the multicast

tree (fewer hops from the source) than to optimize the last hop delay. If MULTI+ works correctly, the difference between these two policies should not be excessive, because the topology-awareness is already embedded in the protocol through TOPLUS. We do not constantly rearrange the multicast tree structure with each arrival, because the longer a peer has been in the system the more likely it is to stay connected [4], and short-term delay minimization optimizations may not pay in the long run in terms of network stability.

4 Testing MULTI+ with King

In this first experiment we test the characteristics of Multicast trees built with MULTI+ using `King`. We use two sets of peers, one randomly (to some extent) generated, and the other formed by a subset of the hosts in the Planet-Lab project. For both of them we use `King` to measure the distance between arbitrary hosts, and the TOPLUS tree with exclusively 8-bit group prefixes in tier-1.

4.1 King performance

We use `King` to measure the distance between two arbitrary hosts. First, `King` gives a good approximation of this distance in most cases, and second, `King` allows one to have an estimation of the distance between two hosts when more accurate methods like `ping` do not work. To see why, we proceed with our first experiment. We take the 328 hosts in Planet-Lab (at the time of the test), and we measure the distance with both `ping` and `King` from a host in Eurecom. Figure 7 shows the relative error in percentage of the distance measured by `King` to the actual distance, measured with `ping`. 80% of the measurements present less that 20% of relative error.

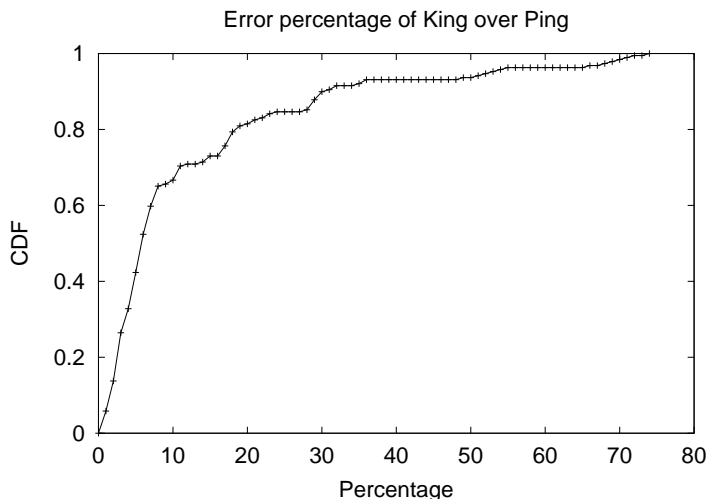


Fig. 7. Error in King measurement compared to Ping (from Eurecom).

But more interesting, with `King` we were able to measure 249 distances to hosts, while `ping` was only able to reach 191 destinations (obviously Figure 7 compares the hosts (171) that both `ping` and `King` were able to reach.) This can be due to a number of reasons, like security policies not allowing some ICMP traffic, hosts disconnected from the network or going under maintenance, etc. `King` is clearly a compromise between the total accuracy of `ping` and the convenience of a more reliable set of hosts used to obtain the measurements, the DNS servers. Furthermore, `King` ease to measure from a single point of the network a matrix of distances among a large set of hosts is more convenient than the need to log into each host to make the measurement with `ping`.

`King` measurement method is based on information contained in requests to DNS servers, thus is an *application level* method. `King` does not induce high stress in the network, like `traceroute` does, but it is clearly heavier than an ICMP packet. Measuring the distance between two hosts can take several seconds. The population in our experiments is constrained by the time we can practically dedicate to measure the inter-host distance matrix. Even

in the assumption that the matrix is symmetric, the time cost of the experiment is still $O(n^2)$. To give an idea, a 300 peer set experiment, dedicating an average of 6 seconds to each measurement, takes almost 3 days to complete. *Active measurement methods' cost limit the extent of our experiments.*

4.2 MULTI+ Performance

We measure how MULTI+ performs in two scenarios. The first one is a randomly generated set of peers, and the second is Planet-Lab. In both tests we measure two different parent selection policies when building the multicast tree: FIFO and proximity-aware. In all cases we test MULTI+ when we do not set a limit on the maximum number of connections a peer can accept, and for a limited number of connections, from 2 to 8 per peer.

We measure the following parameters:

- The percentage of peers that connect to the closest peer in the system when they arrive, and the percentage of peers that are connected to the closest peer once the multicast tree is fully built.
- The percentage of the peers present at arrival time closer to one arriving peer than the parent it actually connects to; the percentage of the peers in the total system, when the full multicast tree is built, closer to one peer than this peer's parent. Those figures *exclude* the peers actually connected to the closest peer, and those connected to the source. They measure how bad is the connection made by MULTI+ when this decision is not optimal.
- The fan-out or out-degree of peers in the multicast tree. Obviously leafs in the tree have 0 fan-out.
- The level peers occupy in the multicast tree. The more levels in the multicast tree, the more delay we incur in along the transmission path and the more subject to errors due to peer failure becomes the transmission.

4.3 Random Peer Set

The set of IP addresses used in this scenario is not purely random. There are two reasons for this: one, we cannot measure the distances in an arbitrarily large set of peers. As explained before, measuring the whole matrix of distances for sets exceeding the 300 peers can represent months of work. Two, with a reduced set of peers, if we allow true randomness, we may finish with hosts exclusively located in tier-1 groups, a scenario for which MULTI+ is not conceived (and would be of no practical use.)

Thus, we choose a set of tier-1 groups, and we randomly generate IP addresses *inside* those tier-1 groups. The tier-1 groups are chosen so that we can find at least other tier-2 groups inside.

The high-level characteristics of the TOPLUS tree obtained are as follows: 19 tier-1 groups, 45 inner-groups and 200 peers.

FIFO parent choice

Table 1. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed.

	No restriction	2	3	4	5	6	7	8
to closest peer upon arrival	30	24	27	28	29	29	29	29
to closest peer in full system	12	9	10	9	12	11	11	11




Fig. 8. Fan-out distribution.

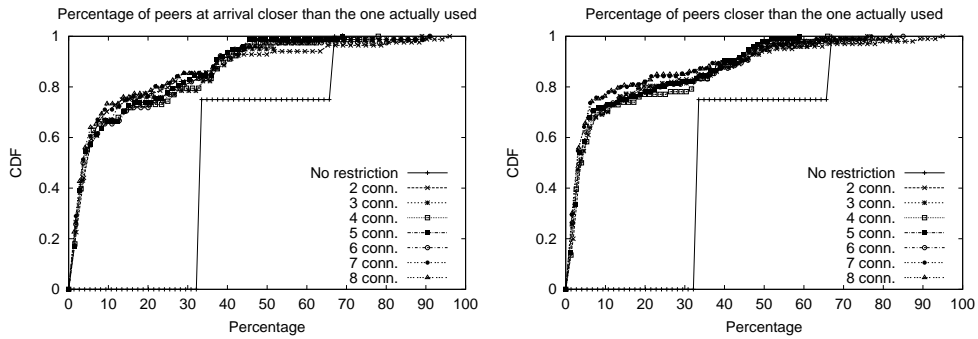


Fig. 9. Percentage of peers found upon arrival (left) and in the full system (right) closer than the one actually used (for those not connected to the source and not connected to the closest peer.)

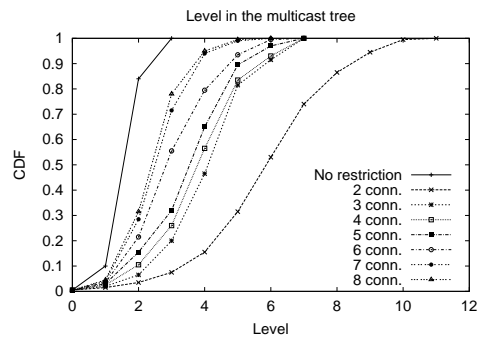


Fig. 10. Level of peers in the Multicast tree.

Proximity-aware parent choice

Table 2. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed.

	No restriction	2	3	4	5	6	7	8
to closest peer upon arrival	31	35	36	40	39	31	32	29
to closest peer in full system	12	18	16	20	18	15	14	12




Fig. 11. Fan-out distribution.




Fig. 12. Level of peers in the Multicast tree.

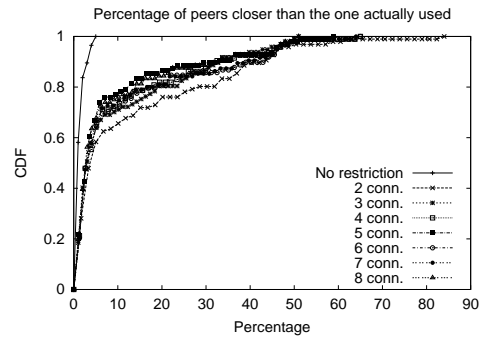



Fig. 13. Percentage of peers found upon arrival (left) and in the full system (right) closer than the one actually used (for those not connected to the source and not connected to the closest peer.)

4.4 Planet-Lab Peer Set

In this scenario we have chosen the hosts in Planet-Lab as peers. There are a number of features that make this set interesting: they are real machines connected today to the Internet; they are more or less distributed around the world (many in the US, but so are most of the IP addresses); and they form a pretty open platform, mainly contributed by universities and first-tier industrial partners, and not obscure ISPs in some corner of the world that may introduce a scarcely representative but very negative bias in the experiment.

The high-level characteristics of the TOPLUS tree obtained are as follows: 30 tier-1 groups, 61 inner-groups and 171 peers.

FIFO parent choice

Table 3. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed.

	No restriction	2	3	4	5	6	7	8
to closest peer upon arrival	42	37	38	38	35	36	37	36
to closest peer in full system	20	27	25	25	25	26	26	25

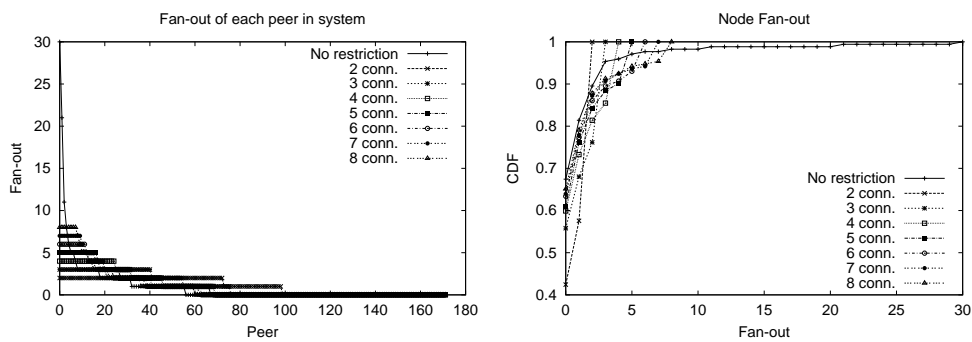


Fig. 14. Fan-out distribution.




Fig. 15. Percentage of peers found upon arrival (left) and in the full system (right) closer than the one actually used (for those not connected to the source and not connected to the closest peer.)




Fig. 16. Level of peers in the Multicast tree.

Proximity-aware parent choice

Table 4. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed.

	No restriction	2	3	4	5	6	7	8
to closest peer upon arrival	42	46	47	45	45	42	43	46
to closest peer in full system	30	30	29	28	27	27	28	29

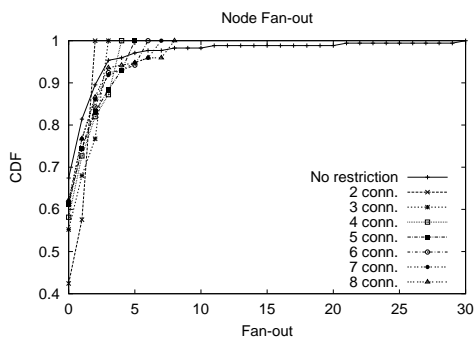


Fig. 17. Fan-out distribution.

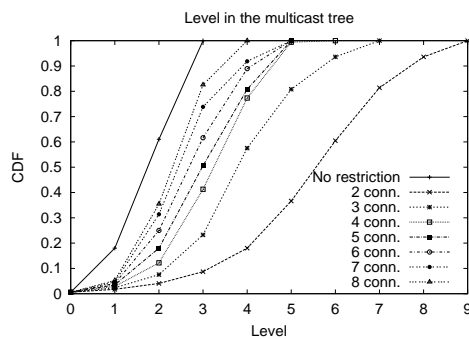


Fig. 18. Level of peers in the Multicast tree.

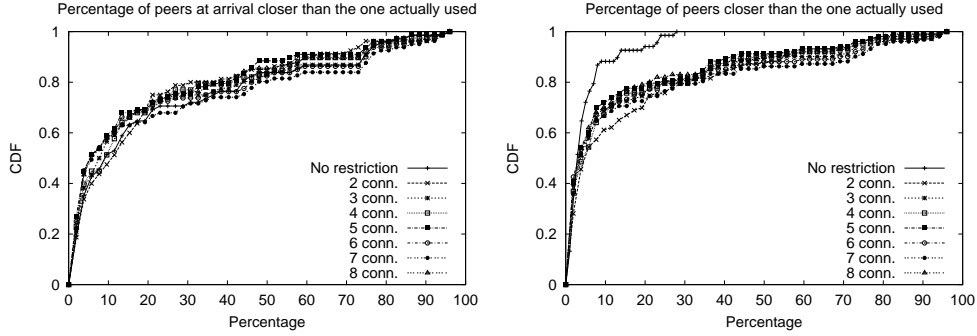


Fig. 19. Percentage of peers found upon arrival (left) and in the full system (right) closer than the one actually used (for those not connected to the source and not connected to the closest peer.)

4.5 Observations

- In all experiments, when there is no restriction in the number of connections, the source receives as many connections as tier-1 groups are in the system (Figures 8 and 14).
- With no connection limits, most peers are organized in a 3-level multicast tree (Figures 10 and 16), which corresponds to most peers being in tier-2 groups in the TOPLUS hierarchy. However, a 4 connections limit per peer allows for a 5-level multicast tree, which are affordable figures. For the 2 connection limit case, the random set tree goes a little out of control with 13 levels, but the Planet-lab experiment achieves 9 levels, actually close to $\log_2(171) = 7.42$.
- Among the non-optimal connections in the Planet-Lab tree, the error (Figure 15, left) can go further than in the random set tree (Figure 9, left), but the number of optimally connected peers is greater in the Planet-Lab tree (Tables 3, 4) than in the random set (Tables 1, 2).
- The proximity-awareness improves the optimality of connections in the Planet-Lab set (Tables 3, 4), but not in the random set (Tables 1, 2). The error in the non-optimal connections is also reduced in the Planet-Lab set with proximity-aware parent selection algorithm (Figure 19, left) compared to the blind one (Figure 15, left). Still, the difference with the blind algorithm is small enough to conclude that MULTI+ itself creates fairly good proximity-aware multicast trees. This is strengthened by the fact that limiting the connections does not affect a lot the optimality of the generated tree.

5 Skitter Coordinate Space Performance

Obviously, the $O(n^2)$ cost of actively measuring the full inter-host distance matrix for n peers limits the size of the peer sets we can use. P2P systems must be designed to be potentially very large, and experiments should reflect this property by using significant peer populations. Some methods ([16, 18, 15, 24]) to map hosts into a M -dimensional coordinate space have been developed. The main advantage is that given a list of n hosts, the coordinates for all of them can be actively measured in $O(Mn)$ time (the distances of the hosts to a set of M landmark hosts, with $M \ll n$). Then, the inter-host distance matrix can be calculated off-line, and the population of our experiment is only limited by CPU power and memory (a far less tight bound).

RIPE Probes. RIPE has deployed a number of machines (probes) that measure the delay among each other. We use the median value of each of these delays measured during the week before our experiments as the distance between their two corresponding machines. We could find one RIPE TTM probe in the Skitter Coordinate space. We found another 15 probes in the 26 bit prefix neighborhood of some IP address in the Skitter Coordinate space. We assume that the delay for the probes is similar to the IP addresses in the same 26 bit prefix, because each of those prefixes represents a network of 64 machines. Euclidean distance D is defined as: $D(x_i, x_k) = \sqrt{\sum_{k=1, \dots, M} (x_{ik} - x_{jk})^2}$, for any two hosts identified by their M -coordinate vectors x_i and x_j .

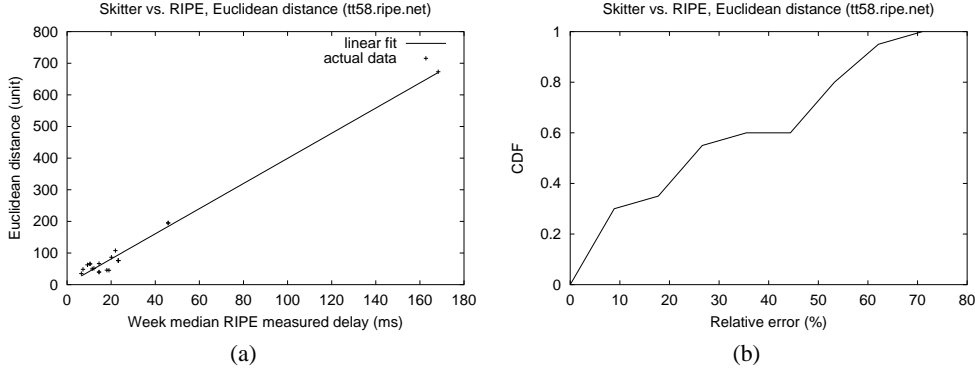


Fig. 20. Euclidean distances compared to measured delay and their relative error to linear fit data.

Another metric as defined in [25] is $D_+(x_i, x_j) = \min_{k=1, \dots, M} (x_{ik} + x_{jk})$.

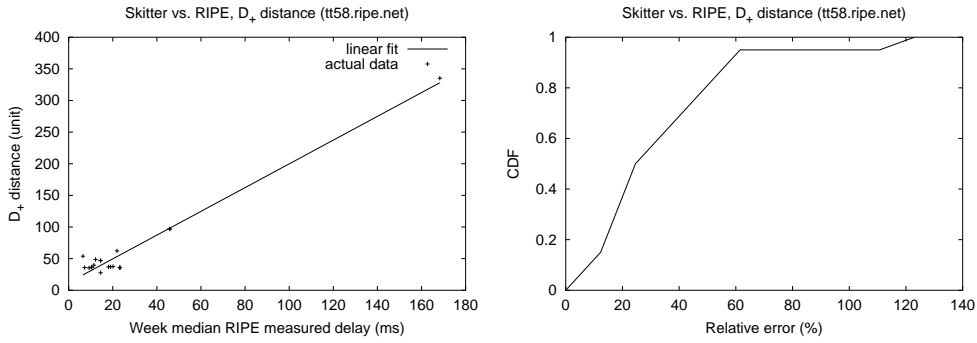


Fig. 21. D_+ distances compared to measured delay and their relative error to linear fit data.

The authors of [25] claim that other metrics than the Euclidean obtain better results, but we did not observe noticeable differences, so we use the Euclidean distance hereafter. We find very good correlation between the measured distances and those calculated with the Skitter coordinates. The RIPE probes should give precise and stable measurements, while other sets may present less accurate results.

5.1 King Measurements

We take a set of 300 IP addresses and we measure their distances to a set of targets using both the Euclidean space based on Skitter Coordinate data and King [11] measurements. We plot King distances on x axis and the corresponding Skitter Coordinate space distances on the y axis, as well as their linear fit. If King were making perfect measurements, the samples should form a line.

In Figure 7 we have shown the relative error of King measurements compared to the information obtained with Ping. We took distances from a host at *Institut Eurecom* to a set of 171 hosts in *Planetlab* [2], thus fairly well-connected machines.

Looking at that graph, one would now expect the coordinates obtained from the Skitter Coordinate space *not* to perfectly follow the distances obtained with King. However, one should expect correlated results for *most* samples.

We measured distances to hosts in the US, Europe, Far East Asia and Australia. The locations were obtained with *NetGeo* [1]. Perhaps the relative error could be improved through weighing samples according to some “subjective” perception. Summarizing, we may say that for most US and European locations, 60% of the calculated distances present at most a 20% relative error. Considering the Far East targets, performance for the same fraction of the distances degrades reaching up to 30-40% relative error. The Australian locations perform similarly.

5.2 USA Targets




Fig. 22. Euclidean distances to 207.138.176.16 and their relative error to linear fit data.




Fig. 23. Euclidean distances to 66.71.191.66 and their relative error to linear fit data.

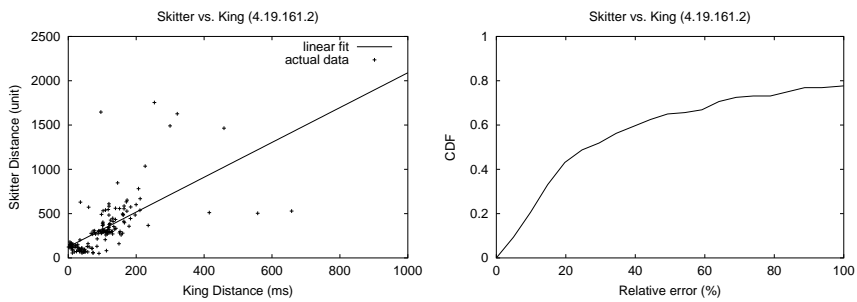


Fig. 24. Euclidean distances to 4.19.161.2 and their relative error to linear fit data.

5.3 European Targets

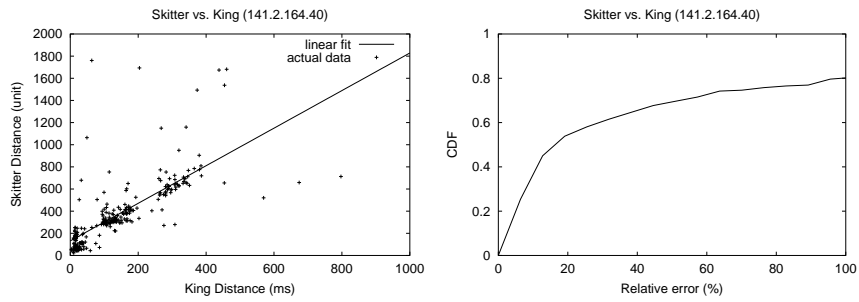


Fig. 25. Euclidean distances to 141.2.164.40 (DE) and their relative error to linear fit data.




Fig. 26. Euclidean distances to 213.76.162.177 (PL) and their relative error to linear fit data.




Fig. 27. Euclidean distances to 138.37.36.200 (UK) and their relative error to linear fit data.

5.4 Far East Targets

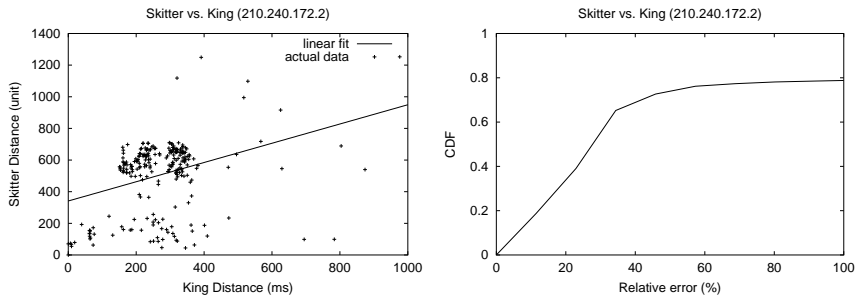


Fig. 28. Euclidean distances to 210.240.172.2 (TW) and their relative error to linear fit data.




Fig. 29. Euclidean distances to 211.174.60.101 (KR) and their relative error to linear fit data.

5.5 Australian Targets

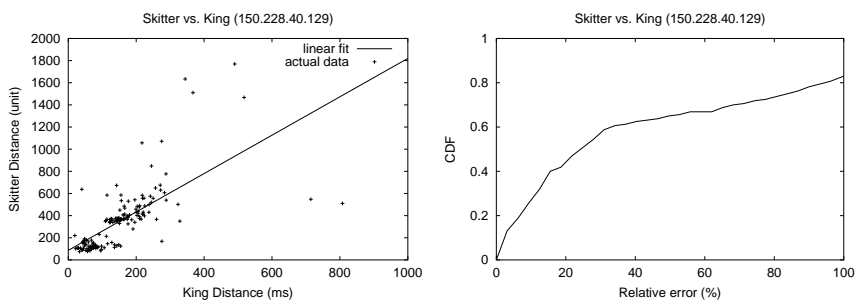


Fig. 30. Euclidean distances to 150.228.40.129 and their relative error to linear fit data.

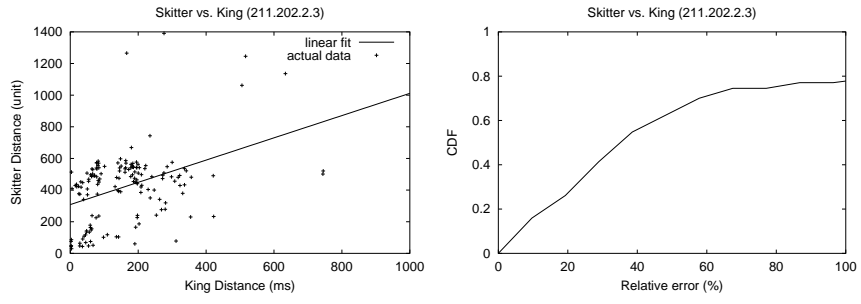


Fig. 31. Euclidean distances to 211.202.2.3 and their relative error to linear fit data.

6 Testing MULTI+ with Skitter Coordinates

6.1 1000 Peer Set

We chose randomly 1,000 peers from the set of IP addresses in the Skitter coordinate space.

The high-level characteristics of the TOPLUS tree obtained are as follows: 47 tier-1 groups, 781 inner-groups and 1,000 peers.

Table 5. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed (FIFO).

	No restriction	2	3	4	5	6	7	8
to closest peer upon arrival	8.4	5.2	5.3	6.7	7.0	7.9	8.0	7.9
to closest peer in full system	6.1	3.5	2.0	4.4	4.9	5.2	5.6	5.5

Table 6. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed (proximity-aware).

	No restriction	2	3	4	5	6	7	8
to closest peer upon arrival	8.4	10.7	12.1	12.3	12.0	12.9	12.0	12.3
to closest peer in full system	6.1	5.9	7.6	7.6	7.7	8.8	7.2	7.5

With the random selection algorithm at most 0.2% of peers connect to the closest peer in the network, so we do not show a table like in the two previous cases.




Fig. 32. Percentage of peers found upon arrival closer than the one actually used (for those not connected to the source.)

We show in Table 6.1 the maximum number of flows going through a TOPLUS group interface. MULTI+ reduces the average and maximum number of flows *per network*, when compared to an arbitrary connection scheme and assuming each TOPLUS group (IP network prefix) represents a physical network.

Table 7. Maximum number of flows through a group interface.

Max. connections	2	3	4	5	6	7	8
FIFO	72	50	46	47	47	42	34
Proximity-aware	45	50	37	27	26	23	18
Random	184	186	197	176	177	191	186




Fig. 33. Percentage of peers in the whole system closer than the one actually used (for those not connected to the source.)

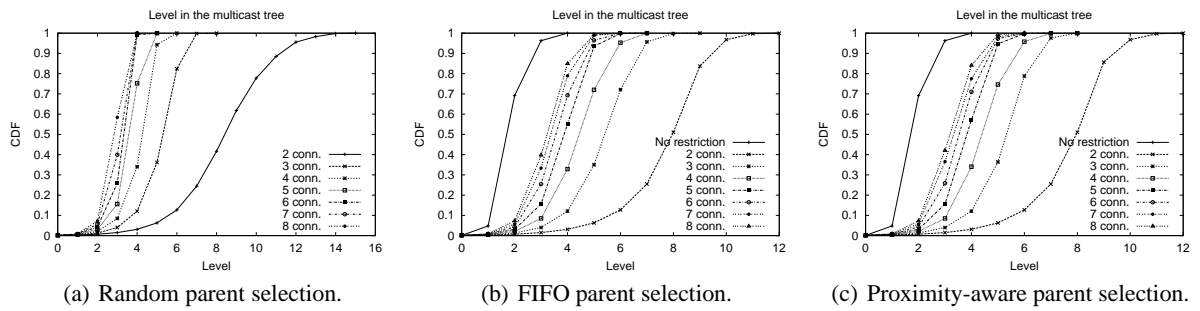


Fig. 34. Level of peers in the multicast tree.




Fig. 35. Latency from root to leaf (in Skitter coordinate units) in the Multicast tree.

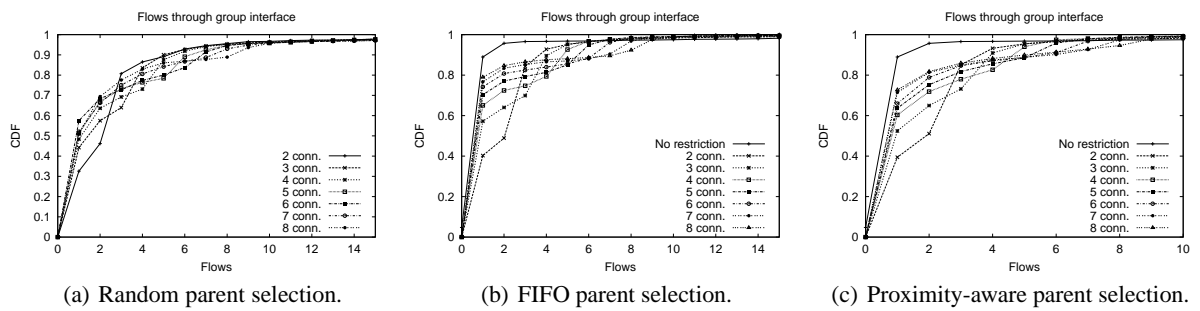


Fig. 36. Number of flows through group interface.

In average, each interface sees 2 flows with MULTI+ (independently of the parent selection policy), while 4 are used with the random algorithm.

Table 8. Average delay from root to leaf.

Max. connections	2	3	4	5	6	7	8
FIFO	2,951(±90)	2,055(±68)	1,706(±71)	1,226(±63)	1,339(±62)	1,256(±62)	1,121(±60)
Proximity-aware	2,319(±68)	1,498(±57)	1,109(±52)	825(±47)	674(±41)	607(±40)	568(±36)
Random	3,493(±106)	2,125(±64)	1,809(±57.27)	1,496(±57)	1,337(±51)	1,208(±49)	1,099(±49)

The average latency from the root of the multicast tree to each destination is smaller for the proximity-aware parent selection algorithm. The more we limit the number of connections a peer can give service to, the better that MULTI+ behaves compared to an arbitrary parent selection algorithm.

6.2 5000 Peer Set

In this experiment we test the characteristics of Multicast trees built with MULTI+ using the Skitter 12-dimensional coordinate space and a set of 5,000 peers. We use the coordinate space to measure the distance between every pair of hosts. In order to make the experiment as realistic as possible, we use a TOPLUS tree with routing tables of reduced size, obtained from the grouping of small and medium-sized tier-1 groups into virtual groups, and this process introduces a distortion in the topological fidelity of the resulting tree [10]. The 5,000 peers are organized into a TOPLUS tree with 59 tier-1 groups, 2,562 inner-groups, and up to 4 tiers. We evaluate the two different parent selection policies described before: FIFO and proximity-aware. We also compare these two approaches with random parent selection. In all cases we test MULTI+ when we do not set a limit on the maximum number of connections a peer can accept, and for a limited number of connections, from 2 to 8 per peer. This limitation also applies to the source. However, in our test we give the source a non-existent IP address, so that our measurements do not get biased by the choice of a given peer. In the test we measure the following parameters, presented here using their CDF (Cumulative Distribution Function):

- The percentage of peers that connect to the closest peer in the system when they arrive (Figure 37).
- The percentage of the peers in the total system, when the full multicast tree is built, closer to one peer than this peer’s parent. Those figures *exclude* the peers directly connected to the source (Figure 38).
- The fan-out or out-degree of peers in the multicast tree. Obviously leaves in the tree have 0 fan-out.
- The level peers occupy in the multicast tree. The more levels in the multicast tree, the more delay we incur in along the transmission path and the more the transmission becomes subject to losses due to peer failure (Figure 39).
- The latency from the root of the multicast tree to each receiving peer (Figure 40).
- The number of multicast flows that go into and out of each TOPLUS group (network) (Figure 41).

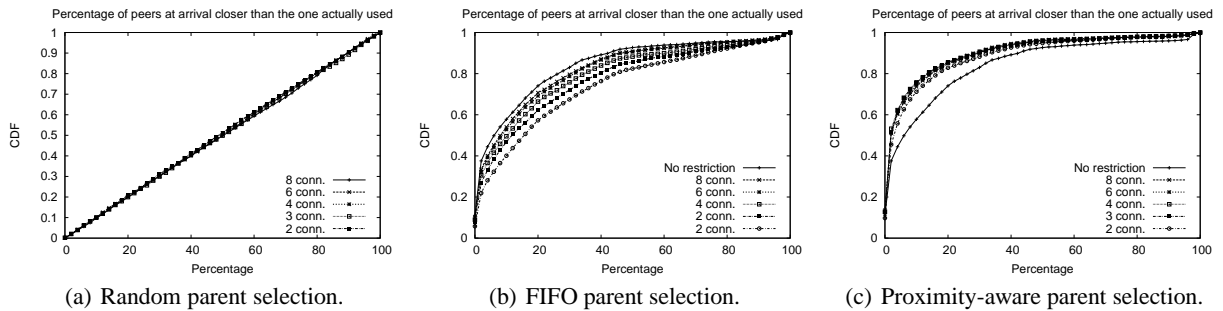


Fig. 37. Percentage of peers found upon arrival closer than the one actually used (for those not connected to the source.)

There are two points that are capital in order to properly evaluate the results of these experiments:

- Many peers (31%) are *alone* in their respective inner groups. Moreover, 65% of the peers in are in inner groups with 3 or fewer peers. This reinforces the impact of the “topology-awareness” introduced by the TOPLUS layout, because connections inside an inner group are rare, and therefore mainly made among *different* inner groups.
- Peers join the system in a completely random order. This means that close peers do not arrive in sequence. In particular, when a peer appears in an inner group where a previous peer exists, the new peer may not find an available connection in the existing peer. To address that problem, one can argue that a peer should reserve a few connections for peers that come in the same inner group, or drop one connection in favor of the newcomer. However, we do not proceed like that, for a number of reasons: first, one cannot tell *a priori* whether peers will appear in a given inner group, and reservation would be a waste of resources; second, given the nature of P2P systems (hosts in edge networks of the Internet), bandwidth is a scarce enough resource itself, specially upstream; and third, given that peers’ connections may be highly volatile, and that the longer a peer has been in the system the more likely it is to stay connected [4], short-term delay minimization optimizations may not pay in the long run in terms of network stability.

Table 9. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed (FIFO).

	2	3	4	5	6	7	8	No restriction
to closest peer upon arrival	5.9	7.9	9.0	9.7	9.8	9.9	10.3	11.1
to closest peer in full system	3.7	5.4	5.9	6.3	6.4	6.6	6.8	7.4

Table 10. Percentage of peers connected to closest peer, depending on the maximum number of connection allowed (proximity-aware).

	2	3	4	5	6	7	8	No restriction
to closest peer upon arrival	9.9	12.9	13.0	13.7	13.7	13.3	12.7	11.1
to closest peer in full system	6.3	8.1	7.8	8.5	8.3	8.2	7.9	7.4

From our experiments we obtain very satisfactory results. From Figures 37 to 41 we draw a number of conclusions:

- Individual peers do not need to support a large number of outgoing connections to benefit from MULTI+ properties: three connections are feasible for broadband users, and the marginal improvement of 8 connections is not very significant.
- The proximity-aware policy performs better than FIFO in terms of end-to-end latency (Figure 40) and connection to closest parent (Figure 38). However, with respect to the number of flows per group (Figure 41) and level distribution in the multicast tree (Figure 39), they are very similar. That is because both trees follow the TOPLUS structure, but the proximity-aware policy takes better decisions when the optimal parent peer has no available connections.
- In Figures 37(c) and 38(c) we can see that having no connection restrictions makes closeness to parent less optimal than having restrictions, for the proximity-aware policy. This is normal, since when we have available connections, a peer’s main goal is to connect as high in the multicast tree as possible. See in Figure 39 how peers are organized in fewer levels, and in Figure 40 how the root-to-leaf latency is better for the unrestricted connection scheme. Still, we can assert that the multicast tree is following (when possible) a topology-aware structure, because most peers connect to nearby parents.
- The random parent selection policy organizes the tree in fewer levels than the other two policies (Figure 39(a)), because connections are not constrained to follow the TOPLUS structure. However those connections are not optimized, and the resulting end-to-end delay performance in any aspect is considerably poorer.




Fig. 38. Percentage of peers in the whole system closer than the one actually used (for those not connected to the source.)

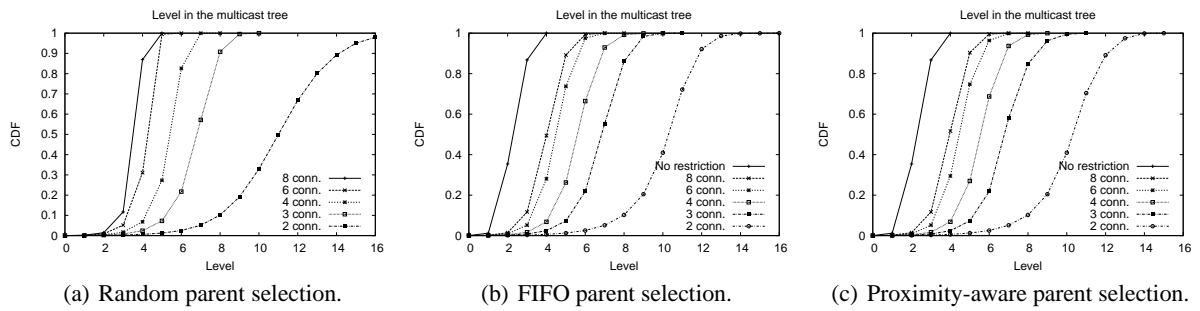


Fig. 39. Level of peers in the multicast tree.




Fig. 40. Latency from root to leaf (in Skitter coordinate units) in the Multicast tree.

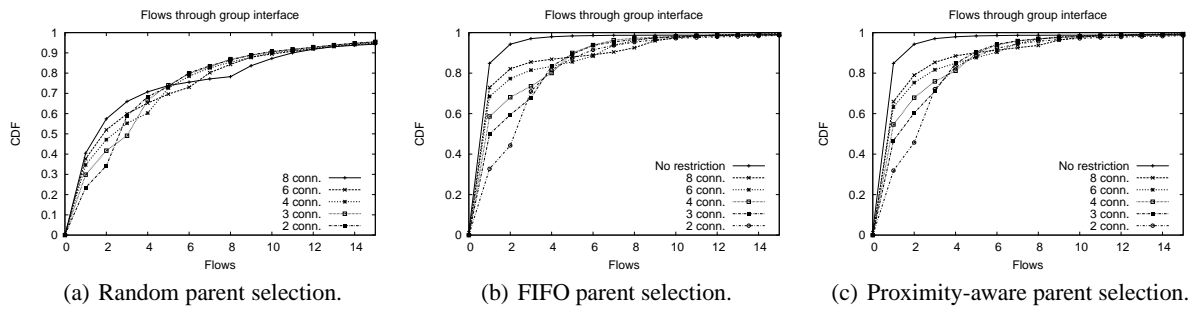


Fig. 41. Number of flows through group interface.

We show in Table 11 the maximum number of flows going through a TOPLUS group interface. MULTI+ reduces the average and maximum number of flows *per network*, when compared to an arbitrary connection scheme and assuming each TOPLUS group (IP network prefix) represents a physical network. On average, each group counts 3 or 4 flows under MULTI+ (independently of the parent choice policy), while 7 are used with the random algorithm. The proximity-aware parent selection clearly improves both the number of flows per group (Table 11) and the end-to-end latency (Table 12). In this 5,000 peer set, even limited to two connections per peer, MULTI+ builds multicast trees with significantly better properties than a random one. Indeed, MULTI+ is conceived for large scale deployments counting many peers, and its properties become asymptotically better the more peers we introduce in the system.

Table 11. Maximum number of flows leaving/entering a group.

Max. connections	2	4	6	8
FIFO	260	145	107	109
Proximity-aware	165	101	39	55
Random	885	907	931	897

Table 12. Average latency from root to leaf.

Max. connections	2	4	6	8
FIFO	3,905	1,641	1,135	965
Proximity-aware	2,871	1,205	797	698
Random	5,339	2,742	2,072	1,646

7 Conclusion and Future Work

We have presented MULTI+, a method to build application-level multicast trees on P2P systems. MULTI+ relies on TOPLUS in order to find a proper parent for a peer in the multicast tree. MULTI+ exhibits the advantage of being able to create topology-aware content distribution infrastructures without introducing extra traffic for active measurement. Admittedly, out-of-band information regarding the TOPLUS routing tables must be calculated offline (a simple process) and downloaded (like many P2P systems today require to download a list of peers for the join process). The proximity-aware scheme improves the end-to-end latency, and using host coordinates calculated offline and obtained at join time (as is done for TOPLUS) avoids the need for any active measurement. MULTI+ also decreases the number of redundant flows that must traverse a given network, even when only few connections per peer are possible, which allows for better bandwidth utilization. As future work, we plan to evaluate the impact of leaving and failing peers on the multicast tree performance.

Acknowledgments

This work could not have been done without the kind help of Prof. Mark Crovella, who provided us with the Skitter coordinate space data of more than 190,000 IP addresses for our experiments.

Institut Eurcom's research is partially supported by its industrial members: Bouygues Tlcom, Fondation d'entreprise Groupe Cegetel, Fondation Hasler, France Tlcom, Hitachi, ST Microelectronics, Swisscom, Texas Instruments, Thales

References

1. "NetGeo-Caida", <http://netgeo.caida.org/perl/netgeo.cgi>.
2. "PlanetLab", <http://planet-lab.org>.
3. S. Banerjee, B. Bhattacharjee, and C. Kommareddy, "Scalable Application Layer Multicast", In *Proceedings of ACM SIGCOMM*, August 2002.
4. F. E. Bustamante and Y. Qiao, "Friendships that last: Peer lifetime and its Role in P2P Protocols" In *Proceedings of 8th Int. Workshop on Web Content Caching and Distribution*, Hawthorne, NY USA, September 2003.
5. M. Castro, P. Druschel, Y. C. Hu, and A. Rowstron, "Topology-aware routing in structure peer-to-peer overlay network", In *International Workshop on Future Directions in Distributed Computing (FuDiCo)*, June 2002.
6. M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "SplitStream: High-bandwidth multicast in a cooperative environment", In *Proceedings of SOSp'03*, New York, USA, October 2003.
7. M. Castro, P. Druschel, A.-M. Kermarrec, and A. Rowstron, "Scribe: a large-scale and decentralized application-level multicast infrastructure", *IEEE Journal on Selected Areas in Communications*, 20(8):1489–1499, October.

8. M. Castro, M. B. Jones, A.-M. Kermarrec, A. Rowstron, M. Theimer, H. Wang, and A. Wolman, "An Evaluation of Scalable Application-level Multicast Built Using Peer-to-peer overlays", In *Proceedings of INFOCOM*, San Francisco, USA, April 2003.
9. M. Freedman and D. Mazieres, "Sloppy hashing and self-organizing clusters", In *Proc. of 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, March 2003.
10. L. Garces-Erice, K. W. Ross, E. W. Biersack, P. A. Felber, and G. Urvoy-Keller, "Topology-Centric Look-Up Service", In *Proceedings of COST264/ACM Fifth International Workshop on Networked Group Communications (NGC)*, pp. 58–69, Munich, Germany, September 2003.
11. K. P. Gummadi, S. Saroiu, and S. D. Gribble, "King: Estimating Latency between Arbitrary Internet End Hosts.", In *Proceedings of 2nd Internet Measurement Workshop 2002*, Marseille, France, November 2002.
12. Y. hua Chu, S. G. Rao, S. Seshan, and H. Zhang, "A Case for End System Multicast", *IEEE Journal on Selected Areas in Communication (JSAC)*, Special Issue on Networking Support for Multicast, 20, 2003.
13. M. M. John W. Byers, Jeffrey Considine and S. Rost, "Informed Content Delivery Across Adaptive Overlay Networks", In *Proceedings of ACM SIGCOMM 2002*, pp. 47–60, august 2002.
14. B. Krisnamurthy and J. Wang, "On Network-Aware Clustering of Web Sites", In *Proc. SIGCOMM 2000*, August 2000.
15. H. Lim, J. C. Hou, and C.-H. Choi, "Constructing internet coordinate system based on delay measurement", In *Proceedings of the 2003 ACM SIGCOMM conference on Internet measurement (IMC-03)*, pp. 129 – 142, Miami Beach, FL, USA, October 2003.
16. E. Ng and H. Zhang, "Predicting Internet network distance with coordinates-based approaches", In *Proceedings of Infocom 2002*, 2002.
17. S. S. M. T. Nicholas J. A. Harvey, Michael B. Jones and A. Wolman, "SkipNet: A Scalable Overlay Network with Practical Locality Properties", In *Proc. of the Fourth USENIX Symposium on Internet Technologies and Systems (USITS '03)*, March 2003.
18. M. Pias, J. Crowcroft, S. Wilbur, and T. H. S. Bhatti, "Lighthouses for Scalable Distributed Location", In *Proc. of 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, March 2003.
19. S. Ratnasamy, M. Handley, R. Karp, and S. Shenker, "A Scalable Content-Addressable Network", In *Proc. ACM SIGCOMM*, 2001.
20. S. Ratnasamy, M. Handley, R. M. Karp, and S. Shenker, "Application-Level Multicast Using Content-Addressable Networks", In *NGC 2001*, pp. 14–29, November 2001.
21. A. Rowstron and P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems", In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pp. 329–350, Heidelberg, Germany, November 2001.
22. S. Shenker, S. Ratnasamy, M. Handley, and R. Karp, "Topologically-Aware Overlay Construction and Server Selection", In *Proceedings of Infocom'02*, New York City, NY, 2002.
23. I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for Internet applications", In *Proc. ACM SIGCOMM*, 2001.
24. L. Tang and M. Crovella, "Virtual Landmarks for the Internet", In *Proceedings of the ACM/SIGCOMM Internet Measurement Conference (IMC-03)*, Miami Beach, Florida, USA, October 2003.
25. S. van Langen, X. Zhou, and P. V. Mieghem, "On the Estimation of Internet Distances Using Landmarks", In *Proceedings of the International Conference on Next Generation Teletraffic and Wired/Wireless Advanced Networking (NEW2AN'04)*, St. Petersburg, Russia, February 2004.
26. Y. Zhu, J. Guo, and B. Li, "oEvolve: Towards Evolutionary Overlay Topologies for High Bandwidth Data Dissemination", *IEEE Journal on Selected Areas in Communications, Special Issue on Quality of Service Delivery in Variable Topology Networks*, Third Quarter 2004.
27. S. Q. Zhuang et al., "Bayeux: An Architecture for Scalable and Fault-tolerant Wide Area Data Dissemination", In *Proc. NOSSDAV'01*, pp. 124–133, June 2001.