

On Global Multicast Networks Using Satellite Unidirectional Links

ACHMAD HUSNI THAMRIN,[†] HIDETAKA IZUMIYAMA,^{††}
HIROYUKI KUSUMOTO^{†††} and JUN MURAI^{†††}

IP Multicast network deployment is still lagging far behind the Internet. Meanwhile, satellite links have been considered as the links that can deliver multicast traffic efficiently because of their broadcast nature. This paper analyzes the deployment of global multicast networks using satellite unidirectional links assuming that the Internet's backbone networks become the UDL feeds, and regional and stub networks become the receivers. The main objective is to understand how and how much satellite unidirectional links benefit the deployment of global multicast networks. We simulate deployment scenarios on several instances of the Internet topology taken from the RouteViews BGP routing table snapshots, where each scenario consists of UDL feeds and receivers placement strategies. The performance of each deployment scenario is measured using the multicast link count, cumulative multicast out-degree, and the expected number of downstreams per UDL receiver. Our simulations demonstrate that such a deployment gives an advantage over deployment using the existing links on the Internet.

1. Introduction

IP Multicast networks have been researched for many years, however the network scale is still far from that of the Internet. Diot, et al.¹⁾ argue that the reason behind the slow deployment of IP multicast network is the current service model and architecture do not provide many features required for a robust commercial implementation of multicast. Even though IP multicast has several problems, we have seen developments in network deployment with the evolution of multicast networks from flat networks, MBONE, to inter-domain networks, and Internet2²⁾. On the multicast infrastructure, Rajvaidia and Almeroth³⁾ find that the infrastructure is slowly but steadily shrinking in terms of the number of advertised addresses; however, the stability of infrastructure that remains is much improved.

Satellite links have long been considered as potential links for delivering multicast traffic because a satellite link has a broadcast nature and covers a large geographic area. However, satellite links are associated with high cost and long delay problems. The development of Link Layer Tunnelling Mechanism (LLTM)⁴⁾ increases the potential of satellite links further by enabling the use of such links as unidirectional links (UDL), instead of point-to-point

links. Using satellite links as unidirectional links not only takes advantage of the broadcast nature of the links but also reduces per-node-costs since the links are shared by many receiving nodes. Meanwhile, the long delay property of satellite links may limit the links use to the applications that are not sensitive to delay, such as multicast file transfer and audio-video multicasting. Global multicast networks using satellite links suit these applications since they are likely to have a small number of — or even a single — multicast sources and a large number of receivers.

A scenario for using satellite unidirectional links for Internet-wide IP multicast network deployment is to use the links as inter-domain multicast links. For example: the Internet's backbone networks setup satellite UDL feeds and their customers (regional or stub ISPs) place receivers at their networks to receive multicast traffic sent via the satellite links. In this scenario, satellite unidirectional links provide shortcuts from multicast sources to receivers, bypassing many terrestrial links and routers in-between. Thus, this scenario creates less multicast states in routers compared to the original network, and therefore preferable.

The above scenario seems to be promising for Internet-wide IP multicast network deployment, given the possibility of bypassing many terrestrial links between multicast sources and receivers. But how does, and how much, adding satellite unidirectional links to the Internet give better multicast networks compared to the original one? How would global mul-

[†] Graduate School of Media and Governance, Keio University

^{††} WISHnet, Inc.

^{†††} Faculty of Environmental Information, Keio University

multicast trees change with the addition of satellite unidirectional links? If an Internet Service Provider (ISP) serves as a UDL feed, how many downstreams it can expect? These are among the questions surrounding the deployment of Internet-wide IP multicast networks using satellite unidirectional links.

This research tries to answer these questions by simulating several scenarios of satellite unidirectional links deployment for Internet-wide multicast network. We first select a set of ASes for UDL feeds by considering their network connectivity and also their international links, then we select UDL receivers based on their network and geographical distances from the feeds. In our simulations, we create synthetic multicast trees on real Internet topology instances. We then introduce satellite unidirectional links to the topologies and compare the resulting multicast trees of each scenarios. We use three metrics to evaluate the benefits: multicast link count, cumulative multicast out-degree, and expected number of downstreams per UDL receiver.

Our study using simulations demonstrates that deploying global multicast networks using satellite unidirectional links benefits the Internet in terms of reducing the number of routers that will need to maintain multicast states, compared to deploying multicast networks using the existing links. Furthermore, who would experience the reduction and how much the reduction would depend on the UDL feeds and receivers placement strategies.

This paper is organized as follows. Section 2 briefly mentions related work in the Internet topology and multicast trees. Section 3 shows our approach to determine feeds and receivers of satellite unidirectional links for Internet-wide multicast networks. Section 4 describes our simulation setup and the performance metrics. Section 5 and Section 6 presents the simulation results. Section 7 concludes this paper.

2. Related Work

There are several work that study the usage of satellite links as multicast path. Almeroth⁵⁾ study the usage of satellite links for The Multicast Backbone (MBone) and showed that satellite links provide lower packet loss compared to terrestrial links. Filali^{6),7)} describe the issues of IP multicast service over satellite-terrestrial hybrid networks and explores how to plan such networks for multicast service by minimizing

the mean cost of multicast trees.

Research in the topology of the Internet^{8)~10)} in recent years provide the research community with the understanding on how the Internet looks like. Dolev and Mokryin¹¹⁾ use the results of these research to create synthetic multicast trees and develop the Fast Algorithm that estimates the number of multicast receivers based on the findings from generated multicast trees.

Chuang and Sirbu¹²⁾ conclude via simulations that the average number of multicast path L to reach n nodes where the unicast path is U is $L/U = \Theta(n^{0.8})$. Adjih, et al.¹³⁾ analyze multicast tree topologies and conclude that the power-law property of multicast tree is likely to hold in real network topologies, even though it is very difficult to infer the validity using the methods in Ref. 12). Chalmers and Almeroth^{14),15)} use a metric to measure multicast efficiency using multicast and unicast hop counts introduced in Ref. 16), which is defined as $\delta = 1 - L_m/L_u$, where L_m is the total number of multicast links in the distribution tree and L_u is the sum of all unicast hops. δ represents the percentage gain in multicast efficiency over unicast.

3. IP Multicast Network Deployment using Satellite UDLs

One of the reasons behind the slow deployment of Internet-wide IP multicast network is the management cost¹⁾, including router states. Satellite links can provide shortcuts from multicast sources to receivers, thus removing multicast states in some middle routers on the terrestrial links. With the broadcast nature and wide-area coverage of satellite links, satellite links are thought to be a viable solution to serve as the backbones of the global IP multicast network.

From this perspective, it is preferable to locate satellite UDL feeds as close as possible to the multicast sources. Placing as many receivers as possible under each link is also preferable, because multicast traffic does not traverse many terrestrial routers. At one extreme, if feeds are to be placed in every networks with possible multicast sources, then we would have many feeds with huge cost implications because a satellite transmission station is considerably expensive. On the other hand, a large number of receivers on a satellite UDL may lead to scalability issues. Both these situations would

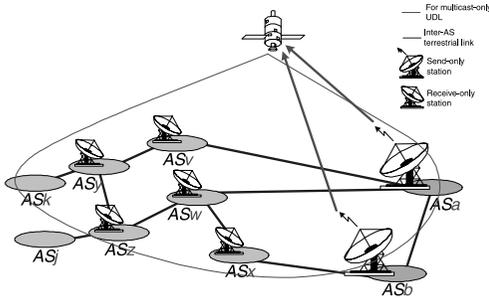


Fig. 1 Internet-wide IP multicast network with satellite UDLs.

limit the number of feasible satellite unidirectional links and the scale of the links. Therefore with these limitations, this research opts to explore multicast efficiency with small number of feeds to deploy an Internet-wide IP multicast network using satellite UDLs.

A plausible deployment scenario is the Internet's backbone networks serve as satellite UDL feeds, and regional and stub networks becomes the UDL receivers, where the terms of backbone, regional, and stub networks are according to Govindan and Reddy¹⁷⁾. In this scenario, backbone networks also have their inter-domain terrestrial links multicast-enabled, while terrestrial links between receiver networks may, or may not, be multicast-enabled. **Figure 1** is an illustration of this deployment scenario. *ASa* and *ASb* are backbone ASes, and they have satellite UDL feeds. ASes with UDL receivers forward multicast traffic from *ASa* and *ASb* to their terrestrial links to reach receivers, including those in *ASj* and *ASk*.

In this section we discuss the above scenario to deploy such network by looking at the Internet topology, and selecting ASes for UDL feeds and receivers. We use the Internet topology data from BGP routing table snapshots provided by *RouteViews*, taking six instances of the routing table snapshots, which are the ones in November 2001, in May and November of 2002, 2003, and also in May 2004.

3.1 Unidirectional Link Feeds

We first investigate UDL feeds candidates for satellite unidirectional link deployment by examining the backbone networks of the Internet. Backbone networks are defined by Ref. 17) based on network vertex degrees. These networks have very high number of degrees compared to the average, as shown by the skewed

distribution of the Internet topology^{8),9)}, and are located at the center of the Internet. But a high vertex-degree network does not automatically constitute an Internet core network¹⁸⁾, therefore a high-vertex degree network must be confirmed whether it is an Internet core network.

We start our examination of UDL feed candidates using the Internet's *dense core*. Subramanian, et al.¹⁸⁾ define the dense core as the largest subset of ASes where every node in a directed graph of N ASes has an in-degree and out-degree of at least $N/2$. Looking at the Internet from many points is necessary to get the directed graph of the Internet and to find its dense core. However, we do not use multiple vantage points to view the Internet topology, thus we relax the requirement of dense core from a directed graph into an undirected graph. We define the *undirected dense core* as the largest subset of ASes where every node in an undirected graph of N ASes has a degree of at least $N/2$. Details on the algorithm to find undirected dense follows directly from Ref. 18) with the graph requirement relaxed into an undirected one.

We map ASes into their geographic locations in the country name granularity using the Internet Registries' WHOIS¹⁹⁾ database, and we check the undirected dense core against the database. The reasons to use geographic locations are: 1. satellites have limited footprints; and 2. intuitively, geographic distance plays a part in a decision to connect two networks using a satellite link. An AS may span several countries, especially for transit ASes, thus mapping the geographic location of an AS in the granularity of country is still coarse and prone to errors. With these situations, we assume that the central location of an AS is in the country name of the AS that is registered to the Internet Registries database, and that would be the geographic location of the satellite UDL feeds or receivers.

In our mapping we notice that there are ASes in the undirected dense core that do not have international or inter-regional link. On the other hand, there are ASes that are not in the undirected dense core but they have such links. For example AS702 and AS703, which belong to UUNet, are connected to AS701 to reach other dense core ASes but both ASes have many international links. We also examine the ASes in undirected dense core and all other ASes down

Table 1 Undirected dense core and international backbone ASes.

BGP instance	num. of AS	size of undir. dense core	num of intl. backbone	undir. dense core in intl. backbone	lowest degree-rank
Nov. 2001	12312	21	24	14	52
May 2002	13298	22	37	18	136
Nov. 2002	14335	23	26	18	61
May 2003	15424	25	29	16	69
Nov. 2003	16413	29	53	25	124
May 2004	17579	31	50	26	117

to the lowest degree-rank AS in the core that is only one hop away from the undirected dense core to find whether these ASes provide national or inter-regional transit.

The definition of backbone network in Ref. 17) in terms of international or inter-regional connectivity is very loose, for example: can a high degree AS with just one international or inter-regional customer be classified as an international backbone? We use the following approach to identify whether an AS is an international backbone network. First, we set the AS in question as the root AS of a graph and then we list the neighbors of the root AS whose degree is lower than the lowest degree AS in undirected dense core, and mark them as first-hop ASes. For every first-hop AS, we list its neighbors whose degree is not larger than itself and mark them as second-hop ASes. We then compare the geographic location of the root AS with every first-hop and second-hop ASes, whether they are located in the same country or not. If we see a significant number of ASes that are not in the same country as the root AS, and the number increases in the second-hop ASes, then we define the root AS as an international backbone AS. Defining *significant* is indeed difficult, but for our purpose we define that the minimum number of both international and inter-regional link is 3. We understand that this definition does not draw a strict line to distinguish international backbone ASes, but we believe that it provides a better definition compared to Ref. 17).

Table 1 summarizes the undirected dense core and international backbones that fit into our definition. We will use these two sets of ASes for the UDL feeds in our simulations.

3.2 Unidirectional Link Receivers

After assessing the candidate ASes for UDL feeds, we now turn to the candidates for UDL receivers. A satellite link generally covers a large geographic area, and all Autonomous Systems having a Point of Presence (POP) within the footprint of the satellite link can be con-

nected each other using the link. But intuitively, ASes within the footprint are not likely to be connected using the satellite link when a terrestrial link is easy to deploy and will provide a better connectivity. We model the decision of an AS to become a UDL receiver using a probability value based on the geographic and network distances of the AS from the UDL feed.

We generate synthetic UDL receivers using the same traversal method as in the previous subsection when we identified international backbones. Our assumptions are an AS may only be a UDL feed or a receiver, and when an AS is a receiver, it has only a single UDL receiver. The ASes found during the traversal are likely to be the customers of the root AS²⁰⁾, the customers of these customers, and so on. Traversals from different feeds may coincide a same AS. In this research, when an AS is traversed more than once, the AS randomly picks its feed among the traversing root ASes.

The traversed ASes form UDL receiver candidates, and each AS has a probability to become a receiver based on its geographic and network distances from its feed. We denote the probability of an AS to become a UDL receiver using a matrix of probabilities where the row and columns are the network and geographic distances from feed, respectively. Each entry in the matrix represents the probability of an AS of a certain network and geographic distance from the feed to become a UDL receiver. We classify geographic distances into three classes: same country, different country but same region, and different region; where regions in our research are Asia, Oceania, North America, South America, Europe, Africa, and Middle-East. The purpose of this geographic distance classification is to model that the probability of an AS to become a UDL receiver of a feed also depends on how far is the distance between the AS and the feed. We limit the network distance of UDL receivers to three AS hops from the feed because we find that almost all of the networks can be reached within three

$$\text{ROW2} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix} \quad \text{COL2-3} = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

Fig. 2 Two of basic scenarios for generating UDL receivers.

AS hops, therefore each scenario is a 3×3 matrix.

Our basic UDL receiver generation scenarios are matrices whose entries are either 1 or 0. These scenarios represent the most number of receivers for each distance type. We use $ROWx$, $1 \leq x \leq 3$ to denote a basic scenario where the UDL receivers are located in x AS-hop distance from the UDL feed, Meanwhile, $COLy$, $1 \leq y \leq 3$ is to denote a scenario where the UDL receivers are located in the same country ($y = 1$), or in a different country but within the same region ($y = 2$), or in another region ($y = 3$). **Figure 2** shows two basic scenario matrices that represent the scenarios where the UDL receivers are the ASes whose network distance is two hops from the feed (ROW2), and ASes that are located in different country but in one region, or located in different region (COL2-3).

4. Simulations and Metrics

We simulate our UDL receiver generating scenarios against synthetic multicast trees on the Internet topologies derived from the BGP routing snapshots of Routeviews. The multicast receivers are randomly placed on the network and the number of multicast receivers are up to 1000. We use two classes of multicast source:

- (1) the root AS has a higher vertex degree than the minimum vertex degree of international backbone networks;
- (2) root AS has a vertex degree less than the type 1.

Multicast sources are also randomly selected from each class.

We first simulate several basic scenarios to know the results if all of the UDL receiver candidates in a scenario are willing to become UDL receivers. After generating synthetic multicast trees, we add satellite unidirectional links to the underlying network and modify the trees. For each multicast receiver, we traverse the path from the receiver up to the root of the tree, and if we find an AS that is a UDL receiver, the rest of the path to multicast source is replaced by the UDL feed and the subsequent ASes upto the multicast source. **Figure 3** illustrates a

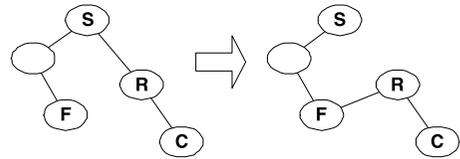


Fig. 3 Multicast path between S and C is modified by a unidirectional link.

unidirectional link, where F is the feed and R is a receiver, that modifies the multicast path between multicast source S and receiver C .

The assumptions used in our simulations are:

- (1) All ASes in the network are multicast-enabled using inter-domain multicast routing protocols.
- (2) Multicast trees are rooted at the multicast source.
- (3) Only a single multicast receiver exists within an AS.
- (4) Multicast routing sets the unidirectional links with the lowest routing cost, hence UDLs are always preferred even though the resulted paths may not be the shortest ones at the AS level.

We also simulate scenarios where not all of the UDL receiver candidates are willing to become UDL receivers. The simulations for each data-set in this research are using 144 combinations of randomly selected unidirectional links (feed and receivers) and multicast trees.

We explore the performance of each scenario using these metrics: multicast link reduction, cumulative multicast out-degree, and expected number of downstreams per unidirectional link. The multicast link count of a multicast tree is the total number of links in the tree. Deploying satellite unidirectional links will change the multicast link count of multicast trees. In counting the multicast link we assume that each inter-domain link is a point-to-point link, except for an unidirectional link. For example, the multicast link count of links between AS A and its two customers B and C is 2. Meanwhile, if a unidirectional link exists where A is the feed while B and C are the receivers, then the link count is 1.

A multicast link count demonstrates the total links in a global multicast tree, but this metric does not show which ASes have the links, thus keep the multicast states for the tree. To know which ASes have the links, we introduce a notion of multicast out-degree distributions on a network, which we call *cumulative multicast out-degree*. Let us define the multicast out-

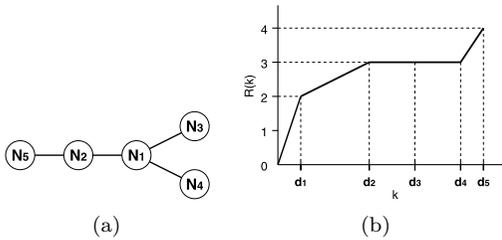


Fig. 4 Multicast tree and its $R(k)$.

degree of a node as the number of outgoing links of the node in a multicast tree. The multicast out-degree hence implies how many copies of multicast packet a node has to send. The cumulative multicast out-degree $R(k)$ is the sum of the multicast out-degree of nodes whose degree-rank is less than or equal to k . That is

$$R(k) = \sum_{i=1}^k m_i$$

where m_i is the number of multicast out-degree of a node whose degree-rank is i . **Figure 4** illustrates a multicast tree and its $R(k)$. Suppose there is a multicast tree whose source is N_5 , and N_3 and N_4 are the receivers. The degree-rank of N_i is d_i , and $d_i < d_{i+1}$. The $R(k)$ graph of this tree is shown in Fig. 4 (b).

In our research, a node in $R(k)$ represents an Autonomous System on the Internet. $R(k)$ can capture the distribution of multicast out-degree in backbone networks, regional, and stub networks because of the power-law property of the Internet as AS level. Considering that the UDL feeds in our research are ASes whose degree-rank is low, $R(k)$ of global multicast trees for low k will be reduced by satellite unidirectional links deployment. The reduction of multicast out-degree in high degree-rank ASes is more important than in feeds because it shows that non-feed ASes enjoy the advantages of the deployment. We use this metric to measure a deployment scenario by comparing the cumulative states in the feeds with the rest of the network.

A satellite unidirectional link is a broadcast link, thus if the link only has a single downstream, it means that the link is not utilized efficiently. The number of expected downstreams per UDL receivers is also a metric for deployment scenarios in our research. We can expect three factors will determine the expected number of downstreams of a satellite unidirectional link: 1. the location of UDL receivers relative to the feed and to stub networks, 2. the number of UDL receivers, and 3. the number and

affinity of multicast receivers.

5. Multicast Link and Out-degree Reductions

In this section we present the simulation results in terms of multicast link and cumulative multicast out-degree reduction. We start with the results of basic scenarios and then give the results in the cases where not all candidate UDL receivers will become receivers.

5.1 Basic Scenarios

Figures 5, 6, 7 show the results of deploying satellite unidirectional links on the May 2004 instance of RouteViews. Figure 5 shows the average multicast link count of multicast trees whose source is in a low vertex degree AS. The feeds of these satellite unidirectional links are in the international backbone. The results for simulations using undirected dense core feeds are displayed in Fig. 6. Figure 7 depicts the results of multicast trees whose source is in a high vertex degree AS. Figure 7 (a) is for international backbone feeds, and Fig. 7 (b) is for undirected dense core feeds. The horizontal axis of these figures is the number of multicast receivers, and vertical axis is the average number of multicast links in the synthesized tree.

We can see that the results with sources in high vertex degree ASes are better than those with sources in low vertex degree ASes because the former are closer to the core of the Internet. Comparing scenarios ROW1, ROW2, and ROW3, i.e., scenarios with different network distances, we find that ROW2 gives better multicast link count reduction among the three scenarios. The reason for better reduction by ROW2 is because the UDL receivers in this scenario are closer to the core of the Internet than those of ROW3, therefore the probability of a UDL receiver to have a downstream is higher than the UDL receivers in ROW3, that are mostly stub networks. Scenario ROW1, which is the closest one to the core, does not give a better result than ROW2 because some of the multicast paths are already joined before reaching ASes in ROW1. This situation also explains why ROW2-3 better reduces multicast link count compared to ROW1-2.

Scenario COL2 (UDL receivers are in the same region but different countries) is the worst performer of all scenarios in our simulations, and the multicast link reduction of undirected dense core feeds is worse than that of the international backbone feeds. On the other hand,

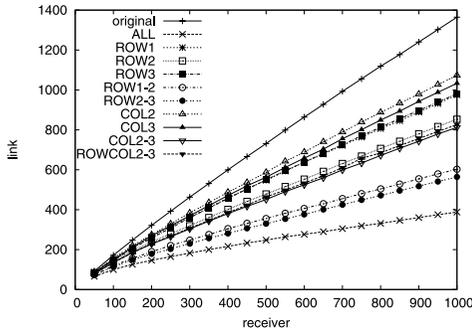


Fig. 5 Multicast link count for international backbone feeds (source in low vertex degree AS).

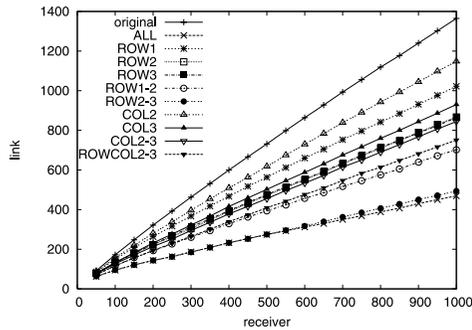


Fig. 6 Multicast link count for undirected dense core feeds (source in low vertex degree AS).

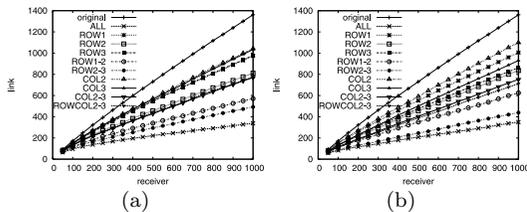


Fig. 7 Multicast link count for trees whose source in high vertex degree AS.

COL3 gives better results for undirected dense core feeds compared to the results for international backbone feeds. These results are the reflections of the feed’s geographic locations. The number of ASes in the undirected dense core feeds that is not located in North America is 11 (Europe 9, Asia 2). This number in the international backbone feeds is 26 (Europe 19, Asia 6, Oceania 1). The proportion of these feeds contributes to the difference of the results between COL2 and COL3. The results from the Routeviews instance in May 2003 (not shown in here) exhibit similar results. In the May 2003 instance, the number of ASes not located in North America in undirected dense core is 5 (Europe 3, Asia 2), compared to 10 (Europe 6, Asia 4) in international backbone.

A likely scenario for satellite unidirectional links for multicast networks is UDL receivers are located more than one-hop away from the feed and not in the same country as the feed. This situation is modeled as a basic scenario matrix whose entries $s_{i,j} = 1, i > 1, j > 1$, where i and j are the row and column of matrix. We denote this matrix as ROWCOL2-3. This scenario gives better results for undirected dense core feeds. In addition, ROWCOL2-3 reduces multicast link better than COL2-3 using the undirected dense core feeds. But on the deployment using international backbone feeds, the results of ROWCOL2-3 are slightly worse than those of COL2-3. This is because international backbone feeds by definition are connected to more networks located not in the feed’s countries compared to undirected dense core feeds. This slightly worse multicast link count in our simulations shows that only a few of networks located in 1 network distance and in different countries than the feeds that become UDL downstreams.

Figures 8 and 9 show the average cumulative multicast out-degree from the simulation results of our deployment scenarios at 750 multicast receivers. Figure 8(a) and Fig. 9(a) are for the international backbone feeds, while Fig. 8(b) and Fig. 9(b) are for the undirected dense core feeds. In these figures a node is an Autonomous System, and the displayed results are for the degree-rank up to 10-thousand ASes. The rest of the ASes are not displayed because they bring only an insignificant change to $R(k)$. We separate the scenarios having a single non-zero valued column or row from our other scenarios for clarity. These figures depict the sources of multicast link count reduction using each scenario in our simulations.

The results show that scenario ROW1-2 that gives a very good multicast link reduction turns out to benefit high vertex degree ASes (satellite UDL feeds), and not the low ones where UDL receivers and multicast receivers are located. These results indicate that multicast link reduction alone is not enough to measure the performance of satellite unidirectional deployment scenarios. Using this metric, we measure how much a scenario reduces the multicast out-degree to non-backbone ASes, i.e., the non-feed ASes whose vertex degree is smaller than any undirected dense core ASes. From the perspective of non-feed ASes, the scenario that gives non-feed ASes the least multicast out-degree

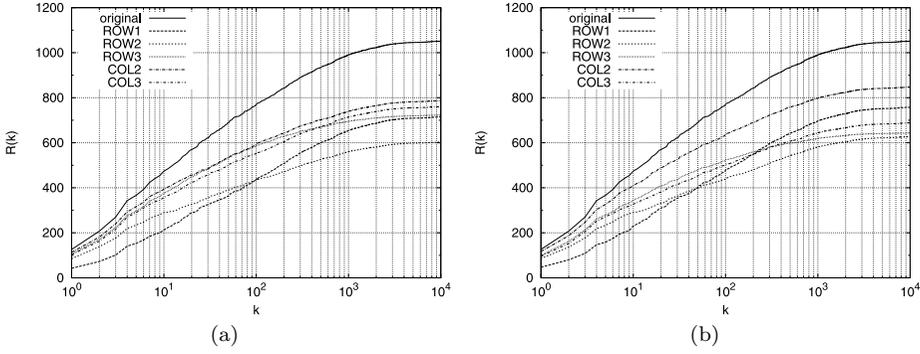


Fig. 8 Cumulative multicast out-degree using COL_x and ROW_x scenarios.

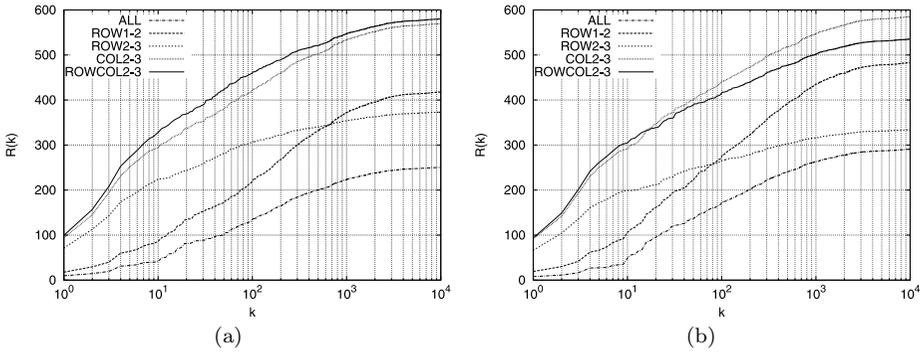


Fig. 9 Cumulative multicast out-degree using other scenarios.

Table 2 Cumulative multicast out-degree $R(k)$, $k > 117$.

International backbone feeds			
scenario	$R(117)$	$R(17579)$	$R'(117)$
original	787	1053	266
ALL	141	251	109
ROW3	597	724	127
ROW2-3	310	374	63
ROWCOL2-3	468	581	113

Undirected dense core feeds			
scenario	$R(117)$	$R(17579)$	$R'(117)$
original	787	1053	266
ALL	177	291	114
ROW3	530	645	115
ROW2-3	269	335	65
ROWCOL2-3	421	536	115

sum is the most preferred. **Table 2** lists the sum of multicast out-degree for such ASes using the May 2004 instance. This table gives the average cumulative multicast out-degree (rounded down to the nearest integer) of the AS with the vertex degree-rank of 117, i.e., the least-vertex degree of the undirected dense core, and $R'(117) = R(17579) - R(117)$, which is the total multicast out-degree in ASes that have less vertex degrees than the least-vertex degree AS of the undirected dense core. Based on this table,

scenario ROW2-3 provides the least multicast out-degree for ASes where UDL receivers are located, much better than scenario ALL that gives the best multicast link count reduction. Scenario ROWCOL2-3 that we believe to be a likely scenario for satellite unidirectional link deployment exhibits good results, and differs with ALL only by a little.

5.2 Scenarios with Smaller Number of UDL Receivers

Basic scenarios set each value in the scenario matrix entries with either 0 or 1. We also simulate the cases where the values of scenario matrices are between the 0 and 1. **Figure 10** display the multicast link count results of several scenarios based on the basic scenarios with 750 multicast receivers. Figure 10(a) depicts the average multicast link count versus the probability p , $0.1 < p < 1$ of basic scenarios, for example 0.2ROW2-3. Figure 10(b) shows the average multicast link count versus the number of UDL receivers of the scenarios in Fig. 10(a). Figure 10(a) and 10(b) show the results of the May 2004 instance. Meanwhile, Fig. 10(c) depicts the average multicast link count of the May 2003 instance.

Figure 10(a) shows that the multicast link

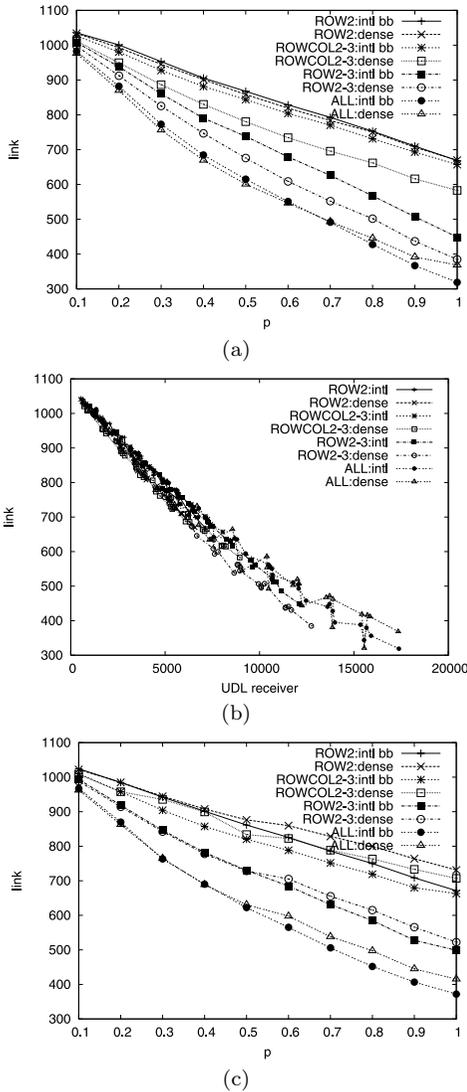


Fig. 10 Multicast link count for smaller number of UDL receivers.

count of ROW2 is linear to p . Other scenarios in the figure exhibit a saturation trend as p increases; this is caused by the number of UDL receivers that approaches the total number of ASes in the network as we increase the p of these scenarios. The total number of UDL receivers of the ROW2 scenario with undirected dense core feeds is 6580, while for ROW2-3 it is 12373. The number for ALL is 17345, which is very close to 17579, i.e., the total number of ASes. We confirm the saturation trend with Fig. 10 (b) where we find that the multicast link count reduction is linear only for small number of UDL receivers.

The multicast link count of the May 2003 in-

stance also exhibits a saturation trend as the number of UDL receiver approaches the total number of ASes. The main difference that we find between the two instances is the international backbone feeds give better multicast link count reduction in the May 2003 instance instead of the undirected dense core feeds as in May 2004 instance. This phenomenon can be explained when we compare the two instances in terms of the number of the undirected dense core and the international backbone relative to the network size. The May 2004 instance has 50 international backbone ASes, which is a big increase from 31 ASes in undirected dense core. On the other hand, the number of AS's difference between the two feed sets in the May 2003 instance is only 4. With the properties of small world networks exhibited by the Internet and the way we traverse the UDL receiver candidates, more feeds doesn't mean an increase in the number of receivers because an AS is more likely to be traversed more than one, thus decreasing the average number of downstreams per unidirectional link.

Table 3 lists the cumulative multicast out-degree when the average total number of UDL receivers is 2500 using scenarios that are based on the basic scenarios. This table shows that scenarios 0.5ROW3 and 0.41ROW3 give the least $R'(117)$ for international backbone and undirected dense core feeds. These scenarios are followed by p ROWCOL2-3 and p ROW2-3. The three basic scenarios, i.e., ROW2, ROW2-3, and ROWCOL2-3, are also among the best performing scenarios in terms of $R'(117)$ (Table 2). Besides that, these three basic scenarios also demonstrate good multicast link count reductions. With these results, we believe that deployment scenarios based on these three scenarios should be prioritized by ISPs that will act as satellite UDL feeds.

6. Downstreams of UDLs

In this section we look at the number of downstreams for each unidirectional links in our simulations, where multicast receivers are sparse and randomly selected. We plot the simulation results to show the relationship between the number of receivers and the number of downstreams for different scenarios. **Figure 11** depict this relationship for multicast trees of the May 2004 instance with 750 receivers. The horizontal axis of Fig. 11 is the number of UDL receivers for each feed while the vertical axis is the

Table 3 Cumulative multicast out-degree of scenarios with avg. 2500 UDL receivers.

scenario	intl. backbone			undir. dense core		
	p	$R(17579)$	$R'(117)$	p	$R(17579)$	$R'(117)$
ALL	0.14	892	229	0.14	901	234
ROW1	0.47	891	266	0.54	890	267
ROW2	0.35	888	228	0.38	884	231
ROW3	0.50	885	193	0.41	877	201
ROW1-2	0.20	906	246	0.22	903	247
ROW2-3	0.21	892	214	0.20	885	217
COL2	0.56	891	215	0.74	905	218
COL3	0.53	881	226	0.40	883	228
COL2-3	0.27	882	217	0.26	877	221
ROWCOL2-3	0.31	883	209	0.28	878	215

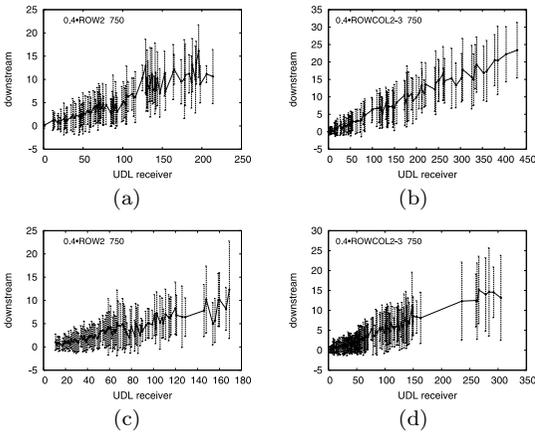


Fig. 11 UDL receiver-downstream relationship.

number of downstreams. The number of downstreams are displayed using their average with their 90% confidence interval error bar. Figure 11 (a) and Fig. 11 (c) show the results using undirected dense core feeds, while Fig. 11 (b) and Fig. 11 (d) are using international backbone feeds.

The figures show that we might expect a linear relationship between the number of UDL receivers and the number of downstreams, given the number of multicast receivers and the probability of an AS to become an UDL receiver, even though the confidence interval is large. We can also see that the maximum numbers of UDL receivers of undirected dense core feeds are higher compared to those of the international backbone feeds. Also, international backbone feeds exhibit a clustering for scenarios p ROWCOL2-3, which is less evident in p ROW2. Further examination reveals that feeds having many UDL receivers are located in Europe, thus this phenomenon is a result of including geographic distance in both UDL feeds and UDL receivers selections.

The average downstreams per UDL receiver

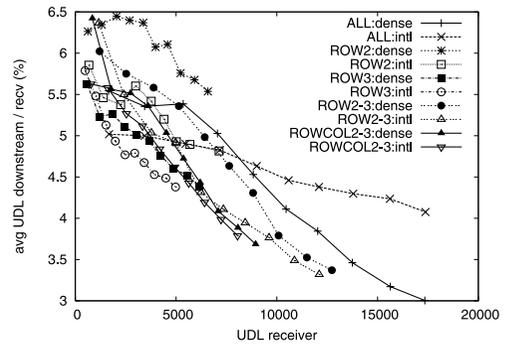


Fig. 12 Percentage downstream per UDL receivers.

(in percent) for the May 2004 instance is depicted in **Fig. 12**. This figure displays the results for various total number of UDL receivers using scenarios in the previous section. From the perspective of ISPs serving as satellite UDL feeds, the higher the percentage is the better because it means that the links are utilized more efficiently. We can view the links are under-utilized if the percentage of downstream per UDL receiver is lower than the number of multicast receiver per total number of ASes. In this case, the number is $750/17579 = 4.27\%$.

This figure shows that the expected number of downstream per UDL receiver is higher for undirected dense core feeds compared to that of the international backbone feeds, which is simply due to the difference of the number of feeds between them. Also, scenario ROW2 of undirected dense core gives the best results except for a small range of UDL receivers. This result indicates that there is a trade-off to reduce multicast out-degree in receiver networks or to increase the downstream per UDL receiver. Among the three scenarios that give the best multicast out-degree reduction in receiver networks, we find that the scenario ROW2-3 gives the best value despite that the value drops below 4.27% at large UDL receivers.

The scenarios displayed in Fig. 12 exhibit a tendency to give lower percentage as the number of total UDL receivers increases. Figure 12 shows that there is a range of UDL receivers (less than 5000) where the deployment scenarios show a less steep percentage decrease, or even an increase for some scenarios, such as ROW2. We believe that this range would be the practical range of UDL receivers for satellite unidirectional links because the percentage decrease is less steep and the three best scenarios ROW3, ROW2-3, and ROWCOL2-3 give the values that are more than the 4.27% threshold.

7. Conclusions

This paper investigates several issues regarding the benefits of deploying satellite unidirectional links for global multicast networks where UDL feeds are located in the Internet backbone networks and the receivers are located in regional or stub networks. We define a set of networks to be the backbone networks based on the number of links and the international connectivity. Our scenarios for deployment considers the network and geographical distances between the satellite UDL feeds and receivers. We simulate several deployment scenarios using randomly selected multicast receivers located sparsely in the Internet and measure their performance. Three metrics are used to measure the performance of a deployment scenario: multicast link count, cumulative multicast out-degree, and the expected downstreams per UDL receiver. These metrics views the deployment scenario from different perspectives: the Internet resource (link) usage, the benefits for non-feed networks, and the efficiency for feed networks.

This paper shows that satellite unidirectional links clearly give benefits to the deployment of global multicast networks in terms of reducing the number of routers that need to maintain multicast states. As far as we know, this is the first paper to demonstrate how much multicast state reductions may be expected and who will experience the reductions when satellite unidirectional links are used for multicast network deployments. Our simulation results show that UDL receivers that are three hops away from the feed give good results, but there are tradeoffs in deploying satellite unidirectional links for global multicast networks. Furthermore, geographic-distance-based deployment scenarios do not play a major role in reducing mul-

ticast states due to how Autonomous Systems are distributed geographically on the Internet.

This paper does not address other issues that will affect the deployment such as delay, routing, and traffic models. The long delay and the bandwidth limitations of satellite links are certainly issues for some applications. Routing protocol can become an issue when there are many routers on a satellite unidirectional link. The deployment scenarios should also be tested with simulated multicast traffic to show the dynamics of the link, e.g. multicast joins and leaves. These issues will be addressed in the future work.

References

- 1) Diot, C., Levine, B.N., Lyles, B., Kassem, H. and Balensiefen, D.: Deployment issues for the IP multicast service and architecture, *IEEE Network*, Vol.14, No.1, pp.78–88 (2000).
- 2) Almeroth, K.: The evolution of multicast: From the Mbone to inter-domain multicast to Internet2 deployment, *IEEE Network*, Vol.14, pp.10–20 (2000).
- 3) Rajvaidya, P. and Almeroth, K.C.: Analysis of Routing Characteristics in the Multicast Infrastructure, *IEEE INFOCOM* (2003).
- 4) Duros, E., Dabbous, W., Izumiyama, H., Fujii, N. and Zhang, Y.: A Link-Layer Tunneling Mechanism for Unidirectional Links, RFC3077, IETF (2001).
- 5) Almeroth, K. and Zhang, Y.: Using satellite links as delivery paths in the multicast backbone (Mbone), *Proc. ACM/IEEE WOSBIS'98* (1998).
- 6) Filali, F. and Dabbous, W.: Issues on the IP Multicast Service Behaviour over the Next-Generation of Satellite-Terrestrial Hybrid Networks, *Proc. IEEE ISCC'2001* (2001).
- 7) Filali, F., Dabbous, W. and Kamoun, F.: Efficient planning of satellite-terrestrial hybrid networks for multicast applications, *ICC 2001*, pp.216–223 (2001).
- 8) Faloutsos, M., Faloutsos, P. and Faloutsos, C.: On Power-law Relationships of the Internet Topology, *SIGCOMM*, pp.251–262 (1999).
- 9) Chen, Q., Chang, H., Govindan, R., Jamin, S., Shenker, S. and Willinger, W.: The Origin of Power-Laws in Internet Topologies Revisited, *IEEE INFOCOM* (2002).
- 10) Medina, A., Matta, I. and Byers, J.: On the Origin of Power Laws in Internet Topologies, *ACM Comp. Comm. Review*, Vol.30, No.2, pp.18–28 (2000).
- 11) Dolev, D., Mokryn, O. and Shavitt, Y.: On Multicast Trees: Structure and Size Estima-

- tion, *IEEE INFOCOM* (2003).
- 12) Chuang, J.C.-I. and Sirbu, M.A.: Pricing Multicast Communication: A Cost-Based Approach, *Telecomm. Systems*, Vol.17, No.3, pp.281–297 (2001).
 - 13) Adjih, C., Georgiadis, L., Jacquet, P. and Szpankowski, W.: Is the internet fractal?, *SODA*, pp.338–345 (2002).
 - 14) Chalmers, R.C. and Almeroth, K.C.: Modeling the Branching Characteristics and Efficiency Gains of Global Multicast Trees, *IEEE INFOCOM*, Vol.1, pp.449–458 (2001).
 - 15) Chalmers, R.C. and Almeroth, K.C.: On the topology of multicast trees, *IEEE/ACM Trans. Netw.*, Vol.11, No.1, pp.153–165 (2003).
 - 16) Chalmers, R.C. and Almeroth, K.C.: Developing a Multicast Metric, *Proc. IEEE GLOBECOM'00*, San Francisco, USA (2000).
 - 17) Govindan, R. and Reddy, A.: An Analysis of Internet Inter-Domain Topology and Route Stability, *IEEE INFOCOM '97*, p.850 (1997).
 - 18) Subramanian, L., Agarwal, S., Rexford, J. and Katz, R.H.: Characterizing the Internet Hierarchy from Multiple Vantage Points, *IEEE INFOCOM*, New York (2002).
 - 19) Harrenstien, K., Stahl, M. and Feinler, E.: NICKNAME/WHOIS, RFC 954, IETF (1993).
 - 20) Gao, L.: On inferring autonomous system relationships in the Internet, *IEEE/ACM Trans. Netw.*, Vol.9, No.6, pp.733–745 (2001).

(Received December 24, 2004)

(Accepted September 2, 2005)

(Online version of this article can be found in the IPSJ Digital Courier, Vol.1, pp.604–615.)



Achmad Husni Thamrin is a research associate at the Graduate School of Media and Governance, Keio University. He received an M.A. in Media and Governance from Keio University in 2002. His research topics are satellite Internet and multicast.



Hidetaka Izumiyama received a B.E. of information science from Tsukuba University in 1988, and a Ph.D. from Keio University in 2003.

Hiroyuki Kusumoto was graduated from the Master's Program of Graduate School of Science, Osaka University in 1985. He received his Ph.D. from Keio University in 1995. He is now an Associate Professor at the Faculty of Environmental Information, Keio University.



Jun Murai is a Professor at the Faculty of Environmental Information, and a Vice-President of Keio University. He received his Ph.D. in computer science from Keio University in 1987. Adjunct Professor at the Institute of Advanced Studies, United Nations University; Instructor at Tokyo University of Art and Music; General Chairperson of the WIDE Project; Vice Chairperson of the Japanese chapter of the Internet Society; and Vice President of the Japanese Internet Association. He is a member of the board of the Internet Society. His research has centered on electronic observation, satellite Internet, multimedia Internet, and mobile and ubiquitous computing.