Evaluation of Low-Distortion Approximation Methods For Overlay Multicast Tree construction

Huan Luo and Khaled Harfoush
Department of Computer Science
North Carolina State University
Raleigh, NC 27606
{hluo,harfoush}@cs.ncsu.edu

Abstract— **In this paper we investigate the performance of distortion-bounded approximation methods for overlay multicast tree construction and compare them to common overlay multicast tree construction techniques such as the minimum spanning tree (MST) and core-based trees. We study the fanout-distortion tradeoff of these techniques on both synthetic and real Internet topologies. The performance results vary not only between different methods but also between different topologies. Our study of the fanout-distortion tradeoff reveals that low distortion approximation methods sacrifice fanout for low distortion suggesting a more careful design in order to balance both mertics. Furthermore, MST outperforms low distortion approximation methods both in distortion and fanout, especially on realistic Internet topologies, even though it was not designed to optimize distortion or fanout.**

I. Introduction

Interactive multi-party communication among a small group is becoming very popular. Examples include video conferencing, remote education and distributed gaming. A large number of such groups may coexist in the Internet and it is desirable to optimize the perceived performance of the group members together with the stress on the Internet resources. While IP multicast was originally proposed for this purpose, it is not supported in today's Internet mainly due to its heavy reliance on router support. Application-level (overlay) multicast is a viable alternative [6], [18], which only requires intelligence at the participating nodes. Typically, overlay multicast requires an overlay tree connecting group members who communicate over this tree. Two performance metrics are important from a user's perspective: (1) the *delay* between communicating nodes or the *delay distortion* defined as the ratio between overlay delay to direct communication delay, and (2) the fanout at these nodes, defined as the out-degree of nodes in the overlay. Minimal delay is critical for realtime applications and a small fanout is needed since nodes are responsible for data copying and forwarding to other nodes and the larger the fanout the more the stress on nodes' resources. Efficient support for multiparty real-time applications at the application level requires an overlay tree, which optimizes/balances the distance and fanout metrics. Furthermore, from a network's perspective, it is desirable to reduce the *stress* on network links.

Constructing an optimal fanout and distance-bounded tree is an NP-complete problem [23], [13] and while many existing overlay multicast tree construction algorithms limit node fanout, none of them target bounded distortion. On the other hand, extensive research has been done on probabilistic metric space approximations, embedding arbitrary graphs into trees with bounded distortion [8], [9], [10], [20], [5], [2], [1], [13], [19], [3], [23]. In this paper we investigate the performance of representative distortion-bounded approximation methods [10], [20] for overlay multicast tree construction and compare them to common overlay multicast tree construction techniques such as the minimum spanning tree (MST) and core-based trees. Both [10] and [20] can achieve the tight distortion bound of $O(\log n)$ [5]. To our best knowledge, this is the first work investigating the application of distortion-bounded algorithms to application level multicast tree construction. Apart from the bounded distortion, these approximation techniques rely on hierarchical partitioning of nodes using a well-defined partition radius, which is convenient for overlay tree construction and suitable for small group communication.

We study the fanout-distortion tradeoff of these techniques on both synthetic and real Internet topologies and show how the performance results vary not only between different methods but also between different topologies. Our study of the fanout-distortion tradeoff reveals that low distortion approximation methods sacrifice fanout for low distortion suggesting a more careful design in order to balance both mertics. Surprisingly, the MST outperforms low distortion approximation methods both in distortion and fanout, especially on realistic Internet topologies, even though it was not designed to optimize distortion or fanout.

The rest of the paper is organized as follows: In Section II we introduce the overlay construction techniques to be investigated. In Section III We analyze the worst-case fanout bound of the approximation methods; In Section IV we introduce the network topologies that we use in our experiments. In Section V we define the performance metrics of interest. In Section VI we provide and comment on the experimental results. We finally conclude in Section VII.

II. Overlay Construction Techniques

We investigate four different techniques: (1) Bartal, (2) Berkeley, (3) MST, and (4) Star. They target different objectives and rely on different constructions. Both Bartal's algorithm [10] and Berkeley's algorithm [20] target an overlay with bounded distortion of $O(\ln n)$. They both rely on clustering nodes into *balls* with some radius, r, but differ in the way r is selected. The value of r is drawn from an *exponential* probability distribution function in Bartal's algorithm and drawn from an *uniform* distribution in Berkeley's algorithm. The Minimum Spanning Tree algorithm, MST, targets an overlay with the minimum sum of the delays between connected overlay nodes. The Star overlay construction is a *core-based* algorithm, which selects one node as the *core* and communication between all nodes is carried through the core.

The input to each of these algorithms is a set of nodes, V, with known pairwise distances (delays). Let Δ be the diameter of the graph; i.e., the largest distance between any pair of nodes. The output of each algorithm is an overlay graph connecting the nodes. Let $B(x, r)$ be the ball including all nodes within radius of r of node $x \in V$ and $|B(x, r)|$ be the size of the ball.

A. Bartal's Algorithm

By combining ideas from [8] and [9], Bartal proposed a graph decomposition technique in [10], with a proven distortion bound of $O(\ln n)$, where n is the number of nodes in the system. The key idea is that each cluster/ball of nodes is recursively partitioned into two non-overlapping clusters/balls. A ball's radius, r, is selected from an exponential probability distribution function, $p(r) = \frac{\lambda^2}{(1-\lambda^{-2})}\frac{ln\lambda}{R}\lambda^{-r/R}$, where $R = \Delta/128$, $\lambda \equiv |B(v, 16R)|/|B(v, R)|$, and v is the center of the ball, which is selected to minimize $|B(v, 16R)|/|B(v, R)|$. Algorithm 1 provides the details of Bratal's construction.

algorithm 1 Bartal(V)

$R = \Delta/128$
Find $v \in V$ that minimizes $|B(v, 16R)|/|B(v, R)|$
choose a radius $r \in [2R, 4R]$ according to exponential $p(r)$
V_1=set of all nodes $\in B(v, r)$
V_2=set of all nodes $\notin B(v, r)$
If ($V_1 \neq \emptyset$) then Bartal(V_1)
If ($V_2 \neq \emptyset$) then Bartal(V_2)

B. Berkeley's Algorithm

Berkeley's Algorithm was proposed in [20] and again relies on graph decomposition into balls but as opposed to Bartal's, it tries to get a *balanced* decomposition of the graph as a ball's radius, r, is selected from a *uniform* distribution function in the range $[\Delta/4, \Delta/2]$. The algorithm works as follows: A random permutation of nodes, π, is fixed and used throughout the process, and also β is chosen uniformly at random in $[1, 2]$. For each i, we compute the set of node clusters, D_i and D_{i+1} as follows. First, set r_i to $2^{i-1} * \beta$. Let S be a cluster in D_{i+1}. A node $u \in S$ is assigned to the first node $v \in V$ closer than r_i to u, according to π. Each child cluster of S in D_i then consists of the set of vertices in S assigned to a single center v. Note that the center v itself need not be in S. Thus one center v may correspond to more than one cluster, each inside a different level $i+1$ cluster. Also note that the radius of each cluster, r_i, is at most 2^i. Algorithm 2 provides the details of Berkeley's construction.

algorithm 2 Berkeley()

```
choose a random permutation π of nodes v1, v2, ···, vn
choose β uniformly at random in [1, 2]
δ = log₂Δ
D_δ = {V}
i = δ − 1
while D_{i+1} has non-singleton clusters do
    r_i = 2^{i−1} ∗ β
    foreach l = 1, 2, ···, n do
        foreach cluster S in D_{i+1} do
            Create a new cluster consisting of all unassigned
            vertices in S closer than r_i to π_s(l)
        end
    end
    i = i − 1
end
return G ;
```

C. Degree-Bounded MST Algorithm

The Minimum Spanning Tree (MST) algorithm targets the minimum overlay *cost* (sum of the delays between connected overlay nodes), which makes it a good candidate for application layer multicast. However, it has been shown in [26] that the MST algorithm results in linear $O(n)$ distortion. Despite its linear distortion bound, the MST algorithm has been found to incur low distortion in practice [4]. In this paper, we use Kruskal's algorithm to construct the MST. Algorithm 3 provides the details of the MST construction.

algorithm 3 MST()

```
order all links nondescendingly by the distance
L = 0        L is the number of included links in the overlay
while L < (n − 1) do
    foreach link (u, v) in the ordered list do
        if (adding (u, v) does not lead to a cycle) then
            Add link (u, v) to overlay
            L = L + 1
        end
    end
end
```

D. The Star Algorithm

The Star overlay construction is a *core-based* algorithm, which selects one node as the *core* and communication be-

TABLE I

MAXIMUM FANOUT FOR BARTAL'S ALGORITHM.

Distributions	p	No. levels (h)	Max fanout (p*h)
sub-linear	2	$0.5log_{64}(n), 0.5log_{32}(n)$	$log_{32}(n)$
linear	2	$[log_{64}(n), log_{32}(n)]$	$2log_{32}(n)$
polynomial	2	$[2log_{64}(n), 2log_{32}(n)]$	$4log_{32}(n)$
exponential	2	$[nlog_{64}(2), nlog_{32}(2)]$	$2nlog_{32}(2)$

tween all nodes is carried through the core. We use the node which has the minimum-maximum distance to all other nodes as the core. This algorithm is used to study the behavior of the case in which one node is highly capable (has excellent Internet connectivity) and acts as an intermediary between the source of communication and all destination nodes. Intuitively the Star overlay construction should result in a better distortion as communication will only hop through one intermediary node (the core). As we show in Section VI, this is not necessarily the case.

III. FANOUT ANALYSIS

Clearly, there is a tradeoff between distortion and node fanout. The higher the node fanout the more likely it is that an overlay construction would achieve a low distortion. The distance between nodes in the *underlay* impacts the fanout and distortion in the resulting overlay. In this section we analyze the expected worst-case fanout under certain underlay *distance distributions*. We choose the following popular distance distribution functions for our analysis: (1) sub-linear ($\sqrt{n}$), (2) linear (n), (3) polynomial (n^2), and (4) exponential (2^n). We do not claim that any of these distance distributions are representative of the distance between Internet nodes. Still, two factors have lead us to the choice of these distributions: (1) the diversity in their properties, and (2) the lack of a representative distance distribution for Internet nodes. Our analysis is intended to clarify and back the simulation results, which we present in Section VI.

A. Analysis of Bartal's algorithm

Bartal's algorithm uses an exponential radius function to cluster nodes and explicitly decomposes a graph into two subgraphs, so each partitioning level always leads to a fanout, p, of 2, and the final fanout depends on the number of partitions, h. In the worst case, a single node is always chosen as center nodes for all partitioning levels leading to a worst case (maximum) fanout of $p.h = 2h$. Notice that using the exponential radius distribution function, there is a high probability for choosing a "small" radius, which produces many small sub partitions. For example, for a linear distance distribution and a radius range of $[\Delta/64, \Delta/32]$, the sub partition size would be $[n/64, n/32]$, which will set $h = [log_{64}(n), log_{32}(n)]$. So the worst case fanout would be in the range $[2 * log_{64}(n), 2 * log_{32}(n)]$, when a single node is chosen as the center at all partitioning levels. The maximum fanout for the considered distance distributions when using Bartal's algorithm are outlined in Table I. As is clear from the table, Bartal's resulting overlay has a logarithmic fanout bound independent of the underlay node distances. The resulting distance distortion will heavily depend on the underlay distances though.

TABLE II

MAXIMUM FANOUT FOR BERKELEY'S ALGORITHM.

Distributions	p	No. levels ($h = log_p n$)	Final fanout(p*h)
sub-linear	[4,16]	$[log_{16}n, log_4 n]$	$4log_4 n$
linear	[2,4]	$[log_4 n, log_2 n]$	$2log_2 n$
polynomial	$[\sqrt{2}, 2]$	$[log_2 n, log_{1.4} n]$	$1.4log_{1.4} n$
exponential	$[n/2, n]$	[1,2]	n

B. Analysis of Berkeley's Algorithm

Berkeley's algorithm uses an uniform radius function in the range of $[\Delta/4, \Delta/2]$ to cluster nodes resulting in a balanced hierarchy and a low fanout. For example, for a linear distance distribution and a radius range of $[\Delta/4, \Delta/2]$, the sub-partition size would be $[n/4, n/2]$, which sets the number of sub-partitions $p = [2, 4]$, and $h = log_p n$. The maximum fanout in this case would be the maximum value for $p.h$; i.e. $4log_4 n$. The maximum fanout for the considered distance distributions when using Berkeley's algorithm are outlined in Table II. From the above analysis, we conclude that the fanout can be (asymptotically) well below n under various distance distributions, except for the exponential distribution case.

IV. NETWORK TOPOLOGIES

The structure of the Internet *underlay*, the number of nodes in the system, and the distance between nodes together with the overlay construction technique all affect the overlay performance. We use five different network topologies to test the performance of the overlay construction techniques discussed in Section II.

Two of these topologies are synthetic and are generated using (1) Brite [27], [28] and (2) GT-ITM [14] topology generators. Brite is *degree-based* and generates a graph with *power law* degree distribution, while GT-ITM is *structure-based* and generates graphs emulating the Internet structure with transit and stub autonomous systems (AS). The difference between the generated Brite and GT-ITM topologies reflects the disagreement in the networking research community about the exact structure of the Internet and its features. The generated BRITE topology has 1000 routers, while GTITM has 600 routers. End hosts in our experiments are attached to randomly chosen routers and we vary the number of end-hosts to study the scalability of the overlay construction techniques.

The other three network topologies were obtained through actual network measurements to obtain the distance between a set of Internet nodes through *Ping* packets: (1) The PlanetLab topology [30], (2) the NLANR topology [29], and (3) the Meridian topology [33]. The PlanetLab topology connects 364 nodes on the PlanetLab infrastructure [30], the NLANR topology connects 117 nodes and the Meridian data set connects 2500 nodes. We randomly select different size subsets of these nodes to study the performance of the overlay construction techniques for different group sizes. Each data point in our graphs (Section VI) is a result of 25 runs leading to a 95% confidence in the results.

V. Performance Metrics

We compare the performance of the overlay construction techniques using the following metrics:

1) *Average Distortion Ratio*, or *distortion* for short, defined as the ratio of the overlay delay to the direct connection delay, averaged on all pairs of nodes
2) *Maximum Fanout*, or *fanout* for short, defined as the maximum degree of the overlay nodes. The fanout reflects the heaviest load on any node in the overlay.
3) *Stress*, defined as the maximum load (copies of exchanged content) on any physical links (underlay link).

In the presented results in Section VI, stress and fanout are measured when the center node of the whole graph is the sender. Also, we do not study the stress for the Internet topologies, PlanetLab, NLANR, and Meridian due to the difficulty in attaining the stress for such topologies.

VI. Simulation Results

In Table III we compare the distortion/fanout tradeoffs for the different overlay construction techniques using the different network topologies when the number of nodes in the system is 100. The results in Table III reveal that: (1) Topologies affect the performance. however, this trend is observed for some techniques more than the others and for some metrics more than the others. For example, the distortion resulting from the Star topology is much higher in the PlanetLab topology than for all other topologies, while the distortion of the Berkeley's algorithm does not vary dramatically for different topologies. (2) Statistical approximation techniques such as Berkeley and Bartal indeed bound the overlay distortion. However, this comes at the cost of relatively high fanout. (3) MST, which is not directly intended to optimize distortion, does lead to better distortion than Berkeley and Bartal, which are designed to bound distortion, especially with realistic Internet topologies (PlanetLab, NLANR, and Meridian). Furthermore, MST, which is also not intended to bound fanout, does lead to better fanout than all other techniques. (4) The Star construction, which seemingly is expected to produce very low distortion at the cost of using a node with very high degree, does *not* produce low distortion especially in realistic PlanetLab, NLANR, and Meridian topologies. The degradation in performance of the Star topology can be attributed to the high variation in node delays in the Internet. In general, synthetic topologies have low variation in distance , while real data sets distance is arbitrary, asymmetric, not complying the triangle inequality and of high variation. For example, the PlanetLab distances range from 0.1 ms up to 1500 ms, while the node distances range between 7 ms and 400 ms for the Brite topology and between 1 ms and 200 ms for the GT-ITM topology. (5) In our experiments, we have found that when the maximum fanout is large, then typically this is due to the presence of a a very small number of nodes with this high fanout but most nodes have typically a low degree. This suggests that if some nodes are able to *take* high fanout due to their superior Internet connectivity, then they can act as supernodes and relieve the group from high distortion.

We next study the scalability of investigated algorithms as the number of nodes in the system (group size) increases. Figures 1 and 2 compare the distortion, fanout, and stress of the different algorithms on the Brite and GT-ITM topologies; and Figure 3 compares the distortion and fanout of the algorithms on the Planetlab, NLANR, and Meridian topologies. Figures 1, 2 and 3 reveal the following: (1) Berkeley's distortion scales better as the group size increases in almost all considered topologies. This however comes at the cost of a slightly higher fanout. (2) The Star algorithm is very sensitive to different topologies. While it leads to the best distortion in synthetic topologies, it leads to high distortion in realistic Internet topologies as in Planelab and NLANR. It is also clear that Star algorithm leads to the worst stress among all other algorithms. (3) The MST algorithm is still performing better than the others for all topologies and all metrics: distortion, fanout, and stress. The performance of Berkeley as the group size increases does not deviate by a lot from the MST algorithm. (4) Again, in our experiments, we have found that when the stress is large, then typically this is due to the presence of a a very small number of links with this high stress but most links have typically a low stress.

Overall, the results suggest that the MST outperforms all existing algorithms including distortion bounded approximation techniques and the Star algorithm. The close similarity between MST and Berkeley performance suggests that Berkeley maybe suitable for application-level multicast due to its simple deployment. The reasons for the MST superiority results from its ability to dynamically adapt its structure based on nodes' distances allowing pairwise communication with minimum overall cost.

To show whether our fanout analysis can match with the simulation, both results are shown in Figure 4, including fanout values for the 3 algorithms on the 5 topologies, where "(B)","(G)","(P)", "(N)", "(M)" represent the Brite, GT-ITM, PlanetLab, NLANR and Meridian topologies, respectively, with 100 nodes in each topology. While there is a difference between the absolute values of simulation results and our analysis, one can see that the analysis is able to capture the relative performance differences between the different algorithms on the different topologies.

VII. Conclusions and Future Work

In this paper we compare the performance of low distortion hierarchical partitioning multicast tree construction techniques to other tree construction techniques such as MST and Star. We investigate the correlation between fanout and distortion and study the scalability of these techniques as the system size increases. To our best knowledge, this is the first work on studying the performance of approximation methods with bounded distortion for application- layer multicast. Our results

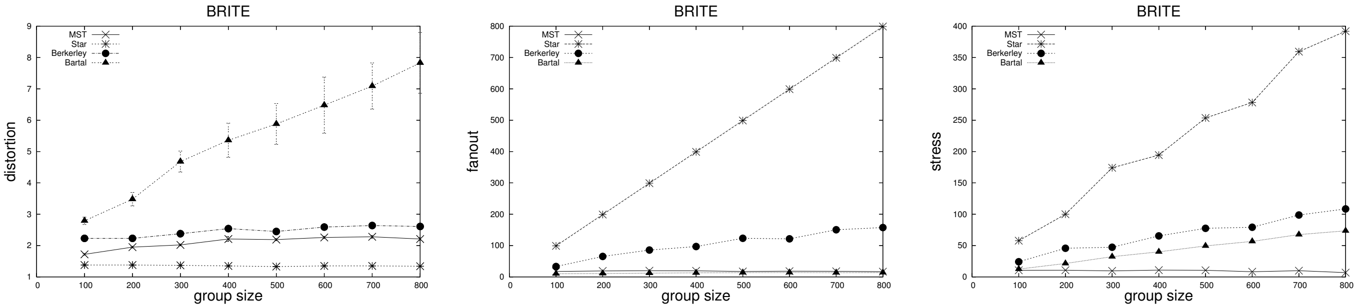

Fig. 1. Performance of overlay construction techniques as the group size increases on the Brite Topology.

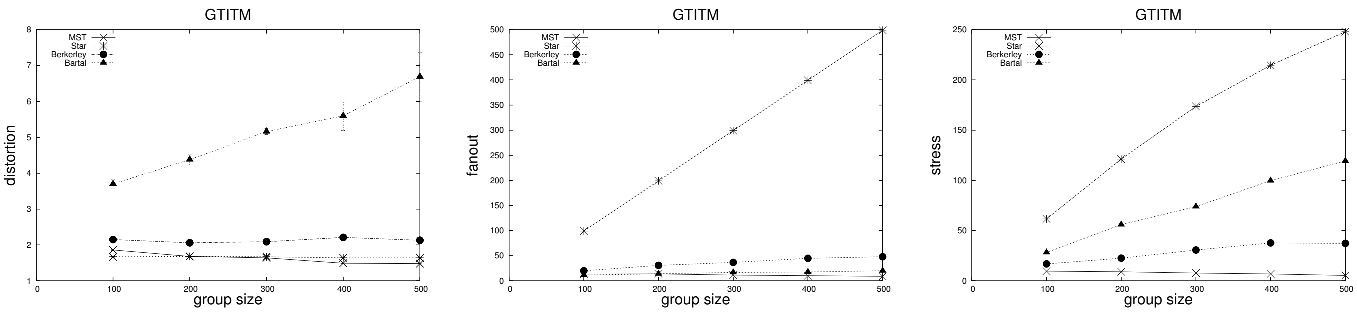

Fig. 2. Performance of overlay construction techniques as the group size increases on GT-ITM Topology.

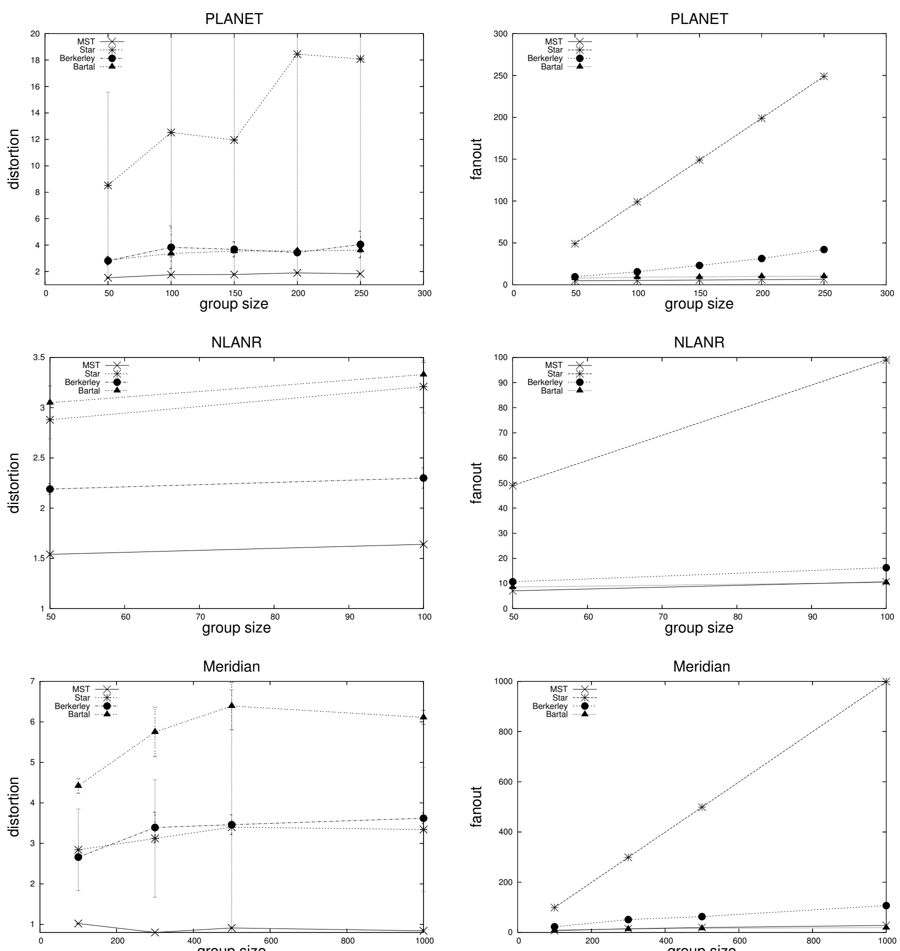

Fig. 3. Performance of overlay construction techniques as the group size increases on the PlanetLab, NLANR, and Meridian topologies.

TABLE III

FANOUT-DISTORTION TRADEOFFS IN TOPOLOGIES WITH 100 NODES.

	Distortion				Fanout			
	MST	STAR	Berkeley	Bartal	MST	STAR	Berkeley	Bartal
Brite	1.61	1.39	2.10	2.84	16.48	99.00	34.40	10.24
GTITM	1.83	1.77	2.12	3.55	12.32	99.00	21.88	11.60
PLANET	1.68	11.61	3.22	3.37	5.20	99.00	15.12	8.88
NLANR	1.65	3.18	2.18	3.11	10.80	99.00	17.08	10.56
MERIDIAN	0.98	2.91	2.89	3.94	7.60	99.00	20.08	10.64

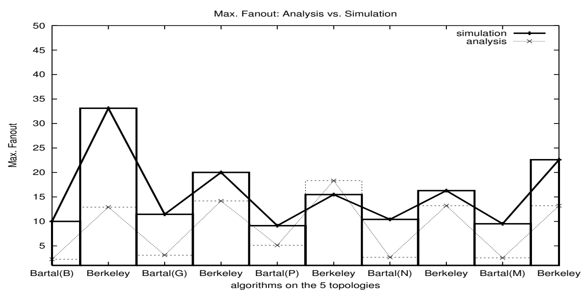

Fig. 4. Maximum fanout: Analysis vs. simulation results.

provide interesting insights on the difference between synthetic and realistic Internet topologies and the difference in performance of the different techniques on these topologies. Our results also suggest that it in order to balance distortion and fanout, it is important to partition nodes evenly in order to get a low fanout, and to wisely select partition centers in order to get low distortion. Further detailed investigation of the tradeoffs between distortion and fanout on different topologies, in a distributed and dynamic setup is part of our future research agenda.

REFERENCES

[1] I. Abraham, Y. Bartal, O. Neiman,*Local embeddings of metric spaces*, Symposium on Theory of Computing(STOC), San Diego,CA,June 2007.

[2] I. Abraham, Y. Bartal, O. Neiman, *Embedding metrics into ultrametrics and graphs into spanning trees with constant average distortion*,Symposium on Discrete Algorithms(SODA),New Orleans, Louisiana,January 2007.

[3] N. Alon,R. M. Karp,D. Peleg,D. West,*A Graph-Theoretic Game and its Application to the k-Server Problem*,SIAM Journal on Computing archive Volume 24 , Issue 1,Feb 1995.

[4] B. Awerbuch, Y. Shavitt.,*Topology Aggregation for Directed Graphs*, IEEE/ACM Transactions on Networking, 9(1), Feb 2000.

[5] M. Badoiu,P. Indyk, A. Sidiropoulos, *Approximation Algorithms for Embedding General Metrics Into Trees*, 18th Symposium on Discrete Algorithms(SODA), New Orleans, Louisiana,January 2007.

[6] S. Banerjee, B. Bhattacharjee, *A comparative study of application layer multicast protocols* , University of Maryland, College Park.

[7] S. Banerjee, B. Bhattacharjee, *Analysis of the NICE Application Layer Multicast Protocol* , UMIACS TR-2002-60 and CS-TR 4380, University of Maryland, College Park, June 2002.

[8] Y. Bartal, *Probabilistic approximation of metric spaces and its algorithmic applications* , proceedings of the 37th IEEE Symposium on Foundations of Computer Science, pp. 184-93 (1996).

[9] Y. Bartal, *On Approximating Arbitrary Metrics by tree Metrics* , Proceedings of the 30th Annual ACM Symposium on Theory of Computing(STOC),Dallas, Texas,May 1998.

[10] Y. Bartal, *Graph Decomposition Lemmas and their Role in Metric Embedding Methods* , European Symp. on Algorithms, Bergen, Norwaym, September 2004.

[11] Y. Bartal, M. Mendel, *Randomized k-Server Algorithms for growth-rate bounded graphs* , In Proc. of the 15th ACM-SIAM Symp. on Discrete Algorithms(SODA), New Orleans, Louisiana,January 2004.

[12] E. Brosh, Y. Shavitt, *Approximation and Heuristics Algorithms for Minimum Delay Application -Layer Multicast Trees* , IEEE INFOCOM, Hong Kong, China, March 2004.

[13] L. Cai,D. G. Corneil,*Tree Spanners*,SIAM Journal on Discrete Mathematics archive Volume 8 , Issue 3,August 1995.

[14] K. Calvert, M. Doar and E. W. Zegura. *Modeling Internet Topology*, IEEE Communications Magazine, June 1997.

[15] M. Charikar, C. Chekuri, A. Goel, A. Guba, S. A. Plotkin, *Approximating a finite metric by a small number of tree metrics* , In IEEE Symposium on Foundations of Computer Science, Page 379-388, 1998.

[16] M. Charikar, C. Chekuri, A. Goel, S. Guha, *Rounding via Trees: Deterministic Approximation Algorithms for Group Steiner Trees and k-median* , Proceedings of the 30th Annual ACM Symposium On Theory of Computing,June 2000.

[17] Y. Chu, S. Rao, and H. Zhang. *A Case for End System Multicast* , Internet End-to-End Research Meeting, ACM SIGMETRICS, Santa Clara, CA,June 2000.

[18] A. EI-Sayed, V. Roca, *A survey of proposals for an alternative group communication service* , IEEE Network, 2003, 17(1):46-51.

[19] M. Elkin, Y. Emek, D. A. Spielman, S. Teng,*Lower-Stretch Spanning Trees*, Symposium on Theory of Computing(STOC), Baltimore, MD, May 2005.

[20] J. Fakcharoenphol, S. Rao, K. Talwar, *A Tight Bound On Approximating Arbitrary Metrics by Tree Metrics* , In Proc. of the 14th ACM-SIAM Symp. on Discrete Algorithms, 257-265, 2003.

[21] S. Hazarika, D. Towsley, *Delay Analysis of Application Level Multicast On Content Addressable Networks* , Proceedings of IEEE Globecom, Dallas, Texas, December 2004.

[22] R. Jothi, B. Raghavachari, *Degree-Bounded Minimum Spanning Trees*, In Proc. 16th Canadian Conf. on Computational Geometry (CCCG), 2004.

[23] G. Kortsarz, D. Peleg. *Generating Low-Degree 2-Spanners* , SIAM Journal on Computing Volume 27, Number 5 pp. 1438-1456, 1998.

[24] Z. Li, P. Mohapatra, *The impact of topology on overlay routing service* , INFOCOM, HongKong, China, March 2004.

[25] N. Linial, E. London, and Y. Rabinovich, *The geometry of graphs and some of its algorithmic applications* , Combinatoric, 15(2):215-245,1995

[26] J. Matousek, *Bi-Lipschitz embeddings into low-dimensional Euclidean spaces*, Comm Math Univ. Carolina, 1990.

[27] A. Medina, A. Lakhina, I. Matta, and J. Byers, *BRITE: An Approach to Universal Topology Generation* , In Proceedings of the International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems- MASCOTS, Cincinnati, Ohio,August 2001.

[28] A. Medina, I. Matta, and J. Byers, *On the Origin of Power Laws in Internet Topologies* , ACM Computer Communications Review, April 2000.

[29] *The NLANR project, http://amp.nlanr.net/AMP/*

[30] *http://planetlab.org*

[31] V. Paxson, *End-to-end Internet Packet Dynamics*, In ACM SIGCOMM, Cannes, French Riviera, France,September 1997.

[32] D. Pendarakis, D. Verma and M. Waldvogel, *ALMI: an Application Level Multicast Infrastructure*, Proceedings of 3rd Usenix Symposium on Internet Technologies & Systems (USITS),San Francisco, California ,March 2001.

[33] B. Wong, A. Slivkins and E. G. Sirer, *Meridian: A Lightweight Network Location Service without Virtual Coordinates*, in ACM SIGCOMM, Philadelphia, PA, August 2005