PREDICTING THE PERFORMANCE IMPACTS OF MULTICAST NETWORK WORKLOADS

Tim Harrison Carey L. Williamson
Department of Computer Science
University of Saskatchewan
Saskatoon, SK, Canada S7N 5A9
e-mail: {harrison,carey}@cs.usask.ca

Keywords: Computer Systems and Networks, Telecommunications, Multicast Routing, Simulation

ABSTRACT

This paper addresses the problem of multicast (*i.e.,* multipoint) routing in computer networks carrying a mix of unicast and multicast traffic. In particular, we derive a closed form mathematical expression to predict the average link usage for multicast calls in a two-dimensional rectangular mesh network. The analytical expression is used to evaluate the relative performance of three simple multicast routing techniques: M-Unicast, Conference Call, and Core-Based Trees. Simulation experiments are then used to assess the accuracy of the analytical model.

The results show that, among the three routing models considered, the Core-Based Trees model has the lowest average links required per multicast recipient. Simulation results show that the analytic model is highly accurate: always within 10%, and often within 5% of the results obtained by simulation. The accuracy of the analytical model degrades only slightly as the network load increases.

1 INTRODUCTION

Multicast is the selective broadcast of a message from a source node S to a group G of M recipients, where the group G can consist of an arbitrary subset of the total nodes in a network. Multicast is distinct from *unicast*, which provides point-to-point (*i.e., one-to-one*) communication, and from *broadcast*, which provides *one-to-everyone* communication.

The goal in multicast transmission is for a sender to only have to send *one* copy of each message, regardless of how many group members there are. In other words, rather than making multiple copies of each message and sending one copy, in turn, to each recipient, the sender relies on the network itself to make the copies. These copies are made as needed at branch points in the *multicast delivery tree* so that one copy is delivered to each recipient. A related goal is that messages should *only go where they are needed* in the network. Again, the sender relies on the network to achieve this: the network should route a message to each member of the group, while trying to avoid routing any messages into portions of the network where no intended recipients reside.

The way in which multicast routing is implemented in a network has a major impact on the efficiency of multicast routing, and on the impact of multicast traffic on other (*i.e.,* unicast) traffic in the network. There are several existing and proposed techniques for routing multicast network traffic (Ballardie *et al.* 1993; Deering 1991; Kompella *et al.* 1993; Moy 1994; Verma and Gopal 1993; Waitzman *et al.* 1988). Three such techniques are M-Unicast, Conference Call, and Core-Based Trees (Ballardie *et al.* 1993). These three are the approaches evaluated in this paper.

The purpose of this paper is to present a performance study of the three multicast routing algorithms mentioned above, using simulation and analysis, in an attempt to assess the performance impact of multicast network traffic workloads. The performance analysis is carried out on simple mesh-based network topologies. While not necessarily representative of realistic network scenarios, the mesh-based networks do provide two key advantages. First, they provide a network scenario in which optimal routing can be computed, providing a baseline for evaluating the performance of other routing heuristics. Second, they provide a simple enough environment to interpret and fully understand the results of the performance analysis.

The key benefit of the model that we present is its ability to predict the performance characteristics (*e.g.,* network link usage) when multicast traffic is added to a primarily unicast network workload. The results predicted by the model are useful in evaluating the scalability of multicast routing algorithms: scalability with respect to network size, and scalability with respect to multicast group size. For example, while routing cost grows linearly with group size in all algorithms considered, the *per recipient* routing cost and the call blocking behaviour can be significantly different.

The remainder of this paper is organized as follows. Section 2 describes the multicast routing algorithms evaluated in the paper. Section 3 presents a brief mathematical analysis of the multicast routing problem, including an analytical model for predicting link usage for multicast traffic. Section 4 presents a simulation experiment designed to evaluate the multicast routing algorithms, and the accuracy of the analytical model. Section 5 concludes the paper.

2 MULTICAST ROUTING MODELS

2.1 Problem Statement and Assumptions

The purpose of a multicast routing algorithm is to find a routing path from a designated source node S in a network to a designated set of recipient nodes in the network. The performance objectives in multicast routing are to route the call as efficiently as possible (*i.e.,* using as few network links as possible), and to keep the call blocking probability as low as possible.

In our work, we assume that the set of multicast recipients is known at the time of call arrival, and remains static for the duration of the call. If all intended recipients can be reached successfully without exceeding the network capacity on any network link, then the call is accepted. Otherwise, the call is blocked (rejected), recorded as such, and removed from further consideration by the network. Note that accepting a call implies that the call must be routed through the network, which reduces the available network capacity on the chosen route for the duration of the call. The available capacity on the network links correspondingly increases whenever a call completes.

We focus our attention only on *greedy* approaches to multicast routing: route the arriving call if a path exists to accommodate the call right now. Otherwise, reject the call. This paper considers three such approaches to multicast routing: M-Unicast, Conference Call, and Core-Based Trees.

2.2 M-Unicast Model

The M-Unicast routing algorithm achieves multicast transmission by using M separate unicast transmissions, one to each recipient. That is, the source makes M copies of each message, and sends one copy to each group recipient using unicast. An example of M-Unicast is electronic mail sent to a mailing list in the Internet: a separate copy is routed to each recipient on the distribution list, regardless of where in the network they reside.

M-Unicast is thus *not* a multicast routing algorithm at all, since it violates the goals of multicast that were outlined in Section 1. Nevertheless, it does provide a useful baseline for comparison: if a proposed multicast routing algorithm A provides performance results that are *worse* than M-Unicast, then clearly A is not a practical multicast routing algorithm.

2.3 Conference Call Model

The conference call routing model uses a specialized conference hub, statically located in the network, as a rendezvous point for the sender and the recipients. This model is used in the telephone system for (voice) conference calls. Participants must dial in to a pre-assigned number on a pre-assigned conference hub (bridge). For economic reasons, conference bridges are usually available only in selected metropolitan areas. However, this fact means that sub-optimal routing can easily result when receivers are clustered elsewhere in the network.

In our version of the conference call model, the source of a multicast call sends *one* copy of a message to the conference hub. The hub node is responsible for replicating the message, and delivering one copy separately to each of the M recipients (using M-Unicast). In our simulator, the number of hubs can be scaled with network size (*e.g.,* 1 hub node in a small 4 x 4 mesh, and 4 equally spaced hub nodes in a large 10 x 10 mesh). In the experiments in this paper, only 1 hub is used, regardless of the network size. The location of the hub node is statically assigned, and only interior nodes of the mesh (never corner nodes or edge nodes) are eligible to be conference hubs.

2.4 Core-Based Trees Model

The Core-Based Trees (CBT) (Ballardie *et al.* 1993) model is similar to the conference call model in that it uses a hub (core) to effect the delivery of multicast messages. However, rather than using a static hub, the CBT model chooses the location of the hub *dynamically* for *each* multicast call, based on where the group members reside. As in the conference call model, the source sends one copy of each message to the hub (core), and the hub is responsible for replicating the message and sending a copy separately to each group member (using M-Unicast).

In our implementation, the location of the hub is determined by averaging the X and Y coordinates of all group participants (source and recipients) and rounding off to the nearest node in the mesh. With this approach, corner nodes in the mesh can never be chosen as hubs. Edge nodes can be chosen as a hub only if a large number of group participants lie along the same edge of the mesh. Thus the vast majority of hub nodes chosen will be interior nodes of the mesh.

2.5 Summary

This section has presented three simple routing techniques for handling multicast calls in mesh-based networks. Note that there is no "link sharing" in any of the routing models[1] described here. That is, we do *not* fully exploit the true potential of multicast as described in Section 1. With link sharing, multicast calls that use common links in routing a multicast group use only one unit of channel capacity for all recipients on those common links, regardless of how many recipients use that link. This "no link sharing" assumption is made to simplify the mathematical analysis presented in Section 3. A full set of performance results for multicast routing algorithms, including a study of the effects of link sharing, is available in (Harrison 1995).

[1] The original CBT model (Ballardie *et al.* 1993) certainly uses link sharing, but we intentionally exclude that feature from our model so that we can make a fair comparison with the other two models.

3 MATHEMATICAL ANALYSIS

This section presents a mathematical analysis of the multicast routing problem. In particular, we derive a closed form expression for the average (*i.e.,* expected value) for the distance between two randomly chosen distinct points in a regular mesh-based network.

3.1 Analysis of Mesh Distances

Consider a two-dimensional rectangular mesh with r rows ($r > 1$) and c columns ($c > 1$). (An N x N mesh is simply a special case where $r = c = N$.) Assuming that the nodes are all unit distance from their neighbours, (*i.e.,* the points reside at the integer coordinates in a Cartesian plane), then what is the average distance between two randomly chosen (distinct) points within the mesh?

This question can be answered in a straightforward (but tedious) fashion using recurrence relations. If $C(i,j,r,c)$ represents the number of links traversed, starting from node (i,j), in order to reach every other possible node in an r x c mesh, using M-Unicast, then $C(i,j,r,c) =$

$$\begin{cases} \frac{rc(r+c-2)}{2} & \text{if } i=1 \text{ and } j=1 \\ C(1,1,r,c) & \text{if } i=1 \text{ and } j=c \\ C(1,1,r,c) & \text{if } i=r \text{ and } j=1 \\ C(1,1,r,c) & \text{if } i=r \text{ and } j=c \\ C(1,1,1,j)+C(1,1,1,c-j+1) & \text{if } r=1 \\ C(1,1,i,1)+C(1,1,r-i+1,1) & \text{if } c=1 \\ C(1,1,r,j)+C(1,1,r,c-j+1)-C(1,1,r,1) & \text{if } i=1 \text{ or } i=r \\ C(1,1,i,c)+C(1,1,r-i+1,c)-C(1,1,1,c) & \text{if } j=1 \text{ or } j=c \\ C(1,1,i,j)+C(1,1,i,c-j+1) & \\ +C(1,1,r-i+1,j)+C(1,1,r-i+1,c-j+1) & \\ -C(i,1,r,1)-C(1,j,1,c) & \text{otherwise} \end{cases} \tag{1}$$

After much algebraic manipulation, this recurrence relation reduces to $C(i,j,r,c) =$

$$\frac{rc(r+c-2(i+j-1))}{2} + rj(j-1) + ci(i-1) \tag{2}$$

Equation 2 is the key to our evaluation of multicast routing algorithms, since once the source node is known (for the M-Unicast algorithm), or the hub node is known (for the Conference Call model and the Core-Based Trees model), Equation 2 predicts the expected link cost of routing to each recipient in a multicast group.

From Equation 2 it is a simple matter to compute the average internode distance D for source node (i,j) in an r x c mesh:

$$D(i,j,r,c) = \frac{C(i,j,r,c)}{rc-1} \tag{3}$$

The average distance between two randomly chosen distinct nodes in the whole network can be obtained by considering each possible source node (i,j):

$$E[D] = \frac{\sum_{i=1}^{r}\sum_{j=1}^{c} D(i,j,r,c)}{rc} \tag{4}$$

In the general case ($r \neq c$), the formula evaluates to $E[D] = \frac{r+c}{3}$. For $r = c = N$, the formula predicts $E[D] = \frac{2N}{3}$.

3.2 Analysis of Multicast Routing Models

The foregoing analysis is exact for the M-Unicast routing algorithm. Thus the average number of links used per multicast call will be a strictly linear function of the group size (since recipients are assumed to be chosen uniformly at random, independently). The slope of this line will be $\frac{r+c}{3} = \frac{2N}{3}$, which depends only on the network size N. The intuition here is simple: the larger the mesh size, the further apart the recipients are (on average), and the greater the incremental cost of routing to a multicast recipient.

An exact analysis of the Conference Call model requires knowing the location of the conference hub(s) in the mesh. With this information, Equation 2 can be used to determine the (expected) cost of routing a multicast call: the source has to reach the hub, and the hub has to reach all of the recipients. This cost function will again be linear in the number of recipients, but may have a slightly higher starting point (because of the cost of connecting the source to the hub), and a lower slope than the M-Unicast algorithm (since a well placed hub is in the interior of the network, and never at an edge or corner node). For example, for 1 hub at location (2,2) in a 4 x 4 network, we expect a slope of 2.13 for the Conference Call model (instead of $\frac{2N}{3} = 2.67$ for M-Unicast), and for 1 hub at location (5,5) in a 10 x 10 network, we expect a slope of 5.05 instead of 6.67.

An exact mathematical analysis of the Core-Based Trees model is not possible, since the hub location is chosen dynamically for each group, and it makes little sense to calculate the "average" core location. Nevertheless, we do still expect the links used per call to be linear in the number of group recipients, and we can place a fairly tight upper bound on the expected slope of the line. For example, the CBT model should do about as well as (but slightly better than) having the core node always fixed at an innermost interior node in the network.

3.3 Summary

This section has presented a mathematical analysis of three simple multicast routing strategies for mesh-based networks. The approaches differ in the presence or absence of a conference hub, and in whether the hub location is static or dynamic. Mathematical analysis (and intuition) suggests that the M-Unicast model will provide the worst performance, the Core-Based Trees model will provide the best performance, and the Conference Call model will provide an intermediate level of performance.

4 EXPERIMENTAL METHOD AND RESULTS

4.1 Simulation Methodology

A set of simulation experiments was conducted to evaluate the performance of the three multicast routing algorithms described in Section 2. The purpose of these experiments was to assess the accuracy of the analytic results derived in Section 3.

The experiments were conducted using *call level simulation* on simple mesh topologies of size 4 x 4, 6 x 6, 8 x 8, and 10 x 10. A separate workload file was constructed for each mesh size, containing 100,000 synthetically generated calls. Workload parameters included the call arrival rate, the mean call duration, the mesh size, the mean multicast group size, and the percentage (10%) of the total calls that are to be multicast calls (the remainder are unicast calls). Call arrivals are assumed to be Poisson, and call durations are exponentially distributed. The number of recipients for each multicast call is drawn from a geometric distribution (a modeling assumption) with a specified mean ($g = 4$). For each mesh size, the exact same workload file is run on each multicast routing algorithm, so that performance differences in the results obtained are due to differences in the algorithms, and not differences in the workload.

In our simulation setup, we can intentionally scale the workload to make the call blocking probability as low or as high as desired. In the experiments reported here, call blocking is intentionally low ($\leq$ 1.0%), so that link usage (*i.e.,* average number of links used for each successfully routed call) is the only performance metric of interest. A more complete description of simulation experiments, simulation validation, call blocking behaviour, and other performance results is available in (Harrison 1995).

4.2 Simulation Results

Figures 1, 2, and 3 summarize the results of our simulation experiments. Figure 1 presents link usage results (both analytic and simulation) for the M-Unicast routing model, for the four different mesh sizes considered. Figures 2 and 3 do the same for the Conference Call and Core-Based Trees models, respectively. In all graphs, the data points show the mean and 95% confidence intervals for the simulation results. Solid lines denote analytically derived values.

M-Unicast Model: Figure 1 shows that the average number of links used per successful call in the M-Unicast model is strongly linear in the number of group recipients. The simulation results (plotted with mean and 95% confidence intervals[2]) for each multicast group size agree very closely with the analytic predictions.

[2] The confidence intervals naturally get larger for large group sizes because our workload (intentionally) generates fewer calls with a large number of recipients.

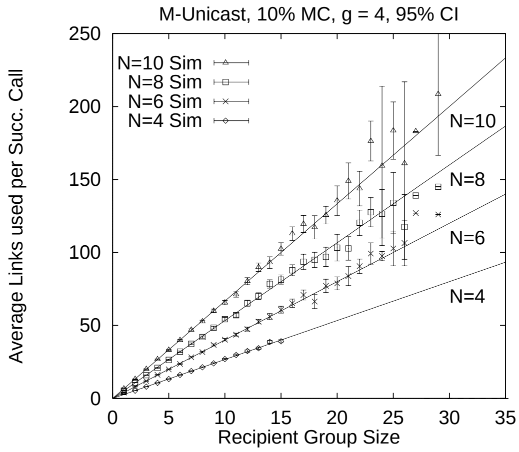

Figure 1: Link Usage Results for M-Unicast Model

Table 1 further confirms these observations. A linear regression fit of the simulation data yields a very high R^2 value (0.98-0.99). The regression shows a slope close to $\frac{2N}{3}$ for $N = \{4, 6, 8\}$. The slope value for the $N = 10$ case is distorted somewhat by the link usage results for several large multicast groups, as evidenced by a relatively large (negative) y-intercept value. For a given recipient group size, larger meshes generally show higher variability in the link usage results because of the greater range of possibilities for the physical placement of those group members.

Table 1: Regression Results for M-Unicast Model

Mesh Size	Slope	y-intercept	R^2
4 x 4	2.65	0.20	0.993
6 x 6	3.97	0.21	0.977
8 x 8	5.37	0.08	0.984
10 x 10	7.24	-4.29	0.981

Conference Call Model: Figure 2 shows the results for the Conference Call model. Again, the line is strongly linear (see Table 2), and again the analytic results match very well with the simulation data.

There are two other interesting observations to note about Figure 2. First, there is a slight "kink" in each line between the data points for 1 and 2 recipients. This kink effect arises in the Conference Call model because calls with 2 or more

Table 2: Regression Results for Conference Call Model

Mesh Size	Slope	y-intercept	R^2
4 x 4	2.08	1.52	0.998
6 x 6	3.33	0.97	0.986
8 x 8	4.05	4.11	0.980
10 x 10	4.97	5.86	0.990

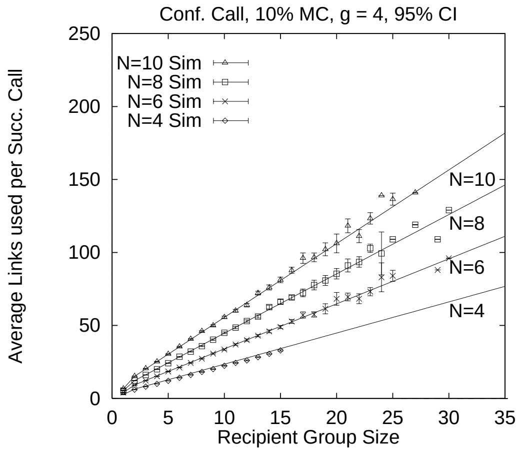

Figure 2: Link Usage Results for Conference Call Model

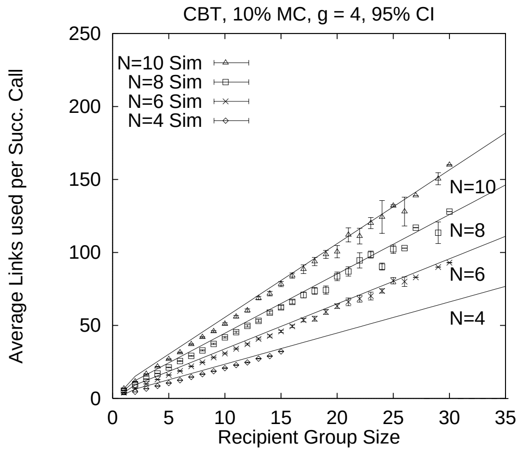

Figure 3: Link Usage Results for Core-Based Trees Model

recipients are multicast calls (which must route from source to hub, and then from hub to recipients), whereas calls with 1 recipient are unicast calls (which need not route through the hub). Thus the change in slope represents the overhead of source to hub routing. This overhead tends to increase in larger meshes (see the y-intercept values in Table 2).

The second observation is that the slopes of the lines in Figure 2 are significantly lower than those in Figure 1 (see the regression values in Tables 1 and 2). The slopes in Table 2 are 16-25% lower than the corresponding (analytically derived) slopes for the M-Unicast model, indicating that the Conference Call model, despite its additional overhead for source to hub routing (note the larger y-intercept value in Table 2 compared to Table 1), is more efficient in its routing of multicast calls.

Core-Based Trees Model: Figure 3 shows the results for the Core-Based Trees model. All lines are strongly linear (see Table 3), and all simulation results match quite closely with the analytical model. Here, however, recall that the analytic prediction is only an upper bound for each line representing simulation data.

As expected, the CBT algorithm performs better than the analytic bound, especially for small multicast groups, where dynamic core placement can lead to significantly better routes than using a static hub. The simulation results approach the analytic bound for large multicast groups since the (expected) location of the dynamic core will be in the middle of the mesh (close to the location of the static hub assumed in the analytic bound). The analytic upper bound is tightest on small meshes since there is little flexibility in dynamic core placement.

The regression analysis in Table 3 confirms the strong linearity in Figure 3, and shows the average slopes observed. While the slopes that result for the Core-Based Trees model are almost identical to those of the Conference Call model (compare Table 2 and Table 3), the Core-Based Trees model has lower y-intercept values, reflecting the lower cost of source-to-hub (core) routing when dynamic core placement is used. Thus the Core-Based Trees model consistently outperforms the M-Unicast and Conference Call routing models, using these performance metrics.

Table 3: Regression Results for Core-Based Trees Model

Mesh Size	Slope	y-intercept	R^2
4 x 4	2.07	0.22	0.998
6 x 6	3.09	0.35	0.994
8 x 8	4.11	0.75	0.995
10 x 10	4.94	2.55	0.986

Effect of Network Load: Finally, we conducted a simulation experiment to determine how rapidly the predictive ability of the analytic model degrades when the network gets congested.[3] We conducted this experiment by artificially scaling the workload[4] to create call blocking probabilities ranging from 0% to 5%. In the latter case, it is known that many calls are blocked. Thus, many links must be saturated, and non-optimal routing must occur. This experiment was conducted only for the M-Unicast model (for which the analytic model is known to be exact when the network has infinite capacity), and only on the 4 x 4 mesh.

Performance results for this experiment are shown in Figure 4. The analytic model is exact (or close to exact) when there is no call blocking, as expected. The link usage results for the simulations begin to diverge further and further from

[3] Note that the analytic model always assumes perfect routing in an infinite capacity network, while real networks end up routing around hot spots, and blocking calls when congestion occurs.

[4] This scaling was done by simply scaling the call durations. All other workload characteristics in the workload file (*e.g.,* call arrival time, source node, group size, recipients) remained the same.

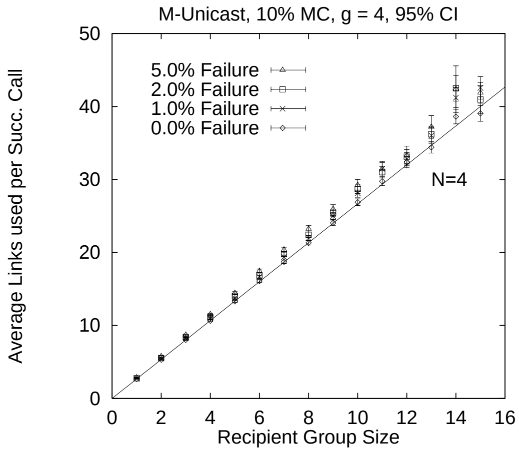

Figure 4: Effect of Network Load on Model Accuracy

the analytic predictions as the call blocking probability increases. However, the analytic model is still within 10% of the simulation results.

Regression analysis (Table 4) shows that the slope of the link usage plot changes only marginally as the workload increases. Thus the analytic model is deemed a useful predictor of the link cost for multicast calls under reasonable levels of network load.

Table 4: Regression Results for M-Unicast Model Under Higher Network Loads

Blocking	Slope	y-intercept	R^2
0.0%	2.68	-0.01	0.999
1.0%	2.86	-0.47	0.998
2.0%	2.84	-0.17	0.995
5.0%	2.87	0.12	0.997

4.3 Summary

This section has presented a simulation study of three simple multicast routing strategies for mesh-based networks. The simulation results are found to match very well with the analytic results predicted in Section 3.2. The analytic results are always within 10% of the simulation results, and diverge from the simulation results only when the network load (and thus the call blocking probability) is extremely high.

5 CONCLUSIONS

This paper has presented a performance study of multicast routing on simple mesh-based network topologies. Closed form solutions are obtained to measure the network-level performance impacts of multicast network traffic workloads, in terms of link usage.

The analytic results are used in conjunction with simulation to evaluate three possible multicast routing techniques: M-Unicast, Conference Call, and Core-Based Trees. Among the three routing algorithms evaluated, the Core-Based Trees model provides the most efficient multicast routing. Throughout the performance analysis, the analytic results and the simulation results are found to agree very closely (0-10%). Simulation results diverge from the analytic results only in highly overloaded networks (*i.e.*, when significant call blocking occurs).

ACKNOWLEDGEMENTS

Financial support for this research was provided by the Natural Sciences and Engineering Research Council (NSERC) through research grants OGP0121969 and IOR152668, and by Telecommunications Research Laboratories (TR*Labs*) in Saskatoon.

REFERENCES

Ballardie, T.; P. Francis; and J. Crowcroft. 1993. "Core-Based Trees (CBT): An Architecture for Scalable Inter-Domain Multicast Routing". In *Proceedings of the 1993 ACM SIGCOMM Conference* (San Francisco, CA, Sept. 13-17), pp. 85-95.

Deering, S. 1991. *Multicast Routing in a Datagram Internetwork*. Ph.D. Thesis, Stanford University.

Deering, S.; D. Estrin; D. Farinicci; V. Jacobson; C. Liu; and L. Wei. 1994. "An Architecture for Wide-Area Multicast Routing". In *Proceedings of the 1994 ACM SIGCOMM Conference* (London, UK, Aug. 31-Sept. 2), pp. 126-135.

Harrison, T. 1995. *Performance Evaluation of Routing Algorithms for Multicast Traffic*. M.Sc. Thesis, Department of Computer Science, University of Saskatchewan (Aug.).

Kompella, V.; J. Pasquale; and G. Polyzos. 1993. "Multicast Routing for Multimedia Communications", *IEEE/ACM Transactions on Networking*, **1**, no. 3 (June), pp. 286-292.

Moy, J. 1994. "Multicast Routing Extensions for OSPF", *Communications of the ACM*, **37**, no. 8 (Aug.), pp. 61-66.

Verma, D.; and P. Gopal. 1993. "Routing Reserved Bandwidth Multi-point Connections", In *Proceedings of the 1993 ACM SIGCOMM Conference* (San Francisco, CA, Sept. 13-17), pp. 96-105.

Waitzman, D.; C. Partridge; and S. Deering. 1988. "Distance Vector Multicast Routing Protocol", RFC 1075.

Tim Harrison is a full-time Research Assistant in the Department of Computer Science at the University of Saskatchewan in Saskatoon, Canada. An M.Sc. graduate of the University of Saskatchewan in 1995, his research interests include multicast routing, MBONE, and ATM.

Carey Williamson is an Associate Professor in the Department of Computer Science at the University of Saskatchewan. His interests include network traffic measurement, workload characterization, congestion control, multicast routing, high speed networks, and ATM.