Mobile multicasting in wireless ATM networks

D.-K. Kim * and C.-K. Toh

Mobile Multimedia and HiSpeed Networking Laboratory, School of Electrical & Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

Multicasting is a desired function in both wired and wireless networks. Currently, multicasting-based applications have pervasive presence and influence in our Internet. Wireless ATM aims to support different traffic types over a high speed wireless network. In particular, traditional approaches towards multicasting in wired networks cannot be directly applied to mobile ATM networks. This includes the mechanism proposed to handle mobility in wired networks. In this paper, we address the issues behind multicasting in wireless ATM networks and the underlying mechanisms to support handoffs of both unicast and multicast connections. Through simulation, we evaluated the performance of optimal core-based tree (CBT), late-joined CBT, optimal source-based tree (SBT), and late-joined SBT during changing multicast host membership conditions. In addition, we evaluated the impact of multicast receivers' migration on the total link cost of the resultant multicast tree. Simulation results revealed that source-based trees incur more total link costs under both dynamic host membership and receivers' migration scenarios. However, it generates less control messages than the distributed CBT and SBT approaches.

1. Introduction

Asynchronous Transfer Mode (ATM) technology has received attention in both the computer networking and the telecommunications industries. Along with it, the area of wireless and mobile communication networks has initiated an enormous amount of activities. There is a trend towards mobile and ubiquitous computing where the users will migrate from fixed hosts running on broadband (ATM) nets to portable devices running over the air medium. To satisfy the needs caused by that movement of interest, existing networks have to be extended and enhanced for wireless. Multimedia applications are the key applications that need and use the large bandwidth provided by ATM. These include: audio–visual conferencing and whiteboard sharing, distance learning, video-on-demand, virtual reality, tele-metering and remote sensing, data distribution (news, weather, stocks), server/router synchronization (DNS/router updates), advertising and location servers and communicating to unknown/dynamic group.

The common characteristic of the applications mentioned above is that they are multidestination applications. In that context, a sender might want to or might have to send the same information to multiple receivers. This increases the attention that must be given to point-to-multipoint connections.

While earlier research has been focused on point-to-multipoint or multicast, it was done only for traditional ATM networks. In the WATM domain, very few researchers have investigated handoff protocols for mobile multicast connections. There are a few points which are considered to be important for implementing wireless ATM multicasting:

(a) one address, (VP/VC) is used per group,
(b) members are located anywhere,
(c) members can join and leave at will,
(d) sender need not be aware of membership,
(e) sender need not be a member,
(f) uniform QoS must be provided to all hosts, and
(g) QoS fulfillment check during route establishment.

These points should be considered when performing a multicast setup.

It is a fact that in wireless and mobile networks the main challenge is the rapidly changing link environment. Multicasting makes the situation even more complex since the number of users on one connection is usually more than two and the changes in topology might be more often. Multicast service in ad hoc networks is even more complex because of the more dynamic nature of them.

2. Wireless ATM

WATM is an extension of the terrestrial ATM system, which allows the last hop of the ATM connection to be carried over a wireless link. As with mobile telephones, WATM uses a similar cellular coverage. Thus, when a mobile host (MH) leaves one cell and enters another, the system must be able to "handover" the connection. The way in which the system handles handovers greatly affects QoS and, hence, users' satisfaction of communication service.

2.1. Wireless ATM components

The WATM system uses the ATM terrestrial backbone as its basis, and is wireless only on the last hop. The mobile

* D.-K. Kim is a visiting research scholar at GeorgiaTech and is a Ph.D. candidate with the Department of Computer Engineering, Seoul National University, Korea.

hosts (MH) each uses a radio-frequency terminal adapter (radio-modem) to communicate with base stations (BS). The BSs connect the wireless link to the terrestrial backbone. There is a proposed standard by the ATM Forum which defines an extended protocol stack (revised physical and data-link layers) for the wireless medium. The BSs are connected to Mobility Enhanced ATM-Switches (MESs) which are part of the ATM backbone. The MESs are loaded with software, which allows them to support handovers and mobility. A collection of base stations which are connected to a common MES is defined as a cluster. The common switch which connects these BSs is known as a cluster switch (CLS). Location, connection, and multicast servers are items, which may also exist in a WATM network.

2.2. Wireless ATM handovers

One of the most important issues facing wireless ATM is providing seamless mobile switching, i.e., handovers. Ideally, there should be minimal traffic disruption during the handover. In addition, there should be no cell loss or out-of-sequence delivery. Finally, due to the transient nature of the radio link, the handover mechanism must support some type of QoS renegotiation. There two different types of handovers, intra- and inter-cluster. Intra-cluster handovers occur when the MH is leaving one cell and entering another, where both cells belong to the same cluster. The switching and processing required for an intra-cluster handover can therefore be localized to the CLS and the two BSs involved. Because the switching is done entirely at the local CLS, it is inherently fast, causing only a minimal disruption to the WATM connection. On the other hand, inter-cluster handovers are a result of a MHs migration from a cell in one cluster to a cell in a different cluster. As the original base station (BSold) and the new base station (BSnew) do not share a common CLS, a different switching point must be located. The inter-cluster handover is handled using partial path reestablishment [7]. This approach attempts to reuse the original connection path, while seeking only to reestablish a short partial path to the new BS. The point at which the new partial path connects to the original path is called the crossover switch (CX). Because of spatial locality, it is very likely that the reestablished path shares most of the virtual paths in the original route. Buffering is used to ensure in-order delivery of all cells. Routing and crossover switch discovery are important factors governing the effectiveness of partial path reestablishment.

2.3. Routing and crossover switch discovery

The ATM Forum has adopted the PNNI (Private Network to Network Interface) protocol as the standard for routing and signaling in wireless ATM [4]. PNNI includes two categories of protocols:

Routing protocol. This protocol communicates the information used to compute paths through the network. It is based on a hierarchical model so that it will scale well with increasing network size. It can also automatically reconfigure itself in networks in which the address structure reflects the topology. PNNI uses the well known link state routing technique and includes QoS metrics when it makes the routing decisions.

Signaling protocol. PNNI signaling is based on a subset of the ATM Forum UNI signaling with some added mechanisms to support source routing and provide for dynamic call setup. It makes use of the information gathered by PNNI routing to dynamically establish, maintain and clear ATM connections at the PNNI interface between two ATM nodes.

The protocol we discuss here adds another module in the PNNI architectural model (figure 1). Specifically a CX partial route determination module is added to the topology layer and it communicates directly with the topology database. Hence, it makes use of the information gathered and propagated by the PNNI routing protocol.

The PNNI specifications do not specify any single required algorithm for path selection. Rather each implementation is free to use whatever path selection is felt appropriate. There are many different methods that can be used for crossover switch discovery, including: loose select, prior path knowledge, prior path optimal resultant, distributed

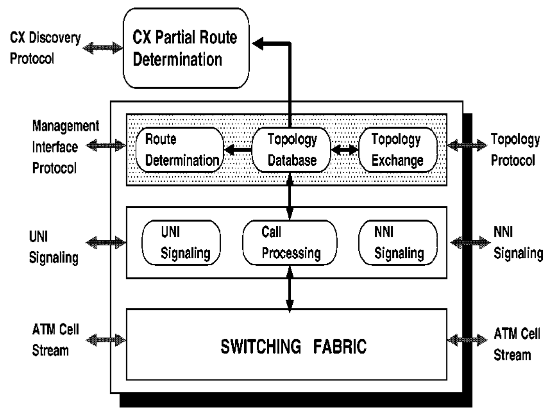

Figure 1. Interaction of CX discovery with P-NNI routing.

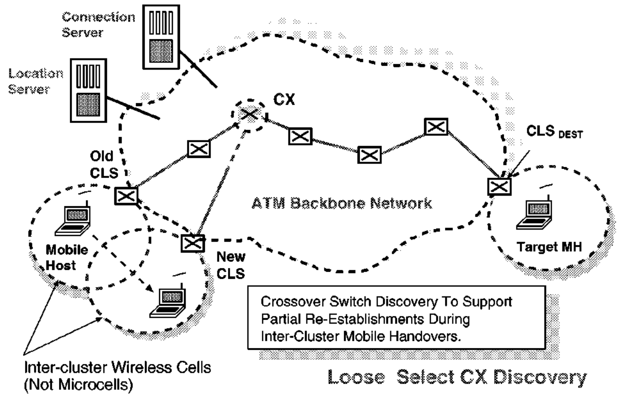

Figure 2. A WATM network structure showing the CX switch.

hunt, and backward tracking [8]. Figure 2 shows the position of the CX in a WATM network. The position of the CX might change depending on the method used, but the important thing is that it will still be on the previous path towards the MH and will try to minimize the hops to the newBS and reuse as much of the old path as possible.

2.4. Mobile ATM connection management

Since ATM and subsequently wireless ATM networks are connection oriented, a virtual circuit has to be established before any actual data transfer can take place. Connection management can be performed in a centralized or a distributed fashion. The former requires all connection requests to be directed to a connection server. The connection server then derives a possible route and interacts with all the network elements in the route (ATM switches, base stations) for possible call admission [3]. Once a route is granted, an acknowledgement is relayed back to the sender node. The second approach uses a hop-by-hop call setup scheme, where setup packets are forwarded from one switch to another with call admissions performed at all the network elements encountered [7]. The former maintains call states at the connection server compared to maintaining them at the switch in the latter.

3. Problems in wireless multicasting

Multimedia applications are often multicast in nature, meaning that a single (or multiple) source is sending to multiple receivers. MBONE, the Virtual Multicast Backbone, was first experimentally deployed in 1992. From its inception, this concept has been well received and allows users to receive audio, video, newsfeeds, etc., from a virtually unlimited number of sources. However, in its current state, multicasting is not a reality since Internet protocols today cannot provide quality of service guarantees. This may be acceptable for an informal videoconference, but in situations where quality must be ensured, a different solution must be used. Thus, the interest is in ATM, which provides guaranteed quality of service.

When extending the multicast concept to the wireless environment, an additional challenge is faced in that any of the senders or receivers could be mobile. When these mobile hosts migrate, any disruption of service can cause an interruption not only to that user, but possibly to the entire multicast group as well. An improper handover may result not only in severe traffic flow interruptions, but may also result in the loss of several group members. The handover protocol is independent of the number of members in the multicast group, meaning that any protocol must be scalable to support multicast groups of varying size. Providing seamless service in this situation is the primary goal of the multicast handover protocol. In addition, having separate handover protocols for unicast and multicast connections can lead to additional system complexity. An algorithm would be needed to decide which handover protocol should be used for each connection and situation. Thus, a handover protocol that can support handovers for both unicast and multicast mobile connections in a WATM system is desirable.

4. Wireless ATM multicasting

Increasingly, multimedia applications are requiring multipoint communications. Point-to-point connections are classified as unicast connections. Multicast connections can be either point-to-multipoint or multipoint-to-multipoint connections. ATM multicasting requires the use of distribution points (DP). A distribution point is an ATM switch, which is able to duplicate cells and forward them to multiple receivers. DPs are identified once a routing algorithm derives the multicast tree. There are two main functions provided by a multicast routing protocol. One function is to derive a multicast tree (based on certain QoS constraints) so that a multicast connection can be established while the other function is to handle problems associated with dynamic multicast groups.

4.1. Multicast routing

Existing multicast routing algorithms may be classified [10] into: (a) Minimum Steiner Tree (MST), (b) Shortest Path Tree (SPT), and (c) Constrained Multicast Tree (CMT). Their characteristics are summarized below:

(a) *Minimum Steiner trees.* In MST, a minimum cost-spanning tree per source-group pair has to be computed. The total cost of a multicast tree is defined as the sum of the cost associated with all its edges. Constructing minimum-cost Steiner trees is a well-known NP-complete problem. Hence, several heuristic algorithms to construct near-optimal Steiner trees, based on global cost information available at the source node or based on local cost information maintained at each node, have evolved. Although MST algorithms attempt to minimize the cost of the tree, they do not consider the end-to-end delay requirement.

(b) *Shortest path trees.* SPT can be further classified into source-based, core-based, and hybrid-based trees. Examples of source-based SPT protocols are DVMRP and MOSPF. While MST minimizes cost, source-based SPT minimizes the end-to-end delay. CBT is an example of a non-source-based SPT. Variations of the CBT, such as GCBT (Group Core-Based Tree) and DCBT (Dynamic Core-Based Tree) have been proposed. Finally, PIM (Protocol Independent Multicast) combines both the earlier schemes to provide added flexibility for applications that demand specific QoS requirements. This, along with HMR (Hierarchical Multicast Routing protocol), falls under the hybrid SPT category. Recently, the RBM (Reservation-Based Multicast) routing protocol which combines the feature of PIM and RSVP (Resource Reservation Protocol) was introduced.

(c) *Constrained multicast trees.* CMT combines SPT and MST to consider both cost and delay during multicast tree computations. This has been motivated by the need to support multipoint multimedia communications. CMT algorithms are source-based and they assume that sufficient information is available at the source node in order to compute for an appropriate multicast tree. While the availability of up-to-date and useful QoS information has been seen as a difficult problem in the past, recent developments of the P-NNI (Private Network-to-Network Interface) routing protocol by the ATM Forum have brought hopes that a scalable and QoS-based routing protocol is possible. This, therefore, provides the necessary QoS information required by existing source-based CMT multicast routing algorithms. The proposed protocol is intended to be applicable to all these trees. In this paper, however, we will restrict our simulation to the Shortest Path Tree (SPT) category, especially the source- and core-based trees.

4.2. Methods of performing multicast

Source-based trees (figure 3) require that the multicast source (sender) construct the entire multicast tree. The source must first know who the receivers are and this is achieved by consulting a multicast server. When the source wishes to establish a multicast connection, it generates the multicast tree by discovering the optimum route between itself and each of the receiving nodes. The source node communicates with all multicast receivers over point-to-multipoint connections.

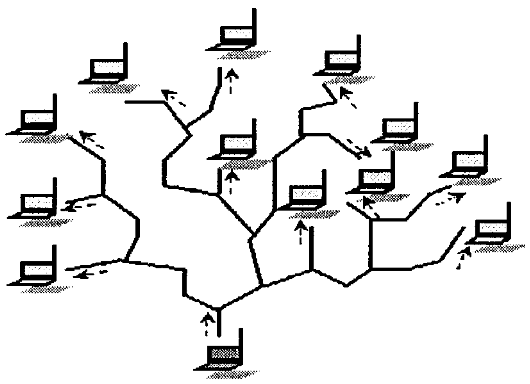

Figure 3. Source-based tree.

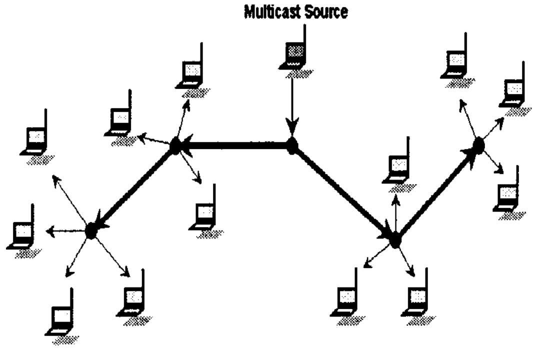

Figure 4. Core-based tree.

Core-based tree (figure 4) uses the shared-tree approach, where all multicast communications will traverse a common “shared-tree”. CBT uses a number of collective hosts as a group. Each group has a “core” ATM switch that performs all multicast routing for all members of the group. Any member of a group wishing to send or receive multicast data must communicate this to its group’s core. The connections between all cores that have members of the multicast group form the CBT backbone, or delivery tree. The CBT approach was designed to be a scalable architecture for multicasting, and thus is most appropriate choice for any large scale network with multimedia multicast applications. Since CBTs route cells from one core in the CBT backbone to another core, routing becomes fast and simple. However, routing in a CBT is generally non-optimal because all routes must use the shared tree, it is likely that many connections will not be the shortest path.

Source-based routing scales O(number of total nodes in tree), while core-based routing scales O(number of groups). This significantly impacts the performance and scalability of the multicast trees.

5. Integrated unicast and multicast handoffs

The handover scheme we consider here is an enhancement of the hybrid handover protocol proposed in [7] that has two distinct parts:

(a) the setting and tearing down of connections and connection rerouting at the CX, and

(b) the discovery of a suitable crossover switch.

By modifying the CX discovery algorithm, handovers of multicast connections can be supported in a seamless manner using partial reestablishment. This facilitates resource reuse and avoids full reestablishment. By doing that we get a much less severe interruption in our traffic flow and we also reduce the potentially long setup time. Of course, such an approach might lead to situations where, after a number of handovers or new members joining or leaving, we may end up with a non-optimal tree.

Handovers are also present in the multicast connections in WATM networks and can be initiated by both the sender and the receiver in a mutually exclusive manner. We propose the use of the same handover protocol and CX discovery as in the unicast situation with some differences. By doing that we avoid having a completely new and separate handover protocol for multicast and, therefore, we provide uniformity and unification in the actions performed by the network elements. Unlike mobile unicast connection handovers, the goal of the enhanced protocol procedures is to ensure that the connection rerouting and the partial path deletion operations will be performed in such a manner such that no problems will occur. Potential problems are:

(a) the deletion of the old path while there are still MHs using connections to the multicast group through it, and

(b) the change in the direction of data flow.

The first problem can be resolved by having the old BS checking the links it maintains and decides if a complete teardown can be done. In the case that there are still connections to the multicast group through it, it will acknowledge the teardown only if the remaining connections use different VP/VCs. If they use the same VP/VC as the station that moved away, then the connection is not released, but a "passive" tear-down-ack message is sent back to the crossover switch (CX) to inform it to retain the existing connection. This, therefore, avoids pruning off valid multicast members that might have been gone otherwise. The second problem is addressed by considering nodes for convergence on the path between the old CLS and the Distribution Point (DP).

Figure 5 shows a WATM mobile multicast connection with one sender and two receivers. When the sender node migrates, an inter-cluster mobile handoff is initiated. By selecting an appropriate crossover switch (which must reside in the path from CLSnew to DPx), the resultant multicast tree does not prune away existing members from the multicast tree and there is no disruption (in terms of flow) to traffic flows to other existing members in the tree.

Apart from handovers, WATM multicast connections are set up or deleted in the network also due to dynamic host membership. This is mostly true for the current applications that use multicasting like multi-party conferencing where users connect and leave at will. The issues for the dynamic host membership have been studied in the ATM Forum's Signaling Sub-working Group under the title "Leaf-initiated join extensions" [1].

In the next section, we will present our simulation approach and results based on the evaluation of the source- and core-based trees under dynamic host membership conditions. Specifically, we will reveal the difference between the total signaling messages required and the link costs for optimal and non-optimal trees after a number of users have joined the multicast group.

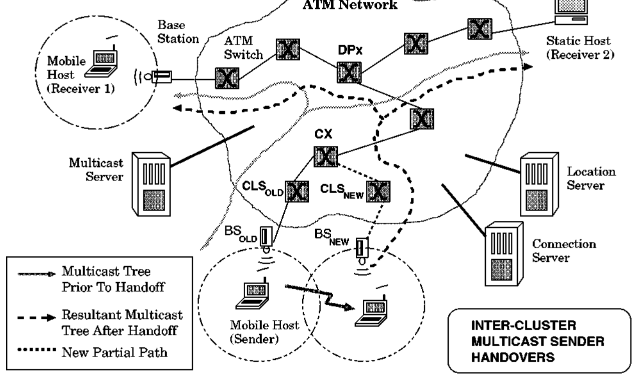

Figure 5. Reconfiguration of the multicast tree after a sender's handover.

6. Simulation of dynamic host membership

To perform simulation, a random graph consisting of 50 ATM switches located at random points of 50×50 coordinates is generated. We adopted a model which is different from Waxman's model in that the mean degree of each node is introduced in the real network. In [2], the connectivities between nodes are determined by the probability function:

$$P_e(u,v) = \frac{ke}{n}\beta \exp\left\{\frac{-d(u,v)}{\alpha L}\right\},$$

where e is the mean link degree of each node; $d(x,y)$ is the Euclidean distance between the nodes; α is a parameter related to the number of connections to nodes from further away; β is another parameter controlling the number of edges from each node; n is the total number of nodes in the graph; L is the maximum possible distance between two nodes. Therefore, ke/n is a scaling factor ensuring that the mean degree of each node remains constant regardless of the network size. In this simulation, k, α and β have values of 25, 0.25 and 0.2, respectively.

By following the above-mentioned rules, a generated random graph whose nodes are fully connected is shown in figure 6. Each node is located at the corresponding logical coordinates. In addition, the links between nodes are logically connected. In other words, the location of each node can be relocated for better visualization. Additionally, each link has its cost which is the distance between two nodes, or the amount of bandwidth depending on communication requirements.

Based on figure 6, we selected five random nodes as the initial receiver-members of a multicast group. We simulated two schemes for generating a multicast tree with some late joiners: Core-Based Tree (CBT) and Source-Based Tree (SBT). In both core- and source-based trees, we established an initial multicast tree consisting of one sender and five receivers. By adding nodes into the multicast group, we

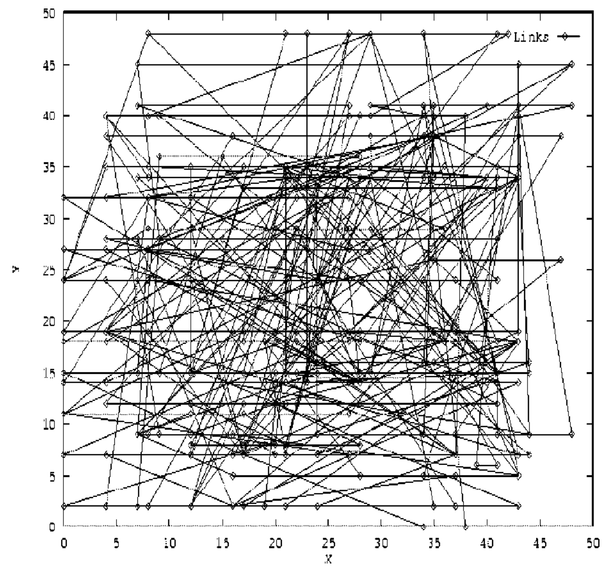

Figure 6. Initial random graph.

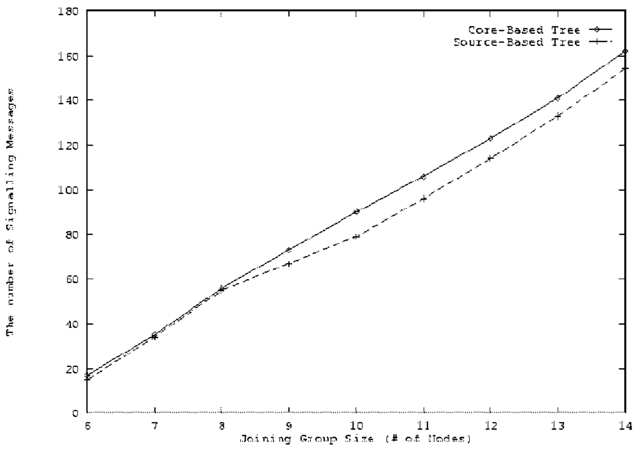

Figure 7. Comparison of total signaling messages over a network.

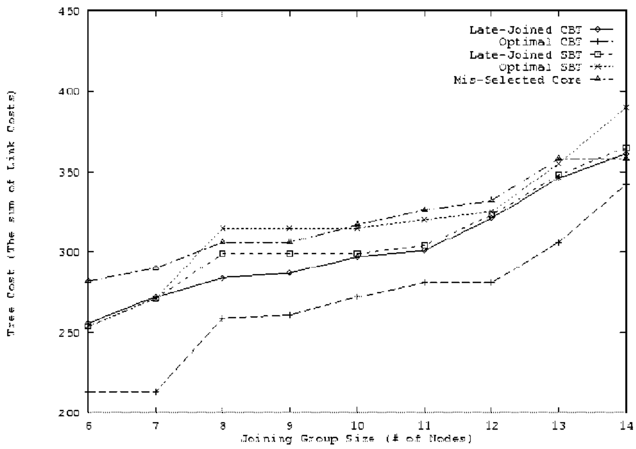

Figure 8. Total link costs over a tree.

applied the receiver-initiated multicast scheme to each tree. All receivers wanting to join the group broadcast their JOIN request packets over a network. All nodes over a tree can respond to the requests generated by late joining receivers. The reply packets can follow the reverse path from the joiners to the nodes over a tree. After receiving multiple replies, the joiner determines which node it will attach to for making the branch of the tree. In this simulation, the joiner selects the shortest path from itself to the nodes generating the replies. Consequently, we can check the number of these signaling messages exchanged. As shown in figure 7, as the number of joining nodes grows, the total number of signaling messages also increases. In this simulation, since a core-based tree has more intermediate nodes than that of a source-based tree, the latter shows better performance from the point of view of signaling messages.

In addition, we measured the total link costs of a tree which is the sum of each link cost, and not the total cost of individual paths from the sender to the receivers. We evaluated the following four cases:

1. *Optimal core-based tree.* After the central server[1] is notified of a new joiner, it can regenerate a new core-based tree. This is the ideal case. However, the signaling overhead involved is high to maintain the tree. In addition, it may be necessary to select a new core.

2. *Late-joined core-based tree.* The receiver can join the group by broadcasting a request message, receiving multiple replies, and determining the best attachment point over a core-based tree in a distributed manner.

3. *Optimal source-based tree.* After the central server is notified of a new joiner, it can generate a new source-based tree. However, the signaling overhead involved is too high to maintain the tree.

4. *Late-joined source-based tree.* The receiver can join the group by broadcasting a request message, receiving the multiple replies, and determining the best attachment point over a source-based tree in a distributed manner.

[1] A central server may be a sender or a multicast server.

In case of the optimal source-based tree, a new joiner can be reached through a new shortest route without traveling through the intermediate nodes over an existing tree. Therefore, when compared with the case of late-joined source-based tree, the optimal source-based tree incurs more total link costs than late-joined source-based tree (figure 8).

However, when a fixed core continues to be maintained for late joining receivers, in the environment where the terminals move around the network, the distance from the switch node which the terminals contact to the core grows larger. Therefore, it results in making the total link costs higher. From this simulation, we can conclude that in the environment where the mobile nodes move around across many switches, the late joined source-based tree can be a candidate for a multicast tree from the point of view of the number of signaling messages and total link costs.

7. Simulation of mobility by multicast receivers

Another important aspect for consideration is the performance of the various multicast routing schemes at times of host mobility. Therefore, we perform another set of simulation to evaluate the total link cost associated with the resultant multicast tree after migration by multicast receivers.

Using the same network model as mentioned earlier, we establish multicast trees based on a subset of a source node and several receiver nodes. The trees are established using each of the multicast routing protocols and receiver nodes are moved away, thereby losing their connectivity to their old base stations and gaining connectivity with the new base stations. A partial path is then established to a node in the multicast tree. The partial path and the crossover switch are selected based on the integrated prior path crossover switch discovery scheme proposed in [9]. The total link cost of the resultant multicast tree is then noted for handoffs caused by subsequent multiple receivers. The results are plotted and compared with those using different multicast routing approaches.

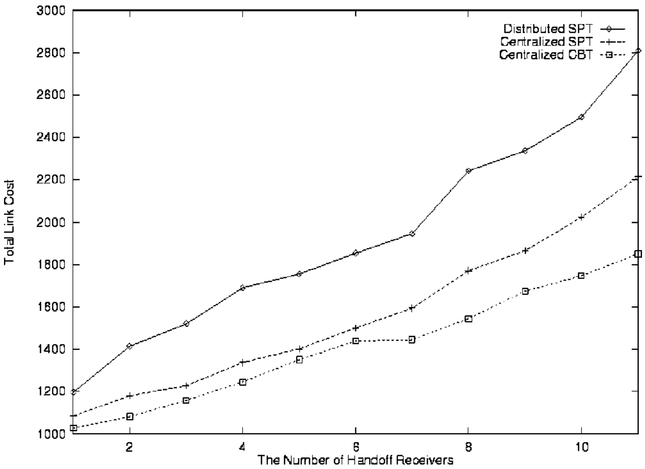

Figure 9. Comparison of resultant multicast tree link cost after receivers' migration.

As shown in figure 9, the CBT approach has the least total link cost when compared to the source-based centralized or distributed SPTs after receivers' migration. The concept of shared delivery trees has resulted in higher link cost efficiency and this is more significant when the multicast group size increases. The centralized SPT performed better than the distributed SPT due to the ability of the centralized node to recompute the next optimal tree.

8. Conclusion

In this paper, we have investigated the impact of dynamic host membership and receivers' migration on the performance of mobile multicast in WATM networks. In the case of both centralized optimal core-based and source-based approaches, one could take into consideration all requirements related to bandwidth and delay when it comes to multicast tree establishment. However, one should also consider the overhead resulting from routing information updates whenever a late joiner wants to participate in an on-going mobile multicast session. In addition, such an approach demands a centralized server, which gives a burden to the server when it comes to fault tolerance and traffic concentration for control messages exchanges.

On the other hand, the two distributed multicastrouting methods proposed construct branches to the multicast tree without the limitation of a centralized server. This results in the reduction on the total link costs of the multicast tree. In terms of delay or bandwidth requirements, a late joiner can determine the best next-hop among multiple replies generated from intermediate nodes. However, in case of distributed core-based tree, if the core is located far away from the sender and receivers, it is still very difficult to obtain an optimal multicast tree. In this paper, we showed the possiblity of applying late-joined source-based tree approach to a multicast environment. We have also examined the performance evaluation of WATM multicasting under the effects of mobility by receivers.

Acknowledgement

This article is an extended version of our previously published article submitted to wmATM'99 workshop, organised by the Wireless Mobile ATM Task Force of Delson Group.

References

[1] R. Bubenik and P. Flugstad, Leaf initiated join (LIF) extensions, ATM Forum Document 94-0325R2 (1994).

[2] M. Doar, Multicast in the ATM environment, Ph.D. thesis, Cambridge University Computer Laboratory (1993).

[3] T.F. La Porta and R.R.M. Veeraraghavan, A distributed control strategy for wireless ATM networks, Wireless Networks (WINET) 1(3) (1995).

[4] Private network–network interface specification v1.0 (PNNI 1.0), ATM Forum specification af-pnni-0055.000 (1996).

[5] C.-K. Toh, A handover paradigm for wireless ATM LANs, in: *Proc. of ACM Sympos. on Applied Computing (SAC)* (February 1996).

[6] C.-K. Toh, Crossover switch discovery for wireless ATM LANs, Mobile Networks and Applications – Special Issue on Routing in Mobile Communication Networks 1(2) (1996).

[7] C.-K. Toh, A hybrid handover protocol for wireless ATM LANs, Mobile Networks and Applications – Special Issue on Wireless ATM 1(3) (1996) 313–334.

[8] C.-K. Toh, *Wireless ATM and Ad-hoc Networks: Protocols and Architectures* (Kluwer Academic, Dordrecht, 1996).

[9] C.-K. Toh, A unifying methodology for handovers of heterogeneous connections in wireless ATM networks, ACM SIGCOMM Computer and Communications Review 27(1) (1997).

[10] W. Yen and I.F. Akyildiz, A Hierarchical Multicast Routing (HMR) protocol, in: *Proc. of IEEE Computer Communications Workshop* (October 1995) pp. 154–156.

Dongkyun Kim received B.E. degree in computer engineering from Kyungpook National University, Taegu, in 1994, and the M.E. degree in the Multimedia and Communication Laboratory from Seoul National University, Seoul, Korea, in 1996. He is now a Ph.D. candidate in the Multimedia and Communication Laboratory and a visiting scholar to Georgia Institute of Technology, Atlanta, Georgia. His research interests are multimedia systems, wireless ad hoc network, mobile computing and routing protocol.

C.-K. Toh was educated at the EEE Department, University of Manchester Institute of Science & Technology, and the Computer Laboratory, University of Cambridge, England, in electrical engineering and computer science. His Cambridge Ph.D. was focused on "Protocol Aspects of Mobile Radio Networks". Dr. Toh founded and chaired the Cambridge Mobile Special Interest Group from 1994–1996. He works on several aspects of mobile networks (routing, protocols, algorithms, mobility management, etc.). Prior to Cambridge, Dr. Toh was a Network Specialist, R&D Engineer and Technical Staff Member, working on multimedia communications over Broadband ISDN. At Cambridge, he was an Honorary Cambridge Commonwealth Trust Scholar and King's College Cambridge Research Scholar. Dr. Toh is the inventor of the Cambridge ad hoc mobile routing protocol, which is currently pending a US patent. He authored the

book on "Wireless ATM & Ad-Hoc Networks", which was published by Kluwer Academic Press. Dr. Toh is an Area Editor on wireless networking for IEEE Communications Survey Journal, a Feature Editor for ACM Mobile Computing and Communications Review and serves as an editorial board member for IEEE Network Magazine and Springer-Verlag Personal Technologies Journal. Dr. Toh is a Member of Sigma Xi Honor Society, New York Academy of Science, American Association for Advancement of Science, IEE, IEEE, ACM, USENIX, a Senior Member of King's College Cambridge, a Fellow of the Cambridge Commonwealth Society and a Fellow of Cambridge Philosophical Society. He serves as a panel member for Wireless Communications Research at the National Science Foundation (NSF) and is listed in International Who's Who of Professionals (1997) and Who's Who in Information Technology (1998). He was with Hughes (USA) leading research work on ad hoc wireless networks, satellite and wireless ATM networks. He was Project Technical Leader for DARPA TTO "Deployable and Adaptive Ad Hoc Mobile Wireless Networks (DAMAN)" program. As of 1998, he joined Georgia Institute of Technology as an Assistant Professor of Electrical and Computer Engineering, directing the Mobile Multimedia and Networking Laboratory. Prof. Toh is a registered Chartered Electrical Engineer (CEng), UK.