Extending Seamless IP Multicast Edge-Coverage Through Mobile Ad Hoc Access Networks

Pedro M. Ruiz[1], Antonio F. Gomez-Skarmeta[1], Pedro Martinez[1], and David Larrabeiti[2]

[1] University of Murcia, Facultad de Informatica, Dept. of Information and
Communication Engineering, Campus de Espinardo,
E-30100 Espinardo (Murcia), Spain
`{pedrom,skarmeta,pma}@dif.um.es`
[2] University Carlos III of Madrid, Dept. of Telematic Engineering,
Campus de Leganes, Avda. Universidad, 30,
E-28911 Leganes (Madrid), Spain
`dlarra@it.uc3m.es`

Abstract. The provision of multicast communications in wireless and wired networks has followed different paths which have led to different solutions. Little has been accomplished to-date in bringing together the traditional IP multicast model used in fixed networks and multicast routing protocols for wireless ad hoc networks. We analyse the provision of an integrated IP multicast service in which mobile hosts can seamlessly participate in IP multicast sessions regardless of the currently underlying network type. We propose a multicast architecture in combination with a new ad hoc multicast routing protocol called MMARP. MMARP nodes are challenged with special IGMP-handling capabilities allowing our solution to combine the efficiency of multicast ad hoc routing protocols with the support of standard-IP nodes without an impairment in the performance of the protocol. Our empirical results demonstrate that such kind of multicast ad hoc access networks offer a good performance when compared with the traditional single-hop wireless multicast access.

1 Introduction

IP Multicast is suited for efficient multipoint communications among a group of nodes. It has emerged as one of the most researched areas in networking. The problem of efficient packet distribution to a specific group of destinations has been researched since the late 80's and most of the routers nowadays support IP multicast routing protocols. The main benefit of IP Multicast is that the bandwidth consumption for group communications is dramatically reduced compared to unicast-based group communications. This is of particular interest for 'all-IP' and 'beyond 3G' mobile networks consisting of a high number of user terminals using applications which are typically interactive, multiparty and bandwidth-avid.

Many projects like the IST project MIND (Mobile IP-based Network Developments)[1] have researched the extension of IP-based radio access networks to include ad-hoc wireless elements within the access infrastructure as a natural evolution towards 'beyond 3G' systems. In this ad hoc fringe, a user terminal employs those of other users as relay points to provide multi-hop paths between mobile nodes and the fixed access network architecture.

The provision of an integrated IP multicast service in such an heterogeneous scenario consisting on traditional IP core networks interconnecting a variety of wireless and wired access networks and technologies is extremely complex. There are specific solutions for wireless ad hoc networks, but the real challenge is their effective and efficient integration with (fixed) IP multicast protocols to achieve a seamless IP multicast service in which group members from any of these network types can take part in the same IP multicast session. Furthermore, mobile nodes should be allowed to move among these types of networks without any service disruption.

To our knowledge, for the specific problem of IP multicast interworking between IP access networks and wireless and mobile ad hoc networks, there are not satisfactory solutions so far. The typical intra-domain IP multicast protocols for fixed networks (i.e. IGMPv2[2] for multicast group membership and PIM-SM[3] for IP multicast routing) are not able to deal with the quick and unpredictable link changes which characterise ad hoc networks. They would consume too much overhead to keep updated distribution paths in such variable topologies. In addition, multicast ad hoc routing protocols like CAMP[4], ODMRP[5], and ADMR[6] among others, incorporate specific functionality which enables them to cope with the particular characteristics of ad hoc networks but they are only suitable for isolated ad hoc networks. These protocols do not provide any means to interoperate with the protocols used in the fixed IP networks and they do not support the attachment of standard IP multicast nodes to the ad hoc extension. In fact, the only few proposals to connect ad hoc networks to the Internet, like the one by Lei and Perkins[7] have only considered the case of unicast traffic.

In this paper we propose an integrated IP Multicast solution for ad hoc network extensions consisting of a novel IP multicast architecture and the Multicast MAnet Routing Protocol (MMARP). MMARP is a new multicast ad hoc routing protocol based on the same basic mechanisms as other ad hoc multicast routing protocols. However, it incorporates additional functionalities to deal with the complexity of supporting traditional IP nodes whilst interoperating smoothly with fixed IP networks. MMARP nodes are able to intercept and process standard IP multicast messages. They further permit standard IP nodes to seamlessly participate in IP multicast communications as they do when attached to a fixed IP network. The novelty of our approach is not only the provision of such an integrated IP multicast solution, but also the way in which the functions are divided among the fixed and ad hoc nodes so that the interworking is achieved without a noticeable impairment in the overall performance.

The remainder of the paper is organised as follows: section II comments on the problems, requirements, and proposed architecture for ad hoc access network extensions. A detailed description of the MMARP protocol is given in section

III. Section IV presents some empirical results. Finally, section V gives some conclusions.

2 Proposed Multicast Architecture

One of the most important design issues in the multicast architecture for seamless IP multicast provision in ad hoc network extensions is the separation of the functions between the different network boundaries. We followed a top-down approach which allowed us to derive the best design options from the particular requirements and related issues of a seamless and integrated multicast solution.

2.1 Requirements

The first step towards an integrated IP multicast solution is the identification of the requirements. As an objective for ad hoc network extensions we seek a trade-off in which at least the following requirements are met:

- Interoperability with IP Multicast mechanisms in fixed networks
- Efficiency, scalability and low signalling overhead
- Resilience and robustness (e.g. several points of attachment to the fixed network)
- Compatibility with inter-domain multicast routing
- Support of seamless moving of terminals among network types

2.2 Problems to solve

Trying to map the traditional IP multicast model into the concrete scenario of ad hoc network extensions, allows us to identify specific problems which need to be solved. According to the IP multicast model for IP multicast hosts, the process of taking part in multicast communications is quite straightforward. When they wish to send multicast traffic they simply use a class-D address as a destination and send the datagrams. When they are interested in receiving multicast traffic, they use the Internet Group Management Protocol (IGMP[6]) to inform their First Hop Multicast Router (FHMR) about the group they wish to join. This simple operation is not automatically supported in ad hoc networks due to some of the problems presented below.

TTL issues. IGMP uses IP datagrams with a time-to-live (TTL) of one hop for the communication between hosts and routers. Thus, by default, only directly connected hosts are able to join multicast groups since IGMP messages are unable to transit a multi-hop ad hoc network fringe.

Multihop nature of MANETs. Packets sent by sources which are more than one hop away will not automatically be received by the FHMR. However, intermediate ad hoc nodes must ensure that these packets reach the FHMR as it is required by most IP Multicast routing protocols (e.g. PIM-SM). The support of standard IP nodes is an issue that requires that ad hoc nodes incorporate capabilities for the interception and processing of IGMP messages since these are the means by which hosts join IP multicast groups in fixed networks. To date, none of the proposed multicast ad hoc routing protocols is able to handle such types of messages.

Flat addressing. An additional issue relates to the differences between the hierarchical addressing architecture which is used in fixed networks and the flat addressing architecture used in ad hoc networks. The problem is that multicast routers usually perform a process called an 'RPF-check' on every incoming packet. This process drops any packet which arrives at an interface which that router would not use to reach the source of the packet.

2.3 Proposed architecture.

There are several alternatives to achieve efficient network layer multicasting support between nodes within the ad hoc network extension and those in the access network. As we showed in [8], the most relevant are basically what we called a tunnel-based approach, and multicast ad hoc fringe. The former is based on the creation of a tunnel between receivers and the access routers. We have selected the multicast ad hoc fringe approach because, as we demonstrated in [8], it is much better in terms of scalability, simplicity and performance.

The key point in our proposed architecture is the idea of confining any new functionality to within the ad hoc fringe, challenging ad hoc nodes with the ability to process standard protocols (i.e. IGMP) to interact with non-ad hoc nodes. This mechanism exploits the anonymous nature of IP multicast because the FHMR does not need to know which node is interested in joining a particular multicast group, but only if there is any. So, when a standard-IP host generates an IGMP Report, internally ad hoc nodes will not need to transport that message. They use the MMARP protocol to create efficiently multicast paths within the ad hoc extension and any of the ad hoc nodes at a single hop from the FHMR will regenerate such an IGMP Report message. This shields the solution from the particular IP multicast routing protocol being used in the fixed network: we can interoperate in the same way with all of them just by sending IGMP Reports. So, our approach does not require any changes in standard IP nodes and routers. Mobile nodes will behave according to the standard IP Multicast model in which there is no requirement for senders and the only requirement for receivers is the use of the IGMP protocol to join multicast groups.

In addition, as the use of Standard-IP mechanisms (e.g. the ARP or IGMP protocols) within ad hoc networks is costly and usually offers limited performance, we propose a specific multicast ad hoc routing protocol called MMARP which incorporates particular path creation mechanisms to support standard-IP messages without an impairment in the protocol's performance. These specific MMARP extensions are described in the next section, whereas the proposed architecture is depicted in Fig. 1.

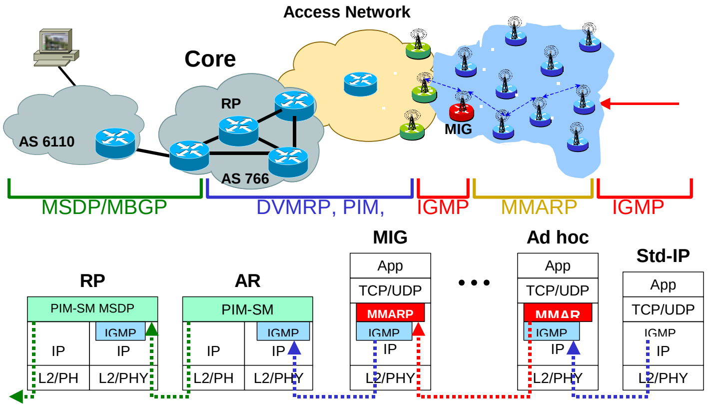

Fig. 1. Proposed multicast architecture

The AR and RP nodes in the figure represent standard multicast-enabled routers running PIM-SM. 'Ad hoc' represent pure ad hoc nodes and 'Std IP' represnts a standard IP multicast-enabled mobile host. The protocol providing efficient paths between the nodes within the ad hoc network fringe is the MMARP protocol presented below. From the point of view of the core network and the AR, the ad hoc fringe is seen just as another BMA subnet (i.e. group membership are being dynamically updated by IGMP Report messages received by the ARs).

3 The MMARP Protocol

The MMARP protocol is especially designed for mobile ad hoc networks (MANETs). It is fully compatible with the standard IP Multicast model and it allows standard IP nodes to take part in multicast communications without requiring any change because MMARP supports the IGMP protocol as a means to interoperate both with access routers and standard IP nodes. The interoperation with access routers is performed by the Multicast Internet Gateways (MIGs) which are the ad hoc nodes situated just one hop away from the fixed network. Every MMARP node may become a MIG at any time. The only difference between a MIG and a normal MMARP node is that the MIG is responsible for notifying the access routers about the groups memberships within the ad hoc fringe. The mechanism allows MMARP to work with any IP multicast routing protocol in the access network and, therefore, it shields the MMARP operation from the protocols performing the intra-domain or inter-domain multicast routing.

For the remaining text we use the terms standard IP source or standard IP receiver to refer to a traditional IP Multicast source or receiver and we use the term ad hoc sender or ad hoc receiver to refer to pure ad hoc nodes.

3.1 Overall operation

MMARP uses an hybrid approach to build a distribution mesh similar to the one used by ODMRP[5]. Routes among ad hoc nodes are established on-demand, whereas routes towards nodes in the fixed networks are maintained proactively. This offers a good trade-off between efficiency, smooth interworking with the fixed network while still having a good protection against link breakages (see Fig. 2). However, the way in which the mesh is created is different from ODMRP due to the special requirements which MMARP nodes have to face. For example, MMARP nodes can participate in the mesh creation process on behalf of standard IP nodes or even on behalf of the access router (AR). In addition, they have behave so that the standard IP multicast model can be preserved (i.e. making all the traffic generated within the ad hoc fringe to be delivered to the AR). These specific differences are explained in the next subsections.

The reactive part consists of a request phase and a reply phase. When an ad hoc node has new data to send, it periodically broadcasts a MMARP_SOURCE message which is flooded within the entire ad hoc network to update the state of intermediate nodes as well as the multicast routes. These messages have an identifier which allows intermediate nodes to detect duplicates and avoid unnecessary retransmissions. When such a message is received by an ad hoc node for the first time, it stores the IP address of the previous hop and rebroadcasts the packet. When one of these messages arrives at a receiver, or at a neighbour of a standard IP receiver, it broadcasts a MMARP_JOIN message including the IP address of the selected previous hop towards the source. When an ad hoc node detects its IP address in an MMARP_JOIN message, it recognises that it is in the path between a source and a destination. It then activates its MF_FLAG (Multicast Forwarder Flag) and rebroadcasts a MMARP_JOIN message containing its previously stored next hop towards the source. In this way, a shortest multicast path is created between the source and the destination. When there are different sources and receivers for the same group, the process results in the creation of a multicast distribution mesh like the one presented in Fig. 2).

The proactive part of the protocol is simply based on the periodic advertisement of the MIGs as default multicast gateways to the fixed network. As the TTL of IGMP messages is fixed at one, the reception of an IGMP Query can be used by ad hoc nodes to detect that they are MIGs and activate its MIG_FLAG. MIGs periodically broadcast a MMARP_DFL_ROUTE message which is flooded to the whole ad hoc network. The reception of this message informs intermediate nodes about the path towards multicast sources in the access network. When the MMARP_DFL_ROUTE message reaches a receiver or a neighbour of a receiver, this node initiates a joining process similar to the one that we have just

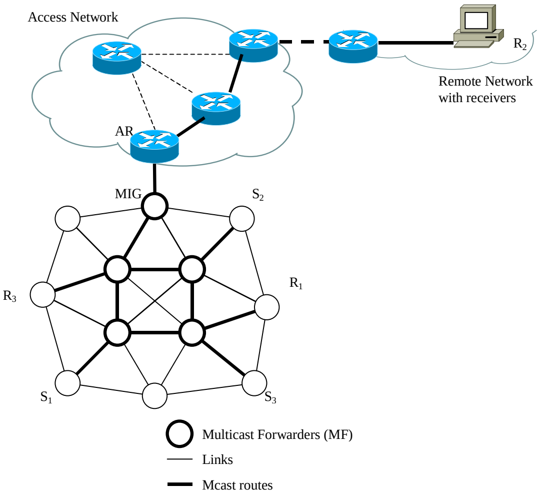

Fig. 2. Multicast mesh after request/reply phase

described for the reactive approach. When the MIG receives the MMARP_JOIN message, it then sends an IGMP Report towards the FHMR, ensuring the IP multicast data from sources in the fixed network reach the destinations within the ad hoc network extension.

The protocol incorporates local repairing mechanisms to overcome link breakages during the creation of the distribution mesh. Whenever a node is unable to deliver a MMARP_JOIN message to its next hop after four retries, it broadcasts a MMARP_NACK message to its one-hop neighbours. Upon the reception of this message, the neighbours use their own route to reach that next hop. Should any of them not know an alternate path, they repeat the process until a path is found. Although this recovery process does not offer optimal routes, it offers a quick recovery before the next topology refresh.

Once the mesh is established, the data forwarding is very simple: data packets addressed to a certain multicast group are only propagated by ad hoc nodes which have their MF_FLAG active for that group. When such a data packet arrives at a node whose MF_FLAG for that group has not expired, it checks that it is not a duplicate and in that case retransmits the packet. In any other case the packet is dropped.

3.2 Support of standard IP multicast protocols

The protocols used by standard IP nodes to perform their basic operation (such as ARP, or IGMP) were designed to operate in BMA (Broadcast Medium Access) networks. However, in multihop ad hoc networks, the link layer has a different semantics. The neighbours of a node are able to receive the frames it sends but it is not guaranteed that they are able to directly communicate among all of them.

In traditional ad hoc routing protocols without explicit support for standard IP nodes this is not a problem because each ad hoc node sends its own source announcement or join message. In order to be compatible with the standard IP multicast model, MMARP nodes in the neighbourhood of a standard IP node have to send MMARP_SOURCE or MMARP_JOIN messages on behalf of the standard IP node. This means that messages generated by standard IP nodes, may be received by all neighbours and processed independently, creating unnecessary paths.

The MMARP protocol has been designed to avoid unnecessary generation of these messages. It includes a field in its header which facilitates the identification of the node which actually triggered the sending of the control message; this allows ad hoc nodes to identify all the MMARP packets which are triggered by a specific standard IP node, independently of the ad hoc neighbour which actually generated it. Thus, ad hoc neighbours of standard IP nodes and intermediate ad hoc nodes are able to detect these types of MMARP_SOURCE and MMARP_JOIN messages as duplicate and avoid the creation of unnecessary paths.

4 Empirical Results

We have set up an indoor 802.11b multicast wired-to-wireless ad hoc network testbed to evaluate the feasibility of our MMARP-based seamless IP multicast approach for wireless ad hoc access networks. Our target is to evaluate the benefits of MMARP-driven infrastructureless ad hoc access networks when compared to traditional single-hop wireless IP multicast in a realistic scenario.

4.1 Testbed description

As it is shown in Fig. 3, the testbed consists of six x86-compatible PCs and a laptop. Different processor and memory configurations are used, since there are not any specific hardware requirements. In fact, all of these PCs are able to support the workload of the experiments. Three out of the six PCs are acting as MMARP-enabled nodes running Red Hat Linux 7.2 with the 2.4.17 kernel. They have a Lucent 802.11b pcmcia card as unique NIC. The nodes labelled as WR (Wireless Router) and WWR(Wired-to-Wireless Router) are PCs running FreeBSD 4.6 OS. The WWR node is equipped with two NICs, one of them being a Lucent WaveLan pcmcia card to provide coverage for the wireless area, while the other one is a 100 Mbps Ethernet NIC. The Wired Router (WR) contains two 100 Mb/s Ethernet NICs, one connected to the WWR and the other one to the rest of fixed networks. The Sender and Receiver nodes are both running Red Hat 7.2 with kernel 2.4.17. The sender is a x86-compatible desktop with a 100 Mb/s Ethernet card whereas the receiver is a laptop PC equipped with a Lucent 802.11b-compatible pcmcia wireless NIC. Wireless 2.422 GHz channel operating at the maximum capacity of 2 Mb/s has been used for the experiment. We have previously checked that this channel was not occupied by any other equipment. All the WaveLan NICs are operated in ad hoc mode.

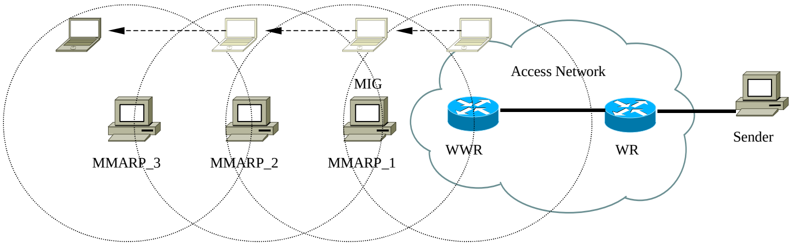

Fig. 3. Topology of the testbed

4.2 Description of the experiments

To assess the effectiveness of our proposal, we have performed two different tests: single-hop IP multicast and multihop ad hoc IP multicast. The former consists of the network depicted in Fig. 3, in which every MMARP node is switched off, so that there is a dedicated IEEE 802.11b wireless link between the receiver and the WWR. The wired part of the network is running the PIM-SM routing protocol to create the multicast path between the source and the WWR, which acts as an IGMP designated router forwarding multicast datagrams to the receiver when it joins the source.

The multihop tests are exactly the same regarding the wired part of the network. However, we deploy a self-organising ad hoc network extension with nodes running MMARP rather than a single-hop link between the receiver and the WWR. The receiver joins the multicast source in the fixed network through this multihop access network.

We use CBR traffic generator to measure the end-to-end bandwidth and packet delivery ratio. This application generates UDP packets with a payload of 900 bytes (i.e. 942 bytes including the IPv4 and UDP headers) which are then accounted at the destination. For each of the tests we have performed several measurements at increasing distances (7m, 15m, 24m, 30m, 42m) between the receiver and the WWR. At each distance, we have repeated the measurements using three different data rates of 100 packets/s (753.6 Kb/s), 50 packets/s (376.8 Kb/s) and 25 packets/s (188.4 Kb/s) respectively. The results of the different trials are described in the next section.

As expected, in our indoor scenario the performance depends not only on the distance but on the node's position as well. This is mainly due to random noise caused by traversing walls, obstacles, etc. The results are calculated as the mean values over quite a huge number of measurements per experiment. In addition, the measurements are performed in the same positions for each distance. So, random noise is expected to be nearly the same in all the trials, not affecting the validity of our experiments.

4.3 Experimental results

To be sure about the cause of the packet losses in our analysis, we empirically checked that there were not packet losses within the wired part of the network. So, all the packet losses perceived at the receiver will occur in the wireless part of the network.

As it is shown in Fig. 4, both approaches are able to deliver the transmitted bandwidth at short distances. For those cases the link-layer contention is low and the signal quality is good enough.

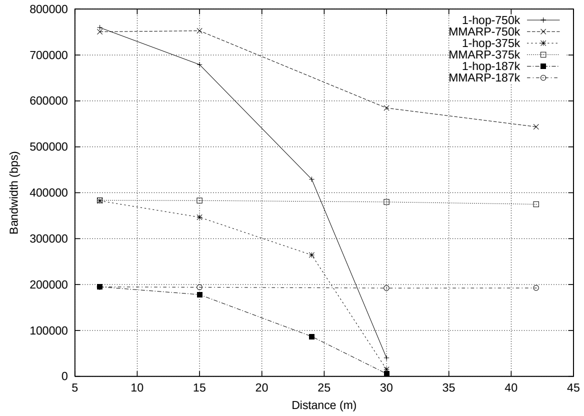

Fig. 4. Effective data rate achieved at increasing distance

In the single-hop trials, due to the degradation of the signal strength with increasing distance, the achieved bandwidth is lower as the distance increases. This is clearly assessed both for the achieved bandwidth and the packet delivery ratio in the '1-hop' cases of Fig. 4, and Fig. 5. These results are basically the expected behaviour as long as it is commonly known that (particularly in indoor scenarios) the signal strength usually decreases at a rate inversely proportional to d^2, d^3 and even in some cases d^4 in really bad indoor conditions. In our case, it is clearly shown that the bandwidth and packet delivery ratios rapidly drop to zero for distances around 30 m and beyond. As expected for the 1-hop dedicated link, given a fixed distance, the difference between the achieved bandwidth and the one being used at the source is bigger at higher data rates.

In the case of the multihop MMARP-based multicast ad hoc access network, it can be noticed that the performance at increasing distances degrades much slower than $1/d^2$. This is because the average distance in each of the intermediate hops is lower than in the single-hop trial. Thus, the mean signal strength is higher and the achieved bandwidth and packet delivery ratio are higher as well.

However, as Fig. 4 shows, MMARP only manages to achieve a 100% delivery ratio at distances at which only one or two of the MMARP nodes are needed.

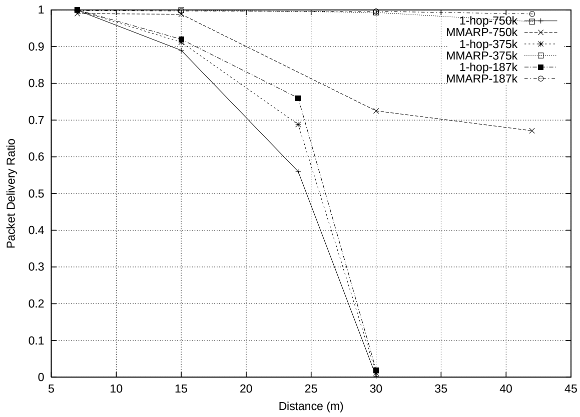

Fig. 5. Packet delivery ratio at increasing distance

At a distances higher to 30m some packet losses come up. These losses are mainly due to the well-known hidden terminal problem which happens among the nodes MMARP_1, MMARP_2 and MMARP_3. As long as IEEE 802.11b does not implement layer 2 acknowledgements of multicast frames (as it does for unicast traffic), each time a collision happens the packet is lost.

It is also particularly noticeable that the trial with the higher bandwidth in the multihop case performs much worse than the others. This is because, due to contention, the effective bandwidth, even in the ideal case, is lower than the 753,6 Kb/s generated by the source. When MMARP_1 receives a packet from WWR and forwards it to MMARP_2, the effective bandwidth is reduced to half of the original. One half of the channel is used for receiving the packet and the other half for sending it. When MMARP_2 sends the packet to MMARP_3, the effective bandwidth is further reduced to a third of the original (in optimal channel conditions). Leaving thus an effective bandwidth of 667 Kb/s (2/3 Mb/s) which is lower than the 753 Kb/s that the source is using.

However, in the trials without that bandwidth limitation the MMARP protocol has demonstrated to be able to deliver mostly 100% of the packets (even in non-optimal channel conditions) without an impairment in the overhead or the scalability of the protocol. The differences in the packet delivery ratio between these two multihop cases are mainly due to the hidden terminal problem. At higher data rates, the probability of two packets actually colliding is higher. However, as the figure shows, the performance has not been severely degraded for that reason. So, it is clear that the real limitation towards multicast ad hoc access networks is mostly the IEEE 802.11b MAC layer, which is known not to be very adequate for ad hoc networks. This demonstrates that, regarding the protocol's behaviour, having a higher number of nodes in the same radio link is not an issue. Only nodes with the MF_FLAG active will forward packets, and only the best of all those nodes would be selected as a forwarder.

5 Conclusions and Future Work

Currently there is not a real solution to seamlessly support efficient IP multicast communications in future heterogeneous wireless scenarios. We present our solution for ad hoc networks extending fixed IP access networks. It consists of a novel architecture and a new multicast ad hoc routing protocol called MMARP. This approach is the first to our knowledge being able to support seamless roaming from multicast nodes (including traditional IP multicast hosts) between traditional IP multicast networks and ad hoc network extensions. In the authors' opinion, in addition to the proposed solution, it is also an important contribution the demonstration through empirical experimentation that this kind of extensions driven by MMARP are able to easily extend IP multicast edge-coverage in a cost-effective way, without an impairment in the overall throughput. The results show that even at distances which the traditional single-hop approach is not able to cover, the multihop option offers more than a 98% packet delivery ratio.

For future work, we are working towards the analysis of the approach in hybrid ad hoc networks (e.g. mixed WLAN, Bluetooth scenarios), and with different layer 2 protocols to improve the performance of the IEEE 802.11b MAC layer.

Acknowledgements

This work has been partially funded by Spanish MCYT by means of the projects ISAIAS(TIC2000-0198-P4-04) and SAM(TIC2002-04531-C04-04).

References

1. IST-MIND Official Web site. [On-line] http://www.ist-mind.org/
2. Fenner, W.: Internet Group Management Protocol, Version 2. IETF Request For Comments, RFC 2236, November, 1997.
3. Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma P, Wei, L.: Protocol Independent Multicast Sparse Mode (PIM-SM): Protocol Specification. RFC 2362, June 1998.
4. Garcia-Luna-Aceves, J.J., Madruga, E.L.: The Core Assisted Mesh Protocol. IEEE JSAC, Vol 17, No. 8, August 1999, pp.1380–1394.
5. Lee, S.-J., Su, W., Gerla, M.: On-Demand Multicast Routing Protocol in Multihop Wireless Mobile Networks. ACM/Kluwer Mobile Networks and Applications, 2000.
6. Jetcheva, J.G., Johnson, D.B.: Adaptive Demand-Driven Multicast Routing in Multi-Hop Wireless Ad Hoc Networks. Proceedings of the 2001 ACM International Symposium on Mobile Ad Hoc Networking and Computing, ACM, Long Beach, CA, (October 2001): 33–44
7. H. Lei and C.E. Perkins. Ad Hoc Networking with Mobile IP. In Proceedings of the Second European Personal Mobile Communications Conference, October 1997, pp. 197–202.
8. Ruiz, P.-M., Brown, G., Groves, I.: Scalable Communications for Ad hoc Extensions connected to Mobile IP Networks. Proceedings of the (PIMRC'2002). Lisbon, September, 2002. Vol. 3, pp. 1053–1057.