MNet – A new multicast approach for the future Internet

Jorge Sá Silva, Sérgio Duarte, Edmundo Monteiro, Fernando Boavida
University of Coimbra, Department of Informatics Engineering / CISUC
Polo II - Pinhal de Marrocos, P-3030-290 Coimbra, PORTUGAL
Tel.: +351 239 790000 Fax.: +351 239 701266 Email: sasilva@dei.uc.pt

Abstract – Multicast communication in the Internet has deserved an increasing attention in the last few years. Nowadays, there are more and more applications that require communication systems with multipoint communication capabilities. Multicast communication reduces both the time it takes to send data to a large set of receivers and the amount of network resources required to deliver such data. This is why Internet Service Providers (ISPs) have a strong need for multicast solutions from vendors of networking products. The question now is how long it will be before multicast becomes a true Internet Service.

This paper presents a new multicast proposal that is being developed at Laboratory of Communication and Telematics of the Informatics Engineering Department of the University of Coimbra. The approach explores the use of Multicast Servers in conjunction with the SSM protocol, addressing some of the SSM limitations, of which the support of dynamic sources and the support of QoS heterogeneity are the most important. After presentation and discussion, evaluation results are analysed, showing that the proposed model is beneficial when compared with the original SSM solution.

I. INTRODUCTION

There is a growing demand for one-to-many multimedia content delivery. Although IP multicast has been available through the experimental MBONE for a number of years, it is just beginning to see support from current ISPs. The use of a set of point-to-point connections to emulate multipoint communication (Figure 1) is still used although it has high cost in terms of network and communication server resources. This may be acceptable in a small number of cases, but increasing demand prevents this type of solution from becoming widespread due to its inefficiency.

This paper presents a project which aims at proposing and studying new multicast management solutions, based on the Source Specific Multicast (SSM) protocol [1] and on Multicast Servers.

The remainder of this paper is organized as follows. Section 2 describes the state-of-the-art of multicast communication and presents the main protocols for its support. Also, in this section, the present and future requirements of the multicast technology are enumerated. The MNet project is described in section 3. Its main purpose is presented along with a description of the MNet platform where some scenarios are being implemented. Section 4 illustrates the application of the MNet multicast model to real environments in order to overcome some multicast limitations. This section also presents some evaluation studies. Conclusions and topics for further work are described in the last section.

II. IP MULTICAST

Although IP broadcast services have been available for some time, currently there are few efficient operational solutions for IP multicast environments. This is manly due to the complexity in the management of multicast communications.

Clearly, there is the need to find solutions to cope with the increasing number of multipoint applications for multimedia content delivery like videoconference, distributed simulation, network games or Internet TV.

Wireless technology offers a natural medium for multicast, where the addition of new members to a multicast group does not involve new costs by the introduction of additional hardware. On the other hand, for wired communications the same process is more complex. In these systems it is necessary to choose a correct developing plan to support bandwidth saving and optimisation of network resources.

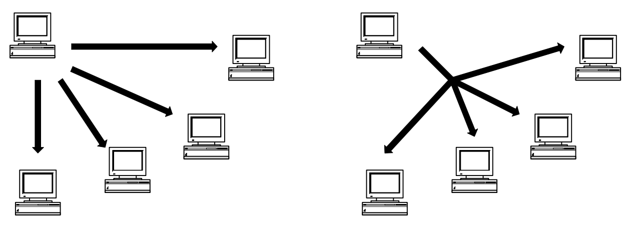

Figure 1 – Set of point-to.point connections vs Multipoint connection

As mentioned before, the use of a set of point-to-point connections to emulate a multipoint system (Figure 1) is

0-7803-7661-7/03/$17.00©2003 IEEE

highly inefficient. Additionally, it proves difficult to implement in large networks: terminal workstations would need to store large amounts of information pertaining to all connections, and this information would have to be updated and maintained in a consistent manner. Moreover, the use of point-to-point connections would result in a large number of control messages inside the network.

In a local area network, due to the usually inherent broadcast capabilities of Ethernet, it is straightforward to implement a multicast system. If, however, there is the need that the multicast group extends to other networks it is necessary to use a multicast routing protocol.

Forwarding of multicast packets is subject to certain risks. If there are several routers on the same physical network and if special care is not taken, they may all relay the packets again and again. In this way, there is the risk of creating not only a multicast loop but also a multicast avalanche, bringing the whole network to a stop as it is quickly filled to capacity. The whole purpose of multicast routing is, precisely, to achieve delivery of multicast packets without loops and without excess transmissions.

Multicast addresses in IPv4 use the class D address space (from 224.0.0.0 to 239.255.255.255). In Internet Protocol version 6 (IPv6) the multicast address space is defined in the range of FF::/8. Each station, in addition to its own unicast address, has a list of all multicast groups that it belongs to. It is also necessary that each router uses the Internet Group Management Protocol (IGMP) [2] to find out if the network has group members for a given multicast address, and a routing protocol for transporting multicast information between different networks and domains.

There are different routing protocols, some using rudimentary techniques such as flooding, and others using more elaborated techniques that rely on source-based trees or shared-based trees algorithms. Work in the multicast area started by developing and refining intra-domain routing protocols. Later, particular emphasis was placed on developing inter-domain multicast routing protocols.

There are two types of multicast routing protocols: Sparse-Mode and Dense-Mode. While the first type uses shared trees with a source rendezvous point, the latter uses algorithms that construct source-based trees. Dense-mode protocols are well suited to Local Area Networks (LANs), while sparse-mode protocols are used when multicast members are disperse over some large set of networks.

The Distance Vector Multicast Routing Protocol (DVMRP) is a distance vector routing protocol, a multicast version of the Open Short Path First (OSPF) protocol. The forwarding process uses the reverse-path-forwarding algorithm.

The Multicast extensions to OSPF (MOSPF) protocol uses a network map constructed using a link state database.

Protocol-Independent Multicast (PIM) is one of the most popular multicast routing protocols. It works both as a Sparse or a Dense protocol. The design and implementation of IGMPv2 usually predate the PIM routing protocol. However, it is possible to imagine a host that ignores IGMP and speaks PIM directly with routers on the local sub-network.

For inter-domain use, PIM requires the Multicast Source Discovery Protocol (MSDP) in order for the group-to-Rendezvous-Point mapping to be advertised to all hosts on the Internet. MSDP runs over the Multicast Border Gateway Protocol (MBGP), which is a set of multicast extensions to support the advertisement of reachability information for multicast routes. This allows an autonomous system to support incongruent unicast and multicast routing topologies, and thus implement separate routing policies for each.

The current solution for IP multicast – based on IGMP to announce hosts interested in receiving multicast information and on PIM-SM /MBGP/MSDP – is too complex and does not solve some of the current needs. Previous solutions like Simple Multicast or Express Multicast [3] were rejected on the premises that they did not solve all IP multicast problems.

There are a number of transport layer protocols that try to improve and assist multicast routing protocols. However, given the limitations of current multicast protocols, it is clear that the solution is not to complement them with a complex set of higher layer protocols.

The SSM draft [1] was finished in 2000 as an interim solution. This protocol supports source-based multicast trees across multiple domains in the Internet. According to the SSM specification, when a host decides to join a multicast group it must specify not only the multicast address but also the source address(es) that it accepts to receive the multicast information from, overcoming the usually any-to-any model. End-hosts use version 3 of the IGMP protocol for IPv4 or version 2 of the MLD protocol for IPv6.

The address range 232/8 has been allocated by IANA [4] for deploying source-specific IPv4 multicast services. For IPv6 the address range is defined in [5].

SSM avoids the flooding approach of dense mode protocols, and the core/rendezvous approach used by sparse mode protocols. When the multicast address and the source address(es) are specified, the join message is routed toward the source, using the a reverse shortest path tree.

The use of MSDP is precluded, as there is no need to run it between domain sources. As described before, the joining host must know beforehand the source addresses of the multicast information that it accepts to receive information from.

III. MNET – A NEW APPROACH FOR MULTICAST SYSTEMS

Although SSM will solve a large number of current multicast drawbacks, it also has some limitations. There is the need for new algorithms that address management, billing, addressing, quality of service and traffic control issues. The use of a large number of transport-layer protocols to complement the deficiencies of IP is not a good solution. Also, as IP is the protocol that supports the universal integration of systems, it is important that it offers the mechanisms that are necessary to efficiently deploy multicasting.

For example, in the case of the Session Description Protocol (SDP) [6] the extensions proposed in this paper are extremely important. SDP is used extensively for multimedia conferencing, for describing multimedia sessions for the purposes of session announcement, session invitation, and other forms of multimedia session initiation. SDP is being used on the Internet multicast backbone (Mbone) and uses IP multicast to support an efficient many-to-many communication.

Numerous extensions have been developed to circumvent SDP's shortcomings. A successor protocol – SDPng – was developed to overcome the current limitations. It is important that IP offers a set of new properties that support the necessary requirements of important protocols such SDP or SDPng.

The MNet project's main purpose is the study of the use of Multicast Servers (MSs) in SSM environments. When a host wants to join a multicast group, the host must identify not the source address(es) but one or a set of MSs. So, the Reverse Shortest Path is constructed from the receiver node to the MS.

MSs are responsible for agglutinating the traffic from the different authorised sources and for routing it to the receivers (Figure 2).

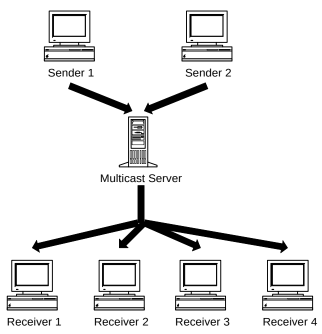

Figure 2 – Multicast Servers

Currently, all the routing algorithms that build multicast trees use only one metric – the distance. The introduction of Quality of Service (QoS) parameters requires the use of other metrics and the support of non-symmetrical links. In Differentiated Service (DiffServ) environments, routing is supported by feedback mechanisms to inform upstream nodes of the downstream properties. This process is even more complex in multicast systems, as different downstream nodes may have different properties.

The SSM protocol, *per se*, does not allow for the QoS heterogeneity of receivers, although it simplifies the path construction from source to receivers. The model proposed by the MNet project – the use of multicast servers – easily supports the use of different levels of QoS, in a DiffServ environment.

Currently, there are three approaches to the support of multicast heterogeneity:

- the source uses a single multicast stream for all the receivers, although they have different needs and properties;
- the source transmits versions of the same information, encoded with different degrees of QoS; each receiver joins the group that carries the desired QoS;
- the source uses a layered encoding approach: a basic layer that provides the basic level of QoS and additional levels, each improving the final signal; depending on their needs, receivers join the appropriate multicast groups.

The MNet approach has the potential to offer QoS heterogeneity in contrast with the first approach, overcomes the inefficient use of network resources and the installation of complex modules in each source that characterises the second approach, and provides a simpler solution when compared with the third approach. In the MNet model, each source sends the information to the MS, where the original signal gives rise to different versions according to the desired levels of QoS. Although the proposed solution is similar to the second approach, it frees the sources of implementing the desired levels of QoS. Also, the MNet solution offers a set of properties that will be described next.

Using a centralised process, MNet leads to non-static systems, supporting dynamic QoS needs of each receiver. This property is especially suited for the future Internet, where the bandwidth occupation in each link will be highly dynamic.

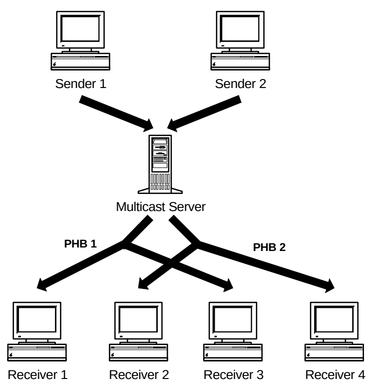

Figure 3 – DiffServ environments using Multicast Servers

The creation of different source trees from a given MS according to desired Per-Hop-Behaviours (PHB) is a simple process (Figure 3). It is only necessary that the MS manage different multicast streams, using different IP multicast addresses.

In an SSM environment, when a host decides to receive a multicast stream, it sends an IGMPv3 message identifying, not only the multicast address, but also a (or a set of) source address(es). The method by which receivers identify source address(es) is still an open issue. In [1] this task is left to the upper layers. As a future work the author proposes the use of HTML. In MNet environments, the upper layers of the system can take care of the identification of MS addresses. This respects the original SSM specification because, by transferring source-identifying mechanisms to the upper layers, flooding techniques are avoided.

At Coimbra University, we are implementing a real testbed based in SSM (Figure 4). The testbed is composed of 5 PCs, one of them acting as a router and the others as terminals. The network platform is built with FastEthernet technology.

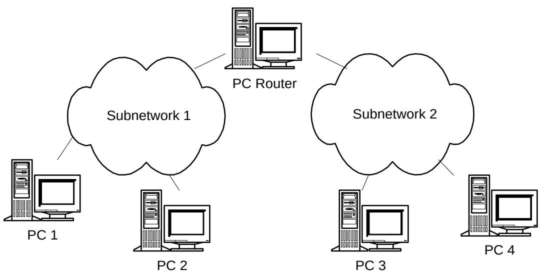

Figure 4 – MNet testbed

We use an IGMPv3 patch for Linux operating system available from Sprint Labs [7]. Also, we use the new multicast API presented in [8]. The changes needed in PIM-SM to support the SSM model were made according to [1] and [9]. Basically, the PIM-SM router needs to be able to skip the shared tree and the Rendevous-Point (RP) and go straight to source specific trees.

PIM-SM already supports source-based trees as, originally, it supports the two types of trees. However, routers in a PIM-SM environment cannot choose between a shared tree and a source-based tree. Receivers always join a PIM shared tree, and may later be switched to a per-source tree by its adjacent edge router. So, it is necessary to start with a source-specific tree.

The connection establishment between each source node and the MS respects the source tree described in the SSM specification [1]. Using the IGMPv3 functionalities, the MS receives the multicast information from only the authorised sources, previously inserted in the local database.

The connection establishment between the MS and each receiver must also respect the SSM specification. When the receiver joins the multicast group it must identify the MS address.

The use of different multicast addresses between the sources and the MS, and between the MS and the receivers is also necessary to prevent that each source sends the information directly to the receivers. The MS, acting as an intermediate node where management tasks can be implemented, must control all the authorised elements and apply the necessary mapping functions.

IV. APPLICABILITY AND EVALUATION STUDIES

The number of Internet applications that require multipoint support is increasing. Internet TV, videoconferencing and distributed simulation are a few examples of applications that require this technology.

According to the evaluation framework presented in [10], the criteria for evaluating a multicast approach include:

- How expensive is the deployment of the approach? Does it require the cooperation of the local network administrator to operate?
- Is it compatible with best current practices of routing IP multicast in the Internet core?
- Which of the Multicast Service Model [11] Core and Desired operations are supported?
- Does it ensure that only one copy of a multicast packet need traverse a network link?
- Does it allow for both multicast transmission and reception?

It is important to distinguish between the original SSM model and the extensions proposed in this paper. With respect to the first criterion, the implementation of an SSM environment requires some changes in the routers, as described earlier, based on [1] and [9]. Additionally, it is necessary to install IGMPv3. The SSM model is compatible with current IP routing, and the Multicast Service Model described in [11] is fully supported. As a network layer protocol, SSM supports multicast transmission and reception, and it offers an efficient use of network resources overcoming the duplication of messages in the links.

The extensions proposed in this paper do not change the evaluation presented above. Also, all of the proposed extensions are transparent to the original SSM model and offer an important set of properties that solve most of the SSM drawbacks.

One of the most important advantages of the MNet extensions is the support of dynamic sources. With the MNet approach, the addition of new sources does not require the building of a new point-to-multipoint tree from the new element to the receivers, and overcomes the complex use of notification protocols of which the Session Announcement Protocol (SAP) [12] is an example. Notification protocols are used to assist the advertisement of multimedia multicast conferences and other multicast sessions.

In the MNet model, when a new authorised source wants to start sending information to the multicast group it needs only to inform the MS (using a unicast address), in contrast with the original SSM model where all the receivers need to be informed. In MNet approach, before the sending of the IGMPv3 message with the set of sources actualised, which can be done periodically, the new source can use the same MS unicast address to transmit the multicast information to the MS.

The use of encrypted announcements is not recommended in SAP environments, as several receivers may not be able to decrypt those announcements. Also, if announcements are being performed via a proxy, then there is the possibility that the proxy is not immediately aware of superseded announcements, which may lead to the relaying of stale announcements.

In SAP environments, it is also necessary to use authentication announcements in order to verify that changes to a session description or removal of a session are allowed, and in order to authenticate the session creator.

Protocols like SSM or the one proposed in [12] do not support the specific addition of selected receivers. This drawback is a limitation to the implementation of some billing algorithms. With the presence of a MS, some encryption techniques and keys, these algorithms can be implemented.

Adaptive applications, that reduce the traffic sent to the network when some congestion problems arise, can offer interesting results but are not well suited to multicast systems. They require feedback information from the network to adapt the traffic to resource availability. In multipoint communications this process is more complex, as it is necessary to construct a considerable number of inverse multicast trees. This mechanism is even more difficult when the system supports dynamic sources – sources that come and go over the duration of a session.

The inherent heterogeneity of the Internet poses several challenges. As most of the multicast applications are QoS-sensitive in nature, it is important that new multicast models support QoS. The MNet approach offers a solution to the problem, as it supports different PHBs in a DiffServ environment, as described above. And, when compared with the original SSM model it does not require the installation of complex modules in each source.

As there is no method of uniquely allocating addresses in the current multicast model, and as the current multicast address space is unregulated, nothing prevents applications from sending data to any multicast address. Members of two sessions will receive each other's data if separate addresses are not chosen. Although in the SSM original model each source is responsible for resolving address collisions, this is not enough and some algorithms like the Multicast Scope Zone Announcement Protocol [13] have been proposed for discovering the multicast administrative scope zones that are relevant at a particular location. The MNet model can contribute to the solution of this problem, as MSs are centralised points that manage the multicast membership.

The proposed model also offers the ability to implement some billing mechanisms. As MSs manage the existing connections, it is easy to implement the registration process and to control which node is authorised to send or receive data.

The validation of members in MNet environments prevents flooding attacks, in which high-rate useless data are transmitted to a specific multicast group.

It is important to compare the original SSM approach with the model being studied in the MNet project. Although the SSM protocol is in a birth phase it is important to anticipate its limitations and to immediately propose possible solutions.

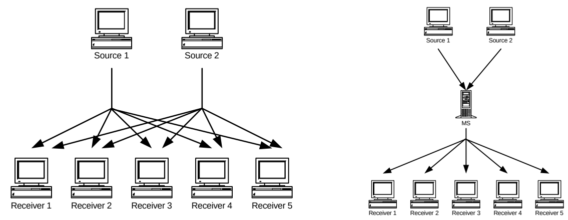

SSM model MNet model

Figure 5 – SMM model vs. MNet model

Figure 5 compares the original SSM model with the MNet model.

Potentially, there are two disadvantages of using the MNet approach:

- data throughput and end-to-end latency may be adversely affected due to the additional level of indirection introduced by MSs;
- MSs can potentially become a bottleneck and a central point of failure.

The ns-2 simulator [14] was selected to evaluate these potential problems.

The network used in our simulation was composed of 8 intermediate nodes (n_i), 3 senders (s_j) and 4 receivers (r_k), as described in Figure 6.

The links had a bandwidth of 5 Mbps with a mean delay of 2 ms per packet. We used exponential on/off applications with a packet size of 48 bytes, a burst time (the average "on" time for the generator) of 100 ms, an idle time (the average "off" time for the generator) of 100 ms, and a rate (during "on" times) of 64 Kbps. The nodes used simple drop tail policies in their queues.

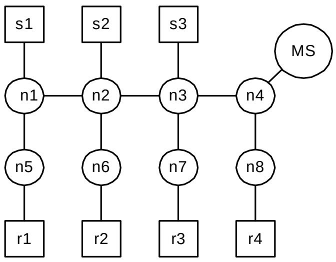

Figure 6 – Network scenario used in ns-2

The next figure compares the number of packets managed by node 4 in 50 ms intervals in the SSM model and in the MNet model.

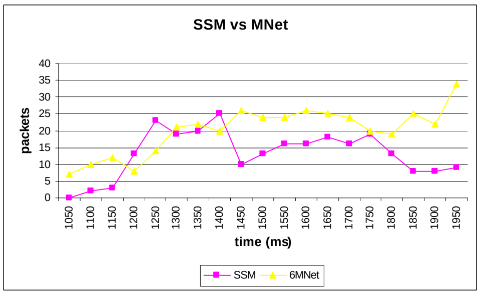

Figure 7 – SSM and MNet comparison using exponential on/off applications

In the case of the SSM model, the multicast packets were sent directly from the senders to the receivers. In the MNet model the packets were routed first to the MS. In this latter model the MS was located in node 4.

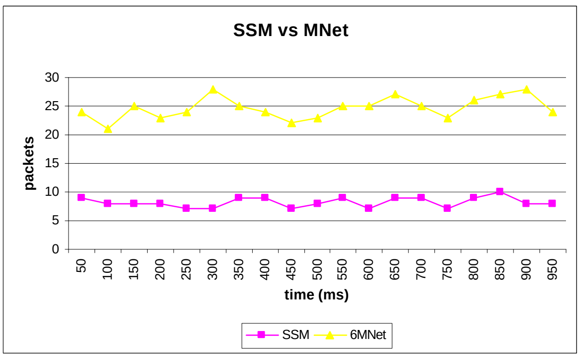

Figure 8 - SSM and MNet comparison using CBR applications

To evaluate more deeply the congestion degree in the MS we also used the same scenario but with 64 Kbps Continuous Bit Rate (CBR) applications. We also introduced some background traffic in order to produce some probabilistic behaviour. The next figure compares the two models.

As both studies show, in the simulation scenario, the overhead introduced by the MS is not significant. Nevertheless, in the case of the necessity of reducing this overhead, the MNet model also proposes the use of multiple MSs. With this approach it is possible to offer not only load sharing and fault-tolerance mechanisms, but also, depending on an optimised choice of the MSs position, to overcome the data throughput and the end-to-end latency.

As it is necessary to guarantee that each destination receives all the multicast information, there are two different approaches:

1) Receivers sharing (figure 9a) – where the sources send the information to all MSs, using a single multicast address. In this scenario each destination must receive the multicast information from only one MS. To support this, it is necessary to identify each MS with a distinct multicast address.
2) Sources sharing (figure 9b) – where all the MSs send the multicast information to all of the receivers (identified by a single multicast address). On the other hand, each source must send the information to a unique MS.

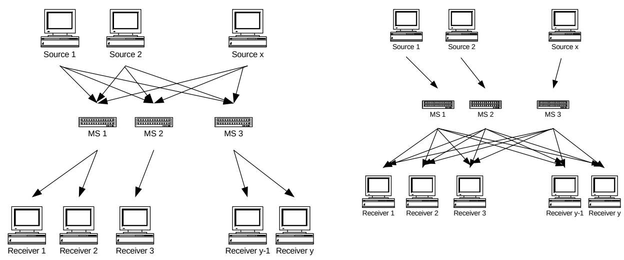

a) Receivers sharing b) Sources sharing

Figure 9 – Multiple MSs

As the "*path length*" (number of hops of the path) increases, the sub-network bandwidth and used node resources also increase. Due to the multicast protocols advantages, a sub-network supporting a large number of hops does not necessarily represent a sub-network with heavy load traffic. However, for a fixed number of MSs, as the number of hops present in the sub-network increases the probability to find larger distances between terminal nodes (source or receiver) and an MS also increases.

The 'sources sharing' approach presents a set of disadvantages. Usually, as the number of receivers is larger than the number of existing sources, the consumption of network resources is accentuated. The next table compares the number of paths in the MNet approach, using the 'receivers sharing' and the 'sources sharing' approaches. The table compares two scenarios: when the number of multicast sources is one, and when all the receivers are also sources of the multicast group.

Table 1 - Quantitative comparison of 'receivers sharing' and 'source sharing' models

	# sources = 1			# sources = # receivers		
	SSM	receiv. sharing	sources sharing	SSM	receiv. sharing	Sources sharing
outS	nR	nMS	1	nR	nMS	1
inMS	-	1	≤ 1	-	nR=nS	≤ nR=nS
outMS	-	≤ nR	nR	-	≤ nR=nS	nR=nS
inR	nS	1	nMS	nS	1	nMS

outS	-	number of output paths per source
inMS	-	number of input paths per MS
outMS	-	number of output paths per MS
inR	-	number of input paths per receiver
nMS	-	number of MSs
nR	-	number of receivers
nS	-	number of sources

As the table shows, the main difference between the 'receivers sharing' approach and the 'sources sharing' approach is a large number of "paths" between the sources and the MSs, or between the MSs and the receivers, respectively.

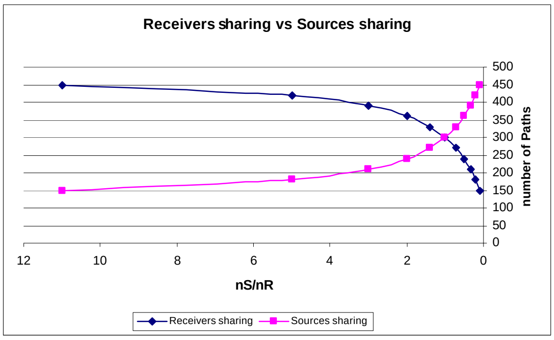

Figure 10 – Number of paths in a subnetwork as a function of nS/nR (case study 1)

As figure 10 presents, as nS << nR the 'receivers sharing' approach offers a reduced consumption of network resources, as expected.

In this case study we used 4 MSs. The initial scenario used 110 sources and 10 receivers, and the final scenario used 10 sources and 110 receivers.

Figure 11 presents a similar case study, the difference being that the sources and the receivers varied from 50 to 1000 and the number of MSs was 8.

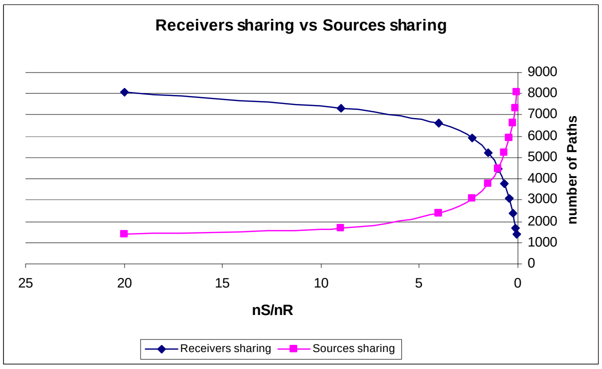

Figure 11 - Number of paths in a subnetwork as a function of nS/nR (case study 2)

V. CONCLUSIONS AND FUTURE WORK

This paper presented a proposal for the use of multicast servers in conjunction with the SSM protocol. The authors believe in the advantages of multicast and in the success of SSM. Nevertheless, SSM still needs additional work in order to overcome some of its limitations. The work described in this paper, carried out in the scope of the MNet project that is being developed at Coimbra University, tries to address some of these needs, of which the support of dynamic sources and the support of QoS heterogeneity are the most important. Other needs, such as authentication, security, management, billing, addressing, and traffic control, can also benefit from the multicast model proposed in the MNet project.

The proposal was presented and discussed in sections III and IV. In addition, the proposal was evaluated by simulation, in order to assess its performance with respect to the original SSM specification.

Usually, the applications that are driving multicast are one-to-many and few-to-few. Up to now, many-to-many applications have not gained popularity nor have they received much attention. The MNet model can easily accommodate these types of applications, anticipating future needs.

As future work, we are also very interested in developing new billing solutions. This is especially relevant if multicast is to be used as a money-making enterprise for commercial companies. We also plan to extend our study to IPv6 environments. IPv6 was designed to overcome the limitations of the current version of IP – IPv4. One of the main reasons for the introduction of IPv6 was the foreseeable exhaustion of the IPv4 address space. The delay observed in the migration to IPv6 is caused by the use of palliatives such as Network Address Translation or Classless Inter-Domain Routing, but these have only postponed the inevitable address exhaustion.

Although the IETF is developing some methods for allocating addresses like (GLOP) [15] or Multicast Address Allocation Architecture (MAAA) [16], currently, a multicast application can pick whatever class D address it chooses. Address collision is, thus, an important issue, as it is possible that two sessions will receive each other's data if distinct addresses are not chosen. IPv6 presents a larger set of multicast addresses, and this, in conjunction with solutions such as the MNet model, can reduce chance of collisions to near zero.

ACKNOWLEDGMENT

The work presented in this paper was partially financed by POSI - Programa Operacional Sociedade de Informação of Portuguese Foundation for Science and Technology and European Union - FEDER.

REFERENCES

[1] Holbrook, H. and Cain B.
"Source-Specific Multicast for IP", Internet Draft 2000

[2] Fenner, W.
"Internet Group Management Protocol, Version 2",
RFC 2236, November 1997

[3] Holbrook, H. and Cheriton, D.
"IP Multicast Channels: Express Support for Large-scales Single Source Applications", Proc. of SIGCOMM 1999

[4] Internet Assigned Numbers Authority
http://www.isi.edu/in-notes/iana/assignements

[5] Haberman, B. And Thaler, D.
"Unicast Prefix Based IPv6 Multicast Addresses", Internet Draft 2001

[6] Handley, M., and Jacobson, V.
"SDP: Session Description Protocol", RFC 2327, April 1998

[7]http://www.sprintlabs.com/Department/IP-Interworking/multicast/linux-igmpv3/

[8] Thaler, D., Fenner B. and Cain B.
"Socket Interface Extension for Multicast Source Filters", Internet Draft 2001

[9] Fenner, B. et all
"PIM–SM: Protocol Specification", Internet Draft 2001

[10] Nickless, B.
"A Survey of IPv4 Multicast Service Implementations", Internet Draft 2002

[11] Nickless, B.
"IP Multicast Service Model", Internet Draft 2002

[12] Handley, M; Perkins, C. and Whelan, E.
"Session Announcement Protocol", RFC 2974, October 2000

[13] Handley, M., Thaler, D., and Kermod, R.
"Multicast-Scope Zone Announcement Protocol",
RFC 2776, February 2000

[14] The VINT project
"NS-2 – The ns Manual", February 2002

[15] Meyer, D. and Lothberg, P.
"GLOP Addressing in 223/8", RFC 2770, 1997

[16] Handley, M., Thaler, D. and Estrin, D.
"The Internet Multicast Address Allocation Architecture",
Internet Draft - 1997