

Policy Aware QoS Inter-domain Multicast Routing

António Costa*, Maria João Nicolau†, Alexandre Santos* and Vasco Freitas*

*Departamento de Informática,
Universidade do Minho, Campus de Gualtar,
4710 Braga, Portugal

Email: {costa,alex,vf}@uminho.pt

†Departamento de Sistemas de Informação,
Universidade do Minho, Campus de Azurém,
4800 Guimarães, Portugal
Email: joao@uminho.pt

Abstract—Future Internet applications are expected to have higher QoS requirements. Therefore, routing protocols must be adapted to find a path which satisfies such requirements. Most multicast applications are QoS sensitive in nature, because they involve the transmission of real-time multimedia data streams.

There are several research projects aiming to extend existing multicast routing protocols with QoS capabilities, or even proposing new QoS aware ones. However key requirements for inter-domain routing like scalability, intra-domain independence, and policy awareness are missing in most of existing multicast routing protocols.

In this paper, it is proposed an inter-domain multicast routing strategy that builds inter-domain unidirectional multicast distribution trees, taking into account multicast specific routing policies, and supporting QoS requirements.

I. INTRODUCTION

As well as for unicast routing, the multicast routing can be treated at two different levels: intra-domain or inter-domain level. The structure of the Internet can be decomposed into end-systems, local networks and finally administrative domains (also called autonomous systems). Different administrative domains are interconnected with routers that are called border routers. The network composed by the highest level of the hierarchy is called the inter-domain level. In this paper we propose a multicast routing strategy which tries to address some of the problems faced by inter-domain multicast routing.

An important issue related with inter-domain routing is the fact that the topology must be considered as an asymmetric one. This is because inter-domain routing policies might be such that two distant domains might not agree on the same transit domain.

Traditional multicast routing protocols use RPF (Reverse Path Forwarding) concept in constructing multicast trees. This concept is based on the idea that an actual delivery path to a node is the reverse of the path from this node to the source. This concept fits well in symmetric environments, but in a routing environment where routing policies are applied, the guarantee that symmetrical path will exist between two network addresses is broken. Therefore reverse-path routing may not be used.

This is also true when considering QoS routing because asymmetries for a specific service quality may result from

the available network resources. Routing in the Internet has so far been based on a best-effort service model, primarily concerned with connectivity. Packets are delivered using a route based on source and destination addresses and typically a single metric is taken into account to make route decisions. Thus routing protocols build routing tables having the goal of minimizing the cost of each path. However, this model is not adequate to satisfy the growing demands of the next generation applications, most of which demand QoS assurances. In order to support a wide range of QoS requirements, routing protocols need to have a more complex model where the network is characterized with multiple metrics, such as bandwidth, delay and loss probability. The basic problem of QoS routing is then to find a path that satisfies multiple constraints.

In this paper it is proposed a multicast routing strategy that aims to contribute with a solution to the above mentioned problems. A new model is proposed for QoS and policy aware inter-domain multicast routing, taken into account link asymmetry.

II. INTER-DOMAIN MULTICAST ROUTING

Currently there are two inter-domain multicast routing frameworks: one based on MBGP/PIM-SM/MSDP protocol suite[1], already in use, and another one based on MASC/BGMP[2], proposed as a near-future solution. In both frameworks, MBGP[3] plays an important role in distributing multicast specific information between neighbor border routers.

In the MBGP/PIM-SM/MSDP framework, inter-domain multicast trees are constructed using PIM-SM (Protocol Independent Multicast - Sparse Mode)[4], also commonly used at intra-domain level. PIM-SM constructs shared and source based unidirectional trees, by using explicit join messages directed to the tree root unicast address. At inter-domain level, join messages are forwarded by border routers using a multicast specific routing information base called M-RIB. Like the standard unicast RIB, M-RIB also contains unicast prefixes and their attributes, but is only intended for multicast usage. This new M-RIB allows a domain to announce the same unicast prefix with different attributes for multicast and unicast purposes. The usage of PIM-SM at inter-domain level,

demands for a new protocol called MSDP (Multicast Source Discovery Protocol)[5]. MSDP is used by border routers to discover active sources in other domains. Each domain Rendez-Vous Point (RP) router announces the available sources inside its domain to other RP border routers

In the MASC/MBGP framework, inter-domain multicast trees are created by BGMP (Border Gateway Multicast Protocol)[6]. BGMP constructs a bidirectional shared tree among border routers, rooted at a root domain. This approach is pointed as more efficient for the inter-domain scenario. The root domain of a group is known in advance, as a result of the hierarchical multicast address allocation strategy proposed in MASC (Multicast Address Set Claim)[7]. With MASC, each domain must previously allocate multicast address blocks, and claim their usage to all other domains in order to avoid address usage conflicts. Successfully claimed multicast address prefixes are then communicated to the MAAS domain server for intra-domain individual distribution, and also injected to other domains through MBGP, as group routes. Group routes are stored in a new multicast specific information base, called G-RIB, kept at border routers. Unlike the other RIBs, G-RIB contains multicast prefixes and their AS Path Attributes, and not unicast ones. They are used by border routers to forward the BGMP control messages when constructing the bidirectional tree. As in other RIBs maintained through MBGP, there may exist multiple entries for the same prefix as a result of different possible paths between domains.

III. RELATED WORK

There are few proposals for inter-domain constrained multicast routing. YAM[8] and QoSMIC[9] aim to discover multiple paths from an existing tree onto a joining node and afterward make a choice based on certain criteria.

YAM builds shared trees that have the capability to provide multiple routes to connect a new node onto an existing tree. It handles dynamic membership and does not require any global network state at routers, but it has excessive communication overhead because it relies on flooding to find a feasible tree branch to connect a new member.

QoSMIC alleviates the flooding behavior, but introduces a new complex element: the Manager Router. QoSMIC uses two different procedures to find a feasible tree: a local search and a multicast tree search. Local search is initiated by the new member router by flooding BID-REQ messages to its neighborhood, with scope controlled by TTL (Time To Live). Any in-tree router that receives a BID-REQ message becomes a candidate router and replies with a BID message, which is unicast to the new router. The BID message collects information about the path on its way that can be used for selection purposes. The multicast tree search occurs at the same time, initiated by a Manager after receiving a M-JOIN request from the new receiver. The Manager sends a BID-ORDER message to a set of in-tree routers, that become candidate routers and reply with BID messages exactly as described for local search procedure.

The Policy Tree Multicast Routing (PTMR) Protocol[10] is a single layer protocol that extends PIM-SM in order to accomplish policy routing. PTMR architecture is characterized by a structure called Policy Tree. Group members receive source location information through PIM RP shared tree, following the PIM-SM routing and forwarding mechanism. If the receiver would like to switch to a source-based tree, it does so according to the PIM-SM mechanism. But, if the source is not located in the same domain, a Policy Tree is constructed by PTMR following the next sequence of actions: 1) the border router of the receiver's domain sends a request message directed to the source; 2) the first-hop router for the source will establish a policy route from it to the administrative domain whose border router issued the request message; 3) the answer to the request message (*Mark Message*) may enter at the receiver's domain through other border router than the one that has sent the request message; 4) in that case, the border router who has received the *Mark Message*, multicasts an announce message through the initial border router with the address of the new border router in the multicast tree; 5) with this information the routers close to receivers join this new border router in order to receive traffic from the source via the Policy Tree. Although PTMR protocol implements policy inter-domain multicast routing, it does not address the problem of supporting QoS requirements.

In QRMP (QoS-Aware Multicast Routing Protocol) [11] two search modes are defined: single path mode and multiple path mode. The routing process starts with the single path mode, attempting to search only the unicast routing path traveled by the join request through the multicast tree. The *join request message* carries the QoS requirements. As it travels, it checks the resource availability of every intermediate node and proceeds only when the node has the required resources. If an intermediate node does not have the required resources it triggers the multiple path mode by sending a *not acknowledge message* to the previous node. Upon receipt the *not acknowledge message* the previous node sends the *join request message* to all neighbor nodes except those from which the join request and not acknowledge messages were previously received. Once a feasible branch is detected an *acknowledge message* is sent back along the branch that triggers to multiple path mode. If more than one acknowledge message arrives at this node, the node will select the best branch and reject all the others. In QRMP, tree construction occurs from new receiver in direction to tree instead of from first in-tree node found backwards to receiver, therefore it does seem adequate to asymmetric topologies. In multiple path mode the join request messages are flooded. Besides, it does not support the establishment of multicast routing policies.

IV. THE BASIC PRINCIPLES OF PAQoSIDMR

The Policy Aware QoS Inter-Domain Multicast Routing Protocol (PAQoSIDMR) aims to build inter-domain unidirectional multicast distribution trees, taking into account multicast specific routing policies and supporting QoS requirements.

New branches are built from in-tree routers to new members in order to properly deal with asymmetries.

In this section, inter-domain multicast routing requirements are presented, highlighting some key aspects of this proposal.

A. Intra-domain independence

A key requirement is to keep, as far as possible, the independence between interior and exterior multicast protocols, in a way similar to unicast.

There is no real need for an inter-domain multicast tree until some host decides to join a group rooted at some external domain. Intra-domain receivers and source(s) join together in an intra-domain multicast distribution tree, constructed by some multicast interior gateway protocol (M-IGP). Each domain is free to choose the best M-IGP according to its own needs, policies, etc.

Domain border routers must simultaneously run an interior and an exterior multicast protocol, in components usually called M-IGP and M-EGP respectively. When a host in the domain joins a group rooted at some other domain, the M-IGP component of the exit router informs the M-EGP component of that join.

B. Multicast policy awareness

At inter-domain level, routing policies are often more important than any other parameters. Also, multicast policies may be distinct from those defined for unicast routing. A transit domain may accept unicast traffic with no restrictions and wish to restrict multicast traffic to some specific multicast prefixes. Therefore, it would be nice to construct multicast inter-domain distributions trees using only multicast specific routing information. In other words, all control messages used for tree construction must be forwarded based on multicast addresses using only multicast routing information. This is one of the main goals of this proposal.

As pointed before, the current multicast frameworks already allow border routers to maintain multicast specific routing information by means of M-RIB and G-RIB. G-RIB is in fact a collection of group routes, containing multicast address prefixes, next-hop and AS Path attributes. Therefore, *Join-Request* messages can be addressed to the multicast group address, and successfully forwarded towards the Root Domain by looking up group routes in the G-RIB.

A little more complex is the addressing of the replies. *Join-Request* messages are answered by some in-tree border router using *Join-Answer* messages. The in-tree border router must send one *Join-Answer* message for each available, policy consistent return path registered. Those replies could be addressed to the unicast address of the originator border router using M-RIB to forward the messages. In this way, both *Join-Request* and *Join-Answer* messages would be forwarded taking into account the multicast specific policies established.

If M-RIB is not available it can be considered the possibility of using G-RIB to forward *Join-Answer* messages too. But in this case *Join-Answer* messages should be addressed to a special multicast address, that is representative of the New

Receiver Multicast Domain. This address should be included in the *Join-Request* message by the originator border router.

C. Dealing with asymmetries maintaining the propose of scalability

As pointed before asymmetries are more likely to happen at inter-domain level for several reasons, like distinct domain routing policies, unbalanced border router resource consumption, etc. So it is not possible to construct the inter-domain multicast distribution tree by ignoring them and assuming symmetric inter-domain links.

The right way to deal with asymmetry is to start a tree construction from its root towards the new leaf member. This solution has been proposed in [10] and has two known drawbacks: a greater join latency caused by excessive control messages and an excessive load on root routers.

In order to avoid root routers overload, in PAQoSIDMR proposal the join requests are handled by the first in-tree border router that receives them, thus relieving the root border router of that task. This enhancement gains extra importance at inter-domain level, if we consider that it can reduce the number of domains involved in tree branch setup, but also reduces the possibilities of finding a feasible path.

If the first in-tree border router fails to connect a new member, a controlled number of retries may be conducted by other in-tree border routers, including, as a last retry, the root border router.

D. Achieving QoS through path probing

New group members may have different QoS requirements and may express them in terms of multiple parameters, like bandwidth, delay and losses. Different alternative inter-domain paths may be evaluated in terms of how well do they fulfill those requirements. This path evaluation can be done either by using static metrics ([8]) or dynamic ones ([9]).

In order to capture dynamic metrics of different available paths, a probing strategy inspired in [9] is used. The border router that handles join requests, sends a limited set of probing messages (*Join-Answer*) back to the requesting border router. The probing is however restricted by current available group routes, after multicast policy enforcement.

However, some resource reservation protocol must still be used to preserve path quality on new tree branches, after successful joins.

E. QoS Information Propagation

The path probing strategy selects a connecting candidate based on the QoS proprieties of the connecting paths. However the QoS proprieties of the tree branch between the tree root and the connecting node should not be ignored. The QoS achieved by the new receiver would be the combination of these two values: the QoS parameters of the new connecting path with the QoS parameters of the tree branch already built between the tree root and the connecting node. For instance, if we consider the end-to-end delay as a QoS metric, and a new receiver asks for an end-to-end delay bellow x , it is not enough

to guarantee that the connecting path between the first in-tree node reached by *Join-Request* message and the new receiver is less than x . We must add this delay to the one verified between the tree root node and the connecting node.

Therefore when the *Join-Answer* message is generated it must include the state of the QoS parameters in the potential connecting node, and all the routers along the path between the connecting node and the new receiver should update these values, before forwarding the *Join-Answer* message.

To accomplish this some new features must be added to the path probing strategy used in PAQoSIDMR:

- new fields were added to the multicast routing table. Per each considered QoS parameter, the in-tree node should know how much QoS is available in terms of the corresponding parameter in the tree branch that ends in that node.
- in order to maintain this information up to date, a new control message was used called *QoS-state control* message. This message should be sent periodically by the tree root towards the multicast tree. Along its multicast path the *QoS-state control* message collects the QoS available for each considered QoS parameter. All in-tree nodes receiving those messages should update the QoS fields in the corresponding routing table entries.

These QoS state fields should be used when a *Join-Request* message arrives in an in-tree node in order to determine if the node is a potential connecting node or not. If the QoS parameters in the *Join-Request* message might be satisfied by the QoS state in that node then the in-tree node might be a potential connecting node and a set of *Join-Answer* messages are generated and sent. If not, the *Join-Request* message must be propagated towards the Root Domain until a new in-tree node is reached.

F. State Information

The state information needed at each router can not be excessive and may not depend on the number and dimension of multicast active groups.

In PAQoSIDMR proposal, in-tree border routers trigger a new tree branch construction after receiving a *Join-Request*, but they must keep a copy of the request together with list of all routes that were used to send *Join-Answer* messages, until branch setup fails or succeeds. For every *Nack* message received, one possible path is eliminated from the list, until it eventually becomes empty. In that case, a failure is detected, and the *Join-Request* message kept is forwarded to next hop.

If an *Ack* message is received, this temporary state information is also discarded, and the in-tree state information that already exists is updated with the new outgoing interface, from which the *Ack* message was received. Every out of tree border router receiving an *Ack* message must create in-tree state information.

V. A QoS AND POLICY-AWARE INTER-DOMAIN MULTICAST ROUTING PROPOSAL

It is assumed that multicast domains are interconnected by border routers running one multicast exterior gateway protocol (M-EGP), in parallel with some multicast interior gateway protocol (M-IGP), like in the unicast scenario. Furthermore it is assumed that multicast and unicast domains have the same administrative boundaries, but different routing policy rules. All border routers have permanent exterior MBGP sessions with neighbor border routers on adjacent domains, and also internal MBGP sessions with all internal border routers in the same domain. We also assume that there are different policy consistent alternatives routes for each destination.

Next sections present a detailed protocol description.

A. Detailed Description

Inter-domain tree construction is only required when a new member in a domain joins a group rooted at some external domain. This also means that any group is initially formed in the interior of a domain called the root domain. Inter-domain tree is always rooted at a border router of the root domain, which is the domain of the first source.

Figure 1 illustrates a successful inter-domain tree branch construction. To avoid unnecessary complexity, we assume that only one source (S) is generating multicast traffic. The multicast address in use was previously allocated to S's domain by MALLOC[12], and all intra-domain receivers join the source S in a multicast distribution tree constructed by one multicast interior gateway protocol (M-IGP). This procedure will not be described here, and is M-IGP dependent.

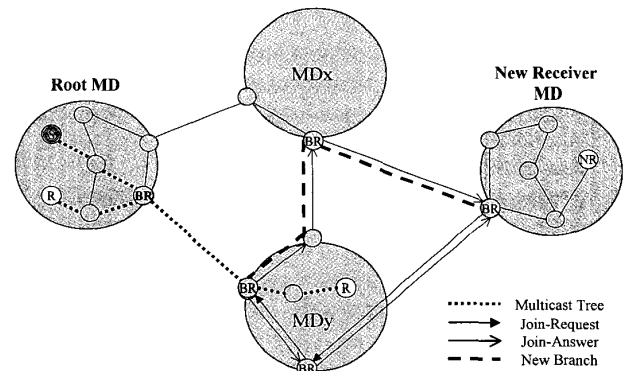


Fig. 1. Inter-domain tree construction

The source domain is designated by root multicast domain (Root MD), because the inter-domain unidirectional tree will always be rooted at one of its border routers (BR), as a result of the multicast address allocation procedure. A new inter-domain tree branch construction is always initiated when a new receiver (NR), belonging to a remote domain (New Receiver

MD), decides to join the group and none of its border routers is already connected.

The M-IGP protocol in use in the New Receiver MD domain will inform its *best exit* border router (BR) that a new receiver host (NR) wants to join the group. The border router, after verifying that none of its internal domain border peers is already connected, creates a *Join-Request* message, including the required QoS parameters in it, and addresses it to the group multicast address. It then issues a lookup in its multicast specific routing information base (G-RIB), and forwards the *Join-Request* towards the Root MD domain, using the forwarding information contained in the appropriate group route.

The *Join-Request* message is forwarded, hop by hop, until it reaches a border router in domain MDy that is already connected to the multicast tree. This router will then initiate the new branch construction by sending *Join-Answer* probing messages back to the New Receiver MD domain. The border router must send one *Join-Answer* for each available, policy consistent, return path registered in its multicast routing information base.

Before forwarding *Join-Answer* messages, all routers must collect dynamic QoS metrics and append them to the message. If the accumulated path QoS metric does not meet the QoS requirements included in the request, the *Join-Answer* must be discarded, and a *NAck* message is sent back to the border router. Those *NAck* messages can be forwarded using unicast routes.

Any *Join-Answer* that reaches the New Receiver MD border router, contains information about a valid feasible tree branch. A selection procedure must then be executed to select one of them according to some criteria, and, finally an *Ack* message establishes the branch, including the necessary state information in each router. The *Ack* message must travel in the exact opposite direction followed by the *Join-Answer* message selected.

1) *Dealing with tree branch setup failures:* The tree branch construction may fail if none of the possible alternative paths can meet the QoS requirements. This situation is illustrated in figure 2.

The in-tree border router of domain MDy that handles the *Join-Request* and initiates the tree branch construction, can also detect branch setup failures. It will always receive an *Ack* message when a feasible path is found, and a *NAck* message for all eliminated alternatives. In the example shown in figure 2, there was only one possible path available, and a *Join-Answer* was sent on that path. However, when the MDz border router handles the *Join-Answer*, it calculates the cumulated QoS metrics and concludes that the path can not meet the requested QoS, and immediately sends a *NAck* message back.

The MDy border router receives the *NAck* and, since it was the only path available, recognizes a branch setup failure. As a consequence, it forwards the original *Join-Request* message towards the Root MD in order to find another in-tree border router capable of initiating another tree branch setup. In the example shown in figure 2, that router is the border router of

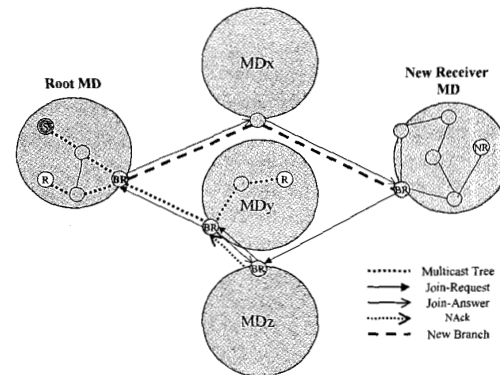


Fig. 2. Inter-domain tree construction - branch setup fail

Root MD domain. This router will try to create a new branch using exactly the same procedure described before.

In extreme situations, with consecutive retry failures, many in-tree border routers may be involved in the join procedure without success. This results in a very large join latency. Note that even in the worst case, when the *Join-Request* reaches the border router of the root domain, the procedure terminates. If there is no possible path a *NAck* message is sent back to the New Receiver MD border router.

In order to control better the join procedure, any router that fails to setup a branch includes its domain identification (AS number) in the *Join-Request* message before forwarding it. This allows all other border router to avoid paths already probed, by filtering routes containing that domain. A retry counter is also included in the *Join-Request* to reduce the number of retries to an acceptable limit. The counter may be initialized with a value of 2 (only first and last routers will try to construct the branch) or greater.

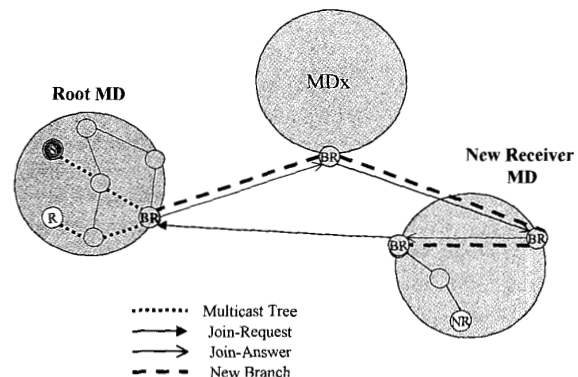


Fig. 3. Answers entering through a different border router

2) *Joining at a different border router:* Because *Join-Request* and *Join-Answer* messages are forwarded in opposite directions and also because asymmetries may exist, it is possible that the New Receiver MD domain best exit router is not the best entry border router. When this situation occurs, *Join-Answer* messages are not received by the same border router that sent the *Join-Request*.

In this situation, illustrated in figure 3, the border router that receives the *Join-Answer* must forward it internally to the peer border router that has originated the *Join-Request* messages.

3) *When a Join-Answer message finds an in-tree node:* As pointed before, the inter-domain routing environment is inherently asymmetrical therefore the PAQoSIDMR strategy builds directed multicast trees. Due to this reason it is probably that a *Join-Answer* message from the first in-tree router found by the *Join-Request* message might cross in its way other in-tree node. In this case, we must choose only one path between those two alternatives, in order to prevent that two different tree branches will meet in one node. There are different alternatives to solve this problem:

- to ignore this potential connecting path and send a *Nack* message to the node that generated the *Join-Answer* messages.
- if the tree branch that ends in this node can meet the QoS requirements of the new receiver, a new set of *Join-Answer* messages could be generated by this second in-tree node, and the previous one canceled. To accomplish this the new fields added in the routing table with the state of the QoS parameters until this in-tree node, are verified.
- if the tree branch that ends in this node can not meet the QoS requirements of the new receiver, but the new connecting path can (at least until reach this node), then we can let the *Join-Answer* message reach the receiver and after that if an *Ack* message is sent towards this connecting path, the previous tree branch should be pruned and replaced by the new connecting path.

In this stage of development we are using the first approach for simplicity. However, in a near future, we intend to evaluate the other two alternatives.

VI. DISCUSSION

In this article we have presented a proposal for inter-domain multicast routing, that builds policy and QoS aware inter-domain unidirectional multicast distribution trees. New members are connected through tree branches that meet some QoS requirements specified by them. Tree construction is requested by new members but initiated from in-tree members, in order to better deal with asymmetries. Alternative paths are probed with messages that collect QoS path metrics, and one of them is selected and established by the new receiver. No assumptions are made about intra-domain choices.

The path probing strategy, previously proposed in various different ways by protocols like YAM, QoSMIC and QRMP, was improved by PAQoSIDMR to address important inter-domain requirements, like policy awareness, intra-domain in-

dependence and AS Path asymmetries. For that reason, new branches are established in a three-way handshake, initiated by new receiver domain border router, with all control messages being forwarded using only multicast routing information. This ensures multicast policy awareness. In order to reduce join time latency and minimize the amount of costly inter-domain paths evaluated, only one request is sent, and thus, only one in-tree node - the first found - conducts a search for a feasible path. The probing messages are sent by that node on a limited set of alternate paths, all in compliance with multicast policies established for its domain. A controlled number of retries, issued by upper in-tree nodes may be issued in case of failure.

Those design options seem more appropriate for the current inter-domain scenario. The PAQoSIDMR proposal is currently being evaluated in Network Simulator (NS)[13].

VII. ACKNOWLEDGMENTS

This work has been partially funded by FCT under the Project QoS II, POSI EEI/10168/98.

REFERENCES

- [1] K. C. Almeroth, "The Evolution of Multicast: From the MBONE to Interdomain Multicast to Internet2 Deployment," *IEEE Network*, pp. 10–20, January/February 2000.
- [2] S. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoglu, D. Estrin, and M. Handley, "The MASC/BGMP architecture for inter-domain multicast routing," in *Proceedings of ACM SIGCOM*, Setembro 1998, pp. 93–104.
- [3] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, "Multiprotocol extensions for BGP-4," Internet Engineering Task Force, Request for Comments 2283, Feb. 1998. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc2283.txt>
- [4] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol independent multicast-sparse mode (PIM-SM): protocol specification," Internet Engineering Task Force, Request for Comments 2362, June 1998. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc2362.txt>
- [5] D. Farinacci, Y. Rekhter, P. Lothberg, H. Kilmer and J. Hall, "Multicast source discovery protocol (msdp)," June 1998, internet draft-farinacci-msdp-*.txt.
- [6] D. Thaler, D. Estrin, and D. Mayer, "Border gateway multicast protocol (bgmp): Protocol specification," November 2000, internet draft-ietf-bgmp-spec-02.txt.
- [7] P. Radoslavov, D. Estrin, R. Govindan, M. Handley, S. Kumar, and D. Thaler, "The multicast address-set claim (MASC) protocol," Internet Engineering Task Force, Request for Comments 2909, Sept. 2000. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc2909.txt>
- [8] K. Carlberg and J. Crowcroft, "Building shared trees using a one-to-many joining mechanism," *ACM Computer Communication Review*, pp. 5–11, 1997.
- [9] M. Faloutsos, A. Banerjee, and R. Pankaj, "Qosmic: Quality of service sensitive multicast internet protocol," in *SIGCOMM*, 1998, pp. 144–153. [Online]. Available: citeseer.nj.nec.com/faloutsos98qosmic.html
- [10] H. Hodel, "Policy tree multicast routing: An extension to sparse mode source tree delivery," *SIGCOMM Computer Communication Review*, vol. 28, no. 2, 1998.
- [11] S. Chen, K. Nahrstedt, and Y. Shavitt, "A qos-aware multicast routing protocol," in *INFOCOM* (3), 2000, pp. 1594–1603. [Online]. Available: citeseer.nj.nec.com/article/chen00qosaware.html
- [12] D. Thaler, M. Handley, and D. Estrin, "The internet multicast address allocation architecture," Internet Engineering Task Force, Request for Comments 2908, Sept. 2000. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc2908.txt>
- [13] K. Fall and K. Varadhan, "The NS Manual," Jan 2001, URL=<http://www.isi.edu/nsnam/ns/ns-documentation.html>.