

Low Cost Quality of Service Multicast Routing in High Speed Networks

Technical Report 13-97 - Edition 1.0

John Crawford and Gill Waters
J.S.Crawford@ukc.ac.uk, A.G.Waters@ukc.ac.uk

Computing Laboratory
University of Kent at Canterbury
Canterbury
Kent CT2 7NF

Telephone: +44 1227 764000
Fax: +44 1227 762811

22nd December 1997

Acknowledgement

This work was supported by the EPSRC (Grant GR/K55837)

Abstract

Many of the services envisaged for high speed networks, such as B-ISDN/ATM, will support real-time applications with large numbers of users. Examples of these types of application range from those used by closed groups, such as private video meetings or conferences, where all participants must be known to the sender, to applications used by open groups, such as video lectures, where participants need not be known by the sender. These types of application will require high volumes of network resources in addition to the real-time delay constraints on data delivery. For these reasons, several multicast routing heuristics have been proposed to support both interactive and distribution multimedia services, in high speed networks. The objective of such heuristics is to minimise the multicast tree cost while maintaining a real-time bound on delay.

Previous evaluation work has compared the relative average performance of some of these heuristics and concludes that they are generally efficient, although some perform better for small multicast groups and others perform better for larger groups.

Firstly, we present a detailed analysis and evaluation of some of these heuristics which illustrates that in some situations their average performance is reversed; a heuristic that in general produces efficient solutions for small multicasts may sometimes produce a more efficient solution for a particular large multicast, in a specific network. Also, in a limited number of cases using Dijkstra's algorithm produces the best result. We conclude that the efficiency of a heuristic solution depends on the topology of both the network and the multicast, and that it is difficult to predict.

Because of this unpredictability we propose the integration of two heuristics with Dijkstra's shortest path tree algorithm to produce a hybrid that consistently generates efficient multicast solutions for all possible multicast groups in any network. These heuristics are based on Dijkstra's algorithm which maintains acceptable time complexity for the hybrid, and they rarely produce inefficient solutions for the same network/multicast. The resulting performance attained is generally good and in the rare worst cases is that of the shortest path tree. The performance of our hybrid is supported by our evaluation results.

Secondly, we examine the stability of multicast trees where multicast group membership is dynamic. We conclude that, in general, the more efficient the solution of a heuristic is, the less stable the multicast tree will be as multicast group membership changes. For this reason, while the hybrid solution we propose might be suitable for use with closed user group multicasts, which are likely to be stable, we need a different approach for open user group multicasting, where group membership may be highly volatile.

We propose an extension to an existing heuristic that ensures multicast tree stability where multicast group membership is dynamic. Although this extension decreases the efficiency of the heuristics solutions, its performance is significantly better than that of the worst case, a shortest path tree.

Finally, we consider how we might apply the hybrid and the extended heuristic in current and future multicast routing protocols for the Internet and for ATM Networks..

Contents

1	Introduction	1
2	Multicast Routing	3
2.1	Networks, Routing and Resource Reservation	3
2.2	Multicast Group Types	5
2.3	Internet Multicast Routing Protocols	6
2.3.1	Internet Group Management Protocol (IGMP)	6
2.3.2	Multicast Shortest Open Path First (MOSPF)	7
2.3.3	Distance Vector Multicast Routing Protocol (DVMRP)	9
2.3.4	Core Based Trees (CBT)	11
2.3.5	Protocol Independent Multicasting (PIM)	12
2.4	Internet Resource Reservation Protocols	14
2.4.1	Resource Reservation Protocol (RSVP)	14
2.4.2	Internet Stream Protocol, Version 2 plus (ST2+)	15
2.4.3	Tag or Label Switching Multicast	16
2.5	B-ISDN/ATM Routing Protocols	17
2.5.1	Private Network/Network Interface (PNNI)	18
2.5.2	Broadband Integrated Services Digital Network (B-ISDN)	20

3	Low Cost Quality of Service Multicasting	22
3.1	The Bounded Delay, Minimum Cost Multicast Routing Problem	22
3.2	Heuristics with an arbitrary delay bound	22
3.3	The CCET Heuristic extended	23
3.3.1	A Worked Example	24
3.3.2	Time Complexity of the CCET Heuristic	24
3.3.3	Pathological Behaviour of the CCET Heuristic	25
3.3.4	When CCET costs increase	27
3.3.5	Multicast Tree Stability and Dynamic Groups	28
3.4	The CST _c Heuristic	28
3.4.1	A Worked Example	28
3.4.2	Time Complexity of the CST _c Heuristic	29
3.4.3	When CST _c costs more than SPT	29
3.4.4	Delay Bound Granularity and Time Complexity	30
3.4.5	Bounds on Tree Cost	30
3.4.6	Multicast Tree Stability and Dynamic Groups	30
3.5	The CSPT Heuristic	31
3.5.1	A Worked Example	31
3.5.2	Time Complexity of the CSPT Heuristic	33
3.5.3	When CSPT costs more than the SPT.	33
3.5.4	Bounds on Tree Cost	33
3.5.5	Multicast Tree Stability and Dynamic Groups	33
3.6	Evaluation Method	34
3.6.1	Benchmark Algorithms	34

3.6.2	Network Models	36
3.6.3	Link Metrics	37
3.7	Evaluation of the Candidate Heuristics	38
3.7.1	Performance Averages	38
3.7.2	Specific Multicast Comparisons	40
3.7.3	Network Load and Multicast Failures	41
3.8	sCSPT, Multicast Tree Stability and Dynamic Groups	42
4	Hybrid Approach	44
4.1	Hybrid Multicast Heuristic	44
4.2	Evaluation of Hybrid Heuristic	45
4.2.1	Performance Averages	45
4.2.2	Specific Multicast Comparisons	45
4.2.3	Network Load and Multicast Failures	47
4.2.4	Multicast Tree Stability and Dynamic Groups	47
5	Application of the Heuristics	48
5.1	Characteristics Required of an Heuristic	48
5.2	Characteristics of the Heuristics Evaluated	48
5.3	Combining Heuristics, Multicast Types and Network Types	50
5.4	Application of the chosen Heuristics in Multicast Routing Protocols	51
6	Conclusions and Further Research	54
6.1	Further Research Issues	55

Glossary

AS	Autonomous System
ATM	Asynchronous Transfer Mode
B-ISDN	Broadband Integrated Services Digital Network
CBT	Core Based Trees
CCET	Constrained Cheapest Edge Tree
CSPT	Constrained Shortest Path Tree
CST _c	Constrained Steiner Tree
DAR	Dynamic Alternative Routing
DNHR	AT&T's Dynamic Non-Hierarchical Routing
DTL	Designated Transit List
DVMRP	Distance Vector Multicast Routing Protocol
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ITU	International Telecommunications Union
LAN	Local Area Network
MOSPF	Multicast Open Shortest Path First
MST	Minimal Steiner Tree
OSPF	Open Shortest Path First
PIM	Protocol Independent Multicast
PNNI	Private Network/Network Interface
QoS	Quality of Service
QoS _R	Quality of Service Routing
RP	Rendezvous Point
RPB	Reverse Path Broadcasting
RSVP	Resource Reservation Protocol
SPT	Dijkstra's shortest path tree
sCSPT	stable Constrained Shortest Path Tree
sCST _c	stable Constrained Steiner Tree
ST2+	Internet Stream Protocol, Version 2 +
TOS	Type Of Service
UNI	User/Network Interface
VCI	Virtual Channel Identity
VPI	Virtual Path Identity
WG	Working Group

Chapter 1

Introduction

Many of the new services envisaged for B-ISDN/ATM high speed networks will require point to multipoint routing. Some of these services, such as interactive multimedia communications, will require real-time bounded delays on data delivery and will consume high bandwidths. Calculation of multicast routes for these types of applications must take account of their conflicting requirements for efficient network usage and real-time delay bounds on data delivery.

The problem of finding arbitrary delay bound low cost multicast routes in networks, where link cost and link delay are different functions, was first published by Kompella, Pasquale and Polyzos in [32]. Since then there have been a number of other proposals for solutions to this problem. Previous evaluation work [46][57][11][58] shows that, on average, these heuristics perform well. However, further detailed analysis and evaluation of some of these heuristics has shown that there is a wide variance in the efficiency of their solutions. Whilst on average one heuristic may be more efficient than another, either for all multicast group sizes or for a particular range of multicast group sizes, there are some multicast group and network combinations where this position is reversed. In particular, we have found that as a multicast group membership changes the heuristic that provides the most efficient multicast solution also changes. The results of our evaluation work indicates that it is difficult to predict which heuristic provides the most efficient solution for any particular multicast/network combination. The variance in the efficiency of the heuristic solutions is wide enough that on occasions Dijkstra's shortest path algorithm (SPT) calculated on delay is more efficient. By selecting two such heuristics that can be efficiently integrated with each other and the SPT algorithm, we propose a hybrid heuristic that produces reasonably consistent and efficient solutions to the multicasting problem, with an acceptable order of time complexity for all possible multicast groups in any network.

Another important characteristic of delay bound low cost multicast routing algorithms is the stability of their solutions, where multicast group membership is dynamic. We have found that, in general, the more efficient the solutions of an heuristic are, the less stable their solutions will be when multicast group membership is dynamic. While this may be an acceptable characteristic for closed multicast groups, such as a private video meeting, which is likely to have stable multicast group membership, it will not be acceptable for open multicast groups, such as a public video lecture where the multicast group membership may be quite volatile. For this reason we propose

an extension to an existing heuristic, so that the solutions it generates are stable, when multicast group membership is dynamic. This extension reduces the cost efficiency of the heuristic, but its performance is still a significant improvement over that of a shortest delay path tree solution.

The rest of this report is organised as follows :-

- Chapter 2 provides an overview of networks, routing and resource reservation. We identify different types of multicast groups and their characteristics, and then describe the multicasting and resource reservation methods used, or proposed for use, in the Internet and B-ISDN/ATM networks.
- In Chapter 3 we define the Bounded Delay, Minimum Cost Multicast Routing Problem. We follow this with explanations of how the heuristics we have evaluated work, and describe their behaviour. We then describe our evaluation method, and introduce a new network model we have used for some of our work. The results of the evaluation of the heuristics is presented. The chapter ends with the introduction of an extension to Sun's heuristic [50] that calculates stable multicast solutions of reasonable efficiency, for dynamic multicast groups. the results of our evaluation of the heuristics we have described.
- Chapter 4 introduces the Hybrid solution to the Low Cost, Bounded Delay Multicast Routing problem. We present the results of the Hybrid evaluation.
- Chapter 5 addresses how the Hybrid heuristic, and the extended "stable" heuristic for dynamic multicast groups, might be applied in networks.
- In Chapter 6 we conclude our work and identify further research that needs to be undertaken in this field.

Chapter 2

Multicast Routing

2.1 Networks, Routing and Resource Reservation

There are two major classes of routing methods to which routing algorithms can be applied: fixed and dynamic. Fixed routing assumes that the topology of the network and traffic requirements are known in advance, and remain constant, so that paths through the network can be calculated and downloaded into the network before they are used. Dynamic routing assumes that the network topology and traffic load are likely to change and may not be globally known throughout the network, and that route calculations have to be performed regularly to accommodate dynamic changes within the network.

Routes through a network can either be calculated at the data source, or by a route server or in a distributed manner by each of the nodes along the path. The process that calculates a path through the network is also responsible for establishing the path in the network. Source based dynamic routing methods require full knowledge of the network topology, and possibly that of traffic loading as well, to perform route calculations. In distributed dynamic routing, paths through the network are calculated at multiple points along the path between the source and destination as data is forwarded. As the topology of the network and the traffic load varies, paths through the network may change. The topology and traffic load information required by each routing point in the network varies depending on the routing algorithm used. Distributed dynamic routing algorithms vary from those that require only local information to those that require global knowledge of the network [51].

All routing methods have advantages and disadvantages. For example, in reliable networks that support permanent connections between end points, route calculation can be performed “off-line” and installed in the network prior to use. There is no necessity for the network itself to support either route calculation or maintenance of the information required to calculate routes. In such networks the switches have low functionality and traffic load for each communication is guaranteed. A disadvantage of such a network and routing method is that all changes to network load and end-to-end connections have to be made “off-line”. The network cannot dynamically adapt to changing usage. At the other extreme networks that support distributed dynamic routing may have

considerable functionality in their routers and require regular exchanges between routers of large volumes of topological and traffic loading information. While such a scheme enables the network to adapt to changes in network usage it requires additional capacity and processing capability to manage its dynamic nature, irrespective of the volume of traffic it is carrying at any time. Source based dynamic routing methods are a compromise between the extremes of fixed routing and distributed dynamic routing. Not all switching points in the network need to know the topology or loading of the entire network, but the route calculating points are responsible for establishing paths through the network. How often this is done depends on the type of network. It would not be very practical method for routing in a datagram network, but would be for a virtual circuit network.

Introducing multicasting to any routing scheme brings additional complexity to the routing method. In the most simple case this may be the extraction and combination of paths between the multicast source and its destinations to form the multicast tree. In the most complex cases it may involve the exchange of multicast destination location data, in addition to link state data, between all routing points in the network, as is the case with MOSPF [35] which floods link state advertisements and multicast group membership throughout the network. The introduction of multicast routing heuristics to dynamic routing methods may also bring with it the problem of multicast tree reconfiguration as destinations join and leave the multicast group. The resource savings made by using a near optimal minimum cost solution may outweigh the multicast tree reconfiguration costs, particularly if the multicast group membership changes frequently.

The main characteristics of networks that have an influence on the choice of multicast routing algorithm to be used by their routing protocols are :-

- Connectionless networks use distributed route calculations, performed by every router in the network, to forward user data across a network. This allows routers to alter paths taken by user data as changes in the network topology or traffic conditions occur, but incurs the cost of path calculation by each router along the path. User data is either carried across the network according to the “best efforts” of the network, or along paths that are maintained using “soft state” [10].
- Supporting resource reservation for user data flows in connectionless networks requires the establishment of “pseudo” paths between the source and destination routers. In order to maintain the ability of the network to re-route user data while “in-transit” these paths are maintained using soft-state. That is, reservations for resources along a user data path have to be maintained by the regular transmission of resource reservation messages along the user path. If these messages cease to be received by intermediate routers along the user path, then the resources allocated to the path are relinquished. Resource reservation messages use the same paths as the user data, so as paths change due to link congestion or failure the resource reservations also change paths.
- Connection oriented networks use source based route calculations, performed at the network ingress router or by a route server, to find paths across the network for user data. These paths are established in the network as “hard state” connections between the source and destinations before any user data is transmitted. This enables switches along the user data path to quickly and efficiently pass incoming data onto the appropriate outgoing links with minimal processing. If links in the network become congested or fail then the source/destination path has to be torn down and a new path calculated and established in the network.

- Resource reservation for user data paths in connection oriented networks is an integral part of the path calculation and establishment process. Once established, user data paths are held in place by hard-state. That is, the path resources remain intact until they are explicitly removed by the user.
- Fixed networks require user data paths to be pre-calculated and downloaded into the network before they are used. The network does not adapt to changes in topology or load variance.
- Resource reservation in fixed networks also needs to be pre-calculated and downloaded into the network.

2.2 Multicast Group Types

Multicasting protocols that are currently implemented or are in the process of development, provide support for two types of group communication.

- **Heavyweight:** The sender knows who the receivers are and is responsible for adding them to the multicast. Applications such as video conferencing and distributed gaming might require this type of service, since it may be important that the sender know who is participating in the multicast. Heavyweight multicasting protocols are likely to use considerable state data to establish and maintain multicast trees [48], but may suffer little dynamic behaviour by their membership.
- **Lightweight:** The sender is unaware of who the receivers are. Applications that may require this type of service are multimedia lectures and other distribution services, where the sender does not need to know the identities of receivers or where they are located. Lightweight multicast trees require little state data for their establishment and maintenance, and their group membership is likely to be volatile.

These contrasting requirements result in different methods of multicast tree construction.

- For a heavyweight multicasts, where all the recipients are known to the sender, the complete multicast tree can be constructed once. Thereafter it should remain fairly stable, provided the group membership is static.
- If a lightweight multicast has all its membership established before a sender starts transmitting, then a complete multicast tree can be constructed as soon as transmission starts. However, the multicast tree is likely to be less stable than a heavyweight multicast, since the multicast group membership is likely to change.
- For both heavyweight and lightweight multicasting the delivery tree may have to be calculated in a piecemeal fashion as receivers join the multicast group one at a time. This may happen in a heavyweight multicast where receivers are required to register with the source before they are allowed to join the multicast, or in a lightweight multicast where a receiver joins the group after the sender has started transmitting user data.

2.3 Internet Multicast Routing Protocols

While multicasting within local area networks (LANs) is generally available the ability to multicast across store-and-forward networks that interconnect LANs has not been available until recently, and is still being developed through the working groups (WGs) of the Internet Engineering Task Force (IETF). Much of the initial multicast development has been based on the original work of Deering and Cheriton [16] who proposed protocols that construct multicast trees rooted at a sending source. However, these source specific multicast solutions do not scale well for large internetworks because of the network resources they consume. For this reason multicast protocols that construct trees shared between a number of sources have been proposed by Ballardie et al [4] and by Jacobson et al [17].

Internet Protocol (IP) Multicasting is the delivery of an IP datagram to a group of hosts, where a host group is identified by a unique IP multicast address. Multicast group addresses are either persistent and well-known administratively assigned IP addresses or transient IP addresses which are assigned to multicast groups for as long as they have members. Delivery of datagrams is “best efforts”, as is the case with unicast IP datagrams. Datagrams are not guaranteed to reach all members of the multicast group, nor is the order of datagram arrival guaranteed.

Membership of a host group is dynamic; hosts may join and leave a group at any time. A host that is not a member of a multicast group may send datagrams to a group and may have no idea which hosts belong to the group. A host group can have no members. IP multicasting is receiver initiated; it is the host receivers that join and leave a multicast group. The sender may be unaware of who the receivers are, or where they are located.

2.3.1 Internet Group Management Protocol (IGMP)

The Internet Group Management Protocol [15] is used by IP hosts and their local routers to join and leave multicast groups. Support for the protocol is a prerequisite for all hosts and routers that support IP multicasting.

The protocol uses two message types:-

1. Host Membership Query (Queries).
2. Host Membership Report (Reports).

Multicast routers use these IGMP messages to find out which multicast groups have host members on the subnetworks to which they are directly attached. Periodically, multicast routers broadcast “Queries” on each of their locally attached subnetworks. Hosts on the subnetworks then reply to the multicast router using “Reports” to indicate to which multicast groups they belong. The protocol uses timers to spread the sending of “Reports” from hosts on an attached subnetwork over short time intervals to avoid implosions of concurrent messages on the subnetwork. This mechanism also reduces the volume of “Reports” returned to the multicast router as hosts that receive a report for a group to which they belong need not send a report themselves. Multicast routers then forward

remotely originated multicast group datagrams onto the subnetworks for which hosts have been registered as group members. If, after several periodic query transmissions on any subnetwork, no host replies are received for a multicast group previously active on the subnetwork, the multicast router assumes that there are no longer any host members for the group on the subnetwork and ceases to forward multicast datagrams to it. When hosts join multicast groups they send a report for the group to the multicast router rather than waiting for a “Query” first. This enables hosts to join groups immediately, rather than waiting for the multicast router to discover that the host wishes to join the group.

There is no explicit message from a host to indicate that it no longer wishes to receive data for a multicast group. The consequence of this is that a multicast router will continue to forward multicast data onto a subnetwork that has no hosts wanting to receive the data. Forwarding will continue until the timers used by the multicast router to wait for Host Membership Reports expire. The router will then know that there are no longer any hosts on the subnetwork that belong to the multicast group and so cease forwarding multicast datagrams. The IETF Inter-Domain Multicast Routing WG is developing future versions of the IGMP that will enable routers to immediately cease forwarding multicast datagrams when there are no longer any hosts on attached subnetworks belonging to the multicast group [6].

2.3.2 Multicast Shortest Open Path First (MOSPF)

The Open Shortest Path First (OSPF) [36] routing protocol is based on the proposal for a new ARPAnet routing protocol, by McQuillan in [34]. The protocol uses Dijkstra’s shortest path algorithm (SPT) to calculate routes in the network to which it is applied. Also known as Link State Routing, protocols based on Dijkstra’s SPT are used in, or have been proposed for use in both connection oriented, e.g. [54], and connectionless networks, e.g. [36] and [39].

The SPT calculation requires complete topological knowledge of the network to which it is applied. In dynamic connectionless networks, routers must periodically flood the state of their incident links to all the other routers in the network so that every router can maintain complete knowledge of the network’s topology. To forward unicast data, a router that receives a datagram uses the link state data and the SPT algorithm to calculate the shortest path, from the router to the datagram’s destination, on which to send the data. The router does not need to know the source of the data it forwards since the path it calculates is rooted at itself. The calculation requires only the destination identity and the network topology to find the forward path for the datagram.

The OSPF protocol is a link state routing protocol that has been designed to run between groups of routers that exchange routing information using a common protocol. Such a group of routers is termed an Autonomous System (AS) and OSPF is classified as an Interior Gateway Protocol (IGP) because it runs within an AS. To reduce the volume of link state data that must be flooded and maintained by routers throughout an AS, OSPF allows the grouping of contiguous networks and hosts into areas, each of which runs a copy of the basic OSPF protocol. Routers inside an OSPF area each maintain their own topological database of the area, which is not visible to routers in other areas. This means that the topological view a router has depends on the area it resides in, and is different from the topological views routers in other areas have. Areas are interconnected either by routers that are present in two or more areas (Area Border Routers) or by networks that are not in

any area, but which are connected to two or more areas by attached routers. The “network” that interconnects areas, which can be either contiguous or virtual, is termed the backbone network. A datagram routed between hosts in different areas will travel along an inter-area path from the local router to the area’s Border Router. It then follows a backbone path between the source and destination areas and finally reaches the destination router along a path from the destination area’s Border Router. Because the SPT calculation is rooted at each router along a path, all paths will be the shortest. Link state data from within each area is summarised and flooded throughout the backbone by the Area Border Routers. These summaries are also advertised internally to areas attached to the Border Routers, thus enabling routers within one area to calculate forwarding paths to destinations in other areas. As we shall see, this hierarchical routing structure impacts multicast route calculation.

For multicast data forwarding, the SPT must be rooted at the multicast source rather than at the forwarding router. To perform route calculation every multicast router in the network must therefore know, in addition to the networks link state data, the location of both the multicast sources and destinations for all multicast groups.

The Multicast Extensions to OSPF (MOSPF) [35] extend the capabilities of OSPF to provide source specific multicasting.

The MOSPF protocol registers multicast destination host group memberships with their local MOSPF router(s) by using the IGMP. The local MOSPF routers then flood the identities of hosts that have registered as receivers of the multicast group on their attached subnetworks to all other MOSPF routers in the area. This flooding process is periodically repeated for all multicast group hosts that remain registered with a local MOSPF router. This mechanism provides every multicast router with the identities of all the destinations for all multicast groups active within an area. MOSPF routers calculate multicast routes “on-demand” when the first datagram is received for a multicast group. This datagram contains the identity of the multicast source; the MOSPF router already has the identities of the multicast destinations and the topology of the network area. From these data the MOSPF router is able to calculate the forwarding path(s) for the multicast datagram, which may be cached for future use as the forwarding path will only need to be re-calculated if multicast destinations either join or leave the multicast group or if there are changes in the network topology. MOSPF supports calculation of multicast trees for each type of Internet Protocol Type of Service (TOS) [2]. However, it should be noted that the TOS feature has been removed from later editions of OSPF [37], and it will be removed from MOSPF [38]. Proposals for QoS OSPF are divided as to whether or not the TOS field should be used to specify QoS requirements [24][64]. Because of this interest TOS may continue to be developed.

Multicasting between MOSPF areas is managed in a different way to OSPF unicast routing. Each area in an AS will flood summaries of multicast group membership into the backbone. However, unlike OSPF, the backbone will not flood these data, nor will it advertise it’s own group memberships, into attached MOSPF areas. This procedure has been adopted to prevent the broadcasting of the locations of all members of all multicast groups throughout the AS, thus avoiding serious problems with multicast scaling. In order to reach multicast destinations that are outside a source’s area, multicast area border routers are designated as “wildcard” multicast routers. This means that they belong to all multicast groups, and so receive user data for all multicast groups. Since the border routers know which multicast groups have members in the areas attached to the backbone they are able to forward user data appropriately.

If a multicast is between a source in one area and a destination in another area, then the multicast tree no longer uses shortest paths, as it would if the user data were unicast. The reason for this is that a router performing the multicast tree calculation needs to know the distance from the source to itself. If the source is in an area remote to a multicast router this distance is not available, as the router only has the OSPF summary link state data describing the distance from the router to the source [35]. MOSPF therefore has to use the reverse path distance to the source for the SPT calculation across a network backbone. All additional forward links used in the tree calculation are also selected on the basis of their reverse path cost.

The distribution to, and storage of multicast source and destination data at every router in the network, coupled with the processing and storage costs of the SPT at each forwarding router for each multicast, impacts the scalability of link state multicasting. For this reason MOSPF is considered to be suitable for intradomain use only.

Some proposals have been made through the Internet-drafts mechanism of the IETF to extend both OSPF and MOSPF to include Quality of Service Routing (QoS) for data flows. Best efforts traffic would be handled in the same manner as it is currently, since it requires no QoS. These proposals have no authority, and the lifetime of such documents is limited to six months. However, the general direction these proposals are taking is of interest here.

In their proposed extensions to OSPF and MOSPF [64], Zhang et al suggest that the link state data exchanged between OSPF and MOSPF routers might include information about a link's available resources. and that data flows would request that specified quantities of these resources be allocated to them. The proposals include ideas on what metrics might be used in route calculations, such as link delay and token bucket depth and rate. There is also a suggestion, that because of the significant increase these extensions would make to the volume of link state data maintained by routers, that explicit route calculation by the data source router might be used instead of hop-by-hop calculation, and that routes should be fixed (pinned) once established. These proposals indicate a move towards a connection oriented approach for QoS.

2.3.3 Distance Vector Multicast Routing Protocol (DVMRP)

The DVMRP [55] is based on refinements by Deering and Cheriton [16] of Dalal and Metcalfe's Reverse Path Broadcasting (RPB) [13] method. It is an Interior Gateway Protocol for use within an Autonomous System, based on distance-vector routing [26].

In reverse path broadcasting a router forwards a data packet if, and only if, it arrives on the shortest path from the router to the data source. The router forwards the data packet on all its outgoing links except the link on which it arrived. By this means all routers in the network receive a copy of the data. For multicast delivery this is a waste of network and router resources since not all routers require the data. DVMRP uses a modified form of RPB (Truncated RPB) in which routers are able to identify attached "child" and "leaf" routers, and to know which "child" and "leaf" routers are members of multicast groups.

"Child" routers are relative to a multicast source and are identified by their distance from the source. In figure 2.1, routers *A* and *B* have to decide which will forward multicast data from *S*




Figure 2.1:

towards C . All the routers periodically exchange distance-vector data indicating their distances from each other. From this data, routers A and B each discover that router A is the closest to the multicast source S and is therefore the “parent” of router C , and hence responsible for forwarding multicast data to C from S . On discovering that router A is the “parent” of C , router B will no longer forward multicast data received from S towards router C . Router D is a “child” of B and router E is a “child” of C . By this method routers will receive only one copy of a data packet, rather than multiple copies as in RPB.

A “leaf” router is a “child” router that no other routers use to reach the multicast source. In figure 2.1, router C is a “child” of A , but it is not a “leaf” router because the attached router E uses it to reach the source S . Routers D and E are “leaf” routers. “Leaf” routers only forward multicast data onto subnetworks that belong to appropriate multicast groups. By this means the multicast tree is truncated at the “leaf” routers.

Later versions of DVMRP have been proposed and implemented and are documented in IETF draft documents [42]. These are “work in progress” documents which it is inappropriate to cite. However, it is appropriate to state that the later versions of DVMRP implement the Reverse Path Multicast method which allows “leaf” and up-stream “child” routers to be periodically pruned from the multicast tree provided that they have no dependent subnetworks belonging to multicast groups. The multicast tree periodically re-grows to discover if non-member subnetworks have subsequently joined a multicast group. The reader is referred to [16] for an expansion of this method.

The DVMRP only calculates a shortest path multicast delivery tree if the delay on all links is the same in both directions; otherwise some or all of the paths in the tree may not be the shortest. DVMRP routers do not need to know the topology of the network, but only the distances between themselves and all other routers in the network.

The periodic pruning and re-growing of multicast tree branches by DVMRP, irrespective of whether or not a branch has multicast group members attached, renders the protocol unsuitable for large scale internetwork multicasting.

2.3.4 Core Based Trees (CBT)

CBT is an architecture, proposed by Ballardie et al [4], for scalable internet multicasting and is based on the work of Wall [56]. The primary motivation for the development of CBT was to significantly improve the scaling factor for a multicasting method over that of the existing IP source rooted multicasting methods. Source rooted multicasting protocols, such as DVMRP and MOSPF, either use high volumes of bandwidth or require the exchange of large amounts of state data, irrespective of the number or distribution of multicast group members. For source rooted multicasting the scaling factor for each method is the product of the number of sources and the number of receivers in the multicast group, since each sender requires a source rooted multicast delivery tree. The objective of CBT is to reduce this factor to the number of receivers in the multicast.

A CBT consists of a primary core router and an ordered list of additional core routers, which are included for robustness, for each multicast group. At group initiation time the additional cores join the primary core to form a central hub of the CBT. The protocol does not define how cores are selected and placed in a network.




Figure 2.2: Multicast user data route in a CBT

Hosts wishing to join a multicast group will use a directory service, such as X.500 [25], to obtain a multicast group's identity and the list of the group's core unicast addresses. Using IGMP the host informs its local CBT capable multicast router that it wishes to join a multicast group, and provides it with the group identity and core list obtained from the directory service. The local multicast router then sends a join request towards a core router for the group. The join request is forwarded at each intermediate router along the path towards the core until it either reaches the core, or reaches an intermediate CBT capable router that is already a member of the CBT for the group. In general the CBT router will reply to the join request by sending a join acknowledgement back along the path towards the hosts local CBT router, thus establishing a branch in the multicast tree. All the CBT routers along the path become non-core routers for the CBT. There are exceptions to this action which occur for a variety of reasons, such as in loop detection. The reader is referred to [3] for a detailed description of CBT.

Sources wishing to send to a particular multicast group address their user data to the CBT core router for the group, and include the multicast group identity in the option field of the IP datagram.

When the user data reaches any of the CBT routers for the multicast group (core or non-core), the core address in the IP datagram destination field is replaced by the multicast group identity from the option field, and the data is multicast across the CBT. Each CBT router will forward the multicast data across all interfaces it has for the group tree (including the path towards the core), except across that on which the data arrived. Figure 2.2 illustrates how user data is sent towards a CBT for a multicast group, and is then multicast to all the group recipients.

The CBT for any multicast group is maintained by the exchange of “keep alive” messages between adjacent routers along tree branches. If a subtree becomes detached from the main tree it can either clear itself down, in which case each router in the subtree will attempt to re-attach itself to the main tree, or the router that became disconnected from its parent may attempt to re-attach the subtree to the main tree via an alternative core router.

The benefits of this multicasting method are that it utilises the underlying unicast routing protocol for the network, and only requires multicast state data to be maintained by the CBT routers for each tree to which they belong. Also, the exchange of “keep alive” messages is limited to only those routers that are in the multicast tree. Unlike DVMRP and MOSPF, CBT uses “hard-state” to maintain the multicast tree. This means that branches in the tree have to be explicitly torn down when they are no longer required. The costs of using CBT are that :-

- All user data traffic has to pass through the core router for a multicast group, thus forming a potential bottleneck.
- As Wall shows [56], the bound on the maximum delay of an optimal centre based tree is twice the shortest path delay. In other words, the delays experienced by CBT multicast deliveries could be up to two times that of an SPT based multicast.
- Because the multicast branch to a receiver is established by sending a join acknowledgement from the existing CBT back along the path along which the join request was received, CBT branches are reverse shortest paths. In networks with asymmetric links the CBT shared tree is therefore not a shortest path delivery tree. This characteristic may add additional delay to the already increased delay incurred by using a shared tree.

2.3.5 Protocol Independent Multicasting (PIM)

The Protocol Independent Multicast method was proposed to address the problems of multicasting to group members that might be sparsely distributed across wide area internets [17]. Like CBT, it was a design goal of PIM to use less resources than those used by source rooted multicasting protocols such as MOSPF and DVMRP. The PIM design also recognised that :-

- High data, rate low latency applications could only be served by source rooted multicasting (which CBT does not support),
- Low data rate, high latency applications with high numbers of sources would save network resources if supported by shared multicast trees,

- Traditional shared trees, such as CBT, may have a problem with multicast traffic concentration at their cores.

PIM supports two modes of multicasting. Dense mode is designed to support applications that have either low latency requirements or where the number of simultaneous senders is such that network performance would be unacceptably degraded by the concentration of traffic on a shared tree. Sparse mode is designed to multicast to a group whose membership is considerably smaller than the number of routers in the network, and which may be widely distributed over a large area and for which high latency is acceptable.

A PIM multicast tree initially consists of a set of rendezvous points (RPs), the addresses of which are associated with a multicast group. How RPs locations are decided has not been specified, but it will either be a configuration responsibility or their addresses may be obtained by hosts and distributed to PIM routers in much the same manner as envisaged for core lists in CBT. A host will join a multicast group by sending an IGMP report message to its local PIM router containing the multicast group identity and RP address. The PIM router will then send a join request towards the RP indicating that it wants to receive user data for the multicast, via a shared tree. As the join message travels towards the RP, intermediate PIM routers establish a path from themselves back towards the host's PIM router using a soft state mechanism and forward the join request to the next PIM router on a path to the RP. This process is repeated at each intermediate PIM router until the RP is reached. At the RP the join request is dropped and a path from the RP to the receiving host has been established. The RP maintains the path to the receiving host by periodically sending it reachability messages down the established path. Sources wishing to send to a particular multicast group address the user data to their local PIM router. The local PIM router then sends a register message (and user data) to all the RPs for the multicast group. The RPs reply to the source using join messages to set up user data delivery paths from the source to the RPs.

If a receiver requires source-specific tree delivery of user data, it first joins the shared tree for the multicast group. After receiving user data from the source the receiver's local PIM router can switch to the source rooted tree. The local PIM router recognises user data for the multicast group from which it obtains the source address, to which it sends a join request. Once the local PIM router starts receiving user data on the source rooted path, it sends a PIM prune towards the RP for the shared tree indicating that it no longer wants to receive user data from the source via the shared tree.

PIM offers much the same benefits as CBT, with the addition of the option to use a source rooted multicast tree if required. PIM, however, maintains the multicast tree using "soft-state". That is, PIM periodically sends refresh messages upstream towards each source to maintain the tree.

2.4 Internet Resource Reservation Protocols

2.4.1 Resource Reservation Protocol (RSVP)

All of the multicast routing protocols described above provide only “best-efforts” delivery of user data to receivers. With the advent of distributed applications such as multimedia conferencing, audio/video multicast delivery and distributed visualisation, “best-efforts” delivery is becoming untenable for some applications. Networks must be able to guarantee a requested quality of service for user data delivery, if the needs of these kinds of applications are to be met. To achieve this objective the Resource Reservation Protocol (RSVP) [63] has been proposed.

Strictly speaking, RSVP does not reserve any network resources for a communication. Rather it is a protocol used to establish router resource reservation state along a communication path between a source and receiver of a “flow” of user data. Sources can always send user data into the network without regard to the resources available for its delivery. This is the primary quality of the Internet Protocol. On the other hand, receivers know what they are capable of receiving and can therefore request the quality of service they want. In a multicast environment, different receivers may be either incapable of receiving, or not require, the same quality of service as other receivers. For these reasons RSVP is a receiver initiated protocol.

In order to reserve network resources, receivers need to know the path that user data takes from the source to reach them, and the transmission characteristics of the data flow that the source will send. The source therefore periodically sends path messages to the receivers that contains a specification of the data flow [41][9] that the source will send into the network. Path messages are carried towards the destinations by whichever routing protocol the network uses. They are not routed by RSVP. As the path messages pass through intermediate RSVP capable routers, they establish soft-state that describes the incoming and outgoing links for the multicast. A path message is then forwarded to the next-hop router for the multicast, and the process is repeated until the receivers router is reached. Intermediate routers along the path do not reserve any resources at this stage, they just establish path state between the source and receiver(s). On the basis of the flow specification contained in the path message and the resources available to a receiver, it replies to a path message by sending a reservation message back along the route the path message arrived on. This reservation message contains the description of the resources the receiver wants reserved along the path. As the reserve messages passes through routers on the return path resources are either reserved for the user data flow, or the reservation is rejected. The reserve messages establish the resource reservation state along the path established by the path messages.

RSVP is much more complex than described here. The protocol supports different styles of reservation requests and the merging of requests where paths for the same group meet. The use of soft-state to maintain reservations allows paths to be re-routed to adapt to changes in topology of the network. The reader is referred to [63] for a detailed description of the protocol.

In addition to a resource reservation mechanism (for example see [9][18][41]) the protocol requires admission control to manage network resources and routing protocols that will select data paths on the basis of requested quality of service [27] [30]. None of the routing protocols described above use admission control or quality of service path selection.

2.4.2 Internet Stream Protocol, Version 2 plus (ST2+)

The Internet Stream Protocol, version 2+, is an experimental protocol that enables applications to establish end-to-end paths, with specified reserved resources, for real-time data flows (streams) between a source and one or more destinations, across an internet [48]. The resource reservations are for single direction data flows, between a sender and any number of destinations, i.e multicast.

Like IP, ST2+ is a network layer protocol and is independent of underlying subnetworks. It uses the same addressing schemes as IP to identify hosts and may coexist with IP in network routers. ST2+ has the facility to encapsulate its messages within IP packets so that they can be transported transparently through IP routers that do not support the protocol.

Unlike IP and the multicasting protocols DVMRP and MOSPF, ST2+ is explicitly connection oriented. User data flows cannot be transmitted between a sender and receiver until a path, with the specified resources allocated, has been explicitly set up between them. The protocol has two components :-

- Path management.
- Real-time data transfer.

ST2+ path management is responsible for setting up, modifying and tearing down the paths used for transmitting data flows. To do this it uses two distinct services :-

- Routing; to select paths from the source to the destination(s).
- Resource Management; to reserve the appropriate resources for the data flow(s).

The ST2+ protocol specification does not define either of these services since they are considered to be external to the protocol, but it does assume their existence. ST2+ also makes assumptions about how these services are provided. The ST2+ setup protocol assumes that the external routing method it uses calculates unicast routes, on a hop-by-hop basis, as paths are constructed. How the routing method calculates which router is the next hop along any path is not defined by ST2+, but it could be by either a simple shortest path next-hop selection process, or by a more complex method based on network resources and a data flows quality of service requirements, such as proposed by Zheng and Crowcroft [65]. The calculation of complete data flow paths by the senders router (source routing), which was supported by earlier versions of the protocol, has been removed from ST2+. Resource management is invoked at each router along a data flow path, as it is constructed, to allocate local resources to the flow. The local resource manager is supplied with a specification which describes the resources required by the data flow. This it checks against the resources available, e.g. buffer space, bandwidth, and rejects the request if enough resources are not available or accepts it otherwise. Once a path is established, ST2+ path management has mechanisms at each router to efficiently switch data flow packets to the next router(s) along the path(s) and to monitor the status of routes. Data flow paths persist for either the lifetime of the data flow or until a transmission failure occurs.

ST2+ supports construction of multicast trees by either senders or receivers or by both senders and receivers.

- For sender construction of data flow paths (or multicast), the source knows the addresses of all the receivers. This information is sent to each ST2+ router along paths, as they are constructed. Hence all ST2+ routers know which receivers they have downstream. Sender initiated multicast can be considered heavy weight because of the potentially large destination lists used during path set up, and because this information is maintained by ST2+ routers within a multicast tree.
- For receiver initiated path construction, the sender sets up a data flow with no receivers. That is, the ST2+ router local to the data flow source creates entries in its database for the data flow, but does not setup any paths. For a destination to initiate receipt of a data flow, it must obtain the identity of the data flow (stream ID) and the IP address of the source. How this information is obtained is not within the scope of the ST2+ protocol specification, but it might be via a directory service, for example. With this information the ST2+ router local to a destination sends a join request towards the data flow source. The join message traverses ST2+ routers on a path back towards the source, until it reaches one that is receiving the specified data flow. This ST2+ router can then set up a path between itself and the destination on which to send the data flow. Receiver initiated multicasts are light weight in comparison to sender initiated ones. Less state data is held at each ST2+ router, since not all routers will necessarily know all of the downstream destinations they are forwarding data flows to.
- Sender and receiver construction of multicast trees is achieved through a combination of both of the sender method and the receiver join method.

The ST2+ protocol is far more complex than the overview given above. The protocol has to deal with a variety of problems that may occur during path set up, such as admission failure. It also has mechanisms to deal with failures in the network. The protocol supports groups of streams (data flows) and has the facility for an application to request modification to the resource requirements of a data flow. The characteristics of the ST2+ protocol that are important to our work are :-

- Multicast trees can be constructed either “en masse” by a single protocol message being sent into the network from the data flow source.
- Multicast trees can be built in a piecemeal fashion by receivers joining a data flow arbitrarily.
- Multicast trees are held in place by “hard state”, that is, they are connection oriented. Path construction is therefore a relatively expensive and time consuming process.

The reader is referred to [48] for a detailed description of the protocol.

2.4.3 Tag or Label Switching Multicast

It is recognised that the growing demand for increasing bandwidth in the Internet can, in some part, be achieved by improving the forwarding performance of routers. In IP networks, each router uses

a network layer routing calculation and data from a packet's header to determine the path on which to forward the packet. The packet header contains far more data than required to calculate the next hop towards its destination, particularly where a stream or flow of data is being transmitted. Working groups within the Internet community are developing Label Swapping (or Tag Switching) architectures to reduce the complexity of the current next hop IP routing method [43][44], and thus improve the forwarding efficiency of network routers.

In essence, label or tag switching is quite similar to cell switching in ATM networks, in that each switch (or router, in this case) maintains a table of incoming labels (or VPI/VCIs, in ATM) which are mapped to outgoing labels. Both architectures consist of two components :-

- Control; the mechanism used to bind network layer routes to tags. For example, in tag switching [43], a switch uses the routing method to identify the next hop router for a packet and requests a tag binding, from the next hop router, for the path. In label swapping [44], upstream and down stream routers agree bindings between labels and streams sent between them.
- Forwarding; the mechanism used to switch tags or swap labels and hence forward packets. A packet destination address is mapped to a tag or label as the packet enters the network. From then on, as the packet passes through the network, each router recognises the tag, or label, on an incoming packet, which it maps with the corresponding entry in its tag or label database. From this mapping the tag or label and forwarding information to the next router is obtained. In the case of multicast, an incoming tag may map to several outgoing tags.

These architectures also support explicit routing, where the entire route is chosen by a single router. The reader is referred to [43] and [44] as a starting point for further explanations of how these architectures might work, and for references to other work in this area. One observation made in [43] that is worth noting is the applicability of these architectures to ATM networks by the implementation of the Tag Switching Control component.

Like ST2+, neither architecture specifies how routes are calculated (for the binding mechanism), nor what routing information needs to be exchanged between routers in order to perform route calculations.

2.5 B-ISDN/ATM Routing Protocols

Asynchronous Transfer Mode (ATM) networks will offer significant improvements over existing local area and wide area network technologies [1]. In particular, they will support the integration of a wide variety of applications that use both voice, moving image video and other high bandwidth communications. The primary features of ATM networks are that they will offer users a guaranteed quality of service based on a collection of additive and non-additive route metrics (e.g. bandwidth and delay) while remaining scalable over wide areas. For these reasons, among others, ATM will be the underlying network technology for Broadband Integrated Services Digital Networks (B-ISDN)

Unlike IP networks, ATM is connection oriented and so requires both a signalling system and

routing methods to establish user calls. To support the features of ATM networks, signalling and routing will be far more complex than those that currently support IP networks.

2.5.1 Private Network/Network Interface (PNNI)

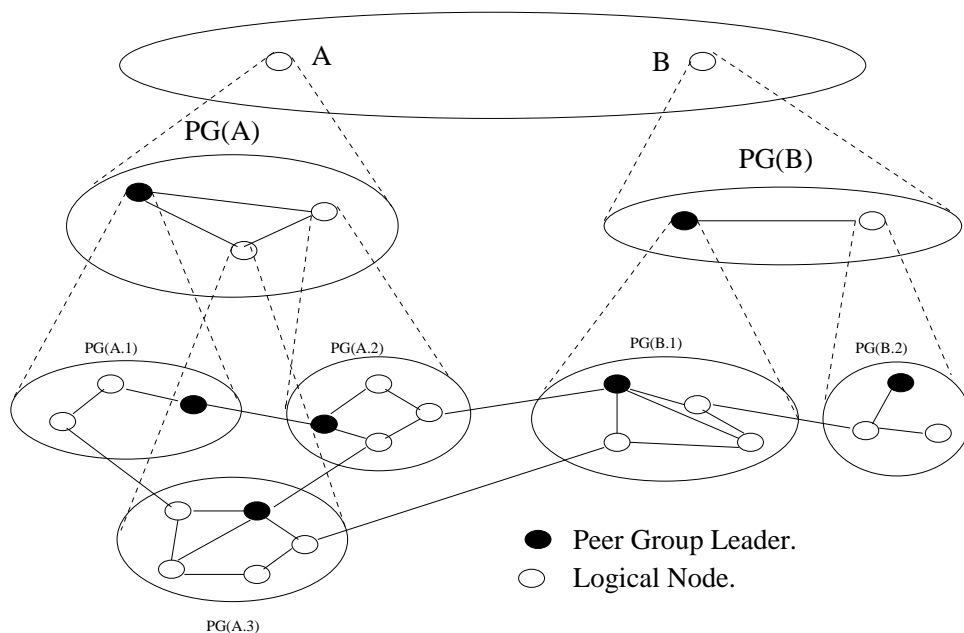


Figure 2.3: A PNNI hierarchy

The PNNI protocol [54] has been specified by the ATM Forum for use between private ATM switches and groups of private ATM switches. The protocol has two parts :-

- Signalling.
- Routing.

PNNI signalling is based on ATM UNI signalling [52][53] and is used to establish both unicast and multicast connections across an ATM network. Version 1.0 of PNNI does not support all the services defined for ATM UNI version 4.0 [53]. In particular, it does not support leaf initiated joins to multicast groups. Multicast destinations can only be added to a multicast group at the explicit request of the sender.

The PNNI routing protocol calculates the paths connections take across a network, and also maintains the data necessary to perform these calculations. Both of these protocols have an influence on the algorithms that can be used by PNNI to calculate multicast delivery trees.

PNNI routing specifies a hierarchical addressing scheme for routing. Within a PNNI switching system, starting from the lowest level (i.e. the switch level), logical nodes are organised as a hierarchy of peer groups. There may be one or more peer groups at each level of the hierarchy.

Within each peer group, one logical node is designated the peer group leader. Peer groups may then become logical nodes within the next higher level peer group of the hierarchy. A higher level peer group may also be a logical node of yet a higher level peer group. This structure, which may be of as many levels as deemed necessary for any particular switching system, is illustrated in figure 2.3.

Each logical node has an identity, which contains the nodes peer group identity. The peer group identity indicates the level of the peer group in the hierarchy and its parentage. Parent peer identities are always shorter than those of their children and the child peer group identity always contains the peer group identity of its parent.

PNNI is a link state routing protocol, where the link state data consists of topology state data and address reachability state data. Logical nodes exchange this data so that they can build databases of the network topology and tables of address reachability. Logical nodes only exchange link state data within their own peer group.

Peer group leaders aggregate and summarise the topology and reachability information of their peer group into a complex node description and flood it as link state data within the next higher layer peer group. Peer group leaders also pass link state data they receive at a parent level down to their child level and flood it to the other child logical nodes in their peer group. By this means destination reachability data and (a virtual) network topology permeates the entire network without the necessity of all link state data being flooded to all nodes in the network.

To set up a user call in an ATM network, PNNI has to perform two tasks :-

- path selection,
- establish call state along the selected path.

ATM users are able to specify a Quality of Service (QoS) and bandwidth requirement for each call they request. For this reason, path selection in PNNI is based on both the users requirements and the resources available within the network. Paths are calculated by the connection source node and established across the network using designated transit lists (DTLs). A DTL is essentially a stack of next node identities used by each node along the path to find out where the next forward node is. DTLs are removed from the stack and replaced by other DTLs as the path moves across peer groups. So, although the source node calculates the entire path, parts of the path are virtual in that they traverse complex nodes in higher level peer groups. The sections of the path that cross higher level peer groups have to be mapped onto the underlying lowest level peer groups (at switch level) as peer group boundaries are crossed. Using call admission control, the process of establishing call state along a selected path confirms whether or not the requested resources are available. If during this second step of the call set-up process, the required resources are not actually available along part of the selected path, the route is unwound to a point at which a new path can be calculated. To minimise the likelihood of happening, the path selection procedure may use a generic call admission control procedure to predict which links are likely to have sufficient resources available, and to use only these links in the path. The PNNI specification does not mandate any particular algorithm for path selection, although it provides an example of an acceptable one. The example algorithm uses a single additive metric to optimize the route it calculates.

User bandwidth and QoS requirements are specified using the UNI SETUP message ATM traffic descriptor and QoS parameter[52][53]. The traffic parameters specify Peak, Sustainable and Burst Cell rates for the communication, while the QoS parameters specify which class of service is required. For each QoS class of service values can be specified for a variety of performance parameters such as Cell Transfer Delay and Variation.

UNI version 3.1 signalling [52] specifies that multicast connections are set up by establishing an initial unicast path between the source node and a destination node. Further destination nodes are then included in the multicast delivery tree by means of “add party” requests. The source node can either wait for each add party request to be acknowledged (serial join) or it may have multiple requests outstanding at the same time (parallel join). The UNI 4.0 specification [53] adds the ability of a receiver to join the multicast group directly.

2.5.2 Broadband Integrated Services Digital Network (B-ISDN)

The origins of Integrated Service Digital Networks (ISDNs) are in the services provided by telephone companies, who recognised the need to integrate their separate voice, data and dedicated network services into a single network. The Broadband ISDN (B-ISDN) recommendations of the International Telecommunications Union (ITU) are being developed to further this aim by supporting a wider range of audio, video and data transfer services within the same network[8].

The original, or Narrowband, ISDN provides the subscriber with a “digital pipe” into an ISDN switch. Connections between ISDN switches are made using either packet switching (data), circuit switching (voice) or non-switched (dedicated) capabilities. The B-ISDN will maintain the concept of a “digital pipe”, but will integrate the packet switching, circuit switching and dedicated capabilities of N-ISDN into one broadband network, B-ISDN. The transfer mode to be used for the B-ISDN is the Asynchronous Transfer Mode (ATM). The B-ISDN architecture will be described in functional terms and so is implementation independent [7].

Telephone networks from which ISDN has grown, in general, use one of three routing architectures :-

- Direct routing; where routes are fixed and pre-established.
- Alternate hierarchical routing; where routes are organised into hierarchies, from local offices, through toll centres and so on, up to regional centres. Trunk routes are provided, beyond the tree structure, as alternative routes to be used when network loading dictates.
- Dynamic two hop alternate routing; where routes are calculated dynamically through more complex network architectures, at call set-up time. Routes are selected on the basis of network loading.

In alternate hierarchical routing, a path between two subscribers follows the lowest level of connectivity in the hierarchy that has the necessary resources available for the call. As resources are consumed on trunks at lower levels, the path is selected from trunks higher up in the network hierarchy. In dynamic two hop alternate routing, networks have a logical link between each pair of

switches, and all switches are equally responsible for routing calls. If a call cannot be established along a direct link between a pair of switches, then an alternative, two link route is used, if one is available. Otherwise the call is blocked. Examples of networks that use dynamic two hop alternate routing include AT&T's DNHR scheme, and DAR which is planned for BT's domestic network. Each of these systems uses a different mechanism for the selection of alternative routes [45]. Recent work has addressed mechanisms for dynamic two hop alternate routing in ATM networks [5].

The implementation of such routing schemes does not preclude the use of arbitrarily connected, multi-hop dynamically routed networks as the transfer scheme for B-ISDN. The ITU recommendations for B-ISDN do not specify, nor imply, the network architecture or how a provider of B-ISDN services should route calls through a network, although as Stallings points out [49], ISDN is evolving from the circuit switching technology of the telephone networks to the packet (or cell) switching technology of broadband networks (such as ATM), as it takes on broadband services. How architectures for B-ISDN/ATM networks evolve compared to ATM networks in the Internet remains to be seen.

Multicast connections in B-ISDN are set up in a manner similar to that of PNNI. A path is first established between a sender and one receiver, whilst indicating that the connection is to be point-to-multipoint. Once this connection set up has become alerting or active, additional receivers can be joined to the connection using "add party" requests. Multiple receivers can be pending at any one time. That is, the sender does not need to wait for the response to any other add party request before issuing another [29]. B-ISDN does not currently support receiver initiated joins.

Like PNNI, the sender's ATM traffic descriptor, broadband bearer capability and QoS parameters are specified using the User/Network Interface SETUP message [28].

There are significant differences between the way B-ISDN and PNNI use their underlying ATM networks. For example, in PNNI networks the user data path is the same as the call control path, which is not the case for B-ISDN. The reader is referred to both the ATM Forum and ITU Recommendations for detailed descriptions of these methods (which are beyond the scope of this report).

Chapter 3

Low Cost Quality of Service Multicasting

3.1 The Bounded Delay, Minimum Cost Multicast Routing Problem

The bounded delay minimum cost multicast routing problem can be stated as follows.

- Given a connected graph $G = \langle V, E \rangle$ where V is the set of its vertices and E the set of its edges, and the two functions: cost $c(i, j)$ of using edge $(i, j) \in E$ and delay $d(i, j)$ along edge $(i, j) \in E$.
- Find the tree $T = \langle V_T, E_T \rangle$, where $T \subseteq G$, joining the vertices s and $M_k, k=1, n \in V$ such that $\sum_{(i,j) \in E_T} c(i, j)$ is minimised and $\forall k, k = 1, n; D(s, M_k) \leq \Delta$, the delay bound, where $D(s, M_k) = \sum_{(i,j) \text{ on path}} d(i, j)$ for all (i, j) on the path from s to M_k in T .

Note that, if the delay is unimportant, the problem reduces to the Steiner tree problem. The addition of the finite delay bound makes the problem harder, and it is still NP-complete, as any potential Steiner solution can be checked in polynomial time to see if it meets the delay bound.

3.2 Heuristics with an arbitrary delay bound

Several heuristics have been proposed that use arbitrary delay bounds to constrain multicast trees. Kompella, Pasquale, and Polyzos [32] propose a Constrained Steiner Tree (CST_c) heuristic which uses a constrained application of Floyd's algorithm. Widyono [62] proposed four heuristics based on a constrained application of the Bellman-Ford algorithm. Zhu, Parsa and Garcia-Luna-Aceves [66] based their technique on a feasible search optimisation method to find the lowest cost tree in the set of all delay bound Steiner trees for the multicast. Evaluation work carried out by Salama, Reeves,

Vinitos and Sheu [46] indicate that Constrained Steiner Tree heuristics have good performance, but are inhibited by high time complexity.

The proposals for Constrained Shortest Path Trees by Sun and Langendoerfer [50], which we abbreviate as CSPT and by Waters [59], which we abbreviate as CCET (Constrained Cheapest Edge Tree), generally have a lower time complexity than Constrained Steiner Tree's but their solutions are not as efficient. In this evaluation work we compare our heuristics against the solutions




Figure 3.1: The example network

generated by the Sun and the Kompella et al heuristics, both of which we describe below. For consistency all the worked examples use the network illustrated in Figure 3.1. The arbitrary delay bound, Δ , is set to 9 for the Kompella et al and Sun heuristics because they both seek solutions with a delay less than Δ . For the Waters heuristic Δ is set to 8 since it seeks a solution where the delay is less than or equal to Δ . It should be noted that the solutions produced by each heuristic in these examples do not illustrate their general performance.

3.3 The CCET Heuristic extended

The CCET heuristic was first published in [57] along with some simple preliminary evaluations. In [59] important variations of the heuristic were introduced and comprehensively evaluated.

The original heuristic and its variant [11] were bound by either the broadcast delay or the multicast delay. Here we extend the heuristics such that they are bound by an arbitrary delay, Δ . The effect this has upon the heuristic is to vary the size of the search space for the multicast tree in the second stage of the process (steps 4 and 5). The greater value Δ has, the larger the search space becomes. The extended procedure for the CCET heuristic is as follows:

1. Use an extended form of Dijkstra's shortest path algorithm, to find for each $v \in V - \{s\}$ the minimum delay, dbv , from s to v . As the algorithm progresses keep a record of all the dbv found so far, and build a matrix $Delay$ such that $Delay(v, k_i)$ is the sum of the delays on edges in a path from s to k_i , whose final edge is (v, k_i) , for each k that is adjacent to v .
2. The arbitrary delay bound is Δ . Set all elements in $Delay(v, k)$ that are greater than Δ to ∞ . The matrix $Delay$ then represents the edges of a directed graph derived from G which contains all possible solutions to a multicast tree rooted at s which satisfy the delay constraint.

3. Now construct the multicast tree T . Start by setting $T = \langle \{s\}, \emptyset \rangle$.
4. Take $v \in V_T$, with the maximum dbv , that is less than Δ , and join this to T . Where there is a choice of paths which still offer a solution within the delay bound, choose at each stage the cheapest edge leading to a connection to the tree.
5. Include in E_T all the edges on the path (s, v) not already in E_T and include in V_T all the nodes on the path (s, v) not already in V_T .
6. Repeat steps 4 and 5 until $V_T = V$, when the broadcast tree will have been built.
7. Prune any unnecessary branches of the tree beyond the multicast recipients.

3.3.1 A Worked Example

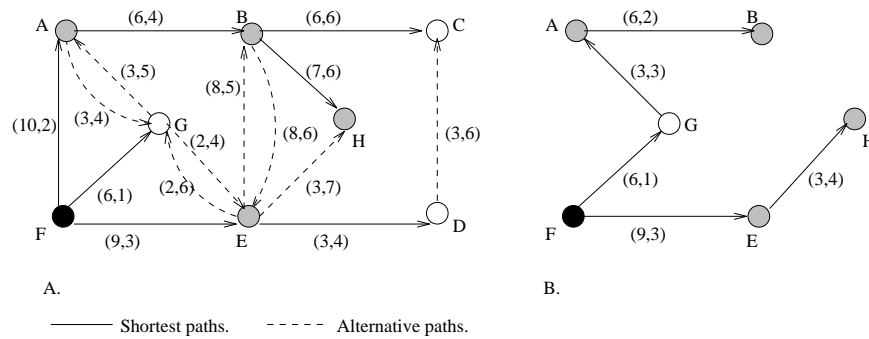


Figure 3.2: The CCET heuristic

To illustrate the working of the heuristic we start with the graph shown in Fig. 3.1. The bracketed parameters for each link indicate $(cost, delay)$. The example finds the multicast route from source F to destinations A, B, E and H.

The application of the extended form of Dijkstra's algorithm pruned to the arbitrary delay bound Δ results in the directed graph shown in Fig. 3.2A where the parameters shown against each link represent the edge cost and total delay from the source F to reach the node at the end of that link. The multicast tree is then constructed starting with $T = \langle F, \emptyset \rangle$. First H is connected to F using the path HE, EF. Node C is connected via the path CD, DE and then node B is connected via path BA, AG, GF. Finally, the edges CD and DE are pruned to give the multicast tree in Fig. 3.2B) with a cost of 27 units and a final delay bound of 7.

A minimum spanning tree, when pruned to the multicast, gives a cost of 20 units and a delay bound of 7.

3.3.2 Time Complexity of the CCET Heuristic

The first stage, determining the directed graph, has the same time complexity as Dijkstra's algorithm, $O(n^2)$. The vertices can be put in delay bound order during the construction of the directed

graph.

In the second stage, building the multicast tree, requires a depth first search from each leaf node to find a path to the source. As the multicast tree grows the search space for each leaf to source node path becomes smaller. The time complexity of the depth first search is $O(\max(N, |E|))$ [22] where N is the number of nodes, and $|E|$ is the set of edges, in the leaf node to source tree. The values of N and $|E|$ depend on the topology of the network, the position of the multicast source node and the arbitrary delay bound. As the network edge density or the arbitrary delay bound increase so do the values of N and $|E|$. In practice, an optimal upper bound can be placed on the arbitrary delay to limit the values of N and $|E|$.

3.3.3 Pathological Behaviour of the CCET Heuristic




Figure 3.3: Example of a Rogue Path

The heuristics first stage constructs a directed graph of paths between the multicast source node and every other node that can be reached within the delay bound. The number of paths between any node and the source offered by this graph depends on the delay bound and the graph density. The higher either of these values are, the more paths available. This graph also contains rogue paths that exceed the delay bound because they include a high proportion of alternative edges. The “cheapest” path in Figure 3.3 between node G and the source includes three alternative edges GH, FB and CD with a delay of 22. If the arbitrary delay bound placed on this multicast were 21 the “cheapest” path is a rogue and would not be detected until the last link, DS, was added to the path.

In the second stage the heuristics extract the bounded delay minimal cost tree from the bounded directed graph constructed in first stage. To do this the heuristics start by constructing a path between the node furthest from the source and the source node, that is within the delay bound. If the arbitrary delay bound is the broadcast bound or less the path between the first node selected and the source will be a minimum delay path, irrespective of cost. If the delay bound is greater than the broadcast delay then this path is not necessarily a minimum delay path. This first path becomes the trunk of the bounded delay minimal cost tree. The heuristics then add the return paths from each node successively closer to the source until the tree is complete. Two characteristics affect how paths join the existing tree:-

- Nodes closer to the source have a greater choice of paths back to the source because of the




Figure 3.4: Rogue Paths and the Pathological Walkabout

slack between their shortest path delay to the source and the delay bound on the multicast.

- As the tree grows the probability of a path joining the tree at a node closer to itself than the source increases.

The combination of these characteristics generally minimise the probability of loops occurring during tree construction. When the tree is young and sparse branches are likely to be close to their shortest paths to the source. As the tree grows, branches are more likely to meet the existing tree sooner.

In multi-cluster networks the furthest node from the source and the source node may be in different clusters. If the network comprises three or more clusters then some clusters may not be on the trunk path from the furthest node back to the source. Under these conditions the two characteristics described do not apply to 'non-trunk' clusters. The available delay slack within a cluster may be high and the existing tree may be remote from the node wishing to attach to it. If the cluster is very dense and has few links on a path to the source the second stage may examine a large number of rogue paths that ultimately break the delay bound before finding the path back to the source. Without low delay slack or a close-by existing tree the search for a path to the source may take a very long time. How this problem can occur is illustrated in Figure 3.4. When constructing the return path from node A to the Source edge BC is chosen instead of BJ because it is the cheapest exit. The heuristic then starts to search the tree rooted at node C for a path back to the source. It is not until this search has been exhausted that the edge BJ is examined.

This behaviour may also apply to single cluster networks where the arbitrary delay bound is very much larger than the network diameter and few edges are removed in the first stage of the heuristic.

3.3.4 When CCET costs increase

The CCET heuristic selects return paths on the basis of the “cheapest” exits from each node, back towards the source, that do not violate the arbitrary delay bound Δ . In some networks this rule can cause multicast trees found by the heuristic to be more expensive than might otherwise be expected.





Figure 3.5: Costs increase as Δ is relaxed. Figure 3.6: CCET more expensive than SPT.




Figure 3.7:

The cost of multicast trees found using the CCET heuristics can increase when the arbitrary delay bound is relaxed. Such a case is illustrated in Figure 3.5. With a delay bound of $\Delta = 5$ the multicast tree will include the edges FC,FG and GH at a cost of 8 units. This happens because the edge BH will have been excluded as it gives node H a delay of 8 units from the multicast source node F. If Δ is increased to 8 the edge BH is included and will be selected as the “cheapest” return route from node H towards F. The multicast tree then becomes FC, CB, and BH at a cost of 9 units.

The cost of solutions found using Dijkstra’s shortest path algorithm can sometimes be cheaper than those found using the Water’s heuristic. The multicast tree found using Dijkstra’s algorithm for the network in Figure 3.6 includes the edges FB,FG and GH, the shortest paths. The cost of this tree is 20 units. If the CCET heuristic is used to calculate the multicast tree with an arbitrary delay bound of $\Delta = 6$ the solution will include edges FB,FG,GA and AH because AH offers the “cheapest” exit back to the source from node H. The cost of this tree is 21 units.

As CCET multicast trees grow their cost difference from the corresponding SPT solution will fluctuate. The addition of a single node to the multicast can cause the CCET solution to change from being cheaper than the SPT to becoming more expensive. Figure 3.7 illustrates the CCET and SPT solutions for a multicast. The multicast source is node F and the delay bound, Δ , is greater than 26. In the first instance the multicast includes only node G. The SPT solution will choose the route FA,AG to reach G at a cost of 13 units. CCET will choose the route FB,BH,HG to reach G at a cost of 11 units. If the multicast grows by the addition of node A the cost of the SPT solution will remain the same since A is already on the path to node G. The CCET solution has to add the link BA, increasing the tree cost to 15, to reach node A.

3.3.5 Multicast Tree Stability and Dynamic Groups

The broadcast tree constructed by the CCET heuristic will be the same for all multicast groups with the same multicast source and arbitrary delay bound. This occurs because the heuristic constructs the broadcast tree using only the multicast source and the arbitrary delay bound. The multicast tree is extracted from the broadcast tree by removing unwanted branches. This means that in a dynamic environment where the multicast trees grows and dies, the broadcast tree only need be recalculated if the topology of the underlying network changes.

3.4 The CST_c Heuristic

The algorithm has three main stages.

1. A closure graph (complete graph) of the constrained cheapest paths between all pairs of members of the multicast group is found. The method to do this involves stepping through all the values of delay from 1 to Δ (assuming Δ takes an integer value) and, for each of these values, using a similar technique to Floyd's all-pairs shortest path algorithm (see [21]).
2. A constrained spanning tree of the closure graph is found using a greedy algorithm. Two alternative selection mechanisms are proposed, one based solely on cost, the other on cost and delay. In our evaluation we use the most efficient of these (cost only) which selects edges for the spanning tree using the function :-

$$f_C = \begin{cases} C(v, w) & \text{if } P(v) + D(v, w) < \Delta \\ \infty & \text{otherwise} \end{cases}$$

where $C(v, w)$ is the cost of a constrained path from node v to node w , $P(v)$ is the delay from the multicast source to node v and $D(v, w)$ is the delay on the path (v, w) .

3. The edges of the spanning tree are then mapped back onto their paths in the original graph. Finally any loops are removed by using a shortest paths algorithm on the expanded constrained spanning tree [31].

3.4.1 A Worked Example

Applying the first stage of the heuristic to the network in Figure 3.1 produces the constrained closure graph illustrated in Figure 3.8A. Note that this graph need not be a complete graph so long as there are paths between every multicast node and the source. Path AF includes node G , path HF includes E and path EF includes G . There is a conflict between the paths HF and EF which will result in a loop occurring in the constrained spanning tree. The other paths have no intermediate nodes.

Figure 3.8B shows the spanning tree obtained from the closure graph using the edge selection function f_C . Expansion of the spanning tree into their original paths results in a graph with a loop




Figure 3.8: The CST_c heuristic

(Figure 3.8C.) which when removed produces the solution in Figure 3.8D. This tree has a cost of 29 units and a delay of 7.

3.4.2 Time Complexity of the CST_c Heuristic

The calculation of the constrained shortest paths during the first stage of the heuristic is the most time consuming, with a complexity of $O(\Delta n^3)$, where n is the number of vertices in the graph [21]. The second stage has a time complexity of $O(m^3)$ where m is the number of nodes in the multicast group. Mapping the closure graph back onto the original graph has a time complexity of $O(mn)$. Loop removal using Dijkstra's algorithm has a time complexity of at most $O(n^2)$. This gives the algorithm an overall time complexity of $O(\Delta n^3)$. However, since it takes $O(\log \Delta)$ bits to represent Δ , the solution is exponential unless Δ is bound.

3.4.3 When CST_c costs more than SPT

In most cases CST_c calculates multicast solutions that are cheaper than those produced by SPT, but it does sometimes generate more expensive solutions. Figure 3.9 illustrates such a case. The multicast is from the source node, F, to the destination nodes, B and D. The arbitrary delay bound is 12. The first stage of CST_c constructs a closure graph from the cheapest constrained paths between the multicast nodes and the source in the underlying graph. From the closure graph CST_c selects




Figure 3.9: CST_c more expensive than SPT

the solution. In the example the multicast solution selected will be the closure graph edges FB and FD at a cost of 22 and a delay of 11. The final stage of CST_c maps the closure graph solution back onto the original graph, providing the solution FA,AB and FC,CD. The SPT algorithm will select paths solely on the basis of the delay from the source to each node. The solution SPT provides is FA,AB and AD at a cost of 21 and delay 5. By chance the SPT has been able to take advantage of the common edge FA, which was not available in the closure graph for CST_c.

3.4.4 Delay Bound Granularity and Time Complexity

As the arbitrary delay bound, Δ , increases the computation time of the heuristic increases. This characteristic can be overcome by decreasing the granularity of Δ through scaling, although this will compromise the accuracy of the results [62].

3.4.5 Bounds on Tree Cost

If the arbitrary delay bound applied to the heuristic is effectively infinite then the solutions produced will be similar to those calculated using the Minimum Spanning Tree (MST) for the Steiner Tree Problem [23].

3.4.6 Multicast Tree Stability and Dynamic Groups

The topology of a CST_c multicast tree may be reconfigured as the tree grows or shrinks. The second stage of the algorithm applies a greedy process to extract the solution from a closure graph that comprises only the multicast nodes. If a node is added or removed from the multicast the closure graph changes, and so the greedy process has a different set of nodes to consider. This may




Figure 3.10: Changing topology of CST_c multicasts

result in a multicast solution with a topology different from its predecessor. Figure 3.10 illustrates an example of such a change. Initially the multicast is between the source node, F , and a single multicast node, D . The closure graph constructed by CST_c will be the single virtual edge FD . When this edge is mapped back onto the original graph we get the solution FG, GM, MD for the multicast. The node B is then added to the multicast, giving the closure graph of virtual edges FD, DB, FB . From this graph the greedy process selects a multicast tree FB, BD . When this solution is mapped back onto the original graph the multicast solution becomes FG, GB, BD . The addition a single node to the multicast has caused edges GM and MD to be removed from the solution and replaced by edges GB and BD .

3.5 The CSPT Heuristic

This algorithm has three steps.

1. Using Dijkstra's shortest path algorithm compute a lowest cost spanning tree to as many destination nodes in the multicast as is possible without any path breaking the arbitrary delay bound, Δ .
2. Use Dijkstra's algorithm to compute a shortest delay path tree to those multicast nodes not reached in the previous step.
3. Combine the lowest cost spanning tree from the first step with the shortest delay path tree from the second step making sure that the delay to any destination node does not break the delay bound, Δ , and that all loops are removed.

3.5.1 A Worked Example

Applying the first step of the heuristic to the network in Figure 3.1 produces the minimum cost path tree illustrated in Figure 3.11A. Node H is not included in this tree because it's minimum cost path has a delay of 8, which breaks the delay bound. Figure 3.11B is the shortest delay path




Figure 3.11: The CSPT heuristic

tree constructed only as far as node H , the multicast node not yet included in the solution. The combination of the minimum cost path tree and the shortest delay path tree will create a loop nodes F and A . For this reason the edge FA is selected in preference to edge GA to give the final solution in Figure 3.11C. This tree has a cost of 31 units and a delay of 6. Loop removal in the

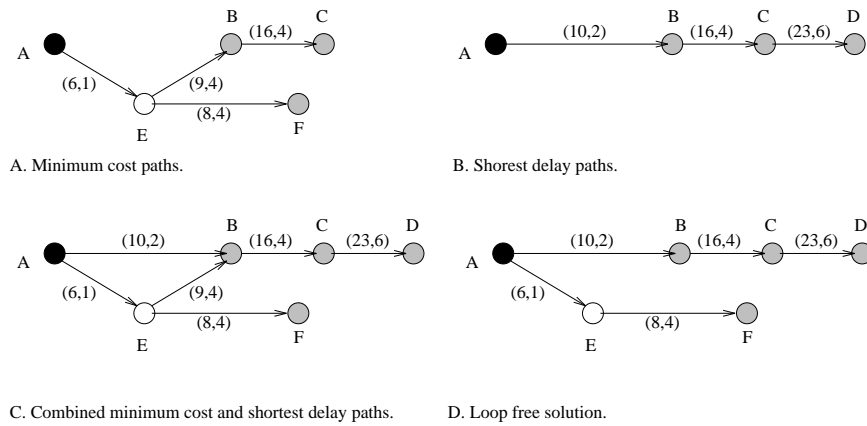


Figure 3.12: Loop removal in the CSPT heuristic

CSPT heuristic is much simpler than it is with the CST_c heuristic. Because steps 1 and 2 both use Dijkstra's algorithm to compute their trees, a loop occurs. The loop can be avoided by selecting, from the loops downstream node, the shortest delay path tree branch in preference to the minimum cost path branch. This will increase the tree cost, but prevents violation of the delay bound. For example, Figure 3.12 illustrates how the lowest cost spanning tree (A) and shortest delay path tree (B) when combined create a loop (C). By choosing the path AB from the shortest delay path tree and ignoring the path EB from the minimum cost path tree we obtain a loop free solution that does not violate the delay bound (D).

3.5.2 Time Complexity of the CSPT Heuristic

Each of the first two steps of the heuristic have the time complexity of Dijkstra's algorithm, which is at most $O(n^2)$. Because these two steps are independent of each other they can be performed in parallel. The last step has a time complexity of $O(n)$.

3.5.3 When CSPT costs more than the SPT.




Figure 3.13: CSPT more expensive than SPT.

For the majority of multicasts CSPT also calculates solutions that are cheaper than those produced by Dijkstra's SPT algorithm. As with CCET, there are also some cases where the cost of solutions found using the SPT algorithm can be cheaper than those found using the CSPT heuristic. In figure 3.13, for a delay bound greater than 8, to connect the multicast nodes A,G and H to the source F, the CSPT heuristic will use the path AB,BF at cost 18 and path GH, HF at cost 13 because they are the shortest paths based on cost between the multicast nodes and the source. This results in a multicast tree of cost 31. The SPT algorithm based on delay will choose the path AG,GH,HF at a cost of 21 to connect all the multicast nodes to the source.

3.5.4 Bounds on Tree Cost

The minimum cost of a CSPT tree will be obtained where the multicast group can be connected using only the minimum cost spanning tree. On the other hand the maximum cost will be incurred if the multicast group can only be connected using the shortest delay path tree.

3.5.5 Multicast Tree Stability and Dynamic Groups

As CSPT multicast trees grow their topologies are prone to reconfiguration if the arbitrary delay bound is less than the delay along the cheapest path to the new destination node. This happens if the heuristic has to add the new node using a shortest delay path, which may require the removal of a cheapest path from the existing tree. In figure 3.14A, a CSPT multicast tree has been constructed to the nodes G and H from the source F. The arbitrary delay bound, Δ , is 7 so the multicast tree is constructed entirely of the cheapest paths FA,AG and FA,AH. If node E is added to the multicast




Figure 3.14: Changing topology of CSPT multicasts

it cannot be connected by the cheapest path (FA,AH,HE) because the arbitrary delay bound will be violated. The heuristic calculates a shortest delay path to node E which comprises the links FB, BG and GE which is added to the tree. The link AG is removed from the tree to prevent a forward loop. The resulting multicast topology is significantly different from its predecessor. The instability of CSPT multicasts can be eliminated by a minor modification to the heuristic. Instead of calculating a solution for the multicast group the calculation includes all nodes in the network, as is the case with the CCET heuristic. The multicast tree is extracted from the broadcast tree calculated in this way. We refer to the modified version of the CSPT as the stable CSPT, or sCSPT. What effect this change has on the efficiency of the multicast solutions is assessed in the evaluation 3.7. Figure 3.31 compares the cost efficiency of multicast trees generated using CSPT and sCSPT. For smaller multicast group sizes the original heuristic produces, on average, more efficient solutions than sCSPT. As the multicast group size increases the performance of the heuristics converges, as expected.

3.6 Evaluation Method

3.6.1 Benchmark Algorithms

The ideal benchmark algorithm to use would be one that produces an optimal delay bound minimum cost multicast tree. To find such a solution requires the enumeration of all spanning trees in the network that are bound by the arbitrary delay. Kompella [31] bases such a process on Kirchoff's method for enumerating spanning trees but removes all the spanning trees that break the delay bound and selects the cheapest tree as the solution. As Cayley's theorem shows [20] this method becomes NP-complete as the network size increases because it has a time complexity of $O(n^{n-2})$. It is only practically useable for very small networks.

In the evaluation of their own algorithms Kompella et al first compare the performance of their most efficient method against the performance of optimal solutions [32]. The evaluations are limited to a few multicast groups sizes in small networks (20 nodes) because of the complexity of finding optimal solutions. They then use their most efficient algorithm as the benchmark for larger evaluations (50 to 100 node networks).

Salama also uses the optimal solution method as his benchmark for small network (also 20 nodes), but uses the bounded shortest path algorithm of Zhu et al [66] as the benchmark for larger networks.

Salama chose this algorithm because in his simulations with small networks the solutions it produced were the closest to the optimal.

Sun used Kompella et al's most efficient algorithm as his benchmark, in much the same way as Kompella used it.

We have chosen two different benchmarks for following reasons. Firstly we use the MST heuristic [23]. The solutions provided by this heuristic are not bound by any arbitrary delay, but we record their delays and use them in the evaluation. Our reasons for this choice is that we can compute large "minimum" cost trees within a reasonable amount of time ($O(n^3)$) and that the solutions of the Kompella et al algorithm which are the most efficient of all the heuristics under evaluation, tend towards the MST solution as the arbitrary delay increases. The MST heuristic works as follows :-

1. Use Floyd's all pairs shortest paths algorithm to construct a complete graph, T , of all the nodes in the multicast group, including the source node, from the original graph, G .
2. Use Prim's minimum cost spanning tree algorithm to construct a minimum cost spanning tree, T' of the closure graph, T .
3. Map the minimum cost spanning tree T' back onto the edges they represent in the original graph G , removing any loops that may occur.

The MST heuristic applied to the graph in Figure 3.1 results in the tree in Figure 3.15B.

Secondly we use Dijkstra's shortest path tree as a benchmark to evaluate the cost savings made by using the various heuristics. We do this for the pragmatic reason that there is no optimal (or near optimal) solution to the problem that can be credibly used in real networks. We consider it more realistic to look at the savings a multicasting heuristic can make in comparison to the shortest path multicasts currently in use [15]. Dijkstra's shortest paths algorithm applied to the graph in Figure 3.1 results in the tree in Figure 3.15A.




Figure 3.15: Benchmark solutions

3.6.2 Network Models

Three models were used for the evaluation of the heuristics. The first is attributed to Waxman [60] and is used to generate single cluster networks such as backbones or autonomous systems. Doar’s model [19] produces network clusters interconnected via a central core network. Although this model is intended to represent hierarchical networks the interconnection between clusters and the network core is rudimentary. It has been retained because it’s use lead to the discovery of rare but extreme behaviour in our heuristics. The third model extends Doars ideas by generating a random backbone network to interconnect the network clusters. This model is a more realistic representation of real networks. We have not yet considered hierarchical networks with more than two levels nor interconnected adjacent area networks as models for our evaluations, although these have been proposed in both unicast and multicast routing protocols for the Internet [36] [35].

The Waxman model randomly distributes nodes over a rectangular coordinate grid. The Euclidean metric is then used to determine the distance between each pair of nodes. Edges are introduced with a probabily that depends on their length. The edge probability is given by:

$$P(\{u, v\}) = \beta \exp \frac{-d(u, v)}{L\alpha} \quad (3.1)$$

where $d(u, v)$ is the distance from node u to u , L is the maximum distance between two nodes and α and β are parameters in the range (0,1). Small values of α increase the density of short edges relative to longer edges, while larger values of β produces higher node degrees. Edge lengths are used to represent delays. This implies that node queue processing delay can be ignored. Edge costs are selected at random from the range $[1..L]$. Figures 3.16 and 3.17 are examples of networks generated using Waxmans model.




Figure 3.16: $\alpha = 0.25, \beta = 0.15$




Figure 3.17: $\alpha = 0.75, \beta = 0.75$

The network model suggested by Doar is based on that of Waxman. Doar introduces a factor, related to the number of nodes in the network, to scale the probability of edges being included. The probability function becomes :

$$P(\{u, v\}) = \frac{k\bar{e}}{n} \beta \exp \frac{-d(u, v)}{L\alpha} \quad (3.2)$$

where k is a scale factor related to the mean distance between two random points, \bar{e} is the mean degree of a node and n is the number of nodes in the graph.

Doar goes further by introducing hierarchical graphs as networks models. These are generated using the modified probability function to generate clusters of networks that can then be connected to a central core network using a fixed number of links. Figure 3.18 is an example of such a network.

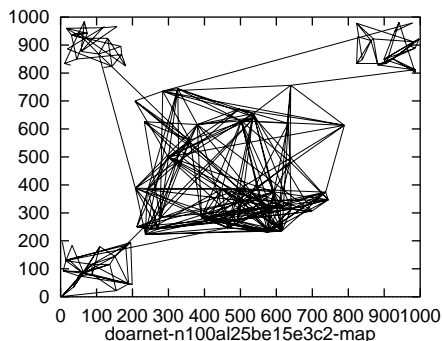


Figure 3.18: $\alpha = 0.25, \beta = 0.15$




Figure 3.19: Crawford model network

The cluster interconnection mechanism proposed by Doar means that the number of links connecting each cluster to the core is predefined and static. We modified this rudimentary method to obtain the third network model. Our network model uses the modified probability function to generate cluster networks. We then use the same function to generate a network with as many nodes as there are clusters. Each node in this network represents one of the cluster networks and its edge lengths are scaled to represent interconnection distances between the cluster networks. This 'backbone' network is then mapped onto the clusters by connecting together nodes selected at random from each cluster. Figure 3.19 illustrates a backbone interconnecting network clusters. In some clusters internet links terminate at different nodes, as illustrated by nodes A,B and D. In other clusters some nodes act as terminators to more than one link, as in clusters C and E. The internal structure of each cluster will otherwise be similar to that in Figures 3.16 and 3.17.

3.6.3 Link Metrics

Arbitrary Delay Bounds

In our network models we use Euclidean edge lengths to represent link delays. These delays have no units of time associated with them. Node queue processing delay is assumed to be negligible, as implied by the network models.

In his simulations of multicasting algorithms Salama [46] uses a similar network model, but assigns delay on the basis of a propagation speed of two-thirds the speed of light in networks distributed over a 3000 by 2400 Km^2 grid. He further assumes that each network node is a non-blocking ATM switch and that node queue processing delay can be ignored when calculating link delay. His simulations use an arbitrary delay bound of 0.03 seconds, as might be required by interactive video and voice applications. This delay was chosen to allow enough time for higher level end-to-end protocols to process transmissions without degrading the required quality of service.

In Distributed Interactive Simulation applications human reaction times may require delay bounds

of 200-300ms. Tightly coupled applications may require bounds as low as 100ms. If these delay are applied over networks of typically 50ms diameter then the range of arbitrary delay bounds required by applications may vary from being very close to the multicast bound up to several times the network diameter.

We choose three arbitrary delay bounds to evaluate the multicast algorithms. The tightest delay bound is the multicast delay. Evaluations using this arbitrary delay bound will indicate the minimum improvement in network utilisation achievable by each heuristic. Our second choice is to use the network diameter as the arbitrary delay bound. This purely arbitrary choice provides an evaluation “mid-point”. As the arbitrary delay bound increases in relation to the network diameter so the maximum improvement in network utilisation for each heuristic will be achieved. Our third delay bound is that of the MST, the lowest cost tree that can reasonably be calculated.

Interpretation of the evaluation results can be adapted to time based link delays by scaling network edge lengths appropriately.

Costs

If a provider charges by data units then only the link cost is required in the heuristics. On the other hand if the provider charges for duration (such as might be with a CBR call) then the link cost might be “link speed (or bandwidth) * duration used”. This has implications for the metrics we use in our evaluations. See the evaluation.

3.7 Evaluation of the Candidate Heuristics

3.7.1 Performance Averages

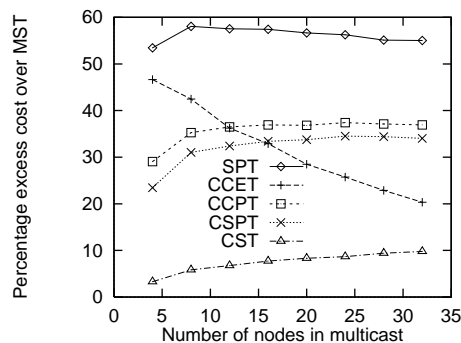


Figure 3.20:

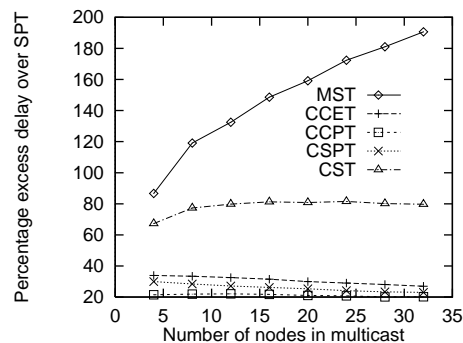


Figure 3.21:

Figures 3.20 and 3.21 illustrate the percentage excess costs and delays of using the four heuristics described above. The excess costs are measured relative to the MST benchmark and the excess delays are measured relative to Dijkstra’s SPT benchmark. The evaluation networks have 35 nodes each and are of low edge density. The arbitrary delay bound, Δ , is set to the diameter of the network.




Figure 3.22:

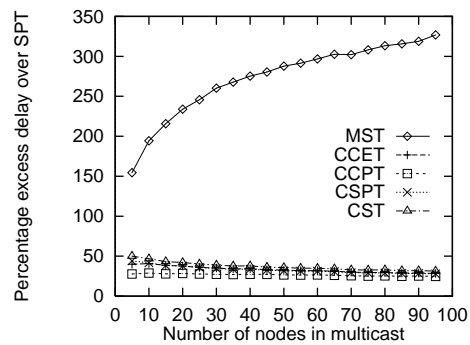


Figure 3.23:

The algorithm of Kompella et al generates multicast solutions that are on average cheaper than the other heuristics, although as the size of the multicast group size increases the Water heuristic's solutions converge with those of Kompella et al. The performance of the CCET heuristic is much better than CSPT and Crawford beyond a multicast of size 20, but is worse for smaller multicasts. The Crawford variant of the CCET heuristic shows poor performance in comparison with that of CSPT. As illustrated in Figure 3.21 Kompella et al and CCET make greater use of the delay bound to find cheaper paths. The performance of the heuristics is confirmed for larger networks in figures 3.22 and 3.23, although as the network size increases so does the percentage excess cost of each heuristic.

It is reasonably clear from the above evaluation results, albeit that they only cover low density single cluster networks, that the Crawford heuristic is in general significantly less efficient than the others. It is also clear that the algorithm of Kompella et al is the most efficient, but is hampered by its high order time complexity.

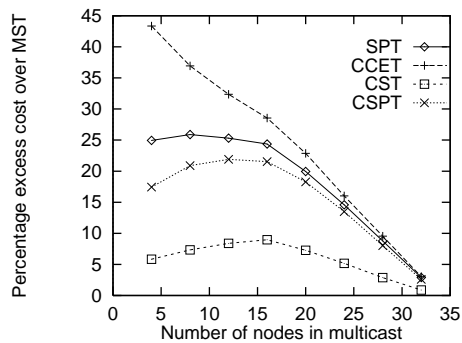


Figure 3.24:

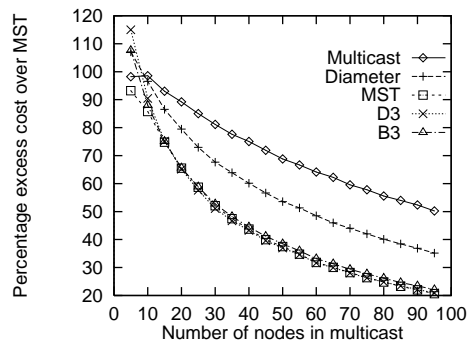


Figure 3.25:

Figure 3.24 illustrates the behaviour of the heuristics when the cost of each link in the network is the same. The CST_c and CSPT heuristics maintain a greater efficiency than SPT, but CCET becomes more expensive. Both CST_c and CSPT attempt to specifically minimise the cost of their solutions by growing their trees from the multicast source out to the destinations, so their behaviour is as expected. The CCET heuristic adds links by working back from each destination towards the source using the cheapest return links, so it has little control over the number of links a path might require. Rather it relies on the likelihood that cheaper links will lead to cheaper paths, irrespective of the number of links in the path. It also relies on the chance of paths meeting

at the soonest possible point. If link costs are all the same the heuristic will choose the first return link from each node on a path back to the source, without minimising the number of links in the path. In this case path cost is a product of the number of it's links and so this is not minimised. Figure 3.25 illustrates the improvement in solution efficiency of the CCET heuristic, as the arbitrary delay bound is increased. We have observed that as the delay bound approaches the MST delay, improvements in solution efficiency become negligible. Figure 3.25 shows how the costs of multicast trees with delays of three times the network diameter ($D3$), three times the broadcast delay ($B3$) and the MST delay almost coincide.

Up to these delay bound limits the number of nodes visited during the tree search in the heuristic's second stage is of $O(< 2n)$, by observation. If the delay bound goes much beyond these limits the heuristic is occasionally prone to very long execution periods which suggests that either N or $|E|$ (or both) can become unacceptably large.

3.7.2 Specific Multicast Comparisons




Figure 3.26:

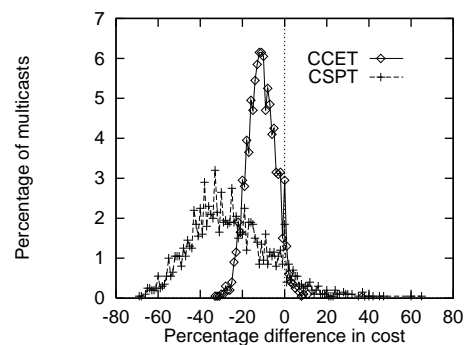


Figure 3.27:

The CSPT heuristic is generally better for smaller multicast group sizes, while the CCET heuristic is more suited to larger multicasts, although this is not always the case. Figure 3.26 illustrates the percentage of times CCET solutions are more expensive than those of CSPT and when the solutions of both CSPT and CCET are more expensive than Dijkstra's SPT. In nearly 5% of the sample's solutions for multicast groups of 95 nodes the solutions generated by CCET were more expensive than those generated by CSPT. Similarly, in 7% of the solutions for multicast groups of 5 nodes the solutions generated by the CSPT solution were more expensive than those generated by CCET. For smaller multicast groups sizes both CSPT and CCET generated some solutions that were more expensive than the shortest paths solutions. Although the CST_c heuristics has generally much better performance than any of the other heuristics, it is also prone to generating some inconsistent solutions, as illustrated in figure 3.28.

This behaviour suggests, as might be expected, that the solutions each heuristic generates depend both on the algorithm and the topology of the network to which they are applied. None of these heuristics can realistically provide the "cheapest" multicast solutions in all networks for all multicast groups.

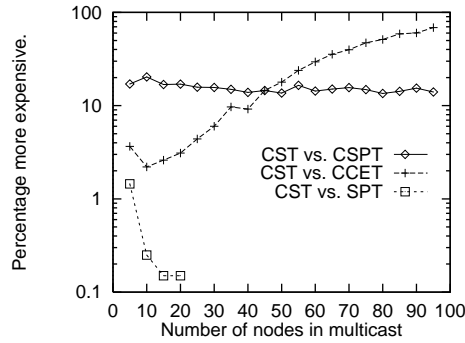


Figure 3.28: Effect of different delay bounds on CCET

3.7.3 Network Load and Multicast Failures

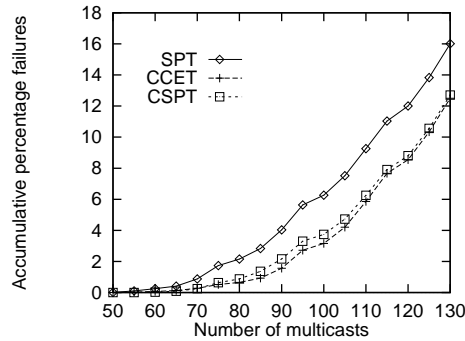


Figure 3.29: Network load and multicast failures

Figure 3.29 illustrates how each heuristic consumes network resources. In this evaluation multicast groups of differing sizes are successively applied to networks and the resources they consume are removed from the network before the next multicast is applied. Each multicast specifies a bandwidth requirement for the connection. Any links in the network that do not have sufficient bandwidth available for the connection are removed before the multicast is applied. The SPT heuristic then finds a multicast solution based solely on delay. All the other heuristics select links that have the highest available bandwidth, provided the delay constraint is not violated. As successive multicasts are applied bandwidth is used up and multicasts begin to fail because either the network becomes disconnected or, as the diameter of the network effectively widens, the delay bound, Δ , is violated. This latter condition applies to the SPT sooner than the other heuristics because it does not necessarily spread the the multicast tree evenly over the network [14]. Shortest path links that centre on the multicast source will become heavily used while those not centred on the source will not be utilised. As the bandwidth of the centred shortest paths is used so the paths become unusable, thus widening the network diameter and hence path delays. The other heuristics spread the paths they use for the multicast tree over a wider area than the SPT because they use the slack between the shortest path delay and the delay bound, Δ to try and minimise tree cost.

3.8 sCSPT, Multicast Tree Stability and Dynamic Groups

As discussed previously, the CST_c and CSPT heuristics reconfigure their multicast tree topology as nodes join and leave the multicast group. This behaviour is also true for the benchmark MST algorithm. Table 3.30 gives the percentage of multicast tree reconfigurations that occurred during the tree growth, and the average percentage of path changes per reconfiguration. These results were obtained by growing a sample of 20 multicasts per network from group size 1 to group size 34 over 200 single cluster 35 node networks. The arbitrary delay bound was set to the network diameter. CSPT and CST_c can be modified, as described previously, to extract their multicast

Heuristic	MST	CST_c	CSPT
% multicast trees reconfigured	89.22	100	67.53
Average % of paths changed per reconfiguration	23.98	20.66	11.56

Figure 3.30:

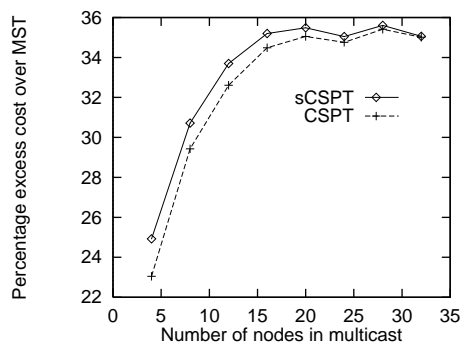


Figure 3.31:




Figure 3.32:

trees from static broadcast trees, although this change affects the cost performance of the multicast solutions. We label the modified versions of these two heuristics sCSPT and sCST_c, respectively. Figure 3.31 illustrates the cost performance for CSPT and sCSPT. The cost difference between the variants is greater for smaller multicasts than larger ones, as would be expected. The actual cost difference, as opposed to the difference in percentage excess over the cost of an MST solution, is on average less than 1%, making the sCSPT an acceptable alternative to CSPT where required. Figure 3.32 illustrates the cost performance for CST_c and sCST_c. In this case, for smaller multicasts, the percentage excess cost over the MST solution is markedly different for each of the heuristic variants. In fact the sCST_c cost performance tends towards that of CCET, as illustrated in figure 3.33. By using a cheapest cost path algorithm to attach nodes to the multicast tree the CSPT and sCSPT heuristics find cheaper multicast trees for sparse multicasts. As the multicast tree grows paths to additional nodes are included in the multicast tree on the basis of either their cost or delay from the source. There is no direct attempt to aggregate routes and so the tree cost steadily increases. Both sCST_c and CCET aggregate cheapest paths to produce efficient solutions for the broadcast case. This means that paths to many nodes are not the cheapest in themselves. The cost saving of the multicast occurs because expensive paths are aggregated to reach more destinations. Consequently, as the broadcast tree is pruned to find solutions for multicast groups of reducing

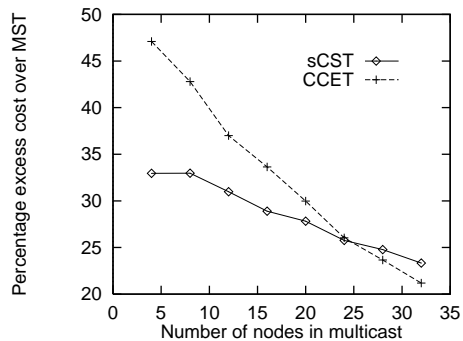


Figure 3.33: Cost performance of sCST_c verses CCET

size so the result is a tree that has an increasing proportion of expensive paths to fewer nodes, i.e. there is less aggregation of paths that are not the cheapest routes to the destination nodes.

The advantage of extracting multicast solutions from static broadcast trees is that the multicast tree does not need to be re-configured as nodes join and leave the group. When a node joins, if it is not already in the solution, its path of attachment to the multicast tree can be easily extracted from a cached broadcast tree. Nodes leaving the multicast are removed only if they are not on a path between the source and another node. Path removal also uses the broadcast tree to trace the route to be removed from the multicast solution.

Chapter 4

Hybrid Approach

4.1 Hybrid Multicast Heuristic

We conclude from our analysis and evaluation work that none of the heuristics we have considered can provide the “cheapest” multicast solutions in all networks for all sizes of multicast groups. They either take too long to find their solutions or are vulnerable to generating unacceptable solutions that depend on the network topology and/or the multicast topology. We propose that by combining heuristics of acceptable time complexity that can be efficiently integrated, the resulting Hybrid [12] will generate solutions that are predominantly cheaper than SPTs for all network topologies, for all multicast group sizes.

Kompella et al generates good solutions but has a high order of time complexity. As suggested by Widjono [62] this can be ameliorated by reducing the granularity of the arbitrary delay bound, but at the cost of compromising the algorithms accuracy. The Crawford variant generates poor solutions overall. We discard the Crawford heuristic because of its poor performance and the Kompella et al algorithm because its time complexity may be too high for practical.

Because the Waters and Sun heuristics generate their most efficient solutions at opposite ends of the multicast group size range their combination as a Hybrid might result in a heuristic of acceptable time complexity that produces solutions of significantly improved efficiency over shortest path trees. The absolute guarantee of minimal efficiency can be made if Dijkstra’s algorithm is included in the Hybrid to cater for the rare instances where both Sun and Waters produce solutions that are more expensive than the shortest path tree.

Integration of the three heuristics is simple. All three calculate Dijkstra’s shortest path tree for delay which is extended for the second stage of the Waters heuristic. The Sun heuristic calculates Dijkstra’s shortest path tree for cost, a task which can be conducted simultaneously with the delay calculation. Once the trees have been obtained for each method their costs can be easily calculated and the cheapest tree selected as the solution.

The Hybrid heuristic procedure is as follows :-

- Execute a modified version of the Waters heuristic that returns the Water’s solution and Dijkstra’s SPT for the multicast (which is calculated in the heuristics first step).
- Execute the first step of Sun’s heuristic to obtain a lowest cost spanning tree to as many destination nodes in the multicast as is possible without any path breaking the arbitrary delay bound.
- If not all of the multicast nodes have been reached in the previous step combine the shortest paths to these nodes from Dijkstra’s SPT with the lowest cost spanning tree, making sure that the delay to any destination node does not break the delay bound. This is the last step of Sun’s heuristic.
- Calculate the cost of the Waters solution.
- Calculate the cost of Dijkstra’s SPT.
- Calculate the cost of the lowest cost spanning tree, with additional SPT paths.
- Select the cheapest tree from the above three steps as the multicast solution.

4.2 Evaluation of Hybrid Heuristic

4.2.1 Performance Averages

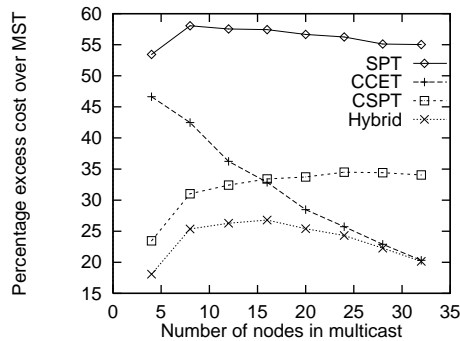


Figure 4.1: $\Delta =$ network diameter.




Figure 4.2: $\Delta =$ network diameter.

Figure 4.1 and 4.2 illustrate the cost performance of the Hybrid heuristic in comparison to CCET and CSPT for single cluster networks of 35 nodes and 100 nodes respectively. The arbitrary delay bound is set to the network diameter, our chosen “mid-point” delay bound. Figures 4.3 and 4.4 show the performance of the heuristic for the tightest delay bound (when $\Delta =$ multicast delay) and the loosest delay bound ($\Delta =$ MST delay bound), heuristic in 100 node single cluster networks. Figures 4.5, 4.6 and 4.7 illustrate the cost performance of the heuristic in hierarchical networks of 100 nodes for the three arbitrary delay bounds we used.

4.2.2 Specific Multicast Comparisons




Figure 4.3: $\Delta =$ multicast delay bound




Figure 4.4: $\Delta =$ MST delay bound

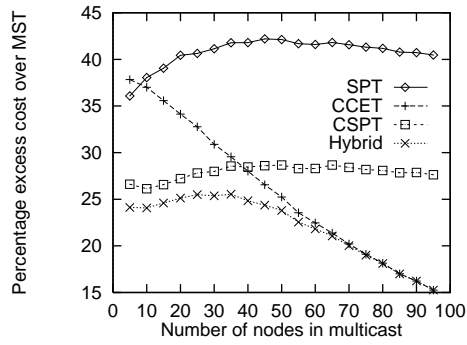


Figure 4.5: $\Delta =$ multicast delay bound.




Figure 4.6: $\Delta =$ network diameter

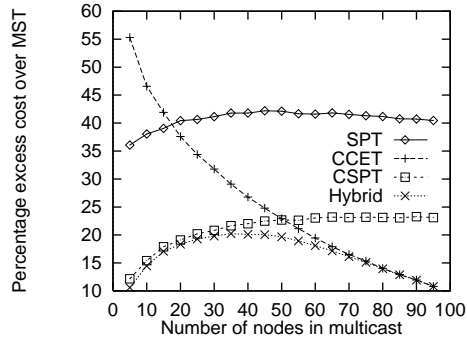


Figure 4.7: $\Delta =$ MST delay bound.

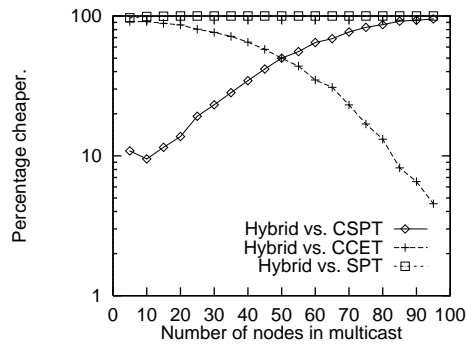


Figure 4.8: Hybrid cost saving




Figure 4.9:

In all examples the Hybrid outperforms or equals both CCET and CSPT; the improvement is particularly noticeable in smaller networks or where the multicast group size is small. This occurs because both CCET and CSPT are their most volatile in these cases whereas the Hybrid is able to select the better solution of either heuristic. As the network size increases there is a greater coincidence between the Hybrid and its constituent heuristics. Figure 4.8 illustrates the percentage of times the Hybrid is cheaper than CCET, CSPT and SPT. The SPT function in the Hybrid only achieves cheaper solutions in 0.33% of the cases.

4.2.3 Network Load and Multicast Failures




Figure 4.10: Network load and multicast failures

As noted in section 3.7.3 the SPT uses up network resources faster than CCET or CSPT. We see, in figure 4.10 that as we would expect the Hybrid consumes much the same resources as its constituent heuristics.

4.2.4 Multicast Tree Stability and Dynamic Groups

Heuristic	Hybrid	Hybrid using sCSPT
% multicast tree reconfigurations	85.13	93.74
Average % of paths changed per reconfiguration	28.79	37.9

Figure 4.11:

The Hybrid heuristic is prone to reconfigure the multicast tree as nodes join and leave the multicast group. This happens for either of the following reasons. The Hybrid may switch the heuristic used for the solution when a node joins or leaves the multicast group, thus calculating an entirely “new” solution for the “new” multicast group. On the other hand, the Hybrid may not need to switch heuristics, but might already be using CSPT which is prone to reconfiguration as the multicast tree grows. Table 4.11 gives the percentage of multicast trees that were reconfigured at least once during their growth, and the average number of path changes per reconfiguration. We also include results for a variant of the Hybrid that uses the sCSPT heuristic (the static version of CSPT), which reduces the number of reconfigurations due solely to the use of CSPT. These results were obtained using the save evaluation networks and multicast groups used in Section 3.8.

Chapter 5

Application of the Heuristics

5.1 Characteristics Required of an Heuristic

To be of practical use for multimedia applications in high speed networks, a bounded delay low-cost multicast routing algorithm should aim to have :-

- Low time complexity, since route calculation may have to be performed in “real-time” at call set-up or as part of a distributed route calculation as user data crosses a network.
- Narrow variance in the efficiency of its solutions to minimise the occurrence of expensive multicast solutions being calculated.
- Uniformity of performance over all possible multicast groups.
- Minimal or no reconfiguration of a multicast tree caused by receivers joining or leaving a multicast group, because of the impact it may have on established user data flows. Switching paths in a multicast delivery tree may increase transmission jitter or cause user data packets in flows, temporarily, to arrive out of sequence. The degree of disruption to existing flows will depend on the scale of any reconfiguration. In this case, reconfiguration does not include the addition of links required to join the new receiver to the multicast tree. It only refers to links that are dropped from an existing solution and replaced by other links, because more efficient paths are to be found as a consequence of a receiver joining or leaving the multicast group.

Heuristics that do not satisfy all of these criteria may still have limited application.

5.2 Characteristics of the Heuristics Evaluated

Our work has identified a number of characteristics that appear to be generally applicable to the low-cost multicast heuristics we evaluated in Chapter 3. These characteristics, which do not apply

to the Hybrid, are :-

- The closer to the optimal, the solution of an heuristic is, the higher its time complexity.
- Single metric shortest path algorithms, such as Dijkstra's SPT, add paths to multicast destinations as delivery trees grow. This is not the case for delay constrained multiple metric algorithms, such as those we have evaluated here. The CCET heuristic, for instance, constructs a broadcast tree which it then prunes back to the multicast tree. The CSPT and CST heuristics have to remove any forward loops from their delivery trees before they can be used as multicast solutions. This characteristic of constructing the complete multicast tree before any path can be used, needs to be taken into account if any of the heuristics evaluated is being considered for implementation in connectionless networks. Unlike routers that use single metric path calculations, routers that construct delay constrained multiple metric paths will not be able to update their user data forwarding tables until the entire multicast tree has been calculated.
- The solutions of the heuristics studied are prone to wide variance. This variance can be so large that, sometimes, the solution provided by a shortest delay path tree is more efficient than that found by using an heuristic. Further, it is not possible to easily predict for which multicast/network combinations this extreme variance occurs.
- The efficiency of each heuristic solution is not uniform across the range of all multicast group sizes. In general some heuristics find efficient multicast solutions for small multicast groups, while others find better solutions for large groups.
- The heuristics we have evaluated select their solutions from search spaces that contain either broadcast trees or multicast trees. Broadcast trees include all nodes in the network. Once selected, a broadcast tree is pruned back so that its leaves are the multicast group. Multicast trees include only those nodes in the multicast group, and the intermediate nodes on the paths from the source to the multicast destinations.
 - Once calculated, heuristics that use broadcast tree based solutions are able to join new receivers to the multicast tree without having to recalculate their solution. Heuristics that use only multicast trees either have to recalculate their solutions to include the new receiver or use some other method to perform the join. Mechanisms for joining new destinations to multicast trees vary from those that use simple shortest delay paths, but ignore cost [33] to genetic algorithms that attempt to find joining paths which balance cost and delay [47].
 - Heuristics that select their solutions from a search space based on multicast trees produce more uniform results over the range of possible multicast group sizes than those that use broadcast trees.
 - Heuristics that select their solutions from a search space based on broadcast trees, that find very efficient solutions for large multicast groups appear to do badly for small multicast groups.
- Combining several heuristics into a hybrid reduces the occurrence of inefficient multicast solutions, which are limited to the worse case performance of a shortest delay path tree. However, the hybrid approach means that the multicast solution often requires reconfiguration as nodes join or leave the multicast group, because there may be jumps from one to another of the constituent heuristics.

- All the heuristics evaluated require up-to-date knowledge of the entire topology of the network to which they are applied. For this reason implementation of the heuristics will require the use of “link-state” routing protocols, such as PNNI. When implementing a protocol to support a particular multicast tree calculation, consideration should be given to the cost of constructing and maintaining the tree in comparison to the network resources the chosen heuristic may save by its use[61].

5.3 Combining Heuristics, Multicast Types and Network Types

As discussed in Chapter 2, we identified two extreme types of multicast groups and three types of network :-

- Multicast group types.
 - Heavyweight; which will tend to have little or no changes of group membership. Closed multicast user groups are heavyweight.
 - Lightweight; which may have considerable changes of group membership. Open multicast user groups are lightweight.
- Network types.
 - Static: connections are reconfigured within the network.
 - Hard state: connections are expensive to reconfigure.
 - Soft state: paths reconfigure themselves at relatively low cost.

We also note that multicast tree reconfiguration will be minimised or eliminated where multicast group membership is (almost) static or a single, broadcast tree based, heuristic is used for the tree calculation. From our observations we conclude that the heuristics evaluated might satisfy a number of different low cost delay bound multicast routing scenarios, as summarised in figure 5.1.

- Heavyweight multicasts are (almost) static and so might best be supported by a hybrid heuristic, since this will calculate the lowest cost solution for the multicast tree. Reconfiguration is only likely to occur if the multicast group membership changes. The level of disruption to user data flows due to any rare reconfiguration caused by nodes joining or leaving multicast groups would need to be assessed.
- A lightweight multicast may have extremely dynamic membership behaviour. For this reason, in a hard state network the hybrid heuristic approach is inappropriate. A single heuristic, based on a broadcast search space, might be the best solution because it will not reconfigure the multicast tree as receivers join or leave the multicast group. The level of risk associated with the occurrence of inefficiency spikes of a single heuristic would need to be assessed.
- A lightweight multicast in a soft state network could be supported by a hybrid heuristic. Although the multicast tree will probably require reconfiguration quite often, the costs of doing so are unlikely to be high.

network type	hard state (e.g. connection oriented)	soft state (e.g. connectionless)	static (e.g. pre-configured)
multicast type			
heavyweight (sender join)	Hybrid	Hybrid	Hybrid or SuperHybrid
lightweight (receiver join)	sCSPT	sCSPT or Hybrid	not applicable

Figure 5.1: Combining heuristics, network types and multicast types

- In the case of static routing or the calculation of permanent connections, such as permanent virtual circuits in ATM networks, we would propose the extension of the hybrid to a super hybrid that includes a larger variety of multicasting heuristics than we have in the hybrid discussed here. It is unlikely that time complexity would be a predominant consideration where multicast solutions are calculated “off-line” and downloaded into the network. If time complexity were a significant contribution to the process of creating or modifying multicast trees, then the hybrid heuristic as previously proposed would suffice.

For the hybrid we would choose the Hybrid, described in Chapter 4, as it generates efficient multicast trees and has an acceptable time complexity. The choice of a single, broadcast tree based heuristic would be sCSPT. This is because it generates fairly uniform solutions of reasonable efficiency, across the entire range of multicasts, and it has an acceptable time complexity. Although sCSPT, like all the single heuristic solutions, is prone to generating spikes of inefficient solutions, the percentage of times it does so is fairly low. The other heuristics have not been chosen either because of their time complexity or because their performance is not sufficiently uniform across all possible multicasts.

5.4 Application of the chosen Heuristics in Multicast Routing Protocols

Clearly, neither of the chosen heuristics will be applicable to routing protocols that do not use some form of link state data for route calculation. The heuristics need to know the topology of the network, even if it is aggregated in places, to calculate their solutions. For this reason implementation of either heuristic in protocols such as DVMRP is not possible. The heuristics cannot be implemented in the shared tree protocols CBT and PIM either, because both these protocols use their underlying unicast routing methods, irrespective of what types they may be, to establish a shared multicast tree. PIM and CBT do not perform any multicast tree calculation; their primary purpose is the reduction of the volume of state data and route computation in comparison to protocols such as DVMRP and MOSPF. However, PIM and CBT could provide the basic framework for shared low-cost multicast trees. Using our techniques would require the cores or RPs to calculate the paths to the receivers (once they have received the join request). Although this is a significant departure from the philosophy of PIM and CBT, we believe it is worth considering. If explicit routing from the cores/RPs were also introduced, the scaling properties of MOSPF and DVMRP would, to some extent, be avoided. The other protocols discussed in Chapter 2, with the

exception of ST2+, Tag Switching and Label Swopping, are all based on link state protocols and so are candidates that, subject to any necessary modifications, might be able to use the Hybrid and sCSPT heuristics. The ST2+ protocol, Tag Switching and Label Swopping, like PIM and CBT, depend on an external unicast route calculation mechanism. However, unlike PIM and CBT, the purpose of these methods is not, primarily, to minimise the volume of state data used by multicast protocols, and so they might also be able to use an external multicast route calculation, such as the Hybrid or sCSPT heuristics.

MOSPF as a single link metric routing protocol, may reserve networks resources, if used with a protocol, such as RSVP. If the protocol were to be extended to cater for QoS, as has been proposed by Zhang et al in [64], then both of the chosen heuristics could be used within MOSPF areas. Although the heuristics could also be applied across the inter-area backbone of MOSPF Autonomous Systems, the calculation of remotely sourced multicast trees would need to be altered to use forward path link state data, rather than the reverse path data currently used. This problem could be overcome by modifying MOSPF to include “come from” link metrics in backbone advertisements, rather than relying on the OSPF backbone summary link state advertisements [36], or by means of explicit route calculation. However, the current multicast paradigm used by MOSPF is receiver initiated, in which destinations may join and leave the multicast group arbitrarily. If the Hybrid heuristic is used under these circumstances, then the multicast tree may be continually reconfigured, which could have a serious impact on the data flow it is carrying. On the other hand, the sCSPT heuristic maintains a stable multicast tree irrespective of the multicast groups join/leave behaviour. For this reason it would be suitable as a route calculation method for a QoS MOSPF network. The Hybrid heuristic might be used in a QoS MOSPF network if features such as explicit routing and pinning [64] were introduced in conjunction with heavyweight multicasting.

The ST2+ protocol specifies both source initiated, destination initiated and source and destination initiated multicast tree construction. Although the protocol does not mandate the routing calculation it uses to construct paths, since it relies on an external unicast route calculation, it could do so. The support of ST2+ for both heavyweight and lightweight multicasting means that it could use a route calculation based on either the Hybrid heuristic (heavyweight, source initiated) or the sCSPT heuristic (lightweight, source or destination initiated). This argument applies to Tag Switching and Label Swopping architectures, although as these support explicit routing, and hence heavyweight multicasting, the Hybrid might be a more appropriate choice. This would be particularly so if the architectures were applied to ATM networks.

The PNNI specification does not mandate how routes through ATM networks are to be calculated, although it does illustrate how they might be by giving a sample route generation algorithm. Route calculation is multiple metric and based on the exchange of link state data, both of which are requirements for delay bound low-cost multicast routing algorithms. However, PNNI (and UNI) specifies that multicast connections are established by first connecting the source to a single destination and then joining additional destinations to the connection by means of an “add party” request. This mechanism is in direct conflict with the characteristics of some delay bound low-cost multicast routing algorithms which reconfigure multicast trees as they grow. They can only avoid this reconfiguration if all the multicast destinations are known when the multicast tree is calculated. For this reason the more efficient of the heuristics we have evaluated are unsuitable for use by PNNI in its current form. This argument applies to B-ISDNs which also construct their multicast connections using an add party mechanism. Given the evolutionary nature of ISDN architectures neither of the heuristics proposed may, initially be of use, particularly if B-ISDN uses dynamic

two-hop alternate routing. On the other hand, if evolving ATM networks use dynamic routing for arbitrarily connected switches, and support for both heavyweight and lightweight multicasting becomes necessary (as is likely), then both the Hybrid and sCSPT heuristics provide an efficient means of calculating low-cost delay bound multicast trees. In the interim, while PNNI and B-ISDN use add party mechanisms for receivers to join multicast trees, if the group membership is known by the sending router prior to the establishment of a multicast tree, the join paths could be taken from a pre-calculated multicast tree. Such a procedure would enable the use of heuristics, such as the Hybrid, in B-ISDN and PNNI networks.

A potential application of the Hybrid in any link state network is in the construction of “permanent” multicast trees for closed user groups, where the sender wants the resources in the network to be reserved, for all receivers, even if they are not active. Receivers may join and leave the multicast as they wish, but their paths from the sender would remain intact.

The Hybrid and sCSPT heuristics could be used in a shared tree protocol, but the savings in state data, required to construct and maintain the tree, would not be as great as that achieved by PIM or CBT. The heuristics require multicast group membership link state data to be maintained by all routers that perform tree calculation, as is the case with both MOSPF and PNNI but not the case for CBT and PIM. However, with an appropriate explicit routing or centre based route calculation, shared trees using either the Hybrid or the sCSPT heuristics need not require as much state data to be used as is required by MOSPF or DVMRP. Where either of the heuristics is used in a shared tree protocol, the arbitrary delay bound, Δ , has to span the diameter of the multicast group rather than the height, as is the case for source based multicast trees. This is because all user data goes via the centre of the tree. Some evaluation work has been carried out to assess the savings that the CEPT heuristic used in a shared tree might have over CBT, which is documented in [40]. This work remains to be carried out for the Hybrid and sCSPT.

In summary, the Hybrid heuristic is more applicable to heavyweight multicasting protocols that support closed multicast user groups. The sCSPT heuristic is better suited to lightweight multicasting protocols, that support open multicast user groups.

Chapter 6

Conclusions and Further Research

We have evaluated the heuristics using both flat networks and networks constructed from clusters interconnected via backbone networks. These have been both densely and sparsely connected, large and small. The multicast groups used range in size from small to large, across the networks. We chose the network diameter as our primary arbitrary delay bound.

We have compared the average excess cost performance of the heuristics against a benchmark (MST). We have also compared the performance for individual multicasts against each other. The variance in the efficiency of the heuristics solutions has also been examined. Finally, we have assessed the efficiency of the heuristics as network load increases, and how often and how much the topology of multicast trees changes as they grow.

Through our evaluation work we have identified problems of time complexity and performance variability in heuristics that have been proposed to calculate low-cost multicast trees that are bound by an arbitrary delay. By combining appropriate heuristics we propose a hybrid (which we call the Hybrid) that produces efficient solutions within an acceptable order of time complexity, for all multicast group sizes. Our evaluations indicate that the Hybrid performs well for both single cluster and hierarchical networks.

An important result of this work, and a departure from current routing solutions, is the integration of several heuristics which are individually prone to “spikes” of inefficiency (as might be expected in an heuristic approach) into a hybrid that generates efficient solutions for all multicast groups.

We have also identified the need to develop multicast heuristics that calculate stable multicast solutions, where multicast group membership is dynamic. We have found that very efficient heuristics suffer more multicast tree reconfiguration than less efficient heuristics, as receivers join or leave a multicast group. If the efficient heuristics are restrained, so that they do not reconfigure their multicast trees, then they produce very inefficient solutions for small multicast groups. In order to provide a reasonably efficient heuristic we constrain Sun’s heuristic [50], to obtain the sCSPT heuristic, that will not reconfigure multicast trees as members join or leave the multicast group.

Both heuristics require full knowledge of the network topology, which makes them more amenable

for implementation in connection oriented networks, where they can be implemented as extensions to a unicast SPT route calculation. However, connection orientation is likely to be the communication paradigm required for reliable quality of service routing.

In assessing how the heuristics can be applied, we have identified two classes of multicast types; heavyweight and lightweight. We have also examined existing and proposed multicasting protocols in order to identify where the Hybrid and sCSPT heuristics can be best applied. The Hybrid heuristic is applicable to heavyweight multicasting for closed user groups, in robust networks, whereas the sCSPT heuristic is suitable for lightweight multicasting for open groups, in less robust networks. With appropriate extensions, the Hybrid is more suited to architectures such as PNNI (and, possibly, B-ISDN and ST2+), and the emerging Internet tag switching and label swapping architectures. The sCPST heuristic is applicable to most multiple metric link state routing protocols, including the potential QoS SR MOSPF.

Finally we have considered how the heuristics could be used in shared tree routing protocols, and consider this a viable application area.

6.1 Further Research Issues

The work to date leads into the more practical issues of implementing the heuristics in real networks, such as QoS SR MOSPF, etc, which is for further study as QoS SR protocols are evolved through the working groups of the IETF.

The application of the heuristics to shared trees is a second area of more practical work. While the use of the heuristics in protocols such as CBT and PIM is precluded, these protocols provide a basic framework from which low-cost delay bound shared trees might evolve. There is a need, in this area, to study how the heuristics might construct shared trees with multiple centres.

A further area of continued research is the study of join/leave mechanisms for the Hybrid such that multicast tree reconfiguration is minimised. Since the problem the Hybrid is attempting to solve is NP-complete, work has been undertaken by Salmon [47] to investigate the use of non-deterministic algorithms to find join paths to existing multicast trees. This approach, which is based on genetic algorithms, may use the search space of an existing Hybrid solution as a starting inwhich to find join paths.

Other application areas that may benefit from this work are the calculation of load sharing paths in networks, and shared trees. These are areas for further research.

Bibliography

- [1] A. Alles. ATM Interworking, May 1995.
- [2] P. Almquist. Type of Service in the Internet Protocol Suite., July 1992.
- [3] A.J. Ballardie. *A New Approach to Multicast Communication in a Datagram Internetwork*. PhD thesis, University College London, 1995.
- [4] A.J. Ballardie, P.F. Francis, and J. Crowcroft. Core based trees. *Computer Communications Review*, 23(4):85–95, 1993.
- [5] R. Bolla, P. Castelli, F. Davoli, and M. Marchese. Dynamic Distributed Two-hop Alternate Routing in ATM Networks. In Demetres Kouvatsos, editor, *Fifth IFIP Workshop on Performance Modelling and Evaluation of ATM Networks*, 21st-23rd July 1997.
- [6] B. Cain, S. Deering, and A. Thyagarajan. Internet Group Management Protocol, Version 3 - draft-ietf-idmr-igmp-v3-00.txt, November 1997.
- [7] CCITT. Recommendation I.121 (04/91) - Broadband Aspects of ISDN.
- [8] CCITT. Recommendation I.325 (03/93) - Reference configurations for ISDN connection types.
- [9] CIP Working Group, editor C.Topolcic. Experimental Internet Stream Protocol, Version 2 (ST-II), October 1990.
- [10] D.D Clark. The Design Philosophy of the DARPA Internet Protocols. In *Proceedings of ACM SIGCOMM'88*, 1988.
- [11] J.S. Crawford. Multicast Routing: Evaluation of a New Heuristic. Master's thesis, University of Kent at Canterbury, 1994.
- [12] J.S. Crawford and A.G. Waters. A Hybrid Approach to Quality of Service Multicast Routing. In Demetres Kouvatsos, editor, *Fifth IFIP Workshop on Performance Modelling and Evaluation of ATM Networks*, 21st-23rd July 1997.
- [13] Y.K. Dalal and R.M. Metcalfe. Reverse Path Forwarding of Broadcast Packets. *Communications of the ACM*, 21(12):1040–1048, December 1978.
- [14] D.W. Davies, D.L.A. Barber, W.L. Price, and C .M. Solomonides. *Computer Networks and their Protocols*. John Wiley and Sons, 1981.
- [15] S. Deering. Host Extensions for IP Multicasting. RFC 1112, Aug 1989.

- [16] S.E. Deering and D.R. Cheriton. Multicast Routing in Datagram Internetworks and Extended LANs. *ACM Transactions on Computer Systems*, 8(2):85–110, May 1990.
- [17] S.E. Deering, D. Estrin, D. Farinacci, V. Jacobson, C-G. Liu, and L. Wei. An Architecture for Wide Area Multicast Routing. *Computer Communications Review*, 24(4):126–135, October 1994.
- [18] L. Delgrosso, C. Halstrick, R.G. Herrtwich, and H. Stuttgen. A Transport Protocol for ST-II. GLOBECOM'92, December 1992.
- [19] J.M.S. Doar. Multicast in the Asynchronous Transfer Mode Environment. Technical Report No. 298, University of Cambridge Computing Laboratory, April 1993.
- [20] Alan Dolan and Joan Aldous. *Networks and Algorithms, An Introductory Approach*. John Wiley and Sons, 1995.
- [21] R.W. Floyd. Algorithm 97: Shortest Path. *Communications of the ACM*, 5(6):345, 1962.
- [22] Alan Gibbons. *Algorithmic Graph Theory*. Cambridge University Press, 1989.
- [23] E.N. Gilbert and H.O. Pollack. Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 16, 1968.
- [24] R. Guerin, S. Kamat, A. Orda, T. Przygienda, and D. Williams. QoS Routing Mechanisms and OSPF Extensions, March 1997.
- [25] S.E. Hardcastle-Kille. X.500 and Domains, 1991.
- [26] C. Hedrick. Routing Information Protocol, 1988.
- [27] M. Hyman, A.A. Lazar, and G. Pacifici. Joint Scheduling and Admission Control for ATIS-based Switching Nodes. In *Proc. ACM SIGCOMM'92*, August 1992.
- [28] ITU-T Study Group 11. Recommendation Q.2931 (02/95) - B-ISDN Integrated Services Digital Network (B-ISDN) - Digital Subscriber Signalling System No. 2 (DSS 2) - User-Network Interface Layer 3 Specification for Basic Call/Connection Control.
- [29] ITU-T Study Group 11. Recommendation Q.2971 (10/95) - B-ISDN Integrated Services Digital Network (B-ISDN) - Digital Subscriber Signalling System No. 2 (DSS 2) - User-Network Interface Layer 3 Specification for Point-to-Multipoint Call/Connection Control.
- [30] S. Jamin, S. Shenker, L. Zhang, and D. Clark. Admission Control Algorithms for Predictive Real-Time Service. In *Proc. 3rd International Workshop on Network and Operating System Support for Digital Audio and Video.*, November 1992.
- [31] P. Kompella, V. *Multicast Routing Algorithms for Multimedia Traffic*. PhD thesis, University of California, San Diego, USA, 1993.
- [32] V.P. Kompella, J.C. Pasquale, and G.C. Polyzos. Multicast Routing for Multimedia Communications. *IEEE/ACM Transactions on Networking*, 1(3):286–292, 1993.
- [33] T Kuzminski. Alternatives for Multicast Routing in ATM Networks. Master's thesis, University of Kent at Canterbury, 1996.

- [34] J.M. McQuillan, I. Richer, and E.C. Rosen. The New Routing Algorithm for the ARPANET. *IEEE Transactions on Communications*, COM-28(5):711–719, May 1980.
- [35] J. Moy. Multicast Extensions to OSPF. RFC 1584, March 1994.
- [36] J. Moy. OSPF Version 2. RFC1583, March 1994.
- [37] J. Moy. OSPF Version 2. RFC2178, July 1997.
- [38] J. Moy. Private e-mail correspondence, 1997.
- [39] D. Oran. OSI IS-IS Intra-Domain Routing Protocol, RFC1142, February 1990.
- [40] S. Parkinson. Multicast Routing in the Internet: Evaluating Proposed Routing Mechanisms. Master’s thesis, University of Kent at Canterbury, 1995.
- [41] C Partridge. A Proposed Flow Specification, September 1992.
- [42] T Pusateri. Distance Vector Multicast Routing Protocol - draft-ietf-idmr-dvmrp-v3-05, October 1997.
- [43] Y. Rekhter, B. Davie, E. Rosen, and G. Swallow. Cisco Systems’ Tag Switching Architecture Overview, February 1997.
- [44] E.C. Rosen, A. Viswanathan, and R. Callon. A Proposed Architecture for MPLS. draft-ietf-mpls-arch-00.txt, August 1997.
- [45] K.W. Ross. *Multiservice Loss Models for Broadband Telecommunications Networks*. Springer-Verlag, 1995.
- [46] H.F. Salama, D.S. Reeves, I. Vinitos, and Tsang-Ling. Sheu. Evaluation of Multicast Routing Algorithms for Real-Time Communications on High-Speed Networks. In *Proceedings of the 6th IFIP Conference on High-Performance Networks (HPN’95)*, 1995.
- [47] M Salmon. Survivable Multicast Routing for ATM Networks. Master’s thesis, University of Kent at Canterbury, 1997.
- [48] ST2 Working Group. Internet Stream Protocol, Version 2 (ST2) Protocol Specification - Version ST2+, August 1995.
- [49] W. Stallings. *ISDN and Broadband ISDN*. MacMillan Publishing Company, 2nd edition, 1992.
- [50] Q. Sun and H. Langendoerfer. Efficient Multicast Routing for Delay-Sensitive Applications. In *Second Internatiopnal Workshop on Protocols for Multimedia Systems (PROMS’95)*, pages 452–458, 1995.
- [51] A.S Tannenbaum. *Computer Networks*. Prentice-Hall International, 1989.
- [52] The ATM Forum Technical Committee. *ATM User-Network Interface Specification, Version 3.1*. The ATM Forum, September 1994.
- [53] The ATM Forum Technical Committee. *ATM User-Network Interface Specification, Version 4*. The ATM Forum, July 1996.

- [54] The ATM Forum Technical Committee. *Private Network-Network Interface Specification Version 1.0*. The ATM Forum, March 1996.
- [55] D. Waitzman, C. Partridge, and S.E. Deering. Distance Vector Multicast Routing Protocol, 1988.
- [56] D. Wall. *Mechanisms for Broadcast and Selective Broadcast*. PhD thesis, Stanford University, 1980.
- [57] A.G. Waters. A New Heuristic for ATM Multicast Routing. In *2nd IFIP Workshop on Performance Modelling and Evaluation of ATM Networks*, pages 8/1–8/9, July 1994.
- [58] A.G. Waters. *Multi-Party Communication over Packet Networks*. PhD thesis, University of Essex, UK, 1996.
- [59] A.G. Waters and J.S. Crawford. Low-cost ATM Multimedia Routing with Constrained Delays. In *Third COST 237 Workshop on Multimedia Telecommunications and Applications*, November 1996.
- [60] B.M. Waxman. Routing of Multipoint Connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622, 1988.
- [61] L. Wei and D. Estrin. The Trade-offs of Multicast Trees and Algorithms. Technical Report 93-560, Computer Science Department, University of Southern California, 1993.
- [62] R. Widyono. The Design and Evaluation of Routing Algorithms for Real-time Channels. Technical report, University of California at Berkeley and International Computer Science Institute”, type = ”TR-94-024, September 1994.
- [63] L. Zhang, S.E. Deering, D. Estrin, S. Shenker, and D. Zappala. RSVP: A New Resource ReSerVation Protocol. *IEEE Network*, 8(17), September 1993.
- [64] Z. Zhang, C. Sanchez, B. Salkewicz, and E. Crawley. Quality of Service Extensions to OSPF, June 1996.
- [65] W. Zheng and J. Crowcroft. Quality-of-Service Routing for Supporting Multimedia Applications. *IEEE Journal on Selected Areas in Communications*, 14(7):1228–1234, 1996.
- [66] Q. Zhu, M. Parsa, and J.J. Garcia-Luna-Aceves. A Source-Based Algorithm for Near-Optimum Delay-Constrained Multicasting. In *Proceedings of IEEE INFOCOM’95*, pages 377–385, 1995.