

# Workshop Note: Artificial Intelligence for Network Routing Problems

Steven Willmott\* and Boi Faltings†

May 18, 1999

## Abstract

Rapidly growing network usage, the deployment of heterogeneous technologies and diverse service requirements are all factors fuelling the drive towards the improvement of routing techniques. This paper aims to give an overview of network routing problems, the current state of technology and previous Artificial Intelligence contributions to the problem area.

It is hoped that this will help stimulate discussion at the AiDIN'99 workshop held at AAAI'99 in Orlando, Florida.

## 1 Introduction

Communications networks require complex resource management mechanisms to stay operational. Network routing techniques are the key element in this resource allocation. A large number of routing technologies have been developed and many are still in service. The diversity of deployed networks and, above all, the rapid pace of technological change means that there remains a continuing need for new solutions. Three of the major driving forces behind the requirements for new routing techniques are:

- The need for Quality of Service guarantees,
- Newly deregulated telecommunication industries,
- Explosive growth in network size and usage.

---

\*Laboratoire d'Intelligence Artificielle (LIA), Department Informatique (DI), Swiss Federal Institute of Technology (EPFL), IN (Ecublens), CH-1015 Lausanne, Switzerland. e-mail: (willmott@lia.di.epfl.ch).

†Laboratoire d'Intelligence Artificielle (LIA), Department Informatique (DI), Swiss Federal Institute of Technology (EPFL), IN (Ecublens), CH-1015 Lausanne, Switzerland. e-mail: (faltings@lia.di.epfl.ch).

Despite numerous predictions of a bandwidth glut ([Smith, 1999] among others), resources used still need to be carefully managed. The increased volumes of data flowing across modern networks mean that mismanagement can very quickly result in bottlenecks and potentially catastrophic cell loss. Above all, the complexity of today's network technology and the operational strain it is put under generates more need than ever for innovative and "intelligent" network management solutions.

Routing problems have been widely studied in Artificial Intelligence and Computer Science in an abstract context. The graph theoretic nature of the underlying problem has often been helpful in abstracting away technical detail. Unfortunately the recent rapid pace of technological advancement and the increasing obscurity of networking terminology is making harder for researchers without the relevant background to make progress in the area. This informal review gives a basic outline of the problems, key network technologies and solution techniques in the area. The final third of the paper then goes on to outline some of Artificial Intelligence (AI) contributions to the study of routing.

Note that this review is exclusively focused on routing problems and we have tried to avoid digressing into related areas (such as network topology design, call scheduling, connection acceptance control capacity planning, etc.). Works containing good reviews for the applications of Artificial Intelligence techniques to network problems in general include: [Liebowitz, J. and Prerau, D. S., 1995], [Kumar and Venkataram, 1997], and [Weihmayer and Velthuijsen, 1998].

It should also be noted that this review is a preliminary version and far from being comprehensive. The classifications used to divide up the technologies and problems are also not the only possible groupings and hence should be treated only as a guideline. We hope, however, that the interested reader will be able to find more detailed information in the referenced material.

## 2 A Generic Problem Description

Perhaps the easiest way to understand the issues in a domain is to relate them to a generic problem definition. The definition given in Section 2.2 is a slight variation on the one given to encourage discussion at the AiDIN workshop. Following sections will then relate characteristics of the most important network types to this model.

### 2.1 Terminology

Before beginning it is perhaps helpful to clarify some of the terminology which will be used throughout this paper:

- **Routing mechanism, routing algorithm, routing method and routing protocol:** these terms are all very similar but have slightly different meanings:

- Routing method, routing algorithm: these terms refer to the computational method applied to finding a route or several routes. In this paper the two terms are treated as interchangeable.
- Routing mechanism: a routing mechanism is the apparatus deployed in the network for routing calls which may include: the routing algorithms, strategies for decomposing the network into more manageable pieces (e.g. a routing hierarchy) and methods for updating network state information.
- Routing protocol: a routing protocol was originally taken to mean the protocol used in a network to exchange network state information used for routing (e.g. RIP messages, hello messages) but is often taken to mean something similar to a routing mechanism.
- **Demand:** a demand is taken to mean a request by a user or system for a data transmission service, be it the sending of a single packet, establishment of a connection, transfer of a file or establishment of a multicast tree.

Note that in some approaches the routing algorithm is strictly distinct from the routing mechanism: the algorithm generates routing solutions, the routing mechanism applies the generated routes when forwarding packets or allocating circuits.

## 2.2 Problem Definition

The problem to be solved is defined in terms of the available infrastructure and the required results. The available infrastructure is a communications network comprising:

- 10; 100; 1000; 10,000; 100,000; 1,000,000; 10,000,000 or 100,000,000 nodes (e.g.: routers, switches).
- Communication links connecting these nodes such that:
  - The links are *evenly* / *unevenly* distributed throughout the network (i.e. varying distributions of links of entering and leaving nodes).
  - All links have *equal* / *very different* capacity, delay, cell loss and other transport characteristics.
  - The network graph is *densely* / *sparsely* connected (i.e. varying the ratio of links to nodes across the network).
- Traffic in the network is:
  - Connection-oriented (or at least individual flows can be identified and routed): resources are allocated on a demand-by-demand basis at set-up time.
  - Connection-less: traffic is managed on a packet-by-packet basis.

Given this generic infrastructure a routing mechanism must ensure that data can travel through the network between arbitrary end points (or arbitrary sets of endpoints in the case of multicast). The resulting network should be able to support some subset of the following:

- Heavily and lightly loaded networks,
- Overload situations,
- Fluctuation of traffic patterns (over the long and short term),
- Prevention of routing oscillations and loops,
- Fast response to resource demands (although this may not mean that a route finding process itself has to be on-line and real time),
- Quality of Service (QoS) guarantees (or at least the high probability of good service),
- The ability to plan for traffic which is booked on the network ahead of the actual demands start time (introduces a temporal aspect),
- Function in a multi-provider environment where some network state information may not be available outside provider domains.
- Multicast and/or point-to-point traffic.

### 3 Today's Networks

The generic problem covers many types of networks (indeed in some cases the requirements may be contradictory). This section groups together certain sets of features to identify combinations which are common in today's networks. There are two main classes of networks: circuit-switched and packet switched. Circuit switched networks are descended from telephone networks and discussed in Section 3.1. Packet switched networks (discussed in Section 3.2) are further sub-divided into two classes: those operating in a connection-less transfer mode (such as IP based Internet) and those operating in a connection-oriented transfer mode (such as ATM networks).<sup>1</sup>

---

<sup>1</sup>Another class of routing problems is in Optical networks where routing involves assigning wavelengths to each path so that paths sharing an edge have non-interfering (different) wavelengths. The goal is minimise the number of wavelengths assigned in the network. Optical networks will not be discussed further in this paper, see [ Ramaswami, R. and Sivarajan K. , 1998] for an overview.

### 3.1 Circuit Switched Networks

Telephone networks have been in use for approximately a hundred years and a large majority of the deployed network infrastructure is still based on the Public Switched Telephone Service (PSTN). The evolution from telephone calls to carry voice to connections carrying data has left the underlying principles essentially unchanged. The first generation of ISDN (Narrowband ISDN), for example, still uses similar multiplexing techniques to the Time Division Multiplexing (TDM) used in telephone networks.

On receipt of a demand for communication between two nodes A and B in the network, a circuit switched network sets up a dedicated circuit between the two nodes. This circuit (variously known as a call or connection) is only removed when the communication ends. During the duration of the communication resources remain reserved. The routing problem has several characteristics which facilitate the construction of routing mechanisms:

- Often all calls have identical basic resource requirements (all telephone calls in a network generally have the same resource needs). This uniformity makes it much easier to collect statistics about call patterns, it also makes calculating routing tables computationally less complex (since, on a given link, two circuits are interchangeable).
- There are many calls. Large call volumes make it easier to collect call statistics since consistently large samples iron out fluctuations due to individual users. This statistical data can be used to construct relatively accurate traffic models.
- There is no interaction in resource use between established circuits. Since calls are served by reserved channels there is no resource sharing after call setup. This means that new calls can be allocated in the network without fear of reducing the quality experienced by users already connected.

Although the introduction of digital services such as ISDN has complicated this picture (the assumption that all calls require equal resources no longer holds for example), circuit switched networks still represent the best understood domain for the routing problem.

The main problems which need to be solved for circuit switched networks are related to the sheer volume of demands they need to handle and to unexpected fluctuations in demand. Briefly there are two main types of routing performed in circuit switched networks - *alternate routing* and *adaptive routing*. For alternate routing, a set of routing tables for the network is generated off-line using optimisation algorithms, shortest path computation and search heuristics. These statically generated routing tables are then distributed to network nodes. Alternate routing essentially provides several alternative route choices for each source-destination pair. The routing algorithm steps through these one by one until it finds the first which can service the demand. If the algorithm runs out of alternatives the call is rejected.

An improvement on the static alternate routing mechanism is to take network congestion into account. The Dynamic Non-Hierarchical Routing (DNHR [Ash and Huang, 1993]) and Real Time Network Routing (RTNR [Ash et al., 1991]) routing algorithms both are called *adaptive*. These algorithms allow a limited search for alternative routes when congestion blocks the pre-calculated routes.

For detailed (though somewhat dated) discussion of the routing issues in circuit switched networks see [Girard, A., 1990].

## 3.2 Packet Switched Networks

In packet switched networks, instead of establishing circuits, data is transferred in discrete packages called datagrams, packets or cells. There are two main classes of packet switched networks: connection-less and connection-oriented. These two types are discussed in sections Section 3.2.1 and Section 3.2.2 respectively.

### 3.2.1 Connection-less Networks

Data transmission in connection-less mode is handled on a per packet basis and in the network there is no explicit relationship established between packets belonging to a single demand.

By far the best known connection-less network is the Internet based on the Internet Protocol (IP). A second connection-less packet switched network is Appletalk. The Internet is also the fastest growing network: in late 1982 there were several hundred hosts connected to the Internet, this number has been growing almost exponentially ever since. According to Bellcore's continuously updating estimates<sup>2</sup> there are now over 56 million hosts connected, an increase of 13 million in the first 4 months of this year.

Along with this large size and rapid growth, comes great diversity in usage. The technology which IP and other connection-less networks are based on remains relatively simple:

- Traffic is generally best effort only (there is no guarantee of timeliness of delivery or even of delivery itself). There is no reservation of resources as in circuit switched networks or connection-oriented networks where the connection is reserved for the duration of the call.<sup>3</sup>
- Packet delivery may be unordered. End systems are responsible for re-assembling and re-ordering the packets to make up the data data sent for each demand (e.g. a gif file, e-mail message, video stream etc.).
- Packets belonging to different users and demands are mixed in the network and share the resources. New demands can affect the quality of transmission experience by other demands.

---

<sup>2</sup>See the web-site at <http://www.netsizer.com/>.

<sup>3</sup>Note, however, that recent years have seen the development of reservation protocols which run over IP networks, such as RSVP [Braden et al., 1997], YESSIR [Pan and Schulzrinne, 1997] and Differentiated Services [Nichols et al., 1997].

Connection-less networks primarily use distributed hop-by-hop routing. Each router (node) in the network has a table which looks something like the one shown in Figure 1. The addresses on incoming packet headers are matched in the table to determine the next hop along the way. The next router reached does a similar look up and so on, until the packet eventually reaches its destination. The task of any routing mechanism is to update the value in the table to ensure that packets reach their destination quickly and without getting stuck in (or causing) congestion.

Address Prefix	Next Hop
128.8.0.0/16	128.8.1.10
128.8.70.0/24	128.8.70.1
128.8.74.0/24	128.8.70.1
128.8.74.0.24	128.8.74.1
128.8.14.0/24	128.8.74.1
.	.
.	.

Figure 1: A simple IP routing table. The router matches the packet address in the left-hand column and identifies the next router to send the packet on to in the right-hand column.

RIP [Hedrick, 1988], OSPF [Moy, 1989], IGRP [Hedricks, 1991] and OSI's IS-IS (dual) [Callon, 1990] are the most commonly used routing protocols in the Internet today, with OSPF [Moy, 1989] the most widely used. These protocols are all known as Interior Gateway Protocols (IGPs).




Figure 2: Three Autonomous Systems linked using EGP and BGP protocols

Large networks such as the Internet are made up of clusters of routers and hosts called “autonomous systems” (AS) which are then linked together to cover the whole network. The definition of an AS is somewhat vague but can be thought of as a set of routers which are linked by an Interior Gateway Pro-

to col (IGP) to do their routing. The ASs themselves are linked together by another type of routing protocol - Exterior Gateway Protocols (EGPs) such as the Exterior Gateway Protocol (EGP [Mills, 1984]) and the Border Gateway Protocol (BGP [Lougheed and Rekhter, 1991]). RIP, OSPF, IGRP and IS-IS (dual) hence work *inside* an AS serving a group of hosts, whereas EGP and BGP work *between* ASs. Connection to the outside world for an AS is through routers which also support an EGP (See Figure 2). EGP was the first EGP protocol but is now being commonly replaced by BGP which was designed to remedy some of EGP's shortcomings and is now on its fourth revision [Rekhter and Li, 1998]. See [Huitema, 1995] for a good description of both EGP and BGP.

Even though the Internet is made up of many millions of nodes, routing problems are divided up by its hierarchical structure. High capacity backbones link Autonomous Systems together giving a two level hierarchy. Domains using OSPF as a IGP can add a third internal level to the hierarchy by breaking an AS up into "areas". The structure and architecture of Internet backbone networks is best illustrated by example:

- A Map of the SWITCH Swiss Academic and Research Network can be found at <http://www.switch.ch/lan/national.html> (<http://www.switch.ch/> for the main Switch site).
- <http://www.uunet.com/lang.en/network/usa.html> shows a map of UUNET US (<http://www.uunet.com/> for the main UUNET site).
- [http://www.eunet.com/maps/global\\_map\\_current.gif](http://www.eunet.com/maps/global_map_current.gif) shows EUNET's global network (<http://www.eunet.com/> for the main EUNET site).

Note that some of these networks also carry traffic other than Internet and data traffic. The subnetworks in ASs connected to the backbone vary greatly in size, capacity, deployed technology and topology. There are currently around 3000 Autonomous systems in the Internet (January 1999, [Govindan et al., 1999]), ranging in size between continent-wide and small university campuses.

A more detailed overview of the protocols used in the Internet can be found in [O'Neill et al., 1998].

### 3.2.2 Connection-oriented Networks

In networks using connection-oriented modes, a logical / virtual connection is established before data is transferred. All packet transfer is carried out in association with the established connection, guaranteeing various service characteristics for the data transfer. Examples of connection-oriented networks are: ATM, Frame Relay, SNA and X.25 networks<sup>4</sup>. The discussion here will focus on ATM networks (since these are the most up-to-date and present the most open problems).

The grouping of packets into connections has two main advantages:

---

<sup>4</sup>Note that some of these network infrastructures can also operate in a connection-less mode.



1. Connection-oriented networks can maintain properties of the data transfer which relate to the connection as whole rather than to individual packets. ATM for example can support a large set of Quality of Service (QoS) parameters, including end-to-end parameters.
2. Once the virtual connection has been established in the network, less processing per packet is required than in a connection-less network. Less processing is required because the routers along the path of the connection retain information about the connection. When packets arrive knowing how to deal with them often reduces to deciding which connection they belong to. The reduction in processing can significantly increase the throughput of the nodes in the network (and hence the speed and capacity). The throughput gains are one reason why ATM-like networks are sometimes referred to as "Fast Packet Switching" networks.

Connection-oriented networks also differ from circuit switched networks in that they are designed for the transfer of diverse data types and allow for much finer resource control (due mostly to their use of packets to effect the actual data transfer). For more details on ATM networks and related technologies see [de Pryker, M., 1995] and [Sexton, M. and Reid, A., 1997]. The main characteristics of these networks are as follows:

- Due to the reductions in packet, processing these connection-oriented networks tend to have high speed / high capacity links - making good network backbones.
- Links tend to have uniform transport characteristics. For example ATM links mostly come in set sizes: 155Mbit/s 622Mbit/s or (in the future) 2480Mbit/s capacity.
- For ATM at least, deployment is still not that wide spread, so ATM networks tend to be relatively small by Internet standards.

ATM networks are used mainly in backbone networks which are very sparse (few if any cycles) which leads to simple routing problems. Many ATM links are also used in a point to point fashion (hence with no routing problem to solve). Despite this restricted deployment the the likelihood is that with future increases in ATM deployment routing techniques for these types of networks will still be of great importance. One of the highest profile ATM deployments to date is in the Singapore One network<sup>5</sup>. [Clarke et al., 1998] describes the "James" European ATM test network which was constructed with participation of major European telecommunication companies to test ATM services. James spanned the length and breadth of Europe with around 20 nodes connected by 34 Mbit/s and 155Mbit/s links. The links to network operators given for IP networks in the previous section are also relevant to ATM deployment since many of the networks employ ATM technology in their backbones.

<sup>5</sup>See the commercial website <http://s-one.net.sg/>.

For networks which offer QoS guarantees operation allocation of a demand is broken down into four closely related activities:

1. Connection acceptance control: deciding whether or not to accept a particular connection request,
2. Route selection: selecting a route for a given demand,
3. Flow control: controlling the flow of packets associated with the connection to ensure that QoS parameters are respected and that the connection does not have a negative impact on the quality experienced by other connections,
4. Allocation of buffer space and bandwidth: reservation of the resources in the network to support the QoS requirements of the demand.

Routing needs to be adaptive to the network state [Lee et al., 1995] making standard strategies based on the optimisation of a routing metric unsuitable. [Lin and Yee, 1993] argues that the large volumes of data being carried make feedback/reactive control systems impractical - only reacting to a congestion problem once it has occurred would still result in a heavy cell loss. Routing in ATM networks is officially based upon the PNNI architecture [ATM-FORUM, 1996] which provides a complete framework for ATM routing. In practice PNNI is often only partially implemented and small ATM networks use static routing generated by human operators.

Many of the tasks performed in these high capacity networks involve setting up Virtual Private Networks (VPNs) which link several customer sites with guaranteed connections. These logical topologies consist of several point to point connections which need to be maintained for long periods of time. An increasingly important area is the integration of IP and ATM networks - principally by carrying IP traffic over ATM networks (see [Laubach and Halpern, 1998]).

## 4 Types of Routing problem

Having reviewed the major types of networks that need to be managed today, we now move on to looking at the major problem areas which divide up the space of routing problems.

### 4.1 On- and Off-line Routing

Routing processes can be run at different times during the operation of the network. The types of routing method can be classified into three broad areas: *off-line*, *adaptive (dynamic)* and *on-line*. Different update strategies are appropriate for different types of networks.

#### 4.1.1 Off-line Routing

This approach relies on statistical estimates of the traffic demands being made in the network (called traffic matrices for circuit-switched networks). Given a representation of the network topology and a traffic matrix, the demands in the traffic matrix are allocated to the network representation by an optimisation algorithm. The allocations made can then be used to generate routing tables which are inserted into the network routers. When a demand arrives in the network it can automatically be allocated a route using the pre-calculated routes stored in the routing tables. Routing methods which are effectively controlled by human operators in this way are often called “static” and routes generated called “pre-computed”.

The primary advantages of this approach are:

- Global information is available during route calculation, hence the theoretical optimum (by some measure) routing table configuration could be calculated.
- When a demand is made in the network, little calculation is needed to determine which route should be taken (this is fixed in the routing tables).
- Off-line, more powerful machines can be applied to the task, the routers in the network require less computational power.

The main disadvantages are that the resulting routing solutions are heavily dependent upon the accuracy of the traffic matrix and that calculations for large networks are generally very complex (often NP-complete by analogy to the multi-commodity flow problem [Garey, M. R. and Johnson, D. S., 1979]). Highly dynamic networks pose problems: solutions produced off-line might be very quickly out of date if the network state diverges significantly from traffic predictions.

This method is often used in telephone networks with network-wide routing calculations being run in the order of every few months. As mentioned in Section 3.1 the large call volumes and uniform traffic types in telephone networks make it possible to generate relatively accurate traffic matrices.

#### 4.1.2 Adaptive / Dynamic Routing

These “pseudo on-line” routing methods also work with fixed routing tables but have a dynamic aspect (and are therefore sometimes labelled “dynamic”). Route calculations are triggered directly by changes in the network state (such as link failure, new hosts being added etc.) and routing tables are updated automatically. For any demand arising, a route already exists (defined by the routing tables) but this may be updated over time by algorithms running continuously in the network. Our term “pseudo on-line” indicates that these techniques fall in between on- and off-line routing techniques.

Examples of this type of mechanism are all the routing protocols mentioned in Section 3.2.1 (RIP, OSPF, BGP etc.) which update routing tables in

IP routers with respect to reachability information and estimated shortest paths through the network. [Ash and Huang, 1993], [Ash et al., 1991] are examples of adaptive routing in circuit switched networks.

Enabling on-line updates to existing routing tables addresses one of the major problems with off-line routing - response to changes in the network state. It should be noted however that the update strategies applied by protocols such as RIP and OSPF are very simple and in general do not take into account congestion (experiments with ARPANET [McQuillan and Walden, 1977] showed that for packet networks this could lead to serious oscillation problems).

#### 4.1.3 On-line Routing

Arguably the most flexible routing mechanisms are those which apply on-line routing. On-line routing means that no pre-determined routes are retained in the network and whenever a demand arises in the network it is routed on-arrival, using the current state of the network to calculate the best route to take.<sup>6</sup> Theoretically on-line routing could even be applied on a packet by packet basis where each packet has its own route chosen on arrival in the network.

This type of routing is often associated with source routing (see Section 4.2.2) although it could be done in a distributed fashion. On-line routing is most responsive to the network state and is important in QoS networks which require several different parameters to be taken into account when evaluating the suitability of a route. It is used almost exclusively in connection oriented networks with resource reservation since many QoS metrics are defined over connections rather than on individual packets. If the network state is very dynamic this may be the only way to find a suitable route (See [Alles, 1995] for discussion of the use of on-line routing for ATM).

The difficulties with on-line routing are that:

- Route calculations need to be made as demands arrive, incurring delay in call setup. Deciding on each route may itself be an NP-complete task depending upon the metrics used [Wang and Crowcroft, 1996].
- Route calculations depend upon extensive network state information which may be rapidly changing and needs to be accessible from throughout the network. For large networks, information updates become very expensive to send and suffer from non-negligible delays.

For these reasons it is an open question as to whether full on-line routing will ever be feasible in large networks.

## 4.2 Distributed and Source Routing

For a given demand, routing choices can be made at various places in the network. There are three main ways of allocating which entities in a network

---

<sup>6</sup>Routing mechanisms which take into account the network state are termed: "state-based".

perform routing tasks: *distributed routing*, *source routing* and *hierarchical routing* (following [Chen and Nahrstedt, 1998]). This section owes much to a more comprehensive review in [Chen and Nahrstedt, 1998].

#### 4.2.1 Distributed Routing

In distributed routing mechanisms (sometimes known as hop-by-hop routing), the route of a packet or connection is constructed distributely by all the nodes along the eventual path. Once a packet (or in connection oriented / circuit switched networks a set-up call) arrives at a node, the node calculates the node it is best to forward the packet to. This calculation could be made on-line or using an existing routing table. The packet continues on its journey from node to node until it arrives at its destination.

The advantages of this approach are its simplicity and scalability. At no time is a global map of the network required to calculate the route - the route is generated by a series of local calculations. The disadvantages are that it is hard to apply complex heuristics in a distributed fashion and that global properties (such as the prevention of loops and adherence to end-to-end constraints such as delay) can be hard to maintain. This routing method is used in IP networks with each router containing a routing table listing the next hops for all possible destinations. On receipt a packet, the router is able to choose which of its neighbouring routers to forward the packet to.

#### 4.2.2 Source Routing

In source routing, the whole route for a connection (source routing is generally not used for individual packets) is made at the node which originates the demand. The route decided upon is then communicated to the other routers in the network in the setup call for the connection or in the packet headers.

Nodes performing source routing require access to non-local network state information to make their routing decisions (since they effectively need to have access to a global view of the network state). Source routing is often used in conjunction with on-line routing. [Alles, 1995] gives the following arguments for the use of source routing in ATM networks:

1. If hop-by-hop routing were used, each node along the path would need to evaluate available QoS across the entire network to determine the next hop.
2. Hop-by-hop routing requires a standard (identical) route determination algorithm to be used at each hop to (help) preclude the danger of routing.

A further reason for using source routing is that some QoS metrics defined for a connection are expressed as end-to-end constraints which cannot be checked in a hop-by-hop fashion (related to Alles' first point).

The main problem with source routing is scalability: as networks become larger it becomes impractical to store and manipulate sufficient network state

information in each node to obtain good routes [Guerin and Orda, 1997]. Source routing nodes will always be working with outdated information due to the propagation time of information updates in the network.

#### 4.2.3 Hierarchical Routing

In hierarchical routing, the network is broken up into clusters of nodes which are then grouped together into larger clusters and so on until the whole network is included. The hierarchy of clusters thus created can be used to distribute aggregated information about the network state to network nodes. Hierarchical routing strategies use the aggregated information to generate routes which become less detailed the further the path gets away from the source node, other nodes in the hierarchy then expand these abstract sections of the path further when necessary. The PNNI routing architecture proposed for ATM [ATM-FORUM, 1996] is the best example of a hierarchical routing mechanism for connection-oriented networks.

Hierarchical routing maintains some of the benefits of source routing in that each node in the network does have a network wide view (although some of the information is aggregated). Route computation and information storage is however shared between network nodes which reduces the scalability problems of source routing. The key challenge with hierarchical routing is being able to make useful routing decisions based on aggregated or imprecise network state information.

### 4.3 Multicast Routing

Most routing techniques address the problem of routing traffic from one single point to another point in the network (point-to-point problems). This is by far the most common requirement for routing mechanisms in current networks. However an increasing number of applications require communication between many users at once (radio or video broadcast being among the most important examples). The resulting “multicast” problem involves finding routes in the network to connect many senders to many receivers (many-to-many communication) and perhaps allowing to users join and leave the group dynamically (hence multicast solutions also require mechanisms to handle this multicast group).

The problems to be solved are significantly different to those in the point-to-point problem. The goal of multicast routing is to reduce the amount of traffic in the network by reducing the number of copies of each data packet that need to be produced. This breaks down into:

- Constructing a multicast tree connecting all senders with all receivers,
- Minimising the cost of the multicast tree. The cost of the tree is usually defined as the sum of the cost of all the network links which form edges in the tree. The cost of a link is in turn defined w.r.t a particular metric such as the available bandwidth, congestion, number of hops etc.

The problem of finding a minimal delivery tree for a many-to-many multicast session is equivalent to the graph theory Minimal Steiner Tree problem which is known to be NP-complete [Hwang and Richards, 1992] and [Hwang et al., 1992]. [Wei and Estrin, 1994] provides a useful discussion of several of the algorithms and heuristics available. Due to need to manage the multicast group much of the difficulty in routing for multicast is also in updating existing (and running) multicast trees when users join and leave rather than generating complete trees from scratch.

There are four main multicast protocols in use in the Internet today: Distance Vector Multicast Routing Protocol (DVMRP [Waitzman et al., 1988]), Multicast extensions to OSPF (MOSPF [Moy, 1994]), Core-Based Trees (CBT [Ballardie, 1997]), and Protocol Independent Multicast (PIM, see [Estrin et al., 1998] for PIM Sparse Mode).

#### 4.4 Multi Authority Routing

All the discussion so far has assumed a free flow of information between the different parts of the network. An increasingly important problem is dealing with the situation where this is not the case. With the deregulation of many telecommunications markets worldwide, the number of actors involved in call setup has increased dramatically. What were once state monopolies controlling everything from end user access down to the copper wires has become several layers of competing firms - service providers, networks providers, brokers etc.

In this environment network resources are owned by many different “authorities” but need to be made to work together to provide a coherent service to the user. The routing task is made very difficult by the fact that individual providers are unwilling to release detailed information about the state, or even topology of their internal network. This is highly sensitive information. A further major problem is that there is little or no infrastructure to support the information exchange or coordination between different service and network providers. In TMN compliant networks the TMN-X interface provides a rudimentary low level interface for synchronising the settings in routers and other network elements<sup>7</sup> In the Internet, providers effectively interact using the BGP protocol and information exchanged by domain internal IGP protocols is hidden. Each provider is in charge of one or more domains and BGP information is used to give and receive an abstract view of connectivity available through other domains. A review of the current state of play in Internet multi authority routing can be found in [Govindan et al., 1999].

The challenge is to find a way of making use of restricted information to make adequate routing decisions when passing through domains controlled by several different authorities.

---

<sup>7</sup>The TMN-X interface is, however, not specifically for routing but supports a more general set of management functions. In the routing context, it is only used to establish static (off-line computed) routing tables to cross organisational boundaries.

## 4.5 Quality of Service Routing

Routing for networks which need to support QoS guarantees is particularly challenging and is becoming more and more important with the rapid increase in applications and users requiring quality guaranteed services. The following difficulties need to be dealt with:

- QoS connections often depend on many different parameters (available bandwidth, delay, jitter, cell-loss ratio etc.). Routing decisions are correspondingly complex (often NP-complete to route just a single demand).
- The number of parameters along with the large number of different applications leads to very heterogeneous traffic types, making traffic prediction difficult.
- QoS connections are very sensitive to the network state, this sensitivity reduces the usefulness of pre-determined routes. As the network state changes, the best route to take changes rapidly with it.
- Quality parameters are almost exclusively applied to the connection / single demand level. Networks which take routing decisions on a packet by packet basis require extra mechanisms to use routing decisions to help maintain quality of service. The worries about RSVP's (the best known proposed resource reservation mechanism for the Internet - see [Braden et al., 1997]) independence from the routing process illustrate this [Huitema, 1995]. Changes in routing decisions at the IP packet level in the network would cancel RSVP reservations, switching traffic out from under the reserved path.

As noted in Section 4.1 and Section 4.2, providing QoS in a network exerts a pull towards on-line and source routing. On-line routing is useful to take into account the current network state. Source routing is useful because decisions need to be made taking into account information from non-local areas of the network (and it makes little sense to continually "re-make" these decisions at every hop).

Both on-line and source routing pose complex problems in information distribution and decision making in the network. The PNNI architecture for ATM is the most comprehensive approach to this problem to date and illustrates the difficulties involved by its very complexity. For further discussion of the impact of supporting Quality of Service on routing methods see: [Chen and Nahrstedt, 1998], [Lee et al., 1995] and [Lin and Yee, 1993].

[Chen and Nahrstedt, 1998] also provides a useful review of the routing techniques which have been developed for Quality of Service networks.

## 4.6 Making Future Reservations

A desirable feature of a network would be to allow users to book resources or connection time in advance of the time they actually need to service, or perhaps



book a regular weekly service. This corresponds to allocating resources not only in the current network but for future network states. The most difficult case for this problem is when call times holding times are not known in advance. If the holding time becomes known on connection time the problem is considerably simplified since this information can be used to predict a clash with service reserved at some time in the future.

## 5 AI Approaches to Routing Problems

Since the previous sections have already given some idea of the current solutions and approaches to routing problem this section outlines some of the work done in Artificial Intelligence which tackles routing problems.

### 5.1 Shortest Path Algorithms

Research into shortest path algorithms cannot be claimed to be primarily an Artificial Intelligence pursuit. However, since shortest path algorithms are so fundamental to much of the work on routing they warrant some discussion.

The number of shortest path algorithms which have been developed and published runs into the hundreds and there are new variants still appearing. The best known algorithms such as Dijkstra's [Dijkstra, 1959] and the Bellman-Ford algorithm [Ford, L. R. Jr. and Fulkerson, D. R., 1962] run in low order polynomial time. The standard Dijkstra algorithm, for example, runs in time  $O(n^2) + O(m)$  where  $n$  is the number of nodes in the network and  $m$  is the number of links in the network. Although the algorithms are low-order polynomial several difficulties arise when applying them to a running network:

1. Routing many demands. For solving off line problems - such as applying an optimisation algorithm to allocating a set of demands to a given network topology, the problem becomes equivalent to the knapsack problem [Vedantham and Iyengar, 1998] (when allocating w.r.t a cost function), or multi-commodity flow. Both of these problems are known to be NP complete [Garey, M. R. and Johnson, D. S., 1979].
2. Metrics for QoS. The choice of link/path evaluation metrics has an impact on path selection complexity. Using two additive or multiplicative metrics causes finding a route for even single demand to become in NP-complete [Wang and Crowcroft, 1996]. Such NP-complete combinations are allowed in ATM where possible metrics include: bandwidth, transfer delay, delay jitter and cell loss ratio.
3. Distribution. All of the polynomial time algorithms were initially designed to run centrally, performing pruning and search based on a complete representation of the network and its state. In a real network, information is highly distributed. Many of the algorithms can be run in a

distributed fashion (most notably: distributed Bellman-Ford used in RIP [Hedrick, 1988]), however the communication cost between network nodes then becomes the major factor in the execution time. Route calculation is typically much slower than demand arrival.

Shortest path algorithms under-pin almost all efforts to solve network routing problems however it is clear that other issues such as information distribution, speed and timing of updates, decision making metrics etc. all need to be solved to produce an effective routing mechanism. [Cherkassky et al., 1994] presents an evaluation of some common shortest path algorithms and [Deo and Pang, 1984] provides a full review of the subject area up until 1984 documenting over 200 algorithms.

## 5.2 Algorithmic Resource Allocation Methods

Much of the AI research effort on scheduling and optimisation techniques in this area has been directed towards problems such as the Travelling Salesman Problem or towards shortest path problems. There would seem to be considerable scope, however, in adapting some of these sophisticated search techniques to solving at least off-line routing problems.

[Mann and Smith, 1996] discusses routing techniques based on Genetic Algorithms and Simulated Annealing. [Mikler et al., 1996] treats routing as a multi-criterion optimisation problem and presents results of applying utility-theoretic heuristics to grid networks. [Frei and Faltings, 1998] describes the application of Constraint Satisfaction (CSP) techniques combined with abstract problem representations to routing with bandwidth constraints (again for off-line problems).

## 5.3 Distributed AI and Agent Based Routing

Since routing involves complex distributed control and information representation, techniques from Distributed Artificial Intelligence have been proposed for addressing routing problems.

[Hayzelden and Bigham, 1998] applies ideas from Brooks' subsumption architecture to develop a functional decomposition for ATM network management which includes methods for route choice between several pre-determined routes. Papers by Susan Conry and colleagues ([Conry et al., 1988], [Conry et al., 1991]) address circuit restoral problems. In the circuit restoral problem, failed connections due to link or node failures in the network need to be re-routed. This is treated as a special case of the off-line routing problem which needs to be solved very quickly (the failed connections form a set of demands which needs to be re-allocated). The approach applied in both papers uses a generalisation of the contract net algorithm. [Clark et al., 1996]

and [Weihmayer and Brandau, 1990] also address circuit restoral problems using distributed planning methods.

The use of adaptive organisations of agents to on-line source routing problems is discussed in [Willmott et al., 1999]. [Willmott and Faltings, 1999] presents the need for organisational techniques for routing in Active Networks. The ATM Forum's PNNI specification can be seen as example of the use of an organisational structure to perform routing tasks in a network. Although not explicitly agent based, the PNNI architecture could certainly be cast as a multi-agent system. Agents and, more specifically, negotiation and CSP techniques have also been applied to the multi authority routing problem [Calisti and Willmott, 1999].

## 5.4 Market Based Approaches

Economically inspired systems are finding applications in many distributed optimisation problems. There have been several applications to routing problems - primarily to the off-line routing problem.

[Wellman, 1992], [Wellman, 1994] use the example of transport planning which has strong similarities to an off-line network routing problem. Loads to be transported between different source-destination pairs in the network are routed over a fixed network topology whilst ensuring that there are no resource conflicts. [Kuwabara et al., 1996] studies the effect of communication delay on using a market based approach to on-line routing. [Gibney and Jennings, 1998] describes a market based system for managing VPC topologies by selecting between a known set of routes. An application which shares bandwidth between a set of users with dynamically changing needs is described in [Yamaki et al., 1996] and [Yamaki et al., 1998]. Another aspect of the use of market based approaches is in bandwidth auctions which enable network operators to sell spare link capacity. Initial investigations into this subject can be found in [Miller et al., 1996].

## 5.5 Biologically Inspired Paradigms

In 1994, [Appleby and Steward, 1994]<sup>8</sup> presented an elegant approach to updating routing tables in telephone networks using very simple mobile agents which communicate using traces left in the environment.

The possibilities for this approach have since been explored by several research groups.

- The work in [Schoonderwoerd et al., 1997] extends the biological metaphor and simplifies the approach,
- [Bonabeau et al., 1998] discusses extensions to increase the amount of computation done at each node.

---

<sup>8</sup>Reprinted in [Hayzelden, A. L. G. and Bigham, J., 1999].

- [Subramanian et al., 1997] applies the technique to IP packet networks and draws in learning techniques.
- [White and Pagurek, 1998] applies the approach to detecting faults and routing using multiple types of ants.

Almost all the applications of “ant based” approaches to routing problems have so far been to adaptive, hop-by-hop routing in telephone or circuit switched networks. There seems to be potential for extensions to other problems however - to data networks (connection-oriented or packet networks) and to off-line routing problems (ant systems have also been used to solve more general optimisation problems, see [Colomi et al., 1991] and [Dorigo et al., 1996]).

## 5.6 Non-Symbolic and Learning Methods

Neural networks are also a good candidate technology for application to routing problems. [Mehmet-Ali and Kamoun, 1993] presents the use of Hopfield models for solving shortest path problems. The off-line routing problem is treated in [Rauch and Winarske, 1988] (using a modification of a neural network travelling salesman algorithm), and in [Hakkinen et al., 1998] (using a network of Potts neurons). Further discussion on the application of Neural Networks to routing problems can be found in [Loofbourrow, 1995]. In general, there seem to be few changes necessary to adopt solutions for constrained optimisation problems such as the Travelling Salesman Problem to routing problems.

There have also been several learning based approaches. Wolpert and colleagues [Wolpert et al., 1999] apply reinforcement learning techniques to routing traffic with congestion avoidance in packet networks. Other reinforcement learning approaches to routing can be found in [Choi and Yeung, 1996], [Marbach et al., 1998] and [Subramanian et al., 1997].

## 6 Looking Towards the Future

Despite the large increases in available capacity, resource allocation in networks is just as critical now as it was 10 or 20 years ago. Rapidly growing network usage, heterogeneous deployed technology and diverse service requirements are all factors fuelling the drive towards improvement of routing techniques.

The main challenge areas appear to be:

- Multi authority domains. As described in Section 4.4 routing problems increasingly have to be solved in environments where much of the detailed network state information is unavailable and / or various network technologies are in use.
- Quality of Service Networks. Section 4.5 outlines the difficulties involved in routing under Quality of Service constraints. These problems are becoming increasingly important with the advent of technology which can

theoretically deliver quality guaranteed services and new applications which rely on these services. This problem arises both:

- in connection-oriented networks such as ATM, which provide the infrastructure for making reservations and keeping to service guarantees but for which the associated routing problems are far from solved,
- and in connection-less networks such as IP which also require new mechanisms to make the required resource management possible.

In both areas there are challenging problems to be solved.

Artificial Intelligence techniques seem to have considerable potential for addressing some of these problems but have yet to really prove their worth. New technologies such as active networks [Tennenhouse et al., 1997] and the new version of the IP protocol (IP version 6 [Mcgregor, 1994]) are likely to increase the scope for flexible controlling networks. This added flexibility (particularly the increased computational power in the network offered by Active Networks) should increase the opportunities for applying AI technology to solve network routing problems.

These advances depend upon a good understanding of the network problems at hand. This understanding is very hard to achieve since today's networks are shrouded in arcane terminology and voluminous documentation. We hope that this introduction at least provides an outline of where the major problem areas lie and what the difficult issues are. Finally, since this paper glosses over many technical details we recommend [Bertsekas, D. and Gallager, R. , 1992] and [Perlman, 1992] as good treatments of more technical and mathematical details involved in network routing problems.

## 7 Acknowledgements

The authors would like to extend their thanks to the other partners in the SPP-ICC IMMune project (of which this work is part). Funding for IMMune from the Swiss National Science Foundation<sup>9</sup> is also gratefully acknowledged. Special thanks also go to Monique Calisti, Ljubica Blazevic and Eric Gauthier for their helpful insights on routing problems and network architectures.

## References

- [Ramaswami, R. and Sivarajan K. , 1998] Ramaswami, R. and Sivarajan K. (1998). *Optical Networks: a practical perspective*. Morgan Kaufmann Publishers, San Francisco.
- [Alles, 1995] Alles, A. (1995). ATM Internetworking. Technical report, Cisco Systems, Inc.

---

<sup>9</sup>Project Number SPP-ICC 5003-45311.

- [Appleby and Steward, 1994] Appleby, S. and Steward, S. (1994). Mobile Software Agents for Control in Telecommunications Networks. *British Telecom Technology Journal*, 12(2).
- [Ash et al., 1991] Ash, G. R., Chen, J. S., Frey, A. E., and Huang, B. D. (1991). Real-Time Network Routing in a Dynamic Class of Service Network. In Cohen, J. W. and Pack, C. D., editors, *Proceedings ITC-13*, pages -. North Holland.
- [Ash and Huang, 1993] Ash, G. R. and Huang, B. D. (1993). An Analytical Model for Adaptive Routing Networks. *IEEE Transactions on Communications*, 41(11):1748–1759.
- [ATM-FORUM, 1996] ATM-FORUM (1996). P-NNI V1.0 - ATM Forum approved specification, af-pnni-0055.000. *ATM FORUM*.
- [Ballardie, 1997] Ballardie, A. (1997). RFC 2189: Core Based Trees (CBT version 2) multicast routing. Status: EXPERIMENTAL.
- [Bertsekas, D. and Gallager, R. , 1992] Bertsekas, D. and Gallager, R. (1992). *Data Networks, 2nd Edition*. Prentice-Hall.
- [Bonabeau et al., 1998] Bonabeau, E., Henaux, F., Geurin, S., Snyers, D., Kuntz, P., and Theraulaz, G. (1998). Routing in Telecommunications Networks with Ant-Like Agents. In Albayrak, S. and Garijo, F. J., editors, *Proceedings Second International Workshop on Intelligent Agents for Telecommunications Applications IATA'98*, pages 61–71. Springer (as LNAI-1437).
- [Braden et al., 1997] Braden, R., Zhang, L., Berson, S., and Herzog, S. and Jamin, S. (1997). Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification. RFC 2205, IETF Network Working Group, Standards track.
- [Calisti and Willmott, 1999] Calisti, M. and Willmott, S. (1999). Supporting Interworking among Network Providers using a Multi-agent Architecture. In *Proceedings of the European Conference on Network and Optical Communications*.
- [Callon, 1990] Callon, R. W. (1990). RFC 1195: Use of OSI IS-IS for routing in TCP/IP and dual environments. Status: PROPOSED STANDARD.
- [Chen and Nahrstedt, 1998] Chen, S. and Nahrstedt, K. (1998). An Overview of Quality of Service Routing for Next-Generation High-Speed Networks: Problems and Solutions. *IEEE Network*, pages 64–79.
- [Cherkassky et al., 1994] Cherkassky, M. B. V., Goldberg, A. V., and Radzik, T. (1994). Shortest Path Algorithms: Theory and Experimental Evaluation. In *Proceedings of the 5th ACM-SIAM Annual Symposium on Discrete Algorithms* , pages 516–525.

- [Choi and Yeung, 1996] Choi, S. P. M. and Yeung, D. Y. (1996). Predictive Q-routing: A Memory Based Reinforcement Learning Approach to Adaptive Traffic Control. In *Advances in Neural Information Processing Systems - 8*, pages 945–951. MIT Press.
- [Clark et al., 1996] Clark, R., Grossner, C., and Radhakrishnan, T. (1996). Consensus and compromise: Planning in cooperating agents. *Journal of Cooperative Information Systems*, 5(1):27–72.
- [Clarke et al., 1998] Clarke, P. L., Cooper, N. J. P., and Sutherland, D. J. (1998). ATM - Next Generation . *BT Technology Journal*, 16(1):106–113.
- [Colorni et al., 1991] Colorni, A., Dorigo, M., and Maniezzo, V. (1991). Distributed Optimisation by Ant Colonies. In Varela, F. and Bourgine, P., editors, *Proceedings of the First European Conference on Artificial Life*, pages 134–142. MIT Press.
- [Conry et al., 1991] Conry, S. E., Kuwabara, K., Lesser, V. R., and Meyer, R. A. (1991). Multistage Negotiation for Distributed Constraint Satisfaction. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-21(6):1462–1477.
- [Conry et al., 1988] Conry, S. E., Meyer, R. A., and Searleman, J. E. (1988). Multistage Negotiation in Distributed Planning. In Bond, A. H. and Gasser, L., editors, *Readings in Distributed Artificial Intelligence*, pages 367–384. Morgan Kaufmann.
- [de Pryker, M., 1995] de Pryker, M. (1995). *Asynchronous Transfer Mode: Solution for Broadband ISDN, 3rd Edition*. Prentice Hall.
- [Deo and Pang, 1984] Deo, N. and Pang, C.-Y. (1984). Shortest-Path Algorithms: Taxonomy and Annotation. *Networks*, 14:275–323.
- [Dijkstra, 1959] Dijkstra, E. W. (1959). A Note on Two Problems in Connection with Graphs. *Numerical Mathematics*, 1:269–271.
- [Dorigo et al., 1996] Dorigo, M., Maniezzo, V., and Colorni, A. (1996). The Ant System: Optimisation by a Colony of Cooperating Agents. *IEEE Transactions on Systems, Man and Cybernetics*, 26:1–13.
- [Estrin et al., 1998] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma, P., and Wei, L. (1998). RFC 2362: Protocol independent multicast-sparse mode (PIM-SM): Protocol specification. Status: EXPERIMENTAL.
- [Ford, L. R. Jr. and Fulkerson, D. R., 1962] Ford, L. R. Jr. and Fulkerson, D. R. (1962). *Flows in Networks*. Princeton University Press, Princeton, NJ.
- [Frei and Faltings, 1998] Frei, C. and Faltings, B. (1998). A Dynamic Hierachy of Intelligent Agents for Network Management. In Albayrak, S. and Garjjo, F. J., editors, *Proceedings Second International Workshop on Intelligent*

- Agents for Telecommunications Applications IATA '98*, pages 1–16. Springer (as LNAI-1437).
- [Garey, M. R. and Johnson, D. S., 1979] Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman and Co.
- [Gibney and Jennings, 1998] Gibney, M. A. and Jennings, N. R. (1998). Dynamic Resource Allocation by Market-Base Routing in Telecommunications Networks. In Albayrak, S. and Garijo, F. J., editors, *Proceedings Second International Workshop on Intelligent Agents for Telecommunications Applications IATA '98*, pages 102–117. Springer (as LNAI-1437).
- [Girard, A., 1990] Girard, A. (1990). *Routing and Dimensioning in Circuit-Switched Networks*. Addison-Wesley.
- [Govindan et al., 1999] Govindan, R., Alaettinoglu, C., Eddy, G., Kessens, D., Kumar, S., and Lee, W. S. (1999). An Architecture for Stable, Analyzable Internet Routing. *IEEE Network*, pages 29–35.
- [Guerin and Orda, 1997] Guerin, R. and Orda, A. (1997). QoS-based Routing in Networks with Inaccurate Information: Theory and Algorithms. In *Proceedings of IEEE INFOCOM'97*.
- [Hakkinen et al., 1998] Hakkinen, J., Lagerholm, M., Peterson, C., and Soderberg, B. (1998). A Potts Neuron approach to Communication Routing. *Neural Computation*, 10:1587–1599.
- [Hayzelden and Bigham, 1998] Hayzelden, A. L. G. and Bigham, J. (1998). Heterogeneous Multi-Agent Architecture for ATM Virtual Path Network Resource Configuration. In Albayrak, S. and Garijo, F. J., editors, *Proceedings Second International Workshop on Intelligent Agents for Telecommunications Applications IATA '98*, pages 45–59. Springer (as LNAI-1437).
- [Hayzelden, A. L. G. and Bigham, J., 1999] Hayzelden, A. L. G. and Bigham, J. (1999). *Software Agents for Future Communication Systems*. Springer Verlag.
- [Hedrick, 1988] Hedrick, C. (1988). Routing information protocol. RFC 1058, IETF Network Working Group.
- [Hedricks, 1991] Hedricks, C. L. (1991). An Introduction to IGRP. Technical report, Laboratory for Computer Science Research, Rutgers University.
- [Huitema, 1995] Huitema, C. (1995). *Routing in the Internet*, P T R Prentice-Hall, Englewood Cliffs, NJ 07632, USA.
- [Hwang and Richards, 1992] Hwang, F. K. and Richards, D. S. (1992). Steiner Tree Problems. *Networks*, 22(1):55–89.



- [Hwang et al., 1992] Hwang, F. K., Richards, D. S., and Winter, P. (1992). *The Steiner Tree Problem*. Elsevier.
- [Kumar and Venkataram, 1997] Kumar, G. P. and Venkataram, P. (1997). Artificial Intelligence Approaches to Network Management: Recent Advances and a Survey. *Computer Communications*, 20:1313–1322.
- [Kuwabara et al., 1996] Kuwabara, K., Ishida, T., Nishibe, Y., and Suda, T. (1996). An Equilibratory Market Based Approach for Distributed Resource Allocation and its Applications to Communication Network Control. In Clearwater, S. H., editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, pages 53–73. World Scientific.
- [Laubach and Halpern, 1998] Laubach, M. and Halpern, J. (1998). RFC 2225: Classical IP and ARP over ATM. Obsoletes RFC1626, RFC1577. Status: PROPOSED STANDARD.
- [Lee et al., 1995] Lee, W. L., Hluchyj, M. G., and Humblet, P. A. (1995). Routing subject to Quality of Service Constraints. *IEEE Network*.
- [Liebowitz, J. and Prerau, D. S., 1995] Liebowitz, J. and Prerau, D. S. (1995). *Worldwide Intelligent Systems: Approaches to Telecommunications and Network Management*. IOS Press.
- [Lin and Yee, 1993] Lin, F. Y. S. and Yee, J. R. (1993). A Real-Time Distributed Routing and Admission Control Algorithm for ATM Networks. In *Proceedings IEEE INFOCOM 12 volume 3*, pages 792–801.
- [Loofbourrow, 1995] Loofbourrow, T. (1995). Neural Networks in Telecommunications. In Liebowitz, J. and Prerau, D. S., editors, *Worldwide Intelligent Systems: Approaches to Telecommunications and Network Management*, pages 163–185. IOS Press.
- [Lougheed and Rekhter, 1991] Lougheed, K. and Rekhter, Y. (1991). RFC 1267: Border Gateway Protocol 3 (BGP-3). Obsoletes RFC1163. Status: HISTORIC.
- [Mann and Smith, 1996] Mann, J. W. and Smith, G. D. (1996). A Comparison of Heuristics for Telecommunications Traffic Routing. In *Modern Heuristic Search Methods*, pages 235–254. John Wiley & Sons Ltd.
- [Marbach et al., 1998] Marbach, P., Mihatsch, O., Shulte, M., and Tsisiklis, J. (1998). Reinforcement Learning for Call Admission Control and Routing in Integrated Service Networks. In *Advances in Neural Information Processing Systems - 10*, pages 922–928. MIT Press.
- [Mcgregor, 1994] Mcgregor, T. (1994). IP version 6: An introduction. Technical Report 94-4, Monash University, Department of Robotics and Digital Technology.

- [McQuillan and Walden, 1977] McQuillan, J. M. and Walden, D. C. (1977). The ARPANET Design Decisions. *Computer Networks*, 1.
- [Mehmet-Ali and Kamoun, 1993] Mehmet-Ali, M. K. and Kamoun, F. (1993). Neural Networks for Shortest Path Computation and Routing in Computer Networks. *IEEE Transactions on Neural Networks*, 4(6):941–954.
- [Mikler et al., 1996] Mikler, A. R., Honavar, V., and Wong, J. S. K. (1996). Analysis of utility-Theoretic Heuristics for Intelligent Adaptive Network Routing. In *Proceedings of AAAI'96*, pages 96–101. AAAI Press.
- [Miller et al., 1996] Miller, M. S., Krieger, D., and Hardy, N. (1996). An Automated Auction in ATM Network Bandwidth. In Clearwater, S. H., editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation.*, pages 96–125. World Scientific.
- [Mills, 1984] Mills, D. L. (1984). RFC 904: Exterior Gateway Protocol formal specification. Updates RFC0827, RFC0888. Status: HISTORIC.
- [Moy, 1989] Moy, J. (1989). The OSPF specification. RFC 1131, IETF.
- [Moy, 1994] Moy, J. (1994). RFC 1548: Multicast extensions to OSPF. Internet Draft 1548, IETF Network Working Group.
- [Nichols et al., 1997] Nichols, K., Jacobsen, V., and Zhange, L. (1997). A two-bit differentiated services architecture for the internet. draft-nichols-diff-svc-arch-01.txt.
- [O'neill et al., 1998] O'neill, A., Tatham, M. C., Carter, S. F., Tsirtsis, G., and Dann, A. J. (1998). An Overview of Internet Protocols. *BT Technology Journal*, 16(1):126–139.
- [Pan and Schulzrinne, 1997] Pan, P. and Schulzrinne, H. (1997). YESSIR: A Simple Reservation Mechanism for the Internet. Technical Report RC 20967, IBM Research, T. J. Watson Research Center.
- [Perlman, 1992] Perlman, R. (1992). *Interconnections: Bridges and Routers*. Addison Wesley.
- [Rauch and Winarske, 1988] Rauch, H. E. and Winarske, T. (1988). Neural Networks for Routing Communications Traffic. *IEEE Control Systems Magazine*, pages 26–31.
- [Rekhter and Li, 1998] Rekhter, Y. and Li, T. (1998). A border gateway protocol 4 (bgp-4). draft-ietf-idr-bgp4-08.txt.
- [Schoonderwoerd et al., 1997] Schoonderwoerd, R., Holland, O., and Bruten, J. (1997). Ant-like agents for load balancing in telecommunications networks. In Johnson, W. L. and Hayes-Roth, B., editors, *Proceedings of the 1st International Conference on Autonomous Agents*, pages 209–216, New York. ACM Press.

- [Sexton, M. and Reid, A., 1997] Sexton, M. and Reid, A. (1997). *Broadband Networking: ATM, SDH, and SONET*. Artech House, Boston/London.
- [Smith, 1999] Smith, J. M. (1999). Programmable Networks: Selected Challenges in Computer Networking. *IEEE Computer Magazine*, 32(1):40–42.
- [Subramanian et al., 1997] Subramanian, D., Druschel, P., and Chen, J. (1997). Ants and Reinforcement Learning: A Case Study in Routing in Dynamic Networks. . In *Proceedings of IJCAI'97*, pages 832–838.
- [Tennenhouse et al., 1997] Tennenhouse, D. L., Smith, J. M., Sincoskie, W. D., Wetherall, D. J., and Minden, G. J. (1997). A survey of active network research. *IEEE Communications*, 35(1):80–86.
- [Vedantham and Iyengar, 1998] Vedantham, S. and Iyengar, S. S. (1998). The Bandwidth Allocation Problem in the ATM network model is NP-Complete. *Information Processing Letters*, 65:179–182.
- [Waitzman et al., 1988] Waitzman, D., Partridge, C., and Deering, S. E. (1988). RFC 1075: Distance vector multicast routing protocol. Status: EXPERIMENTAL.
- [Wang and Crowcroft, 1996] Wang, Z. and Crowcroft, J. (1996). Quality-of-Service Routing for Supporting Multimedia Applications. *IEEE Journal on Selected Areas in Communications*, 14(7).
- [Wei and Estrin, 1994] Wei, L. and Estrin, D. (1994). The Trade-offs of Multicast Trees and Algorithms. In *Proceedings of the 1994 Conference on Computer Communications and Networks*.
- [Weihmayer and Brandau, 1990] Weihmayer, R. and Brandau, R. (1990). Co-operative Distributed Problem Solving for Communication Network Management. *Computer Communications*, 13(9):547–557.
- [Weihmayer and Velthuisen, 1998] Weihmayer, R. and Velthuisen, H. (1998). Intelligent Agents in Telecommunications. In Jennings, N. R. and Wooldridge, M., editors, *Agent Technology Foundations, Applications and Markets*, pages 201–217. Springer Verlag and UNICOM UK.
- [Wellman, 1992] Wellman, M. P. (1992). A General-Equilibrium Approach to Distributed Transportation Planning. In *Proceedings AAAI'92*, pages 282–289.
- [Wellman, 1994] Wellman, M. P. (1994). A Market-Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems. *Journal of Artificial Intelligence Research*, 1(1):1–23.
- [White and Pagurek, 1998] White, T. and Pagurek, B. (1998). Towards Multi-swarm Problem Solving in Networks. In *Proceedings: International Conference on Multi Agent Systems*, pages 333–340.

- [Willmott and Faltings, 1999] Willmott, S. N. and Faltings, B. (1999). Active Organisations for Routing. In TBA, editor, *Proceedings of the First International Working Conference on Active Networks.*, page TBA. Springer Verlag, Lecture Notes in Computer Science Series.
- [Willmott et al., 1999] Willmott, S. N., Frei, C., Faltings, B., and Calisti, M. (1999). Organisation and Co-ordination for On-line Routing in Communications Networks. In Hayzelden, A. L. G. and Bigham, J., editors, *Software Agents for Future Communication Systems*. Springer Verlag.
- [Wolpert et al., 1999] Wolpert, D. H., Tumer, K., and Frank, J. (1999). Using Collective Intelligence to Route Internet Traffic. *Advances in Information Processing Systems*, 11.
- [Yamaki et al., 1996] Yamaki, H., Wellman, M. P., and Ishida, T. (1996). A Market-Based Approach to Allocating QoS for Multimedia Applications. In *Proceedings: International Conference on Multi Agent Systems*, pages 385–392.
- [Yamaki et al., 1998] Yamaki, H., Yamauchi, Y., and Ishida, T. (1998). Implementation Issues on Market-Based Qos Control. In *Proceedings: International Conference on Multi Agent Systems*, pages 357–364.