

Active Organisations for Routing

Steven Willmott¹ and Boi Faltings¹

Laboratoire d'Intelligence Artificielle, Department Informatique,
Swiss Federal Institute of Technology, IN (Ecublens), CH-1015 Lausanne, Switzerland.
{willmott,faltings}@lia.di.epfl.ch

Abstract. Communications networks require increasingly complex resource management to stay up and running. This is particularly true in networks which aim to provide some guaranteed quality of service (either explicitly as in ATM and other connection-oriented architectures or implicitly as in a smoothly running IP network).

The resulting increased complexity of routing procedures needs to be handled in a coordinated and flexible manner. This control could well be provided by customisable control programs in the network which rely on the computational capabilities provided by active nodes. Not only will control programs need to act independently and autonomously but they will also need to coordinate their actions with each other to ensure that decisions in the network are taken in a coordinated and consistent manner.

This paper presents a framework for organising groups of control programs for routing tasks in a network. The organisation is able to adapt its own structure over time as the state of the network changes.

Keywords: Routing, organisation, coordination, distributed artificial intelligence.

1 Introduction

Resource management and routing are network management problems which require careful control in today's communications networks. Despite numerous predictions of a bandwidth glut ([11] among others), bandwidth use still needs 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. The problems are particularly acute for networks which aim to provide any kind of quality guarantees:

1. Connection-oriented network architectures such as TDM, SDH, SONET and ATM aim to guarantee Quality of Service (QoS) on a connection by connection basis. Making route calculations involves taking into account large amounts of link state information.

2. In packet-based networks such as IP, routing is based on shortest path algorithms and there is typically a single main route available given for each source-destination pair (the estimated shortest). One of the principle aims of packet network operators is to keep the call rejection rate low whilst ensuring that accepted customers experience high levels of service. The advent of flow identification, as proposed for IPv6 [RFC2460] and per flow or per application routing (which may be possible with active network technology) would enable far more flexible resource allocation.

Good allocation strategies can in both cases dramatically reduce the amount of over-capacity required to ensure smooth running.

This paper discusses the use of customisable control programs running in active network nodes to control network routing. To accomplish this control programs need not only local control and information but also to coordinate with each other throughout the network. We present notions of organisation drawn from work in Distributed Artificial Intelligence (DAI) and Management Science as an approach to this problem. The main focus of this paper is on bandwidth adaptive organisations which change structure to match network resource availability.

2 Organised Routing?

Active networks provide the means to insert (possibly arbitrary) control programs and decision logic into individual network nodes. This supports the key aim behind much of today's active networks research: enabling *custom user control programs to be added into the network on a user by user or application by application basis*.

Arguably, for many applications there will also be an additional need for broader control to ensure that these control programs deliver a coherent final result across the network. Additionally the computations carried out at individual nodes will require (possibly non-local) network state and policy information for execution. Routing and resource allocation is perhaps chief amongst applications requiring wider coordination and state information. Routing algorithms often need to display non-local characteristics, such as prevention of routing loops and avoidance of bottlenecks.

It directly follows that resource allocation processes running in network nodes need to be coordinated in their actions and have a clearly defined way of accessing distributed network state information - *they need to be organised*.

3 Active Nodes

Active networks are an enabling technology for the deployment of more intelligent and flexible network management schemes [12]. In the context of network routing problems active nodes in a network need to provide the following facilities:

- A computational environment which executes control programs operating on the routing process (we assume this environment can execute arbitrary control code but there should be scope for restricting this).
- Access for control programs to a restricted set of primitives controlling node and link resources (a virtual instruction set).
- A mechanism for updating the control programs present in the computation environment. Ideally this mechanism would be of the *programmable switch* type [12], acting as a “back door” for uploading new control programs into nodes.

These together provide for logical (or actual) mobility of control programs, information, routing policies and inter controller relationships between nodes in the network. The work presented here is being applied in two domains: ATM networks and IP packet networks. The following two subsections outline the types of active nodes required for each.

3.1 Active ATM Nodes

In the case of an ATM network, controllers (the routing processes) do not directly manipulate the packet flow since routing decisions are made on a connection by connection basis. Once routes have been chosen they are set up in the switching tables of intermediate ATM nodes. Controllers instead control the application of route decision policy and how this is coordinated with other switches, making decisions on a per-demand basis.

“Active” ATM nodes for our purposes therefore need to provide access to the primitives which control route selection for connections and to signalling processes used for connection set-up. In ATM networks there seems to be less scope for the active network “capsule” approach to passing code to nodes in packet headers. In general, parameters and settings for a whole flow are declared at connection set-up time - leaving much less flexibility for actions to be applied to individual packets.¹

3.2 Active Packet Nodes

In packet networks, routing algorithms have direct influence at the packet forwarding level. In fact a strong branch of Active Networks research (advocated in [13] for example) is based upon the idea that the packet is the fundamental unit *for control* and *to control*.

It is, however, difficult to see how effective routing and resource control could be achieved at the individual packet level (although for other network management functions this may well be the best level for control). Abstraction away from individual packets by aggregation of traffic into flows, groups of flows or other groups is often seen as essential for useful resource planning (see, for example, the efforts to provide flow identification in IPv6 [RFC2460]). Routing

¹ Active packet headers might however be used profitably for making VBR and ABR services more controllable.

tables in IP networks route packets in real time but stay relatively static, these tables are then updated by routing protocols such as RIP [RFC1058] and OSPF [RFC1131] which run in the background. IP routing tables therefore correspond to packet aggregation based on *destination address*. In active networks there may be several useful criteria for aggregation in the real time routing mechanism (rather than simply by destination) such as by source and destination, application or packet priority.

An active packet node in the packet network for our purposes in the routing context would need to provide:

- Access to the on-line routing mechanism principally so that it can be updated by routing mechanism operating in the background (e.g. not on a packet by packet basis).
- A mechanism for flow identification.
- Per-flow routing capabilities.

The last two points allow much finer control of routing and are now often seen as essential for good resource management and QoS support in packet networks. Active networks is an important enabling technology for these two properties (and, of course, the first).

4 Building Organisations

There is a large body of work in both Management Science and DAI on how to apply organisational theory to distributed computational systems. [5] gives an AI perspective, [3] and [9] give interesting Management Science viewpoints and a collection of papers covering market based systems can be found in [2]. This work is complemented by the general trend in the network research community towards decentralisation and the use of hierarchies [10] and delegation [6]. The PNNI framework under development by the ATM Forum [1] is perhaps the most advanced and best known use of organisations in network architecture to date.

A useful organisation needs to provide the following:

- **Information Organisation:** Dividing up and providing access to information. For routing, this corresponds to representing the information required for making routing decisions (the link states and topology etc.).
- **Control Organisation:** Ensuring distributed control decisions lead to coherent actions being carried out throughout the network. In routing, this corresponds to: 1) de-limiting where routing decisions are taken (hop by hop? at the source?) and 2) avoiding bottlenecks, congestion, oscillations, etc. and making sure the correct reservations are made.

Organisations are made up of two types of component: units and relations. Units represent, for example; company departments, employees, sites or (here) areas of a network. Relations describe the relationships between units (such as superiority, parent, child, peer etc.) and define the organisational structure. The

examples given in this paper are all spatially distributed hierarchies. However, much of the discussion is relevant to other types of distribution (functional, by authority), and organisation ([3] for example discusses heterarchies, hierarchies and markets).

5 Static Organisations

In a static organisation the composition of units and existing relations between them remain fixed over time. The following two sections describe a static routing organisation in terms of controllers (units) and structure (formed by relations).

5.1 Control Programs

Control programs represent organisational units and perform management tasks in local areas of the network. For the routing task, control programs require the following types of information:

- Routing policies and algorithms.
- Continuously updated local information about the network state. This forms the basic input for solving routing tasks.
- Responsibilities to other controllers elsewhere in the network (for example their superiors). This corresponds to control programs knowing their place in the organisation.

Both the local information and the algorithms/policies may change over time to reflect the changing network state and management control. Controllers have knowledge of a limited area of the network and information about what lies outside this area is obtained via the organisational links. Control programs in a static organisation perform functions at two levels:

- **Local:** Executing local routing tasks given demands for routes and their local state information.
- **Organisational:** Cooperating with other controllers in the organisation to execute non-local routing tasks. These may be tasks arising in the controller's own area for which non-local information or control is needed, or tasks arising elsewhere (which need information or action in control program's own sphere of influence).

5.2 Control Structure

Relations are applied to compose many local control programs into an organisational structure which spans the whole network. Figure 1 shows a hierarchical structure with three levels of organisational units.

The lowest level controllers (one level above the individual network nodes) each have a local viewpoint and hold information about the network state in that area. Actions are coordinated in accordance with the relations between the units,

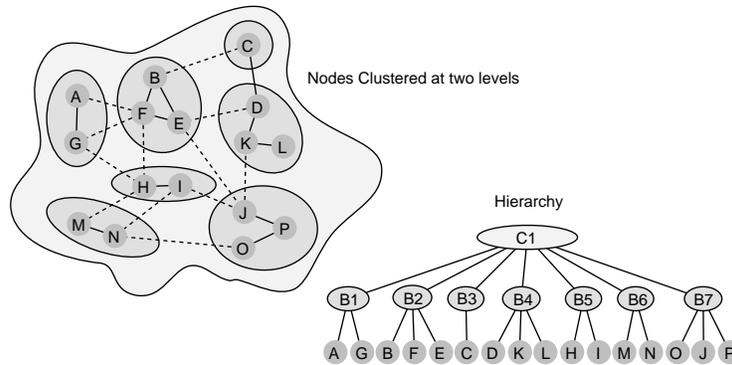


Fig. 1. The network nodes on the left are clustered at two levels,

e.g.: peer to peer: B2 communicates directly with B4 to find out about connectivity to node L, or hierarchically: controller C1 mediates between controllers in level B (all the Bx) to perform routing tasks. This control structure stays fixed over time (although it may be updated for the physical addition/removal of nodes for example).

6 Adaptive Organisations

There are many different organisational structures (even if restricted to hierarchies) and no single organisation is appropriate for all tasks. There are clear arguments for allowing organisations to adapt over time, this is particularly the case when the environment they operate in is dynamic or there may need to be ad-hoc re-organisation (due to failures for example).

There has been some preliminary work on adaptive organisations in the Distributed Artificial Intelligence (DAI), literature with [8] and [7] the most useful examples². The key requirement behind adaptive organisations is that the controllers in the organisation have some *representation of their place in the organisation*. The controllers can then apply a set of adaptation rules to decide when to change this representation (and inform other controllers of the changes). The following sections present an organisation which adapts to bandwidth availability over time to illustrate the idea and utility of adaptive organisations.

6.1 Resource Summarisation at Different Levels of Abstraction

[4] introduces a clustering scheme for structuring network graphs. The network is divided up into equivalence classes according to connectivity at a specified available bandwidth. The regions created are called Blocking Islands³. Figure 2

² We refer the reader to [15] for further references.

³ Please note that the clustering techniques and their applications are subject to patent protection.

gives an idea of the structures created by this clustering approach. A single network graph including nodes A to M and connected via links of varying residual capacities is clustered into regions at 6Mbits/sec and 9Mbits/sec.

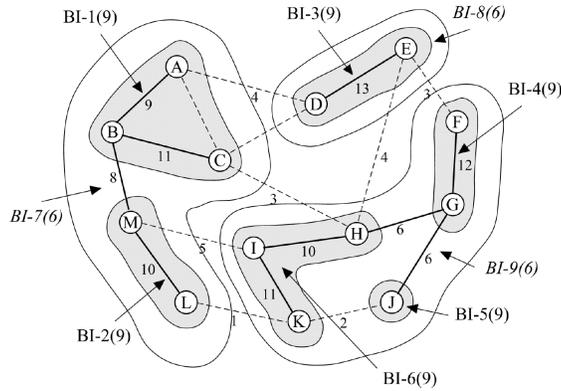


Fig. 2. Two layers of blocking islands cluster a network at different levels of abstraction. The light grey regions ($BI-1(9)$ - $BI-6(9)$) cluster equivalence classes of nodes reachable at 9Mbits/sec. The larger regions ($BI-7(6)$ - $BI-9(6)$) cluster groups of nodes which can be reached at 6Mbits/sec. Dashed lines represent communication links with less than 6Mbits/sec free capacity, solid lines represent links with 6Mbits/sec or more free.

The sets of blocking islands generated for one bandwidth requirement are *unique, identify bottlenecks* (the inter-regional links) and highlight the *existence and location of routes* at a given bandwidth level. If two nodes are clustered in the same blocking island at a given bandwidth level there *must exist at least one route between them* - furthermore all links which form part of the path *lie inside this blocking island*.

Applying this clustering technique several times for different bandwidth requirements represents bandwidth bounded connectivity in the network at different levels of abstraction. Changes in the available bandwidth on the links can cause splitting or merging of regions.

6.2 A Bandwidth Adaptive Routing Organisation

The hierarchy generated by the resource summarisation in Section 6.1 can be used to build an organisation. Control programs in the network nodes gather and hold information for each of the regions (blocking islands above). The simplest mapping is to designate one piece of control code responsible for each region both in the abstract and ground space (node-level). Control programs retain links to their neighbouring (peer) and parent/child regions (or at least to the controllers of these regions). Through the clustering scheme in the section above, this structure is then related directly to the bandwidth available and changes over time as resources are allocated and de-allocated.

This organisation is applied to performing routing tasks in the network. More specifically, we are currently applying this to allocating CBR demands in ATM networks based on a source routing model.⁴ Controllers at the lowest level of abstraction perform routing tasks on real network nodes whilst controllers at higher levels coordinate the efforts of their subordinates (at lower levels of abstraction).

The useful properties of the clustering scheme identified in [4] apply to any convex metric. However, bandwidth appears to be the most useful of these, primarily because many other QoS parameters (such as delay, jitter etc.) depend heavily upon available bandwidth [14]. Having the organisation adapt to the available bandwidth means that bandwidth information for routing decision making is *already implicit* in the information structure (before any routing algorithm has even been executed).

6.3 Control Programs

Control programs in an adaptive organisation are generalised versions of those used in static organisations (see Section 5). Controllers now require an additional meta-level of information:

- Metrics for evaluating the need to change the organisation structure. These metrics form the update rules of the organisation.

In a bandwidth adaptive organisation these metrics and update rules are based upon bandwidth availability in the network (and in this work arising out of the clustering techniques described in Section 6.1). As a result of this additional meta-level control, programs may now act at three levels:

- **Local:** Executing local routing tasks given demands for routes and their local state information.
- **Organisational:** Cooperating with other controllers in the organisation to execute non-local routing tasks.
- **Meta Organisational:** In updating the organisation structure by applying the update rules.

It is important to note that the update rules for this new meta-level must be embedded within the individual control programs themselves. Updates of the organisational structure cannot realistically be controlled by external processes (since network failures could cause all local adaption to stop and in large networks external adaption is a complex problem in itself if solved centrally).

⁴ Since this paper's main aim is to discuss organisation issues, the routing schemes related to the organisational structures are not discussed here - these are described more precisely in [15].

6.4 Control Structure

Figure 3 shows two clusterings of controllers for a grid network. The first (left hand side) is the starting state with the controllers covering initially defined regions. For a static organisation this configuration persists and remains fixed over time. The right hand side shows an example adapted control structure for the same grid network. Nodes clustered at the top level (sharing the darkest regions, such as B and C) have high bandwidth connectivity available and regions only connected at low levels (such as A and B) are only reachable at low bandwidth.

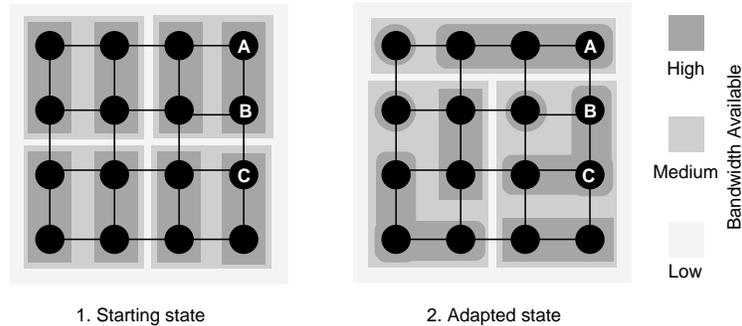


Fig. 3. A grid network with a starting organisation as shown on the left. Over time the structure changes for a bandwidth adaptive organisation (as shown on the right). For a static organisation the initial state would be preserved.

Finding a route between neighbouring nodes B and C is quick in the adaptive organisation since both nodes lie in the same local area of control. The same task takes longer in the static organisation because the two nodes happen to be clustered only at the highest level of abstraction. This difference reflects the representation of the ready availability of resources between nodes B and C in the adaptive organisation, which is something the static organisation does not capture. The situation is reversed for communications between nodes A and B and the adaptive organisation clusters these at the top level of the hierarchy whereas the static organisation can make a decision at the most local level. The extra effort required in the adaptive organisation reflects the fact that A and C are connected only by paths which are *resource critical* which may mean that they should be dealt with by an entity with a broader view of the network. Routing traffic on a critical link may have wider consequences for the rest of the network (for instance it may unnecessarily disconnect two regions of the network).

7 Active Organisations for Routing

An organisation in an active network forms part of the environment control programs on a given node execute in, it defines:

- What information is available at execution time (information organisation),
- What the wider context for the execution of the program is (control organisation). The relationships with other entities in the network may restrict the possible outcomes of computation, may determine, counteract or cause non-local effects.

In our work on routing problem, network state information is managed within the organisation. Each controller holds local information and is able to query other controllers to obtain non-local information. The organisation also defines how a routing problem is solved: resource allocation decisions are finally taken by those controllers responsible for the lowest (link) level but where reservations are made is partly controlled from higher up in the hierarchy. The coordination structure provided by the hierarchy ensures that local resource reservations hang together to give complete routes and that load is evenly distributed to avoid congestion.

The adaptive organisations described in the previous sections are also “active” in the sense that its control programs are able to represent the state of the organisation explicitly and update the organisation itself. Updates of code and information, as well as relationships etc., do not only come from network operators or users of individual applications but from within the network. Controllers on each node have some autonomy and influence over their own status in the organisation. Controllers may also have influence over controllers on other nodes (at lower levels of abstraction).

7.1 Importance of Active Network Developments

The requirements laid out in Section 3 clearly show the need for active network technology to support the work presented in this paper. Essentially, control programs need to be logically or actually mobile. As the organisational structure, network state and management policy (specific routing algorithms for example) change, the control programs in the network also need to adapt (or be replaced) dynamically. The adaptivity of the organisation requires considerable flexibility in the network nodes.

8 Status of Work

The work on adaptive organisations outlined in this paper is still in its preliminary stages. The node execution environments, control programs, communication mechanisms and adaption algorithms have recently been completed. What is still lacking is a generator for traffic scenarios and extensive testing. Current work is based on an ATM network model but under certain assumptions should also be applicable to packet-based networks (see Section 3).

9 Conclusions

Increasingly intelligent network management schemes, particularly for resource management, are vitally important for the smooth running of future networks. This not only true in connection-oriented networks, such as ATM, but also in packet-based networks where careful resource management is required to improve the ratio between *potential load* and *available capacity* (e.g. minimising the amount of over-capacity required).

Coordinating the actions of on-line control programs throughout the network goes hand in hand with this need for better resource management, leading to interesting questions for active networks research:

- How to facilitate this coordination?
- How to prevent the potentially wide diversity of injected programs interacting catastrophically in the network? (even if none of them are violating security instructions).

We introduce notions from the field of Distributed Artificial Intelligence on organisations and discuss the use of organisations for routing tasks. The paper contains two main threads of argument:

1. Organisations are important in ensuring that active networks behave coherently when executing many different user/system injected control programs. Programs executing at nodes require both information organisations (to be able to perform useful tasks) and control organisations (to ensure coherent behaviour).
2. Organisational structures in networks are heavily dependent upon active network techniques to provide flexible computation at network nodes and mechanisms for the dynamic update of control programs. This is particularly true of adaptive organisations.

To help illustrate these points the paper also presents a bandwidth adaptive organisation scheme based on control programs which update themselves, their network information state and their organisation structure dynamically. The control programs coordinate with each other to ensure coherent execution of the resource management tasks.

9.1 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 Foundation⁵ is also gratefully acknowledged. Thanks also go to Monique Calisti and Christian Frei for helpful comments on earlier drafts.

⁵ Project Number SPP-ICC 5003-45311.

References

1. ATM-FORUM. P-NNI V1.0 - ATM Forum approved specification, af-pnni-0055.000. *ATM FORUM*, 1996.
2. S. H. Clearwater. *Market Based Control: A paradigm fo Distributed Resource Allocation*. World Scientific, Singapore, 1996.
3. M. S. Fox. An Organisational View of Distributed Systems. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-11(1):70–80, 1981.
4. C. Frei and B. Faltings. A dynamic hierarchy of intelligent agents for network management. *Workshop on Artificial Intelligence in Distributed Information Networks (held at IJCAI'97)*, 1997.
5. L. Gasser. DAI Approaches to Coordination. In N. M. Avouris and L. Gasser, editors, *Distributed Artificial Intelligence: Theory and Praxis*, pages 31–51. Kluwer, 1992.
6. G. Goldszmidt and Y. Yemini. Distributed management by delegation. In *Proceedings of the 15th International Conference on Distributed Computing Systems (ICDCS'95)*, pages 333–341, Los Alamitos, CA, USA, May30 June-2 1995. IEEE Computer Society Press.
7. F. Guichard and J. Ayel. Logical Reorganisation of DAI Systems. In *Proceedings of the ECAI-94 Workshop on Agent Theories, Architectures and Languages (ATAL'94)*, pages 118–128. Springer Verlag (as Lecture Notes in Artificial Intelligence 890), August 1994.
8. T. Ishida, L. Gasser, and M. Yokoo. Organization Self-Design of Distributed Production Systems. *IEEE Transactions on Knowledge and Data Engineering*, 4(2):123–134, April 1992.
9. T. W. Malone. Modeling Coordination in Organisations and Markets. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 151–158. Morgan Kaufmann, 1988.
10. M. R. Siegl and G Trausmauth. Hierarchical Network Management: a Concept and its Prototype in SNMPv2. *Computer Networks and ISDN Systems*, 28(4):441–452, February 1996.
11. J. M. Smith. Programmable Networks: Selected Challenges in Computer Networking. *IEEE Computer Magazine*, 32(1):40–42, January 1999.
12. David L. Tennenhouse, Jonathan M. Smith, W. David Sincoskie, David J. Wetherall, and Gary J. Minden. A survey of active network research. *IEEE Communications*, 35(1):80–86, January 1997.
13. David L. Tennenhouse and David J. Wetherall. Towards an active network architecture. *Computer Communication Review*, 26(2), April 1996.
14. Z. Wang and J. Crowcroft. Quality-of-Service Routing for Supporting Multimedia Applications. *IEEE Journal on Selected Areas in Communications*, 14(7), 1996.
15. S. N. Willmott, C. Frei, B. Faltings, and M. Calisti. Organisation and Coordination for On-line Routing in Communications Networks. In A. L. G. Hayzelden and J. Bingham, editors, *Software Agents for Future Communication Systems*. Springer Verlag, 1999.