

# An Agent Future for Network Control?

*Steven Willmott, Monique Calisti*

*In an age of rapidly increasing network complexity and diversity, the idea of “bringing intelligence to the network” is becoming more of a necessity than “future work”. Since the early 1990’s, agent technology has often been proposed as a way of achieving this more effective, robust and above all autonomous network control. This article provides a brief tour of current trends in network development and discusses the potential for agent based solutions to some of the most pressing communications network problems.*

## 1 Introduction

From a network engineer’s point of view, a network is a complex system requiring complicated management under very trying domain constraints. To a Distributed Artificial Intelligence researcher, a network is a highly distributed, complex and challenging environment for the application of intelligent systems ([Lewis95], [Maes94], [Weihmayer/Velthuijsen 94]). The idea of distributing communication network control and management tasks by deploying “smart”, “cooperative” and “autonomous” entities in network infrastructures has thus received considerable attention from both the Distributed Artificial Intelligence (DAI) and the Communications Network (CN) communities.

As networks become increasingly complex and difficult to control, the ideal of a distributed, intelligent network management and control system is becoming more and more of a necessity. Furthermore, new software and network technologies are revolutionising what can be deployed in the network

and even what we think of as the network itself. Despite the lack of deployed systems, these trends make an “agent future for network management” seem closer than ever.

This article does not aim to replicate the useful surveys already completed in this area. Instead, the aim is to give a briefer overview of the research field which balances the traditionally separated CN and DAI viewpoints. Rather than going into detail on individual research efforts we review the area by:

- Identifying the current trends which suggest that agent technology may play an increasingly important role in network control (Section 2).
- Highlighting three key areas which might benefit most from agent technology: multi provider environments (Section 3.1), resource management (Section 3.2) and communications integration (Section 3.3).
- Discussing the necessary steps for the deployment of agent systems in future communications networks (Section 4.).

Those readers interested in more detailed accounts of previous work should find the following surveys useful starting points:

- [Kumar/Venkataram 97], [Weihmayer/Velthuijsen 98] and a recent volume of collected works [Hayzelden/Bigham 99] all give useful DAI perspectives.
- [Martin-Flatin/Znaty 2000] gives an overview of existing network management paradigms which places work on agents in a Network Management context. More specific works on software agents for management operations can be found in previous proceedings of the IATA<sup>1</sup> and DSOM<sup>2</sup> workshops.

### 1.1. “Agent” Terminology

One of the unfortunate side effects of the separation of work between the DAI and CN communities is confusion over terminology – particularly surrounding the term “agent”. Many similar terms (for example SNMP agents, mobile agents, “intelligent” agents, agents, BDI agents) are used for different purposes by the two communities. In this article we follow the agent definition given in [Jennings/Wooldrige 98]. This defini-

1. Intelligent Agents for Telecommunications Applications.
2. Distributed Systems: Operations & Management.

*Steven Willmott* is a researcher in the Artificial Intelligence Laboratory at EPFL Lausanne. His primary research interests centre on the control of distributed systems using Distributed Artificial Intelligence techniques – with particular focus on agent coordination, organisation and language. In 1999 he chaired the 3rd workshop on Artificial Intelligence in Distributed Information Networking (AiDIN’99 held at AAI in Orlando, USA). Aside from his academic activities he is also heavily involved in the FIPA agent standardisation effort as the editor of the 1999 FIPA Agent Message Transport Specification.

*Monique Calisti* has received her master degree in Electrical Engineering at the University of Bologna (Italy). After having received the pre-doctoral School diploma in Communication Systems at EPFL, she joined the Artificial Intelligence Laboratory of EPFL at the end of 1997. She is currently PhD student and her primary area of interest is the interaction of distinct telecommunication network providers. The focus is on the allocation of multi-provider service demands with Quality of Service requirements. The use of Distributed Artificial Intelligence techniques for enabling a flexible and automatic provider-to-provider paradigm represents one of the main topic under investigation. She is also an active member of the FIPA agent standardisation group as the editor of the 1999 FIPA Content Language Library Specification.

tion has strong DAI roots and proposes that an agent is an entity which is:

- *Situated* in some environment.
- *Autonomous*, in the sense that the system can act without direct intervention from others (humans or other software processes).
- *Flexible*, which is further broken down into three properties: *Responsive* (perceives its environment and responds to changes in a timely fashion), *Pro-active* (exhibits opportunistic, goal-directed behaviour) and *Social* (able to interact with humans or other artificial agents).

This definition might be considered consistent with that of an “intelligent agent” in the CN community. The reason for dropping the “intelligent” is that the notion of intelligence is difficult to pin down. Often, this becomes subdivided into arbitrary “types of intelligence” or “levels of intelligence”. The agent definition may also be loose enough to incorporate some of the agents not considered “Intelligent” by the CN community. It should also be noted that this definition considers mobility as an optional property of an agent (or of a piece of code) rather than something which defines a whole class of agent; helping avoid problems such as classifying “intelligent mobile agents”. This definition, we feel, is particularly useful since it reinforces the view of an agent as a piece of software disposing a fundamental set of properties. Entities displaying more or less of these properties can then be considered more or less “agent-like” (see Section 4).

## 2 Trends in Network Development

The main factors behind the increasing interest in agent technology for network control can be divided into two categories: *application pull* – the need for innovative solutions to increasingly urgent network problems and *technology push* – the development of new techniques which make agent deployment a real possibility.

There are three main factors which are generating potential need for agent based solutions are:

- *Market liberalisation*: The deregulation of telecommunications markets has forced major changes to the roles, business models and operational practices of network operators and service providers alike (in fact until the early 1990s Network Operators and Service Providers were often one and the same). Competition is fierce and has kick started an industry-wide drive for efficiency.
- *Rapidly changing technology*: The number and diversity of deployed network technologies is continually growing. This diversification is creating a complex heterogeneous network infrastructure and serious technological challenges in providing uniform and coherent services. There is often not enough time for the industry to develop guidelines before those guidelines are already obsolete. Standard bodies (such as the ITU, ISO, ANSI, ATM Forum and IETF) are having to catch up with common practice rather than setting the agenda.
- *Increasing flexibility in usage requirements*: With market liberalisation and increasing customer demand comes a need

for flexible service deployment. Networks need to be adapted to provide what customers are demanding, cope with fluctuations in usage and handle the introduction of new multimedia services (such as video, audio, Internet telephony and e-commerce related communications).

These three factors together are combining to produce very complex network architectures and requirements. Issues of scalability, reliability, security and interactions between services are increasingly replacing any other concerns network operators may have had. Section 3 picks out some key areas where agent technology may be able to play a leading role in solving some of the most pressing network control problems.

Until recently, many agent applications have remained nothing more than small pilot projects in the research laboratory. One key reason behind this is that the necessary network architecture for agent deployment was just not available. This is changing. There are three main areas of technology push:

- *Mobile Agents*: The utility of mobile agents and mobile code for network control has been a recurring theme since the early 1990s. This paradigm is now beginning to gain wider acceptance in the CN community (see works such as [Baldi et al. 97] and [Breugst/Magedanz 98] for example). Consequently, the likelihood that agent capable platforms might be supported by future networks is increasing.
- *Programmable Networks*: Researchers in the relatively new field of Active Networks [Tennehouse et al. 97] argue for programmable networks which can receive and execute code on time scales down to single packet arrival. Programs can be downloaded to a router using a “backdoor” mechanism or injected into the network in the headers of individual data packets. Either way, this type of programmable network would greatly increase the scope for the deployment of agent based network control services into the network infrastructure.<sup>3</sup>
- *Standardisation*: Ongoing standardisation efforts within bodies such as OMG and FIPA are providing standard interaction mechanisms for agent based software. These efforts to provide interoperability for agent applications are a key factor in enabling the use of agent technology for a large range of tasks – including network related applications (also see Section 4).

Above all, agent technology is maturing as a software development paradigm. Developing environments and standards is becoming available. This trend is likely to build confidence in agent techniques and allow more wide-spread experimentation.

## 3 Key Application Areas

Agent technology has been proposed for a large number of network related areas (publications easily run to the 100s).<sup>4</sup> This section picks out three areas of network control in which agent technology may have real potential to make a difference.

3. Some CISCO Systems routers in fact already include Java Virtual Machines, however their interfaces have not yet been made available to network engineers.

### 3.1 Multi Provider Environments

Market liberalisation and increasing demands for the allocation of services which span several networks are pushing every network provider to evolve the way they interact with peer operators. In order to understand what “to evolve” means, several factors have to be taken into account.

- *Heterogeneity*: What were once state monopolies controlling everything from end user access down to the copper wires have become several layers of competing firms – service providers, networks providers, brokers, etc. Distinct networks can be based on different technologies and can deploy different network management platforms. This implies heterogeneity also of the information models used in different networks, i.e., different Management Information Databases (MIBs).
- *Distribution of resources*: Network resources can be owned by many different “authorities”, that need to be made to work together to support advanced services spanning several domains (Virtual Private Networks crossing different networks for example). This task is even more delicate for networks which aim to provide any kind of Quality of Service guarantees, since individual providers are unwilling to release detailed information about the state or topology of their internal network.
- *Flexibility*: Currently, many aspects of the interaction between distinct networks, are statically fixed by contracts (number and available capacity of links connecting one network domain to another, prices, etc.) and many steps of the interaction are regulated by human operators via fax, e-mail, etc. This makes the overall inter-interopability process very slow (several months can pass before effective inter-domain network configuration changes take place) and inefficient.
- *Robustness*: A further major problem is that there is little or no infrastructure to support robust information exchange or coordination between different service and network providers. In the best case, TMN compliant networks use standard TMN-X interfaces which provide a rudimentary low level interface for synchronising the settings in routers and other network elements. A common database, the Shared Management Knowledge (SMK), allows the visualisation of a minimal amount of information that needs to be shared for the interaction. Even in this case however, humans are responsible for supervising and controlling the interaction.

Considering these aspects, what seems more suitable for future networks is a management solution based on static and/or mobile software entities, collecting network state information and which have the ability to directly invoke effective changes to switch controllers, without the interaction of a

4. Note also that many “network supported” agent applications are also being developed (such as agent based information services, web auction houses, information filters etc.) which have greater or lesser contact with the network infrastructure. This work is not treated here. See sources such as proceedings of the PAAM series of conferences (<http://www.practical-applications.co.uk/PAAM/>) for this type of application.

human operator (see [Posladt et al. 99], [Corley et al. 98], [Calisti/Faltings 99] and several works in [Hayzelden/Bigham 99]). Software agents have strong potential since they can be distributed, intelligent, expert, heterogeneous, self-learning and dynamic [Bigham et al. 99]. As concrete examples of potential agent usage in the multi-provider framework, (see Figure 1) software agents acting on behalf of every network operator could:

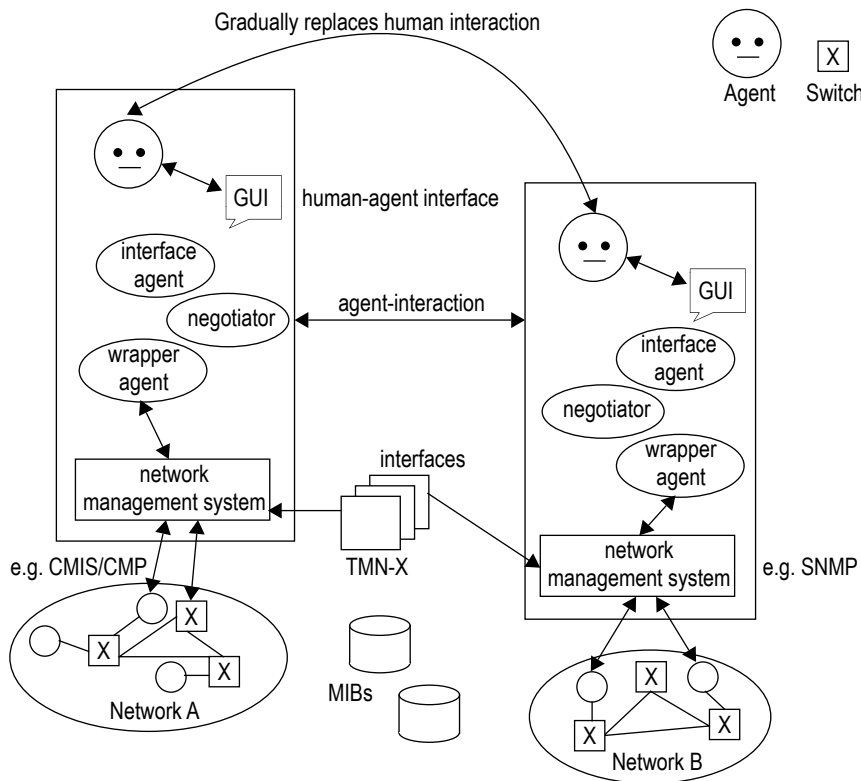
- Reduce the need of human interventions and communications.
- Abstract from technical details, such as SNMP primitives or CMIS/CMIP terms, and translating them into a more understandable form for human operators.
- Automate the control of switches and routers, i.e., active routing.
- Provide automatic service negotiation with both peer operators and final end-users.

In [Calisti et al. 99] a multi-agent paradigm for the automatic allocation of inter-domain service demands is defined. Here one of the main challenges 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 (see [Calisti/Faltings 99] for more concrete results).

### 3.2 Resource Management

Despite predictions of bandwidth glut ([Smith 99] amongst others), network resource management remains a very challenging area. In the United States, backbone deployment is a race against time and user demand, fluctuations and routing errors can have catastrophic results<sup>5</sup>. With the rapid rise of the Internet as an essential business tool there are also concerns about the potentially serious effects of prolonged periods of poor or reduced service. This factor, above all, is driving companies to demand 1) improvements in the overall service quality the Internet provides and 2) the minimisation of potentially damaging periods of poor service. These demands need to be met with increasingly sophisticated techniques for resource management (such as the efforts now going on under the umbrella term of Traffic Engineering [Bhaniramka et al. 99]) both at the backbone level and at the IP network level. These efforts correspond to *controlling resource allocations* in the network to improve the use of the available infrastructure. This is either done using reservation protocols (such as RSVP [rfc2205]) or in an across the board fashion. Agent technology has already been applied to several problems in this area:

- *IP routing*: there have been various approaches to routing problems but perhaps amongst the most intuitively appealing are those based on the use of mobile agents to mimic “ant like” behaviour. First proposed in [Appleby/Steward 94] and [Schoonderwoerd et al. 97] and continued by several other research groups, this approach simulates the trail laying behaviour of social insects such as ants in simple mobile agents. Individual ants migrate around the network laying and reinforcing trails on their chosen paths. Packets
- 5. Such as the almost global outage due to an single erroneous router configuration on March 25, 1997 [Labovitz et al. 98].



**Fig. 1:** “Agentification” of the future networks: traditional management tasks and typical human interactions are carried out by software agents.

(or calls) can then be routed along the strongest reinforced trails which are biased to be the shortest and encounter the least congestion (ants leave stronger signals when arriving at a destination more quickly than their counterparts). This approach has recently been adapted to IP networks [Subramanian et al. 97] and developed to produce more general protocols [Chen et al. 99].

- *Bandwidth commerce:* is a newly emerging model for network resource management which is based on owners of network infrastructure selling spare capacity in open markets. This trade is already happening<sup>6</sup> and is currently carried out by human operators bidding for bandwidth goods. Agent based interactions could potentially provide much more flexible interaction using standard agent interfaces, automatic negotiation, bargaining over bundles of goods (which involves complex reasoning) and, not least, saving the patience of human operators. Projects that are addressing these possibilities include MACH<sup>7</sup>. Market based resource allocation has also previously been tried for off-line allocations in work such as [Gibney/Jennings 98] and [Wellman 94].
- *Connection-oriented networks:* work including [Hayzelden/Bigham 98] and [Willmott et al. 99] has shown how agent systems could be used to control resources in back-

6. See <http://www.rateexchange.com/>, <http://www.band-x.com/>, and <http://www.interxion.com/> for example.

7. <http://liawww.epfl.ch/~calisti/MACH/mach.html>

bone networks (based on ATM or other connection oriented technologies). These methods are particularly applicable for networks where large amounts of state data is generally needed to make routing decisions. To fully automate the network it would seem a logical progression to allow agents managing IP network resources to communicate with agent counterparts charged with managing backbone resources.

In general terms, network resource management can be described at many levels of granularity (from the routing of a single packet and the routing of a flow up to the implementation of network operator allocation policies). Agents with properties such as those described in Section 1.1 are more appropriate at the higher levels of this description.

Figure 2 shows a two tier model often applied in control problems (also adopted in [Hayzelden/Bigham 98] and our own work). Control is divided into two systems: an on-line system responsible for fast time scale allocations (packet route decisions for example) and a background system which monitors, controls and updates the faster on-line system. The on-line system makes most of the day to day routing allocations, however the supervisory system would intervene if (for example): failures occur, operator policies change, traffic congestion appears to be building or certain types of traffic need to be specially treated. For this application agents appear to be particularly suitable for the supervisory style of control system since they can:

- monitor and react to the environment – hence pro-actively deal with undesirable traffic patterns in the network,
- provide control in localised areas of the network using only local information,
- communicate with each other to ensure that a more coherent, global resource allocation policy is enforced.

### 3.3 Communications Integration

There is a clear trend towards providing the user with services rather than network access or bandwidth. Furthermore, users are increasingly demanding that information services (such as news, e-mail, fax, telephone etc.) are integrated seamlessly. These demands require multiple services provided by various network technologies to be coupled together effectively. The types of integration required can be broadly classified into two types: *interface integration* and *network integration*.

- *Interface integration:* the integration of many network services requires complex coordination between network infrastructure, end devices and services. Agent based approaches have already been tested for this type of problem (see [Abu-hakima et al. 96] for example which uses a purpose built LAN). The integrated network service should ideally:

- Allow the addition of new services, network technologies and end devices (such a new pager) dynamically – delivering the communications service over the newly added medium when appropriate.
- Group together existing services to make them appear as on “virtual service”, for example integrating voice mail, fax and e-mail by delivering the messages arriving over all three media in whichever of the three formats is currently most appropriate.

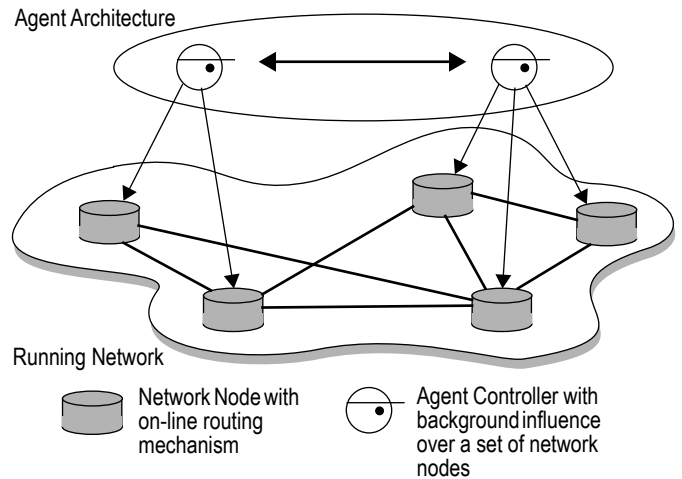
The top level of this service integration is human-machine interaction which is in turn supported by coordination in the network infrastructure to carry out the required services. The key advantages of agents here are in their pro-activity and flexible interaction with the environment. Agents enable the integration of humans and diverse hardware or software entities by adapting their behaviour to individual preferences, characteristics of users and characteristics of the network hardware [Fipa 98].

A concrete example of this kind of integration is the effort to provide a Virtual Home Environment (VHE) for 3rd generation mobile phone systems (UMTS). The aim is to have mobile phone users presented with the same options, services and interfaces wherever he or she is in the world and whichever mobile phone provider is currently providing these services. Both static [Lloyds/Pearmain 99] and mobile agent (the EU ACTS “On the Move”<sup>8</sup> project for example) approaches have been proposed for this problem.

- *Network integration*: as the number of deployed network technologies grows, providing homogeneous services requires abstraction from technological details and standardised models for communication. Additionally, different functional parts of the infrastructure may be owned by different companies with, for example, service providers leasing bandwidth from network operators.

Applying agents to integrate heterogeneous networks and network technologies has been proposed both within the CN community (in the TINA framework<sup>9</sup>) and within the DAI community (with the FIPA agent network management model, Section 7 of the 1997 FIPA Specification [Fipa 97]). In the TINA architecture, software entities interact with both humans and physical network devices, communicating over a distributed execution environment. The FIPA architecture further encompasses the notion of different authorities owning different levels and parts of the network and looks to address the question of establishing end-to-end services over several (separately owned) networks (hence similar to the issues discussed in Section 3.1).

This trend towards integration in all directions looks set to continue and is perhaps one of the most challenging problems networks of the future will have to face. In this area the concept of agent middleware which bridges the technological and architectural gaps in current systems seems to have great potential (see [Poslad et al. 99]). Agents provide a means of:



**Fig. 2:** The routers in the network each have an on-line allocation mechanism. Agents communicate with each other to resolve longer term allocation problems, occasionally intervening in the on-line system’s operation.

- Abstracting from the technological idiosyncrasies of different network technologies to improve their interoperation.
- Enabling richer and more flexible interaction between both user and system (user network service access) and system and system (automatically exchanging tasks between different agents to customise service delivery).

#### 4 The Agent Future?

The three key application areas discussed in Section 3 cover a large part of the communications network infrastructure, however this is not intended to advocate the use of “agents everywhere”. The type of software agent which fits the definitions given in Section 1.1 would arguably be inappropriate for tasks which:

- *Required vary fast repetitive processing*: the utility of using agents is generally in providing flexible execution behaviour to function correctly in a dynamic environment. By its very nature this type of processing is likely to be less efficient for highly constrained, repetitive tasks (such as packet forwarding).
- *Required rapid and precise information exchanges*: particularly in cases where agents communicate using standard agent communication languages such as KQML [Finin et al. 93] and FIPA ACL [Fipa 97], the flexibility in agent communication may be problematic. For many network tasks, highly constrained, concise protocols are the best way of exchanging information.<sup>10</sup>
- *Need to execute on very low specification devices*: agents may well be pieces of software of substantial complexity and not be able to run everywhere. This obstacle is gradually being removed by smaller footprint agent platforms and more performant network devices.

8. <http://www.sics.se/~onthemove/>

9. <http://www.tinac.com/>

10. This is not to say agents cannot also employ these protocols, however DAI purists might argue these do not completely fill the role of agent communication.

Passing down the network stack and to operations which need to be carried out at faster and faster time scales, one would expect control software to have less and less of the features listed in Section 1.1. However, this change is likely to be a continuum rather than a sharp break and deciding where the dividing line between “agent” and “non-agent” will perhaps become somewhat academic (Figure 3).

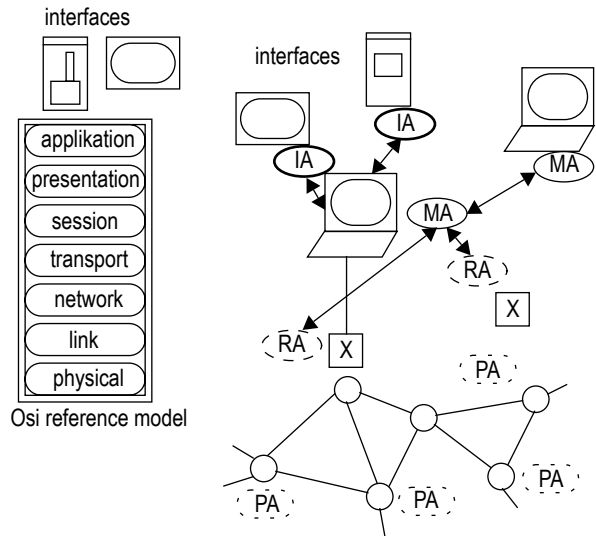
In summary, entities near the bottom of the stack (such as an SNMP agent or deployed service code) may be mobile and/or have very limited tasks whereas entities higher in the stack (such as top level managers or user interface agents) may begin to have properties which a DAI researcher might find agent-like. The two tier resource management model given in Section 3.2 illustrates this idea: agents are applied as layers of supervisory systems controlling layers of increasingly constrained and optimised on-line systems. This model is analogous to what already goes on in networks today and is characterised by [Musliner et. al. 95] as *intelligent reasoning about real-time processes*.

Aside from determining where in the network agents should be deployed, there are also wider considerations which need to be addressed before agent technology can realise its full potential in communication networks. The future deployment of agent technology rests critically on building increased cooperation between the traditionally separate DAI and CN communities. Apart from the terminology problems already mentioned, the division has created other obstacles to development:

- Continuing (to some extent justified) scepticism on the part of communications network engineers as to the utility and suitability (in terms of security, robustness, speed of operation etc.) of agent technology. This has resulted in a lack of tested practical solutions and many approaches which have never made it beyond the test bed stage.
- The biggest stumbling block for DAI researchers has perhaps been the technological complexity of the networks being studied. It would be fair to say that several of the promising methods developed by DAI researchers in the past have met with little success due to failings in the starting assumptions about the network domain.
- Agent solutions which have been proposed by the CN community have remained very simple and not leveraged some of the more powerful techniques developed by the DAI community.

There are indications that this collaboration is increasing and that the interests of the two communities are growing together. The papers presented at the Smartnet<sup>11</sup> and DSOM<sup>12</sup> workshops this year, for example, include a significant number of agent related papers. Furthermore, the continuing interest in agent technology within the OMG and FIPA standards bodies for example (both of whom list many member companies heavily involved in communication networks including: British Telecom, France Telecom, Nortel, Motorola and many others) is encouraging. The European Union AgentLink project<sup>13</sup> is

11. <http://www.cs.ait.ac.th/~ca/smartnet99/>  
 12. <http://www.tik.ee.ethz.ch/dsom99/>



**Fig. 3:** In an “agentified” model of the network software entities operate at different levels. Physical Agents (PA) might control specific network elements (as simple input-output sensors for example). Resource Agents (RA) might invoke changes in the switches and routers using information coming from both the higher and lower levels in the network. Mediator Agents (MA) might be more sophisticated entities needing to be able and inter-operate with other entities by using a common agent language. Finally, Interface Agents (IA) could translate from agent languages to more human understandable information. The lower down through the layers an entity resides, the less sophisticated it is and the less developed its “agent properties” might be considered to be - the decreasing sophistication is illustrated by an increasingly dashed line.

also contributing to this collaboration with a special interest group dedicated to the application of agent technology to telecommunications networks.

### 5 Conclusions

Having covered three areas which might greatly benefit from the application of agent technology and discussed some of the provisos in its application, we can conclude by advancing three main reasons for believing that future network developments may include the deployment of agents:

1. *Need for innovation:* Increasing competition, technological complexity and usage requirements are all contributing to increased strain on network infrastructure. This push is making innovative solutions (and potentially agent solutions) to network problems vital for ensuring continued good service.
2. *Technological feasibility:* Agent technology is maturing as a software paradigm. Alongside the increasing availability of development environments it is increasingly likely that deployed network equipment will in the future be able to support the computational needs of agents.
3. *Increase industry openness:* As the continuing collaborations within FIPA, OMG and many European projects show,

13. <http://www.agentlink.org/>

the communications network industry is increasingly open to experimentation with agent based solutions.

Together, these reasons suggest that there could be a slow agentification at least of the upper layers of the network infrastructure – little by little – agents may begin to appear in our networks. The technologically dynamic communications industry is however known for its frequent changes of tack so only time will tell.

Due to the limited space available this article can only give a brief overview of the subject area. We hope to have included enough pointers to literature (in particular see the survey articles referenced in Section 1) to serve as a useful starting point for further reading.

#### Acknowledgements

The authors would like to extend their thanks to the other partners in the SPP-ICC IMMUNE project (which partly funds this work). Funding for IMMUNE from the Swiss National Science Foundation<sup>14</sup> is also gratefully acknowledged.

#### References

- [Abu-hakima et al. 96]  
S. Abu-hakima, S. Liscano, and R. Impey. Cooperative Agents that Adapt for Seamless Messaging in Heterogeneous Communication Environments. In Proceedings of the AAAI- 96 Workshop on Cooperative Information Agents - Portland Oregon. AAAI Press, 1996.
- [Appleby/Steward 94]  
S. Appleby and S. Steward. Mobile Software Agents for Control in Telecommunications Networks. *British Telecom Technology Journal*, 12(2), 1994.
- [Baldi et al. 97]  
Mario Baldi, Silvano Gai, and Gian Pietro Picco. Exploiting code mobility in decentralized and flexible network management. In Proceedings of the First International Workshop on Mobile Agents, Berlin, Germany, April 1997.
- [Bhaniramka et al. 99]  
P. Bhaniramka, W. Sun, and R. Jain. Quality of service using traffic engineering over mpls: An analysis, March 1999. draft-bhani-mpls-te-anal-00.txt.
- [Bigham et al. 99]  
J. Bigham, L. Cuthbert, A. Hayzelden, and Z. Luo. Multi-agent system for network resource management. *Lecture Notes in Computer Science*, 1597:514-526, 1999.
- [Braden et al. 97]  
R. Braden, L. Zhang, S. Berson, and S. Herzog, S. and Jamin. Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification. RFC 2205, IETF Network Working Group, Standards track, September 1997.
- [Breugst/Magedanz 98]  
M. Breugst and T. Magedanz. On the Usage of Standard Mobile Agent Platforms in Telecommunication Environments. In S. Trigila et al., editor, Proceedings of 5th Int. Conference on Intelligence in Services and Networks (IS&N), Lecture Notes of Computer Sciences 1430, Intelligence in Services and Networks: Technologies for Ubiquitous Telecom Services, pages 275-286, Antwerp, Belgium, May 1998. Springer Verlag.
- [Calisti/Faltings 99]  
M. Calisti and B. Faltings. A multi-agent paradigm for the Inter-domain Demand Allocation process. DSOM'99, Tenth IFIP/IEEE International Workshop on Distributed Systems: Operations and Management, 1999.
- [Calisti et al. 99]  
M. Calisti, C. Frei, and B. Faltings. A distributed approach for QoS-based multi-domain routing. AiDIN'99, AAAI-Workshop on Artificial Intelligence for Distributed Information Networking, 1999.
- [Chen et al. 99]  
J. Chen, P. Druschel, and D. Subramanian. A new approach to routing using dynamic metrics. In Proceedings of INFOCOM99, 1999.
- [Corley et al. 98]  
S. Corley, M. Tesselaar, J. Cooley, and J. Meinkeohn. The application of intelligent and mobile agents to network and service management. *Lecture Notes in Computer Science*, 1430:127-??, 1998.
- [Finin et al. 93]  
Tim Finin et al. Specification of the KQML Agent-Communication Language - plus example agent policies and architectures, 1993.
- [FIPA 97]  
Foundation for Intelligent Physical Agents (FIPA). FIPA Agent specification 1997. Technical report, FIPA, homepage - <http://www.pa.org/>, October 1997.
- [FIPA 98]  
Foundation for Intelligent Physical Agents (FIPA). FIPA Agent specification 1998. Technical report, FIPA, homepage - <http://www.pa.org/>, October 1998.
- [Gibney/Jennings 98]  
M. A. Gibney and N. R. Jennings. Dynamic Resource Allocation by MARKET-Base Routing in Telecommunications Networks. In S. Albayrak and F. J. Garijo, editors, Proceedings Second International Workshop on Intelligent Agents for Telecommunications Applications IATA'98, pages 102-117. Springer (as LNAI-1437), 1998.
- [Hayzelden/Bigham 98]  
A. L. G. Hayzelden and J. Bigham. Heterogeneous Multi-Agent Architecture for ATM Virtual Path Network Resource Configuration. In S. Albayrak and F. J. Garijo, editors, Proceedings Second International Workshop on Intelligent Agents for Telecommunications Applications IATA'98, pages 45-59. Springer (as LNAI-1437), 1998.
- [Hayzelden/Bigham 99]  
A. L. G. Hayzelden and J. Bigham. Software Agents for Future Communication Systems. Springer Verlag, April 1999.
- [Jennings/Wooldrige 98]  
N. Jennings and M. (eds) Wooldrige. Agent Technology Foundations, Applications, and Markets. Springer/UNICOM, February 1998.
- [Kumar/Venkataram 97]  
G. P. Kumar and P. Venkataram. Artificial Intelligence Approaches to Network Management: Recent Advances and a Survey. *Computer Communications*, 20:1313-1322, 1997.
- [Labovitz et al. 98]  
C. Labovitz, A. Ahuja, and F. Jahanian. Experimental study of internet stability and wide-area backbone failures. Technical Report CSE-TR-382-98, University of Michigan Department of Electrical Engineering and Computer Science, December 16, 1998.
- [Lewis 95]  
L. Lewis. AI and Intelligent Networks in the 1990s and into the 21st Century. In J. Liebowitz and D. Prerau, editors, *Worldwide Intelligent Systems*. IOS Press, 1995.
- [Lloys/Pearmain 99]  
S. Lloys and A. Pearmain. Multi Agent System for Establishing 'Virtual Home Environments in the Convergence of Fixed and Mobile Telecommunications Networks. In S. Willmott and S. Abu-Hakima, editors, Proceedings of the 3rd Workshop on

14. Project Number SPP-ICC 5003-45311. See INFORMATIK/INFORMATIQUE 3/1999, p. 29.

- Artificial Intelligence for Distributed Information Networks, held at AAAI'99, pages 78-83. AAAI Press, 1999.
- [Maes 94]  
Pattie Maes. Agents that Reduce Work and Information Overload. *Communications of the ACM*, 37(7):31-40, July 1994.
- [Martin-Flatin/Znaty 2000]  
S. Martin-Flatin and J. P. Znaty. Two Taxonomies of Distributed Network Management Paradigms. In S. Erfani and P. Ray, editors, *Emerging Trends and Challenges in Network Management*. Plenum Press, New York, NY, USA, March 2000.
- [Musliner et al. 95]  
D. J. Musliner, J. A. Hendler, A. K. Agrawala, E. H. Durfee, J. K. Strosnider, and C. J. Paul. The Challenges of Real-Time AI. *IEEE Computer*, 28(1), January 1995.
- [Poslad et al. 99]  
S. Poslad, J. Pitt, R. Mamdani, A. Hadingham, and P. Buckle. Agent-Oriented Middleware for Integrating Customer Network Services. In A. Hayzelden and J. Bughma, editors, *Software Agents for Future Communications Systems*. Springer Verlag, 1999.
- [Schoonderwoerd et al. 97]  
Ruud Schoonderwoerd, Owen Holland, and Janet Bruten. Ant-like agents for load balancing in telecommunications networks. In W. Lewis Johnson and Barbara Hayes-Roth, editors, *Proceedings of the 1st International Conference on Autonomous Agents*, pages 209-216, New York, February 5-8 1997. ACM Press.
- [Smith 99]  
J. M. Smith. Programmable Networks: Selected Challenges in Computer Networking. *IEEE Computer Magazine*, 32(1):40-42, January 1999.
- [Subramanian et al. 97]  
D. Subramanian, P. Druschel, and J. Chen. Ants and Reinforcement Learning: A Case Study in Routing in Dynamic Networks. In *Proceedings of IJCAI'97*, pages 832-838. 1997.
- [Tennenhouse et al. 97]  
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.
- [Weihmayer/Velthuijsen 94]  
Robert Weihmayer and Hugo Velthuijsen. Application of Distributed AI and cooperative problem solving to telecommunications. In *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, pages 378-402, Seattle, WA, July 1994.
- [Weihmayer/Velthuijsen 98]  
R. Weihmayer and H. Velthuijsen. Intelligent Agents in Telecommunications. In N. R. Jennings and M. Wooldridge, editors, *Agent Technology Foundations, Applications and Markets*, pages 201-217. Springer Verlag and UNICOM UK, 1998.
- [Wellman 94]  
M. P. Wellman. A Market-Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems. *Journal of Artificial Intelligence Research*, 1(1):1-23, 1994.
- [Willmott et al. 99]  
S. N. Willmott, C. Frei, B. Faltings, and M. Calisti. Organisation and Co-ordination for Online Routing in Communications Networks. In A. L. G. Hayzelden and J. Bigham, editors, *Software Agents for Future Communication Systems*, pages 130-159. Springer Verlag, 1999.