Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Summer 1995 No. 37 été 1995

An official publication of CSCSI, the Canadian Society for Computational Studies of Intelligence
Une publication officielle de la SCEIO, la Société canadienne pour l'étude de l'intelligence par ordinateur

Special Issue
Artificial Intelligence and the Information Superhighway

Edition speciale
L'intelligence artificielle et l'autoroute de l'information

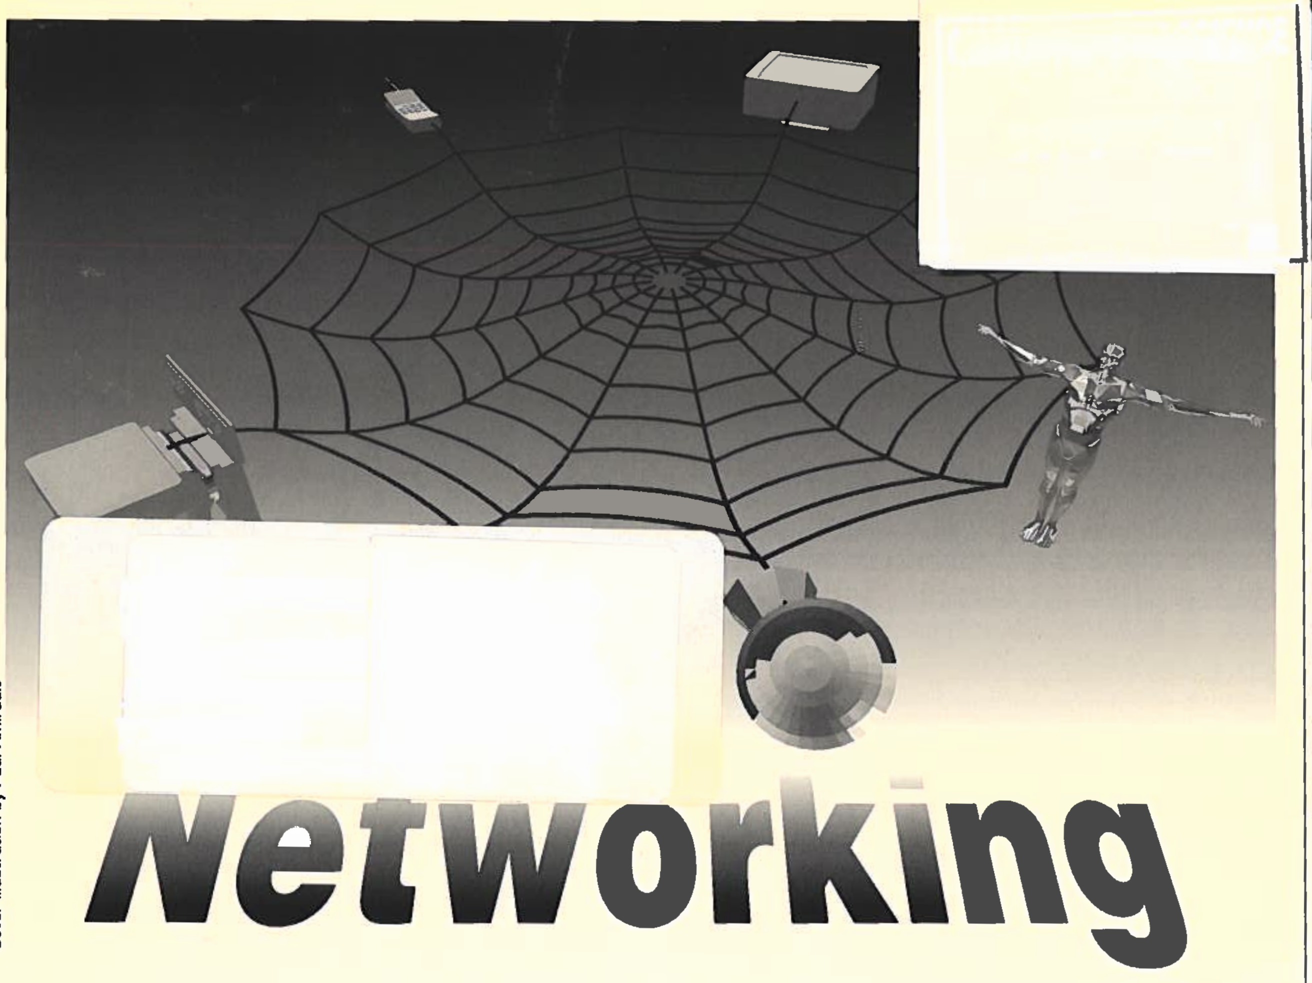

Cover illustration by Paul Amirault

Retrospective

Contents

CSCSI/SCEIO news: Call for nominations, and annual general meeting, and treasurer's report. 5
Fifth Generation Society to be formally constituted. 9
Strong Canadian role possible in machine translation and space station. 11
University of Ottawa and Cognos start major joint project *(Doug Skuce)*. 13
COMPINT conference in Montreal features AI and 5G *(Lois Carson)*. 14
Best-paper award established for Canadian AI Conference. 15
New Bindings. 15
Cartoon *(P.S. Mueller)*. 15
Research in Japan: A visiting researcher at the Institute for New Generation Computer Technology *(Randy Goebel)*. 19
Directory of AI graduate programmes in Canadian universities. 25
New books and journals: Book reviews, and abstracts from *Computational Intelligence*. 31
Recent technical reports. 35
Forthcoming conferences, and calls for papers. 40
Letters to the Editor. 42
All-purpose form. 43

Cover graphic, contents from Issue Number 6, December 1995
Canadian Artificial Intelligence Newsletter

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Summer 1995 | No. 37 | été 1995

Canada's National AI magazine.

Canadian Artificial Intelligence is published three times a year by the Canadian Society for Computational Studies of Intelligence (CSCSI). *Intelligence Artificielle au Canada* est publiée trimestriellement par la Société canadienne pour l'étude de l'intelligence par ordinateur (SCEIO). Canadian Publications Mail Product Sales Agreement No. 507032.

ISSN 0823-9339

CSCSI Executive

Contents / Contenu

Contents		Contenu
Communications	2	**Communications**
Feature Articles		**Gros Titres**
The Information Highway: Reality, Myth and the Role of AI — *Sue Abu-Hakima*	4	L'Autoroute de l'information: Réalité, mythe, et le rôle de l'intelligence artificielle — *Sue Abu-Hakima*
From Current to Future Telepresence Technologies (Did the Interstate System Kill Route 66?) — *Gerald M. Karam*	8	Les technologies de la téléprésence d'aujourd'hui et de demain (Le système Interstate a-t-il détruit la route 66?) — *Gerald M. Karam*
Artificial Intelligence and Distributed Information Networks in the United States — *Jon R. Wright*	18	L'intelligence artificielle et les réseaux d'information distribués aux États-Unis — *Jon R. Wright*
An Agent-Based Multimedia Service Environment — *Michael Weiss, Tom Gray, Aurora Diaz*	22	Un environnement de service multi-média basé sur les agents — *Michael Weiss, Tom Gray, Aurora Diaz*
Students' Corner	28	**Le coin des étudiants**
Towards a Cognitively Plausible Model for Quantification — *Walid Saba*		Vers un modèle de quantification cognitivement plausible — *Walid Saba*
PRECARN Update	31	**Nouvelles de PRECARN**
Book Reviews	33	**Critiques de livres**

Canadian Artificial Intelligence welcomes submissions on any matter related to artificial intelligence. Please send your contribution, electronic preferred, with an abstract, a photograph, and a short bio to:

Dr. Peter Turney
Co-Editor, *Canadian Artificial Intelligence*
Institute for Information Technology
National Research Council Canada
M-50 Montreal Road
Ottawa, Ontario, Canada
K1A 0R6 or — peter@ai.iit.nrc.ca
Telephone (613) 993-8564 Fax (613) 952-7151

Advertising rates are available upon request.

Arlene Merling
Managing Editor, *Canadian Artificial Intelligence*
Alberta Research Council
3rd Floor, 6815 - 8 Street N.E.
Calgary, Alberta, Canada
T2E 7H7 or — merling @arc.ab.ca
Telephone (403) 297-2608 Fax (403) 297-2505

Book reviews and candidate books to review should be sent to:

Graeme Hirst
Canadian Artificial Intelligence
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4 or — gh@ai.toronto.edu

Recycled / Recyclable

COMMUNICATIONS
COMMUNICATIONS

Special Issue: Artificial Intelligence and the Information Superhighway

Sue Abu-Hakima

Welcome to this special issue on the information superhighway. Over the last few years we have seen many articles on this topic in the popular media. In the first article in this issue I will explore the reality, myth, and the role of artificial intelligence on this so-called highway. In the article I present a historical perspective, the players, the potholes, and the silver lining for AI.

The second article is on the Telepresence project and was written by Gerald Karam, the former Director of the project who was, until recently, a Professor at Carleton University's Department of Computer and Systems Engineering. He recently accepted a position at AT&T Bell Laboratories in New Jersey. Telepresence was a three-year pre-competitive project, aimed at designing and field testing advanced media systems in the workplace to gain insight into sociological and engineering issues. It was mostly funded by the Province of Ontario through TRIO and ITRC. The project included collaborative work between the Universities of Toronto and Carleton as well as Bell Canada and Newbridge Networks Corporation. Three commercial products are being spun off from the project: Applied Silicon International's Video Vise, the Active Desk from Arnott Design Group, and Corel Video from Corel Corporation. Thus far, AI did not have a role in the Telepresence project; however, opportunities for its application are explored in this interesting article.

The third article is on AI and Distributed Information Networks in the United States from Jon R. Wright at AT&T Bell Laboratories. The article reviews the current state of the public networks (used for telephony and computing). It then reviews the role of expert systems and software agents in these networks. AI in network and service provisioning is also described. Finally, the role of AI in new services is discussed.

The fourth article in this special issue is from Michael Weiss and his colleagues at Mitel Corporation. It describes the Multi-agent Architecture for Networking Applications. This architecture has been applied in the collaborative pre-competitive project between Mitel, NRC, and Phoneix in a CANARIE-funded project entitled the Multimedia Help Desk. The Help Desk project which was recently concluded successfully consisted of a set of distributed resources. The resources at the Mitel site in Kanata and at the NRC site in Ottawa were linked by OCRINet (a high speed network).

We would be most interested in your feedback on this special theme issue. Please send us electronic mail at suhayya@ai.iit.nrc.ca with your comments.

AGENDA

CSCSI Annual Meeting

August 23, 1995, at IJCAI-95, 2:00 - 4:00 p.m.

(Room number to be announced in the final IJCAI program)

1. Acceptance of minutes of 1994 annual meeting
2. Treasurer's report
3. Membership update
4. Conferences and workshops
5. Canadian AI Magazine - going electronic?
6. Other business

Minutes of the 1994 CSCSI Annual Meeting

Date: May 18, 1994

Location: AI/GI/VI '94 Conference
Banff, Alberta, Canada

Present: Approximately 30 members

The meeting started at 5:15 p.m.

The Minutes of the previous meeting (1993) were accepted unanimously.

Treasurer's Report: The Treasurer's report was discussed. It was generally agreed that the Society should break even on the conference, and that the magazine continue to be the main expense. Motion to accept the treasurer's report as distributed. Proposed: Robert Holte. Seconded: René Elio. Motion passed unanimously.

Membership: Janice Glasgow reported the following figures from CIPS:

Total Membership July 1992	460
Total Membership July 1993	482
New members that period	50
Total Membership April 1994	470
New Members 1993	
January - March	12
April - June	2
July - September	14
October - December	9
New Members 1994	
January - March	11
April	4
Current membership (April 1994)	470
Dropouts	50
New Members	38

It has been raised by a number of people that efforts have to be undertaken to maintain and increase the membership. The perennial issue of CIPS's ineptitude at soliciting renewals was raised by several participants. Janice Glasgow suggested that in the mailings for AI' 96 the Society Membership form be enclosed with the Conference Registration form. Stan Matwin volunteered to work on an on-line, ftp-able version of the form.

IJCAI: In the absence of Renato DeMori, Janice Glasgow appealed for suggestions for Canadian representation (invited speakers), as well as for ideas on fund-raising. Gordon McCalla is in charge of IJCAI tutorials. He asked for suggestions for tutorials by Canadian presenters in either official language.

AI' 94 Conference Report: René Elio suggested that more precise budget guidelines will be helpful for future conferences (e.g., how much to spend on external advertising). She also recommended more specific guidelines for division of tasks between the conference chair and the program chair. She raised the question of whether the different workshop fees should be unified. Several people argued for the usefulness of a poster session in future conferences. Inconclusive discussion on the benefits of combining such a session with the conference reception followed. AI '96 will likely be held in the Toronto area (again jointly with GI/VI), and the new executive will search for the Program Chair.

Magazine: Janice Glasgow reiterated that the magazine is the main benefit for members, as well as the main expenditure in the budget. A number of ideas was discussed informally:

- an on-line version (Robert Holte)
- offering the magazine to non-members for a fee
- asking IRIS to distribute the magazine
- looking into distribution options alternative to Canada Post
- soliciting support for the magazine from Precarn

New Executive: Motion that the new executive be :
President - Stan Matwin, University of Ottawa
Past President - Janice Glasgow, Queen's University
Vice-President - Allan Jepson, University of Toronto
Secretary - Fred Popowich, Simon Fraser University
Treasurer - Peter van Beek, University of Alberta
Magazine Editors - Peter Turney and
Suhayya Abu-Hakima, NRC
Proposed: Nick Cercone. Seconded: Peter Patel-Schneider
Motion passed unanimously.

Computational Intelligence: On behalf of the Executive Board, Nick Cercone reported that there are approximately 340 subscriptions to the journal, of which under 100 are individual, the rest institutional. He strongly suggested that the benefit of reduced subscription ($45 instead of $90) is more aggressively advertised to the membership. The turnaround has decreased from 18 months to 6 months. The authorship of the papers is approximately 30% Canadian.

The meeting adjourned at 6:15 p.m.

S. Matwin, Secretary

The Information Highway: Reality, Myth, and the Role of AI

Sue Abu-Hakima

Sommaire

Depuis quelque temps, les medias portent une attention énorme (certains diraient démesurée) à l'Autoroute de l'Information. On se propose dans cet article de séparer la réalité du mythe entourant l'autoroute. On y présente d'abord quelques notes historiques, avec quelques mots sur là où l'intelligence artificielle a été appliquée et les succès obtenus. Suit une discussion sur les dernières directions (le dernier cri) dans l'application de l'intelligence artificielle pour l'autoroute. On conclut l'article avec une discussion sur les nids de poule, les intervenants (possiblement ceux qui décident des tendances) et la doublure en satin de l'intelligence artificielle.

Introduction

For some time now, we have had a tremendous amount of media attention (some may call it hype) regarding the information highway. This article is intended to separate the reality from the mythology of the highway. Some historical notes are first discussed including some thoughts on where artificial intelligence has been applied and its success. This is followed by a discussion of what are the latest directions (fads) in applying AI for the highway. A discussion of the potholes, players (possibly the trend setters), and the silver lining for AI concludes the article.

Historical Notes

Let us first discuss the term "information highway." It is unclear to this author how the "information highway" we have had for decades can be so blatantly ignored in a rush to embrace the new mythology of yet another highway. For over 100 years, we have had the telephone which is itself a conduit for information. In its early years, it was simply used for communication in the form of voice call discussions between people. It provided people with the opportunity to interact remotely, speak at great length about the happenings of the day, and to stay "in touch" with each other across continents. Several decades later when computers were invented, a need for communicating data across telephone networks was established and modems were invented. Not much later, the Arpanet, a data network of servers, was set up by the US Department of Defence as a means for its contractors to communicate with each other. This was established 25 years ago and has since evolved into the network we know as the Internet [1]. The Internet now has over 30 million users worldwide and has proven to be a very powerful communication network that has evolved to carry voice and video as well as data (text).

The Internet, like the telephone, allows people to interact with each other across the globe. People use the Internet for electronic mail or participation in various newsgroups. Over the last few years we have also seen the emergence of the world wide web (www) which links together servers globally over the Internet and allows users to interact with these servers through graphical hypermedia interfaces, such as Mosaic or Netscape. Netscape is a commercial alternative to the public domain Mosaic with more features for users. Currently on the Internet, one can browse (or surf) the web, download voice or video clips, send electronic mail to any of the world's continents, participate in newsgroups, market services, and buy or sell a variety of products. Entire communities are also putting together networks called FreeNets which are attracting people who normally shy away from computers to participate in a variety of discussions.

Thus, given the long history of steady evolution in communication technology, why has there been so much media hype recently about the information highway? From a technology perspective, the information highway can simply be viewed as another step in the evolution of communication networks driven by the requirement for more bandwidth needed for the rapid exchange of video, voice and data. My hypothesis goes beyond this to debate that several events and advances have converged to drive the media hype beyond the simple evolution of technology. An event in 1991 that may seem only vaguely related was quite important. This event, of course, was the live broadcasting of the Gulf War. In real-time we were shown video-game-like missiles shot back and forth between the two sides on a far continent through live satellite and telecommunication broadcasts. This event spawned further discussion of the new world order and the reality of the global village which also coincided with the breakup of the former Soviet Union. Around the same time, the deregulation of the Cable and Telecommunications Industries in the USA continued. Suddenly, telephone carriers were buying out cable companies and quickly entertainment media giants were formed (examples include USWest and Time Warner, and McCaw and AT&T).

Another major technological push was the availability of products for putting together fiber optic high speed networks linked together by switches that could packetize voice, data, and video. These switches are based on SONET and Asynchronous Transfer Mode (ATM) standard protocols which allow networks to switch high bandwidth information (e.g. video) in real-time. At the same time, consumers have been getting more and more technology savvy through their use of interactive electronic games from Sega and Nintendo, as well as the Internet and the www from their home personal computers, and the other type of ATM machine - namely the automated teller which provided them with cash-on-demand. I must remind the reader that this convergence has occurred over five years and intensified in the last two (note that the number of Internet users has been growing at the rate of 20% per month). Thus, the term "information highway" was coined by the media as a means of explaining this connected universe of information embodied by Internet that was now at people's fingertips.

Past Applications of AI in the Information Highway

All this activity raises the question of where artificial intelligence fits in? AI has not really had a leading role on the information highway. However, AI has been actively pursued in a number of activities connected with the infrastructure of the information highway. For example, in the telecommunications industry which provides the global networks for the highway, AI has been limited to secondary applications. These include the diagnosis of complex equipment and the filtering and interpretation of the thousands of network alarms. AI techniques in these applications include model- and fault-based diagnosis (the earlier forms of which were known as expert systems), neural networks, and classification learning. Another technology in AI has been the development of natural language (NL) parsing as front ends to warehouses of information or databases. Natural language allows users to construct queries in a restricted form of English to search and retrieve information when their queries are parsed. Full Natural Language has not been achieved in AI and many believe that it will be unachievable and that constrained NL is more likely to produce successful results.

AI Successes on the Information Highway

It is first important to attempt to identify what a success could mean for an AI researcher or developer. For a company, an AI success is tied to the bottom line - that is, did the company make a profit from the AI project? There have been successes for a variety of companies as exemplified by the AI success stories that are written up in this magazine on a regular basis. One could generalize that most of the successes that have been deployed in AI on the information highway have been expert systems for diagnosis or sales support in the telecommunications industry.

A researcher would view an AI success as the formulation and publication of a new theory. There have obviously been new theories put forward in AI as is apparent by how prolific researchers have been in AI conferences and publications. However, there are no outstanding theories in AI that have been advanced for the information highway. Rather, there have been some relevant incremental publications on a variety of aspects of AI that include machine learning, diagnosis, distributed AI, and intelligent agents.

Another measure of success for a researcher is the publication of a doctorate thesis on a particular topic. Again, in relation to the information highway there have been no significant publications of theses but I anticipate that many will be generated before the year 2000.

A final measure for successful AI research is the receipt of grant money from funding agencies such as NSERC in Canada or ARPA in US. There has been quite a significant amount of grant money awarded in recent years for research in the information highway but again this is not specifically targeted at the application of AI.

There have not been any significant successes in AI on the information highway. Thus, the conclusion for researchers applying AI to the information highway is that there is plenty of opportunity over the coming years to do outstanding research in this area.

New directions for AI on the Information Highway

One question that both researchers and developers of AI technology for the information highway ask is what are the hot areas to work in? One of the main areas to work in these days is two related subfields, distributed AI (DAI) and intelligent agents. These two approaches are natural candidates as they address the distributed nature of information networks. There is research to perform on the side of the ideal architectures and mechanisms of intelligent agents that can scour the highway for information for users. In addition, agents that can negotiate with other agents on behalf of a user (taking into account all the privacy and security issues) are a real challenge to develop. Furthermore, agents that can help users deal with heterogeneous environments rather than artificial homogeneous ones are badly needed.

Another area that has had considerable activity but has not yet fulfilled its promises on the information highway is machine learning. There are a number of problems that machine learning needs to address. Machine learning has had some success in intelligent alarm filtering for diagnosis of networks. Another promising area includes the intelligent filtering and classification of information for users so that their preferences are taken into account by the tools that assist them in their search for information on the www.

Another area that has been active is that of intelligent human computer interaction (HCI). This area includes work that ranges from natural language parsing to the encoding and use of multimedia knowledge in driving user interfaces. This is an important area for the information highway given the emphasis on addressing important user issues in comprehensibility.

A promising area is the use of AI in network and fault management. The networks that will support the information highway will understandably be quite complex due to their heterogeneous nature. These networks will require sophisticated fault prediction and management techniques that include the use of hybrid AI approaches that encompass machine learning, intelligent agents, sophisticated fault and model-based diagnosis algorithms as well as intelligent HCI for both end users and network administrators. One should not forget that the hybrid AI solutions will most likely be embedded AI solutions where AI forms a small part of the overall solution which requires it interworking with traditional software engineered systems.

Potholes in the Information Highway

As described above, AI holds a lot of promise along with the information highway. However, there are pot holes that practitioners should be wary of. One of these pot holes is formed through the unrealistic expectations of users due to media hype. Consumers are being told that they will soon be able to order their groceries, their favourite films, receive individually tailored electronic newspapers, and do all their banking through their remote control. The production of such hype is irresponsible without presenting some of the pros and cons of these information highway conveniences. Will consumers really be as excited if we tell them that we have not really come up with solutions for the possible breaches of privacy and security these services come bundled with? Also, people interpret the hype from their point of view, not from that of technocrats, so a mismatch of what is delivered versus what is expected will surely occur. An age old problem in computing, not only in AI, that we need to continuously remind ourselves of is that we often ignore naive end-user needs and preferences and develop technology as it suits us. Thus, we need to develop intelligent human-computer interfaces that match user expectations. One example out of today's information highway is that many Mosaic and Netscape users are somewhat frustrated with their interfaces as they are not as intuitive or very useful for quickly browsing information as some of the techies building them believe they are.

Another unpredictable pothole is regulation [3]. In Canada, the CRTC has recently had a series of committee hearings on the information highway and come up with a set of recommendations for the government to base policies on [2]. These recommendations include regulation and deregulation points for the cable and telecommunications (wired and wireless) industries. Instead of simplifying the current scene and facilitating the convergence of these industries so that consumers may be provided with services from their wildest dreams, it looks like regulation will be a bottleneck for some years to come, with the consumer expected to pay higher fees for existing services so that they may support any new services. Consumers are a fickle bunch; fees are not what they are looking for. One only has to recall Rogers Cable introducing a new fee structure through negative marketing for new cable channels. There was such a negative backlash that pushed many consumers to seek direct broadcast satellite service and do away with cable that such negative marketing is being made illegal.

Another threat that is on the horizon for consumers from the telephone carriers is their wish to charge based on per-call use. This would imply that users would be charged for local calls on per-minute usage thus raising their local telephone rates. The CRTC has approved this billing scheme for business users but not yet for consumers. Such levies will likely cause another battle between the consumers, the regulators, and the providers of the information highway since it will affect consumers surfing the Internet off their telephone line modem connection. It is important to understand the volatility of the marketplace and that consumers have the ultimate power to accept or reject new services or charges for old services.

Still another pothole is that many of the media giants in cable, telecom, or in entertainment tend to ignore user needs in the pursuit of higher profits. This is apparent in the push to provide end-to-end fiber to the home for the provision of new services. In the rush to compete, it looks like several providers will bring fiber to the home. This in itself is a very expensive venture and in some ways may be unnecessary. The assumption is that users everywhere want high bandwidth connections that bring in videos and games to their living rooms. What the giants are ignoring is that users have flocked to the Internet and are in many ways becoming addicted to interacting with other people and downloading multimedia information off the www. The Internet does not currently entail expensive cable charges and offers interactive services. Thus, why is the assumption that consumers want to be couch potatoes the premise by which many of the media giants are developing their strategies? [3]

Another pothole is that currently in Canada (and I am certain globally), many companies are competing rather than cooperating to develop the information highway infrastructure and services. I believe Canada would have a much stronger presence worldwide if some of these companies get together and work in a pre-competitive manner on some of the technologies and services. NRC has had some success in bringing together companies that are competitors so that they may work on pre-competitive R&D. A natural area for this in AI is the use of intelligent agents. Most of the players agree that this technology will have an important role in the future of the information highway. There exists an opportunity for companies to get together and work collaboratively with researchers to standardize agent technologies and mechanisms. Companies could then apply the technology in a manner unique to their market segments. The cloak and dagger approach to technology development is outdated and goes against the grain of the openness of the information highway.

A final pothole is that of largely ignoring the security issues and the wide potential for fraud and theft of electronic dollars. Asynchronous Transfer Mode allows the high speed transmission of packets of information worldwide. Global ATM networks will eventually (in 10 years or so) be used for global electronic commerce. This implies that credit card, banking debits and credits, and other financial information will be traveling the highway. In fact, there are now service providers on the Internet that require credit card numbers in exchange for goods and services, so the concern is immediate. Have we really mastered the technologies we need to keep such information secure and private? I suspect that we have not, and we should make security a priority.

Strategies of the Various Players

The players on the information highway are the telephone carriers (such as Bell Canada and Stentor - the consortium of provincial telephone operating companies), the cable providers, the telephone network inter-exchange carriers (such as Unitel, Sprint, etc.), the cellular network providers (such as Bell Mobility, CanTel, etc.), and last but not least the information technology developers (software houses). If we are to believe the hype, all the networks are converging. This implies that services that were previously only available by telephone will soon be available through cable (an interesting example of this is the provision of Internet service through cable which was recently announced by the cable companies in Canada and is already available in the US). However, there is a small detail that will delay full convergence - the scramble for deregulation of the cable and telephone industries to allow them to compete in each other's back yards has not been successfully concluded by the CRTC. For at least three years, the telephone companies will not be able to compete with cable companies and in addition the telephone companies are expected to provide the infrastructure for the cable companies to compete. AI has a role to play in network management applications in the converging networks. The time lag in deregulation provides the community with a window of opportunity for new R&D.

Another set of players are the content providers which include media giants, movie studios, telcos, and a variety of computer companies. To date, the assumption has been that a significant portion of the content will be home shopping network, videos-on-demand,and games-on-demand for home entertainment (more couch potato services?!). However, as emphasized by the information highway recommendations by the CRTC, the government wants many of its services provided electronically to Canadians nationwide. Thus, information structuring, search, retrieval, browsing, filtering, and presentation tools must be researched and developed for consumers to surf government information. This information will likely require intelligent interactive interfaces (multimedia question and answer). In Canada, there is room for growth in the industry providing these technologies.

Finally, there will be new players on the scene experimenting with technologies we thought impossible. A good example of this is the Internet VideoPhone which promises to be one of those new technologies that will break an area wide open since it permits the use of the Internet for international telephone calls.

The Silver Lining

The silver lining is apparent by observing the interactivity that has exploded on the Internet this year. One amazing piece of folklore is that when the war in Yugoslavia broke out and the TV was down and the radio was down - the Internet remained up and allowed people some contact with the rest of the world.

People need to be interactive. They do so by sharing problems with people like themselves, starting newsgroups, falling in love, etc. over the Internet. The lesson here is that for all the people out to make the megadollars on the information highway - do not forget the basic human needs for social contact. Not all the people on the highway will get off on the AI-based Virtual Reality or Negroponte's Daily Me; they will most likely prefer the human contact at the other end of the net.

We the researchers and developers of AI technology have a responsibility and a great opportunity to provide these users with adaptable intelligent tools that facilitate interaction. AI has not had much success in information highway applications but there are tremendous opportunities in developing hybrid solutions of AI with other technologies. We also have a responsibility to minimize the hype so that we may avoid the inevitable backlash which we have felt before.

Acknowledgments

I would like to thank Ken Grigg of Mitel Corp., as well as Innes Ferguson and Peter Turney of IIT at NRC for their helpful comments on the first draft of this article.

References

[1] Jim Carroll and Rick Broadhead. Canadian Internet Handbook. Published by Prentice Hall Inc. (1994)

[2] Information Highway Advisory Council. Approved Recommendations to Date. CRTC (Draft Document) from Industry Canada (1995)

[3] Joseph Kraemer. The Realities of Convergence: a perspective on how to avoid being roadkill on the global information highway. Thought Leadership Series document from EDS Management Consulting Services (1994)

Sue Abu-Hakima *is the Group Leader of the Seamless Personal Information Networking (SPIN) Group at the Institute for Information Technology at NRC. She has been a researcher at NRC since 1987. She had previously worked at BNR starting in 1982. She has Doctorate and Masters degrees in AI from Carleton University's Department of Systems and Computer Engineering. She also has a B.Eng. from McGill University's Department of Electrical Engineering with specialties in Communications and Computers. She is on the Editorial board of the International Journal of Human-Computer Studies and is chairing an IJCAI-95 workshop on AI in distributed information networks.*

From Current to Future Telepresence Technologies (Did the Interstate System Kill Route 66?)

Gerald M. Karam

Résumé

La téléprésence est l'art de rendre possible la proximité sociale peu importe les distances géographiques ou temporelles, par l'intégration des technologies de l'ordinateur, de l'audio-visuel et des télécommunications. Le Projet de Téléprésence de l'Ontario (OTP - Ontario Telepresence Project) a mis au point une variété de logiciels et d'équipements pour supporter des applications de téléprésence, comme: (1) le TMS (Telepresence Media Space), un système de communications synchrones audio-visuelles intra-site et inter-site gérées par ordinateur employant des caméras et des moniteurs placés stratégiquement, (2) Postcards, un outil de contrôle de l'arrière-plan utilisant des clichés des utilisateurs pris à faible fréquence (basé sur Portholes de Xerox), et (3) la gestion audio-visuelle de salle. Cet article décrit l'approche du OTP et son expérience dans la transition des technologies de télécommunication actuelles aux technologies futures.

On y fait de nombreuses observations sur des applications telles que: réunions point à point et multipoints; bureau ouvert virtuel; espaces publics partagés; et séminaires à distance. On discute ensuite du rôle des agents intelligents dans l'amélioration de la transparence des équipements comme d'une étape importante dans le succès futur de ces applications.

Abstract

Telepresence is the art of enabling social proximity despite geographical or temporal distances through the integration of computer, audio-visual, and telecommunications technologies. The Ontario Telepresence Project (OTP) has constructed a variety of software and hardware systems to support telepresence applications, including: (1) the Telepresence Media Space (TMS), a system for computer-mediated intra-site and inter-site synchronous audio/video (A/V) communications employing strategically placed cameras and monitors, (2) Postcards, a background awareness tool using low frame rate snapshots of users (based on Xerox's Portholes), and (3) room level A/V management. Beginning with conventional 112 Kb/s H.261 dial-up A/V codecs for long haul connections, OTP moved its applications to a trial 45 Mb/s ATM network for long haul A/V and data connections and explored a number of uses of its technologies in the context of this emerging telecommunications capability. This paper describes the OTP approach and experiences in the transition from current telecommunications technology to future technologies. Numerous observations are made for end user applications such as: point-to-point and multipoint meetings; virtual open office; shared public spaces; and remote lecturing. It then discusses the role of intelligent agents to enhance the seamlessness of the applications as an important step in their future success.

Keywords: *ATM; multimedia; distributed systems; computer-supported-cooperative-work; telepresence; computer networks; analog and digital video; intelligent agents.*

1. Introduction

Telepresence is the art of enabling social proximity despite geographical or temporal distances. It is a set of computer, audio-visual, and telecommunications technologies, which are carefully integrated to permit people to work together using technology as an intermediary [1, 4]. In the activities of the Ontario Telepresence Project (OTP), [13] we have constructed a variety of software and hardware systems to support telepresence applications and it was the mandate of the project to deploy these applications in the field to evaluate their acceptance and use by people in an office environment. The applications that were developed included: (1) the Telepresence Media Space (TMS) — a computer-mediated intra-site and inter-site synchronous audio/video (A/V) communications employing strategically placed cameras and monitors, (2) Postcards — a background awareness tool using low frame rate snapshots of users (based on Xerox's Portholes [6]), and (3) the Desk Area Network (DAN) — for room level A/V management.

In the last phase of OTP research, several of the applications were re-engineered to function over a 45 Mb/s ATM link between the OTP lab at Carleton University in Ottawa, Canada, and the OTP lab at University of Toronto, in Toronto, Canada. As a result, we had an opportunity to "live in the future" and study the ways in which current and future technologies for collaboration might be combined and interact with one another. As our original long distance A/V link was handled by conventional 112 Kb/s dial-up video (H.320 compliant), this change to broadband A/V transmission was akin to moving from an old 4-lane blacktop highway, such as the famed "Route 66," to a brand new, high-speed, state-of-the-art Interstate highway, such as I-40 that snakes its way across the American southwest, bypassing much of the roadway originally serviced by Route 66.

Connoisseurs of American highway history will know that the Interstate system did indeed kill Route 66, but this did not happen quickly, nor did it happen completely. The

OTP experience with both current and future A/V transport mechanisms replays much of the history of transition from Route 66 to interstate highways, and the analogies serve as historical reminders for the telecommunication service providers, application developers, and new businesses that will spring up around the so-called "Information Superhighway."

The OTP trials over conventional and ATM networks have led to some understanding of the issues that will make or break synchronous collaborative applications, and in many cases, raised more questions than were answered. Of significant importance is the requirement for intelligence in the telepresence applications in order to serve the needs of individual users in a natural and seamless manner. We have made some rudimentary steps in this direction, but more work is needed, and in particular the exploitation of research in intelligent agents.

2. Telepresence Technologies

The goals of the Ontario Telepresence Project were two-fold: (1) to study methods of deployment for telepresence applications in workplace environments, the way in which telepresence technology is incorporated into workpractice, employ iterative, user-centered design, and to recommend design goals and strategies for telepresence technologies based on this work; and (2) develop telepresence technologies; evaluate them in experimental situations, and hopefully in field studies in a workplace environment, and then recommend engineering trade-offs. The philosophy of application development was to use off-the-shelf components to create a possible future workplace.

The engineering phases of the project produced a number of applications that were tested, deployed for use within the project, and in some cases in field study sites. The internal use of applications often provided a good first order evaluation because the physical organization of OTP had been established to require the use of collaborative tools. Specifically, half of the project was located in Toronto, and the other half was located in Ottawa — a separation of some 270 miles; furthermore, in Toronto some members of the project were located on the second floor of the building that housed the project, and some on the fourth floor — this distance was sufficient to make a casual "walk-by" rather tiresome. Communications among staff members was fostered by splitting the development and management teams horizontally, thus the applications engineers and systems engineers were split across sites, and the various managers who coordinated the research effort were similarly divided. In short, OTP was a testbed for its own tools, by design.

2.1 The Telepresence Media Space (TMS)

The phrase "media space" captures the idea that audio and video devices (sources, sinks) and transmission are managed collectively for a site (a "local area," such as a floor or a building). This is different from typical notions of desktop video conferencing through a workstation or conventional conference room video conferencing. OTP uses the concept of a media ecology, in which A/V is not locked in a box or room for a specific purpose, but in fact exists where and when it is needed, to best serve the objective of maintaining social proximity. For example, a meeting between three people in a private office, including one remote visitor, is not well served by a four square inch image of the visitor on someone's workstation that is facing away from the meeting area. Not only is the video surrogate (camera, visual display, microphone, and speaker) of the visitor improperly placed in the room, but the image is too small to maintain a "presence" in the room; thus, the visitor quickly becomes a disenfranchised member of the meeting. The OTP approach is to use a separate video surrogate for each distinct social function (such as face-to-face meeting, small group meeting, visitor "dropping in", etc.). This multiplicity of functions and A/V devices thus demands a more sophisticated approach to managing devices, connections, and services.

The essence of the TMS is a site of A/V devices controlled by an audio/video private branch exchange (A/V PBX). The A/V PBX is composed of software servers running on a host computer, that control a number of devices, most notably the A/V switching fabric. The architecture of the TMS is shown in Figure 1. The key components are a set of conventional off-the-shelf devices, as follows: (1) a node — that includes one or more cameras, monitors, microphones, speakers, associated with a client computer, that executes the TMS client applications; (2) an Ethernet LAN that

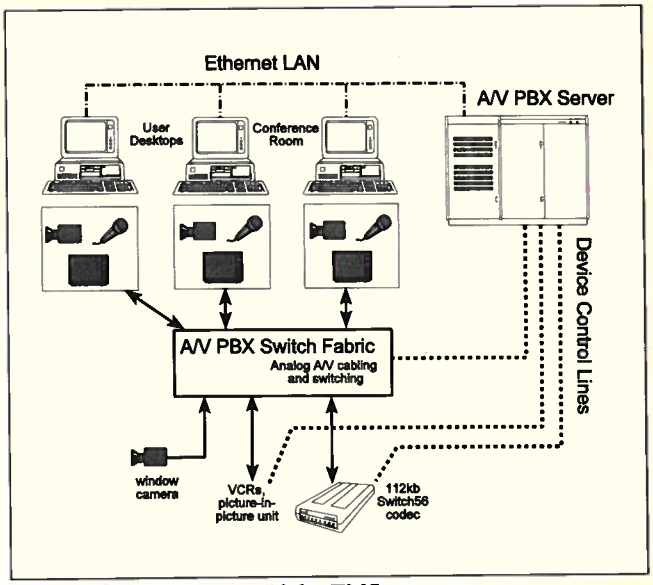

Figure 1. Architecture of the TMS

interconnects the client computers and the A/V PBX host computer; (3) an analog A/V network connecting all A/V devices (most with bi-directional circuits) to the A/V switching fabric (one or more analog A/V switches); (4) various VCRs, audio mixers, preamplifiers, and a picture-in-picture unit (used for multipoint meetings); and (5) an H.320 compatible dial-up-video codec operating over two

Switch56 (2 x 56 Kb/s) or two ISDN (2 x 64 Kb/s) data lines provided by the public switched telephone network.

The TMS provides a number of different functions to the users of the system. Some of the main services are summarized below. Note that these functions are described first in the context of single site, and then later extended to a multi-site model.

Point-to-point and multipoint meetings. Users can contact one another from a node (which includes both conference rooms and desktops), by using the TMS client application shown in Figure 2. User names appear in a list, that is used to select a person with whom to communicate, and then the contact button is used to establish the connection. A door state is used to denote a user's accessibility to others; there are four states: (1) door open — others may enter the user's virtual space without explicit confirmation; (2) door ajar — others must "knock" before entering and receive a positive confirmation before a connection is established, however short segments of video-only (glances or Postcards) could be used without explicit confirmation; (3) door closed — like door ajar, except that no video segments are available; and (4) door locked — no connections are permitted. A user's current accessibility is shown by an icon in the upper left corner of the application and the accessibility of other users are shown by icons beside their names in the name list. Once a connection is established, a user can add other users into a multipoint meeting by highlighting their name and using the contact button again. Multipoint meetings are accomplished by using a picture-in-picture unit to divide the video screen into four quadrants, where each participant (up to four) occupies one of the quadrants. The audio signals are mixed so that an audio source receives a combination of all audio sources except it's own.

Shared video devices. Audio/Video devices can be shared by users for a variety of purposes. For example, VCRs can be controlled and shared by users (although this was not

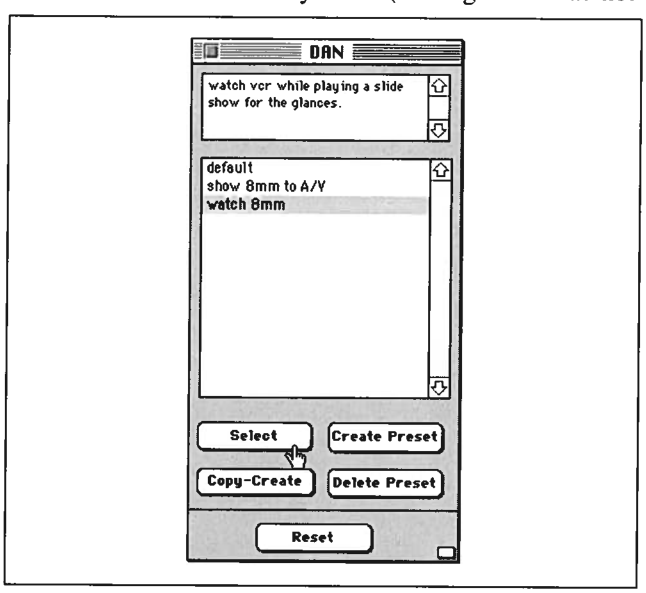

Figure 2. TMS Client Application

built into the simple interface of Figure 2). A video tape could be placed in a common VCR and users can share the VCR to watch the tape as a background activity, and yet still be able to make and accept foreground connections. The simplest and most unassuming shared device is the window camera. In this case, a camera is pointed through a window that shows a live outdoor scene to the community. Users, particularly those in windowless offices, can select the window camera as a background shared device, and thus use their video monitors to gain a sense of the weather or time of day, and share a common view with members of the group. The devices button on the user interface gives access to shared devices.

Integrated desk-level and room-level interactions. The TMS does not distinguish between nodes in private offices and those in conference rooms (except for names). This allows users to participate in conference room meetings from their desks, either as active members (asking questions, presenting information), or as passive members (listening only, in which case the conference room nodes can be treated as shared devices). Furthermore, users may enter a conference room from different logical positions. For example, the client application in Figure 2 shows two positions: "csri conference (back)" and "csri conference room"; in the first case, the user enters and views from the back of the room, and in the second case, from the front of the room. In this example, there are at least two separate A/V lines into the room. Once inside a room, a DAN can be used to move a remote visitor around the room (the DAN is discussed later).

The A/V PBX supports all of these functions through its client-server architecture: (a) clients running on user or room workstations converse with the A/V PBX server to declare ownership of devices, share devices, control devices, and establish connections between devices; (b) the server maintains a database of the devices and possible interconnections, user characteristics, and state information for devices and users; and (c) the server constantly updates each user's client regarding state changes in all users (e.g., door state, login state, etc.).

The TMS supports connections to "the outside world" through two approaches: (1) connections to other TMS sites and (2) connections to foreign sites. In both cases, a variety of short and long haul A/V communications technologies can be used. For a foreign site, the most popular type of connection is through the outside line feature (shown as "codec-telep" in the user list of Figure 2), in which the user selects an outside line, and is then able to dial an explicit phone number using a conventional H.320 dial-up video codec. For a connection to another TMS site, the A/V PBX servers communicate signalling information over the Internet — this includes the exchange of state information, connection establishment, and accessibility. Intersite A/V links are provided transparently through dedicated analog A/V for nearby sites or dial-up video for distant sites. Once one or more TMS sites have been interconnected (by configuration

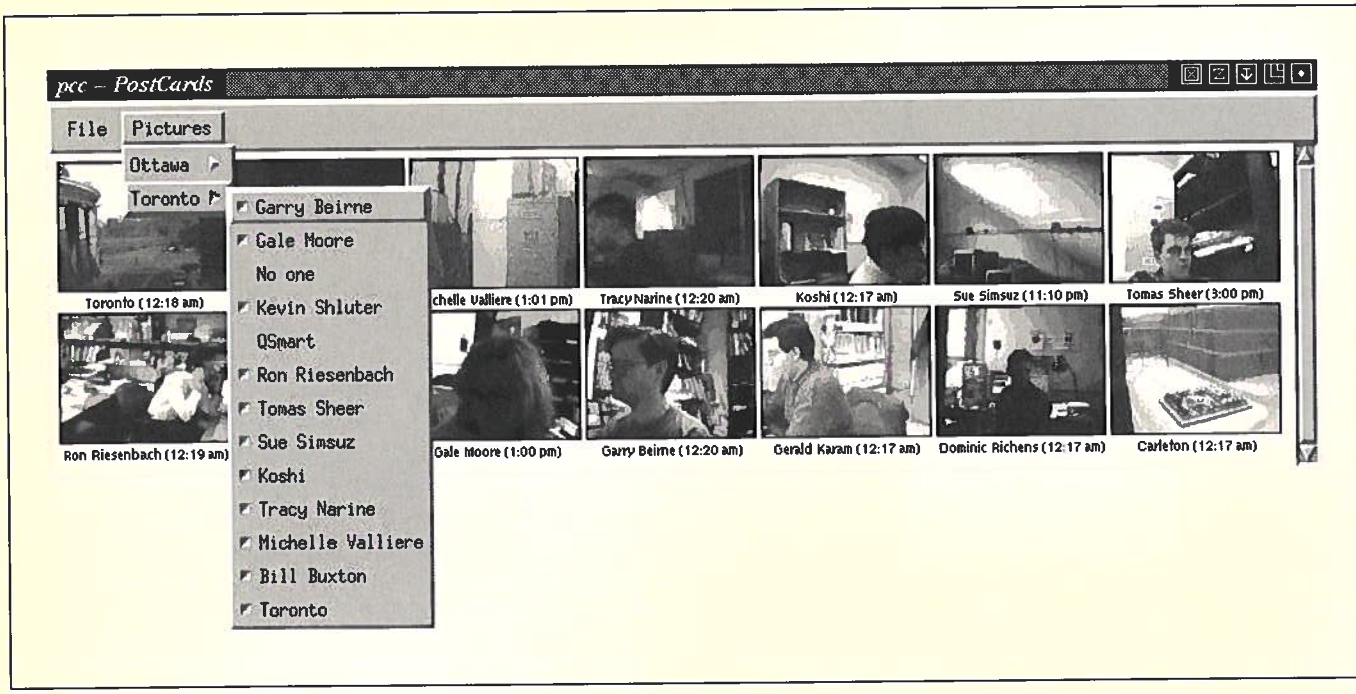

Figure 3. New Postcards system implemented at OTP

files), the user names and rooms become seamlessly integrated. For example, Figure 2 shows users from Ottawa (e.g., "gerald karam") and users/rooms from Toronto (e.g., "garry beirne," "csri conference room"). A user at one site can then contact a user at another site, as easily as a user at the same site. Multisite A/V PBX servers communicate as a fully connected network; while this is not the most effective way for them to function, it was relatively simple to construct.

2.2 Postcards

The original Xerox Portholes system [6] supports background distributed group awareness by providing a palette of continuously updated (about once every five minutes), small "postage stamp" images of persons that are part of the user's community. A community refers to those users that are capable of providing images to a group of one or more sites, that in turn exchange and distribute these images. Technically, this system had a narrow capability for a limited model of use: (1) it was composed of clients connected to a server that acquired image stills through a frame grabber and distributed them to clients on a periodic basis; (2) servers at different sites could exchange images to build a distributed community, but the software was not sufficiently reliable nor efficient to support more than two sites; (3) each user's image was grabbed by a server and interchanged with other servers regardless as to whether or not a user wished to see the image; (4) clients had to poll the server for new images; (5) users could not control who would have access to their image; and (6) anyone with a copy of the Portholes client and Internet access to the server, could connect to the server and view images — even if they were not part of the community.

The new Postcards system implemented at OTP (and shown in Figure 3) revamped both the social model of use and the engineering implementation, in order to provide an efficient, reliable solution that was more sensitive to the realities of a typical workplace. The social model embedded in the behavior of the tool is universal reciprocity augmented by the concept of public and non-public images. Universal reciprocity states that "I am capable of seeing you, if and only if you are capable of seeing me;" thus, any users that want to be able to view other members of the community must make their image available for distribution and if they choose not to do so, then they will be prevented from viewing the images of the community. Images, for which distribution is governed by universal reciprocity, are non-public images. If users (including common images sources, such as the window camera) wish their image to be available to anyone, without the enforcement of universal reciprocity, then they would declare their images to be public; public images can be viewed by anyone, including those outside of the community.

The re-engineering of the application was significant. A Postcards deployment consists of three major components: the clients (one per user), the servers (usually one per "site," or local area of potential clients, but possibly serving only an intermediary store-and-forward function), and the frame grabbers (one or more allocated to one server in a site). Permitting more than one frame grabber per site, allows common framegrabbers to be used, or frame grabbers in personal workstations to be used, or any combination; the OTP deployment uses only a common framegrabber per site. Servers communicate in a spanning tree in order to forward images to clients that are at different sites (since a

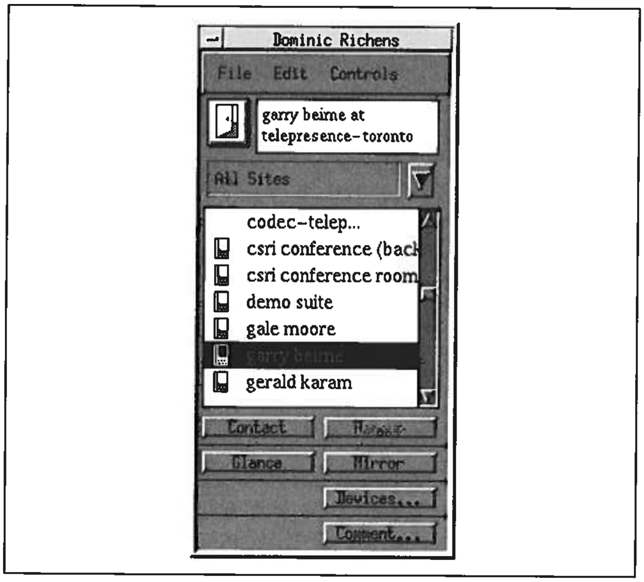

Figure 4. Desk Area Network application

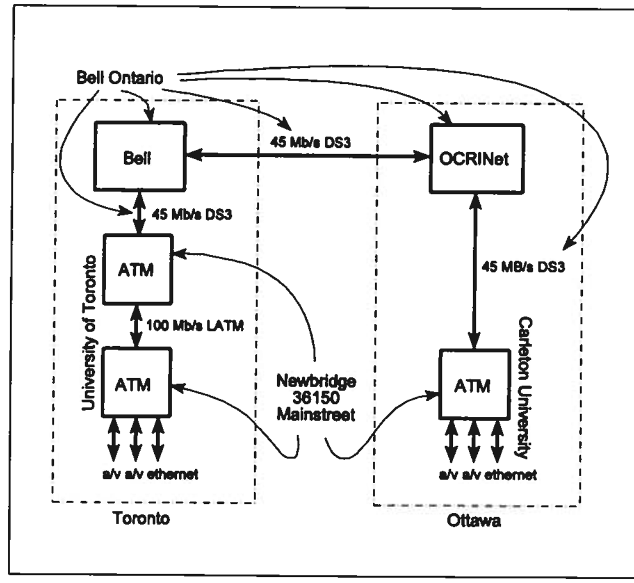

Figure 5. ATM Network

user's community may be spread over a number of sites), thus image transmission is reduced (an image is transmitted only once between two servers). Furthermore, an image is grabbed and shipped only if there is a client that is specifically requesting the image. Each server maintains a database of images available for subscription (to be viewed), images being requested for subscription, and which users are permitted (public/non-public) to subscribe to an image. Every change to the database information about a user at one site is propagated to all other sites through the servers; i.e., it is a distributed database with a copy maintained at each site.

2.3 Desk Area Network (DAN)

The Desk Area Network application, shown in Figure 4, allows local control of many A/V devices in a room; this would suit an office with several video surrogates and devices such as VCRs, or a typical conference room. The current version provides for local switching of the A/V sources and sinks so that the room can be configured easily for any particular social use. For example, an office could be set to video tape both ends of a meeting involving local and remote visitors, or be quickly reconfigured to a simple face-to-face meeting. Previously defined presets are used to specify known or standard configurations of device interconnections. The devices are wired to a small analog A/V switch (such as an 8 x 8 cross bar switch) that is controlled by the DAN application from the user's workstation. TheA/V trunks linking the room to the outside world (typically the A/V PBX in a TMS) are attached to the A/V switch, and form part of its set of connection alternatives.

Further objectives for the DAN include: control of devices such as VCRs, and pan/tilt cameras; automated meeting establishment among several rooms (typically conference rooms); and context sensitive ("reactive") room and device configuration (e.g., placing a paper under a document camera would automatically cause the camera to be routed to the remote visitors). Several aspects of these have been explored as long term research projects by OTP; Karam, McLeod, and Boersma have developed software infrastructure for distributed conference control [8] and Cooperstock et al [5] and have developed reactive room concepts and prototypes.

2.4 Telepresence over ATM

The deployment of the TMS and Postcards to function in an ATM environment was accomplished without significant re-engineering of the software; the DAN was not affected by this transition, but was used extensively for experiments in remote lecturing over the ATM network. This easy conversion was achieved in part, because of the flexibility of the A/V PBX server design and in part due to the choice of ATM equipment. The ATM network, as illustrated in Figure 5, provided three logical connections between the two OTP sites over a dedicated 45 Mb/s DS-3 line: (a) two full-duplex analog A/V circuits, and (b) an Ethernet link. The key elements of the network were the termination equipment — two loaned Newbridge 36150 four-port ATM switches that were equipped with two A/V codecs, an Ethernet interface card and an outgoing circuit card. The A/V codecs provided JPEG compressed near NTSC quality video transmission (scalable according to the peak allocated cell rate) and stereo CD quality audio channels.

The remainder of the network provided the path between the terminal equipment: from the University of Toronto campus, Bell Canada donated a DS-3 channel to the Ottawa area where it connected to OCRINet, an experimental ATM network connecting numerous Ottawa companies, universities and government labs. Carleton University, as one of the OCRINet nodes, provided the DS-3 connection to the Ottawa termination on the Carleton campus. Permanent virtual circuits were created between the two terminal points for the duration of the trial; nominally, the maximum data rates were established as follows: full duplex video circuits -

18 Mb/s; full duplex, stereo audio circuits - 1 Mb/s; and private Ethernet LAN - 1 Mb/s, for a total capacity of 39 Mb/s (2 x 18 + 2 x 1 + 1). This represented most of the available payload capacity of the DS-3 connection.

The TMS was configured to treat the A/V circuits as "very long," dedicated analog links interconnecting the two sites. Therefore, in order to establish A/V connections over the ATM network, only the configuration files of the A/V PBX servers needed modification so that they would search the direct links between sites before resorting to links provided by the dial-up video codecs. Control messages exchanged between servers were routed over the private LAN; therefore, signalling times were only marginally longer than those experienced for local connections. There is nothing unique about this configuration except for the use of the ATM network for A/V transmission and the private LAN. The same configuration was also used for a fully analog link between sites located in two nearby buildings that share a campus-wide LAN.

The Postcards application functions in exactly the same manner as it would in a regular Internet environment, except that the Toronto and Ottawa Postcards servers would exchange images and information over the private LAN. Unfortunately, due to technical delays and the limited time and resources available for the ATM trial, Postcards was never made to operate over the private LAN.

There were several goals for the experimentation of telepresence over ATM; all of these were addressed to some degree before the intercity DS-3 service was discontinued: (1) demonstrate single wire integration of OTP multimedia signals (audio, video, data, and control) over ATM; (2) demonstrate the use of multiple full-duplex channels of A/V in telepresence applications; and (3) demonstrate the integration of current A/V technology (such as analog, and dial-up video codecs), and future technology (A/V transmission over ATM). A number of specific experiments and uses of the telepresence applications were devised in order to illustrate as many of these points as possible. The active term for the trial was about two to three months, some of which was unfortunately over the Christmas holiday season. The intercity DS-3 service continued unexpectedly for an extra two and a half weeks beyond the formal trial and allowed more information and experience to be collected.

3. Experiences

3.1 On Route 66

In the 18 month period from January, 1993 to July, 1994, the TMS sites in Toronto and Ottawa helped to integrate the workpractice of the two groups (about eight people in Toronto and about six people in Ottawa) but only as autonomously running systems; the transparent multisite capability was not in full operation until the end of this period. The same system was also deployed for a field study site connecting users in two cities; it used a single server that controlled a remote switching fabric and dial-up video codec to enable transparent connections between users in the two cities. From July, 1994, the multisite, dial-up video TMS facility was in operation and continued until it was superseded by the ATM transmission service in October, 1994.

This basic TMS system was used for a variety of formal meetings: (1) the engineering team met weekly in a "room-to-room" meeting format; (2) the user interface group met weekly at times, and sporadically at others; (3) infrequent management meetings were held; (4) regular formal seminars were conducted during about a twelve month period (into 1994); (5) formal demonstrations and meetings with interested industrial partners and visitors were held; and (6) even an OTP Halloween party was held over the video link, to explore its use in a purely social situation. It was also used for informal person-to-person connections.

The utility and limitations of 112 Kb/s dial-up video were well understood through our many hours of experience in a variety of situations. The blurry image and annoying time delay (about one to two seconds round-trip) exacerbated larger meetings, and made many social interactions (jokes, interjection, banter) very awkward. Nonetheless, it allowed work to proceed much more smoothly, and inexpensively, than if we had been forced to do without. From the Toronto OTP site to the Ottawa OTP site, about 131 connections were recorded over at 365 day period (November, 1993 to October, 1994); connections in the reverse direction were unfortunately not recorded, however it is reasonable to expect that the number of connections from Ottawa to Toronto was about the same. Nonetheless, this number of connections represents fairly low usage. There are a number of reasons: (1) the time to set-up a call was anything from 35-80 seconds because the dialling required to establish the video link usually required more than twenty seconds; (2) the quality of the image and time delay were workable, but suppressed a real feeling of presence; (3) the time period included December, 1993 and August, 1994, in which most staff were on holidays; and (4) there were staff fluctuations due to some competing projects. During some of this time period (about 30 days in October to November) the ATM capability was up, so there was little demand for the dial-up video codec. Further analysis is presented in the next section where the dial-up service is compared to the ATM service.

In late 1993, we had adapted the original Xerox Portholes to operate in the TMS environment. The work group of about twenty people found it very useful to keep track of when people in the other city were around, and this facilitated explicit meetings using the dial-up video codec. However its limited social model and inability to scale up, necessitated the development of the Postcards system.

3.2 On the Interstate

The basic point-to-point connections over the ATM network were an enormous improvement over the previous dial-up video mechanism. The number of ATM video connections in a 97 day period of operation from Toronto OTP to Ottawa OTP was 294. There was most likely a similar number of connections in the reverse direction, although there is no log data to corroborate this. The higher traffic was due to: (1) a spur of learning, experimentation, and demonstration early on; (2) increased communications to plan ATM experiments; (3) the ease of establishing connections because there was no codec dialling time (set-up time was typically five to fifteen seconds); and (4) the high quality of the image and (virtually) delay-free communications conveyed a much greater sense of presence and made both social and business communications effortless. The traffic level may have been artificially lower because: (1) the number of OTP staff were declining; (2) a three week Christmas holiday season was part of the trial period; and (3) the last 18 days of traffic occurred after OTP had officially ended.

We observed that the ATM service had a much higher proportion of brief contacts because they were very easy to do, and were as much effort and more satisfying than placing a conventional speed-dial telephone call. Also, the relative call traffic volume showed that the ATM service was used more intensively, even though the number of people involved had declined over the peak number of users of the dial-up video service. Some connections were much longer as users simply let the length of the meetings fit the workpractice. This differs from the dial-up video links that were not as convenient to use, did not provide the same quality of interaction, and calls were typically focused on a specific purpose. Cost was not really an issue for the dial-up video link, even though regular phone rates (the cost of two long distance phone calls) were incurred.

A number of multipoint meetings were held (at least 22 of 294 calls) originating from the Toronto OTP site; in some cases they involved users connected by analog and ATM video intersite links, and in other cases also involved dial-up video intersite links. We found that in the first case, there was no difference between an all-local user multisite meeting, and a meeting involving ATM A/V link users (at times we had two ATM A/V link participants since two A/V links were available). As there was no significant time delay in the ATM A/V transmission, interactions proceeded smoothly. Multipoint meetings involving dial-up video participants, proved to be a little more difficult, as the majority of the participants had no time delay, and the dial-up video participants had poorer picture quality and suffered the usual annoying time delay.

In a virtual open office model, a point-to-point or multipoint connection is left on for an extended period of time (e.g., for half or the entire business day), in order allow distant collaborators the opportunity to share a virtual open office. Two offices (one in Ottawa and one in Toronto) used by three of our collaborating engineers were connected in this way for about 30 days. They reported 29 positive interactions and six distractions resulting from the open office. This data is being analyzed by our social science researchers, however, overall, the engineers were impressed by the convenience and were satisfied with the experience. The engineers reported that since the virtual open office has been discontinued, their productivity has been impaired as the typical "quick question, quick answer" that was afforded by the shared office, is now handled by many exchanges of short email messages.

The Telepresence Tunnel was a background full duplex A/V connection between a public space at the Ottawa OTP site and a public space at the Toronto OTP site. Students, faculty, and the general public were able to walk up to their end of the Telepresence Tunnel and readily communicate with peers at the other end. Its purpose was to understand how people in physically widely separated buildings, but with some common background would socialize using audio and video, once given the opportunity. It would also illustrate how video conferencing novices would react to interactive video, and perhaps to future publicly available video communication services.

Typically, the Telepresence Tunnel was in operation 24 hours a day for eight weeks (subject to some interruptions), although for about three weeks there were final exams at both sites, followed by the Christmas break. Over a 30 day period, images were sampled at five minute intervals at both ends of the Telepresence Tunnel and formed into MPEG movies. Analysis of this data from the Carleton site was done by examining individual frames for activity; where the same person (or group of people) remained at the tunnel for more than two frames (at least five minutes), then it was assumed a conversation was taking place with someone at the other end. Where there was only a single sample of a person (under five minutes), no strong conclusion could be made regarding that person's activities. Based on the assumption that a real conversation was at least five minutes, there were 37 recorded conversations. Of these, more than one-quarter were almost 30 minutes, which suggest extended conversations between people who were familiar with one another, in some way. The conversations exceeding 30 minutes all appeared to be between people who knew each other; and were often quite long: 95-100 minutes, 45-50 minutes, and 40-45 minutes.

We ran three evaluations of a remote teaching experience in a three week period. In all cases, the lecturer was at the Toronto OTP site and the audience was at the Ottawa OTP site in the conference room. The DAN was present at both ends to easily reconfigure the rooms for different phases of the lecture. The goal was to assess the impact of image size, image quality, sense of presence, and the effect of having two full duplex A/V channels rather than the traditional one. In the most successful configuration, the two ATM A/V links were used. A document camera was used by the lecturer for paper overheads and was transmitted over one ATM A/V link; the lecturer's hand gestures could be seen in the document camera. In the lecture room, the document camera was displayed on multiple monitors around the room, and the lecturer's image, transmitted over the second ATM A/V link, was displayed using a projection television. In the lecturer's room, a fixed position image of the audience was sent over one ATM A/V link and displayed on a projection television, and a roaming image (guided by an assistant in the lecture room) was displayed on a nearby monitor. During the formal presentation, the assistant used the remote control on the roaming camera to pan and zoom around the room as a lecturer might do naturally while giving a talk. During question period, an assistant in the lecture room used a remotely controlled camera to zoom-in on students asking questions; the discussion flowed easily because there was no time delay and the image quality was very high.

This configuration gave the most satisfying result to both the audience and the lecturer because: (1) the audience saw the full size image of the lecturer at all times, and saw good interaction with the presentation devices; (2) the lecturer saw a large image of the audience and thus always had a feeling of their presence in the lecturing room; and (3) the lecturer got a more intimate view of the audience members during the talk, due to the constantly roaming camera, and thus could get a better feeling for how the material was being received.

4. Integrating Intelligent Agents

Intelligent agents provide important opportunities for the growth and sophistication of telepresence applications, particularly in the future as communications bandwidth becomes plentiful and less expensive, and when audio/video end-user equipment is common place. A few suggested roles for intelligent agents are described below, in the context of the telepresence work, in the hopes that it will open up further areas of exploration for the Artificial Intelligence community.

A common use of intelligent agents, especially for end-user tasks is the personal assistant that aims to manage day-to-day activities [17] such as filtering email and scheduling meetings. In telepresence applications, such agents have additional information sources available to help with user functions, specifically, audio and video data. For example, the meeting scheduling agent can be adapted to identify dynamic meeting opportunities by using the Postcards application. Postcards is used as a background motion detector to help recognize when a small group of colleagues are alone in their offices, and available for a meeting. Thus, if a user wanted to have a quick meeting, the user's agent could work in the background to assemble the meeting and notify the user (or automatically engage the meeting) as soon as the participants were available. The inputs to the agent could be conventional data such as public schedules, but it also includes the historical video information provided by Postcards. Beyond the simple video images, audio information and telephone information could be tracked by agents to ensure that a potential participant is not on the phone, or speaking with someone in the office (but off camera). While it may seem intrusive for one person to have such close access to another, it may be perfectly acceptable for an impartial third party (like a agent) to have such information, solely for their job function. This is not unlike the analogy of the "Butler" or "Maid" who would see much private information, but use this knowledge only in the context of their work.

A second example of a personal assistant is a meeting organizer. Jones and Edmonds [17] describe a set of agents to help in a collaborative meeting application. Beyond the conference control agents that they describe (much of which is support by the emerging ITU Generic Conference Control standard discussed in [17]), there is a great need to organize a user's documents and applications that are shared during a meeting. For example, each user will need documents available, and tools to share those documents; documents may become updated during the meeting, and the "state" of the meeting, is the people, connections, media, documents, and tools used during the meeting. An agent should handle the assembly of the documents needed, tools for sharing, and collection of documents after the meeting is closed, including a record of what is needed for the next time the meeting is called (if appropriate), including persistent state information from the meeting. The personal agent must have knowledge of the user's contribution to the meeting, the available collaborative facilities, and the other collaborators at the meeting. A true human assistant would make sure that all of the meeting materials were available; ensure that the proper audio/visual and communications aids were on hand, and anticipate the impact on the user's contribution based on the identity and function of the other collaborators.

Related to the meeting organizer agent is the room adapter agent. When a user enters a room equipped for collaborative meetings (e.g., a private office or a conference room), the room should adapt the A/V configuration to the user. The reactive room is a first attempt at this type of dynamic room reconfiguration [16]. In the reactive room, the audio-video devices are configured based on certain user actions (such as placing a document under a document camera), or a user identification from an "active badge." This can be extended to an agent that learns and understands a user's actions in different job functions. To become truly sophisticated, such agents would have to learn and adapt such as the calendar apprentice agent described by Mitchell et al. [17]. For example, if a user always had meetings with a remote visitor and local visitors at a different table in the office space, then camera positions could be switched automatically when such a meeting was started (knowing that there were multiple visitors could have been picked up from a calendar or agenda, or even by using live video processing). For a conference room, a user's personal preferences for camera angles, A/V devices, visitor locations, and meeting functions, could be handled by the agent.

Apart from the personal assistants just described, there are network agents that would be used to manage the rich set of telecommunications resources that are needed for the cost-effective utilization of the high data bandwidth used in telepresence over ATM. Weihmayer and Velthuijsen [17] outline how agents are used for network diagnosis, configuration and management. Agents would be especially helpful to manage the video bandwidth needed for telepresence applications; different collaborative functions demand widely varying qualities and rates of video imagery — from Postcards to virtual open offices, to casual desktop meetings, to high end video conferencing. A fixed pool of bandwidth for a company's internal and external communications must be managed on a dynamic, anticipatory basis to ensure that the right image quality and data rate is available when needed. This will require agents both in end-user equipment, at main switches, and at customer- premises/ service provisioner boundaries.

5. Related Work

Media spaces, and in particular, desktop video conferencing systems, are a growing area of application for business technology [9]. From the research domain are a variety of systems that are exemplified by the following implementations: the Media Space project at Xerox PARC [3] and related work at Rank Xerox EuroPARC, Bellcore's Touring Machine [2] and the Internet MBone [10]. The Xerox work defined some of the early notions of a media space that are part of the TMS model, and some elements of the software were derived from this effort. The Touring Machine application combined analog and digital A/V transmission to perform some of the media space functions of TMS; however, its model of operation was different, it never employed ATM technology for transmission, and nor was it deployed as an integrated media space for arm's length organizations. The Internet community advocates an alternate approach to video conferencing that supports large scale broadcasts of fully digital audio/video over the multicast backbone or "MBone." Philosophically, the MBone uses a model in which A/V is widely distributed to a large number of participants over the Internet; this is vastly different that the media space model suggested by both TMS and the Touring Machine.

In standards development is the Multipoint Communications Service (MCS) [1], that will support telepresence applications, to the extent that it provides a useful model of multipoint applications (such as Postcards, or the TMS servers); the MCS standard leaves the transmission of motion video and real-time audio to existing transmission standards (e.g., H.261 for motion video). For example, we have developed and demonstrated a Network Service Interface (NSI) for telepresence applications using an implementation of the original MCS protocol (on which the standard is based).

In the commercial domain [9], there are both all-digital desktop video products and analog/digital products. On the digital side are application sharing/video conferencing systems like Northern Telecom's VISIT, AT&T's Vistium, and PictureTel's LIVE PCS that use ISDN for digital A/V transmission, and document sharing/file transfer. These tools generally use video-in-a-window on the workstation and have small blurry images when transmitting over ISDN lines. The cost is usually high because of the use of ISDN lines per station, and the need for specialized digital A/V hardware. Systems that use analog for local site transmission and shared A/V codecs for long distance transmission (similar to the TMS approach) include products from Target Technologies and Datapoint. The focus of these implementations is also basic desk-to-desk communications. Finally, an emerging analog/digital product is Corel Video! from Corel Corp., that will have a more flexible connection model and better suited to support a media space like TMS than other similar products; this is not surprising as it represents a commercialization of the TMS environment.

6. Remarks

Before the ATM experience, OTP had demonstrated the utility and characteristics of the synchronous A/V media space in supporting workplace collaboration. The advent of telepresence applications over ATM showed us the future could be much better than we had imagined and yet left us realizing that there was so much more to be explored. The experience gained in combining current and future communications technologies, validated much of our design approach, as well as showed the degree to which they could interact today. It also suggested some possible models for deployment in the near term as ATM transmission services become available in metropolitan and long haul networks. Furthermore, our experience opens up more opportunities for greater sophistication and seamless behaviour through the use of intelligent agents for personal assistance and network resource management.

The small stretches of the ATM Interstate on which we drove just made us beg for more — it was fast, smooth, and let us accomplish our travels much more conveniently than the old road. Route 66 served us well, and continues to serve us now since the ATM Interstate is still in its infancy and exists only in small segments. While the completion of the ATM Interstate is, to some degree, inevitable, it will probably take many years due to some of the same challenges that plagued the US Interstate Highway System. To justify the roadway, we will need services and traffic, and our experience suggests that these are not easily defined nor understood without considerable experimentation and study with the real user community.

Acknowledgments

This work represents many hours of hard work and commitment by the staff and researchers of the Ontario Telepresence Project, in particular Bill Buxton, Gale Moore, Ron Riesenbach, Dick Dillon, Ross Jacobs, Tom Milligan, Tracy Narine, Chris Passier, Dominic Richens, Jingping Xie, Garry Beirne, Maciek Kozlowski, and Tabatha Holtz. The ATM phase of the project was a success only because of the gracious help from our industrial collaborators: Bell Canada seconded Michelle Valliere to us for a year, provided funding and the intercity DS-3 service; Newbridge Networks provided the ATM switch gear loan, through the kind support of Jim Mackie, Richard Zwiep, and others who made this a success. Also, OCRINet, Carleton University, and University of Toronto provided much assistance in this phase. Finally, we are indebted to the Province of Ontario and our industrial partners and collaborators for funding and in-kind support throughout the project. Many thanks to Sue Abu-Hakima who provided many editorial suggestions.

References

[1] ITU Study Group 8. T.122: Multipoint Communications Service for Audiographics and Audiovisual Conferencing Service Definition. ITU Telecommunications Standardization Sector, October 1993.

[2] Bellcore Information Networking Research Laboratory, Touring Machine System," Comm. of the ACM, Vol. 36, No. 1, Jan. 1993, pp. 68-77.

[3] S. Bly, S. Harrison, and S. Irwin, "Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment," Comm. of the ACM, Vol. 36, No. 1, Jan. 1993, pp. 28-47.

[4] W. Buxton, "Telepresence: Integrating Shared Task and Person Spaces," Proceedings of Graphics Interface '92, Vancouver, B.C., May 1992.

[5] J.R. Cooperstock, K. Tanikoshi, G. Beirne, T. Narine, and W. Buxton, "Evolution of a Reactive Environment," Proc. of CHI'95: Conf. on Human Factors in Comp. Syst., Denver, May 1995.

[6] P. Dourish and S. Bly, "Portholes: Supporting Awareness in a Distributed Group," Proc. of ACM Conf. on Human Factors in Computer Systems CHI'92, Monterey, CA, May 1992.

[7] R.M. Jones and E.A. Edmonds, "A Framework for Negotiation," in CSCW and Artificial Intelligence, J. Connelly and E.A. Edmonds, eds., London:Springer-Verlag, 1994, pp. 13-22.

[8] G.M. Karam, B. McLeod, and G. Boersma, "A Network Services Interface for Telepresence Applications," IEEE Intl. Symp. on Autonomous Decentralized Systems, April 1995.

[9] G. Koprowski, "Desktop Videoconferencing: The Profit Picture," Reseller Management, November 1994, pp. 106-112.

[10] M. Macedonia and D. Brutzman, "MBone Provides Audio and Video Across the Internet," IEEE Computer, Vol. 27, No. 4, April 1994.

[11] P. Maes, "Agents that Reduce Work and Information Overload," Comm. of the ACM, Vol. 37, No. 7, July 1994, pp. 31-40.

[12] T. Mitchell, R. Caruana, D. Freitag, J. McDermott, and D. Zabowski, "Experience with a Learning Personal Assistant," Comm. of the ACM, Vol. 37, No. 7, July 1994, pp. 81-91.

[13] R. Riesenbach, "The Ontario Telepresence Project," CHI'94 Conf. Companion, ACM Conf. on Human Factors in Computing Systems, pp. 173-174, Boston, MA, Apr. 1994.

[14] S. Schooler and S. Casner, "An Architecture for Multimedia Connection Management," IEEE 4th Comsoc International Workshop on Multimedia Communications, Monterey, CA, April 1992, pp. 271-274.

[15] Q. Scott and S.C. Kelly, Route 66 — The Highway and its People, Norman OK: University of Oklahoma Press, 1988.

[16] J. Steinbeck, The Grapes of Wrath, New York, NY: Penguin Books, 1939.

[17] R. Weihmayer and H. Velthuijsen, "Dsitributed AI and Cooperative Systems for Telecommunications," in Worldwide Intelligent Systems, J. Liebowitz and D.S. Prerau, Eds., IOS Press, 1995, pp. 227-260.

***Gerald M. Karam** is a Member of Technical Staff at AT&T Bell Laboratories in Murray Hill, NJ, and an Adjunct Professor in the Dept. of Systems and Computer Engineering at Carleton University, Ottawa, Canada. Previously he was an Associate Professor at Carleton University for 8 and 1/2 years. His current interests include software design and analysis, concurrent systems, real-time systems, and telecommunications software. An Ottawa native, he received his Ph.D. in Electrical Engineering from Carleton University in 1987, his M.Eng. (EE) from Carleton in 1984, and his B.A.Sc. (EE) from the University of Ottawa in 1982. From 1993-1994 he was also the Director of Engineering for the Ontario Telepresence Project, within TRIO (Telecommunications Research Institute of Ontario) and ITRC (Information Technology Research Centre); TRIO and ITRC are Province of Ontario Centres of Excellence. In this position he provided technical direction to the engineering aspects of the project and managed a group of five full-time research engineers and several graduate students. Dr. Karam is a member of IEEE and ACM.*

Artificial Intelligence and Distributed Information Networks in the United States

Jon R. Wright

Introduction

Les comptes rendus des applications d'intelligence artificielle présentent souvent les choses vues de l'intérieur: on parle en termes de la recherche actuelle en IA et on utilise une terminologie issue de l'expérience de recherche. Bien que cela encourage la communication à l'intérieur de la communauté de recherche, ça peut être une barrière pour ceux qui travaillent sur les applications, qui pourraient ne pas reconnaître le lien entre des problèmes commerciaux spécifiques et les nouveaux développements en IA.

Dans le même ordre d'idée, la plupart des champs d'application ont leur propre technologie, et le jargon qui s'y rattache peut être difficile à pénétrer pour le spécialiste en IA. La révolution technologique actuelle associée à l'autoroute de l'information en est un bon exemple. En fait, la poussée de technologie nouvelle se fait à un tel rythme que les spécialistes en télécommunications ont de la difficulté à suivre. Pourtant il existe clairement des possibilités où l'on pourrait appliquer les techniques et les outils d'IA existants, tout autant que des problèmes de recherche intéressants.

Pour comprendre quelles sont ces possibilités, nous allons présenter les choses vu de l'extérieur, en essayant d'expliquer pourquoi les développements en IA sont importants pour l'industrie des télécommunications. Il y a trois grandes tendances à considérer: (1) la complexité grandissante des réseaux amenée par la nouvelle technologie, (2) la réduction des coûts due aux forces du marché, et (3) l'incertitude du marché face aux nouveaux services.

Introduction

Reviews of applications of artificial intelligence frequently provide an inside out perspective, organizing the discussion in terms of ongoing research in AI and using terminology derived from the research experience. While this fosters communication within the research community, it can be a barrier to those in the applications world, who may not recognize the connection between specific business problems and new developments in AI.

By the same token, most application areas have their own supporting technology and the accompanying jargon can be difficult for the AI specialist to penetrate. The ongoing technological revolution associated with the well-publicized information superhighway is a good example. In fact, the rush of new technology is happening at such at rate that specialists in telecommunications have difficulty keeping up. Yet there are clearly opportunities for applying existing AI tools and techniques, as well as some interesting research problems.

To understand where the opportunities lie, we will organize our discussion from the outside-in, trying to explain why developments in AI are important for the telecommunications industry. There are three broad trends to consider: (1) increasing complexity of the networks brought about by new technology, (2) cost reduction due to marketplace forces, and (3) market uncertainty regarding new services.

The State of Today's Public Carrier Networks

Public carrier networks in the United States are today substantially more complex than any corporate local area network and becoming more complex all the time. To give an example, today's basic cable and pay-per-view movie services are supported by an interoffice network of fiber optic cable that uses a technique called direct supertrunking. Direct supertrunking uses broadcast video signals. It works well and is economical.

However, the demand for services, especially interactive services, will soon exceed the capacity of networks that depend solely on direct supertrunking. Consequently, there will be a need for video signals to be switched rather than broadcast. Hence the anticipated introduction of a new switching technology called asynchronous transfer mode (ATM) that combines switched, high capacity throughput with the flexibility to transport many different kinds of data. For reasons of economy, it seems inevitable that direct supertrunking solutions and ATM networks will coexist for some time.

This situation, in which newer networking technology is overlaid on an older network, is characteristic of most public carriers in the United States. The mixing of old and new technology makes network management ever more complex. Each new technology requires an investment in infrastructure which requires development and enhancement of operations support systems, training of personnel, and so on.

At the same time, the new competitive environment puts an emphasis on cost reduction to be accomplished by reengineering. In telecommunications, reengineering has most often meant centralization of some work functions coupled with automation of others (Bernier, 1995). As in the case of corporate networks, adoption of the client-server model is playing an important role in the reengineering process (Titch, 1995). Client-server architectures are important because they provide the means by which work functions can be automated and distributed into peripheral systems.

The final element is the uncertainty with respect to the new markets. One reason for mentioning ATM networks is that they are widely thought to be the vehicle for providing multimedia services (Roy, Kuthyar, and Katkar, 1994). On the other hand, it is not clear what kinds of multimedia services the market prefers (Burt and Lund, 1994). Consequently, there is much experimentation and a pressing need to maintain flexibility as market demands change and new technology becomes available.

The Role of Expert Systems

Expert systems have become a well accepted technology in the telecommunications world, projected by some analysts to be a key enabler of future success (Bernier, 1995). Part of the reason for their acceptance is that they have proved to be an effective method for automating work functions. In effect, they expand the circle of functions that can be effectively automated.

For our purposes, an expert system is a software application that uses rules as the main programming technique. The telecommunications industry in the United States has been using this technology effectively for some time. The first deployed system was AT&T's ACE system (Vesonder et al., 1983), an expert analyzer of cable faults, which is still working in many sites today.

It's hard to find a telecommunications company in the United States that has not deployed some kind of expert system. Each year, the conference on Innovative Applications of Artificial Intelligence, held simultaneously with AAAI's national conference, contains reports of successfully deployed expert systems in telecommunications. Focusing on published accounts clearly underestimates the extent to which expert systems have had an impact on the business because not all development teams are inclined to publish their results. Expert systems have been applied to areas where reegineering has been most vigorously pursued, primarily network management and maintenance.

Expert systems also have a simplifying effect on the work environment. A good example is the StarREP system which has been deployed in several major U. S. companies and was discussed in an invited address at IJCAI-93 (Benda, 1993). StarREP uses production rules to encode methods and procedures established by telecommunications companies for interacting with customers. Most importantly, however, StarREP relieves employees from having to know how to access and locate information in a large collection of legacy systems. The legacy systems — database systems, testing systems, etc. — contain information or are capable of accessing information relevant to a customer's telecommunications services. StarREP is successful in part because it takes the edge off a very complex work environment.

Software Agents

If expert systems are an accepted technology today, software agents are perceived by many industry analysts to be crucial to future success (Williamson, 1995). Agents are attractive in the telecommunications world for several reasons. First, they fit easily into the increasingly ubiquitious client-server architectures. Second, they place low demands on computational resources, or, at least, take advantage of those resources that are available. Third, their modularity suggests that they can be deployed incrementally, offering the possibility of accruing benefits early on and adjusting as experience is gained.

There are several competing notions of what software agents are. General Magic has developed a language for communicating agents called telescript that has a distinctly object-oriented flavor (White, 1994). In this approach, agents are special kinds of objects with associated methods and data that can be transmitted over a network and executed remotely. The technique is called remote programming. The metaphor promoted by General Magic is the electronic marketplace, a sort of networking version of the shopping channel. There are trials planned with a variety of U. S. companies.

The notion of agent as software object contrasts with the view of agents in the AI community in which agents can both reason and perform actions (for example, Hayes-Roth, 1995). However, these are not necessarily incompatible views, and successful deployment of agent-as-object systems could pave the way for reasoning agents at some later time.

AI in Network and Service Provisioning

Network and service provisioning are very crucial parts of the telecommunications business, and AI has already had an important impact in this area. Provisioning is the business process that makes products available to customers. The term applies both to equipment and to services. Most telecommunications manufacturers work on a build-to-order or assemble-to-order basis. Complex equipment such as switching machines are generally customized for a specific application and not bought off the shelf. Customizing this equipment is essentially a configuration problem.

Since the deployment of R1/XCON by the Digital Equipment Corporation in the early 1980's (Mcdermott, 1982), configuration problems have been recognized as candidates for the application of AI methods. In fact, AT&T switching machines have been configured since the mid 1980's with a configurator application called DOPS (Digital Ordering and Planning System) that uses an essentially rule-based approach. However, we are also starting to see new approaches to the configuration problem. For example, the PROSE configurator (Wright, et al., 1993) uses a description logic called CLASSIC to represent product knowledge. In addition, the Expert Solutions Platform, reported at IAAI-94 and developed by AT&T's Global Business Communications Systems unit, configures PBX equipment using a constraint satisfaction engine.

Services add another layer of complexity on the provisioning problem. One of the important innovations in the United States over the past ten years is something called the Advanced Intelligent Network or AIN (Russo, et al., 1991). Essentially, AIN introduces centralized databases and peripheral processing elements into an existing network without affecting the way network connections are established (Crabill and Kukla, 1994). This allows processing for services to be separated from the switching systems themselves. AIN opens up entire new worlds for service processing.

In the United States, services are implemented in the AIN using application-specific languages called decision graphs or call scripts (Morgan, et al., 1991). However, the world of service provisioning is growing ever more complex, and new techniques extending what we do today will be required in the future.

In effect, a service is nothing more than a collection of database transactions and commands called switch translations. The trick is to organize them in the correct way. Some complex services, such as inbound wide area telephone service (IN-WATS), require a distribution of actions over many network elements and peripheral processors. In the future, there clearly will be increased emphasis on providing solutions rather than simply individual services or pieces of equipment. It seems likely that a business solutions approach will require configuration of equipment and services together. This is a much tougher and more challenging problem than what we see today.

New Services

Finally, AI has an important role to play in the development of new services on the information superhighway. Perhaps the most comprehensive view of what new services will be like in the future is provided by AT&T's Vision-2001, a concept developed in late 1991. The vision includes providing affordable communication between people and machines in an effortless, easy to use way, using any medium that is natural and convenient. Implementing such a vision requires advances in any number of disciplines. In this context, we will mention just one — speech technology.

While there is still much uncertainly regarding the market for new services, we can predict with certainty that speech processing technology will become more prominent. Speech is still thought by many to be the preferred mode of human-machine communication.

Speech processing technology is important in many applications today, for example, in voice messaging systems. However, recent advances in speech recognition, text-to-speech synthesis, speaker verification, and language identification and translation indicate that the number of practical applications of these technologies are increasing.

A recent issue of the AT&T Technical Journal (Tschirgi, 1995) outlines the current state of the art in speech processing and discusses the implications for new services. One important aspect of speech processing technology is that, while it has set rather high goals, there have been important applications fall out along the way. For example, currently a feature known as automatic call-type processing is being deployed in the United States (Juang, Perdue, and Thompson, 1995). This feature uses speech recognition to indicate whether a caller wants to make a collect, calling-card, third-number, person-to-person, or operator-assisted call.

Another example of a important feature on the way is wireless voice digit dialing. This is particularly useful in situations where the caller has both hands and eyes busy, as in using a cellular telephone while driving an automobile.

Conclusion

Along with the increasing complexity, competitive pressures, and marketplace uncertainty, there are also significant opportunities on the emerging information superhighway. Developments in AI can and should play an important role in this new market. AI has a long history of working on complex problems and finding ways to simplify them. As we have discussed, effective ways of dealing with complexity are sorely needed.

At the same time, we know from experience that solutions developed in the research lab do not automatically translate into marketplace successes. We should maintain a good grasp on what can be accomplished today and choose our applications carefully. This is not incompatible with keeping a forward looking vision or working on really hard problems, only that we keep a healthy respect for both research and applied work.

Automatic call processing and wireless voice digit dialing provide good examples of successful technology transfer. Both applications involve innovations in speech technology, but their success is partly due to the thoughtful analysis of the application domain and careful matching of the technology to the needs of the customer. Such successes lay the groundwork for future work and support in speech technology. This is a lesson that we should bear in mind as AI attempts to find its place in the information superhighway.

References

Benda, M. (1993). Applied AI adventures worth writing home about. In Proceedings of the 13th International Joint Conference on Artificial Intelligence, San Mateo, Ca: Morgan Kaufmann, 747-750.

Bernier, P. (1995). The center and the edges. In Stephen Titch (Ed.), The Network Management Puzzle, Telephony (supplement), March 6, 1995, 12-17.

Burt, D. and Lund, B. (1994) Updated OSSs will make the multimedia network fly. Telephony. 226(16), 33-36.

Crabill, E. F. and Kukla, J. A. (1994). Service processing systems for AT&T's intelligent network, AT&T Technical Journal, 73(6), 39-47.

Juang, B. H., Perdue, R. J., and Thomson, D. L. (1995). Deployable automatic speech recognition systems: advances and challenges. AT&T Technical Journal, 74(2), 45-56.

Hayes-Roth, B. (1995). An architecture for adaptive, intelligent systems. Artificial Intelligence, 72(1-2), 329-365.

McDermott, J. (1982). R1: a rule-based configurer of computer systems. Artificial Intelligence, 19(1), 39-88.

Morgan, M. J., Cosky, M. J., Gruenenfelder, T. M., Holmes, T. C., and Raack, G. A. (1991). Service creation technologies for the intelligent network. AT&T Technical Journal, 70(3-4), 58-71.

Russo, E. G., Abdel-Moneim, M. T., Sand, L. L., Shaw, L. J., Taska, P. D., and Wojcik, R. J. (1991). Intelligent network platforms in the U. S. AT&T Technical Journal, 70(3-4), 26-43.

Titch, S. (1995). The first pieces. Telephony. 235(16), 5-12.

Tsichirgi, J. E. (Ed.) (1995). Voice and Audio Processing, Special Issue of the AT&T Technical Journal, 74(2).

Vesonder, G. T., Stolfo, S. J., Zielinski, J. E., Miller, F. C., and Copp, D. H. (1983). ACE: An expert system for telephone cable maintenance. In Proceedings of the 6th International Joint Conference on Artificial Intelligence, San Mateo, Ca: Morgan Kaufmann, 116-120.

Wright, J. R., Weixelbaum, E. S., Vesonder, G. T., Brown, K. E., Palmer, S. R., Berman, J. I., and Moore, H. H. (1993). A knowledge based configurator that supports sales, engineering, and manufacturing at AT&T Network Systems, AI Magazine, 14(3), 69-80.

White, J. E. (1994). Telescript technology: the foundation for the electronic marketplace. General Magic White Paper, Mountain View, Ca.: General Magic, Inc.

Williamson, J. The picture emerges. (1995) Telephony. 235(16), 22-27.

Jon R. Wright *is a Distinguished Member of Technical Staff in the Software Technology Center at AT&T Laboratories, Murray Hill, NJ, where he has been involved in applications of AI to telecommunications for nearly 15 years. Some of the successfully deployed and currently working systems include applications in network management, network testing and maintenance, and network provisioning. He received a Ph.D. in Cognitive Psychology from Rice University in 1978.*

An Agent-based Multimedia Service Environment

Michael Weiss, Tom Gray, Aurora Diaz

Introduction

Cet article décrit nos expériences dans l'utilisation de la technologie des agents pour la conception d'applications multi-média distribuées. Au coeur de la conception des systèmes de télécommunications de demain se trouvent la création et l'ajustement sur mesure des services. Ces systèmes sont fondés sur le paradigme que les systèmes de communication distincts dans les entreprises aujourd'hui (pour la voix, les données, la sauvegarde de l'information, et les rencontres face- à-face) convergeront éventuellement en un système de communication unique. Il formera la base informationnelle de l'entreprise et subordonnera tous les processus de communication à l'intérieur de l'organisation, formant ainsi une part intrinsèque de l'ensemble de ses processus d'affaires.

Ces systèmes seront inévitablement multi-vendeurs et hétérogènes. Les systèmes sous ce paradigme convergent, que l'on désigne généralement comme "middleware", ne peuvent pas être simplement construits sur le modèle transactionnel classique, qui est à la base des standards en développement comme CORBA (Common Object Request Broker Architecture). L'objectif de MANA (Multi-Agent Architecture for Networking Applications) est de supporter la transparence de la technologie, du lieu et des services pour les applications multi-média distribuées, tout en rencontrant les contraintes de fiabilité et de temps réel. Un des buts de l'architecture est de garder locales les interactions de base. Un autre but est d'améliorer l'utilisation des ressources (autant les ressources humaines que l'équipement) pendant l'opération du système.

Introduction

This paper describes our experiences in using agent-based technology for the design of distributed multimedia applications. At the heart of designing tomorrow's telecommunication systems is the creation and customization of services. These systems are based on the paradigm that the distinct communication systems in the enterprise today (for voice, data, information storage, and face-to-face meetings) will eventually converge into a single communication utility. It will form the information base of the enterprise and will subsume all the communication processes within the organization, thus forming an intrinsic part of its overall business processes.

These systems will inevitably be both multi-vendor and heterogeneous. No single vendor will be able to supply all component applications and applications must interwork seamlessly with legacy and other vendors' applications. This means that these new systems must be able to function in an environment that is not only vaguely defined but also constantly evolving. They also have to operate in the face of unexpected contingencies or failures in hardware, software and even human resources. They must be sufficiently adaptive to detect failures and function predictably while continually reconfiguring their resources to meet changing end user demands.

Systems under this convergent paradigm, which are generally referred to as middleware, cannot be constructed solely on top of the classical transactional model, which is the basis for evolving standards such as CORBA (Common Object Request Broker Architecture). The objective of the Multi-Agent Architecture for Networking Applications (MANA) is to support technology, location, and service transparency for distributed multimedia applications, while meeting reliability and real-time constraints.

One goal of the architecture is to keep feature interactions local. Feature interactions occur in collaborative systems because of different user expectations of services and the need to support service customization. The large number of features in modern service environments leads to an explosion of interactions between them (Velthuijsen, 1993). MANA contains the number of feature interactions by explicitly representing a context for each agent-to-agent relationship. Contexts effectively limit the number of services between which harmful interactions may occur as a new service is introduced or an existing one is changed. Our approach amounts to minimizing feature interactions by the way the agent network is set up, that is, by design. Most traditional service code must be tested against a global interaction table of all hypothetical interactions. We use a declarative approach to decouple service definition from service provisioning.

The second goal is to improve the use of resources (both human and equipment) during system operation. Whenever the availability of a resource decreases or increases, the information used by the associated resource allocation mechanism is updated. Over time, the application generates knowledge for the more efficient use of its resources. MANA is not so much concerned with the specific realization of the resource allocation mechanism, but rather with the management of uncertainty within the organization that arises from the dynamically changing behavior of its resources. When it is continuously informed about the health of its resources, the application can adapt its allocation strategies and optimize the exploitation of its resources.

Architecture

MANA is based on the notion of an organizational network of agents that cooperate to provide services to the external and internal customers of an enterprise. It provides a standard structure that allows many applications and users to share and exploit the resources of an enterprise. MANA is not specialized to the computing or telephony domain. It tries to marry these domains under the broader domain of business process enabling.

Enterprise Modeling

In MANA, the enterprise is broken down into business units that perform functionally related enterprise activities such as procurement or marketing. Business units are mapped to groups of agents cooperating toward the common goal of that unit, which is a subgoal of the overall organizational goal. However, MANA does not dictate the organizational form of the enterprise. It can model and operate in any type of organization (such as matrix).

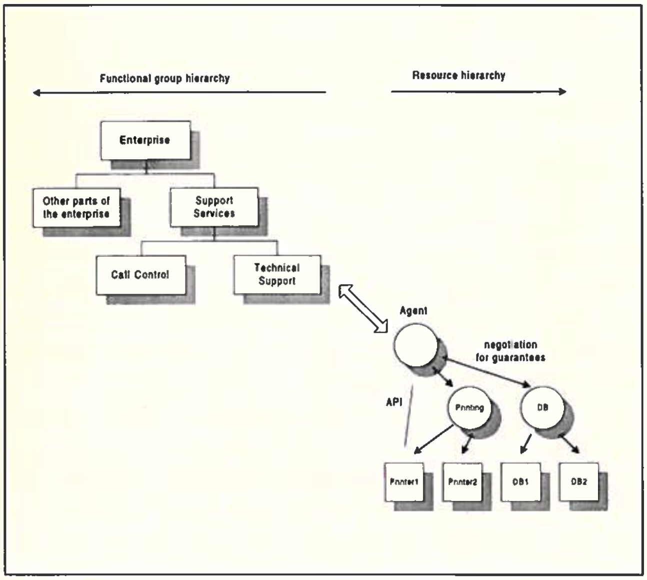

Figure 1. Enterprise model

Each business unit of the enterprise is headed by a supervisor agent that has the right to command its subordinate agents and is also accountable for work done in its business unit. When an agent is asked to accomplish a goal, the requesting agent does not specify how this is to be done. The asked agent will know which resources it can use to perform tasks towards the goal and any organizational procedures and policies that constrain its execution.

The definition of the organizational structure of the enterprise and its resource constraints is supported by the Tool for Enterprise Definition (TED), which is described elsewhere (Abu-Hakima et al., 1995). The organizational structure of the enterprise is defined at the levels of enterprise, functional groups, services, basic services, resource types, and resources. Functional groups contain the services (defined as a sequence of tasks) that comprise the activities of a business unit. The organizational chart is mapped to a hierarchy of functional groups that describes the service provisioning dependencies among the functional groups themselves and the external world.

Basic services provide an interface to a common set of tasks, such as speech recognition. A basic service is associated with resources required to perform its tasks through their resource type. A resource type groups physical resources with the same application programming interface and tasks. Resources may be physical devices (speech recognizer, fax), applications (database, email), or human beings (technical support specialists). For each resource type, TED allows the service designer to define the tasks that the resource type can perform, and attributes such as availability and cost.

Figure 1 shows the functional group hierarchy for an example enterprise that provides a customer help desk service. It also shows the links to the resource hierarchy for a typical functional group agent.

Model of Service Provisioning

Each MANA agent provides services that contribute to the fulfillment of the overall objectives of the enterprise. The success of a MANA-based system is dependent on each agent performing according to its role and providing the services expected of it. We model the service-provisioning dependencies between agents in terms of commitments as defined by (Fikes, 1982). These *commitments* or mutual guarantees are established at set-up time based on the input from TED. The following description is necessarily brief. A more detailed account of the service provisioning model is reported in (Weiss et al., 1995).

The set-up process is supported by two policies: jurisdiction and accountability. *Jurisdiction* describes the reporting structure of the enterprise. This relationship forms a hierarchy of agents. Each agent has one and only one jurisdictional superior or supervisor. A supervisor agent can have any number of subordinate agents. *Accountability* or unity of command specifies which agent is in charge of providing a particular service. Like a manager who has the responsibility of coordinating the efforts of his/her subordinates, this agent is usually different from the agent actually providing the service.

Essentially, in setting up the network of commitments between agents, the approach taken by MANA is to *decouple* service definition from service provisioning. Each definition-provisioning relationship represents a precisely defined *context* for agent interaction. This limits the number of services between which feature interactions may occur. As in Carnot (Huhns and Murindar, 1994) we *isolate* dynamic interactions to agents which have the knowledge to deal with them. While symptoms may occur at any level, only the levels which have direct knowledge of the service intent can be expected to correct them. This isolation of failure improves system reusability.

Design

At this stage of development, the network of commitments between MANA agents is set up manually. Our current design efforts have focused on the run-time system to demonstrate the viability of the envisioned approach. This section describes the agent template, a generic form of a MANA agent that facilitates the definition of agents. It then outlines the design of the run-time components of a MANA agent used to define and implement commitments.

A MANA agent is partitioned into three major areas: a set-up area, an act area, and an information area. The *set-up area* is defined when the agent is specified using TED. It defines the agent's capabilities and contains mechanisms that enable the agent to set up its act area. The *act area* dictates the agent's run-time behavior when interacting with other agents. It provides a goal resolution component, controls a session for each goal that is processed, and allocates the resources to fulfill a goal. The *information area* is used as a log of unexpected events for error handling as well as for posting any information "interesting" to other agents.

The act area is used during agent operation. It contains the information and knowledge needed by the agent to carry out its goals, including its *capabilities* (what the agent can do) and *capacities* (how much the agent can do). This knowledge is represented in terms of goals the agent can achieve, the alternatives in achieving these, and the resources it can use. The act area is created at set-up time and is composed of the following: goal resolution, session control, and resource allocation.

The *goal resolution* component of the act area contains the capabilities of the agent. It has the knowledge of how to achieve the different goals the agent may be asked to do. At run-time, the goal resolution component also contains the current operating context of the agent. This context, a set of sessions, is defined in terms of the goals the agent is working on, what options in achieving these goals the agent has chosen to pursue, feedback on how these current tasks are progressing, etc. Note that an agent is able to process multiple goals concurrently and a new session is created for each incoming goal. As session ids are passed with goals, the agent can trace results returned from other agents.

Goals are decomposed into tasks. At the task-level, *servants* carry out the activities of agents. Servants have the knowledge to perform specific tasks, including how to interface with resource owners. For example, a speech generator servant can perform the following tasks: generate-speech, allocate-device, and free-device. In general, there are several alternatives to perform a certain task. A variety of speech recognizers may be available, each with a different recognition quality and usage cost. The resource allocation component is used to select the appropriate resource, and thus the servant, to execute a task. A second type of servant knows how to send a goal to another agent; they map the task to a goal in another agent. For example, a servant notifies a technical service agent of an incoming call via a handle-call goal and passes the id of the trunk on which the call has been received as a parameter to this goal.

During set-up, agents *negotiate* for the type, quantity, and quality of service required to fulfill their commitments. The agents supplying services interpret the requirements of the

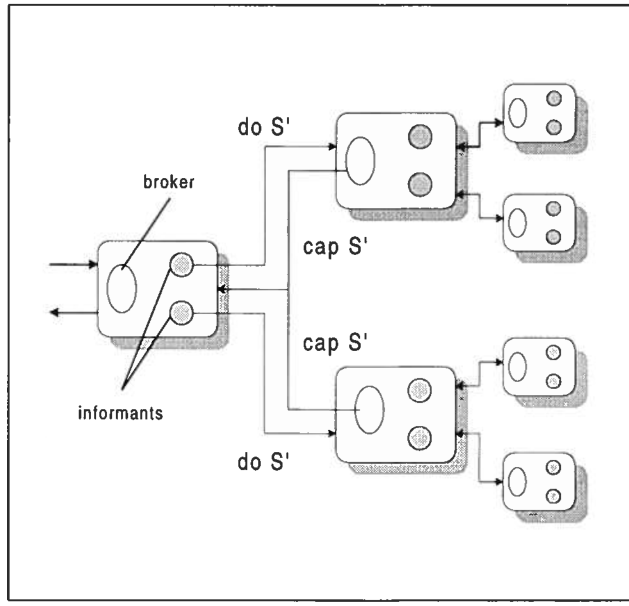

Figure 2. Mapping of commitments

requesting agents and translate them into the type and quantity of resources and the policy by which they will be selected at run-time. Resources and policies are represented within the requesting agents as a set of objects which give them the authority to meet their commitments. For each commitment, the objects passed are: a broker, informants for each resource, servants to drive each resource type, and the policy for run-time selection. Note that regarding resource allocation, *every* agent is a resource to some other agent.

In our model, *informants* are local representations of external resources that an agent can access. The role of the broker is to select the informant that represents the resource that most closely matches service requirements at run-time. It does so by applying the *selection policy* (such as "next resource in sequence" or "lowest loaded re-source first") provided by the resource agent. Note that the selection is performed entirely within the requesting agent.

Figure 2 illustrates how commitments are mapped onto informants and brokers. If the agent to the left commits to perform a service S (do S), it must negotiate for the provision of a service S' (subgoal). As a result of the negotiation, informants are created for each agent that can provide S'. Each commitment is shown as a double-link along which motivation (do S') and authority (cap S' for both capability and capacity) are provided. The motivation is fulfilled by the broker inside the requesting agent by selecting the most appropriate authority and thus supplying agent.

For example, Least Cost Routing (LCR) is a mechanism for reducing the cost of making calls. It is based on the idea of selecting the cheapest carrier for each call based on the time of the call, its destination, etc., and the discounts and special rates offered by the various carriers. Carrier A may offer discount rating after 6 p.m.; carrier B may offer 60% off on all calls made Saturdays until the end of the year; and carrier C may offer special tariffs on the most frequent long-distance calls inside Canada, to the US, and to overseas locations. Specifically, the LCR must be able to accommodate quickly to the ever changing rules of the carriers.

In a network of agents for LCR, one agent would be assigned the responsibility of selecting the carrier for each outgoing call and one agent would represent each of the carriers. The carrier agents are represented by informants inside the agent that selects the carrier. This agent also controls a number of servants for routing the call to a carrier once a selection has been made. When a call comes into the carrier selection agent, this agent matches the guarantees from the carrier agents to the requirements of the call, and selects the cheapest call meeting them. In this set-up, rules can be changed inside the carrier agents without impacting the routing mechanism. Only the guarantees inside the informants representing the carriers must be updated.

It is important to note that the guarantees supplied at set-up time by the agent owning the resource may not accurately reflect the current "health" of the resource owner during system operation. For this reason, in MANA, guarantees in the informants can be updated at run-time based on *feedback* about congestions, failures, etc. that is provided by the resource agents. Feedback allows the system to learn and to optimize its resource allocation policies over time. For example, if the availability of a component decreases, the frequency with which it will be selected by the broker is lowered in the selection policy.

Previous work done in collaboration with Carleton University has investigated the use of market-based mechanisms for resource allocation (Dawes and Kalab, 1993). This work demonstrated that distributed agents bidding for tasks will become specialized at providing certain tasks and ignoring others. The rate at which these agents specialize depends on the resource constraints imposed, the initial cost of the resources, the initial efficiencies at providing a specific task, and the rate at which efficiencies change when contracts are won or lost. The market model described is similar to software agent-based markets foreseen in future public networks. The interactions between the agents requesting and providing services can be viewed as abstractions of software service agents. The results indicate that bidding algorithms and contract protocols can be used to efficiently assign tasks to distributed agents which have become specialized at performing them.

Helpdesk Application

To demonstrate the system, a prototype of a help desk application was built. This application was chosen because the result can be compared with an existing implementation that uses conventional technology. It is also sufficiently complex to exercise the intended functionality of the architecture. The physical set-up of the help desk application is shown in Figure 3. The hardware and software components are distributed onto two sites: one part is located at MITEL, the other at the National Research Council (NRC). The sites are linked via an experimental ATM switch across Ottawa, the OCRInet (Ottawa-Carleton Research Institute Network).

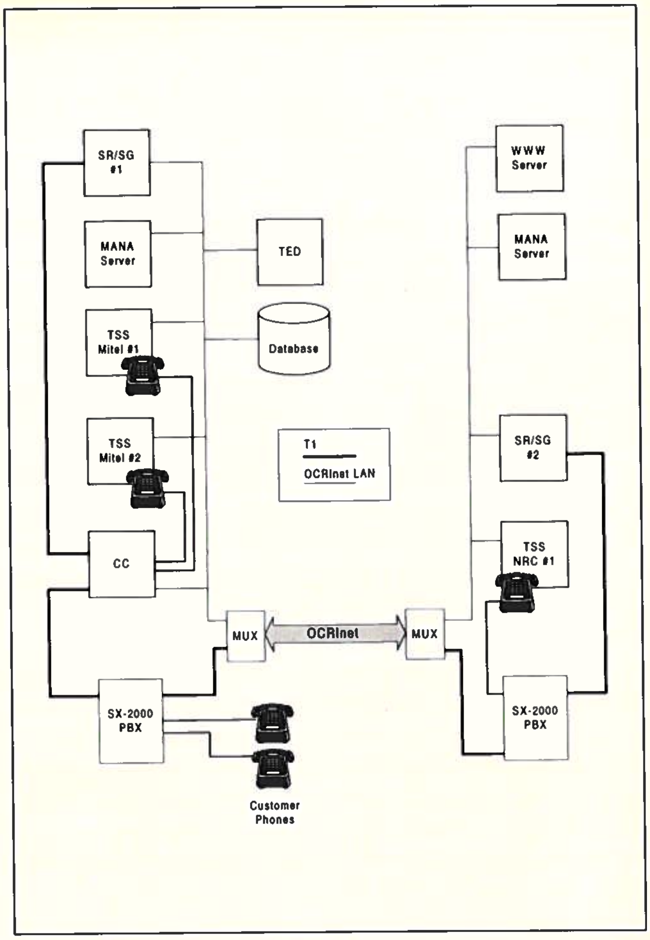

Figure 3. Help desk application

The following is a typical application scenario. The incoming lines are monitored by call control (CC). When a customer call is received, a speech generator (SG) is activated that replies to the customer with a greeting. The customer identification and trouble ticket number are then received by a speech recognizer (SR), verified against a database and routed to one of the available technical support specialists (TSS) located at MITEL and the NRC. A screen pops up at the site of the TSS that provides information on the customer based on the customer identification. The system handles any database requests required to provide this information. In particular, it supplies product documentation needed to answer the customer's problem through a hypertext interface (WWW). The set-up also demonstrates how reliability (that is, the ability to provide the required quality of service) can be built into the application. Note that the speech recognition/generation functionality at MITEL is duplicated at the NRC. Either one of the SR/SG can be used when the basic service of speech generation/recognition is invoked.

A particular capability of our system is that services can be created and customized in TED on the fly. From the service definition in TED, agents are created for each functional group that execute the services associated with them. In our demonstration, we start with just the resources (SR, SG, phones, databases, etc.), and are able to define a complete help desk application in a matter of minutes. This demonstrates the real flexibility of our approach. Once the resource interfaces have been provided, application development time can be dramatically reduced and the application can be customized at any time by the user when his/her requirements change to include (remove) resources into (from) the system. Models of resources and basic services can be sold as "off-the-shelf" components to be integrated into a new system. Third party developers can supply models in addition to those already provided with the system.

Related Work

Several world-wide initiatives are under way to provide seamless communication networks. ROSA (RACE Open Services Architecture) is the definition of an open architecture for integrated broadband communication services (Oshisanwo et al., 1992). Its goal is to provide network services in the form of enterprise capabilities to all users. ROSA is based on an object-oriented approach. Beyond supporting encapsulation, MANA's agent-based approach also describes the semantics of agent interaction to provide autonomous, flexible, and robust operation.

Another initiative is the Telecommunications Information Networking Architecture (TINA) consortium of Bellcore, British Telecom and NTT, Japan (Barr et al., 1993). TINA uses interactive objects on a common distributed platform to model network services in a manner similar to ROSA. The Negotiating Agent approach described in (Velthuijsen, 1993) organizes objects hierarchically with higher-level objects given the ability to negotiate in order to manage interactions between services. This is complementary to MANA's approach to minimize potential feature interactions. All MANA agents negotiate as necessary in trying to satisfy their commitments. It is not clear though whether TINA can model an enterprise as flexibly as MANA.

Common Object Request Broker (CORBA), Distributed Computing Environment (DCE), Distributed System Object Model (DSOM), and related initiatives provide location transparency for object-oriented systems, but are not designed for incorporation in a dynamically changing environment where interactions among people are as important as between devices or applications. MANA is intended to be compatible with these standards, as well as the ROSA and TINA initiatives, such that applications that are based on them can be included in a MANA system.

Moffett et al. (Moffett et al., 1991) have worked on the representation of policies in distributed systems. Management policies are defined as the plans of an organization to meet its goals. Their purpose is to define the goals of an organization and to allocate the resources required to achieve them. In order for policies to be queried and manipulated they are represented as objects in a distributed computer system that influence the actions of agents within the organization. Actions cannot be performed until the two preconditions of motivation and authorization are met, where motivation is defined as the wish of the agent to carry out an action, and authorization as the legitimately obtained power (including access rights) of the agent to carry out the action.

Conclusion

This paper has described an agent-based service environment for business process enabling and the allocation of distributed resources in an enterprise that meets the reliability and real-time constraints of collaborative broadband systems. In our Multi-Agent Architecture for Networked Applications (MANA), organizational goals are translated into policies. For a given set of service requirements, MANA applies those policies to automatically create a network of commitments between agents. In this framework, agents represent either physical devices, applications, human beings, or organizational units.

A major result of our work concerns the definition of agents and the negotiation of commitments. We found that off-line negotiation is more suitable for real-time systems than mechanisms that match agents at run-time. However, our architecture still provides sufficient flexibility by negotiating multiple commitments for a shared resource.

In summary, our contributions are:

- MANA provides a framework for decoupling service definition from provisioning through the separation of policies and agent implementation.
- MANA uses off-line negotiation of service guarantees to meet the real-time constraints of broadband collaborative applications.
- MANA describes services at multiple levels of abstraction, allowing clients to specify their requirements in a language they understand.
- MANA services can be created and customized within an executing system.
- MANA incorporates a learning mechanism to deal with changes to the environment, such as resource failures.

The first application of MANA is a prototype of a help desk which integrates speech recognition, speech generation and client database resources across a network to assist a user calling for service guidance about MITEL products.

The next steps in our work will focus on the design of languages for the definition of service requirements and negotiation. Another major venue is to investigate mechanisms for performance analysis that measure how well services are provided as basis for the renegotiation of guarantees. Also we will more deeply explore the relation to organizational theory, in particular in regard to knowledge-generating, learning organizations, and the work on enterprise modeling by (Fox et al., 1994).

Acknowledgments

This work is the result of a team effort between MITEL, the National Research Council of Canada, and Phoenix. We acknowledge the contributions of Suhayya Abu-Hakima,

Ed Bijman, Richard Deadman, Innes Ferguson, Anthony Lander, Eliana Peres, Peter Perry, Debbie Pinard, Nancy Stonelake, and Mark Van Gulik.

References

Abu-Hakima, S., Ferguson, I., Stonelake, N., Bijman, E., and Deadman, R., A help desk application for sharing resources across high speed networks using a multi-agent network architecture, Workshop on AI in Distributed Information Networks, International Joint Conference on Artificial Intelligence, 1995.

Barr, W.J., Boyd, T., and Inoue, Y., The TINA initiative, IEEE Communications Magazine, 70-76, March 1993.

Dawes, N.W. and Kalab, P., Competitive software agents providing communications services, Technical Report SCE-93-20, Carleton University, 1993.

Fikes, R.E., A commitment-based framework for describing informal cooperative work, Cognitive Science, 331-347, 1982.

Fox, M. and Grunninger, D., Ontologies for enterprise integration, submitted to Conference on Cooperating Information Systems, 1994.

Huhns, M.N. and Murindar, P.S., Automating workflows for service provisioning: integrating AI and database technology, IEEE, 1994.

Moffett, J.D. and Sloman, M.S., The representation of policies as system objects, Conference on Organizational Computing Systems, 117-184, 1991.

Oshisanwo, A.D., Chapman, M.D, et al., The RACE open services architectures project, IBM Systems Journal, 4:691-710, 1992.

Velthuijsen, H., Distributed artificial intelligence for runtime feature-interaction resolution, Computer, 48-55, August 1993.

Weiss, M., Gray, T., and Diaz, A., A service environment for distributed multimedia applications, Workshop on AI in Distributed Information Networks, International Joint Conference on Artificial Intelligence, 1995.

__Michael Weiss__ holds a PhD in computer science and a MS in electrical engineering. He recently joined the Strategic Technology group of MITEL as a software designer. He has done research on distributed artificial intelligence, collaborative design and object-oriented modeling.

__Tom Gray__ has worked in the development of communication products for the last twenty years. He has participated in the evolution of these systems from electronic control of electromechanical relays to the current convergence of fibre optic, microprocessor, and advanced software technology.

__Aurora Diaz__ is a research officer with the Institute for Information Technology of the National Research Council. Her work has focussed mainly on system architectures and data modelling. She has an M.Sc. in Artificial Intelligence and a B.Sc. in Mathematics and Computer Science.

Call for Papers: "Machine Learning"

The Autumn 1995 issue of ***Canadian Artificial Intelligence*** magazine will be a Special Theme issue devoted to Machine Learning. If you are interested in contributing a paper on this topic, please note the deadline for receipt of papers is **September 1, 1995.**

The Winter 1996 issue of ***Canadian Artificial Intelligence*** magazine will be a Special Theme issue devoted to Diagnosis. Contributions for the Winter 1996 issue of ***Canadian Artificial Intelligence*** are due **December 20, 1995.**

If you intend to submit a paper, please inform the editors in advance, so that space can be reserved for your paper. Papers that are not on the above items are also welcome, as are news items, or conference reviews.

Please send your contribution, electronic preferred, with an abstract, photograph, and short biography to:

Dr. Peter Turney
Co-Editor, *Canadian Artificial Intelligence*
Institute for Information Technology
National Research Council Canada
M-50 Montreal Road
Ottawa, Ontario, Canada
K1A 0R6 or — peter@ai.iit.nrc.ca
Telephone (613) 993-8564 Fax (613) 952-7151

STUDENTS' CORNER

Students: This is your chance to discuss your plans for your PhD or Master's Thesis!

Towards a Cognitively Plausible Model for Quantification

Walid Saba

Résumé

Le but de cet article est de démontrer que les quantificateurs dans les langues naturelles n'ont pas un sens de fonction de vérité fixe comme on l'a longtemps maintenu en sémantique logique. Nous suggérons plutôt que les quantificateurs peuvent être mieux représentés comme des procédures d'inférence complexes hautement dynamiques et sensibles au contexte linguistique, et comme des contraintes de temps et de mémoire. Nous proposons un ensemble d'exigences que, nous croyons, un modèle de quantification plausible doit rencontrer. Nous suggérons ensuite quelques étapes computationnelles pour rencontrer ces exigences.

Abstract

This article argues that quantifiers in natural languages do not have a fixed truth-functional meaning as has long been held in logical semantics. Instead, quantifiers can be best be modeled as complex inference procedures that are highly dynamic and sensitive to the linguistic context, as well as time and memory constraints. Proposed are a set of requirements that a cognitively plausible model for quantification must meet. We then suggest some steps towards, computationally, meeting these requirements.

Introduction

Virtually all computational models of quantification are based on some variation of the theory of generalized quantifiers (Barwise and Cooper, 1981) and Richard Montague's (1974) "proper treatment of quantification" (henceforth PTQ).

1. Context-Dependency in Model-Theoretic Semantics

Using the tools of intentional logic and possible-worlds semantics, PTQ models were able to cope with certain context-sensitive aspects of natural language, by devising interpretation relative to a context, where the context was taken to be an index denoting a possible-world and a point in time. In this framework, the intention (meaning) of an expression was taken to be a function from contexts to extensions (denotations). For example, the term president is ambiguous, unless interpreted relative to a place and a point in time. The tools of intentional logic also facilitated, for the first time, a systematic and mathematically elegant approach to the "intentionality problem." Montague's system was perhaps most known for his ingenious use of possible-worlds to handle intentional expressions of the form "John is looking for a unicorn." Such an expression is ambiguous, unless interpreted relative to a possible-world (thus allowing the possible interpretation where John is in a make-believe world where there are unicorns, and John is seeking one). In what later became known as indexical semantics, Kaplan (1979) suggested adding other "coordinates" defining a speaker, listener, location, etc. Using the extended model, he showed how an utterance such as "I called you yesterday" expressed a different content, whenever the speaker, the listener, or the time of the utterance changed.

2. Quantifier Ambiguity

Although model-theoretic semantics were able to cope with certain context-sensitive aspects of natural language, the intentions (meanings) of quantifiers, however, as well as other functional words, such as sentential connectives, are taken to be constant. That is, such words have the same interpretation regardless of the context (Forbes, 1989). In such a framework, all quantifiers have their meaning grounded in terms of two logical operators$\forall$(for all), and $\exists$ (there exists). Consequently, all quantifiers are modeled, indirectly, by two simple logical connectives: negation and either conjunction or disjunction. From this standpoint, quantifier ambiguity is often reduced to "scoping" ambiguity, a problem that has been extensively studied by formal semanticists (Cooper, 1983; Le Pore and Garson, 1983; Partee, 1884) and computational linguists (Harper, 1992; Alshawi, 1990; Pereira, 1990; Moran, 1988). The problem can be illustrated by the following examples:

1. Every student in CS404 received a course outline.
2. Every student in CS404 received a grade.

The syntactic structures of (1) and (2) are identical, and thus according to Monatgue's PTQ both sentences would have the same logical translation. Hence, the translation in PTQ of the second sentence would incorrectly state that students in CS404 received different course outlines. Instead, the desired reading is one in which "a" has a wider scope than "every" resulting in a translation that asserts that there is a single course outline for the course CS404, an outline that all students received. Clearly, such resolution depends on general knowledge of the domain; typically students in the same course outline receive the same course outline, but different grades. Due to the strict compositionality requirement, PTQ models cannot cope with such inferences, since that entails determining the meaning of the whole using contextual background knowledge beyond the text under interpretation. Consequently, a number of syntactically motivated rules for resolving scoping ambiguities are typically suggested. In general, these rules suggest an "ad hoc" (Moran, 1988) "semantic" ordering between functional words, that is, an ordering fitted to a particular set of examples.

Such an approach is problematic, in our opinion, for the following reasons:

a) the indexical semantics approach to modeling context is problematic since, in general, it is not at all clear what set of indices to include in a specific situation (see Cresswell, 1974).

b) the scope of quantifiers, as argued by Zeevat (1989), is usually given by "the linguistic context rather than by the linguistic rule that is responsible for their appearance in a sentence."

c) partitioning a priori quantifier ambiguity into scoping and other "types" of ambiguity inevitably leads to a labeling problem (Corriveau, 1995), and to ad hoc rules that are too specialized to address the general problem of quantification.

What we suggest, instead, is that quantifiers be treated as ambiguous words whose interpretation is always contextual and highly dependent on time and memory constraints. In order to justify the role of time and memory constraints in the disambiguation of quantifiers, consider the following examples:

3. Cubans like rum more than vodka.
4. Students in CS404 are not allowed to work in groups.

Our intuitive reading of (3) suggests that we have an implicit "most" preceding the sentence, while in (4), it is most likely that the intended interpretation is "all students in CS404 are not allowed to work in groups." We argue that such inferences are dependent on time constraints and constraints on short-term memory (STM). That is, since the set of students in CS404 is a much smaller set that the set of "Cubans," it is conceivable that we are able to perform an exhaustive search over the set of all students in CS404 to verify the proposition in (4) within some time and memory constraints. However, in (3), we are most likely performing a "generalization" based on few examples that are currently activated in STM. The basic idea here is that a quantifier is essentially a complex inferencing procedure that is dependent not only on the linguistic context, but on time and memory constraints. In order to further argue this point, consider the following examples, where an additional complexity is added due to the role of discourse:

5. This room is full over-ambitious accountants.
   Everyone works at least 14 hours.
6. John's report on Japanese professionals is remarkable.
   Everyone works at least 14 hours.

It seems probable that, through a combination of disambiguation and reference resolution, "everyone" will be interpreted as "each accountant" in (5), but as "most Japanese professionals" in (6), since the set of Japanese professionals has an indeterminate cardinality for most readers.

From the above examples, it reasonable to assume that the cardinalities of the quantified sets (concepts) are directly related to requirements on processing time and working memory. Such a claim is also in agreement with psychological experiments (Freeman and Stedmon, 1986) on children age 5-12 suggesting that quantification is often dependent, among other things, on the size of the quantified concept. Johnson-Laird (1994) has also argued recently that an appropriate model must allow sets of possibilities to be considered in a "highly compressed way," and allow for the distinction between implicit and explicit information. The suggestion here is that we tend to process sentences like (3) using generic quantification, due to processing limitations of working memory.

Requirements for a Cognitively Plausible Model

Although an argument for a context-dependent, and cognitively plausible model for quantification can be made, such a model must satisfy certain requirements, which we group into two sets of requirements:

a)"Inferencing" requirements.

b) "Cognitive plausibility" requirements.

1. Inferencing requirements

The inferencing requirements are such that valid inferences that are typically captured by the extensional, fixed, and truth-functional model of quantification, must remain to be valid in a cognitively plausible model. Thus, the fixed extensional meaning of quantification that is currently the model, must be "subsumed" as a very special case in the new model. In this respect we agree with the argument made in (Fauconnier, 1994) that logic usually abstracts a solution that could work rather well in some idealized and controlled situation (referred to as "models" in formal logic), but fails in modeling the idiosyncratic nature of human languages.

Our goal therefore is to develop a model that would preserve a number of formal properties that are generally attributed to quantifiers, such as the following: (for a detailed account of such formal properties, see (Peres, 1990)):

1. Persistence (left increasing monotonicity)
2. Right increasing monotonicity
3. Left decreasing monotonicity
4. Reflexivity
5. Symmetry

According to the symmetry property, for example, if "some man from Detroit owns a 1965 Corvette" is true, then the set of "those who own a 1965 Corvette" must include "some man from Detroit." In addition to these formal properties, a cognitively plausible model must allow for such inferences to be made:

every PQ => most PQ	most PQ => many PQ
many PQ => several PQ	several PQ => some PQ
some PQ => a/an PQ	a/an PQ => at least one PQ

Thus, for example, if we know "every student in CS404 has a PC" to be true, then we can definitely say that "most students in CS404 have a PC" is true, etc.

2. Cognitively Plausible Requirements

Besides the inferencing requirements, we require that an appropriate model for quantification must adequately account for the following:

time constraints
constraints on working memory
linguistic context

Steps Towards a Cognitively Plausible Model

The model that we propose is a time-constrained model, where "time" is an ever present factor in determining the final interpretation (i.e., given more or less time, the reader would reach a different interpretation of the same piece of text). Moreover, we also devise a rather non-trivial memory model, in order to capture ambiguities that are due to limitations on working memory. Inferencing in the disambiguation of quantifiers is performed in a time-constrained fashion, using a specified time limit, a decay rate, and information regarding working memory contents (to address some discourse-related problems), and working memory (and STM) capacity. Initially, a quantifier will be assigned a default function that is specialized using various parameters regarding the linguistic context and information regarding time and memory constraints.

Conclusion

We have tried to argue that quantifiers do not have a fixed truth-functional meaning as has long been held in logical semantics. We suggest instead a cognitively plausible approach to quantification where quantifiers are considered to be ambiguous, and where their meaning is always contextual, and highly dependent on time and memory constraints. Such a model must meet two sets of requirements, namely the requirement that the model must preserve valid inferences in the model-theoretic approach. Moreover, the model must account for time and memory constraints, thus meeting our "cognitive plausibility" requirements.

References

Alshawi, H. (1990). Resolving Quasi Logical Forms, Computational Linguistics, 6(13), pp. 133-144.

Barwise, J. and Cooper, R. (1981). Generalized Quantifiers and Natural Language, Linguistics and Philosophy, 4, pp. 159-219.

Barwsie, J. and Etchemendy, J. (1993) Model-Theoretic Semantics, In M. I. Posner (Ed.), Foundations of Cognitive Science, MIT Press, pp. 207-244.

Charniak, E. (1983). Passing Markers: A Theory of Contextual Influence in Language Comprehension, Cognitive Science, 7(3).

Cooper, R. (1983). Quantification and Syntactic Theory, D. Reidel, Dordrecht, Netherlands.

Corriveau, J.-P. (1995). Time-Constrained Memory, to appear, Lawrence Erlbaum Associates, NJ.

Cresswell, M. J. (1973). Logic and Languages, Methuen & Co. Ltd.

Fauconnier, G. (1994). Mental Spaces - Aspects of Meaning Construction in Natural Language, Cambridge University Press.

Forbes, G. (1989). Indexicals, In D. Gabby and F. Guenthner (Eds.), Handbook of Philosophocial Logic, Volume IV, D. Reidel. pp. 463-490.

Freeman, N. H. and Stedmon, A. J. (1986). How Children Deal with Natural Language Quantification, In I. Kurcz, et al. (Eds.), Knowledge and Language, North Holland. pp. 21-48.

Harper, M. P. (1992). Ambiguous Noun Phrases in Logical Form, Computational Linguistics, 18(4), pp. 419-465.

Johnson-Laird, P. (1994). Mental Models and Probabilistic Thinking, Cognition, 50, pp. 198-209.

Kaplan, D. (1979). On the Logic of Demonstratives, The Journal of Philosophical Logic, 8, pp. 81-98.

Le Pore, E. and Garson, J. (1983). Pronouns and Quantifier-Scope in English, Journal of Philosophical Logic, 12, pp. 327-358.

McCarthy, J. (1980). Circumscription - A Form of Nonmonotonic Reasoning, Artificial Intelligence, 13, pp. 27-40.

Montague, R. (1974). Formal Phislosophy: Selected Papers of Richard Montague, edited with an introduction by Richard Thomason, Yale University Press.

Moran, D. B. (1988). Quantifier Scoping in the SRI Core Language, In Proceedings of 26th Annual Meeting of the ACL, pp. 33-40.

Partee, B. (1992). Syntactic Categories and Semantic Type, In M. Rosner and R. Johnson (Eds.), Computational Linguistics and Formal Semantics, Cambridge University Press, pp. 97-126.

Partee, B. (1984). Quantification, Pronouns, and VP-Anaphora, In J. Groenedijk et al. (Eds.), Truth, Interpretation, and Information, Dordrecht: Foris.

Pereira, F. C. N. (1990). Categorial Semantics and Scoping, Computational Linguistics, 16(1), pp. 1-10.

Pereira, F. C. N. and Pollack, M. E. (1991). Incremental Interpretation, Artificial Intelligence, 50, pp. 37-82.

Peres, J. (1990). Basic Aspects of the Theory of Generalized Quantifiers, In M. Filgueiras et al (Eds.), Lecture Notes in Artificial Intelligence, Volume 476 (NLP), Springer Verlag, pp. 141-157.

Quillian, R. (1968). Semantic Memory, In M. Minsky (Ed.), Semantic Information Processing, MIT Press.

(continued on page 32)

PRECARN UPDATE

NOUVELLES DE PRECARN

Jean-Claude Gavrel

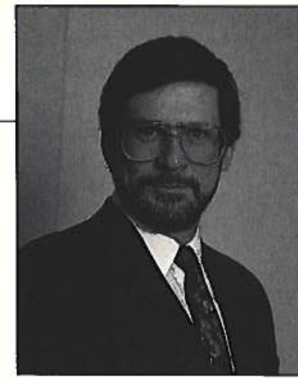

PRECARN Welcomes Third President

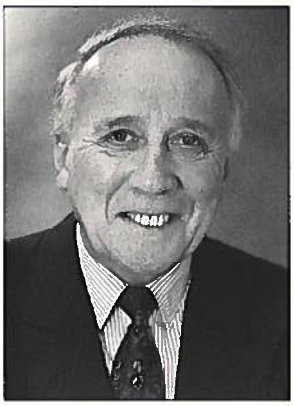

Harry Rogers, President, PRECARN

When Mac Evans announced that he would be moving to take on the challenges of President of the Canadian Space Agency, PRECARN went on a nation-wide search to find its third full-time President. No one could have imagined that the search would succeed in attracting someone as prominent as Harry Rogers. In many ways, the job is a perfect match for Rogers. It fits his philosophy on building Canada's technical strength, his practical on-the-job experience, and his personal goals. "I'm a firm believer that any country must have policies and programs that fit its commercial and industrial structure," Rogers explains. "Canada's industrial structure is dominated by a few large resource companies, large US branch plants and a huge number of small- and medium-sized Canadian-owned businesses. When skills and financial capacity are limited, collaboration among companies is critical. Networking is therefore one of our premier strategies." Before joining PRECARN on July 1, Rogers was Deputy Minister in the federal government's Office of Federal Economic Development in Ontario. He also served as Deputy Minister, Regional Industrial Expansion and Secretary, Ministry of State for Science and Technology. Rogers' private sector experience includes a variety of management and senior management posts at Xerox of Canada, Xerox Corp., and the Ford Motor Company.

Phase 2 Projects Approved

As a result of a Request for Research Proposals issued in March, the following 15 projects were approved by the Board of Directors for funding:

Concept Definition Program: A Low Cost, High Speed Camera Orienting Device Led by Canpolar East, this project is interested in finding a cost-effective technology for orienting a video camera very quickly.

Visual-Vestibular Immersive Display (VIVID) The VIVID study, led by Aastra Aerospace, involves the assessment of sensation of self-motion in humans caused through tilt while viewing a virtual reality display and development of specifications for a prototype system.

Discovery of Control Rules for Operator Actions Lobbe Technologies proposes to analyze data from an industrial process control application to determine how rough sets theory could be used to improve the control.

A Concept Definition Project for a Voice Activated Maintenance Management System Led by GasTOPS, this project is interested in finding out if current, state-of-the-art in voice recognition technology is capable of being used in a noisy industrial environment for interacting with maintenance data management systems.

Enhanced Sensing and User Interaction for Mobile Robots This project, led by Cyberworks Ltd., will try to combine technologies from an IRIS project at McGill with a low-to-medium cost commercial robotic platform from CyberWorks.

Feasibility Studies for Long Term Program Telerobotics in Compliant Environments Led by l'Institut de recherche d'Hydro-Québec (IREQ), this project involves research in control, modeling, and computer vision to develop robust teleoperation and autonomous execution of tasks in oscillatory and contact situations.

System for Intelligent Process Operation (SIPO) The SIPO Project, under the leadership of Ontario Hydro, proposes to develop an automated, intelligent, operator assistant for the operation of continuous process plants under emergency conditions.

ACROBAT: Autonomous Climbing Robots and Associated Technologies The team, led by Ontario Hydro, will conduct research into autonomous climbing robotic technologies. The ultimate goal will be to design, build, and test a prototype climbing robot.

Application of Advanced Database Technologies to Dynamic Evaluation of Situation Models Led by MacDonald Dettwiler, the goal of this project is to develop advanced decision support tools that will train the machines to think like the users.

The Development of Intelligent Sensing Systems for Oil and Mining Industries This project, led by the Alberta Research Council, addresses a need for accurate sensing and intelligent use of the resulting information to increase productivity in mining and extraction.

Short-Term Programs Development and Applications of New Technology for Improved Assessment of Patient Outcomes Led by Western Clinical Devices, the team proposes to develop a prototype robotic system to assist in the assessment of patient outcomes in cases of shoulder surgery.

Research into Learning Techniques to Aid Automated Detection of Environmental Hazards using Remotely Sensed Imagery The project, led by MacDonald Dettwiler, will extend the performance capabilities of an oil-slick detection system based on the analysis of Synthetic Aperture Radar (SAR) data from the European ERS-1 remote sensing satellite.

A Prototype Video Conferencing User Interface Including Sound-Source-Directed Control of Audio and Video Acquisition Under the leadership of MPR Teltech, this Project intends to develop a prototype system to control video and audio acquisition using sound localization technology.

Virtual Environments for Remote Operations "Virtual Environments" are being applied in areas such as: simulation, training, architecture, and design. Spar Aerospace will explore emerging and promising applications of environments that exist but are otherwise inaccessible to humans.

Intelligent Assessment of Rotary Machinery Components Using Filter Debris Analysis This project, led by Tektrend International Inc., proposes to develop new recognition techniques for analyzing the accumulated wear debris contained in lubricating oils entrapped in filters of helicopters.

For additional information on these projects, please contact Dr. Jill Sanders at the PRECARN office.

PRECARN's Vice-President Moves On

If you are a regular subscriber to this article, you will have seen the name Jean-Claude Gavrel, PRECARN's Vice-President. Jean-Claude announced in June that he has accepted the position of Vice-President, R&D Management, with the Centre de recherche informatique de Montréal, effective late-August.

Jean-Claude has made outstanding contributions to the success and evolution of PRECARN. He came on-board at a time when his project management skills were essential to the start-up of the Phase 1 projects. He is also largely responsible for the successful redesign of the Phase 2 Program, which is about to get underway. PRECARN wishes him the very best in his new position.

To obtain more information on PRECARN, please contact Dr. Jill Sanders or Lise McCourt at PRECARN, Telephone: (613) 727-9576, Facsimile: (613) 727-5672 or E-Mail: sanders@precarn.ca or mccourt@precarn.ca.

Towards a Cognitively Plausible Model for Quantification

(continued from page 30)

Shastri, L. (1989). Default Reasoning in Semantic Networks: A Formalization of Recognition and Inheritance, Artificial Intelligence, 39(3), pp. 283-355.

Shastri, L. and Ajjanagadde, V. (1993). From Simple Associations to Systematic Reasoning: A Connectionist Representation of Rules, Variables, and Dynamic Bindings Using Temporal Synchrony, Behavioral, & Brain Sciences, 16, pp. 417-494.

Zeevat, H. (1989). A Compositional Approach to Discourse Representation theory, Linguistics and Philosophy, 12, pp. 95-131.

***Walid Saba** is a part-time Ph.D. student at Carleton University and a consultant with the Expert Systems Development Group at AT&T Bell Laboratories, Middletown, NJ. He can be reached at walid@eagle.hr.att.com, or walid@research.att.com.*

BOOK REVIEWS

Computational organization theory Kathleen M. Carley and Michael J. Prietula (editors) (Carnegie Mellon University) Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, xvii+318 pp; hardbound, ISBN 0-8058-1406-X, US$89.95

Reviewed by
Stephen Regoczei
Trent University

We live in the age of organizations, in the sense that organizations—governments, businesses, and non-profits—dominate public life. Organizations can be looked upon both as AI systems and as artificial life forms (Regoczei, 1993). So given their significance and their artificial character, it is clear that organizations are good candidates for study with the tools of computational modeling. The systematic study of organizations is one of the sciences of the artificial (Simon, 1981). Carley and Prietula have brought together some typical current examples of research in organizational theory using simulation models. This is a useful contribution. The models are interesting and suggestive of promising further work. Yet in their current form, they fall short of our intuitive expectations. Something is missing.

The book contains thirteen different papers on topics such as organizational structure, decision making, organizational learning, and team coordination. In spite of the seemingly wide range of the material, the approaches are unified by a certain spirit: the spirit of SOAR. They also share the vision of organizations as agglomerations of agents, tasks, and structures. The agents carry out tasks, and they communicate and cooperate in patterns defined by the organizational structure. That is all. Those of us who study real-life organizations, and specific administrative practices within organizations may feel short-changed by this simplicity. This model could just as well describe a box full of bouncing gas molecules. The agents are molecules, the task is bouncing, the structure is given by the walls of the box, and communication is achieved by collisions. Is this really all there is to organizations? Does this match our intuitions of what organizations are really like?

I can't help it, but a phrase like "computational organization theory" reminds me of anthills—more precisely, Edward O. Wilson's theory of anthills. This has been the case ever since I read Robert Wright's "Three Scientists and Their Gods." Wright juxtaposes the views of Ed Fredkin on cellular automata, Wilson on sociobiology, and Kenneth Boulding on the hierarchy of systems. After being exposed to Wright's book, one sees computational automata and emergent phenomena everywhere. This is not surprising. Ed Fredkin, according to Wright, thinks that the whole world is a cellular automaton—a vast simulation experiment, the outcome of which even God cannot predict. Within such a mixed domain of discourse, we have to distinguish an anthill that is there because the ants made it from a programmer-created piece of software that simulates the behaviour of, and indeed is, a finite-state automaton. The ants are natural and the programmer's creation artificial—but I confess I find the division blurred. The clincher is that corporations and software packages form a continuum that is hard to partition. Is a corporation more like a natural or an artificial object? The analogy is with plastic flowers and artificial light. The software is like a plastic flower, but the corporation is real artificial life. What is a realistic model of an organization? It certainly shouldn't be like that of the proverbial grotesque example of modeling where the engineer proclaims, "Let us assume that the chicken is a sphere." No, chickens are not very sphere-like. The point is that organizations are not merely mass/energy flows; they have a psycho-social cognitive component. This cognitive component is distinguishable from the people that work for the organization. Because of this cognitive component, organizations cannot be satisfactorily modeled as if they were merely physical systems, such as the solar system, or molecules bouncing around in a box. The cognitive component needs to be modeled using knowledge-based AI techniques.

The organization is also a natural system: it grows, evolves, and maintains itself in the sense of autopoiesis and homeostasis. This would be modelable using artificial life techniques. The point is that traditional mathematical techniques cannot be used to adequately model the cognitive and a-life-like features of organizational behaviour. Verkama et al. acknowledge the importance of AI and a-life techniques and then conclude that although the newer techniques are "appealing, the value of the traditional mathematical approach to organizational problems should not be underestimated." They then feel free to continue with what only could be described as the doing of "pure applied mathematics."

Nevertheless, the current book goes far beyond, for example, the Jay Forrester style of system dynamics. The shortcoming of Forrester models is that they treat the organization as if it were a physical system of interacting variables devoid of any psycho-social components. There are several crucial issues here. The model does not consist of interactive components but rather of a restricted set of quantifiable properties. In other words, variables. But variables model the quantifiable properties and not the behaviour of entities themselves. This presents a conceptual difficulty that is hard to overcome. In fact, it is the analog of object-oriented versus variables-oriented programming. It is an asset of the book that the modeling is mechanism-oriented and not variable-oriented (p. xiv).

Historically speaking, computers and organizations have been closely linked. IBM's business success is a proof of that. Yet computational modeling of organizational behaviour preceded the invention of programmable electronic computers. Much of operations research can be construed as the modeling of organizational processes. The motivation behind general systems theory was largely that of the modeling of organizations as open systems. In light of this historical path, we would like computational organization theory to offer assistance with current concerns such as TQM and BPR. Adding another point: information systems within organizations are now computer-based. Surely a computational organization theory should accept the notion that databases constitute working models of the corporation. Yet none of these topics are even mentioned in the book.

What does it mean to model organizations? More precisely, if we manage to create a model of some aspect or component of organizational behaviour, such as decision making, hierarchical structure, or investment strategy, can we then claim that we have modeled organizations? In terms of the perspectivist parable of the six blind men and the elephant, if one of them thinks that the elephant is a spear, and he manages to model spears in general, did he then thereby model the elephant? Suppose we have a mathematical technique that seems to fit some aspect or feature of organizations. Can we then claim that we have modeled the organization?

Organizations are abstract agents/actors. They are artificial life forms. They are certainly not identical with the groups of people who work for the organization. This is a version of the cliché that systems are greater than the sum of their parts. The corporation employs not a person, but an employee role. Organizations are cognitive agents. In fact, they are the most successful distributed AI systems ever constructed. This is an unpalatable thought: while the Lisp hackers hacked, the business types implemented actual AI. That is certainly one reason why the study of organizations is one of the sciences of the artificial in the sense of Herbert Simon. In other words, an organization is an artificial computational automaton. Now, to create a computer model of the organization, we are limited in the types of automata available to us today. Organizations are not Turing machines; they are not even von Neumann processors. We need a host of new computational automata: Hayek machines, such as a market that computes price and demand levels, Darwin machines that survive in the face of competition, Wilson machines that are like the anthills that run on pheromones, and Dawkins machines that reconfigure themselves through the propagation of memes. The point is that an organization cannot be realistically modeled merely by cellular automata or by a Forrester-style system dynamics model. Organizations have a discourse component; they are laden with talk, text, concepts, and corporate cultures.

Several of the authors in the book are well aware of these issues and shortcomings. Crowston, for example, discusses genetic algorithms as applied to the evolving of novel organizational forms. But then he quotes Hannan and Freeman to the effect that we normally "distinguish hospitals, prisons, political parties, universities, stock exchanges, coal mines, and fast-food chains". (p.24) Crowston then wonders what a cross between a hospital and a coal mine would look like, or how it would be represented. He does not know, and probably nor will many of the readers of the book. The main issue here is that the computational models are not yet conceptually rich enough to plausibly model the range of organizations such as given by the above list.

Modeling knowledge within organizations proves to be elusive. This point is made by several of the contributions. Yet it is clear to organizational theorists that corporate cultures do make a difference—and even determine whether an organization succeeds or fails. Empirically, the knowledge routines that govern the internal functioning of an organization prove to be difficult to capture or specify.

The direction of the research described in the book seems promising, but the reader gets that "we're-not-there-yet" feeling. What is missing? Richard Cyert wrote the preface to the book and I think that he pinpointed the missing ingredient. He should have the last word of this review. "These [research] movements are all in the right direction. The important caution ... is that all work in computational organization theory must continue to pay attention to the real world. We must not commit the fatal error of taking all of our problems from the literature. Live organizations must be studied, and the results of those studies can then be utilized in simulated models of organizations."

References

Regoczei, S. "Reprogramming the software that runs on organizations: Social agents as knowledge-based computational automata." AAAI-93 Workshop: AI and Theories of Groups and Organizations: Conceptual and Empirical Research, Washington, DC, 1993.

Simon, Herbert Alexander. The sciences of the artificial, 2nd edition. MIT Press, 1981.

Stephen Regoczei *teaches courses at Trent University in AI, modeling, artificial life, and information systems in organizations. He also studies administration as a profession. He still can't quite decide whether administrators are indispensable professionals or harmful parasites.*

BRIEF NOTES

Prime Time freeware for AI, issue 1-1 *Mark Kantrowitz (editor)* Sunnyvale, CA: Prime Time Freeware, 1994, 220 pp and two CD-ROMs; ISBN 1-881957-11-X, $60.00

Prime Time freeware is a large package of free AI software chosen from the AI Repository at Carnegie Mellon University. It covers all areas of AI, AI programming languages, and utilities. The entire package amounts to 5 gigabytes, compressed onto two CD-ROMs, along with a book that explains how to use the package and lists and indexes all the software in the package. The publisher plans to issue a new edition annually.

The Journal of Artificial Intelligence Research: Volume 1, August 1993-June 1994 *Steven Minton (executive editor)* San Francisco, CA: Morgan Kaufmann Publishers, 1994; paperbound, ISSN 1076-9757, ISBN 1-55860-347-6; no price listed.
[Also published online by the AI Access Foundation; *http://www.cs.washington.edu/research/jair/home.html*]

The Journal of Artificial Intelligence Research is one of the first fully refereed scientific journals whose primary mode of publication is on the Internet. Papers are submitted electronically, refereed, revised, and then immediately published in the Usenet group *comp.ai.jair.papers*. The papers are also placed in an ftp site. A printed version of the journal will appear each year. The first volume, now printed, contains the 13 papers accepted in the first 10 months of the journal's existence.

BOOKS RECEIVED

Books marked with a + in the list below are scheduled for review in a future issue. Reviewers are still sought for books marked with a *. Readers who wish to review books for *Canadian Artificial Intelligence* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4, or send electronic mail to gh@cs.toronto.edu or gh@cs.utoronto.ca. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

Authors and publishers who wish their books to be considered for review in *Canadian Artificial Intelligence* should send a copy to the book review editor at the address above. All books received will be listed, but not all can be reviewed.

Cognitive systems engineering for user—computer interface design, prototyping, and evaluation *Stephen Andiole and Leonard Adelman* (Drexel University and George Mason University) Hillsdale, NJ: Lawrence Erlbaum Associates, 1995, xiii+272 pp; hardbound, ISBN 0-8058-1244-X, US $59.95

Machine translation: An introductory guide *D. Arnold, L. Balkan, R. Lee Humphreys, S. Meijer, and L. Sadler* (University of Essex) Oxford: NCC Blackwell, 1994, viii+240 pp; hardbound, ISBN 1-85554-246-3, US $49.95, paperbound, ISBN 1-85554-217-X, US $19.95

Readings in human-computer interaction (second edition): Toward the year 2000 *Ronald M. Baecker, Jonathan Grudin, William A.S. Buxton, and Saul Greenberg (writers and editors)* (University of Toronto; University of California, Irvine; University of Toronto; and University of Calgary) San Francisco: Morgan Kaufmann Publishers, 1995, xxi+950 pp; paperbound, ISBN 1-55860-246-1, US $49.95

Natural language processing: A Paninian perspective *Akshar Bharati, Vineet Chaitanya, and Rajeev Sangal* (Indian Institute of Technology, Kanpur) New Delhi: Prentice-Hall of India, 1995, xviii+220 pp; paperbound, ISBN 81-203-0921-9, Rs 95.00, US $20.00

Computational phonology: A constraint-based approach *Steven Bird* (University of Edinburgh) Cambridge University Press (Studies in natural language processing, edited by Branimir K. Boguraev), 1995, xv+203 pp; hardbound, ISBN 0-521-47496-5, US $44.95

Trajectories through knowledge space: A dynamic framework for machine comprehension *Lawrence A. Bookman* (Sun Microsystems Laboratories) Boston: Kluwer Academic Publishers (The Kluwer international series in engineering and computer science: natural language processing and machine translation, edited by Jaime Carbonell), 1994, xx+271 pp; hardbound, ISBN 0-7923-9487-9, US $89.95

Language and learning for robots *Colleen Crangle and Patrick Suppes* (Stanford University) Stanford, CA: Center for the Study of Learning and Information (CSLI lecture notes 41), 1994, xxi+276 pp; distributed by the University of Chicago Press; hardbound, ISBN 1-881526-20-8, US $49.95, paperbound, ISBN 1-881526-19-4, US $21.95

***Artificial intelligence and creativity: An interdisciplinary approach** *Terry Dartnall (editor)* (Griffith University) Dordrecht: Kluwer Academic Publishers (Studies in cognitive systems, edited by James H. Fetzer, volume 17), 1994, xiv+455 pp; hardbound, ISBN 0-7923-3061-7, US $130.00

Artificial intelligence: Theory and practice *Thomas Dean, James Allen, and Yiannis Aloimonis* (Brown University, University of Rochester, and University of Maryland) Redwood City, CA: The Benjamin/Cummings Publishing Company, 1995, xxi+563 pp; hardbound, ISBN 0-8053-2547-6, no price listed

Deixis in narrative: A cognitive science perspective *Judith F. Duchan, Gail A. Bruder, Lynne E. Hewitt* (State University of New York at Buffalo) Hillsdale, NJ: Lawrence Erlbaum Associates, 1995, xix+552 pp; hardbound, ISBN 0-8058-1462-0, US $89.95; paperbound, ISBN 0-8058-1463-9, US $39.95

Natural language computing: An English generative grammar in Prolog *Ray C. Dougherty* (New York University) Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, xlvi+349 pp and 3.5-inch disk for IBM PC or compatible; hardbound, ISBN 0-8058-1525-2, US $69.95; paperbound, ISBN 0-8058-1526-0, US $36.00

Global perspectives on the ecology of human-machine systems: volume 1 [of an unnamed two-volume set] *John Flach, Peter Hancock, Jeff Caird, and Kim Vicente (editors)* (Wright State University, University of Minnesota, University of Calgary, and University of Toronto) Hillsdale, NJ: Lawrence Erlbaum Associates (Resources for ecological psychology series, *edited by Robert E. Shaw, William M. Mace, and Michael T. Turvey*), 1995, xvi+412 pp; hardbound, ISBN 0-8058-1381-0, US $69.95; paperbound, ISBN 0-8058-1382-9, US $36.00; two-volume set, no ISBN listed, hardbound, US $99.95, paperbound, US $52.00

Local applications of the ecological approach to human-machine systems: volume 2 [of an unnamed two-volume set] *Peter Hancock, John Flach, Jeff Caird, and Kim Vicente (editors)* (University of Minnesota, Wright State University, University of Calgary, and University of Toronto) Hillsdale, NJ: Lawrence Erlbaum Associates (Resources for ecological psychology series, *edited by Robert E. Shaw, William M. Mace, and Michael T. Turvey*), 1995, xviii+470 pp; hardbound, ISBN 0-8058-1379-9, US $75.00; paperbound, ISBN 0-8058-1380-2, US $39.95; two-volume set, no ISBN listed, hardbound, US $99.95, paperbound, US $52.00

Neural networks in computer intelligence *LiMin Fu* (University of Florida, Gainesville) New York: McGraw-Hill (McGraw-Hill series in computer science), 1994, xix+460 pp and 5.25 inch diskette; hardbound, ISBN 0-07-911817-8, cdnUS $89.95

Limits to parallel computation: P-completeness theory *Raymond Greenlaw, H. James Hoover, and Walter L. Ruzzo* New York: Oxford University Press, xiii+311 pp; hardbound, ISBN 0-19-508591-4, cdn$81.50

Hybrid problems, hybrid solutions *John Hallam (editor)* (University of Edinburgh) Amsterdam: IOS Press and Tokyo: Ohmsha (Frontiers in artificial intelligence and applications, volume 27), 1995, vii+239 pp; paperbound, ISBN 90-5199-216-5 (IOS) and 4-274-90040-1-C3000 (Ohmsha), no price listed

+Expertise and technology: Cognition and human-computer cooperation *Jean-Michel Hoc, Pietro C. Cacciabue, and Erik Hollnagel (editors)* (CNRS-University of Valenciennes; CEC Joint Research Centre, Ispra, Italy; and Human Reliability Associates) Hillsdale, NJ: Lawrence Erlbaum Associates (Expertise: Research and applications series, edited by Robert R. Hoffman, Michelene T.H. Chi, K. Anders Ericsson, and Gary A. Klein), 1995, xvi+289 pp; hardbound, ISBN 0-8058-1511-2, US $45.00

Design and production of multimedia and simulation-based learning material *Ton de Jong and Luigi Sarti (editors)* (University of Twente and Istituto per le Tecnologie Didattiche, Consiglio Nazionale delle Ricerche) Dordrecht: Kluwer Academic Publishers, 1994, xvi+253 pp; hardbound, ISBN

Syntaxe et sèmantique des quantifiers flottants *tous* et *chacun*: Distributivitè en sèmantique conceptuelle *Marie-Odile Junker* Geneva: Librarie Droz (Langue et cultures 28), 1995, 183 pp; paperbound, ISBN 2-600-00057-7, no price listed

***Applications of neural networks** *Alan F. Murray (editor)* (University of Edinburgh) Dordrecht: Kluwer Academic Publishers, 1995, xii+322 pp; hardbound, ISBN 0-7923-9442-9, US $110.00

***The artificial life route to artificial intelligence: Building embodied, situated agents** *Luc Steels and Rodney Brooks (editors)* (University of Brussels and Massachusetts Institute of Technology) Hillsdale, NJ: Lawrence Erlbaum Associates, 1995, ix+288 pp; hardbound, ISBN 0-8058-1518-X, US $69.95; paperbound, ISBN 0-8058-1519-8, US $29.95

***Introduction to knowledge systems** *Mark Stefik* (Xerox Palo Alto Research Center) San Francisco: Morgan Kaufmann Publishers, 1995, xx+870 pp; hardbound, ISBN 1-55860-166-X, US $69.95

AI '96

Canadian Artificial Intelligence Conference

AI '96 is the eleventh biennial conference on Artificial Intelligence sponsored by the Canadian Society for Computational Studies of Intelligence. It will be held at the **Holiday Inn on King, Toronto, Ontario** during **21-24 May 1996** in conjunction with **Vision Interface** and **Graphics Interface**.

Contributions are invited that present original, unpublished results in all areas of Artificial Intelligence. They should be sent to the program chairman, Gordon McCalla. Papers must be received by: **31 October 1995**.

Submitted papers must not exceed 5,000 words in length, including abstract and bibliography. Theoretical and position papers will be judged on their originality and contribution to the field of AI, and applied papers on the importance and originality of the application. To help in the review process, authors should list, in decreasing order of relevance, 1 to 3 of the following keywords:

applications	learning	reasoning (indicate subarea)
search	cognitive modelling	knowledge representation
planning	architectures	knowledge acquisition
perception	robotics	language understanding
problem solving	theorem proving	neural nets/connectionism
other (please specify)		

Authors should submit four (4) complete copies of the paper in hardcopy form, for review by members of the program committee. Acceptance depends on the overall merit and significance of the reported research, as well as the quality of the written presentation. Each copy of the paper must include a cover page, separate from the body of the paper, which includes, in order, (1) title of the paper, (2) full names, postal addresses, phone numbers, and email addresses of all authors, (3) an abstract of no more than 250 words, and (4) keywords to classify the paper for review purposes. As a condition of acceptance, the author or a co-author must present the paper at the conference. If the paper is being submitted to other conferences, either verbatim or in essence, authors must clearly indicate this on the cover page.

Notification of acceptance or rejection will be mailed to the first author by **31 January 1996**. Camera ready copy of accepted papers is due **28 March 1996**. Each paper will be allotted up to eight (8) pages in the conference proceedings, formatted using 12pt LaTeX or equivalent.

http://ai.iit.nrc.ca/cscsi/conferences/ai96.html

For further information contact the Program Chair:
Gordon McCalla

Program Chairman

Gordon McCalla
Dept of Computer Science
University of Saskatchewan
Saskatoon, SK S7N 5A9
email: mccalla@cs.usask.ca
Phone: 306-966-4902
Fax: 306-966-4884

Program Committee

Marie-Michelle Boulet Université Laval
Renato de Mori McGill University
Mark Fox University of Toronto
Lev Goldfarb University of New Brunswick
Jim Greer University of Saskatchewan
Howard Hamilton University of Regina
Graeme Hirst University of Toronto
Rob Holte University of Ottawa
Dekang Lin University of Manitoba
Jim Little University of British Columbia
Joel Martin National Research Council
Paul McFetridge Simon Fraser University
Bob Mercer University of Western Ontario
Evangelos Milios York University
Peter Patel-Schneider AT&T Bell Labs
Fei Song University of Guelph
Andre Trudel Acadia University
Qiang Yang University of Waterloo

Established in 1983, Applied AI Systems, Inc. (AAI) is the oldest Artificial Intelligence company in Canada. The emphasis of our business is on real world applications of cutting edge intelligent systems technology. AAI's long-term commitment to the "paying respect to science" approach is now reaping benefits. The company is international in approach, with its members constantly traveling the globe to meet other researchers and practitioners, and participating in all major conferences, workshops and symposia in related fields.

Research & Development Areas: ---Artificial Intelligent Robots --- Autonomous Systems --- Interface Agents & other "Intelligent Agents or Soft Robots" --- Evolutionary Systems & Artificial Life --- Neural Networks --- Fuzzy Logic --- Real Time Knowledge-Based Systems (RTKBS) --- Integration of fuzzy logic software with RTKBS --- Speech Recognition --- Service Provided to Clients: --- Custom System Development & Implementation --- Technical Research Reports --- Sales of SA Robots & AI Products --- Training --- Consulting ---

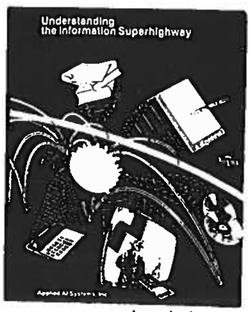

Understanding the Information Superhighway

NEW REPORT from Applied AI Systems, Inc., nearly 400-page book on the Information Superhighway: an up-to date analysis of the complexities of information highways, in addition to a detailed description of the published mergers and deals which are shaping the technology. Knowing how to navigate this highway will be crucial in government, business, education, entertainment, shopping, and banking. The report also provides a comprehensive look at leading-edge information technology, essential for discovering how the information superhighway will shape new products and business opportunities. It contains over 50 illustrations & 11 tables. $995.00 CDN

Intelligent Micro-Robots — IS Robotics — Exploring The Future

for Industrial Robot Prototyping and ALife Experiments. ISR's robots are based on the principle of Subsumption Architecture developed by Prof. Rodney Brooks of the MIT AI Lab. Now available, R3: high performance robot for ALife experiments, T1: highly maneuverable tracked robot platform, T3: T1 based human/animal interaction system, Pebbles: sealed tracked robot for hazardous environments, Genghis-II: dexterous six-legged walking robot, Nexes: advanced six-legged robot, and Hermes: durable six-legged small robot.

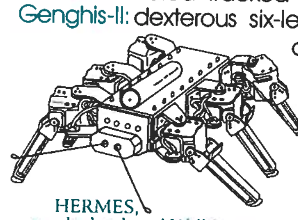

HERMES, recently developed Walking Robot

RESEARCH APPLICATIONS: --- Autonomous Agents --- Co-operative Systems -- Exploration & Surveillance --- Dynamic Environmental Planning --- Find & Fetch Tasks --- Hardware Transport Platforms --- Evolution Experiments Using Genetic Algorithms --- Sensor Transport in Hazardous Environments, & many more.

Inexpensive Miniature Robot !

Khepera is developed by Ecole Polytechnique Federale de Lausanne in Switzerland, mostly for research institutions. Because *Khepera* is small and inexpensive, researchers can conduct experiments with multiple robots to test algorithms for distributed collective intelligence in a small area.

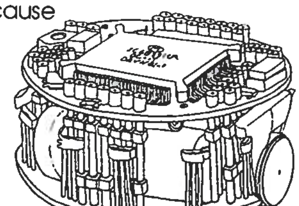

Size: Diameter 55 mm. Height 30 mm. Weight: Appx. 70 g. Processor: Motorola 68331, RAM: 256 Kbyte, ROM: 256 or 512 Kbyte. Possible applications: obstacle avoidance, wall following, target search, embodied learning algorithm, collective behavior (swarm intelligence), various Artificial Life (ALife) experiments.

--------- Small Size ---- Modularity ---- High Computational Power ---- Easy to Control ----- Programmable Highly Intelligent Operations ---- Inexpensive ---------

Khepera® Developed by: EPFL lami Produced by: forelec

For further information on pricing or ordering, please contact Applied AI Machines & Software, Inc.

Applied AI Machines & Software, Inc.
Suite 504, Gateway Business Park,
340 March Road, KANATA
Ontario, Canada K2K 2E4
Phone: + 1 613 592 7729 Fax: + 1 613 592 9762
E Mail: 73051.3521@compuserve.com

The Premier PC Dictation Program

DragonDictate now available for Windows!

Create everything for business from complex documents to spreadsheets. Enter text, make changes, format, and print your work - entirely by voice! You can use it to control Windows functions, accessories, and applications by voice

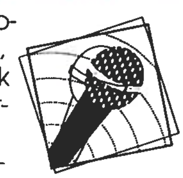

Features include: --- Direct application control and dictation --- True hands-free computing --- Excellent word recognition with minimal training --- Customizable vocabulary lets you use whatever words you want --- 120,000-word backup dictionary with speaker-independent acoustic models for quick recognition --- Dynamic adaptation adjusts to your voice and vocabulary, and environment --- Compatibility with most standard 16-bit multimedia sound cards --- Experience an easy, productive way to work in your Windows applications. Let DragonDictate put you on speaking terms with your PC!

DRAGON Dictate® *for Windows*™

For further information on consultation or research and development, please contact Applied AI Systems, Inc.

Applied AI Systems, Inc.
Suite 500, Gateway Business Park,
340 March Road, KANATA
Ontario, Canada K2K 2E4
Phone: + 1 613 592 3030 Fax: + 1 613 592 2333
E Mail: 71021.2755@compuserve.com