

NAVAL POSTGRADUATE SCHOOL

Monterey, California



VIRTUAL SUPPLY CHAIN RE-INTERMEDIATION THROUGH MULTI-AGENT SYSTEMS

by

Mark E. Nissen

September 2002

Approved for public release; distribution is unlimited.

Prepared for: Office of Naval Research

20021126 102

NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000


RADM David R. Ellison, USN
Superintendent

Richard Elster
Provost

This report was prepared through the Reimbursable-Research Program at the Naval Postgraduate School.

Reproduction of all or part of this report is authorized.

This report was prepared by:



Mark E. Nissen
Associate Professor
Graduate School of Business and
Public Policy

Reviewed by:



Douglas A. Brook
Dean
Graduate School of Business and
Public Policy

Released by:



D. W. Netzer
Associate Provost and
Dean of Research

REPORT DOCUMENTATION PAGE			Form Approved	
OMB No 0704-0188				
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2002		3. REPORT TYPE AND DATES COVERED Technical Report
4. TITLE AND SUBTITLE VIRTUAL SUPPLY CHAIN RE-INTERMEDIATION THROUGH MULTI-AGENT SYSTEMS			5. FUNDING	
6. AUTHOR(S) Mark E. Nissen				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Graduate School of Business and Public Policy Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER NPS-GSBPP-02-002	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words.) Intelligent agent technology offers good potential to integrate supply chain processes and other business-to-business exchanges more closely than current Web technologies and without the rigid inflexibility associated with EDI applications. Viewing respective procurement and order fulfillment processes of buyer and seller as an integrated whole along the supply chain, we identify opportunities for virtual supply chain re-intermediation through multi-agent systems. Using a proof-of-concept agent federation, called the Intelligent Mall, we examine the feasibility and performance implications of this approach for a software supply chain. A long-lived, autonomous agent federation is observed performing delegated supply chain tasks as prescribed and in an intelligent, context-sensitive and policy-conforming manner. The kind of agent-integrated supply chain investigated through this work appears to be inherently scalable yet customizable to the level of a specific individual in the organization. This approach to supply chain re-intermediation also appears to generalize well to a wide variety of corporations, universities and other enterprise forms, and the intelligent mall examined in this report appears to offer good potential for improved process performance. This report summarizes exploratory results from experience with the Intelligent Mall and outlines a number of lessons learned and guidelines for implementing multi-agent systems. The conclusions drawn from this investigation are important for executives and managers considering leading adoption of agent technology, and the research agenda is intended to help guide and stimulate continued work along these lines.				
14. SUBJECT TERMS			15. NUMBER OF PAGES 31	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	

SN 7540-01-280-5800

Standard Form 298 Rev. 2-89)
Prescribed by ANSI Std Z39-18

ABSTRACT

Intelligent agent technology offers good potential to integrate supply chain processes and other business-to-business exchanges more closely than current Web technologies and without the rigid inflexibility associated with EDI applications. Viewing respective procurement and order fulfillment processes of buyer and seller as an integrated whole along the supply chain, we identify opportunities for virtual supply chain re-intermediation through multi-agent systems. Using a proof-of-concept agent federation, called The Intelligent Mall, we examine the feasibility and performance implications of this approach for a software supply chain. A long-lived, autonomous agent federation is observed performing delegated supply chain tasks as prescribed and in an intelligent, context-sensitive and policy-conforming manner. The kind of agent-integrated supply chain investigated through this work appears to be inherently scalable yet customizable to the level of a specific individual in the organization. This approach to supply chain re-intermediation also appears to generalize well to a wide variety of corporations, universities and other enterprise forms, and the intelligent mall examined in this report appears to offer good potential for improved process performance. This report summarizes exploratory results from experience with The Intelligent Mall and outlines a number of lessons learned and guidelines for implementing multi-agent systems. The conclusions drawn from this investigation are important for executives and managers considering leading adoption of agent technology, and the research agenda is intended to help guide and stimulate continued work along these lines.

ABOUT THE AUTHOR

Mark E. Nissen is Associate Professor of Information Systems and Management at the Naval Postgraduate School and Young Investigator for the Office of Naval Research. His research focuses on the investigation of knowledge systems for enabling and managing change in areas such as process innovation, electronic business and knowledge flow. He has been investigating knowledge systems to innovate processes in various domains for over a decade, and he has been developing and experimenting with multi-agent systems for supply chain dis/re-intermediation for several years. His current research focuses on the phenomenology of knowledge flow in the large enterprise. Mark's publications span the information systems and acquisition fields, with recent and forthcoming reports in journals such as *MIS Quarterly*, *Journal of Management Information Systems*, *Decision Support Systems*, *Journal of Information Technology Management*, *Journal of Engineering Valuation and Cost Analysis*, *Acquisition Review Quarterly* and *National Contract Management Journal*. He has also recently published his first book, entitled *Contracting Process Innovation*, and he received the 2000 Menneken Faculty Award for Excellence in Scientific Research, the top research award available to faculty at the Naval Postgraduate School. Before his information systems doctoral work at the University of Southern California, he acquired over a dozen years' management experience in the aerospace and electronics industry and served as a direct-commissioned Supply Officer in the Naval Reserve.

TABLE OF CONTENTS

SUPPLY CHAIN DISINTERMEDIATION.....	1
MULTI-AGENT APPLICATIONS.....	3
AGENT-BASED SUPPLY CHAIN RE-INTERMEDIATION	5
The Supply Chain Process	6
INTELLIGENT AGENT FEDERATION.....	9
Process Design	9
The Intelligent Mall Application.....	10
RESULTS, LESSONS AND GUIDELINES FROM THE INTELLIGENT MALL EXPERIENCE.....	13
Results of The Intelligent Mall Experience	14
Lessons Learned from The Intelligent Mall Experience.....	14
Lesson 1 – agents still require programming.....	14
Lesson 2 – knowledge workers know a lot	15
Lesson 3 – even simple agents can be effective	15
Lesson 4 – complexity compounds	15
Guidelines for Implementing Multi-Agent Enterprise Systems.....	16
Guideline 1 – beware of fad-branding.....	16
Guideline 2 – start with the process	16
Guideline 3 – differentiate and integrate people and machines.....	17
Guideline 4 – don't expect too much from tools	17
Guideline 5 – start simple and experiment incrementally	18
CONCLUSIONS	18
REFERENCES	19
INITIAL DISTRIBUTION LIST	24

LIST OF FIGURES

Figure 1 Agent Framework	5
Figure 2 Integrated Supply Chain Process	6
Figure 3 Intelligent Mall User Interface.....	11
Figure 4 Intelligent Mall Animation.....	12

SUPPLY CHAIN DISINTERMEDIATION

Supply chain management (Gebauer *et al* 1998, Kambil 1997, Lee and Billington 1995, Mabert and Venkataramanan 1998, Monczka *et al* 1998, Nissen 1997, Porter and Millar 1985, Swaminathan *et al* 1998) represents a critical competency in today's global business environment, and a number of effective practices (e.g., just-in-time deliveries, electronic data interchange (EDI), supplier inventory management) are employed to improve the competitiveness and efficiency of enterprises around the world. As a long-practiced form of business-to-business (B2B) electronic commerce, for example, our two decades of experience with EDI suggest commercial processes employed by buyers and sellers must be mutually-compatible in order for business exchanges and transactions to occur effectively, and for rapid purchase and responsive order fulfillment, buyer and seller processes must also be closely integrated (Sokol 1996).

However, not only can EDI be expensive, this technological approach requires a relatively stable and predictable commercial environment to effect supply chain integration. Such conditions are increasingly unlikely in the fast-paced, hypercompetitive business environment of today (see D'Aveni 1994), and modern enterprises now also require flexibility to quickly change vendors and adapt to dynamic economic conditions and diverse commercial practices (see Ciborra 1996, Davidow and Malone 1992). Indeed, the static notion of a fixed supply chain is rapidly giving way to more-dynamic concepts such as *supply webs* and *networks* (see Laseter 1998, Porter 2000, Ross 2000), in which the topology of buyers and sellers changes frequently. Commerce through EDI technology does not satisfy the flexibility requirement well.

To address such limitations and shifts, many firms are replacing EDI with more flexible technology. For instance, moving to Web-based support for commercial transactions (e.g., electronic catalogs, virtual malls and storefronts, online industry exchanges) is now commonplace across numerous industries. Indeed, the Internet offers good potential to "revolutionize procurement" and related commercial activities (Gebauer *et al.* 1998). However, much of the capability for supply chain integration is being lost during the transition from EDI to Web technology. Whereas EDI effectively compels buyers and sellers to integrate their supply chain processes, Web-based supply chain technologies are noticeably one-sided; that is, the latter sites and applications are predominately developed for *either* the buyer or

seller, but not both. Commerce though current Web technologies does not satisfy the integration requirement well.

We also note many calls for disintermediation of the supply chain (e.g., Gellman 1996). The idea—becoming particularly manifest during the first phase of reengineering (see Cypress 1994), with its focus on process streamlining—is to remove the bureaucracy or "middle man" associated with procurement in the enterprise. At present, procurement in most large corporations, universities, government agencies and other major enterprises must pass through a separate, specialized organization devoted solely to the functions of buying products and services and managing suppliers, which many view as "non value added." Thus, disintermediation can be viewed in terms of process redesign focused on eliminating intermediaries (e.g., procurement or procurement departments, market brokers, exchange agents) from the enterprise supply chain.

Building upon work by Barbuceanu and Fox (1993) and others (e.g., Collins et al. 1998, Gini and Boddy 1998, Mehra and Nissen 1998, Rodriguez-Aguilar et al. 1998, Walsh et al. 1998, Wurman et al. 1998), we argue intelligent agents offer excellent potential and capability to complement EDI and current Web technologies and address their shortcomings. Further, through collaborative problem solving enabled by multi-agent systems, we explain how this emerging technology offers potential to substitute federations of intelligent agents for the bureaucratic and expensive intermediaries now employed along most major enterprise supply chains. This is the concept of *virtual supply chain re-intermediation*. First, we disintermediate the supply chain by removing the expensive and bureaucratic "middle man." Then, to make-up for lost knowledge and expertise, we re-intermediate, virtually, with a federation of intelligent agents (i.e., a multi-agent system).

In this report, we investigate opportunities for multi-agent system technology to re-intermediate the enterprise supply chain. By building upon considerable research in this domain, we draw from experience gained through implementation of a multi-agent system in a supply chain environment to articulate how virtual supply chain re-intermediation is technically and operationally feasible in the context of a major, modern enterprise. For additional insight, we describe the demonstrated, proof-of-concept agent federation implemented to re-intermediate the supply chain. And we then generalize from this discussion to

develop insights, lessons and guidelines for both executives and researchers interested in agent technology.

We first provide a high-level overview of multi-agent applications.

MULTI-AGENT APPLICATIONS

In this section, we provide a high-level overview of extant agent applications, with particular emphasis on a framework to relate them to other, more-conventional classes of information technology. Drawing from Nissen (2001), it is informative to group extant agent applications into four classes: 1) information filtering agents, 2) information retrieval agents, 3) advisory agents, and 4) performative agents. Briefly, most information filtering agents are focused on tasks such as filtering user-input preferences for e-mail (e.g., Maes 1994, Malone et al. 1987), network news groups (Sycara and Zeng 1996), frequently asked questions (Whitehead 1994) and arbitrary text (Verity 1997). Information retrieval agents address problems associated with collecting information pertaining to commodities such as compact disks (Krulwich 1996) and computer equipment (uVision 1997), in addition to services such as advertising (PriceWatch 1997) and insurance (Insurance 1997). We also include the ubiquitous Web indexing robots in this class (see Etzioni and Weld 1995, Hsinchun et al. 1998) along with Web-based agents for report writing (Amulet 1997), publishing (InterAp 1995) and assisted browsing (Burke et al. 1997). Agents for technical information delivery (Bradshaw et al. 1997) and information gathering (Knobloch and Ambite 1997) are not Web-based per se, but they perform a similar function.

A third class of agents is oriented toward providing intelligent advice. Examples include recommendations for CDs (Maes 1997), an electronic concierge (Etzioni and Weld 1995), an agent "host" for college campus visits (Zeng and Sycara 1995) and planning support for manufacturing systems (Maturana and Norrie 1997). Agents for strategic planning support (Pinson et al. 1997), software project coordination (Johar 1997) and computer interface assistance (Ball et al. 1997) are also grouped in this class, along with planned support for military reconnaissance (Bui et al. 1996) and financial portfolio management (Sycara et al. 1996). Performative agents in the fourth class are generally oriented toward functions such as business transactions and work performance. Examples include marketplaces (Rayport and Sviokla 1994) for agent-to-agent transactions (e.g., Chavez and Maes 1996), agent auction environments (e.g., Rodriguez-Aguilar et al. 1998) and agent system designs for negotiation (e.g., Bui

1996), in addition to performance of knowledge work such as automated scheduling (Sen 1997, Walsh et al. 1997), cooperative learning (Boy 1997) and automated digital services (Mullen and Wellman 1996).

The intelligent supply chain agents developed through this present research are probably best categorized in the fourth group above (i.e., performative agents). But they have been designed to build upon agent work in other categories as well—for instance exhibiting behaviors such as information filtering and retrieval—and their use can be accomplished through simulation (i.e., in an advisory role) as well as enactment (i.e., the performative role). Thus, intelligent supply chain agents extend and have similarities with examples from each of the four classes above.

To further describe and differentiate intelligent supply chain agents, we again draw from Nissen (2001) and integrate the agent-taxonomy work of Franklin and Graesser (1996) with a three-dimensional, agent-capability structure from Gilbert et al. (1995) to develop the analytical framework presented in Figure 1. In this framework, we use the same intelligence and mobility dimensions developed by Gilbert but substitute the new dimension *collaboration* in lieu of autonomy/agency. This follows the presumption of agent autonomy stressed by Franklin and Graesser. For purpose of discussion, we have annotated this three-dimensional space with one, relatively "pure" exemplar from each agent-capability dimension. For example, many expert system applications are quite extensive in terms of formalized, expert-level intelligence, but they are not traditionally designed to operate on foreign hosts, nor do they generally collaborate with other expert systems to jointly solve problems. Similarly, remote programming of the sort enabled by Java, Telescript and Odyssey equip programs to execute on foreign machines, but these procedural applications are not generally endowed with the capability for intelligent inference, nor are they usually thought of in terms of collaborative processing. Likewise, parallel processing has an explicit focus on collaborative problem solving between multiple, parallel processors, but this problem solving is usually focused more on procedural processing than intelligent reasoning, and execution on foreign hosts is rarely envisioned. Clearly exceptions exist for each class (e.g., distributed AI (see Bond and Gasser 1988), intelligent Java agents (see Neuenhofen and Thompson 1998), others), but these three exemplars should convey the basic concepts associated with each dimension.




Figure 1 Agent Framework

(Adapted from Nissen 2001)

Notice the annotation for intelligent supply chain agents (labeled "ISCA" in the figure). It occupies a position notionally in the middle of this three-dimensional agent-capability space. Although the capabilities of intelligent supply chain agents are not as extensive as any of the three exemplars from above along any *particular* dimension, none of the exemplars from above combine even two of these three capability dimensions associated with intelligent agents. This adds to the challenge of agent development work—particularly where intelligent problem solving must be coordinated among a federation of autonomous agents—but it serves to enable a new set of powerful capabilities (esp. collaborative problem solving) that prove to be quite effective and useful for providing both inter-organizational integration and flexibility to complex processes such as supply chain management.

AGENT-BASED SUPPLY CHAIN RE-INTERMEDIATION

In this section, we discuss agent-based supply chain re-intermediation. This present work builds upon and extends research by Mehra and Nissen (1998) to implement and demonstrate an agent development environment and by Nissen and Mehra (1998) that explores process redesign enabled by intelligent-agent technology. We first discuss the supply chain as a unifying process that combines the procurement activities of a major buyer with the order fulfillment activities of a commercial vendor. We then examine agent-based supply chain re-intermediation for a specific, operational supply chain process.

The Supply Chain Process

Two primary processes are involved with the supply chain: 1) customer purchasing and 2) vendor order fulfillment. Other processes such as customer financing, vendor project management and logistics can also be viewed as supply chain activities, but these are subsumed by our collective use of the terms *procurement* and *order fulfillment*. Customer procurement and vendor order fulfillment are traditionally viewed as separate, intra-organizational process activities. But a strong case can be made for viewing such activities *together*, as an integrated, inter-organizational supply chain process. This reflects the integrated focus of emerging procurement models, such as the one used by Gebauer et al. (1998) to discuss the revolutionary potential of Internet- and Web-based procurement.




Figure 2 Integrated Supply Chain Process

We examine and discuss one specific supply chain instance and use it for context to implement and investigate an agent federation in the section that follows. On the buyer side, the purchasing process pertains to work done by the procurement department at a university on the West Coast. As a large institution, this enterprise is subject to the full complement of purchasing policies and rules that govern the procurement activities of most major enterprises (e.g., corporations, government agencies, universities). On the seller side, the order fulfillment process pertains to work done by the product and licensing department of a leading software development company on the East Coast. This company is a leader in software for its product market and maintains an active research and development activity that drives frequent software

introductions, updates and releases. Therefore, it maintains the kind of rapid product evolution that has been problematic for procurement in large enterprises such as the university. The high-level process delineated in Figure 2 depicts the integration of the user, university procurement department and commercial software contractor. This process is described in greater detail by Nissen and Mehra (1998) with respect to agent-enabled redesign opportunities.

The process begins with a user in the organization identifying a need and determining his or her preliminary software requirements. A market survey follows with the market information (e.g., products, capabilities, companies, prices) used to complete a (paper-based) procurement request (PR) form. This form is submitted to the procurement department for processing, in which a buyer verifies the form (e.g., in terms of completeness, required documentation such as sole-source justification, adequate budget) and then researches some potential sources for procurement (e.g., existing contracts, approved-vendor lists, small/disadvantaged-business lists) in addition to the sources identified through the market survey. Requests for quotation (RFQs) are generally issued next, and vendor-prepared quotations are analyzed by the buyer, whom then summarizes the information for review and source selection by the user. A purchase order is subsequently issued, and the transaction is complete when software is delivered to the user and payment is made to, and deposited by, the vendor.

Not apparent from the figure is the underlying knowledge, expertise and information required for people to perform the purchasing and order fulfillment processes depicted. For example, the user must know how to conduct a market survey and have access to information listing alternative sources of supply, as well as an understanding of the basic procedures for university procurement and information pertaining to the specific PR form used. We note most major corporations, universities and other large enterprises have no shortage of like procurement procedures and forms. Similarly, the buyer must possess thorough knowledge of the applicable procurement policies and rules, know how to review the purchase request (e.g., what constitutes completeness, when to request additional information) and have current information pertaining to alternative sources of supply—sources which often appear, combine, shift direction and disappear frequently in software markets. This, too, reflects much corporate and university procurement, as laws and regulations for inter-state and international commerce can be equally complex, and product markets may be even more dynamic and unpredictable. The buyer must also have access to one or more

suppliers' systems to be able to post the RFQ (e.g., by telephone, fax, e-mail, EDI, electronic bulletin board) and requires knowledge of the procedures required for quotation analysis and source selection. Access to and understanding of the receiving and payment systems and procedures is also necessary to complete the transaction. And of course, vendor personnel must understand the policies, procedures and systems associated with software quotation, order fulfillment and billing activities. These kinds of knowledge, expertise and information can be used to specify and enable autonomous and flexible behaviors of an intelligent agent federation, but we find some process activities lend themselves to agent performance more than others.

As guidelines for process analysis and design, two classes of supply chain activities appear to offer good potential for agent performance: 1) commercial exchange activities and 2) procedural knowledge activities (see Hudson 1998). Commercial exchange activities necessarily cross organizational boundaries and are enabled by network technologies (e.g., local area networks, wide area networks, standard network protocols). Provided the element of exchange (e.g., information, influence, payment) can be delivered via network media (i.e., electronically), agents—as well as EDI and Web technologies—can be employed to perform most exchange activities. Unlike other electronic approaches, however, agent autonomy enables greater process flexibility and faster adaptation to change than EDI, and agents performing activities of *both* buyer and seller processes enable closer supply chain integration than made possible through most extant Web technologies.

Procedural knowledge is also important for agent-based process integration, for such knowledge is often readily accessible and sufficiently formalized in the organization to be used directly for specification of agent behaviors. Just like a competent person (say in the context of a new employee trainee, or an experienced knowledge worker performing occasional or unusual duties) will review procedures to guide work performance and decision making, agents can similarly draw from procedural knowledge to determine appropriate behaviors and guide decisions. In contrast, other processes (e.g., not involving commercial exchange or procedural knowledge) that require common sense, tacit knowledge and human social skills are best performed by people as opposed to machines at present. With this, we focus attention on supply chain activities associated with commercial exchanges and involving procedural knowledge application as candidates for possible performance through an intelligent agent federation.

INTELLIGENT AGENT FEDERATION

In this section, we describe an agent-based supply chain implementation. Drawing in part from Nissen (2001), we briefly describe the process design and then outline the feasibility and performance implications of this agent-based process design.

Process Design

Design begins with the process itself, reflecting our emphasis on the organization as well as agent technology. Recall the integrated software supply chain process from above. Of the sixteen activities delineated in Figure 2, we employ the guidelines from above to designate eight of them (highlighted in bold in the figure) for performance by intelligent agents. For reference, these include completing the PR form, procurement department verification, researching alternative supply sources, issuing RFQs, quotation preparation by vendors, quotation analysis by the customer, and both issuing and filling the ensuing order by the respective supply and vendor organizations. Although some commercial exchange activities such as invoicing, payment and funds deposit also clearly lend themselves to electronic intermediation, our greater interest lies in the activities not currently performed through such traditional technologies as EDI. Indeed, EDI support of such activities is now so routine and well understood, we no longer even consider it to represent "research."

Alternatively, many of the other activities in the figure not selected for agent performance are currently too difficult for agent-based performance, given the current state of agent-design knowledge and -development technology. Hence these are retained as human process tasks. For example, developing an agent to reliably anticipate and identify a user's software needs would be very difficult to accomplish today. Notice this activity involves neither commercial exchange nor procedural knowledge (i.e., as outlined in our guidelines above). Indeed, it is likely only the user herself could specify the knowledge and behaviors required for such an agent, and only if the knowledge used to identify such software needs could be made explicit. Thus, we see this process design calls for human knowledge workers and machine agents to share responsibilities for process performance. Rationale for assignment of the other supply chain process activities (i.e., delegated to agents or retained by people) follows similar reasoning.

The Intelligent Mall Application

As the name implies, The Intelligent Mall employs a shopping mall metaphor for supply chain re-intermediation. Notice we change the *supply chain* metaphor—with its connotation of knowing all trading partners *ex-ante*—and introduce the *mall* concept instead. In the mall, shoppers are not expected to know in advance which shops exist or what products they offer for sale. Neither are the shops expected to know which other shops are selling like products or with which shoppers they will interact. In other words, the supply chain represents a special case of the mall in which trading partners are established before initiating procurements.

The mall metaphor makes for a more robust commercial environment than the supply chain model and requires more intelligent agent behaviors. The mall representation presented in Figure 3 shows three (virtual) shops in the upper-left window. This mall is truly virtual, in that the "shops" do not reside in any single physical location. Indeed, the shops do not necessarily exist physically at all. What does exist is a shop agent created to represent some vendor interested in commercial participation through this medium. We refer to this application as an "intelligent mall"—as opposed to the more common "virtual mall" name—because it offers more than just virtual shopping; that is, every entity in this environment possesses (artificial) intelligence, and the agent federation is performative, offering the capability of autonomously shopping on behalf of customers as well as selling on behalf of vendors.

To "open" a shop, a vendor needs network access to the agent server (we provide a Web interface for this) and a Web browser or a Java virtual machine. Any practical number of shops can co-exist in this mall. Thus, the application offers good potential to scale well to large enterprise supply chains. Not apparent in the figure is a special agent called "Host." Shop agents register their products and services (e.g., catalog items, prices, availability, terms and conditions), and the Host maintains both a "White Pages" and a "Yellow Pages" directory for the mall.




Figure 3 Intelligent Mall User Interface

Requirements for creating a shopper agent are similar to those pertaining to the shops above. Like their shop counterparts, shopper agents represent individuals wishing to participate in the mall and register with the Host. A shopper agent can be specialized to reflect the preferences and priorities of its principal. For example, one shopper agent can be specialized to shop based on my preferences for juice over coffee and bagels over doughnuts, whereas another instance (i.e., identical clone) can be specialized to reflect someone else's preferences. And, clearly, each shopping agent can be instantiated with a unique shopping list of items to buy. Other knowledge and information—such as user preferences and budget restrictions, product requirements and need dates, and consumer heuristics like price comparison—are formalized through rules for the agents.

Figure 3 also depicts input screens for one shop agent and a shopper. The keeper of Shop-2 uses the form on the upper-right to update items, prices and other information normally expected in a product catalog. Here juice is added to the catalog of items listed for sale in this particular shop. The shopper principal similarly specifies items to be purchased using the form on the bottom-right. This user can either type-in items or select from a list of products registered with the Host (see form on bottom-left). Again,

any practical number of different shopper agents can be created to shop for any practical number and quantity of various items in The Intelligent Mall.




Figure 4 Intelligent Mall Animation

Figure 4 presents a still-shot of some animated shopping activities along with a Java interface for Shop-3. When a user has specified the items to be purchased as above, each shopper agent sends messages to all shops carrying each item in The Intelligent Mall to see which vendors wish to sell the items on its shopping list. This message action is equivalent to advertising the intent to procure items and posting RFQs. Shop agents wishing to sell the posted items will reply directly to the requesting agents and respond to the corresponding "RFQ" messages separately. Each shopper agent analyzes the quotes and determines a preferred source for every item on the shopping list, based on user preferences and vendor information received.

Each shopper agent then performs some planning to determine which shops to visit, establish the order in which to visit them, and consolidate multiple items purchased from each shop into a single trip. Notice the shopper agent in Shop-2 is inverted. Such behavior is explicitly required by "regulation" in this

mall environment when juice is purchased, for instance. Although humorous, perhaps, this behavior serves to visibly demonstrate the intelligent agents can be specified to conform to various procurement regulations (e.g., governing pharmaceutical imports, military exports, interstate intellectual property rights). Although difficult to see without animation, the other shopper agent is currently "jumping" up and down, which represents the local custom when conducting business in Shop-1. Similar to the regulation-conforming behavior from above, this context-specific behavior serves to demonstrate the intelligent agents can also be made sensitive to local customs and variations (e.g., contracting for software programming service in India as opposed to Taiwan). The number and kinds of such behaviors that can be specified through intelligent agents are practically unlimited.

When all items have been purchased from a given shop, the shopper agent proceeds to the "exit," where messages are exchanged with the "cashier" to make payment and issue a receipt. Each shopper then moves to the next shop on its route, purchases the corresponding items, and continues in this fashion until all items have been purchased or no vendors are available to sell items still on the list. This latter condition may often obtain in The Intelligent Mall. Although an item may be available when the agent *plans* its shopping route, there is no guarantee the item will not subsequently be purchased by another shopping agent, or even that a shop agent somewhere in the world will not "close" before the shopper reaches its network destination. Moreover, we can instruct an agent to be persistent and *not return* until all items on the shopping list have been purchased. The dynamic and emergent behaviors associated with this distributed application—particularly where intelligent, autonomous, persistent agents are free to roam the network and conduct their business as they see fit—are very rich and can be made quite realistic.

RESULTS, LESSONS AND GUIDELINES FROM THE INTELLIGENT MALL EXPERIENCE

In this section, we summarize and generalize from experience with The Intelligent Mall. We outline key results and lessons learned from the study and then develop high-level guidelines for the manager interested in pursuing virtual supply chain re-intermediation through multi-agent systems.

Results of The Intelligent Mall Experience

Results thus far from our experience with The Intelligent Mall focus principally on feasibility. Having constructed The Intelligent Mall and observed its operation over a period of several months, one can see the multi-agent system approach to virtual supply chain re-intermediation is technically feasible. Indeed, the mall application has been allowed to run continuously for several days without problems of system instability, agent indecisiveness or other maladies that can affect such distributed, intelligent computing applications (see Wooldridge 1998). Although our mall agents remain relatively simple (individually) at this early research stage, they serve to demonstrate agents can represent diverse users, autonomously find and collaborate with other agents to purchase goods, and effectively conform to regulations and customs imposed by various supply chain contexts and trading situations. Technically, we anticipate straightforward extension to incorporate more complex agent behaviors.

Operationally, The Intelligent Mall also shows how an agent federation can be specified and implemented to conform to complex procurement policies and regulations. Although we implement and study only a small subset of knowledge and behaviors required for authorized university procurement, the agents demonstrate they can follow rules specific to enterprise software purchases, broadly advertise procurement intentions, competitively solicit multiple quotations without bias, perform product/price evaluations and like requirements for B2B exchange. And recall the animated agents inverting themselves and jumping up and down. At the level of agent objects and methods, nearly any specific, definable behavior can be implemented for conformance to idiosyncratic enterprise policies, procedures and practices.

Lessons Learned from The Intelligent Mall Experience

Although still at an early stage of investigation, we have already learned a number of lessons from the intelligent mall experience. These lessons become particularly apparent as we implement more sophisticated agent behaviors and complex decision processes and as we continue moving out from the laboratory.

Lesson 1 – agents still require programming

Despite the very high level of abstraction associated with intelligent agents, ultimately every software agent requires executable computer code to function. And notwithstanding the facility of agent

specification made possible through graphical development environments, online debugging facilities and other productivity enhancers, agent decisions and behaviors are implemented through objects, rules, methods and messages. Many managers appear to naively think agent development somehow obviates them from the details and complexities of object-oriented programming.

Lesson 2 – knowledge workers know a lot

The power and potential of intelligent agents become manifest as they are employed to perform and assist with knowledge-work activities in the enterprise. But even relatively low-level and straightforward activities performed by workers in the organization can require considerable skill and effort to capture and formalize the corresponding knowledge required for effective performance by autonomous agents. As we learned through our past three decades of experience with expert systems, not every knowledge task is equally suited to automation, and practical results require focus on a relatively narrow task domain. At this stage of agent technological maturity, designers need to be satisfied with relatively simple behaviors (i.e., low levels of knowledge) for any particular agents and let people address the hard enterprise supply chain problems.

Lesson 3 – even simple agents can be effective

No single agent needs to be an "Einstein" in order to be effective, and intelligent, complex agents are difficult to understand, predict and control (just like most people). The problem-solving power and enterprise potential of intelligent agents derive from the ability of agent *federations*—even comprised of relatively simple agents—to collaboratively address knowledge work. We learned to limit each agent in size and scope to no more than required to solve specific problems, yet to equip agents with a robust architecture and capability for finding and communicating with one another. This lesson may also have knowledge-management implications outside the artificial environment of intelligent agents (e.g., in organizations of people; see Davenport and Prusak 1998).

Lesson 4 – complexity compounds

When compared with other software applications of similar size, even a relatively simple intelligent agent will appear more complex by the very nature of its mobility, communication patterns and attempts to emulate human expertise. Multiple, intelligent agents working collaboratively to solve problems are an

order of magnitude more complex than a single agent, and anticipating their collective and emergent behaviors becomes much more difficult. When multiple instances of multiple intelligent agents interact in the same environment—some collaborating, others specialized to compete—the complexity compounds again. This compounding continues as we then distribute multiple agents on different processors across the network, have them react to dynamic changes in the environment (e.g., new shops entering the mall) and expect the machine agents to share process activities with human knowledge workers.

Guidelines for Implementing Multi-Agent Enterprise Systems

Based in part on the results to date and lessons above, we offer the following five guidelines for developing multi-agent systems in the enterprise.

Guideline 1 – beware of fad-branding

Any current discussion of intelligent agents runs the risk of misperception. This risk is particularly high when one mixes-in characterizations such as "radical process redesign," "artificial intelligence," "B2B" and "the Web." Hearing such "buzzwords," many enterprise executives will brand the agent advocate as simply following the latest fad and quickly dismiss his or her concepts and ideas as speculative or vacuous. Recall the hyperbole and false expectations surrounding AI and expert systems in the early Eighties and the reengineering phenomenon of the early Nineties, not to mention current attention now paid to knowledge management and the fact that agents employ object-oriented techniques and operate via Web in a client/server environment. Executives and others have substantial reason to be skeptical. The first guideline is to be prepared for such skepticism and have an animated agent demonstration ready to show people you're serious.

Guideline 2 – start with the process

The multi-agent systems technology used to enable The Intelligent Mall application is just that: technology. If we learned anything from our reengineering experience last decade, it is that process innovation requires more than just technology (Hammer 1990). Rather, effective redesign requires integration of process analysis, organizational design and human resource enablers of change (Davenport 1993), preferably employed in combination (Stoddard and Jarvenpaa 1995), in addition to technology. We find the same holds true with intelligent agents. Analysis and design should begin with the process to be re-

intermediated, integrated or otherwise redesigned through agent technology. Whatever the focus of agent technology, managers should first identify the associated process and make a business case (e.g., return on investment, faster time to market) for use of the technology. In this respect, Guideline 2 reflects the kind of standard recommendations found in introductory information systems (IS) textbooks that address implementation of any new technology. Again, intelligent agents are just that: new technology.

Guideline 3 – differentiate and integrate people and machines

We employ guidelines gleaned from current agent research to help identify supply chain activities that would lend themselves to performance by intelligent agents over people, and vice versa. Because the emerging technology associated with multi-agent systems remains relatively primitive at this stage, it is probably unrealistic to expect an agent federation to perform *all activities* associated with an enterprise process of any size or significance. Recall we delegate only eight of sixteen software supply chain activities to agents in the intelligent mall application, retaining the balance for people to perform. Thus, people and machine agents will necessarily have to work together and share process responsibilities. It is important not to delegate to machines process activities better performed by people.

Guideline 4 – don't expect too much from tools

The current state of the art in agent development tools and technology to support multi-agent systems is nowhere near that available to support other areas of information systems (e.g., object-oriented analysis and design, decision support systems, enterprise resource planning, computer-supported cooperative work). Although a limited number of agent-development tools (e.g., Agent Builder (1998), Agent Development Environment (1998), Via (1998) and others) are beginning to emerge from beta testing, their functionalities remain quite limited. And many research tools such as JAFMAS (Chauhan and Baker 1998) and RETSINA (Sycara and Pannu 1998) are probably only usable by computer science Ph.D.s. Notwithstanding the considerable benefits of using such "shell like" tools in terms of abstraction, architecture, prototyping and debugging—certainly with respect to starting with a blank Java or C++ edit screen—implementation of a multi-agent system like The Intelligent Mall may require considerable effort, over many months or even years, from a doctoral-level researcher or experienced agent-tools consultant.

Guideline 5 – start simple and experiment incrementally

Intelligent agents represent a relatively unfamiliar abstraction in IS design, and the emergent behaviors of cooperating agents in a federation are difficult to anticipate, and even interpret, in many cases. One approach to managing this situation is to begin with a small application and experiment with it until the corresponding behaviors are understood. Then add onto this simple application incrementally, continuing to experiment along the way. This approach may be at odds with first-mover advantages and executive pressure to "get the agents online quickly," but when you think about the implications of having autonomous, intelligent, persistent software entities roaming the marketspace and transacting business on behalf of the enterprise, this call for caution may warrant attention.

CONCLUSIONS

Intelligent agent technology offers good potential to integrate supply chain processes more closely than current Web technologies and without the rigid inflexibility associated with EDI applications. Viewing the respective procurement and order fulfillment processes of buyer and seller as an integrated whole along the supply chain, we identify opportunities for virtual supply chain re-intermediation through multi-agent systems and develop a proof-of-concept agent federation to examine the feasibility and performance implications of this approach. Launching and operating long-lived, autonomous, intelligent agents at several locations across both local and wide area networks, and observing the agents performing their delegated supply chain tasks as prescribed and in an intelligent, context-sensitive and policy-conforming manner, one can see how agent-based supply chain integration is technically and operationally feasible.

As agent-development tools and technologies continue to emerge and take form, the kind of agent-integrated supply chain investigated through this research appears to be inherently scalable yet customizable to the level of a specific individual in the organization. An intelligent agent approach to supply chain re-intermediation also appears to generalize well beyond the specific university procurement process studied here, across to a wide variety of corporations, government agencies, universities and other enterprise forms. Further, the positive results of this investigation should provide considerable first-mover advantages to the enterprises able to harness this new, powerful technology. Through the new knowledge and developmental guidelines uncovered through implementation and examination of the Intelligent Mall,

executives interested in attaining such advantages are now a big step closer to having these interests become manifest.

REFERENCES

- Agent Builder. Developed by Reticular Systems. World Wide Web address: <http://www.reticular.com> (1998).
- Agent Development Environment. Developed by Gensym. World Wide Web address: <http://www.gensym.com> (1998).
- Amulet. Amulet online description. Internet address: <http://www.amulet.com> (1997).
- Ball, G., Ling, D., Kurlander, D., Miller, J., Pugh, D., Skelly, T., Stankosky, A., Thiel, D., Van Dantzich, M. and Wax T. "Lifelike Computer Characters: The Persona Project at Microsoft," in J. Bradshaw (Ed.). *Software Agents*. AAAI Press: Menlo Park, CA (1997).
- Barbuceanu, M. and Fox, M.S., "The Design of COOL: A Language for Representing Cooperation Knowledge in Multi-Agent Systems," WWW: <http://www.ie.utoronto.ca/EIL/ABS-page/ABS-intro.html> (1993).
- Bond, A.H. and Gasser, L. (Eds.) *Readings in Distributed Artificial Intelligence*. Morgan-Kaufman: San Mateo, CA (1988).
- Boy, G.A. "Software Agents for Cooperative Learning," in J. Bradshaw (Ed.). *Software Agents*. AAAI Press: Menlo Park, CA (1997).
- Bradshaw, J.M., Dutfield, S. Benoit, P. and Woolley, J.D. "KAoS: Toward an Industrial-Strength Open Agent Architecture," in J. Bradshaw (Ed.), *Software Agents*. AAAI Press: Menlo Park, CA (1997).
- Bui, T. "Intelligent Negotiation Agents for Supporting Internet-based Competitive Procurement," working paper (1996).
- Bui, T., Jones, C., Sridar, S. and Ludlow, N. "Decision Support for Reconnaissance Using Intelligent Software Agents," Naval Postgraduate School research proposal (1996).
- Burke, R.D., Hammond, K.J. and Young, B.C. "The FindMe Approach to Assisted Browsing," *IEEE Expert* (July/August 1997), pp. 32-40.
- Chauhan, D. and Baker, A.D. "JAFMAS: A Multiagent Application Development System," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents* (1998), pp. 100-107.
- Chavez, A. and Maes, P. "Kasbah: An Agent Marketplace for Buying and Selling Goods," *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology* London, UK (April 1996).
- Ciborra, C. (1996), "The Platform Organization: Recombining Strategies, Structures, and Surprises" *Organization Science*, Vol. 7, No. 2, March-April, pp. 103-117.
- Collins, J., Youngdahl, B., Jamison, Sc., Mobasher, B. and Gini, M., "A Market Architecture for Multi-Agent Contracting," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN (1998), pp. 285-292.

- Cypress, H.L. "Reengineering – MS/OR Imperative: Make Second Generation of Business Process Improvement Mode Work," *OR/MS Today* 21:1 (1994), pp. 18-29.
- D'Aveni, R.A. *Hypercompetition: Managing the Dynamics of Strategic Maneuvering* New York, NY: Free Press (1994).
- Davenport, T.H. *Process Innovation: Reengineering Work through Information Technology* Boston, MA: Harvard Press (1993).
- Davenport, T.H. and Prusak, L. *Working Knowledge: How Organizations Manage what they Know*. Harvard Business School Press: Boston, MA (1998).
- Davidow, W.H. and Malone, M.S. *The Virtual Corporation* Harper Business Press: New York, NY (1992).
- Etzioni, O. and Weld, D.S. "Intelligent Agents on the Internet: Fact, Fiction, and Forecast," *IEEE Expert* (August 1995), pp. 44-49.
- Franklin, S. and Graesser, A. "Is It an Agent or Just a Program? A Taxonomy for Autonomous Agents," in *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages* Springer-Verlag: New York, NY (1996).
- Gebauer, J., Beam, C. and Segev, A. "Impact of the Internet on Procurement," in: M. Nissen, K. Snider and D. Lamm (Eds.), *Acquisition Review Quarterly* Special Issue on Managing Radical Change 5:2 (1998), pp. 167-184.
- Gellman, R., "Disintermediation and the Internet", *Government Information Quarterly* 13:1 (1996), pp. 1-8.
- Gilbert, D., Aparicio, M., Atkinson, B., Brady, S., Ciccario, J., Grosof, B., O'Connor, P., Osisek, D., Pritko, S., Spagna, R., and Wilson, L. "IBM Intelligent Agent Strategy," working paper, IBM Corporation (1995).
- Gini, M. and Boddy, M., Workshop on Agent-Based Manufacturing," conducted at the Autonomous Agents '98 Conference, Minneapolis, MN (1998).
- Hammer, M. "Reengineering Work: Don't Automate, Obliterate" *Harvard Business Review* 68:4 (July-August 1990), pp. 104-112.
- Hsinchun, C., Yi-Ming, C., Ramsey, M. and Yang, C.C. "An Intelligent Personal Spider (Agent) for Dynamic Internet/Intranet Searching," *Decision Support Systems* 23 (1998), pp. 41-58.
- Insurance. Insurance online description. Internet address: <http://www.dmmatters.co.uk> (1997).
- Hudson, J. *Software Agents and the Defense Information Infrastructure: Reengineering the Acquisition Process* Master's Thesis, Systems Management Department, Naval Postgraduate School (1998).
- Insurance. Insurance online description. Internet address: <http://www.dmmatters.co.uk> (1997).
- InterAp. "InterAp Assigns Intelligent Agents to the Web," *PCWeek* (12 June 1995).
- Johar, H.V. "SoftCord: an Intelligent Agent for Coordination in Software Development Projects," *Decision Support Systems* 20 (1997), pp. 65-81.
- Kambil, A., "Doing Business in the Wired World", *Computer* 30:5 (May 1997), 56-61.

- Knobloch, C.A. and Ambite, J.L. "Agents for Information Gathering," in J. Bradshaw (Ed.), *Software Agents*. AAAI Press: Menlo Park, CA (1997).
- Krulwich, D. *An Agent of Change*. Andersen Consulting Center for Strategic Technology Research (1996).
- Laseter, T.M., "Six Enduring Principles for Integrating the Supply Web", *Purchasing* 125:5 (8 October 1998), pp. 26-31.
- Lee, H.L. and Billington, C., "The Evolution of Supply Chain management Models and Practice at Hewlett-Packard", *Interfaces* 25:5 (1995), pp. 42-63.
- Mabert, V.A. and Venkataramanan, M.A., "Special Research Focus on Supply Chain Linkages: Challenges for Design and Management in the 21st Century", *Decision Sciences* 29:3 (Summer 1998), pp. 537-552.
- Maes, P. "Agents that Reduce Work and Information Overload," *Communications of the ACM* 37:7 (July 1994), pp. 30-40.
- Maes, P. "Pattie Maes on Software Agents: Humanizing the Global Computer," *Internet Computing* (July-August 1997).
- Malone, T.W., Grant, K.R., Turbak, F.A., Brobst, S.A and Cohen, M.D. "Intelligent Information-Sharing Systems," *Communications of the ACM* 30:5 (1987), pp. 390-402.
- Maturana, F.P. and Norrie, D.H. "Distributed Decision-making Using the Contract Net Within a Mediator Architecture," *Decision Support Systems* 20 (1997), pp. 53-64.
- Mehra, A. and Nissen, M.E. "Case Study: Intelligent Software Supply Chain Agents Using ADE," *Proceedings from the AAAI Workshop on Software Tools for Developing Agents* (1998).
- Monczka, R.M. Petersen, K.J., Handfield, R.B. and Ragatz, G.L., "Success Factors in Strategic Supplier Alliances: The Buying Company Perspective", *Decision Sciences* 29:3 (Summer 1998), pp. 553-577.
- Mullen, T. and Wellman, M.P. "Market-based negotiation for digital library services," Second USENIX Workshop on Electronic Commerce (November 1996).
- Neuenhofen, K. and Thompson, M. "A Secure Marketplace for Mobile Java Agents," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents* (1998), pp. 212-218.
- Nissen, M.E. "The Commerce Model for Electronic Redesign," *Journal of Internet Purchasing WWW*: <http://www.arraydev.com/commerce/JIP/9702-01.htm>, (July 1997).
- Nissen, M.E. and Mehra, A. "Redesigning Software Procurement through Intelligent Agents," *Proceedings from the AAAI Workshop on AI in Reengineering and Knowledge Management* (1998).
- Nissen, M., "Agent-Based Supply Chain Integration," *Journal of Information Technology Management* 2:3 Special Issue, Electronic Commerce in Procurement and the Supply Chain (2001), pp. 289-312.
- Pinson, S., Louca, J.A. and Moraitis, P. "A Distributed Decision Support System for Strategic Planning," *Decision Support Systems* 20 (1997), pp. 35-51.
- Porter, A.M., "The Virtual Corporation: Where Is It?" *Purchasing* 128:4 (23 March 2000), pp. 40-48.
- Porter, M. and Millar, V. "How Information Gives You Competitive Advantage," *Harvard Business Review* (July/August 1985), pp. 149-160.

PriceWatch. PriceWatch online description. Internet address: <http://www.pricewatch.com> (1997).

Rayport, R.E. and Sviokla, J.J. "Managing in the Marketplace", *Harvard Business Review* (November-December. 1994).

Rodriguez-Aguilar, J.A., Martin, F.J., Noriega, P., Garcia, P. and Sierra, C., "Competitive Scenarios for Heterogeneous Trading Agents," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN (1998), pp. 293-300.

Ross, A.D. "A Two-Phased Approach to the Supply Network Reconfiguration Problem," *European Journal of Operational Research* 122:1 (1 April 2000), pp. 18-30.

Sen, S. "Developing an Automated Distributed Meeting Scheduler," *IEEE Expert* (July/August 1997), pp. 41-45.

Sokol, P. *From EDI to Electronic Commerce: A Business Initiative* McGraw-Hill: New York, NY (1996).

Stoddard, D.B. and Jarvenpaa, S.L. "Business Process Redesign: Tactics for Managing Radical Change," *Journal of Information Management Systems* 12:1 (1995), pp. 81-107.

Swaminathan, J.M., Smith, S.F. and Sadeh, N.M., "Modeling Supply Chain Dynamics: A Multiagent Approach", *Decision Sciences* 29:3 (Summer 1998), pp. 607-632.

Sycara, K., Pannu, A., Williamson, M. and Zeng, D. "Distributed Intelligent Agents," *IEEE Expert* (December 1996), pp. 36-46.

Sycara, K. and Pannu, A. "The RETSINA Multiagent System: Towards Integrating Planning, Execution and Information Gathering," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents* (1998), video presentation and summary pp. 350-351.

Sycara, K. and Zeng, D. "Coordination of Multiple Intelligent Software Agents," to appear in *International Journal of Cooperative Information Systems* (1996).

uVision. Internet address: <http://www.uvision.com/> (1998).

Verity. Verity online description. Internet address: <http://www.verity.com> (1997).

Via. Developed by Kinetoscope. World Wide Web address: <http://www.kinetoscope.com> (1998).

Walsh, W.E., Wellman, M.P., Wurman, P.R. and MacKie-Mason, J.K. "Some Economics of Market-based Distributed Scheduling, Submitted for publication (1997).

Whitehead, S.D. "Auto-faq: An Experiment in Cyberspace Leveraging," *Proceedings of the Second International WWW Conference* 1 (1994), pp. 25-38.

Wooldridge, M., "Pitfalls of Agent-Oriented Development," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN (1998), pp. 385-391.

Wurman, P.R., Wellman, M.P. and Walsh, W.E., "The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents," in: K. Sycara and M. Wooldridge (Eds.), *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN (1998), pp. 301-308.

Zeng, D. and Sycara, K. "Cooperative Intelligent Software Agents," Carnegie Mellon University technical report no. CMU-RI-TR-95-14 (March 1995).

INITIAL DISTRIBUTION LIST

<u>Agency</u>	<u>No. of Copies</u>
Defense Technical Information Center 8725 John J. Kingman Rd., STE 0944 Ft Belvoir, VA 22060-6218	2
Dudley Knox Library, Code 013 Naval Postgraduate School Monterey, CA 93943	2
Research Office, Code 09 Naval Postgraduate School Monterey, CA 93943	1
Shu Liao Code GB/Lc Naval Postgraduate School Monterey, CA 93943	1
Mark E. Nissen Code GB/Ni Naval Postgraduate School Monterey, CA 93943	2