Research in Enterprise Integration

Mark S. Fox

Department of Mechanical and Industrial Engineering, University of Toronto
4 Taddle Creek Road, Toronto, Ontario M5S 3G9 CANADA
tel: +1-416-978-6823 fax: +1-416-971-2479 internet:msf@ie.utoronto.ca
http://www.ie.utoronto.ca/EIL/

Abstract

This paper provides a brief overview of the research underway in the Enterprise Integration Laboratory at the University of Toronto.

Introduction

In order to compete in the 1990s, industry must adopt what has come to be known as world class behavior. That is "products and process that never fail, shortest cycle time in the industry, competitiveness independent of volume, leadership in defining industry wide manufacturing excellence, and leadership in the development of the best people" [Hansen 91]. It has become increasingly apparent that limits in achieving world class behavior lie not just in labour or capital, but the availability/accessibility of information and the ability to effectively coordinate both decisions and actions.

A basic tenet of organization theory is that an organization's size can be managed by decomposing it into smaller units. While decomposition may reduce the number of entities a manager must manage directly, it introduces two problems: an information problem and a coordination problem. The information problem arises from the natural tendency of organizational units to erect barriers among units thereby reducing information flow; having the right information at the time becomes more difficult. The coordination problem arises from the need to coordinate the decisions and actions of each unit. Getting a product rapidly to market, for example, depends on countless, highly interdependent decisions by hundreds or thousands of individuals, from senior executives to production workers, spread around many engineering centers and manufacturing facilities. It is clear that in order to meet organizational goals in an increasingly competitive environment, decision making must be coordinated in a tight and at the same time flexible manner.

Management has seized upon the concept of Enterprise Integration as a solution to this problem. Many believe that by integrating the enterprise world class behavior will be achieved. The question is, what does it mean to integrate an enterprise? Hansen has stated that "a successful implementation exhibits five principles of leadership in the management of people and technology" [Hansen 91]:

"When people understand the vision, or larger task, of an enterprise and are given the right information, the resources, and the responsibility, they will 'do the right thing'."

"Empowered people - and with good leadership, empowered groups - will have not only the ability but also the desire to participate in the decision process."

"The existence of a comprehensive and effective communications network ..."

"The democratization and dissemination of information throughout the network in all directions irrespective of organizational position ..."

".. Distributed decision making. Information freely shared with empowered people who are motivated to make decisions

will naturally distribute the decision-making process throughout the entire organization."

An organization that encompasses these principles will exhibit a change in how individuals relate; from a hierarchical style to peer-to-peer; interactions among individuals and groups are managed by dialogue and negotiation. Consequently, knowledge will be of greater importance than rank. While we may intuitively understand the meaning of enterprise integration, operationalizing this intuition is the problem. Integrating the enterprise simply means that each unit of the organization will have access to information relevant to its task and will understand how its actions will impact other parts of the organization thereby enabling it to choose alternatives that optimize the organization's goals. Based on this definition, one can envision many ways of achieving integration. For example, having better access to information, or better communication of organization goals, or by providing better incentive structures.

Based on these principles, there are two components to achieving integration: 1) the design of better organization structures and behaviors, and 2) the use of information technologies to provide access to information, support decision making and aid in action execution. It is the latter path, which has come to be known as the Computer Integrated Enterprise (CIE), which is the focus of this report.

It is believed that a Computer Integrated Enterprise (CIE) will be able to rapidly respond to a changing market place by using computers to access information and integrate decision-making and the performance of activities. Although enterprise computing and communication technologies are evolving rapidly, the cornerstone of these technologies continues to lag in development: software. But it is not that we lack software per se, in fact there are "mountains" of software that purport to support individual decision making, but that we lack the theories and methods that would enable computers to provide access to information and support decision making across the organization. It is this theory and method deficit that places CIE beyond our grasp.

The Enterprise Integration Laboratory (EIL) research enables businesses to develop, manufacture, sell, deliver and support products and services with unprecedented speed, flexibility, quality and economy. This is achieved through the application of business practices and technologies that create a business infrastructure enabling the dissemination of information, coordination of decisions, and management of actions to and among people and systems within the organization and outside of it. EIL research explores the creation of Enterprise Integration concepts in a bi-directional manner, in that it is simultaneously theory and application driven; an underlying philosophy to this research is that solving real problems leads to breakthrough research. The theories that are being explored include: coordination theory, common sense enterprise modeling and agent-based enterprise information architectures. Applications include enterprise design, concurrent engineering and integrated supply chain management. The following sections summarize our work in these areas.

Concurrent Engineering: Design-in-the-Large

http://www.ie.utoronto.ca/EIL/DITL/design.html

We tend to visualize design as a problem solved by the solitary activity of an individual. However, there are actually two design problems: "Design-in-the-Small" (DITS) and "Design-in-the-Large" (DITL). Design-in-the-Small describes the activities of a single engineer or small group who select components, determine parameters, perform analyses, etc. In reality, this represents only a fraction of the overall design task. Most of the effort, and expense, is in "Design-in-the-Large".

Many artifacts, including those produced by the project's participating corporations, are so complex that the designs require the

efforts of many engineers. This means that the artifact must be functionally and/or physically decomposed, and the pieces divided amongst groups of engineers. The major issue in "Design-in-the-Large" is how to achieve high levels of collaboration among engineers. That is, how each engineer's design task can be managed so that it integrates well with the results of others. Collaboration among design teams is necessary because each part of a design constrains the others. A change in one part has a "domino effect" on other portions of the design. Unfortunately, changes occur frequently during the course of design, so that each engineer must continually revise his work.

We are approaching the DITL problem by using knowledge-based information technologies to enhance the degree of awareness, understanding, cooperation, and coordination among engineering team members. There are two factors that are critical to the success of this project. First is the unintrusive acquisition of design information and decisions; it is clear that acquiring design information/decisions is very difficult - most engineers do not "design with" computers. If information technology is to be a design process participant, then we must address the barriers to the adoption of technology by end users. Consequently, a major focus on this research is on the role of "electronic engineering notebooks" in the capture, storage and dissemination of design information and decisions, and on their role in integrating and managing design decisions and processes.

The second factor is the types of collaboration services to provide that would both aid the design process and entice engineers into using the technology. Towards this end, we focusing on two services: 1) the creation of an Integrating Knowledge Base (IKB) with supporting access and maintenance functions to provide engineers with the ability to find and/or be informed of information and knowledge of relevance to their task, and 2) the creation of a Collaboration Management System (CMS) that models and manages the technical problems that arise at the interface among engineering groups. For example, the impact that design decisions made by one group affect another, and the requirements that are not satisfied because they fall in between their interests.

Recent Work

We conducted a survey of coordination problems and how engineers spend their time which uncovered some surprises. For example, engineers spend at least 55% of their time searching for information or documenting information others may want. As a whole activities which involve coordination were found to occupy approximately 70% of an engineer's time in collaborative design, and included all but Problem Solving and Thinking. This lead to the observations, that information system tools that successfully aid in information acquisition and distribution will have a significant impact on engineering design productivity [Crabtree et al. 93a] [Crabtree et al. 93b] [Crabtree et al. 96].

Our next study explored the process of engineering at Spar Aerospace to understand how its done, and how it should be done, so that we could ascertain where in the process information systems tools may be best applied. The V model for systems engineering was identified [Fox & Salustri, 94]. The V-model includes concepts such as concurrency, life cycle analysis and risk management. Shared information and coordination appear to be critical to its success.

Concurrent to the above studies, we constructed an initial version of the Electronic Engineering Notebook and studied its efficacy by means of usability experiments [Louie 95]. The EEN provided a book-like look and feel with embedded functions such as: drawing, spread sheets, PIM and email. Information is entered using a pen and domain specific icons indicating: requirements, constraints, objectives, variables, etc. could be associated with the text and used to index information.

Usability studies were conducted with three versions of the EEN: large screen PC-based, large screen slate-based, and small screen Newton-based. It was shown that small

screen PDAs perform poorly as EENs because of screen size and pen limitations. With larger screens, it was found that EENs could be substituted for regular lab notebooks.

A major component of our effort is to provide company-wide access to engineering design information and knowledge. In pursuit of this, we have extended the TOVE Enterprise Model to include ontologies for: parts, assemblies, features, parameters, versions, revisions, requirements, and constraints. This has been implemented as part of the IKB [Lin et al., 96].

We have constructed a model of concurrency that provides for the communication of information to those who "need to know" when ever changes are made to the model [Gupta et al. 96].

An important information access tool for engineers is the capability to retrieve prior designs that are similar to what is currently needed. A requirement-based design case-based retrieval system has been developed that retrieves designs that partially satisfy design requirements [Bilgic & Fox 96].

Access to the IKB is provided across the Internet using World Wide Web standards. Anyone on the Web, with access privileges, may browse the IKB and make changes.

Integrated Supply Chain Management

http://www.ie.utoronto.ca/EIL/iscm-descr.html

The supply chain is a network of suppliers, factories, warehouses, distribution centers and retailers through which raw materials are acquired, transformed and delivered to the customer. Supply chain management is the strategic, tactical and operational level decision making that optimises supply chain performance. The strategic level defines the supply chain network, i.e., selection of suppliers, transportation routes, manufacturing facilities, production levels, warehouses, etc. The tactical level plans and schedules the supply chain to meet actual demand. The operational level executes plans. Tactical and operational level decision making functions are distributed across the supply chain.

In order to optimise performance, supply chain functions must operate in an integrated manner. But the dynamics of the enterprise and the market make this difficult; materials do not arrive on time, production facilities fail, workers are ill, customers change or cancel orders, etc. causing deviations from plan. In some cases, these events may be dealt with locally, i.e., they lie within the scope of a function. In other cases, the problem cannot be "locally contained"; modifications across many functions are required. Consequently, the supply chain management system must coordinate the revision of plans/schedules across supply chain functions.

The Integrated Supply Chain Management (ISCM) project addresses coordination problems at the tactical and operational levels. It is composed of a set of cooperating, intelligent agents, each performing one or more supply chain functions, and coordinating their decisions with other agents - this is called a Logistical Execution System (LES). The focus of our research is on the development of 1) a theory of coordination that allows agents to cooperatively manage change, 2) a theory of agent problem-solving that enables agents to cooperate with other agents in their exploration of solutions, and reason in an "anytime" manner, and 3) a theory of agency and support tools that enable users to build multi-agent systems with minimal programming effort.

Our approach views problem-solving as a constraint satisfaction/optimisation process where agents influence each other's problem solving behaviour through the communication of constraints. Coordination occurs when agents develop plans that satisfy their own internal constraints but also the constraints of other agents. Negotiation occurs when constraints, that cannot be satisfied, are modified by the subset of agents directly concerned. One of the main thrusts of this research is to investigate the use of constraints, their specification and

relaxation (i.e. modification) as a means of coordination and negotiation.

The objectives of the project are to:

- Develop a sharable representation of supply chain knowledge
- Identify an appropriate decomposition of supply chain functions and encapsulate into agents
- Develop a general constraint-directed problem solver that can be employed in every functional agent
- Develop an incremental, "anytime" model of constraint-directed problem solving for each functional agent so that it can provide rapid responses to unplanned for events
- Extend each function oriented agent so that it is able to answer more questions within its functional domain.
- Develop protocols, strategies and tools for the: communication of information, coordination of decisions, and management of change within multi-agent environments.

Recent Work

We have redefined the functions that agents in the supply chain should perform. We reviewed the functions provided by MRP systems and visited companies in order to understand their supply chain problems. Our conclusion is that current MRP system functional decompositions are inflexible, knowledge poor and algorithm poor. The following are the agents that we are developing: Logistics, Transportation Management, Order Acquisition, Resource Management, Scheduling and Dispatching.

We have incorporated and extended the TOVE Enterprise Modelling ontology as the basis for representing supply chain information and knowledge [Fox and Gruninger 94]. Enhancements have been in the areas of inventory and transportation. See the section on Enterprise Modelling for more details.

We have developed a Generic Agent Shell that supports agent construction in a more principled way, providing several layers of reusable services and languages for: agent communication, specification of coordination mechanisms, services for conflict management, services for information distribution, common sense reasoning, time, action, causality, etc. and integration of legacy application programs. It provides standard methods for:

- inter-agent communication: KQML/KIF.
- inter-agent coordination - a finite state language for coordination protocol specification [Barbuceanu & Fox 95a, 95b, 95d, 95e].
- agent knowledge representation - Description Logic [Barbuceanu, 1993].
- integrating a domain specific problem solver, e.g., translating information between the shell and the problem solver.

We have developed a unified theory of constraint-based scheduling and it is being used as the problem solver for the scheduling, resource management, dispatching, logistics and transportation agents [Davis & Fox 93][Davis 94]. The significance is that a single solver can be used for all of these agents. See the section on Coordination Theory for more information.

We have developed a theory of coordination as constraint relaxation that is being incorporated into the logistics and factory agents [Beck & Fox 94a, 94b] [Beck 94].

Enterprise Engineering

http://www.ie.utoronto.ca/EIL/enteng.html

The goal of this project is to conduct research leading to the creation of an information system to support Enterprise Design (also known as business process reengineering) and Execution. An enterprise design environment allows for the exploration of different designs or models of an enterprise along various perspectives such as efficiency, cost, quality and agility [Fox et al. 94]. Enterprise execution is concerned with monitoring the actual performance of the enterprise as specified by the model.

The Enterprise Engineering Project has the following objectives:

- Creation of integrated generic enterprise models encompassing representations for processes, time, causality, resources, products, quality, cost. etc [Fox and Gruninger 95].

• Formalize the knowledge found in Enterprise Engineering perspectives such as Time-based Competition, Quality Function Deployment, Activity-Based Costing, Quality Circles, Continuous Improvement, Process Innovation, and Business Process Re-Engineering. By formalize, we mean the identification, formal representation and computer implementation of the concepts, methods and heuristics which comprise a particular perspective. This not only enables a precise formulation of the intuitions implicit in practice, but it is also a step towards automating the execution of certain tasks involved in enterprise engineering.

• To formalize the intuition of design perspectives, we have introduced the notion of advisors [Gruninger and Fox 94c]. The tasks for each advisor fall into three main groups: evaluation, analysis, and guidance. For each advisor, we define their tasks, purpose, and responsibilities, and represent these tasks using the appropriate ontology. Each advisor is rigorously characterized by these tasks, including a specification of what an advisor is analyzing and in what way that they guide the designer by proposing different alternatives.

• Integrate the knowledge into a software tool that will support the enterprise engineering function by exploring alternative organization models spanning organization structure and behaviour. The Enterprise Engineering system allows for the exploration of a variety of enterprise designs. The process of exploration is one of design, analysis and re-design, where the system not only provides a comparative analysis of enterprise design alternatives, but can also provide guidance to the designer. The tool must therefore incorporate the enterprise models in order to support the integrated modelling and design of the enterprise.

• Provide a means for visualizing the enterprise from many of the perspectives mentioned above. The process of design is performed through the creation, analysis and modification of the enterprise from within each of the perspective visualizations.

Enterprise designs constructed with the tool will beexecutable. That is, portions of the enterprise design may be down-loaded for execution by the run-time system.

Recent Work

The following advisors have been constructed:

• An advisor for Activity-based Costing has been designed and a prototype constructed. The ABC advisor, given a TOVE enterprise model, is able to derive costs for activities, processes and products based on ABC principles. A major accomplishments of this work to date has been the formalization of ABC - which to date has been "fuzzy" at best and providing a solution to the problem of "when to stop accumulating costs" in a large activity model [Tham et al. 94].

• An advisor for ISO9003 has been designed and a prototype constructed. The ISO9000 advisor, given a TOVE enterprise model, is able to determine whether the model complies with ISO9003 requirements. The advisor provides a formal axiomatisation of ISO principles [Kim & Fox 94] [Kim et al. 95].

• An advisor for Time-based Competition has been designed and a prototype constructed.

• An advisor has been constructed that formalizes the business process reengineering heuristics of Hammer [Atefi 95].

We have an initial implementation of a visualization environment for enterprise design called Oak. This is a graphical user interface tool for building, browsing, editing, and visualizing a ROCK knowledge base. It currently provides support for dynamic construction of enterprise models from user-defined libraries and an interactive Prolog query interface to the underlying knowledge base.

In conjunction with our BPR Advisory Group, we have drafted a requirements document for the tools that we will design to support business process reengineering. We have developed a framework for characterizing the process of business process reengineering; this framework has been used to specify where in the BPR process is information technology useful as well as the nature of information technology to support BPR. This framework can be used in two ways (see

http://www.ie.utoronto.ca/EIL/grpdoc/frmwrk.html). First, we will be using it to specify a set of requirements on the software tools that the Enterprise Integration Laboratory will be designing to support BPR [Gruninger & Fox 95].

We can also use the framework as a means of evaluating existing tools and identifying the stages of the BPR process where they provide the most support. In this way, we can identify those stages of the BPR process that are not supported by existing tools. To assist in this endeavor, we are compiling a library of BPR tools and summaries of their capabilities (see http://www.ie.utoronto.ca/EIL/tool/BPR.html). The tool descriptions will be gathered from practitioners of BPR and vendors of the tools. This tool repository will be made available on the World Wide Web.

Common Sense Enterprise Modeling

http://www.ie.utoronto.ca/EIL/comsen.html

An Enterprise Model is a computational representation of the structure, processes, information, resources, goals and constraints of a business, government activity, or other organisational system. It can be both definitional and descriptive - spanning what should be and what is. The role of an enterprise model is to achieve model-driven enterprise design, analysis and operation.

A number of issues exist concerning the design of Enterprise Models [Fox 93]. Reusability is concerned with the large cost of building enterprise-wide data models. Is there such a thing as a generic, reusable enterprise model whose use will significantly reduce the cost of information system building? A second issue is the Consistent Usage of the model:

Given the set of possible applications of the model, can the model's contents be precisely and rigorously defined so that its use is consistent across the enterprise? A third issue is model Accessibility: Given the need for people and other agents to access information relevant to their role, can the model be defined so that it supports query processing, both surface and shallow? Lastly, there is the Selection issue: How do I know which is the right Enterprise Model for my application?

It is our belief that the issues of reusability, consistent usage, accessibility and selection can best be addressed by taking a more formal approach to enterprise modelling. By formal, we are not referring to analytical models as found in Operations Research, but to logical models as found in Computer Science. Towards this end, the TOVE ontology [Fox 92] [Fox et al. 93] has been developed. An ontology [Gruber 93] is a formal description of entities and their properties; it forms a shared terminology for the objects of interest in the domain, along with definitions for the meaning of each of the terms.

The goal of the TOVE (TOronto Virtual Enterprise) project is to create an enterprise ontology that has the following characteristics: 1) provides a shared terminology for the enterprise that every application can jointly understand and use, 2) defines the meaning (semantics) of each term in a precise and as unambiguous manner as possible using First Order Logic, 3) implements the semantics in a set of Prolog axioms that enable TOVE to automatically deduce the answer to many "common sense" questions about the enterprise, and 4) defines a symbology for depicting a term or the concept constructed thereof in a graphical context.

Our approach to engineering ontologies begins with defining an ontology's requirements; this is in the form of questions that an ontology must be able to answer. We call this the *competency* of the ontology. The second step is to define the terminology of the ontology - its objects, attributes, and relations. The third step is to specify the definitions and constraints on the terminology, where possible. The specifications are represented in First Order Logic and implemented in Prolog. Lastly, we test the competency of the ontology by "proving" the competency questions with the Prolog axioms. The TOVE ontologies constitute an integrated enterprise model, providing support for more powerful reasoning in problems that

require the interaction of multiple ontologies.

The TOVE ontologies have been implemented on top of C++ using the ROCK knowledge representation tool from Carnegie Group. The axioms within these ontologies have been implemented using Quintus Prolog and integrated with the ROCK knowledge base.

Recent Work

We have developed the following ontologies:

- activity ontology that spans activity, state, time and causality [Fox & Gruninger 94] [Gruninger and Pinto 95].
- resource ontology [Fadel et al. 94] [Fadel 94].
- traceability ontology to support quality reasoning [Kim & Fox 94] [Kim et al. 95].
- cost ontology to support activity-based costing [Tham et al. 94].
- organisation ontology spanning structure, roles, communication, authority and empowerment [Fox et al. 95].
- product ontology which includes features, parameters, assemblies, versions, revisions, requirements and constraints [Lin et al. 96a 96b].

We have developed a methodology for the definition and evaluation of ontologies [Gruninger & Fox 94a, 94b, 95].

We have built a World Wide Web interface to our ontology database.

Coordination Theory

http://www.ie.utoronto.ca/EIL/Scheduling.html

http://www.ie.utoronto.ca/EIL/coord.html

Conventional wisdom states that scheduling problems exhibit such a richness and variety that no single scheduling method is sufficient to solve all of them. Indeed, over the last fifteen years, Artificial Intelligence research into scheduling has explored a variety of methods for the scheduling of factories, space shuttle refurbishment, auto assembly lines, army logistics, etc. and have demonstrated a surprising degree of effectiveness. The question is, is AI, like Operations Research doomed to create a plethora of methods whose scope is limited, or does there exist a more general theory that underlies the AI approach?

Consider the Microboss and Gerry systems. Each uses a different method for assigning resources over time to activities. Microboss' "constructive" method starts with an empty schedule and iteratively selects and assigns resources to each activity over time, and backtracks whenever a dead-end (i.e., infeasibility) is reached. Gerry's "repair" method starts with a completed scheduled containing unsatisfied (i.e., broken) constraints and iteratively selects and reassigns resources to an activity over time that satisfy as many unsatisfied constraints as possible. Search in both cases ends whenever the "best" schedule is found that satisfies the constraints. The repair method can be augmented by using simulated annealing to find better schedules. On the surface they appear quite different, but a careful examination of their approach reveals that they are equivalent in all but a few details.

A major focus of our research is the development of a unified theory for constraint-directed scheduling. The primary components of the theory are a problem topology, textures and search policy. A problem topology is represented by a constraint graph where nodes are activities and arcs are constraints among the activities, e.g., temporal and resource. Constraints contain both predicates and optimization criteria, and propagation methods. A problem topology my be altered by problem reformulations, such as abstraction and aggregation. Problem textures are fundamental measures of constraint graphs that indicate decision complexity, uncertainty and elasticity. A search policy defines the process of committing and retracting search decisions.

A second focus of our coordination research is on the use of constraint-directed reasoning to coordinate the decisions of multiple agents, where constraints are used as basis for communication, influencing, coordination and negotiation in large multi-agent

distributed problems. This research is based on earlier research that explored how constraint graphs can be used to construct an "approximate topology" of a problem space, the development of problem space complexity measures for focusing search in the problem space, trade-offs between precision and uncertainty in decision making, and reformulation of the problem space structure in order to increase the efficiency of problem.

Recent Work

Our work has focused initially on the creation of a "generic scheduling shell" which is able to model and reason about constraints in order to solve complex scheduling problems. The technology underlying this agent is an integration of both the generative and iterative improvement constraint-directed reasoning techniques explored at CMU and NASA-Ames, respectively. Two versions of the shell have been created. Version 1 of ODO, our scheduling shell, was created as part of Gene Davis's masters thesis [Davis & Fox 93] [Davis 94]. It demonstrated the concept of a combined generative and iterative improvement system which share representations, propagation techniques, and constraint graph texture measurements. The choice of constraints, propagation, textures, and backtracking is left to the user to choose from a library we provide. A number of experiments were performed demonstrating ODO's performance as equivalent to both Microboss from CMU and GERRY from NASA. We also demonstrated that iterative improvement does little to improve good schedules that are generatively created.

Another direction of our work in coordination theory has investigated the role of constraint relaxation as a basis of agent coordination/negotiation. Chris Beck's masters thesis [Beck & Fox 94] [Beck 94] investigated heuristic methods for identifying minimum cost relaxations of constraints. Experiments demonstrated that this approach found the same minimal solution on Freuder's satisfiability problems with lower computational complexity, and the technique was applicable to constraint optimisation problems.

We have developed a method for building robust schedules, i.e., schedules that require minimal cost changes in response to stochastic events such as resource inavailability [Gao 95].

We have successful used ODO as the basis of the Logistics, Transportation scheduling, Factory Scheduler, Resource Manager and Dispatcher agents in the Logistical Execution System.

Enterprise Information Architectures

http://www.ie.utoronto.ca/EIL/infoarch.html

This research focuses on strategies for the automated acquisition and distribution of information from/to agents in a distributed environment, and the management of inconsistency that arises at the juncture of an agent's local data/knowledge bases and enterprise data/knowledge bases in distributed multi-agent information systems.

The Enterprise Information Architecture (EIA) consists of a distributed environment of mediator agents - called Information Agents (IA). An IA services a number of agents (functional and other IA-s) by providing them with a layer of shared information storage and services for managing it. Functional agents register their capabilities and needs with an Information Agent. Information Agents receive information from functional agents, reason about its relevance to other agents and distribute it to those agents for which they consider it relevant, and in the form that is easiest to understand. If distributed information ever becomes inconsistent, the Information Agent mediator will alert all receivers. If functional agents supply contradictory information, the mediator will apply strategies for solving the conflict and reinstall consistency. Communication takes place using the KQML (Knowledge Query and Manipulation Language) and KIF (Knowledge Interchange Format) languages produced by the (US) Advanced

Research Projects Agency Knowledge Sharing Initiative project, in which we are participating.

Recent Work

- Designed and implemented one of the first generic, application independent Information Agents, providing solutions to basic problems of intelligent information infrastructures [Barbuceanu & Fox 94b, 94c, 94d].
- Developed general solutions for the problem of relevance-based information distribution by facilitator-type agents in multi-agent environments.
- Constructed a general model for managing information inconsistency in multi-agent systems [Barbuceanu & Fox 95d].
- Developed solutions for information translation among partially overlapping ontologies, to support communication in multi-agent systems.
- Developed a model for the role of Information Agents as providers of intelligent information infrastructure services in multi-agent information and control systems.

Acknowledgements

The following members of the Enterprise Integration Laboratory have contributed enormously to this research: Katy Atefi, Nirmal Baid, Mihai Barbuceanu, Chris Beck, Taner Bilgic, John Chionglo, Robert Crabtree, Andrew Davenport, Eugene Davis, Fadi Fadel, Mahmud Gani, Hong Gao, Mike Gruninger, Jacek Gwidzka, Peng Hu, Henry Kim, William Leizerowicz, Jinxin Lin, Jim Louie, Bruce Marrett, Robert McDonald, Fil Salustri, Donald Tham, Yin Zhan.

We also acknowledge the contributions of our sponsors: BHP Research, Carnegie Group Inc., Digital Equipment Corp., ILOG SA, Manufacturing Research Corporation of Ontario, Micro Electronics and Computer Research Corp., National Science and Engineering Research Council, Numetrix Ltd., Quintus Corp., Spar Aerospace, Technical University of Berlin, and Toyo Engineering.

Publications

[Barbuceanu 93] Barbuceanu, M., (1993), "Toward Integrated Knowledge Modelling Environments", Knowledge Acquisition, Sep pp. 245-304.

[Barbuceanu & Fox 94a] Barbuceanu, M., and Fox, M.S., (1994), "The Architecture of a Generic Agent for Collaborative Enterprises", EIL working paper, nov. 1994.

[Barbuceanu & Fox 94b] Barbuceanu, M., and Fox, M.S., (1994), "The Information Agent - An Infrastructure Agent Supporting Collaborative Enterprise Architectures", Proc. of Third Workshop onEnabling Technologies: Infrastructure for Collaborative Enterprises, Morgantown, West Virginia, IEEE Computer Society Press.

[Barbuceanu & Fox 94c] Barbuceanu, M., and Fox, M.S., (1994), "The Information Agent: An Infrastructure for Collaboration in the Integrated Enterprise", in Proc. of Cooperative Knowledge Based Systems, University of Keele, Staffordshire, UK, M. Deen (ed), DAKE Centre, Univ. of Keele, 1994. http://www.ie.utoronto.ca/EIL/papers/abstracts/1.htm

[Barbuceanu & Fox 95a] Barbuceanu, M., and Fox, M. S., (1995), "COOL - A LAnguage for Describing Coordination Models in Multi-Agent Systems", Proceedings of the International Conference on Multi-agent Systems, San Francisco CA.

[Barbuceanu & Fox 95b] Barbuceanu, M., and Fox, M.S., (1995), "Conflict Management with the Credibility/Deniability Model", in S. Lander (ed) Proc. of AAAI Workshop on Models of Conflict management in Cooperative Problem Solving, Seattle, WA, AAAI Technical Report, AAAI Press.

[Barbuceanu & Fox 95c] Barbuceanu, M., and Fox, M. S., (1995), "An Agent Building Shell", Proceedings of CASCOM-95, Toronto: IBM Canada, to appear.

[Barbuceanu & Fox 95d] Barbuceanu, M., and Fox, M.S., (1995), "The Architecture of an Agent Building Shell", Proceedings of the IJCAI-95 Workshop on Agent Theories, Languages and Architectures , Montreal Canada.

[Barbuceanu & Fox 95e] Barbuceanu, M., and Fox, M.S., (1995), "An Agent-based Architecture for Agile Manufacturing", Proceedings of the IJCAI-95 Workshop on

Intelligent Manufacturing, Montreal Canada.

[Beck 94] Beck, J.C., (1994), "A Schema for Constraint Relaxation with Instantiations for Partial Constraint Satisfaction and Schedule Optimization", MSc. Thesis, University of Toronto. (also released as a technical report EIT-94-3). http://www.ie.utoronto.ca/EIL/profiles/chris/msc.thesis.abstract.html

[Beck & Fox 94a] Beck, J.C., and Fox, M.S., (1994), "Mediated Conflict Recovery by Constraint Relaxation", Working Notes of AAAI Workshop on Conflict Management, Seattle, USA. http://www.ie.utoronto.ca/EIL/profiles/chris/wcm-aaai94.abstract.html

[Beck & Fox 94b] Beck, J.C., and Fox, M.S., (1994), "Supply Chain Coordination via Mediated Constraint Relaxation", Proceedings of the First Canadian Workshop on Distributed Artificial Intelligence, Banff, Canada. http://www.ie.utoronto.ca/EIL/profiles/chris/cwdai.abstract.html

[Bilgic & Fox 96] Bilgic, T and M. S. Fox, "Constraint-Based Retrieval of Engineering Design Cases: Context as Constraints", Artificial Intelligence in Design '96, Stanford University, June 24--27, 1996, (to appear in S. Gero and F. Sudweeks (eds) AI in Design '96, Kluwer Academic Publishers). http://www.ie.utoronto.ca/EIL/public/ aid96/cbret1.html

[Crabtree et al. 93a] Crabtree, R., Baid, N., and Fox, M.S., (1993), "Design Engineering: Problems in Coordination", in Proceedings of the JSME/ASME Design Theory Workshop, Tokyo Japan. http://www.ie.utoronto.ca/EIL/papers/abstracts/4.html

[Crabtree et al. 93b] Crabtree, R., Baid, N., and Fox, M.S., (1993), "An Analysis of Coordination Problems in Design Engineering", in Proceedings of the International Conference on Engineering Design, The Hague, August 93. http://www.ie.utoronto.ca/EIL/papers/abstracts/37.html

[Crabtree et al. 96] Crabtree, R., Fox, M.S., and Baid, N., (1996), "Case Studies of Coordination Activities and Problems in Collaborative Design", accepted for publication in Journal of Engineering Design.

[Davis 94] Davis, G., (1994), "ODO: A Constraint-based Scheduler Founded on a Unified Problem Solving Model", M.A.Sc. Thesis, Enterprise Integration Laboratory Technical Report TR-EIL-94-1, Industrial Engineering Dept., University of Toronto.

[Davis & Fox 93] Davis, G., and Fox, M.S., (1993), "ODO: A Constraint-Based Scheduling Shell", Workshop on Production Planning, Scheduling and Control, IJCAI-93, Chambery, France, 29 Aug 93.

[Fadel 94] Fadel, F., (1994), "A Resource Ontology for Enterprise Modelling", M.A.Sc. Thesis, Enterprise Integration Laboratory Technical Report TR-EIL-94-1, Industrial Engineering Dept., University of Toronto.

[Fadel et al. 94] Fadel, F., Fox, M.S., and Gruninger, M., (1994), "A Resource Ontology for Enterprise Modelling", Proceedings of the Third Workshop on Enabling Technologies-Infrastructures for Collaborative Enterprises, IEEE Computer Society Press, pp. 117-128. http://www.ie.utoronto.ca/EIL/public/res_ontology.ps

[Fox 92] Fox, M.S., (1992), "The TOVE Project: A Common-sense Model of the Enterprise", Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, Belli, F. and Radermacher, F.J. (Eds.), Lecture Notes in Artificial Intelligence # 604, Berlin: Springer-Verlag, pp. 25-34. A version has been reprinted in the Proceedings of the International Conference on Object Oriented Manufacturing Systems, Calgary Manitoba, pp. 176-181, and in Proceedings of the Twelfth International Conference on Artificial Intelligence, Expert Systems, Natural Language Scientific Conference, Avignon France, pp. 23-30.

[Fox 93] Fox, M.S., (1993), "Issues in Enterprise Modelling", Proceedings of the IEEE Conference on Systems, Man and Cybernetics , Le Toquet France, IEEE. A revised version appeared in Information and Collaboration Models of Integration, S.Y. Nof (Ed.), NATO ASI Series #E259, Kluwer Pub, 1994.

[Fox & Gruninger 94] Fox, M.S. and Gruninger, M., (1994), "Ontologies for Enterprise Integration", Proceedings of the 2nd Conference on Cooperative Information Systems, Toronto, Ontario. http:// www.ie.utoronto.ca/EIL/public/onto_eil.ps

[Fox & Gruninger 95] Fox, M.S. and Gruninger, M., (1995), "Integrating Ontologies for Reengineering", IFIP WG5.7 Working Conference on Reengineering the Enterprise, Galway, Ireland, April 95.

[Fox & Salustri 94] Fox, M.S., and Salustri, F., (1994), "A Model of One-Off Systems Engineering", Proceedings of the Workshop on AI and Systems Engineering, Menlo Park CA: American Association for Artificial Intelligence. http:// www.ie.utoronto.ca/EIL/papers/abstracts/33.html

[Fox et al. 93] Fox, M., Chionglo, J.F., and Fadel, F.G., (1993), "A Common Sense Model of the Enterprise", Proceedings of the 2nd Industrial Engineering Research Conference , pp. 425-429, Norcross GA: Institute for Industrial Engineers. http://www.ie.utoronto.ca/EIL/papers/ abstracts/14.html

[Fox et al. 94] Fox, M.S., Gruninger, M., and Zhan, Y. (1994), "Enterprise Engineering: An Information Systems Perspective", Proceedings of the 3rd Industrial Engineering Research Conference, Norcross Georgia: Institute of Industrial Engineers, pp. 461-466.

[Fox et al. 95] Fox, M.S., Barbeceanu, M., Gruninger, M. (1995), "An Organisation Ontology for Enterprise Modelling: Preliminary Concepts for Linking Structure and Behaviour", Fourth Workshop on Enabling Technologies-Infrastructures for Collaborative Enterprises, (West Virginia University), IEEE Computer Society Press. http://www.ie.utoronto.ca/EIL/public/org.ps

[Gao 95] Building Robust Schedules using Temporal Protection- An Empirical Study of Constraint Based Scheduling Under Machine Failure Uncertainty, Technical Report TR-EIL-95-2, Industrial Engineering Dept., University of Toronto.

[Gruninger & Fox 94a] Gruninger, M., and Fox, M.S., (1994), The Role of Competency Questions in Enterprise Engineering ,Proceedings of the IFIP WG5.7 Workshop on Benchmarking - Theory and Practice, Trondheim, Norway. June 94. http://www.ie.utoronto.ca/EIL/public/ competency.ps

[Gruninger & Fox 94b] Gruninger, M., and Fox, M.S., (1994), "The Design and Evaluation of Ontologies for Enterprise Engineering", Workshop on Implemented Ontologies, European Conference on Artificial Intelligence 1994, Amsterdam, NL. http://www.ie.utoronto.ca/EIL/public/ onto_ecai94.ps

[Gruninger & Fox 94c] Gruninger, M., and Fox, M.S., (1994), "An Advisor-based Architecture for Enterprise Engineering", Workshop on Artificial Intelligence, AAAI-94 in Business Process Reengineering , Seattle.

[Gruninger & Fox 95a] Gruninger, M., and Fox, M.S., (1995), "The Logic of Enterprise Modelling", Proceedings of the Industrial Engineering Research Conference, Atlanta GA: Institute for Industrial Engineering.

[Gruninger & Fox 95b] Gruninger, M., and Fox, M.S., (1995), "Methodology for the Design and Evaluation of Ontologies", Workshop on Basic Ontological Issues in Knowledge Sharing, IJCAI-95. http:// www.ie.utoronto.ca/EIL/public/ method.ps

[Gruninger & Pinto 95] Gruninger, M., and Pinto, J.A. (1995), "A Theory of Complex Actions for Enterprise Modelling", Working Notes AAAI Spring Symposium Series 1995: Extending Theories of Action: Formal Theory and Practical Applications, Stanford. http:// www.ie.utoronto.ca/EIL/ public/complex.ps

[Gupta et al. 96] Gupta, L., Chionglo, J., and Fox, M.S., (1996), "A Constraint Based Model of Coordination in Concurrent Design Projects" , to appear in Proceedings of WET-ICE '96. http://www.ie.utoronto.ca/EIL/ DITL/WET-ICE96/ProjectCoordination/ WETICE96_ProjectCoordination.fm.html

[Gwizdka et al. 96] Gwizdka, J., Fox, M.S., and Louie, J., (1996), "EEN: A Pen-based Electronic Notebook for Unintrusive Acquisition of Engineering Design Knowledge", to appear in Proceedings of WET-ICE '96. http://www.ie.utoronto.ca/ EIL/DITL/WET-ICE96/EEN/ EEN_WetIce96.html

[Hansen 91] Hansen, W.C. The Integrated Enterprise. In Foundations of World-Class Manufacturing Systems: Symposium Papers. National Academy of Engineering, 2101 Constituion Ave, N.W., Washington DC, 1991.

[Hu & Fox 94] Hu, P., and Fox, M.S., (1994), "Managing Transportation by Constraint Heuristic Search", Working Paper, Enterprise Integration Laboratory, University of Toronto, Dec. 1994.

http://www.ie.utoronto.ca/EIL/public/transport.ps

[Kim & Fox 94] Kim, H., and Fox, M.S., (1994), "Formal Models of Quality and ISO 9000 Compliance: An Information Systems Approach", American Quality Congress (AQC) Conference, American Society for Quality Control, Milwaukee WI: American Society for Quality Control, pp. 17-23. http://www.ie.utoronto.ca/EIL/papers/abstracts/20.html

[Lin et al. 96] Lin, J., Fox, M.S., Gupta, L., and Leizerowicz, W., (1996), "Knowledge Network: An Information Repository with Services for Managing Concurrent Engineering Design", to appear in Proceedings of WET-ICE '96. http://www.ie.utoronto.ca/EIL/DITL/WET-ICE96/IKB/ikb.html

[Kim et al. 95] Kim, H., Fox, M.S., and Gruninger, M., (1995), "An Ontology of Quality for Enterprise Modelling", Proceedings of the Fourth Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises, IEEE Computer Society Press.

[Louie 95] Louie, J., (1995), "Electronic Engineering Notebook: Design and Analysis", M.A.Sc. Thesis, EIL Technical Report EIL-TR-95-1, Industrial Engineering Dept., University of Toronto, Canada.

[Tham et al. 94] Tham, D., Fox, M.S., and Gruninger, M., (1994)A cost ontology for enterprise modelling Third Workshop on Enabling Technologies-Infrastructures for Collaborative Enterprises, IEEE Computer Society Press, pp. 197-210. http://www.ie.utoronto.ca/EIL/papers/abstracts/30.html