Tilburg University

Cooperative information systems

Mylopoulos, J.; Papazoglou, M.

Published in:
IEEE Expert

Publication date:
1997

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Mylopoulos, J., & Papazoglou, M. (1997). Cooperative information systems. *IEEE Expert*, 12, 28-31.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Cooperative Information Systems

John Mylopoulos, University of Toronto
Michael Papazoglou, University of Tilburg

Traditionally, information systems consisted of databases and files storing large amounts of data, applications programs performing useful update or reporting tasks, and user interfaces for data entry or retrieval. Construction of such systems required analysis of an organizational setting, design of databases and applications programs, and implementation that depended on database and programming technologies. Moreover, such constructions were generally tailored to the features and needs of the customer organization, resulting in development projects that have been notorious for their underestimated costs, late delivery dates, and reported failures.

This picture is changing rapidly, partly because the software industry is maturing, making greater use of off-the-shelf components and generic solutions, and partly because of the onslaught of the information revolution. These changes have resulted in a new set of demands for information services that are integrated, open in their architecture, and global in their scope. Application domains that are driving the need for new information services include manufacturing, electronic commerce and banking, training, education, and environmental management.

Building the next generation of information systems that meet these demands poses two major challenges. The first is to develop tools and products that can bring together incompatible information sources and software. Software must operate across heterogeneous computing platforms, working in much the same way across global and public networks as it does in a local network. Access to such widely distributed resources is currently hampered not only by the presence of propriety or legacy systems, but also by the rapid expansion of information networks. Information networks accommodate a wide range of users with a broad range of needs and backgrounds. New technologies and tools must help these users to locate, retrieve, abstract, correlate, combine, and process data, to solve a variety of complex tasks.

The second challenge is to develop technologies that permit continuous enhancement and evolution of current massive investments in information sources and systems. Such technologies must offer an appropriate infrastructure that allows not only development but also evolution of distributed information systems. This infrastructure must support the conversion of large numbers of independent multivendor databases, knowledge bases, and application software into dynamic and highly connected cooperative components. These components must run over distributed information networks and address a variety of highly complex applications.

We call systems that meet these challenges *cooperative information systems*.^{1,2} Such systems have received considerable attention and research since the early '90s—in particular, the launch of an international journal, an ongoing conference, an international foundation, and numerous meetings and workshops (see the "Support for CIS Research" sidebar). This special issue will present some of the better research, thereby

promoting the topic and giving *IEEE Expert* readers a glimpse of the possible future of information systems.

The shape of information systems to come

What might these information systems of the future look like? Many people working in the area agree that such systems will consist of reconfigurable aggregations of disparate software components. Some components will come with their own information sources (databases, formatted or unformatted files, and the like), while others will be capable of performing useful transformations, such as translations or filtering or retrieval of information. A promising approach for building such aggregations of components is the *software information agent*.²

Each agent gains associated information, expertise, or both on a specific topic by drawing on information from various sources, including other information agents. Information-agent capabilities typically include interprocess-communication mechanisms and services, such as naming, translation, information discovery, syntactic/semantic-reconciliation, partial integration, distributed query processing, and transaction management. The agents' capabilities also include a host of other activities necessary to enable information sharing and cooperation between diverse, distributed, and incompatible information systems.

To accomplish a complex task, an information agent uses an incremental processing style, which might include dynamically and opportunistically recruiting other agents for executing tasks. This capability assumes that an information agent can invoke the functionality of other agents and that, more importantly, no a priori order of subtask execution exists.

This new computing paradigm extends the classic client-server model to a more dynamic and versatile architecture. Individual information agents help users effectively perform complicated tasks—for example, medical diagnosis, engineering design, banking decision making—efficiently using the most appropriate information and computing resources (whether processing, knowledge, or data) available in large computer networks. Moreover, the agents will be in a position to dynamically change their role in the

client-server relationship: at times they will be recipients of information services from other application-supporting agents, and at other times they will be providers.

Meeting the challenges

What are the research challenges of building cooperative information systems, and what research areas will help us meet them?

First, cooperation among agents (whether organizational, human, or software agents) presupposes that they can communicate and use each others' capabilities, or interoperate.

Moreover, the cooperating agents must be able to coordinate their activities, even if the cooperating agents are distributed and coordination is offline.

Support for CIS Research

In the early 90s, a series of workshops were launched on various aspects of cooperative information systems, sponsored by organizations such as the US National Science Foundation, the European Union, and the Canadian and the Australian Research Councils. As a result of these activities, the First International Conference on CIS met in Rotterdam, The Netherlands, in May 1993. Subsequent conferences met in Toronto in May 1994, Vienna in May 1995, Brussels in May 1996, and Kiawah Island, South Carolina, in June 1997. The *International Journal of CIS* has run successfully since March 1992.

The conference and the journal embrace research activities and original work from areas that contribute to CIS, such as distributed database systems, multiagent systems, computer-supported cooperative work, and distributed environments.

In 1995, the International Foundation on CIS was formed; it includes 20 prominent CIS researchers and educators. The foundation promotes CIS research and education at the international level and has been the major sponsor of CIS conferences since 1996.

For more information about the *International Journal of CIS*, access <http://www.spc.co.uk/wspc/Journals/ijcis/ijcis.html>. For more information about the International Foundation on CIS, access <http://www.ece.sc.edu/Labs/HIIT/CoopIS/IFCIS.html>.

In this issue

The articles in this special issue on cooperative information systems cover these categories:

Cooperative information services

Advanced information services should provide facilities for retrieving information through a form of approximate query evaluation. Approximation involves relaxing the query so that it will retrieve a subset of the answer of the original query. In "Cooperative Answering through Controlled Query Relaxation," Terry Gaasterland presents a theoretical foundation for query relaxation. She also discusses various relaxation techniques, in the context of deductive databases and deductive query processing.

"Query-Free Information Retrieval," an outstanding application article by Peter E. Hart and Jamey Graham, describes a system that supports query-free information retrieval from legacy databases. Hart and Graham present a completed project where the developed system has been experimentally deployed and is in use. The authors describe the AI techniques they adopted to develop the system; they illustrate its operation; and they explain how they evaluated it.

General and visionary

In "The Conceptual Basis for Mediation Services," Gio Wiederhold and Michael Genesereth articulate a compelling vision of how future information management should be done. The centerpiece of their vision consists of a generic architecture for interfacing users and application programs with information sources, such as legacy databases, knowledge bases, files, Web-based information sources, and data-generating computational processes (such as simulations). The mediation architecture consists of modules called *mediators*, which provide value-added information services, such as selecting relevant source material or providing statistical summaries. Communication protocols among mediators, information sources, and applications might range from as low-level as distributed messaging to as high-level as knowledge sharing.

Another research challenge is that these systems must offer state-of-the-art information services. Such services include cooperative query processing, where the system attempts to retrieve information for an intended, rather than a posed, query. For instance, if the user asks "Who won the presidential elections of 1998?" a cooperative system responds, "There were no elections," rather than a mere "No one." Moreover, information services must include generalized database-query facilities that can retrieve information from disparate sources through approximate (or similarity-based) matching, can discover generic rules that characterize a collection of raw data (data mining), and can support the creation of metadata repositories, which store descriptions of other information sources.

A final research challenge arises because both organizational and computing environments are highly volatile and change continually. Accordingly, cooperative information system technologies must grapple with the problem of change management, by having highly flexible architectures, by monitoring the organizational and computing environment to measure change, and by having built-in mechanisms for coping with it.

So, the technical challenges of building cooperative information systems fall into four broad categories:

- *Interoperation.* This category covers topics such as generic, open architectures, distributed object management, network-centric computing, compartmentalized applications, factoring out global control from individual components, integration of user and subsystem communication, communication protocols, translation mechanisms, data-integration mechanisms, semantic metadata repositories, knowledge sharing, and blackboard architectures.
- *Coordination.* This category's topics include computer-supported collaborative work, synchronous and asynchronous sharing, virtual workspaces, performers and customers, concurrency control, transaction management, mediation architectures, workflow systems, AI planning, multiagent technologies, intelligent scheduling, self-describing systems, and reflective architectures.
- *Information services.* Topics include cooperative querying, information retrieval, data mining, metadata management, data ware-

houses, information brokering, knowledge sharing, knowledge-level communication protocols, access to heterogeneous and distributed databases, information access on the Web, and application-specific subareas such as virtual enterprises, digital libraries, and electronic commerce.

- *Change management.* Topics include changes dictated by technology or organizational objectives, constraint enforcement, schema evolution, database-view updates, AI theories of action, truth-maintenance systems, constraint satisfaction, versions and configurations, impact analysis, risk assessment, business-process reengineering, and enterprise integration.

The advances that will lead to cooperative information system technologies will not come from any single research area of information technology. Knowledge-based systems, distributed systems, and database systems have each matured to the point that, although further enhancements are desirable, the greatest leverage for technological advancement should come from their integration into a seamless technology for building and managing cooperative information systems.

Databases, for example, can contribute data-management and engineering techniques, especially for distributed or heterogeneous database systems, data warehousing, and efficient implementations of information sharing. AI, on the other hand, has been traditionally preoccupied with the representation, acquisition, and use of knowledge. AI methods and techniques that are particularly relevant for cooperative information systems include data mining, distributed problem solving and multiagent systems, planning, scheduling, and negotiation. Object-oriented systems and techniques will also contribute to the development of appropriate tools and methodologies. Other related research areas include computer-supported collaborative work, distributed computing, and enterprise integration.

This special issue

We selected the three articles in this special issue from the proceedings of the first three International Conferences on Cooperative Information Systems.³⁻⁵ The articles cover aspects of information services (see the sidebar, "In this issue"). We asked the authors to extend and update their conference sub-

missions, and the resulting articles went through peer review.

We thank the reviewers of the articles submitted to the special issue, and the authors for their submissions and patience in following the publication demands of a special issue.

References

1. M.L. Brodie and S. Ceri, "On Intelligent and Cooperative Information Systems," *Int'l J. Intelligent and Cooperative Information Systems*, Vol. 1, No. 2, June 1992, pp. 249-290.
2. M.P. Papazoglou, S.C. Laufmann, and T.K. Sellis, "An Organizational Framework for Cooperating Intelligent Information Systems," *Int'l J. Intelligent and Cooperative Information Systems*, Vol. 1, No. 1, 1992, pp. 169-202.
3. M.P. Papazoglou and G. Schlageter, eds., *Proc. First Int'l Conf. Intelligent and Cooperative Information Systems*, IEEE Computer Society Press, Los Alamitos, Calif., 1993.
4. M.L. Brodie, M. Jarke, and M.P. Papazoglou, eds., *Proc. Second Int'l Conf. Cooperative Information Systems*, Univ. of Toronto Press, Toronto, 1994.
5. S. Laufmann, S. Spaccapietra, and T. Yokoi, eds., *Proc. Third Int'l Conf. Cooperative Information Systems*, Univ. of Toronto Press, 1995.

John Mylopoulos is a full professor of computer science at the University of Toronto. His research interests include knowledge representation and conceptual modeling, covering languages, implementation techniques, and applications. He has worked on the development of requirements and design languages for information systems, the adoption of database-implementation techniques for large knowledge bases, and the application of knowledge-base techniques to software repositories. He is principal investigator for both a national and a provincial Centre of Excellence for Information Technology. He received his PhD from Princeton University. He received the Canadian AI Society's first Outstanding Services Award and was corecipient of the most influential paper award at the 1994 International Conference on Software Engineering. He is a fellow of the AAAI and an elected member of the VLDB Endowment Board. He has served on the editorial board of several international journals, including the *ACM Transactions on Software Engineering and Methodology*, the *ACM Transactions on Information Systems*, the *VLDB Journal*, and *Computational Intelligence*. Contact him at the Dept. of Computer

Science, Univ. of Toronto, 10 King's College Rd, Toronto, ON M5S 3H5, Canada; jm@cs.toronto.edu.

Michael Papazoglou is a full professor and the director of the Infolab at the Univ. of Tilburg in the Netherlands. His scientific interests include cooperative information systems, heterogeneous database systems, object-oriented systems, distributed computing, digital libraries, and electronic marketing and commerce. He is the founding editor and coeditor in charge of the *International Journal of Cooperative Information Systems*. He serves on several committees and advisory boards for journals such as *Information Systems*, *AI Tools*, *Pattern Recognition and AI*, and *Applied Knowledge Engineering*. He has served as general and program chair for conferences such as the International Conference on Tools for AI, the International Conference on Cooperative Information Systems, the International Conference on the Entity Relationship Approach, the International Conference on Distributed Computing Systems, and the International Conference on Data Engineering. He has written or edited eight books and approximately 100 journal articles and refereed conference papers. He is a senior and Golden Core member of the IEEE, and a Distinguished Visiting Lecturer for the IEEE in the Asia/Pacific region. Contact him at the Infolab, Tilburg Univ., PO Box 5000 LE, Tilburg, The Netherlands; mikep@kub.nl; <http://infolabwww.kub.nl/infolab/people/mikep>.