

PROGRAM

3rd IFCIS INTERNATIONAL CONFERENCE ON



AUGUST 20-22, 1998 • NEW YORK CITY, NEW YORK

CoopIS '98



Sponsored by
The International Foundation on Cooperative Information Systems
New Jersey Institute of Technology


IEEE
COMPUTER
SOCIETY


IEEE

CONFERENCE AT A GLANCE

TH, AUG 20th		FR, AUG 21st		SA, AUG 22nd	
8:25- 8:30	Welcome				
8:30- 9:30	Invited Talk JENNIFER WIDOM	8:30- 9:30	Invited Talk AVI SILBERSCHATZ	8:30-10:00 Soho/Herald (Chair: Piccinelli) Gramercy (Chair: Adali)	S11: WF-II Hiramatsu Frank Luo S12: TR-II Kuo Lee Reddy
9:30-10:00	Coffee	9:30-10:00	Coffee	10:00-10:30	Coffee
10:00-12:00 Soho/Herald (Chair: Ozsu) Gramercy (Chair: Luo)	S1: WBIS Kerr Gruser De Rosa Huck S2: DD-I Chen Thiran Hasselbring Beyer	10:00-12:00 Soho/Herald (Chair: Schek) Gramercy (Chair: Gal)	S5: HDIS Vidal Ahmed Critchlow Moro S6: DD-II Baião Park Hung Josifovski	10:30-12:00	PANEL 2 Brodie Stanley Mylopoulos Dayal Raschid
12:00- 1:30	Lunch	12:00- 1:30	Lunch		
1:30- 3:00 Soho/Herald (Chair: Raschid) Gramercy (Chair: Schmitt)	S3: AG Tesch Chen Au S4: TR-I Vogler Schwarz Zhang	1:30- 3:00 Soho/Herald (Chair: Halper) Gramercy (Chair: Ahmed)	S7: MDB Norrie Lee Peng S8: MO Theodorakis Moulton Calvanese		
3:00- 3:30	Coffee	3:00- 3:30	Coffee		
3:30- 5:00	PANEL 1 Jarke Sheth Rusinkiewicz Ouksel Woo	3:30- 5:00 Soho/Herald (Chair: Dayal) Gramercy (Chair: Critchlow)	S9: WF-I Piccinelli Carlsen Joeris S10: II Schmitt Motz Adali		
6:00- 9:00	RECEPTION	6:00-11:00	BANQUET & SHOW		

Technical Papers

Session 1: Web-Based Information Services

(TH 10:00-12:00, Soho/Herald)

Chair: T. Ozsu

Information Services for the Web: Building and Maintaining Domain Models

S. Kerr, A. Gal, J. Mylopoulos

Wrapper Generation for Web Accessible Data Sources

J.-R. Gruser, L. Raschid, M. E. Vidal, L. Bright

Materializing the Web

M. De Rosa, T. Catarci, L. Iocchi, D. Nardi, G. Santucci

Jedi: Extracting and Synthesizing Information from the Web

G. Huck, P. Fankhauser, K. Aberer, E. Neuhold

Session 2: Distributed Databases I

(TH 10:00-12:00, Gramercy)

Chair: Y. Luo

Query Evaluation for Distributed Heterogeneous Relational Databases

Y. Chen, W. Benn

Interoperation of Independent, Heterogeneous and Distributed Databases: Methodology and CASE Support: the InterDB Approach

Ph. Thiran, J.-L. Hainaut, S. Bodart, A. Deflorenne, J.-M. Hick

A Generative Communication Service for Database Interoperability

W. Hasselbring, M. Roantree

Protecting the Quality of Service of Existing Information Systems

K. S. Beyer, M. Livny, R. Ramakrishnan

Session 3: Agents

(TH 1:30-3:00, Soho/Herald)

Chair: L. Raschid

Scheduling Non-Enforceable Contracts Among Autonomous Agents

T. Tesch, K. Aberer

Dynamic-Agents for Dynamic Service Provisioning

Q. Chen, P. Chundi, U. Dayal, M. Hsu

A Recipe Structure for Plan Execution in Dynamic Worlds

S. Au, J. Liang, N. Parameswaran

Session 4: Transactions I

(TH 1:30-3:00, Gramercy)

Chair: I. Schmitt

Using Multiple Mobile Agents for Distributed Transactions

H. Vogler, A. Buchmann

Execution Dependencies in Transaction Closures

K. Schwarz, C. Türker, G. Saake

A Novel Timestamp Ordering Approach for Co-existing Traditional and Cooperative Transaction Processing

Y. Zhang, Y. Kambayashi, Y. Yang, C. Sun

Session 5: Heterogeneous Distributed Information Services

(FR 10:00-12:00, Soho/Herald)

Chair: H. Schek

A Meta-Wrapper for Scaling up to Multiple Autonomous Distributed Information Sources

M. E. Vidal, L. Raschid, J.-R. Gruser

Management of Work in Progress in Relational Information Systems

R. Ahmed, U. Dayal

Meta-Data Based Mediator Generation

T. Critchlow, M. Ganesh, R. Musick

Data Access Services and Automatic Generation of Cooperating Interfaces

G. Moro, A. Natali, C. Sartori

Session 6: Distributed Databases - II

(FR 10:00-12:00, Gramercy)

Chair: A. Gal

Towards an Inductive Design of Distributed Object Oriented Databases

F. Baião, M. Mattoso, G. Zaverucha

A Client Group-Server DBMS Architecture and Inter-Client Communication Caching Schemes in the WAN Environment

K. Park, H. Kang

A Study of Least Privilege in CapBasED-AMS

P. C. K. Hung, K. Karlapalem, J. W. Gray III

Calculus-Based Transformations of Queries over Object-Oriented Views in a Database Mediator System

V. Josifovski, T. Risch

Session 7: Multidatabases

(FR 1:30-3:00, Soho/Herald)

Chair: M. Halper

OMS Connect: Supporting Multidatabase and Mobile Working through Database Connectivity

M. C. Norrie, A. Palinginis, A. Würzler

Minimization of Resource Consumption for Multidatabase Query Optimization

C. Lee, C.-H. Ke, J.-B. Chang, Y.-H. Chen

A Dynamic and Adaptive Cache Retrieval Scheme for Mobile Computing Systems

W.-C. Peng, M.-S. Chen

Session 8: Modeling

(FR 1:30-3:00, Gramercy)

Chair: R. Ahmed

Context in Information Bases

M. Theodorakis, A. Analyti, P. Constantopoulos, N. Spyrtos

Context Mediation on Wall Street

A. Moulton, S. E. Madnick, M. D. Siegel

Information Integration: Conceptual Modeling and Reasoning Support

D. Calvanese, G. De Giacomo, M. Lenzerini, D. Nardi, R. Rosati

Session 9: Workflows

(FR 3:30-5:00, Soho/Herald)

Chair: U. Dayal

Distributed Workflow Management: The TEAM Model

G. Piccinelli

Action Port Model: A Mixed Paradigm Conceptual Workflow Modeling Language

S. Carlsen

Managing Evolving Workflow Specifications

G. Joeris, O. Herzog

Session 10: Information Integration

(FR 3:30-5:00, Gramercy)

Chair: T. Critchlow

Merging Inheritance Hierarchies for Database Integration

I. Schmitt, G. Saake

Propagation of Semantic Modifications to an Integrated Schema

R. Motz, P. Fankhauser

A Flexible Architecture for Query Integration and Mapping

S. Adali, C. Buji

Session 11: Workflows II

(SA 8:30-10:00, Soho/Herald)

Chair: G. Piccinelli

Interworkflow System: Coordination of Each Workflow System among Multiple Organizations

K. Hiramatsu, K.-i. Okada, Y. Matsushita, H. Hayami

Integration of Statecharts

H. Frank, J. Eder

Cooperative Design for 3D Virtual Scenes

Y. Luo, R. Galli, M. Mascaro, P. Palmer

Session 12: Transactions II

(SA 8:30-10:00, Gramercy)

Chair: S. Adali

Using Constraints to Manage Long Duration Transactions in Spatial Information Systems

D. Kuo, V. Gaede, K. Taylor

Stable Transaction Management for Preserving the Global Integrity Constraints in Multidatabase Systems

K. Lee, S. Park

Reducing the Blocking in Two-Phase Commit Protocol Employing Backup Sites

P. K. Reddy, M. Kitsuregawa

Thursday 8/20/98:
RECEPTION
Marriott Marquis, 9th Floor Promenade
6pm-9pm

The CoopIS 98 Best Paper Award will be announced at the Reception:

**Information Services for the Web: Building and Maintaining
Domain Models**

S. Kerr, A. Gal, J. Mylopoulos

Friday 8/21/98:
BANQUET
Les San Culotte
347 West 46th Street
(Between 8th and 9th Avenues)
6:00pm-7:45

Friday 8/21/98:
BROADWAY SHOW
Les Miserables
Imperial Theater
249 West 45th Street
(Between 7th and 8th Avenues)
8:00pm-11:00

General Chairs
James Geller
New Jersey Institute of Technology, Newark, NJ USA

Frederick H. Lochovsky
University of Science & Technology, Clear Water Bay, Kowloon, Hong Kong

Program Chairs
Americas:

John Mylopoulos
University of Toronto, Ontario Canada

Europe, Middle East, & Africa:

Arne Sølvberg
The Norwegian University of Science and Technology, Trondheim, Norway

Asia & Pacific Rim:
Kyu-Young Whang
KAIST, Daejeon, South Korea

Panels Chair
Michael Brodie
GTE Laboratories Inc., Waltham, MA USA

Publicity Chair
Michael Halper
Kean University, Union, NJ USA

Finance Chair
Bonnie McKellar
Mobius Management Systems, Inc., Rye, NY USA

Local Arrangements
Priscilla Rasmussen
ARCS, New Brunswick, NJ 08901 USA

Invited Talk:
**Cooperative Information Integration and Exchange through
Semistructured Data**
Jennifer Widom
Stanford University

Semistructured data is data that may not conform to a fixed schema: the data may be irregular or incomplete, and its structure may change frequently and dynamically. The research area of managing semistructured data emerged about 5 years ago in response to two primary applications: (1) rapid integration of data from heterogeneous information sources, where it is difficult to impose a uniform, fixed structure over the diverse information; (2) management of information obtained via the World-Wide Web, where consistency and regularity are the exception rather than the rule. The recent emergence of XML (eXtensible Markup Language) as an expressive data representation and exchange format for the Web is particularly relevant and exciting to researchers in semistructured data, since the data models studied and implemented for semistructured data management are very similar to XML. We will review past and ongoing research in semistructured data management, explain how semistructured data facilitates cooperative integration and exchange of information, and discuss the potential impact of the close relationship between semistructured data and XML.

Jennifer Widom is an Associate Professor in the Computer Science and Electrical Engineering Departments at Stanford University. She received her Bachelors degree from the Indiana University School of Music in 1982 and her Ph.D. from Cornell University in 1987. She was a Research Staff Member at the IBM Almaden Research Center before joining the Stanford faculty in 1993. Professor Widom's research interests include heterogeneous and semistructured database, knowledge-base, and information systems, data warehousing, and active database systems. She has served on numerous editorial boards and program committees, and she has published widely in Computer Science.

Invited Talk:
Gestalts: Internet Databases and Scalable Integration
Avi Silberschatz
Bell Labs

(The described work is done jointly with Raghu Ramakrishnan.)

The Web today is essentially a collection of HTML objects, and the primary mode of querying is through HTML links. However, large *data collections*—not just individual HTML pages or objects—are becoming increasingly accessible over the Web, and traversing HTML links is inadequate as a retrieval and query mechanism. We argue that enhancing our model of the Web to include *collections of tables* as well as a *collection of HTML objects* is essential. This is a natural extension of distributed databases, but with quantum shifts in the degree of autonomy and number of nodes, and supporting it offers numerous challenges for the database research community. As networks become the backbone of information exchange, how we meet these challenges can determine whether database technology continues to be an integral part of information management or becomes “roadkill on the information superhighway.”

We discuss the following issues and how they are affected by the new environment: (1) Creating and managing integrated data repositories and sharing collections across an organization, (2) Query optimization and execution, (3) Change notification and management.

We also propose a novel “Publish-Register-Subscribe” framework as a solution to one of the problems that we identify, namely data integration on the Web: (1) Consumers can publish “container descriptions” for desired data, and search for and register into existing containers, (2) Producers can publish descriptions of data sources, and register their sources into published containers. In this approach, the task of creating a complex network of information is gracefully distributed by enabling each producer to contribute to one or more containers without knowledge of other sources of data.

Abraham Silberschatz (Ph.D., The State University of New York at Stony Brook) is the Director of the Information Sciences Research Center at Bell Laboratories, Murray Hill, New Jersey. Prior to joining Bell Labs, he was an endowed professor in the Department of Computer Sciences at the University of Texas at Austin. Dr. Silberschatz is a recognized researcher, educator, and author. His research interest include database systems, operating systems, and distributed systems. His writings have appeared in numerous ACM and IEEE publications and he is co-author of two well known textbooks: *Operating System Concepts* and *Database System Concepts*. Dr. Silberschatz is a Fellow of the ACM.

Panel 1:
**Cooperative IS Manifesto: Components of Solutions, Emphasis and
Recipe for Success**

Matthias Jarke
University of Aachen
Amit Sheth
University of Georgia
Aris Ouksel
University of Illinois at Chicago
Marek Rusinkiewicz
MCC
Carson Woo
University of British Columbia, Canada

In recent editions of the CoopIS conference series, the term "Cooperative Information Systems" has been seen in a purely technological interpretation, implying openness/interoperability of system architectures, technologies to enable and control coordination and collaboration, more explicit modeling of information semantics, the efficient and safe implementation of these concepts, etc. But both the business models of successful companies as well as other research communities that study cooperation (e.g., CHI and CSCW) suggest that we need more than technology for the Cooperative ISs to succeed. In particular, we need to recognize that systems also serve individual and social purposes in work practice.

In a broader agenda that can consist of technologies for coordination and collaboration; modeling of people, processes, and organizations; domain-specific system semantics, behavior and capabilities; and so on, the panel seeks to address the following questions:

- What are the key components of CoopIS that we have ignored in the past, but are critical for our future success as a vibrant scientific and engineering community?
- What is the right balance of technology vs. non-technical issues for CoopIS researchers and practitioners?
- How is our role distinguished from those of related communities, and what uniquely defines us?

Panel 2:
**Cooperative Information Systems: An Industrial Reality and a
Research Fiction**

Michael L. Brodie
GTE Laboratories Incorporated
Jory Stanley
SAP America
Umeshwar Dayal
Hewlett-Packard Laboratories
John Mylopoulos
University of Toronto
Louisa Raschid
University of Maryland

Prior to automation, people cooperated to achieve business objectives within business processes such as selling a product, manufacturing a product, delivering a service, and billing for goods and services. Computers were introduced to support, augment, or replace humans in these business processes. Until the early 1990s, the primary application of computing to business processes was to implement business functions or tasks. The cooperative aspects were automated only for the most trivial forms of cooperation (e.g., data transfer, invocation of well-defined transactions). Competitive pressures have led to increasing demands for the automation of cooperation beyond simple transactions to sophisticated cooperation and collaboration across the entire business process. Indeed, many requirements exceed expectations that had been made of humans.

Various research communities, including those for Databases, Information Systems, Operations Research, and Artificial Intelligence, have, for some time, addressed many approaches to non-trivial cooperation, collaboration, and interoperation between automated functions. However, few if any have been reduced to practice. On the other hand, the industrial community in directing its focus at specific application domains, such as supply chain, manufacturing, human resources, finance, and retail sales, has developed real cooperative information systems (CoopIS) far beyond those conceived of in the research communities. The panelists are asked to respond to the assertion that CoopIS are already fully operational in industry, in many domains, with a next generation in the works, while the research community has yet to grasp the nature of this new paradigm. As a basis for this debate, the panel will discuss the extent to which sophisticated CoopIS are currently in production and the features envisaged for the next generation of commercial CoopIS. The panel will also review current research approaches to CoopIS in terms of the models and features currently being developed and the research challenges being faced. Finally, the panel is asked to identify capabilities that are critical for CoopIS and that are, as yet, far from reality.