科技资料

8th International Conference on Data Engineering

Eighth International Conference on

Data Engineering

February 2-3, 1992 *Tempe, Arizona*

Sponsored by
IEEE Computer Society Technical Committee on Data Engineering

Worldwide
Information
Systems
Bull

With corporate support by
Bull Worldwide Information Systems

IEEE Computer Society Press
Los Alamitos, California
Washington • Brussels • Tokyo

Published by the
IEEE Computer Society Press
10662 Los Vaqueros Circle
PO Box 3014
Los Alamitos, CA 90720-1264

IEEE Computer Society Press Order Number 2545
Library of Congress Number 91-76564
IEEE Catalog Number 92CH3097-3
ISBN 0-8186-2545-7 (paper)
ISBN 0-8186-2546-5 (microfiche)
ISBN 0-8186-2547-3 (case)

Additional copies can be ordered from

IEEE Computer Society Press
Customer Service Center
10662 Los Vaqueros Circle
PO Box 3014
Los Alamitos, CA 90720-1264

IEEE Service Center
445 Hoes Lane
PO Box 1331
Piscataway, NJ 08855-1331

IEEE Computer Society
13, avenue de l'Aquilon
B-1200 Brussels
BELGIUM

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama
Minato-ku, Tokyo 107
JAPAN

Editorial production: Anne Copeland MacCallum and Lisa O'Conner
Cover art: Joe Daigle Studio Productions
Printed in the United States of America by Braun-Brumfield, Inc.

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Message from the General Chairperson

I would like to welcome you to Phoenix and the Eighth International Conference on Data Engineering (ICDE-8). ICDE is back in the United States after its first successful overseas conference in 1991 in Kobe, Japan. Given the strategic position that computer technologies play in the Phoenix area, this is a very exciting place to host ICDE-8. We hope that you will find both the tutorials and technical program useful and stimulating, and that you will take advantage of this opportunity to engage in social and technical dialog with your colleagues from other parts of the world. The submissions we received for this conference from so many different countries makes ICDE-8 a truly international conference.

An international conference of this size and diversity requires much effort and coordination from many people and organizations. It is simply impossible to recognize all persons who have contributed their efforts to make ICDE-8 what it is today. Nevertheless, we would like to acknowledge the invaluable support we have received from the following individuals. As General Chairperson, my primary responsibility is coordinating various tasks carried out by other willing and talented volunteers. I would like to take this opportunity to express my sincere appreciation to Dr. Forouzan Golshani, Program Chair; Carol Chapman, his assistant; and his 60-member program committee for doing an excellent job in selecting 69 high-quality papers for presentation at the conference. Forouzan made my job easy. Thanks are due to Tom Howell for arranging excellent tutorials. Special thanks go to Lana Ruch for taking care of local arrangements, to Julie Mumpy for awards, to our publicity chair Mansur Samadzadeh, to our European coordinator, Herbert Weber, and our industry coordinator, Bob Horton, to our registration chair Oris Frieson, to our treasurer Richard Snodgrass, and for publications, Chenho Kung. I would like to express my thanks to the ICDE steering committee for their vote of confidence.

I would also like to express my appreciation to both Professor Stefano Ceri and Dr. Laszlo for accepting our invitation to be the keynote speakers. Both of them are highly active and productive researchers who have made significant contributions to many areas of computer science and engineering.

Special mention must be made to Bull Worldwide Information Systems; without their generosity, the conference would not be as successful as we expect. Finally, I would like to thank all the attendees of ICDE-8, and hope that you will find time to enjoy both the conference and the city of Phoenix to the utmost extent.

Nick Cercone
General Chairperson
Vancouver, 1992

Message from the Program Chairperson

First and foremost, I would like to acknowledge the generous financial support that Bull Worldwide Information Systems has provided to this conference. I would also like to thank numerous employees of Bull for their assistance in various capacities. Several other companies, including American Express, Honeywell and Intel, have made valuable contributions to the organization of this event.

Following the tradition of previous meetings, this eighth occurrence of IEEE International Conference on Data Engineering begins with a keynote address from Dr. Laszlo Beladey, a leader for many years in various R & D centers. Dr. Stefano Ceri, known for his numerous achievements in the database field, will deliver the other keynote address.

During the planning phase, our hope was to make this conference a meeting place for industry, academia and research centers. Thus, the theme of technology transfer seemed appropriate. As a first step, we organized a plenary session with a highly qualified panel of experts, each representing a particular segment of the computer science community in general and databases in particular.

The number of papers submitted to the 1992 conference was consistent with the figures from previous years. Each submission was reviewed by at least three referees. When there was no consensus, as many as five people read the paper to determine its acceptability. A total of 213 papers were submitted, out of which 69 were selected for inclusion in the proceedings. This represents an acceptable rate of less than one third. In addition, two panel proposals were accepted for inclusion in the program.

On behalf of the Program Vice Chairs, members of the program committee, and the entire Phoenix crew, I welcome you to Phoenix and hope that you will enjoy this conference.

Forouzan Golshani
Program Chairperson

ICDE-8
Committee Members

General Chairperson:
Nick J. Cercone
Simon Fraser University, Canada

Program Chairperson:
Forouzan Golshani
Arizona State University

Program Vice-Chairpersons:
Francois Bry
ECRC, Germany

Ahmed Elmagarmid
Purdue University

Koichi Furukawa
ICOT, Japan

Yao-Nan Lien
AT&T Bell Laboratories

Witold Litwin
University of Paris, France

M. Tamer Ozsu
University of Alberta/GTE Labs

D. Scott Parker
UCLA

Tony Pizzarello
Bull Information Systems

Marek Rusinkiewicz
University of Houston

Gunter Schlageter
Fern University, Germany

Susan D. Urban
Arizona State University

Kyu-Young Whang
Korea Advanced Institute of Science and Technology, Korea

Tutorials Chair:
Tom Howell
Bull Information Systems

Local Arrangements:
Lana Ruch
American Express

Awards:
Julie Mumpy
Bull Information Systems

Treasurer:
Richard Snodgrass
University of Arizona

Publications:
Chenho Kung
University of Texas at Arlington

European Coordinator:
Herbert Weber
Dortmund University, Germany

Industrial Coordinator:
Bob Horton
Honeywell

Publicity:
Mansur Samadzadeh
Oklahoma State University

Registration:
Oris Friesen
Bull Information Systems

Steering Committee:
C.V. Ramamoorthy
University of California, Berkeley

P. Bruce Berra
Syracuse University

Benjamin W. Wah
University of Illinois

John Carlis
University of Minnesota

Joseph E. Urban
Arizona State University

Tadao Ichikawa
Hiroshima University, Japan

ICDE-8
Program Committee Members

Divy Agrawal
Elisa Bertino
Alejandro Buchmann
Sharma Chakravarthy
Arbee Chen
Suzanne Dietrich
Weimin Du
Ramiz Elmasri
Oris Friesen
Arif Ghafoor
William I. Grosky
Jiawei Han
Ki-Joon Han
A.R. Hurson
Keith Jeffrey
Genevieve Jomier
Larry Kerschberg
Wolfgang Klas
Ramamohanarao Kotagiri
Vijay Kumar
James Larson
Dik Lee
Gerard Levy
Leszek Lilien
Akifumi Makinouchi
Rainer Manthey
Dennis McLeod
Nobyuyoshi Miyazaki
Marguerite Murphy
Erich J. Neuhold
Shojiro Nishio
Maria Orlowska
Sylvia Osborn
Jon F. Parks
Sakti Pramanik
Raghu Ramakrishnan
Krithi Ramamritham
David Reiner
Felix Saltor
Amit Sheth
L. Tanca
Shalom Tsur
Patrick Valduriez
Gerhard Weikum
James Wheeldreyer
Kazumasa Yokota
Phillip Yu

ICDE-8
Referees

Amr El Abbadi
Serge Abiteboul
Seth Abraham
D. Agrawal
Ishfaq Ahmad
Romas Aleliunas
Eric Amiel
Jean-Marc Andreoli
Constantin Arapis
P. Aristides
William Armstrong
Hideki Asoh
G. Barbian
K. Barkel
Ken Barker
Raj Kumar Batra
Petra Bayer
Aime Bayle
W. Benn
F.A. Berkel
T. Berkel
Eliza Bertino
Bharat Bhargava
Alan Bier
Jose A. Blakeley
Alan Blei
C. Bosch
M.W. Bright
D.P. Bradshaw
Omran Bukhres
Frank Calliss
Michael Carey
Qiming Chen
T.S. Chen
Kam-Hei Cheng
Y.H. Chin
Jan Chomicki
Panos K. Chrysanthis
Andrea Clematis
Enrique Cortes-Rello
Peter Creasy
Bogdan Czejdo
Peter Dadam
Asit Dan
Ajoy Kumar Datta
Danco Davcev
Andrew Davison
Umeshwar Dayal
Lois Delcambre
Nevenka Dimitrova
Klaus Dittrich
G. Dong
Weimin Du
Christoph Eick
Margaret Eich
H.-D. Enrich
P. Fankhauser
Abdel Farrag
J. Feldkamp
Mariagrazia Fugni
Tetsuya Furukawa
Ghassan Z. Gadah
S. Gala
S. Garg
Dimitrios Georgakopoulos
A. Geppert
Uwe Geuder
Vittoria Gianuzzi
Alan Goerner
James E. Gray
Mingsen Guo
Ki-Joon Han
Brian Hansche
Lilian Harada
Theo Harder
J.A. Harland
John V. Harrison
Christoff Hasse
Abdelsalam Heddaya
Sandra Heiler
Mary Lou Hines
Ki-Hyung Hong
Eui Kyeong Hong
Ryosuke Hotaka
Yong Hu
Jan-Nong Huang
Stephen Huang
Yannis E. Ioannidis
L. Iowin
Sushil Jajodia
Jing Jing
A. Johnsey
D. Jonscher
Bellosto Marie Jose
Yahiko Kambayashi
Hyunchul Kang
G. Karabatis
Anton Karadimce
Kazuhiko Kato
Hirofumi Katsuno
Kyoji Kawagoe
Thomas F. Keefe
David Kemp
G. Kiernan
Hyoung-Joo Kim
Jeonghee Kim
Jin-Ho Kim
Kyung-Chang Kim
Myung-Joon Kim
S. Kim
Yasushi Kiyoki
H. Knolle
R. Krieger
Eva Kuhn
C.H. Kung
T.V. Lakshman
Jean-Marie Larcheveque
Paul Larson
Tim Leask
C. Lee
K.C. Lee
Sang Ho Lee
Alexandre Lefebvre
Fred Lewis
Thomas Little
Ling Liu
M.L. Lo
Darrell Long
Wei Lu
David E. Maier
Roberto Maiocchi
Frank Manola
A. Mansfield
R. Manthey
James Marland
M. Marmann
Yoshifumi Masunaga
S. Mazumdar
Claudia Bauzer Medeiros
Rajiv Mehrotra
Alberto Mendelzon
A. Merchant
R. Meyer
Tatsuo Minohara
D. Mishra
Daniel P. Miranker

Referees (Continued)

Takao Miura
Nobuyoshi Miyazaki
B. Monkeberg
Yukihiro Morita
Kunihiko Moriya
D. Morley
Louse Moser
Eliot Moss
Jim Mullen
Eisuke Muroga
Peter Muth
Jacques Noye
A. Ohori
Natsuki Oka
Ryuichi Oka
Toni Ordi
Seog Park
Young-Chul Park
Bill Perrizo
Randal Peters
Niki Pissinou
Calton Pu
Bhaskar Purimetla
Ghassan Z. Qadah
Abbas Rafii
Thomas C. Rakow
David Reiner
C. Rich
Jim Richardson
Ri Hotaka
Tore Risch
Wagner Roland
P. Sander
Ludger Schaefers
A. Scherer
Donovan Schneider
M. Scholl
G. Scholz
F.A. Schreiber
J.W. Schultz
Hirohisa Seki
Timos Sellis
Arunabha Sen
Hans Sfreibel
Dennis Shasha
Wen-Gong Shieh
Xiaojun Shen
D. Sherman
Phillip Sheu
Kohji Shibano
Chia-Shiang Shik
Shinji Shimozyo
Toramatsu Shintani
Andrea Sikeler
Mukesh Singhal
H.W. Six
Sang H. Son
E.A. Sonenberg
S.M. Sripada
Michael Stonebraker
W. Stucky
Tzong-An Su
V.S. Subrahmanian
V. Sugumaran
Junping Sun
Kazem Taghva
Toshihisa Takagi
Makoto Takeyama
Milinda Tambe
Katsumi Tanaka
Takehisa Tanaka
Yuzuru Tanaka
Ernest Teniente
James Thom
Gomer Thomas
Jesper L. Traff
Markus Tresch
Wei-Tek Tsai
Volker Turau
A. Tzreili
Haruyasu Ueda
E.A. Unger
Vibhavasu Vuppala
Chen-Wan Wang
Shigeru Watari
Horst F. Wedde
Erhard Welker
John A. Wertz
K.Y. Whang
J.K. Wight
W. Wilkes
Andreas Wolf
M.H. Wong
Gene Wuu
S. Xavier
Shi-Nine Yang
Kazumasa Yokota
Jong P. Yoon
Masatoshi Yoshikawa
Li Yu
Takashi Yukawa
P. Zabback
Carlo Zaniolo
Aidong Zhang
Jianhua Zhu
M. Ziane
Gerhard Zimmerman
Justin Zobel

Table of Contents

Track B

Knowledge Based Systems
Chair: M. Samadzadeh

Object Oriented Databases
Chair: P. Gargiulo

Panel

Track C

Database Tools
Chair: G. Occhini

Data Integrity
Chair: A. Sheth

Interoperability
Chair: G. Jomier

Wednesday, December 5, 1992

Plenary Panel Session
Technologe Transfer
Moderator: A. Pizzarello
Panelists: J. Browne, M. Diethelm, and R. Stanley

Track A

Replicated Data Management

Performance Analysis
Chair: A.L.P. Chen

Storage Management
Chair: N. Bourbakis

Track B

Logic and Databases
Chair: S. Dietrich

Object Management and Versioning
Chair: D. Carver

Deductive Databases
Chair: K. Ramamritham

Track C

Query Processing
Chair: D. Davcev

Relational Database Techniques
Chair: E. Neuhold

Track B

Integrity Constraints
Chair: R. Brice

Temporal Databases
Chair: D. Cohen

Track C

Query Management
Chair: S. Tsur

Database Modeling and Design
Chair: S. Pramanik

Algorithms I

Chair: Leszek Lilien

A SPANNING TREE TRANSITIVE CLOSURE ALGORITHM

Shaul Dar * †

H. V. Jagadish *

† University of Wisconsin
Madison, Wisconsin

* AT&T Bell Laboratories
Murray Hill, New Jersey

ABSTRACT

We present a new transitive closure algorithm that maintains a spanning tree of successors for each node rather than a simple successor list. This spanning tree structure promotes "sharing" of information across multiple nodes and hence leads to more efficient algorithms. We suggest an effective relational implementation of the spanning tree storage structure, and show how blocking can be applied to reduce the I/O cost of the algorithm. The algorithm can handle path problems in addition to simple transitive closure. The algorithm not only establishes the fact that two nodes are connected but also implicitly records one path between the two. For extremal path problems, the algorithm can record the best path between every pair of nodes.

We describe how the spanning tree storage structure can be updated incrementally in response to changes in the underlying graph. We also demonstrate how compression techniques can be applied in conjunction with the closure computation.

1. INTRODUCTION

Transitive closure is regarded as an important functionality of future database systems [1, 10, 16, 22], and considerable research has been devoted to devising algorithms for computing the transitive closure of a database relation [3, 4, 7, 13, 14, 19, 20, 25, 26].

All of the algorithms mentioned above use the set of successors of a node as the basic unit of storage and manipulation. In this paper we explore the benefits of maintaining structure in the successor list. We show that by keeping a successor spanning tree for each node, rather than a flat list, we can promote sharing of partial successor lists, and reduce the I/O cost of the algorithm, especially when there are many alternative paths between nodes in the graph. The spanning tree structure can be utilized for the computation of path problems, as well as reachability, and can make effective use of selection conditions specified in the query. An important advantage of this structure is that it stores a path from the root node to every node reachable from that root node. For extremal path problems, such as shortest path, the best path between the root node and its successors can be recorded. Another important advantage of the successor spanning tree structure is that it lends itself to incremental updates, in response to node and edge additions and deletions.

The rest of the paper is organized as follows. In Section 2, we briefly review some current techniques for computing transitive closure. In Section 3 we present the new algorithm using successor spanning trees to compute reachability in a directed acyclic graph. We suggest an efficient relational implementation for the successor spanning tree structure, and show how blocking can be applied to reduce the I/O cost of the algorithm. We also demonstrate that the algorithm can utilize selections effectively. In section 4 we present a modification of the basic algorithm for path computations, and a variation of this algorithm that can be used to record extremal path information. In Section 5 we study the performance of the spanning tree transitive closure algorithm both analytically and through simulation. In Section 6 we discuss how the spanning tree storage structure can be updated incrementally, and describe a compression scheme that can be used to minimize the cost of storing the computed transitive closure. We conclude the presentation in Section 7.

2. BACKGROUND

Let G be a directed graph with n nodes and e edges. We denote the transitive closure of G by G^*. We say that node j is a *successor* of node i if $<i,j> \in G^*$. We call j an *immediate successor* of i if $<i,j> \in G$. Given the set of immediate successors for each node in G, we wish to construct the complete successor set for each of these nodes. Two important families of algorithms to construct these successor sets are based on the matrix and graph representations of G.[1]

2.1 Matrix-Based Algorithms

Given an $n \times n$ adjacency matrix of elements a_{ij}, with a_{ij} being 1 if there is an arc from node i to node j, and 0 otherwise, the Warshall algorithm [27] computes the transitive closure of the given graph as follows:

$$\forall_{j=1}^{n} \forall_{i=1}^{n} \text{ process } a_{ij}$$

"Processing" of an element a_{ij} involves examining whether a_{ij} is 1, and if it is, then making every successor of j a successor of i. Thus, the Warshall algorithm computes closure by "processing" every element of the matrix exactly once, column by column from left to right, and from top to bottom within a column. In [4] the processing order was modified to promote "blocking", that is, the processing of a block of successor lists at one time,

1. There is also a family of iterative algorithms that has been studied extensively [7, 13, 26]. We do not discuss these here since they are not relevant to the spanning tree algorithm.

0-8186-2545-7/92 $3.00