Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis

409

A. Buchmann O. Günther T.R. Smith Y.-F. Wang (Eds.)

Design and Implementation of Large Spatial Databases

First Symposium SSD '89 Santa Barbara, California, July 17/18, 1989 Proceedings



Springer-Verlag Berlin Heidelberg New York London Paris Tokyo Hong Kong

Editorial Board

D. Barstow, W. Brauer, P. Brinch Hansen, D. Gries, D. Luckham, C. Moler, A. Phueli, G. Seegmüller, J. Stoer, N. Wirth

Editors

Alejandro P. Buchmann GTE Laboratories, Inc. 40 Sylvan Road, Waltham, MA 02254, USA

Oliver Günther Forschungsinstitut für anwendungsorientierte Wissensverarbeitung (FAW) Universität Ulm Postfach 2060, D-7900 Ulm, FRG

Terence R. Smith Yuan-Fang Wang Department of Computer Science, University of California Santa Barbara, CA 93106, USA

5-13-1

CR Subject Classification (1987): A.0, E.1--2, E.5, F.2.2, H.2.3--2, H.2.8, U.2.1, H.3.5

ISBN 3-540-52208-5 Springer-Verlag Berlin Heidelberg New York ISBN 0-387-52208-5 Springer-Verlag New York Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in other ways, and storage in data banks. Duplication of this publication or parts thereof is only permitted under the provisions of the German Copyright Law of September 9, 1965, in its version of June 24, 1985, and a copyright fee must always be paid. Violations fall under the prosecution act of the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1990. Printed in Germany

Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr. 2145/3140-543210 – Printed on acid-free paper

Dedicated to the Memory of

Markku Tamminen

1945 - 1989

BIBLIOTHEQUE DU CERIST

Preface

This book contains the proceedings of the First Symposium on the Design and Implementation of Large Spatial Databases (SSD '89), which was held with about 175 participants at Santa Barbara, California, on July 17 and 18, 1989.

The great interest in this symposium seems to be representative of the growing interest in spatial data management in general. There are numerous spatial applications in geography, computer vision, robotics, computer-aided manufacturing, and environmental information systems, and these applications require more powerful data management tools than the ones available today.

Research in spatial data management requires expertise in these application areas *and* in various fields within computer science, such as database management, data structures and algorithms, computational geometry, solid modeling, and computer vision. Experts from the application areas have to cooperate with computer scientists in a highly interdisciplinary field to obtain systems that are both practical and at the cutting edge of today's computer science.

It is our hope that this symposium also served as an opportunity to bring together people from these various disciplines and to establish closer connections between these fields.

We would like to thank NASA, the Environmental Protection Agency, the Oak Ridge National Laboratory, and the U.S. Geological Survey for their generous support. Thanks also to ACM, the IEEE Computer Society, and to Springer-Verlag for their cooperation. The National Center for Geographic Information and Analysis here at Santa Barbara has been supportive of this project from the beginning. Thanks to the members of the program committee for returning their reviews promptly under great time pressure. And thanks to Sandi Glendinning for taking care of our local arrangements.

General Chair;

Oliver Günther, FAW-AI Laboratory, University of Ulm, FRG

Program Chair:

Alejandro Buchmann, GTE Laboratories Inc., Waltham, Massachusetts, USA

Symposium Committee:

Renato Barrera, University of Maine, USA Stavros Christodoulakis, University of Waterloo, Canada Umesh Dayal, Digital Equipment Corporation, Cambridge, Massachusetts, USA Hans-Dieter Ehrich, University of Braunschweig, FRG John Estes, NASA and University of California, Santa Barbara, USA Christos Faloutsos, University of Maryland, USA Andrew Frank, University of Maine, USA Sylvia Osborn, University of Western Ontario, Canada Ralf H. Güting, University of Dortmund, FRG Klaus Hinrichs, University of Siegen, FRG Alfons Kemper, University of Karlsruhe, FRG Hans-Peter Kriegel, University of Bremen, FRG Raymond Lorie, IBM Almaden Research Center, California, USA Frank Manola, GTE Laboratories Inc., Waltham, Massachusetts, USA Jack Orenstein, Object Design, Inc., Burlington, Massachusetts, USA Hans-J. Schek, ETH Zürich, Switzerland Timos Sellis, University of Maryland, USA Terence Smith, University of California, Santa Barbara, USA Markku Tamminen, Helsinki University of Technology, Finland Yuan-F. Wang, University of California, Santa Barbara, USA Eugene Wong, University of California, Berkeley, USA

Contents



Data Structures

Invited Talk:
7 \pm 2 Criteria for Assessing and Comparing Spatial Data Structures
J. Nievergelt, ETH Zürich, Switzerland
The Fieldtree: A Data Structure for Geographic Information Systems
A. U. Frank, R. Barrera, University of Maine, USA
A Full Resolution Elevation Representation Requiring Three
Bits per Pixel
C. A. Shaffer, Virginia Polytechnic Institute, USA
System and Performance Issues
The DASDBS GEO-Kernel: Concepts, Experiences, and the
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step A. Wolf, ETH Zürich, Switzerland
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step A. Wolf, ETH Zürich, Switzerland
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step A. Wolf, ETH Zürich, Switzerland Performance Comparison of Point and Spatial Access Methods H. P. Krianel, M. Schiwiztz, P. Schneider, P. Sconger
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step A. Wolf, ETH Zürich, Switzerland Performance Comparison of Point and Spatial Access Methods HP. Kriegel, M. Schiwietz, R. Schneider, B. Seeger, University of Performance Comparison of Point and Spatial Access Methods
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step A. Wolf, ETH Zürich, Switzerland Performance Comparison of Point and Spatial Access Methods HP. Kriegel, M. Schiwietz, R. Schneider, B. Seeger, University of Bremen, FRG 89
The DASDBS GEO-Kernel: Concepts, Experiences, and the Second Step A. Wolf, ETH Zürich, Switzerland Performance Comparison of Point and Spatial Access Methods HP. Kriegel, M. Schiwietz, R. Schneider, B. Seeger, University of Bremen, FRG Strategies for Optimizing the Use of Redundancy in Spatial Databases

Geographic Applications

Invited Talk:
Tiling Large Geographical Databases
M. F. Goodchild, University of California, Santa Barbara, USA
Extending a Database to Support the Handling of Environmental
Measurement Data
L. Neugebauer, University of Stuttgart, FRG
Thematic Map Modeling
M. Scholl, A. Voisard, INRIA, Chesnay, France
Quadtrees
Invited Talk:
Hierarchical Spatial Data Structures
H. Samet, University of Maryland, USA
Distributed Quadtree Processing
C. H. Chien, T. Kanade, Carnegie-Mellon University, Pittsburgh, USA 213
Node Distribution in a PR Quadtree
CH. Ang, H. Samet, University of Maryland, USA

Modeling and Data Structures

An Object-Oriented Approach to the Design of Geographic Information Systems P. van Oosterom, J. van den Bos, University of Leiden, The Netherlands . . 255 A Topological Data Model for Spatial Databases M. J. Egenhofer, A. U. Frank, J. P. Jackson, University of Maine, USA ... 271

A Well-Behaved File Structure for the Storage of Spatial Objects	
M. W. Freeston, European Computer-Industry Research Center,	
Munich, FRG	287

Spatial Reasoning

Invited Talk:

The Design of Pictorial Databases Based upon the Theory
of Symbolic Projections
SK. Chang, E. Jungert, Y. Li, University of Pittsburgh, USA

BIBLIOTHEQUE DU CERIST

Data Structures

BIBLIOTHEQUE DU CERIST

7 ± 2 criteria for assessing and comparing spatial data structures

Jurg Nievergelt ETH Zurich and UNC at Chapel Hill

Abstract

Spatial data structures have evolved under the influence of several forces: 1) Database technology, with its emphasis on modeling and logical organization; 2) the long history of data structures developed in response to requirements from other applications; and 3) the recent rapid progress in computational geometry, which has identified typical queries and access patterns to spatial data. Rather than attempting a comprehensive survey of many spatial data structures recently developed, we aim to identify the key issues that have created them, their common characteristics, the requirements they have to meet, and the criteria for assessing how well these requirements are met. As a guideline for tackling these general goals, we begin with a brief history and recall how past requirements from other applications have shaped the development of data structures. Starting from the very early days, five major types of applications generated most of the known data structures. But the requirements of these applications do not include one that is basic to spatial data: That objects are embedded in Euclidian space, and access is mostly determined by location in space.

We present six specifically geometric requirements spatial data structures must address. Sections 3, 4, 5 discuss the mostly static aspects of how space is organized, and how objects are represented and embedded in space. Sections 6, 7, 8 consider the dynamic aspects of how objects are processed. We differentiate three types of processing, of increasing complexity, that call for different solutions: common geometric transformations such as translation and rotation; proximity search, and traversal of the object by different types of algorithms. Together with the general requirement of effective implementability, we propose these seven criteria as a profile for assessing spatial data structures. This survey leads us to two main conclusions: 1) That the current emphasis on comparative search trees is perhaps unduly influenced by the great success balanced trees enjoyed as a solution to the requirements of older applications that rely on single-key access, and 2) that spatial data structures are increasingly of the 'metric' type based on radix partitions of space.

Affiliation of author: Jurg Nievergelt (jn@inf.ethz.ch, jn@cs.unc.edu) Informatik, ETH, CH-8092 Zurich, Switzerland and Dept.Computer Sci., Univ. of North Carolina, Chapel Hill, NC 27514, USA