

D. E. Knuth

Cc 01-606

Axioms and Hulls

BIBLIOTHEQUE DU CERIST

Springer-Verlag

Berlin Heidelberg New York

London Paris Tokyo

Hong Kong Barcelona

Budapest

Series Editors

Gerhard Goos
Universität Karlsruhe
Postfach 69 80
Vincenz-Priessnitz-Straße 1
W-7500 Karlsruhe, FRG

Juris Hartmanis
Department of Computer Science
Cornell University
5149 Upson Hall
Ithaca, NY 14853, USA

Author

Donald E. Knuth
Department of Computer Science, Stanford University
Stanford, CA 94305-2140, USA

62.31

CR Subject Classification (1991): I.3.5, F.2.2, G.2.2

1991 Mathematics Subject Classification: 52B55, 52C05, 68Q25, 52B40, 05C20

ISBN 3-540-55611-7 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-55611-7 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 any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1992
Printed in Germany

Typesetting: Camera ready by author/editor
Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.
45/3140-543210 - Printed on acid-free paper

Preface

A FEW YEARS AGO some students and I were looking at a map that pinpointed the locations of about 100 cities. We asked ourselves, "Which of these cities are neighbors of each other?" We knew intuitively that some pairs of cities were neighbors and others were not; we wanted to find a formal mathematical characterization that would match our intuition.

Our first solution was rather complicated. We decided to say that points p and q of a given set S are neighbors if the set V_p of all points in the plane that are closer to p than to any other point of S is adjacent to the set V_q of all points that are closest to q . Another way to state this condition is so say that V_p and V_q have a common boundary point t ; point t is then equidistant from p and q , and every point between p and t belongs to V_p , while every point between q and t belongs to V_q . After several more minutes we realized that the key fact was the existence of a circle running through points p and q (centered at t), with no points of S inside the circle.

We began to look for an algorithm that would find all neighboring pairs of points $\{p, q\}$, according to our criterion. But time ran out; our meeting had to break up, and we went our separate ways. I wonder what would have happened if we had had more time to explore the problem on our own, before learning that it was a famous problem in computational geometry.

Leo Guibas's office was next to mine, and we soon learned from him that points p and q are neighbors in S by our definition if and only if the line segment pq belongs to the so-called Delaunay triangulation of S . I had never encountered that branch of geometry before, and I hadn't had time to read much of the fast-growing literature of computational geometry. After all, I had never promised to write a book about such things, and the other topics that had kept me going for 30 years were already proving to be more than adequate to occupy an entire lifetime! Furthermore I knew that my geometric intuition was rather poor; algebra and logic have always been much easier for me than visualization. I have absolutely no ability to understand 3-dimensional objects until I have built physical models to represent them.

Leo gave me a reprint of [38], a paper that explains (among other things) how to compute the Delaunay triangulation of n points in $O(n \log n)$ steps, using simple primitive operations and elegant data structures. I knew I didn't have time to read it, but I was immediately fascinated. Here was the clarity I was looking for, a paper that provided a bridge between algebraic and geometric intuition. I still had difficulty, however, understanding some of the proofs, which relied on diagrams that illustrated particular cases. I craved a proof of the algorithm that could be checked by a computer.

So I spent a pleasant afternoon developing a set of five axioms from which algorithms could be devised for a simpler problem - to find the convex hull of points in the plane. These algorithms for convex hull did not depend directly on the coordinates of the points; they relied only on testing whether or not a point p lies to the left or right of a line from q to r , when p , q , and r are given. Moreover, the algorithms could be proved correct purely from the axioms. I didn't need to draw any diagrams, although of course I constructed the proofs with images in mind.

A new set of axioms defines a new “universe.” We can experience some of the fun of Star Trek adventurers when we explore the properties of a given list of logical rules. Of course, many sets of axioms are contradictory or immediately seen to be trivial; others are too “random” to be interesting. But some sets, like the axioms for group theory or for projective geometry, lead to immensely fruitful territory. I soon realized that the five axioms I had stumbled across were leading to a potentially rich theory, worthy of further study.

I wrote three pages of notes, to show to Leo and to keep on file for future recreation. Then I noticed a few more things, and decided to write a short paper on the subject. I had learned in the meantime that the mathematical systems defined by my axioms were equivalent to a special class of oriented matroids, so I began to refresh my memory of that subject. During a long plane ride to Singapore I played with the ideas some more, and found that “vortex-free tournaments” have many nice properties intimately related to four axioms of my five. A few days later I proved a theorem that appears in section 6 of these notes, while sitting next to a lake in Singapore’s exotic Botanical Gardens.

Every time I wrote up one topic, other interesting questions would suggest themselves; one step led naturally to another and then another, often with unexpected ties to other branches of mathematics and computer science. Soon the manuscript for my “short paper” had grown to more than 50 pages, but I knew that I was not yet half done. I had become thoroughly hooked on the subject—constantly aware that I was supposed to be doing other things, yet unable to resist mathematical beauty. I was nowhere near a natural boundary at which I could terminate these initial explorations. After supervising about 30 Ph.D. dissertations, I felt like I was embarking on a second thesis of my own.

Finally I completed the first 16 sections of the present notes, which are devoted entirely to axioms for the computation of convex hulls in the plane. I had learned a lot, but I knew that this was merely a beginning; I had investigated those axioms only to prepare myself for the real goal, which was to study the more complex problem of Delaunay triangulation. The latter problem requires axioms for another primitive operation, namely to test whether a point p lies inside the circle that passes through three other points $\{q, r, s\}$.

With a mixture of excitement and trepidation, I turned to the Delaunay triangulation problem and found to my great relief that the same five axioms would work for that problem as well, with only minor extensions. Thus, all of my warmup exercises had turned out to be immediately applicable to the problem I had hoped to solve. Hurray! The remaining work went quickly.

I have tried to write these notes in such a way that readers may share the fun I had during an exhilarating voyage of discovery. Let’s face it: Research is thrilling. Instead of merely presenting the facts, I have tried here to give a reasonably faithful account of the questions I asked and the answers I liked, following closely the chronological order in which the work was done. I did not know what would appear in section $n+1$ of these notes when I was exploring the topics discussed in section n . The final section contains some of the questions that seem to demand answers next. Calls for further exploration of the terrain can be heard from many directions.

Of course much of this work parallels the research of others; the theory developed here is merely an introduction to the vast and beautiful subject of oriented matroids, which will surely continue to provide inspiration to many more generations of researchers in mathematics and computer science. The most remarkable thing about oriented matroids is perhaps that they are extremely interesting even in the simplest, lowest-dimensional cases.

Section 15 of the present notes can be read separately, because it is largely independent of the rest of the material. It discusses “parsimonious algorithms,” a notion that is a natural outgrowth of any axiom-based approach to algorithm design: We say that an algorithm is *parsimonious* if it never makes a test for which the outcome could have been anticipated from the results of previous tests, with respect to a given set of axioms. Algorithms that meet this condition are robust, in instructive ways that may prove to be important in practice, although stronger types of robustness are achievable in the convex hull and Delaunay triangulation problems.

This book may, incidentally, be interesting to typography buffs as well as to computer scientists, because of the rapid turnaround time provided by Springer-Verlag. It is the first publication to use the final revision of the Computer Modern typefaces, released two weeks ago. I made the arrowheads longer and stronger, so that they will not disappear so easily on xerox copies; and I redesigned a few of the letterforms, such as \mathcal{I} , \mathcal{T} , and δ . There is also an improved method for digitization at low resolution. The new characters do not affect any of the line breaks or page breaks made by \TeX , because they fit into exactly the same size boxes as the old ones did. Everyone now using the Computer Modern fonts of 1986 should soon be able to install the 1992 fonts in their place, at little or no cost. I promise not to change them again.

— Donald E. Knuth
Stanford, California
April 1992

Summary. A CC system is defined to be a relation on ordered triples of points that satisfy five simple axioms obeyed by the “counterclockwise” relation on points in the plane. A CCC system is a relation on ordered quadruples, satisfying five simple axioms obeyed by the “incircle” relation. In this monograph, the properties of these axioms are developed and related to other abstract notions such as oriented matroids, chirotopes, primitive sorting networks, and arrangements of pseudolines. Decision procedures based on the CC axioms turn out to be NP-complete, although nice characterizations of CC structures are available. Efficient algorithms are presented for finding convex hulls in any CC system and Delaunay triangulations in any CCC system. Practical methods for satisfying the axioms with arbitrarily degenerate data lead to what may well be the best method now known for Delaunay triangulations and Voronoi diagrams in the Euclidean plane. The underlying theme of this work is a philosophy of algorithm design based on simple primitive operations that satisfy clear and concise axioms.

Contents

Introduction	1
1 Axioms	3
2 Independence	6
3 Interior triple systems	9
4 Vortex-free tournaments	11
5 Pre-CC systems and CC systems	16
6 An NP-complete problem	19
7 Fitting tournaments together	23
8 Reflection networks	29
9 Enumeration	35
10 Oriented matroids	40
11 Convex hulls	45
12 Another algorithm	47
13 Comparison of algorithms	52
14 Degeneracy	55
15 Parsimonious algorithms	61
16 Composition of CC systems	68
17 The incircle predicate	69
18 A generalized Delaunay algorithm	73
19 Incircle degeneracy	82
20 Generalization to higher dimensions	86
21 Historical remarks	93
22 Open problems	97
Acknowledgments	99
Bibliography	100
Index	105