POINTS, SPHERES, AND SEPARATORS

A Unified Geometric Approach to Graph Partitioning

Shang-Hua Teng

August 1991
CMU-CS-91-184

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Submitted to Carnegie Mellon University
in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.

This research was sponsored in part by the National Science Foundation under Contract CCR-9016641.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing official policies, either expressed or implied, of NSF or the U.S. government.

School of Computer Science

DOCTORAL THESIS
in the field of
Computer Science

A UNIFIED GEOMETRIC APPROACH
TO GRAPH PARTITIONING

SHANG-HUA TENG

Submitted in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy

ACCEPTED:

Gary L. Miller	12/13/91
MAJOR PROFESSOR	DATE
R. Ruby	12-13-91
DEAN	DATE

APPROVED:

Paul Chrystall	17 December 1991
PROVOST	DATE

Preface

Geometry is full of visual imagination and concrete intuition. Such imagination and intuition is of great value not only for the research worker, but also for anyone who wishes to study and appreciate the results in geometry.

In this thesis, a unified geometric approach is presented for graph partitioning – a fundamental problem in computer science that has important applications in numerical analysis, VLSI design, computational geometry, complexity theory, and other fields. The main ingredient in obtaining this unified approach is a novel geometrical characterization of graphs that have small *separators*, where a separator of a graph is a relatively small subset of vertices whose removal divides the rest of the graph into two disconnected pieces of approximately equal size.

The characterization is based on elementary geometric concepts such as points, balls, cubes, and spheres. More specifically, a new class of geometric graphs, *overlap graphs*, is proposed. This class has the following properties:

1. In two dimensions, planar graphs are special cases of overlap graphs.
2. In d dimensions ($d \geq 2$), any finite subgraph of the infinite d-dimensional grid is an overlap graph.
3. Every overlap graph of n vertices in d dimensions has an $O(n^{(d-1)/d})$ separator.

At the time of this writing, this is the first time that a class of graphs has been proposed with these three natural properties. The proof that planar graphs are special cases of overlap graphs relies on recent deep theorems by Andreev and Thurston characterizing all planar graphs in a geometric fashion. A consequence is a new geometric proof of a classical theorem of Lipton and Tarjan that every planar graph has an $O(\sqrt{n})$-separator. This is another beautiful illustration of the use of geometry in understanding combinatorial concepts.

Moreover, the class of overlap graphs in d dimensions includes as a special case the class of k-nearest neighborhood graphs – an important class of graphs from computational geometry, statistical analysis, and image understanding. Therefore, I am able to give the first proof that all k-nearest neighborhood graphs of n points in d dimensions have $O(k^{\frac{1}{d}} n^{\frac{d-1}{d}})$-separators. It is worthwhile to mention that the above separator bound for k-nearest neighborhood graphs is also optimal in terms of k. The overlap graph also contains various other classes of graphs defined geometrically, including sphere packing graphs and graphs which can be drawn with finite resolution.

In addition, the overlap graph includes (almost) all graphs associated with finite element and finite difference methods in numerical analysis and scientific computing. Graphs from finite element and finite difference methods are defined geometrically. They are in general meshes of elements in a fixed dimensions (typically 2, 3 or 4 dimensions), that are *well shaped*, i.e., with bounded aspect-ratio and/or without any small angle nor obtuse angle. Therefore, the new separator result can be used not only in Nested Dissection for generating an almost optimal elimination ordering for the sparse Cholesky factorization but also for devising efficient sequential and parallel algorithms for

mapping sparse matrices onto various distributed processor/memories parallel systems to reduce communication cost and to achieve load balance.

Furthermore, all the arguments are based on geometric properties of embedding. The separator bounds come with randomized linear time and processor-efficient randomized NC algorithms.

Graph partitioning is a fundamental subroutine in applying divide and conquer paradigm for developing efficient sequential and parallel algorithms for various problems in computer science and scientific computing. Using the new separator result, algorithms are developed for solving several problems in computational geometry, that are more efficient than the multi-dimensional divide and conquer algorithms of Bentley. These new results on separators lead to a new way to design efficient parallel algorithms for geometric problems in fixed dimensions.

Acknowledgments

Gary Miller is the first theoretician I ran into by chance, a random chance. That was back in the spring of 1986 in SAL 246, the Xerox room of Computer Science Department at the University of Southern California where I was a student working on parallel architectures. He showed me his office, and we had quite a long conversation – not about parallel architectures, not about theory, but about how to encode Chinese characters efficiently in computers. I had never talked so much and so "bravely" in front of any professor before – Confucianess is very hard to overcome, I must admit! I really enjoyed our conversation but did not expect that three days later Gary found me and asked whether I was interested in working with him. "I don't know, Professor Miller", I answered. But I accepted his pile of papers... By the time I called him "Gary" instead of "Professor Miller", I became his student.

Gary was a great source of inspiration and encouragement. He taught me the value of research, of independent thinking, and of cooperation. He spent tremendous time and effort on me and was always ready to discuss my research. He made many contributions to my work, and this thesis is mainly joint work with him and Steve Vavasis (Cornell University).

I would also like to thank the rest of my thesis committee, Alan Frieze, Ravi Kannan, and Danny Sleator for their invaluable help in preparing the thesis and more. Alan not only made the major contribution in our joint work "On the complexity of computing the diameter of a polytope", but also helped me in my own work on various aspects. He was a valuable person to be around; Ravi was always ready to share his broad knowledge and deep insight; Danny introduced me to online computing and gave me a lot of advice, help, and guidance on my work on "Dynamic scheduling on parallel machines" (joint with Anja Feldmann and Jiři Sgall). He was a very good friend and willing to talk about everything.

In addition to my thesis committee, many individuals deserve special mention and thanks. Len Adleman introduced me to complexity theory and cryptography. Ming-Deh Huang, who treated me more as a colleague than as a student, taught me a lot about number theory during our joint work. Merrick Furst introduced the notion of VC-dimension to me. Richard Anderson, Eric Bach, Guy Blelloch, Seymour Ginsburg, Kai Hwang, Erich Kaltofen, Ming-Yang Kao, Mike Luby, John Reif, Steve Rudich, Larry Ruzzo, and Doug Tygar were among many others who have gone out of their way to teach me.

I am very grateful to all my co-workers from whom I have learned so much. They are Mike Atallah, Anja Feldmann, Alan Frieze, Hillel Gazit, Ming-Deh Huang, Rao Kosaraju, Larry Larmore, Gary Miller, Jiři Sgall, Steve Vavasis, and Bin Wang. Steve Vavasis, in particular, made many invaluable contributions to this thesis.

It was such an experience to go through graduate programs in two universities. Someone may say that I wasted one year or so. But I got the chance to have twice as many wonderful friends who greatly enhanced my years as a graduate student. From USC to CMU, to name just a few: Kireeti Kompella, Hillel Gazit, Saffi Naor, PS Tseng, Joydeep Ghosh, Dan Tian, Gouzhou Dong, Stephen Kurtzman, Norman Sadeh-Koniecpol, Bin Wang, Jianwen Su, Xiaoyang Wang, David Applegate, James Aspnes, Allan Heydon, Jay Siplestein, Steve Guattery, John Mount, Prasad Chalasani,

Kim Wagner, Benjamin Pierce, Qingming Ma, Bennet Yee, Jeffery Jackson, Qi Lu, Jim Stichnoth, Anja Feldmann, Dafna Talmor, Margaret Reid-Miller, Thomas Stricker, Sean Smith, Todd Wilson, Manpreet Khaira, Weiming Shen, Ye-Yi Wang, Jiři Sgall, Thomas Mathies, Dennis Grinberg, Brad Chen, Andrew Tomkins and many many others. I would like to express my special gratitude to Prasad Chalasani, Steve Guattery, Manpreet Khaira, John Mount, and Kim Wagner for carefully reading the draft of this thesis and providing invaluable comments which greatly improved the presentation of this thesis.

My deepest gratitude goes to Esther Ashby of USC, Sharon Burks and Anna Marie Mackuliak of CMU for their great effort to make my years (and many other students') as a graduate student smooth and pleasant. I would also like to thank Mary Lou and Ted Southworth, my host family in Los Angeles, who took care of me during my first year in LA, and introduced me to the cultural of this country, including the game of baseball.

Although their contributions to my graduate-school career and thesis are indirect, I would like to mention Ms. Bai Feng and Ms. Qiao-Xiang Li, my high school math teachers who introduced to me the fantasy of mathematics.

This thesis is dedicated to my parents, Gui-Xing Li and Zhan-Hong Teng for Twenty-seven years of love. Mom and Dad, especially in their most difficult years (1967–1975), came strong to support us and teach us the value of life. I wish they were here to share with me what is also belong to them. I owe a considerable debt of gratitude to my wife, Suzie, for her encouragement, understanding, and support. To my brothers Shang-Dong and Shang-Feng, for all the enjoyable time we shared, my deepest thanks.

Finally, I would like to thank the School of Computer Science at Carnegie Mellon University and Department of Computer Science at University of Southern California for their supports. I would also like to thank Xerox PARC for offering me my first job which certainly speedups the writing of this thesis, and to John Gilbert, Frances Yao, Marshall Bern, and Danny Sleator for making it possible. I would like to thank Marshall and John, especially for giving me permission of using some of the fine mesh figures they have collected.

The geometric characterization and separator results of this thesis represent a joint work with Gary Miller and Steven Vavasis. An extended abstract of this joint work will appear in the *32nd Annual Symposium on Foundations of Computer Science, IEEE.*

Contents

1 Introduction 1

2 Some Simple Geometric Graphs 7
- 2.1 Grid graphs in d dimensions . . . 7
- 2.2 Nearest neighborhood graphs . . . 8
- 2.3 Sphere packing graphs . . . 8
- 2.4 Graphs that can be drawn in a civilized manner . . . 9
- 2.5 Graphs of simplicial complex with bounded aspect-ratio . . . 9

3 Neighborhood Systems 11
- 3.1 Definitions . . . 11
- 3.2 Kissing numbers . . . 12
- 3.3 Two basic properties of neighborhood systems . . . 13
- 3.4 Neighborhood systems with bounded density . . . 16

4 Overlap Graphs 17
- 4.1 Graphs over neighborhood systems . . . 17
- 4.2 Inductivity of overlap graphs . . . 18
- 4.3 Why overlap graphs . . . 20
  - 4.3.1 Geometric embedding . . . 20
  - 4.3.2 Why they are special cases . . . 20
- 4.4 Why not dilation graphs . . . 21

5 Sphere Separators 23
- 5.1 Surface area of a sphere . . . 23
- 5.2 Density functions on a unit sphere . . . 24
- 5.3 Intersection number . . . 25
- 5.4 Why spheres but not hyperplanes . . . 26
- 5.5 Center points and a continuous separator theorem . . . 27
  - 5.5.1 Conformal maps and stereographic projections . . . 28
  - 5.5.2 Center points . . . 28
  - 5.5.3 A continuous separator theorem . . . 29
- 5.6 An algorithm for computing a sphere separator . . . 30

6 Separator Theorems 31
6.1 Separating intersection graphs . . . 31
6.1.1 The initial idea . . . 31
6.1.2 Local density functions . . . 32
6.1.3 Putting local density functions together . . . 33
6.1.4 Using norms: Vavasis's idea . . . 34
6.1.5 Relating to the intersection number . . . 34
6.1.6 Some basic lemmas . . . 35
6.1.7 Bound the total volume . . . 36
6.2 Separating overlap graphs . . . 38
6.2.1 A density function for overlap graphs . . . 38
6.2.2 The total volume . . . 38
6.2.3 A vertex separator from a continuous one . . . 40
6.3 Lower Bounds . . . 41

7 Graphs in Minkowski Spaces 45
7.1 Metric spaces . . . 45
7.2 Overlap graphs in Minkowski spaces . . . 46
7.3 Some basic geometric properties . . . 47
7.3.1 Linear transformation and the small separator property . . . 48

8 Center Points and Point Divisions 51
8.1 Introduction . . . 51
8.2 The complexity of TESTING CENTER . . . 54
8.3 The complexity of DIVISIBILITY . . . 55
8.3.1 A connector: the transportation problem . . . 56
8.3.2 Computing a Radon point . . . 59
8.4 Point division versus center points . . . 59
8.5 Approximating center points . . . 60
8.5.1 Computing an approximate center point in fixed dimensions . . . 60
8.5.2 A theory of Vapnik and Chervonenkis . . . 62
8.5.3 The probability of ϵ-good-samples . . . 63
8.5.4 Testing ϵ-center points . . . 63
8.5.5 An improved algorithm . . . 64
8.5.6 A deterministic algorithm . . . 66
8.6 Point divisions on the integer lattice . . . 67
8.6.1 An upper bound . . . 67
8.6.2 A lower bound . . . 68
8.6.3 On $4n$ points in the planar integer lattice . . . 69

9 Separator Based Divide and Conquer 71
9.1 The neighborhood query problem . . . 71
9.1.1 A separator based search structure . . . 72
9.1.2 A parallel construction . . . 75
9.2 Constructing intersection graphs . . . 76

9.2.1 A separator based algorithm 77
9.2.2 A deterministic algorithm 78
9.3 Constructing overlap graphs 79
9.4 Constructing k-nearest neighborhood graphs 80
9.5 Testing k-neighborhoodness 81

10 Final Remarks **83**
10.1 Great rings of the unit sphere 83
10.1.1 Great rings 84
10.1.2 The total area of great rings 85
10.2 The number of random bits and a deterministic algorithm 86
10.3 An elementary proof of the separator results 89
10.4 Open questions 90

List of Figures

1.1 Separators . . . 3
1.2 Hypercube and Butterfly . . . 4
1.3 The $2 \times m \times m$ grid . . . 5
1.4 A subgraph of a 12-nearest neighborhood graph in 2-space . . . 5

2.1 Grid graphs in 1, 2, and 3 dimensions . . . 7
2.2 A sphere packing graph . . . 8
2.3 Airplane wing (Barth and Jespersen) . . . 9
2.4 Simplices in 2 and 3 dimensions . . . 10
2.5 Lake Superior (Joe) . . . 10

3.1 A 2-neighborhood system . . . 11
3.2 Range of possible kissing numbers in fixed dimensions . . . 12
3.3 Kissing arrangements in 2 dimensions and 3 dimensions . . . 13
3.4 $\tau_d = A_d(1)$ . . . 13
3.5 The set of balls that touch p . . . 14
3.6 $||q_i - q_j|| \geq ||p - q_i||$ . . . 14
3.7 A neighborhood system whose density is 2 . . . 16

4.1 An intersection graph . . . 17
4.2 Rule for deciding edges in dilation graphs . . . 18
4.3 How to decide an edge in an overlap graph . . . 18
4.4 Intersection graph yields star graph . . . 18

5.1 GC and its p_{GS}, q_{GS} . . . 24
5.2 A sphere separator . . . 25
5.3 Why not hyperplane separators . . . 27
5.4 Why sphere separators . . . 27

6.1 Influential area of a kingdom and its density . . . 32
6.2 The common area of S and $2 \cdot B$. . . . 33
6.3 G_2 in one dimension and G_4 in two dimensions . . . 42
6.4 A 3-line graph . . . 43

7.1 An example of Minkowski distance . . . 46
7.2 Cubes and co-cubes . . . 47

7.3 The problem with unbounded aspect-ratio 48
7.4 All symmetrical convex figures can be mapped between the cube and the co-cube . . 49
7.5 The common area of S and $2 \cdot B$. 50

8.1 The Structure of Center Points 53
8.2 A 4-pencil 55
8.3 A 2 by m Transportation Problem 56
8.4 The change of ϵ_i and s_i 66

10.1 GC and its p_{GS}, q_{GS} 84
10.2 The great ring induced by B_i 85

Chapter 1

Introduction

In the past three decades, computer scientists have identified several algorithmic paradigms – general methods that are useful to design efficient algorithms not just for one or two problems, but for a wide variety of problems. Some of those paradigms such as dynamic programming, divide and conquer, branch and bound, *et cetera*, have been described in standard university text books (Aho, Hopcroft and Ullman [2], Knuth [54], Cormen, Leiserson and Rivest [18]). These general methods are usually very simple and intuitive at a high level, but they yield a large number of counterintuitive results that are fascinating from a purely mathematical viewpoint. For example, using divide and conquer one can multiply two matrices, two polynomials, or two integers faster than the "ordinary" high school algorithms (Strassen [81]; Schonhage and Strassen [77, 2]); one can solve a large classes of linear systems in $O(n^{1.5})$ arithmetic operations (George [35]; Lipton, Rose, and Tarjan [59]); one can sort optimally in $O(n \log n)$ time (Knuth [54]; Batcher [6]); and one can solve various geometry problems in almost linear time (Bentley [8]).

Despite the simplicity and generality on the high level of these methods, the details of each application may be very different. They depend crucially on the structure of each individual problem. For instance, each example above (namely sorting, matrix multiplying, linear system solving, and geometry computing) has its own distinct way of dividing a problem when using divide and conquer. How to systematically apply a general paradigm to a particular application is thus a major issue for algorithm design. To some extent, those algorithmic paradigms are too general to be discussed precisely. As noted by Bentley [8], the generality, ironically, is quite an obstacle to the systematic application of those paradigms.

In this thesis, I take a "mixed approach" by examining general paradigms within some well defined application domains. A more "middle of the road" paradigm is presented that can be precisely specified and yet can also be used to solve many problems in its domain of applicability. The general paradigm I will study is divide and conquer – one of the most widely used problem solving techniques [2]. The application domains I will mainly focus on are graph problems, especially those with important applications in numerical analysis and computational geometry.

At the heart of this paradigm is a notion of a small separator, roughly speaking, a relatively small set of vertices or edges whose removal divides the rest of the graph into two disconnected pieces of approximately equal size. An important contribution of this thesis is a novel geometric characterization of graphs with small separators. By taking advantage of the underlying geometric structure, I also develop an efficient algorithm for finding such a small separator.

The separator results come with some new applications, epsecially in numerical analysis and computational geometry. In numerical analysis, the new separator result can be used not only in Nested Dissection for generating an almost optimal elimination ordering for the sparse Cholesky factorization but also for devising efficient sequential and parallel algorithms for mapping sparse matrices onto various distributed processor/memories parallel systems to reduce communication cost and to achieve load balance [59, 39, 72, 35, 36, 42]. In computational geometry, a separator based divide and conquer paradigm can be developed and used to solve several geometry problems. This paradigm outperforms a commonly used one in computational geometry, the multi-dimensional divide and conquer of Bentley [8], on various applications. It provides a new way to design efficient parallel algorithms for geometry problems in fixed dimensions.

Divide and conquer

Divide and conquer solves a problem by partitioning it into a set of independent subproblems, finding solutions for the subproblems recursively, and then combining the solutions for the subproblems into a solution for the whole. In order to apply divide and conquer efficiently, it is necessary that all subproblems are of the same type as the original one, but substantially smaller. The cost of combining solutions of subproblems should also be low. There is, in general, a trade-off between the size of the subproblems and the cost of combining.

Associated with each application of divide and conquer is a tree, a *partition-tree*, whose root represents the whole problem, the children of the root denote the top level subproblems, and other nodes denote the subproblems which are recursively defined in a top-down fashion during the process of divide and conquer. It is interesting to point out that, in contrast to the way that the partition tree is defined, the computation in a divide and conquer algorithm is usually performed in a bottom-up fashion – computation associated with the leaves is first performed, then the parents of leaves and so on till the root is reached. This is where parallelism can be exploited: computation at each level of a partition tree can be performed in parallel. The height of the partition tree together with the parallel time of combining solutions of subproblems thus represents the parallel time complexity of a divide and conquer algorithm. Divide and conquer is an effective paradigm not only for sequential but also for parallel computation [37, 68, 73].

Graph partition and small separators

When specialized to graph problems, divide and conquer partitions the whole problem by deleting a subset of vertices or edges. The subproblems are those associated with connected components in the resulting graph. The process of deleting edges and vertices is called *graph partition*, and the subset of vertices and edges removed are respectively called a *vertex-* and an *edge-separator*.

Usually, the cost of combining is determined by the *size* of separators, i.e., the number of vertices or edges removed. The depth of a partition tree depends on the quality of separators used in divide and conquer. Therefore, to achieve high efficiency, the size of subproblems and size of separators need to be balanced. This motivates the following definition:

Definition 1.1 (Separators[1]) *A subset of vertices C of a graph G with n vertices is an $f(n)$-*

[1] In this thesis, we are mainly concern with vertex separators. Edge separators can be defined similarly.

separator *that* **δ-splits** *if* $|C| \leq f(n)$ *and the vertices of* $G - C$ *can be partitioned into two sets* A *and* B *such that* $|A|, |B| \leq \delta n$ *and there is no edge between* A *and* B, *where* f *is a positive function and* $0 < \delta < 1$ *(See Figure 1.1).*

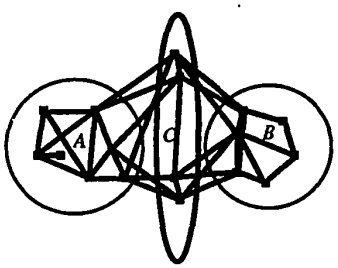

Figure 1.1: Separators

For simplicity, by saying a graph G has a *small separator*, we mean that there is a constant δ and a sublinear function f, such that G has an $f(n)$-separator that δ-splits. We say a class of graphs has an $f(n)$-*separator theorem* if there exist constants $\delta < 1$ and $\beta > 0$ such that every graph in the class has a $\beta f(n)$-separator that δ-splits.

Small separator results for a family of graphs closed under the subgraph operation immediately leads to divide and conquer recursive algorithms for a variety of applications. In general, the efficiency of such algorithms depends on a δ bounded away from 1 and an $f(n)$ that grows slowly (e.g., $f(n) = n^\alpha$ for some constant $0 < \alpha < 1$).

Previous results on separators

Clearly, not every graph has a small separator. An example is the complete graph of n vertices. One may quickly observe that the complete graph has $\frac{n(n-1)}{2}$ edges, and conjecture that all "sparse" graphs, or more specifically, all graphs with bounded degree have small separators.

A beautiful construction of Thompson[2] [85] refutes such a conjecture. Using an information transfer argument, he showed that each network which is capable of computing the discrete Fourier transform on N elements in T steps has no separator of size $\Omega(N/T)$. His result implies that neither butterfly graphs, nor shuffle-exchange graphs, nor hypercube graphs (see Figure 1.2) have a separator of size $o(n/(\log n))$.

Then which families of graphs do have small separators?

Two of the most well-known families of graphs that have small separators are trees and planar graphs. It is known that a tree has a single vertex separator that 2/3-splits. Lipton and Tarjan [60] proved that every planar graph has a $\sqrt{8n}$-separator that 2/3-splits. Their result improved an earlier one by Peter Ungar [90]. Some extensions of their work have been made [21, 64, 31, 33], and separator theorems have also been obtained for graphs with bounded genus [38, 46] and graphs with bounded excluded minor [5]. In particular, Gilbert, Hutchinson, and Tarjan showed that all graphs

[2] As a matter of fact, an earlier result of Erdös, Graham, and Szemerédi [24] has shown that for c a large enough constant, almost all graphs of cn edges cannot be separated into small components without removing $\Omega(n)$ vertices.

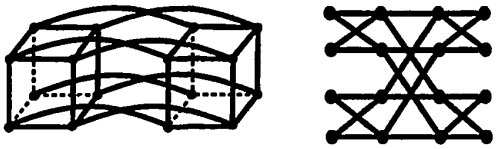

Figure 1.2: Hypercube and Butterfly

with genus bounded by g have an $O(\sqrt{gn})$-separator, and Alon, Seymour, and Thomas proved that all graphs with an excluded minor isomorphic to the h-clique have an $O(h^{3/2}\sqrt{n})$-separator.

Interestingly, all characterizations above are combinatorial as well as their proofs!

Classical applications of separator results

Perhaps, the most classical application of small separator results is Nested Dissection – a widely technique for solving a large classes of sparse linear systems. This approach was pioneered by Alan George [35] who designed the first $O(n^{1.5})$-time nested dissection algorithm for linear systems on regular grids using the fact that the $\sqrt{n} \times \sqrt{n}$ grid has a $\sqrt{n}$-separator. His result was extended to planar linear systems by Lipton, Rose, and Tarjan [59]. Gilbert and Tarjan [39] examined several variants of the nested dissection algorithms, and Pan and Reif [72] demonstrated that nested dissection can be implemented efficiently in parallel.

Small separator results have found fruitful applications in VLSI design (Leiserson [56]; Leighton [55]; Valiant [92]) and efficient message routing (Fredrickson and Janardan [27]). They have also been used in proving several complexity-theoretic results (Paterson [71]; Lipton and Tarjan [61]; Li [57]; Teng[82]), and certainly have been used to design efficient graph algorithms such as parallel construction of breadth-first-search trees (Pan and Reif [72]) and depth-first-search trees (Aggarwal, Anderson, and Kao [1]; Kao [51]; Smith [79]; He and Yesha [45]), testing graph isomorphism (Gazit [32]), parallel evaluating expressions (Gazit, Miller, and Teng [34]; Miller and Teng [65]), and approximating NP-complete problems (Lipton and Tarjan [61]).

A new challenge

The development of computational geometry and numerical analysis calls for deeper understanding of the separator property for graphs embedded in fixed dimensional spaces, especially in 2 and 3 dimensions. Although the planar separator theorem is applicable to many interesting families of graphs embedded in 2-space, several natural classes of graphs in 2-space, including the class of k-nearest neighborhood graphs, are neither planar, nor have bounded genus, nor have bounded excluded minor. In general, none of the above separator theorems are useful for graphs in 3-space.

Let us first look at some simple examples. Figure 1.3 shows a $2 \times m \times m$ grid in 3 dimensions. It is not hard to see that this graph has a separator of size $2\sqrt{n}$ that 1/2-splits, where $n = m^2$. But it has genus $\Omega(\sqrt{n})$ and a minor isomorphic to the $\sqrt{n}$-clique. To see this, we identify vertices at the top layer by row and the bottom layer by column. The resulting graph is a complete bipartite graph with $\sqrt{n}$ vertices on each side. Therefore, the $2 \times \sqrt{n} \times \sqrt{n}$ grid has genus at least $\Omega(\sqrt{n})$ and a minor isomorphic to the $\sqrt{n}$-clique. Thus none of the separator results in the previous section

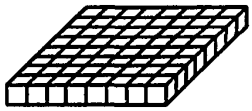

Figure 1.3: The $2 \times m \times m$ grid

can be directly applied to obtain a small separator.

Similar problems also occur in 2-dimensions. If we shift the top layer of the above graph to right by 1/2 unit and to front by 1/2 unit, and project the top layer to the bottom one, we then obtain a graph shown in Figure 1.4.

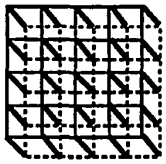

Figure 1.4: A subgraph of a 12-nearest neighborhood graph in 2-space

This graph is a subgraph of a 12-nearest neighborhood graph. Clearly, it has an $O(\sqrt{n})$-separator that 1/2-splits. But it also has genus $\Omega(\sqrt{n})$ and a minor isomorphic to the $\sqrt{n}$-clique, because this graph, combinatorially, is the same graph as the $2 \times \sqrt{n} \times \sqrt{n}$ grid graph.

Motivated by the potential application of separator results in numerical analysis, especially in domain decomposition and the finite element method, several groups of researchers (Vavasis [94]; Miller and Thurston [67]; Miller and Vavasis [69]) have proposed classes of graphs that can be embedded in d dimensions and that have $O(n^{\frac{d-1}{d}})$ separators. For the applications mentioned in the last section, $d = 2$ and $d = 3$ are the interesting cases, in which case the bounds are $O(n^{1/2})$ and $O(n^{2/3})$ respectively.

All of those earlier classes of graphs have the disadvantage that, when specialized to two dimensions, they do not contain all planar graphs. This is a serious drawback because the earliest and best-known separator result is Lipton and Tarjan's theorem [60] that all planar graphs have $O(n^{1/2})$ separators. As a matter of fact, those classes do not even contain any graph with unbounded degree and consequently they do not contain simple geometric graphs such as k-nearest neighborhood graphs. Moreover, the characterizations of those classes are quite complicated.

This thesis

In this thesis, a unified geometric approach is presented for graph separators. A new class of geometric graphs, *overlap graphs*, is proposed. This class is defined based on elementary geometric concepts such as points, balls, cubes, and spheres, and it has the following properties:

1. In two dimensions, planar graphs are special cases of overlap graphs.
2. In d dimensions, a finite subgraph of the infinite d-dimensional grid graph is an overlap graph.
3. Such a graph of n vertices in d-dimensions has an $O(n^{\frac{d-1}{d}})$ separator.

To my knowledge, this is the first time that a class of graphs has been proposed with these three very natural properties. The proof that planar graphs are special cases of overlap graphs relies on recent deep theorems by Andreev and Thurston characterizing all planar graphs in a novel geometric fashion. A consequence is a new geometric proof of a classical theorem of Lipton and Tarjan that all planar graphs have an $O(\sqrt{n})$-separator. This is another beautiful illustration of the use of geometry in understanding combinatorial concepts.

Moreover, the class of overlap graphs in d-space includes as a special case the class of k-nearest neighborhood graphs. Therefore, each k-nearest neighborhood graph of n points in d-space has an $O(k^{\frac{1}{d}}n^{\frac{d-1}{d}})$-separator. It is worthwhile to mention that the above separator bound is also optimal in terms of k. In addition. overlap graphs include some earlier classes of graphs associated with the finite element method, such as simplex graphs with bounded aspect-ratio (Miller and Thurston [67]) and density graphs (Miller and Vavasis [69]). It also contains various other classes of graphs defined geometrically, including sphere packing graphs and graphs that can be drawn with finite resolution. All the arguments are based on geometric properties of embedding. The separator bounds come with randomized linear-time and randomized NC algorithms.

This thesis is organized as follows: Chapter 2 shows examples of some classes of geometric graphs. The notion of a *k-neighborhood system* is introduced in chapter 3, along with some of its basic properties. In chapter 4, three classes of graphs are defined based on neighborhood systems. They are respectively called intersection graphs, dilation graphs, and overlap graphs. Each type starts with a set of balls (one per vertices) and extracts a graph from these. I will show that intersection graphs and overlap graphs are "sparse" while the class of dilation graphs contains complete graphs as special cases. I will also show that the class of overlap graphs contains as special cases all planar graphs, all k-nearest neighborhood graphs, and some earlier classes of graphs associated with domain decomposition and finite element methods. In chapter 5, a notion of a sphere separator for neighborhood systems is introduced. Two cost functions associated with a sphere separator: the weighted surface area and the intersection number is discussed. It is shown that the intersection number is closely related with the size of separators for intersection graphs. Also presented in this chapter are the continuous separator theorem of Miller-Thurston [67] and an algorithm for computing a sphere separator with a low cost. In chapter 6, the main separator results of this thesis are presented and proved. A matching lowerbound on the size of separator of overlap graphs is also given. In chapter 7, the notion of neighborhood systems is generalized from Euclidean space to a more general family of metric spaces known as *normed linear spaces* or *Minkowski spaces.* All the earlier separator results will be generalized to intersection graphs and overlap graphs in the more general setting. In chapter 8, some complexity-theoretic results about the structure of center points and point divisions are presented. Efficient sequential and parallel algorithms are given for computing center points in fixed dimensions. In chapter 9, a divide and conquer paradigm based on the new separator results is presented and used in designing efficient algorithms for various geometry problems in fixed dimensions. In chapter 10, I will bound the number of random bits required in the randomized algorithm for computing a small separator of an overlap graph, based on which I will derive a deterministic polynomial time algorithm for computing a small separator. Finally, I will conclude this thesis by listing a set of open questions.

Chapter 2

Some Simple Geometric Graphs

Geometric graphs, as the name suggests, are those which are defined by the relation of a set of elementary geometric entities such as points, lines, hyperplanes, spheres, cubes, etc. Often they are naturally defined. They have rich structure and find importance in various branches of science and art. Their simple geometric character makes them easy to visualize and reason about.

2.1 Grid graphs in d dimensions

Perhaps, the simplest class of such graphs is the class of grid graphs. They are the underlying graphs for the finite difference method in numerical analysis.

To construct a d-dimensional $m \times \cdots \times n$ *grid graph*, we take as vertices those points $(x_1, \ldots, x_d)$ of $\mathbb{R}^d$ all of whose coordinates are nonnegative integers with $x_i < m$ for all $1 \leq i \leq d$, and join each point by an undirected line segment to each point which is a unit distance away. These connecting segments, which represent the edges of the grid graph, run parallel to one of the coordinate axes of $\mathbb{R}^d$. Grid graphs in 1, 2, and 3 dimensions are illustrated in Figure 2.1.

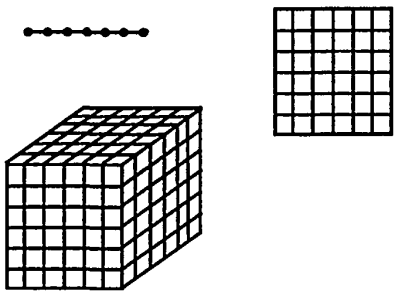

Figure 2.1: Grid graphs in 1, 2, and 3 dimensions

It is not hard to see that a d-dimensional $m \times \cdots \times m$ grid graph has an $n^{\frac{d-1}{d}}$-separator that 1/2-splits, where $n = m^d$. To split the graph evenly, we just choose an $i : 1 \leq i \leq d$, and remove all points with $x_i = \lceil n/2 \rceil$. It is also not hard to see that the above separator size is tight for grid graphs.

2.2 Nearest neighborhood graphs

Another interesting class is the *nearest neighborhood graphs* – an important class of graphs from computational geometry.

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$. For each $p_i \in P$, let $N(p_i)$ be a closest neighbor of p in P, where ties are broken arbitrarily. Similarly, for any integer k, let $N_k(p_i)$ be the set of k nearest neighbors of p_i in P; here too ties are broken arbitrarily.

Definition 2.1 *A* **k-nearest neighborhood graph** *of $P = \{p_1, \ldots, p_n\}$ in $\mathbb{R}^d$, is a graph with vertices $V = \{1, ..., n\}$, and edges E,*

$$E = \{(i,j) | p_i \in N_k(p_j) \text{ or } p_j \in N_k(p_i)\}.$$

Notice that each d-dimensional grid graph is basically a $2d$-nearest neighborhood graph defined by the set of grid points. The only exceptions occur on the boundary.

2.3 Sphere packing graphs

We can also define graphs based on a set of spheres rather than points, where each sphere S has a center p and a radius r, and S is the set of points in $\mathbb{R}^d$ whose Euclidean distance away from p is r. A sphere partitions $\mathbb{R}^d$ into three subsets, $\text{ext}(S)$, $\text{int}(S)$, and S itself, respectively, those points in the exterior of S, in the interior of S, and on S.

Let a *sphere packing* in d dimensions (Conway and Slone [19]) be a set of spheres having disjoint interiors. Each sphere naturally packing defines a graph, called its *sphere packing graph*, where each vertex represents a sphere; There is an undirected edge connecting two vertices if their corresponding spheres have a common point. Figure 2.2 shows an example of sphere packing graphs in two dimensions.

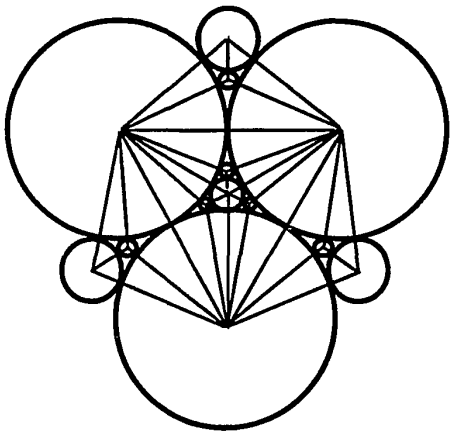

Figure 2.2: A sphere packing graph

2.4 Graphs that can be drawn in a civilized manner

Sometimes, a geometric graph is not directly given by points and segments in a space. Instead it can be graph obtained from a drawing. For example, planar graphs can be viewed this way: they are graphs which can be drawn in 2-space with no two edges crossing. Other graphs may be those that can be drawn in a predefined resolution – in a "civilized manner".

For each pair of positive reals $r > 0$ and $s > 0$, a graph G can be drawn in $\mathbb{R}^d$ in an (r, s)-*civilized manner* if its vertices can be embedded in $\mathbb{R}^d$ so that

1. The length of each edge is $\leq r$.

2. The distance between any two points in at least s.

This class of graphs was originated by Peter Doyle and Laurie Snell [22]. Also see Vavasis [94].

2.5 Graphs of simplicial complex with bounded aspect-ratio

This section presents a class of graphs motivated by the finite element method [80] – a collection of numerical techniques for approximating a continuous problem by finite structures. To approximate a continuous function, the finite element method subdivides the domain (a subset of $\mathbb{R}^d$) into a *mesh* of polyhedral *elements* (Figure 2.3), and then approximates the continuous function by a piecewise polynomial on the elements.

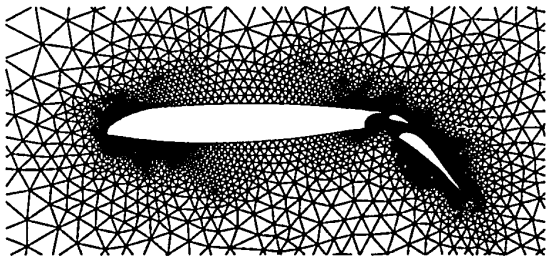

Figure 2.3: Airplane wing (Barth and Jespersen)

Each element is a d-dimensional *simplex* which is the convex hull of $(d+1)$ affinely independent points in $\mathbb{R}^d$, e.g., a triangle in 2 dimensions and a tetrahedra in 3 dimensions (Figure 2.4). A mesh of d-dimensional simplices is called a d-dimensional *simplicial complex* which is a collection of d-dimensional simplices closed under subsimplex and intersection [80, 9, 67] (see Figure 2.3).

Associated with each simplicial complex K is natural graph, the 1-skeleton of K. For example, the 1-skeleton of each 2-dimensional simplicial complex is a planar graph. Conversely, every planar graph can be embedded in the plane such that each edge is mapped to a straight line segment (Fáry [25]; Tutte [86, 87]; Thomassen [84]; Fraysseix, Pach, and Pollack [26]).

To properly approximate a continuous function, in addition to the conditions that a mesh must conform to the boundaries of the region and be fine enough, each individual element of the mesh

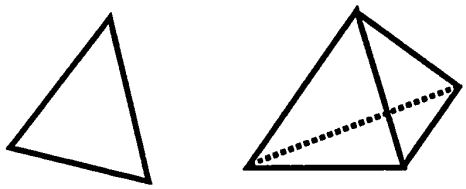

Figure 2.4: Simplices in 2 and 3 dimensions

must be *well shaped*, i.e., they must have bounded *aspect ratio*[1]. The aspect ratio is defined as follows.

Let the *diameter* diameter(S) of a d-simplex S be the maximum distance between any pair of points in S. Define the *aspect-ratio* to be

$$\alpha = \frac{\text{diameter}(S)}{\sqrt[d]{\text{volume}(S)}},$$

where volume(S) denotes the d-dimensional volume of S. An example of a simplicial complex with bounded aspect ratio is illustrated in Figure 2.5. Another one is shown in Figure 2.3.

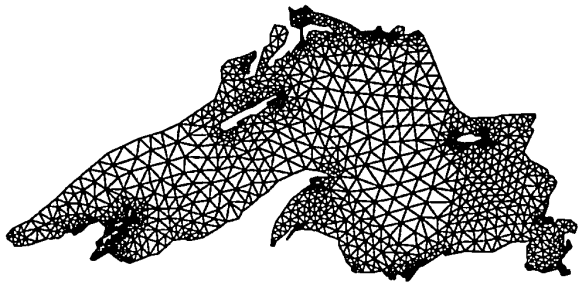

Figure 2.5: Lake Superior (Joe)

[1]The condition of bounded aspect ratio is also closely related with some other well shaped conditions of an element, namely, the *no small angle* condition and the *no obtuse angle* condition (Bern, Eppstein, and Gilbert [9])

Chapter 3

Neighborhood Systems

Although they appear very different, graphs defined in Chapter 2 have a lot in common. These similarities motivate a unified definition, where each example class can be shown as a special case. The general definition allows us to study various graph-theoretical properties of those graphs as a whole. Especially, it enables us to develop a unified method to prove small separator theorems. The basic notion behind the general definition is neighborhood systems.

3.1 Definitions

A *neighborhood* of a given point $p \in \mathbb{R}^d$ is a closed compact set containing p (p is called the *center* of its neighborhood). The type of neighborhood concerned in this chapter is the *Euclidean neighborhood* which is a closed ball of certain radius centered at p. More general neighborhoods will be discussed in Chapter 7.

A *neighborhood system* $\Xi = \{B_1, \ldots, B_n\}$ is a finite collection of neighborhoods. Let p_i be the center of B_i $(1 \leq i \leq n)$. $P = \{p_1, \ldots, p_n\}$ is called *centers* of Ξ. For any integer k, Ξ is a *k-neighborhood system* if for all $1 \leq i \leq n$, the interior of B_i contains no more than k points from P. A 2-neighborhood system in two dimensions is illustrated in Figure 3.1.

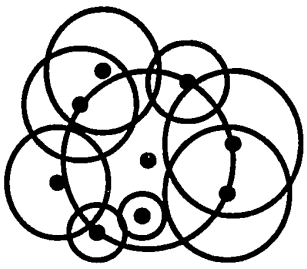

Figure 3.1: A 2-neighborhood system

The following notation will be used throughout the thesis. For each positive real α, suppose B is a ball of radius r in $\mathbb{R}^d$, then $\alpha \cdot B$ denotes the ball with the same center as B but radius αr.

d	τ_d
1	2
2	6
3	12
4	24-25
5	40-46
6	72-82
7	126-140
8	240
9	306-380
10	500-595
11	582-915
12	840-1416
13	1130-2233
14	1582-3492
15	2564-5431
16	4320-8313
17	5346-12215
18	7398-17877
19	10668-25901
20	17400-37974
21	27720-56852
22	49896-86537
23	93150-128096
24	196560

Figure 3.2: Range of possible kissing numbers in fixed dimensions

3.2 Kissing numbers

It is handy to precisely specify the constant term in the separator theorem. To this end, I will use a concept called the *kissing number* in d dimensions, denoted by τ_d, which is the maximum number of nonoverlapping unit balls in $\mathbb{R}^d$ that can be arranged so that they all touch a central unit ball [19]. It is easy to see that $\tau_1 = 2$ and $\tau_2 = 6$ (see Figure 3.3 for τ_2). A little imagination or even by looking at the arrangement of apples in a supermarket, will convince you that $\tau_3 \geq 12$ (also see Figure 3.3). But is 12 is the precise number? Can it be 13? Do not think this is an easy problem – even Isaac Newton and David Gregory disagreed. Back in 1694, Newton believed the answer was 12, while Gregory thought that 13 might be possible.

The correct answer to this question is now known to be 12. Various proofs appeared in the literature (For citations, we refer readers to [19]). The kissing number becomes more difficult to evaluate in higher dimensions. The best bounds currently known on τ_d, are shown in Figure 3.2. It is quite surprising that we know the kissing numbers in 8 and 24 dimensions but no of other dimensions above 3, partly, because the optimal arrangement in 8 and 24 dimensions is unique [19].

Although there is no explicit formula known about kissing numbers, τ_d can be bounded from above and below by the following inequalities.

$$2^{0.2075..d(1+o(1))} \leq \tau_d \leq 2^{0.401d(1+o(1))}$$

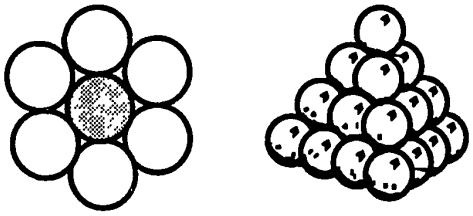

Figure 3.3: Kissing arrangements in 2 dimensions and 3 dimensions

The first inequality was given by Kabatiansky and Levenshtein [49] and the second one by Wyner [96].

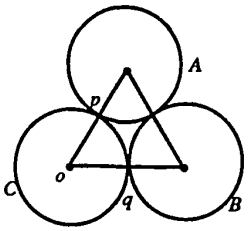

Figure 3.4: $\tau_d = A_d(1)$

For each positive real δ, define $A_d(\delta)$ to be the maximum number of points that can be arranged on a unit $(d-1)$-sphere (the boundary of a unit d-ball), such that the distance between each pair of points is at least δ. As shown in Figure 3.4, if A, B, and C are unit balls such that A and B touch C at points p and q, respectively, then $||p-q|| \geq 1$ if and only if A and B are not overlapping. Hence,

Lemma 3.1 $A_d(1) = \tau_d$.

3.3 Two basic properties of neighborhood systems

Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system in d dimensions. For each $p \in \mathbb{R}^d$, define $\text{density}_\Xi(p)$ to be the number of neighborhoods in Ξ that contain p.

The following lemma bounds the maximum possible density of a k-neighborhood system.

Lemma 3.2 (Density Lemma) *For each k-neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ in d dimensions, for each $p \in \mathbb{R}^d$, $\text{density}_\Xi(p) \leq \tau_d k$.*

Proof: We first prove the lemma for the case when $k = 1$, namely, no ball contains the center of other balls in its interior.

Without loss of generality, let $\{B_1, \ldots, B_t\}$ be the set of all balls that contain p. Further, assume p is on the boundary of all balls $B_1, \ldots, B_t$ (See Figure 3.5), for otherwise, B_i can be replaced by a ball C_i centered at p_i with radius $||p_i - p||$ and the assumption of the Lemma is still satisfied because $C_i \subset B_i$.

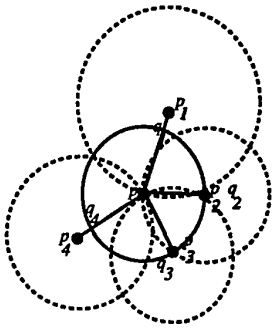

Figure 3.5: The set of balls that touch p

Let $\delta = \min\{||p - p_i|| : 0 \leq i \leq t\}$. By proper linear dilation, assume $\delta = 1$. Let S_p be the sphere centered at p with radius $\delta = 1$. Let q_i be the intersection of the ray pp_i with the sphere S_p. We claim that for each pair $i, j \in \{1, ..., t\}$, $||q_i - q_j|| \geq \delta$.

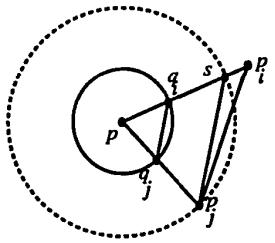

Figure 3.6: $||q_i - q_j|| \geq ||p - q_i||$

Without loss of generality, assume $||p - p_i|| \geq ||p - p_j||$. Let s be a point on the ray pp_i such that $||p - s|| = ||p - p_j||$ (see Figure 3.6). It follows $||p - p_i|| = ||p - s|| + ||s - p_i||$.

By the triangle inequality, we have $||s - p_i|| + ||s - p_j|| \geq ||p_i - p_j||$.

Because $p_j \notin B_i$, and the radius of B_i is $||p - p_i||$, we have $||p - p_i|| \leq ||p_i - p_j||$. Thus $||p - s|| \leq ||s - p_j||$.

By the similarity of triangles Δpq_iq_j and Δpsp_j, we have $||q_i - q_j|| \geq ||p - q_i|| = \delta$. By Lemma 3.1, $t \leq \tau_d$, completing the proof of the lemma when $k = 1$.

We now prove the lemma for any $k > 1$. Without loss of generality, assume $B_1, \ldots, B_t$ contain p. Define a subset Q of $\{p_1, \ldots, p_t\}$ by the following procedure. Initially, let $P = \{p_1, \ldots, p_t\}$ and $Q = \emptyset$.

while $P \neq \emptyset$

1. Suppose q is the point in P with the largest $||q - p||$, let $Q = Q \cup \{q\}$;
2. Let $P = P - \text{int}(B_q)$, (where B_q stands for the closed ball centered at q).

Because no ball contains more than k points from $\{p_1, \ldots, p_t\}$ in its interior, we have $m = |Q| \geq \lceil t/k \rceil$.

We now show that for all $q \in Q$, $\text{int}(B_q) \cap Q = \{q\}$. Suppose $Q = \{q_1, \ldots, q_m\}$ such that for all $i < j$, q_i is put Q in the above procedure before q_j. Notice that for all $j > i$, $q_j \notin \text{int}(B_{q_i})$. Also because $||q_i - q_j|| \geq ||q_i - p|| \geq ||q_j - p||$, we have for all $i < j$, $q_i \notin \text{int}(B_{q_j})$. So $\text{int}(B_q) \cap Q = \{q\}$. Thus $m \leq \tau_d$ which implies $t \leq \tau_d k$. □

The following lemma gives an upper bound on the number of balls in a neighborhood system which intersect a given ball.

Lemma 3.3 (Ball Intersection I) *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a k-neighborhood system in d dimensions. Then for each ball $B \in \mathbb{R}^d$ (with center p and radius r), and for each α,*

$$|\{i : B_i \cap B \neq \emptyset \text{ and } p_i \notin \alpha \cdot B\}| \leq A_d(\frac{\alpha - 1}{\alpha})k,$$

where p_i is the center of B_i $(1 \leq i \leq n)$.

Proof: This lemma can be proved in a fashion similar to the Density Lemma. So, I present a high level argument.

Again, we first prove the lemma in the case when $k = 1$. Without loss of generality, assume $B_1, \ldots, B_t$ is the set of all balls with centers not in $\alpha \cdot B$ that intersect B. Also, by proper dilation, assume $r = 1$. Let $\delta = \min\{||p - p_i|| : 0 \leq i \leq t\}$. We know $\delta \geq \alpha r$. By an elementary geometrical argument, it can be shown that the set of balls $B_1, \ldots, B_t$ induces a set of balls $C_1, \ldots, C_t$ with the following properties:

1. The center o_i of C_i is on the sphere of radius δ centered at p.
2. C_i kisses B; and
3. $\{C_1, \ldots, C_t\}$ forms a 1-neighborhood system.

For each i, let q_i denote the common point of C_i and B. By a simple geometrical argument, we can show that for each pair $1 \leq i, j \leq t$, the distance between q_i and q_j is at least $\frac{\alpha-1}{\alpha}$, and hence $t \leq A_d(\frac{\alpha-1}{\alpha})$. So the lemma holds when $k = 1$.

Using a similar construction as in the proof of Lemma 3.2, we can prove the lemma for all positive integer $k \geq 1$. □

Similarly,

Lemma 3.4 (Ball Intersection II) *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a k-neighborhood system in d dimensions. Then for each ball $B \in \mathbb{R}^d$ (with center p and radius r),*

$$|\{i : B_i \cap B \neq \emptyset,\ r_i \geq r,\ \text{and } p_i \in 2 \cdot B\}| \leq A_d(\frac{1}{2})k,$$

where p_i is the center and r_i the radius of B_i $(1 \leq i \leq n)$.

The following is an interesting consequence of the Density Lemma.

Corollary 3.5 *The degree of all k-nearest neighborhood graphs in $\mathbb{R}^d$ is bounded above by $\tau_d k$ and this bound is tight.*

Proof: For each point p_i in a given set P, let B_i be the largest ball centered at p_i such that the interior of B_i contains no more than k points from P. Notice that if (p_i, p_j) is an edge in a k-nearest neighborhood graph of P, then either $p_i \in B_j$ or $p_j \in B_i$. By Lemma 3.2, we have the degree of k-nearest neighborhood graphs is bounded above by $\tau_d k$. □

3.4 Neighborhood systems with bounded density

Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system in d dimensions. Define the *density* of Ξ, denoted by density(Ξ), to be the maximum density over all points in $\mathbb{R}^d$.

By the Density Lemma, each k-neighborhood system in d dimensions has density bounded by $\tau_d k$. However, as shown in Figure 3.7, there is a neighborhood system with density equal to 2, which has a ball whose interior contains n centers.

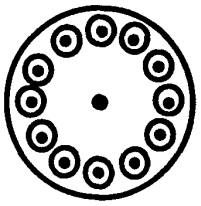

Figure 3.7: A neighborhood system whose density is 2

Also, by the Density Lemma, if a d-dimensional neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ can be partitioned into g disjointed k-neighborhood systems $\Xi_1, \ldots, \Xi_g$, then $density(\Xi) \leq \tau_d g k$. It is interesting to know whether the converse of the above statement is true, namely, whether each neighborhood system with density μ can be partitioned into $O(\mu)$ 1-neighborhood systems. Althought, using the following lemma (Lemma 3.6), we can show that each neighborhood system with density μ can be partitioned into $O\left(\mu + \frac{\log n}{\log \mu - \log(\mu - 1)}\right)$ 1-neighborhood systems.

Lemma 3.6 (Ball Intersection) *For each positive integer μ, for each d-dimensional neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ with density(Ξ) $\leq \mu$, for each d-dimensional ball B with radius r, $|\{i : B_i \cap B \neq \emptyset \text{ and } r_i \geq r\}| \leq 3^d \mu$.*

Proof: Without loss of generality, let $B_1, \ldots, B_t$ be the set of all balls in Ξ of radius at least r, that intersect B. For each $i : 1 \leq i \leq t$, if p_i, the center of B_i, is in $2 \cdot B$, let B_i' be the ball of radius r centered at p_i; if p_i is not in $2 \cdot B$, let p_i' be the point common of the ray pp_i and the boundary of $2 \cdot B$, and let B_i' be the ball centered at p_i' and of radius r. In either case, $B_i' \in B_i$ and B_i' intersects B, and if B_i is replaced by B_i', the density of the resulting neighborhood system does not increase. Therefore, without loss of generality, assume for all $i : 1 \leq i \leq t$, B_i has radius r and with center $p_i \in 2 \cdot B$. Therefore, each ball B_i $(1 \leq i \leq t)$, is contained in the ball $3 \cdot B$. Because the density of the neighborhood system is bounded by μ,

$$\sum_{i=1}^{t} \text{volume}(B_i) \leq \mu \text{volume}(3 \cdot B),$$

which implies $t \leq 3^d \mu$. □

Chapter 4

Overlap Graphs

In this chapter, the class of overlap graphs are defined based on neighborhood systems. The goal is to obtain a geometric characterization of graphs with small separators. In addition, simplicity and generality are also the main objectives.

4.1 Graphs over neighborhood systems

Three classes of graphs are defined based on neighborhood systems. Each class defines the same set of vertices when given a neighborhood system, i.e., one vertex per ball. The only difference is the rule of determining when to put an edge between two neighborhoods.

Given a neighborhood system $\Xi = \{B_1, \ldots, B_n\}$, the simplest rule to decide edges is probably to connect two neighborhoods whenever they intersect. The graph so defined is called the *intersection graph* of Ξ. The intersection graph of a 2-neighborhood system in 2-dimension is illustrated in Figure 4.1.

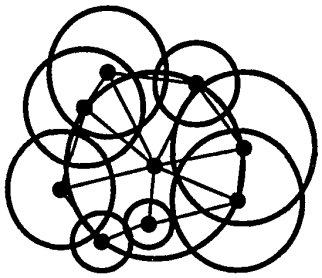

Figure 4.1: An intersection graph

An alternative rule is to connect two neighborhoods B_i and B_j if for a predefined positive real $\alpha \geq 1$, the pair of balls $(\alpha \cdot B_i)$ and $(\alpha \cdot B_j)$ have a common point (see Figure 4.2 for the rule). The resulting graph is called the *α-dilation graph* of Ξ. Clearly, each intersection graph is the 1-dilation graph of its neighborhood system.

A minor modification of the rule for dilation graphs yields overlap graphs. The new rule connects two neighborhoods B_i and B_j by an edge whenever the α dilation of the smaller ball

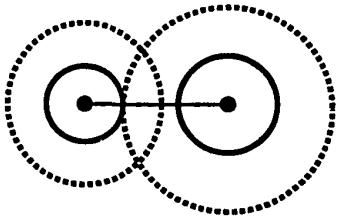

Figure 4.2: Rule for deciding edges in dilation graphs

intersects the larger one (see Figure 4.3), or equivalently, if $(B_i \cap (\alpha \cdot B_j) \neq \emptyset)$ and $((\alpha \cdot B_i) \cap B_j \neq \emptyset)$. The graph so defined is called the *α-overlap graph* of Ξ.

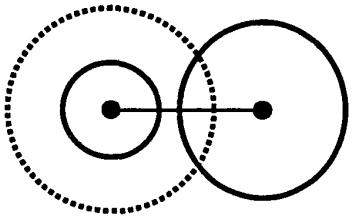

Figure 4.3: How to decide an edge in an overlap graph

4.2 Inductivity of overlap graphs

The first question one might ask is whether an overlap graph is "sparse" (e.g., having a linear number of edges). If the maximum degree of overlap graphs were bounded, then the number of edges would also be bounded. However, it is not hard to see that some 1-intersection graphs may have unbounded maximum degree. For example, each star graph is isomorphic to the intersection graph of some 1-neighborhood system (see Figure 4.4).

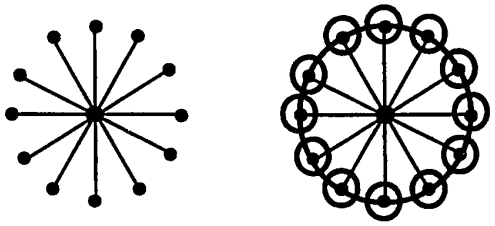

Figure 4.4: Intersection graph yields star graph

For any integer δ, a graph is *δ-inductive* if its vertices can be numbered such that each vertex has at most δ edges to higher numbered vertices. Clearly, a δ-inductive graph with n vertices has at most $(\delta \cdot n)$ edges.

For example, every tree is 1-inductive and each planar graph is 5-inductive. To see the later fact, observe that each planar graph has at least one vertex of degree less than 6 (by Eular's formula). So a 5-inductive numbering can be obtained by assigning the smallest number to such a vertex and inductively numbering other vertices. The following lemma will be used to show the intersection graph of a k-neighborhood system is $2A_d(1/2)$-inductive.

Lemma 4.1 *Let $\Xi = \{B_1, \ldots, B_n\}$ be a d-dimensional k-neighborhood system. Each ball in Ξ intersects at most $2A_d(1/2)$ balls in Ξ of larger or equal radius*[1].

Proof: Let $P = \{p_1, \ldots, p_n\}$ denote the set of centers of Ξ. Let $\Gamma_>$ be the set of all balls in Ξ with the same or larger radius, that intersect B_i. Partition $\Gamma_>$ into two subsets $\Gamma_> = \Gamma_>^{(1)} \cup \Gamma_>^{(2)}$, where

$$\begin{aligned} \Gamma_>^{(1)} &= \{B_i \in \Gamma_> : p_i \notin (2 \cdot B_1)\} \\ \Gamma_>^{(2)} &= \{B_i \in \Gamma_> : p_i \in (2 \cdot B_1)\} \end{aligned}$$

By Lemma 3.3, $|\Gamma_>^{(1)}| \leq A_d(1/2)$, and by lemma 3.4, $|\Gamma_>^{(2)}| \leq A_d(1/2)$. □

Consequently,

Theorem 4.2 *The intersection graph of a d-dimensional k-neighborhood system is $2A_d(1/2)k$-inductive.*

Theorem 4.3 *For each d, there is a constant c_d that only depends on d such that the α-overlap graph of a d-dimensional k-neighborhood system is $(c_d\alpha^d k)$-inductive.*

Proof: Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a k-neighborhood system with centers $P = \{p_1, \ldots, p_n\}$. Let r_i be the radius of B_i. For each i, let $\Gamma_{>\alpha}(i) = \{B_j : r_j \geq r_i \ \& \ B_j \cap (\alpha \cdot B_i) \neq \emptyset\}$

The theorem follows from the claim that $|\Gamma_{>\alpha}(i)| \leq O(\alpha^d k)$. To prove the claim, partition $\Gamma_{>\alpha}(i)$ into two subsets,

$$\begin{aligned} \Gamma_{>\alpha}^{(1)}(i) &= \{B_j \in \Gamma_{>\alpha}(i) : p_j \notin ((\alpha+1) \cdot B_i)\} \\ \Gamma_{>\alpha}^{(2)}(i) &= \{B_j \in \Gamma_{>\alpha}(i) : p_j \in ((\alpha+1) \cdot B_i)\} \end{aligned}$$

By Lemma 3.3, we have $|\Gamma_{>\alpha}^{(1)}(i)| \leq A_d(1/(\alpha+1))$.

Without loss of generality, assume that the radius of each ball in $\Gamma_{>\alpha}^{(2)}(i)$ is r_i. For otherwise, we just shrink each balls in $\Gamma_{>\alpha}^{(2)}(i)$ to one of radius r_i, and the resulting ball still intersects $\alpha \cdot B_i$ because its center is in $(\alpha+1) \cdot B_i$. Moreover the shrinking operation preserves the k-neighborhoodness. It follows that all balls in $\Gamma_{>\alpha}^{(2)}(i)$ are completely contained in the ball B of radius $(\alpha+2)r_i$ centered at p_i. It follows from the Density Lemma, the total volume of balls in $\Gamma_{>\alpha}^{(2)}(i)$ is bounded from above by $\tau_d k$ times the volume of B, and hence,

$$\tau_d |\Gamma_{>\alpha}^{(2)}(i)| V_d r_i^d \leq \tau_d k V_d (\alpha+2)^d r_i^d,$$

which implies $|\Gamma_{>\alpha}^{(2)}(i)| \leq \tau_d k(\alpha+2)^d$, completing the proof of the claim from which follows the lemma. □

Similarly,

Theorem 4.4 *For each d, there is a constant c_d, such that the α-overlap of a neighborhood system with density μ is $O(c_d\alpha^d\mu)$-inductive.*

[1] This bound is not quite tight. I believe the real bound is $A_d(1/2) + \tau_d$.

4.3 Why overlap graphs

One of the most valuable aspects of the class of overlap graphs is that it enables us to give an unified geometric characterization of graphs with the small separator property. This class has a simple definition, yet it contains not only many of naturally defined geometric classes, including grid graphs, sphere packing graphs, k-nearest neighborhood graphs, and graphs associated with finite element methods, as special cases, but also some combinatorially defined graph classes, most notably, the class of planar graphs, as a special case. Furthermore, each graph in this class has a small separator (see Chapter 6).

4.3.1 Geometric embedding

To precisely define what we mean that the class of overlap graphs contains another class as a special cases, we introduce a notion of geometric embedding of a graph. An *embedding* of a graph $G = (V, E)$ in $\mathbb{R}^d$ is a mapping $\pi : V \to \mathbb{R}^d$. Naturally, an edge (v, u) is mapped to the line segment $(\pi(v), \pi(u))$. A graph $G = (V, E)$ is *k-embeddable* in $\mathbb{R}^d$ if there is an embedding π of G such that for all $(u, v) \in E$, $B_u \cap B_v \neq \emptyset$, where B_u is the largest ball centered at $\pi(u)$, that contains no more than k points from $\{\pi(w) : w \in V\}$ in its interior. Similarly, $G = (V, E)$ is *(α, k)-embeddable* in $\mathbb{R}^d$ if there is an embedding π of G such that for all $(u, v) \in E$, $(B_u \cap (\alpha \cdot B_v) \neq \emptyset)$ and $((\alpha \cdot B_u) \cap B_v \neq \emptyset)\}$.

We say a grah G_1 is a *spanning graph* of another graph G_2 if G_1 can be obtained from G_2 by deleting edges. So, a graph G is (α, k)-embeddable if it is a spanning subgraph of the α-overlap graph of a k-neighborhood system.

Notice that the small separator property is preserved under spanning subgraphs.

4.3.2 Why they are special cases

Grid graphs

For each point $(i_1, \ldots, i_d)$ where $0 \leq i_j < n$, let $B_{(i_1,\ldots,i_d)}$ be the ball centered at $(i_1, \ldots, i_d)$ of radius 1. Clearly, $\Xi = \{B_{(i_1,\ldots,i_d)} : 0 \leq i_j < n, 1 \leq i \leq d\}$ is a 1-neighborhood system and the grid graph is the intersection graph of Ξ. Therefore, the d-dimensional $m \times \cdots \times m$ grid graph is the intersection graph of a 1-neighborhood system in d-space.

Sphere packing graphs

Since each sphere in a sphere packing shares no common interior point with other spheres, the set of all balls associated with spheres forms a 1-neighborhood system and therefore, each sphere packing graph in d-space is the intersection graph of a 1-neighborhood system in d-space.

Planar graphs

The proof that planar graphs are special case of overlap graphs relies on the following recent deep theorem of Andreev and Thurston [3, 4, 83] characterizing all planar graphs in a novel geometric fashion.

Theorem 4.5 (Andreev and Thurston) *Each triangulate planar graph G is isomorphic to a 2-dimensional sphere packing graph.*

Consequently,

Corollary 4.6 *Every planar graph is 1-embeddable in 2-space.* □

The k-nearest neighborhood graphs

Given a set of points in $\mathbb{R}^d$, let B_i be the largest ball centered at p_i, whose interior contains no more than k points from P. By the definition of k-nearest neighborhood graph, there is an edge between balls B_i and B_j only if either p_i is in B_j or p_j is in B_i and hence the line segment between p_i and p_j is completely covered in $B_i \cup B_j$. Therefore, Each k-nearest neighborhood graph in d dimensions is k-embeddable in d-space.

Graph that can be drawn in a civilized manner

Let π be a (r,s)-civilized drawing of G in $\mathbb{R}^d$, namely, it maps each vertex v of G to the point $\pi(v)$. Let B_v be the ball centered at $\pi(v)$ with radius s. Clearly, $\Xi = \{B_v : v \in G\}$ is a 1-neighborhood system. For each edge $(u,v) \in G$, $||\pi(u) - \pi(v)|| \leq r = (\frac{r}{s})s$. Thus, $(r/s) \cdot B_u$ contains $\pi(v)$, and therefore, if a graph G can be drawn in $\mathbb{R}^d$ in a (r,s)-civilized manner, then it is $(r/s,1)$-embeddable in d-space.

Density graphs

The following class of graphs is defined by Miller and Vavasis [69].

Let G be an undirected graph and let π be an embedding of its nodes in $\mathbb{R}^d$. We say π is an embedding of *density* α if the following inequality holds for all vertices v in G. Let u be the closest node to v. Let w be the farthest node from v that is connected to v by an edge. Then

$$\frac{||\pi(w) - \pi(v)||}{||\pi(u) - \pi(v)||} \leq \alpha.$$

In general, G is a *density graph* if there exist a π and $\alpha > 0$ such that π is an embedding of density α.

Simply from the definition, each d-dimensional α-density graph is $(\alpha,1)$-embeddable in d-space.

4.4 Why not dilation graphs

During this research, after showing that all intersection graphs have small separators, we tried to prove that all dilation graphs also have small separators. As a matter of fact, we did once find "such a proof" (not surprisingly dilation graphs were called overlap graphs at that time). After a couple of months of enjoyment, we found the following counter-example. Geometric intuition, sometimes can be misleading! But fortunately, our work was not wasted. After regrouping our thoughts, we defined the new class of overlap graphs, and happily found that our old proof, the one which was wrong, was in fact a correct proof for the new class.

Lemma 4.7 *For each $\alpha > 1$, there is a 1-neighborhood system in d-dimensions whose α-dilation graph is a complete graph.*

Proof: Let us first consider $d = 1$. Let $P = \{p_0, ..., p_{n-1}\}$ where $p_0 = 0$, $p_i = \alpha^i$, $1 \leq i \leq n-1$. Let B_i be the 1-dimensional ball centered at p_i with radius $\alpha^i - \alpha^{i-1}$ for $1 \leq i \leq n-1$ and B_0 be the ball with center p_0 and radius α. It is easy to check that $\Xi = \{B_0, B_1, \ldots, B_{n-1}\}$ forms a 1-neighborhood system.

To show the α-dilation graph of Ξ is a complete graph, we need to prove for all $0 \leq i < j \leq n-1$, $\alpha \cdot B_i$ intersects $\alpha \cdot B_j$. Notice that the radius of $\alpha \cdot B_j$ is $\alpha(\alpha^j - \alpha^{j-1}) = \alpha^j(\alpha - 1) > \alpha^j$. The last inequality is true because we assume $\alpha > 1$. But the distance between p_0 and p_j is α^j, so $\alpha \cdot B_j$ contains p_0, and thus all points from $\{p_0, p_1, \ldots, p_{j-1}\}$. Clearly, this construction can be easily embedded in any dimension. □

Therefore, although an overlap graph is always "sparse", a dilation graph can be very "dense". However, the dilation graphs over the following "balanced" neighborhood systems enjoy many properties (such as sparsity and small separator properties) of overlap graphs. This fact will be used to design efficient algorithm for constructing an intersection graph given its neighborhood system (See Chapter 9).

A neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ is *β-balanced* for some positive real $\beta \geq 1$, if for all $1 \leq i, j \leq n$,

$$\frac{1}{\beta} \leq \frac{\text{radius}(B_i)}{\text{radius}(B_j)} \leq \beta$$

Lemma 4.8 *For all $\alpha, \beta > 1$, if G is an α-dilation graph over a β-balanced neighborhood system, then maximum degree of G is bounded from above by $O(\alpha^d \beta^d k)$.*

Proof: Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a β-balanced neighborhood system with centers $P = \{p_1, \ldots, p_n\}$. Let r_i be the radius of B_i. Let $r = \min_i r_i$ and $r' = \max_i r_i$.

Suppose B is a ball with radius $3\alpha r'$ centered at p_i. Let Γ be the set of all balls from Ξ which is the neighbor of B_i in the dilation graph G of Ξ. It is not hard to see that each ball in Γ is completely contained in B. By Density Lemma, the total volume of balls from Γ is bounded from above by $\tau_d k$ times the volume of B. Notice that the total volume of balls in Γ is at least $|\gamma| V_d r^d$. So

$$|\gamma| V_d r^d \leq \tau_d k V_d (3\alpha r')^d,$$

which implies that $|\gamma| \leq \tau_d (3\alpha\beta)^d k$. □

Chapter 5

Sphere Separators

Each $(d-1)$-sphere S separates int(S) from ext(S): any segment connecting a point in int(S) and one in ext(S) must intersect S. In analogue to vertex separators in graph theory, S is called a *sphere separator* in d-space. More specifically, for a set of points $P = \{p_1, \ldots, p_n\}$ in $\mathbb{R}^d$ and a constant $0 < \delta < 1$, S *δ-splits* P if both $|\text{int}(S) \cap P| \leq \delta n$ and $|\text{ext}(S) \cap P| \leq \delta n$.

Each vertex separator of graph has a cost, e.g., the number of vertices in the separator, and most applications require separators with a small cost. In the same spirit, we will assign cost to sphere separators. The usefulness of sphere separators with a small cost will be demonstrated in Chapter 9. One of the goals of this chapter is to investigate the relationship between the cost of sphere separators and the size of vertex separators of overlap graphs and intersection graphs.

5.1 Surface area of a sphere

A natural cost is the surface area of a sphere. The surface area can be either weighted or unweighted. The *unweighted* area of a sphere in d-space is given by

$$\frac{2\pi^{d/2} r^{d-1}}{\Gamma(d/2)}$$

(see [40] pp 68), where

$$\Gamma(x) = \int_0^\infty e^t t^{x-1} dt, \qquad \text{for } x > 0$$

is the *gamma-function*. The gamma-function satisfies

$$\Gamma(n) = (n-1)! \qquad \text{for all positive integer } n.$$

In the weighted case, there is a real valued nonnegative function $f(x)$ defined on $\mathbb{R}^d$ such that f^k is integrable for all $k = 1, 2, 3, \ldots$. Such an f is called a *density function*. The surface area of a $(d-1)$-dimensional sphere S is then given by

$$\text{Area}_f(S) = \int_{v \in S} (f(v))^{d-1} (dv)^{d-1}$$

Also defined is the *total volume* of the function f, denoted by Total-Volume(f), which is given by

$$\text{Total-Volume}(f) = \int_{v \in \mathbb{R}^d} (f(v))^d (dv)^d$$

Notice that the unweighted surface area is just the special case when $f = 1$ at all points in $\mathbb{R}^d$.

5.2 Density functions on a unit sphere

Density function can also be defined on the surface of a sphere. To be consistent with the discussion of subsequent sections, we focus on the unit d-sphere. Suppose U_d is a unit d-sphere in $\mathbb{R}^{d+1}$ and f is a real valued nonnegative function defined on the surface of U_d such that f^k is integrable for all $k = 1, 2, 3, \ldots$. We call f a *density function* of U_d. The *total volume* of f is defined to be

$$\text{Total-Volume}(f) = \int_{v \in U_d} (f(v))^d (dv)^d$$

A *great sphere* of U_d is the intersection of U_d with a hyperplane passing through the center of U_d. The *weighted area* of a great sphere GS of U_d is given by

$$\text{Area}_f(GS) = \int_{v \in GS} (f(v))^{d-1} (dv)^{d-1}$$

Let $\text{avg}(f)$ be the average area over all great spheres of U_d. The following proposition[1] follows simply from the Hölder's inequality [43].

Proposition 5.1 *Suppose f is a density function on a unit d-sphere U_d, then*

$$\text{avg}(f) = A_{d-1} \left((\text{Total-Volume}(f))^{\frac{d-1}{d}} \right),$$

where A_d stands for the surface area of U_d.

Proof: Each great sphere GS can be uniquely identified with the pair of points p_{GS} and q_{GS} on U_d which lay on the normal to GS. (see Figure 5.1 for an example in 2 dimensions).

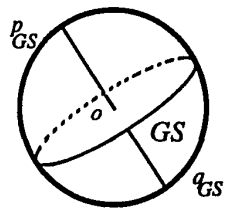

Figure 5.1: GC and its p_{GS}, q_{GS}

For each $x \in U_d$, let $GS(x)$ be the great sphere associated with x. We can write $\text{avg}(f)$ as:

[1] A more complicated proof of this proposition was given in [67] which made use of the Cauchy-Schwartz inequality – a special case of the Hölder inequality. A proof based on Hölder inequality first appeared in a full version of [69].

$$\begin{aligned}
\text{avg}(f) &= \frac{1}{A_d}\left(\int_{x\in U_d}\left[\int_{v\in GS}(f(v))^{d-1}(dv)^{d-1}\right](dx)^d\right)\\
&= \frac{A_{d-1}}{A_d}\left(\int_{z\in U_d}(f(z))^{d-1}(dz)^d\right)\\
&\leq A_{d-1}\left(\int_{z\in U_d}(f(z))^d(dz)^d\right)\\
&= A_{d-1}\left((\text{Total-Volume}(f))^{\frac{d-1}{d}}\right),
\end{aligned}$$

The second equality follows from the observation that for each point z on U_d, the set of point $\{x : z \in GS(x)\}$ forms a great sphere. By choosing $g = f^{d-1}$, $h = 1$, $p = \frac{d}{d-1}$ and $q = d$, the next inequality follows from the Hölder inequality which states (see [43]):

Hölder Inequality:
Suppose g and h are functions suitably integrable on a measurable set P and for positive real numbers p, q such that $1/p + 1/q = 1$. Then the following inequality holds

$$\int_P |g \cdot h| \leq \left(\int_P g^p\right)^{\frac{1}{p}} \cdot \left(\int_P h^q\right)^{\frac{1}{q}}.$$

□

5.3 Intersection number

Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system in d-space. Each $(d-1)$-dimensional sphere S partitions Ξ into three subsets (see Figure 5.2)

- Ξ_I: those that are completely in the interior of S;
- Ξ_E: those that are completely in the exterior of S;
- Ξ_O: those that intersect S.

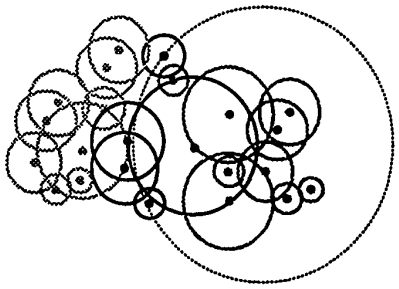

Figure 5.2: A sphere separator

The cardinality of Ξ_O is an important cost of S. This number is called the *intersection number* of S, denoted by $\iota_\Xi(S)$.

Notice that the removal of Ξ_O splits Ξ into two subsets: Ξ_I and Ξ_E, such that no ball in Ξ_I intersects any ball in Ξ_E and vice versa. To see this, assume that there is a $B_i \in \Xi_I$ intersects $B_j \in \Xi_E$. It follows that the line segment between p_i and p_j, the centers of B_i and B_j, respectively, is complete covered by $B_i \cup B_j$. Since p_i is in the interior of S and p_j the exterior, the line segment between p_i and p_j must intersect S. So either B_i or B_j intersects S, contradicting the assumption that neither $B_i \in \Xi_O$ nor $B_j \in \Xi_O$. In analogy to separators in graph theory, Ξ_O can be viewed as a separator of Ξ.

The follow lemma gives a simple a bijection between separators of a neighborhood system Ξ and separators of the intersection graph of Ξ.

Lemma 5.2 *If S is a sphere that δ-splits the centers of a neighborhood system Ξ, then Ξ_O defined above is an $\iota_\Xi(S)$-separator that δ-splits the intersection graph of Ξ.* □

Now the question of whether the intersection graph of a neighborhood system Ξ has a small separator can be reduced to the one of whether Ξ has a sphere separator S with a low intersection number. The reduction is straightforward but it opens up a new geometric approach to study graph separators. As a part of the main separator results, I will prove the following theorem which is interesting on its own right.

Theorem 5.3 (Sphere Separator Theorem) *For each neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ of density μ, there is a sphere S with intersection number $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$ that $\frac{d+1}{d+2}$-splits Ξ.*

As a matter of fact, I will prove the main separator theorem by relating the intersection number to the surface area of some properly selected density function, and using the continuous separator theorem of Miller and Thurston (see Section 5.5).

5.4 Why spheres but not hyperplanes

Perhaps the simplest way to split a set of points in d-space is to use a $(d-1)$-dimensional hyperplane. Like a $(d-1)$-dimensional sphere, a $(d-1)$-dimensional hyperplane H partitions $\mathbb{R}^d$ into three subsets, H^+, those that are above H, H^-, those below H, and H itself. For a set of points $P = \{p_1, \ldots, p_n\}$ in $\mathbb{R}^d$ and a constant $0 < \delta < 1$, H *δ-splits* P if both $|H^+ \cap P| \leq \delta n$ and $|H^- \cap P| \leq \delta n$. The smallest such δ is called the *splitting ratio* of H.

Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system in d-space. For each $(d-1)$-dimensional hyperplane H, let $\iota_\Xi(H)$ denote the number of neighborhoods of Ξ that intersects H. Similar to Lemma 5.2

Lemma 5.4 *If H is a hyperplane separator that δ-splits the centers of a neighborhood system Ξ, then H induces an $\iota_\Xi(H)$-separator that δ-splits the intersection graph of Ξ.*

It seems that hyperplanes enjoy the same basic properties of spheres. This is not surprising, because from conformal mapping viewpoint, a $(d-1)$-dimensional hyperplane is just a degenerated $(d-1)$-sphere. Then why don't we use hyperplane as separators?

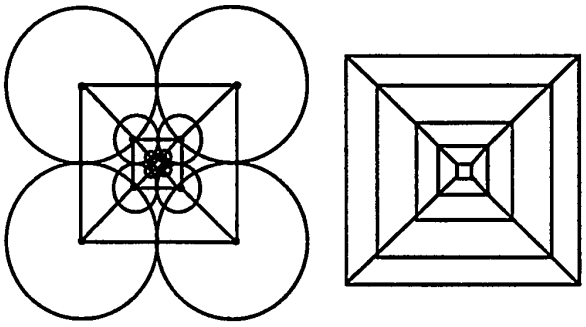

Figure 5.3: Why not hyperplane separators

One problem with hyperplanes as separators is that there is a 1-neighborhood system such that no hyperplane with intersection number $o(n)$ has a constant splitting ratio. Such an example is given in Figure 5.3.

However, by Theorem 5.3, the above neighborhood system should have a sphere separator with intersection number $O(\mu^{1/d} n^{\frac{d-1}{d}})$ that $\frac{d+1}{d+2}$-splits Ξ. A sphere separator 1/2-splitting Ξ is illustrated in Figure 5.4, whose intersection number is $O(n^{\frac{d-1}{d}})$.

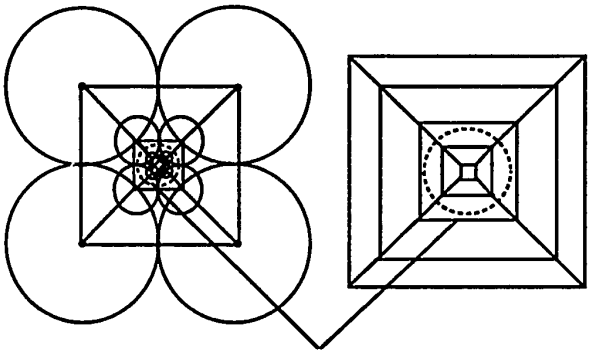

Figure 5.4: Why sphere separators

5.5 Center points and a continuous separator theorem

It is not hard to compute a sphere which evenly splits a set P of n points in d-space. One can start with any point in $\mathbb{R}^d$ as a center, and grow a sphere till it contains at least $\lceil n/2 \rceil$ points from P. However, chances are such a sphere intersects too many balls in the given neighborhood system. To prove each neighborhood system with density μ has a sphere separator of intersection number $O(\mu^{1/d} n^{\frac{d-1}{d}})$, I will show that there is a set of infinite (or finite) sphere separators, whose average intersection number is $O(\mu^{1/d} n^{\frac{d-1}{d}})$. This type of averaging argument was inspired by Proposition

5.1 and first used by Miller and Thurston [67].

5.5.1 Conformal maps and stereographic projections

To use Proposition 5.1 to construct a set of infinite sphere separators, we map points of $\mathbb{R}^d$ onto the unit d-sphere U_d in $\mathbb{R}^{d+1}$ centered at the origin o. Notice that the density function f is mapped to a new density function f' on U_d in order to ensure

$$\text{Total-Volume}(f) = \text{Total-Volume}(f').$$

To preserve the density function in every dimensional simultaneously, the map is required to be *conformal* in the sense that it preserves angles [47, 23, 83]. Clearly, rigid motions and dilations are conformal maps. Another conformal map is the stereographic projection which maps $\mathbb{R}^d$ plus infinity onto U_d [47, 67]. The composition of any two conformal is also conformal, so we can apply a sequence of conformal maps to obtain a desirable conformal map.

The conformal map has another nice property: the pre-image of each great sphere GS of U_d itself is a $(d-1)$-sphere in $\mathbb{R}^d$. Moreover, the interior and the exterior of the pre-image sphere are mapped to the two hemispheres of U_d defined by GS, respectively.

It simply follows from Proposition 5.1 that

Lemma 5.5 *Suppose f is a density function of $\mathbb{R}^d$ and π is a conformal map from $\mathbb{R}^d$ to U_d. Let*

$$SS(\pi) = \{S : S \text{ is a pre-image of some great sphere } GS, \text{ i.e., } GS = \pi(S)\}$$

and $avg_\pi(f)$ be the average area over all spheres in $SS(\pi)$. Then,

$$\text{avg}_\pi(f) = A_{d-1}\left((\text{Total-Volume}(f))^{\frac{d-1}{d}}\right).$$

Suppose $P = \{p_1, \ldots, p_n\}$ is a set of points in $\mathbb{R}^d$ and $Q = \{q_1, \ldots, q_n\}$ is the set of points on U_d such that q_i is the image of p_i $(1 \leq i \leq n)$ for a conformal map π from $\mathbb{R}^d$ to U_d. It follows from the above discussion of conformal map that for each great sphere GS of U_d, if S is the pre-image of GS, then S δ-splits P iff GS δ-splits Q.

Therefore, to ensure that each sphere of $SS(\pi)$ δ-splits P for some constant $0 < \delta < 1$, the conformal map π need to satisfy the condition that all hyperplanes passing o δ-split Q. A useful concept in constructing such a conformal map is a notion of a center point.

5.5.2 Center points

Suppose P is a finite set of points in $\mathbb{R}^d$. For each $0 < \alpha < 1$, a point $c \in \mathbb{R}^d$ is an α-center point of P if every hyperplane containing c α-splits P. Each $\frac{d}{d+1}$-center points is called a *center points* of P, and the set of all center points is denoted by center(P). The balanced separation property of a center point makes it very useful for designing efficient divide and conquer algorithms, especially in computational geometry [97, 66, 68], decision theory [53], and numerical analysis [67, 68, 69, 59].

Given a set of points $P \subset \mathbb{R}^d$, the question of whether P has a center point is always affirmative. It follows from Helly's Theorem:

Theorem 5.6 (Helly's Theorem) *Suppose $\mathcal{K}$ is a family of at least $d+1$ convex sets in $\mathbb{R}^d$, and $\mathcal{K}$ is finite or each member of $\mathcal{K}$ is compact. Then if each $d+1$ members of $\mathcal{K}$ have a common point, there is a point common to all members of $\mathcal{K}$.*

Lemma 5.7 (Center Points) *For each set $P \subseteq \mathbb{R}^d$, $\text{center}(P) \neq \emptyset$.*

Proof[2]: Let HS be the set of all open half-spaces in $\mathbb{R}^d$ which contain more than $\lfloor \frac{d}{d+1}|P| \rfloor$ points of P. We want to show that

$$\text{center}(P) = \bigcap_{h \in HS} h \neq \emptyset.$$

By Helly's theorem, it is sufficient to show that for each $h_1, \ldots, h_{d+1} \in HS$, $\bigcap_{i=1}^{d+1} h_i \neq \emptyset$. Note that

$$\bigcap_{i=1}^{d+1} h_i = \mathbb{R}^d - \bigcup_{i=1}^{d+1} (\mathbb{R}^d - h_i) \supseteq P - \bigcup_{i=1}^{d+1} (\mathbb{R}^d - h_i)$$

Note also

$$|\bigcup_{i=1}^{d+1} ((\mathbb{R}^d - h_i) \cap P)| \leq \sum_{i=1}^{d+1} |((\mathbb{R}^d - h_i) \cap P)| < (d+1)\lceil \frac{1}{d+1}|P| \rceil < |P|$$

Hence, $P - \bigcup_{i=1}^{d+1} (\mathbb{R}^d - h_i) \neq \emptyset$. □

5.5.3 A continuous separator theorem

In [67], Miller and Thurston prove the following result:

Lemma 5.8 *Suppose P is a set of points in $\mathbb{R}^d$. Then there is a conformal map π such that o is a $(d+1)$-dimensional center point of Q, the image of P.*

Proof: Such a conformal map can be obtained by first applying the stereographic projection and then conformally moving a centerpoint to the origin o. In fact, a center point can be moved to o by one rotation followed by a dilation. For details of the proof, we refer readers to Miller and Thurston [67]. □

From the above lemma and Proposition 5.1, Miller and Thurston derived the following continuous separator theorem when weighted surface area is used as the cost function of a sphere.

Theorem 5.9 (Miller and Thurston) *Suppose f is a density function on $\mathbb{R}^d$ and P a set of n distinct points in $\mathbb{R}^d$. Then there is a sphere S which $(\frac{d+1}{d+2})$-splits P such that*

$$\text{Area}_f(S) = O\left((\text{Total-Volume}(f))^{\frac{d-1}{d}}\right)$$

[2]We present this proof for the sake of completeness. The similar proofs can be found in many previous works, e.g. [20].

5.6 An algorithm for computing a sphere separator

The above discussion and the proof of Miller and Thurston [67] suggest the following randomized algorithm for computing a sphere separator with a small area.

ALGORITHM Sphere Separator

1. Compute a conformal map π from $\mathbb{R}^d$ to U_d, such that the origin o is a center point of $\pi(P)$;
2. Randomly choose a great sphere GS of U_d;
3. Using the inverse of π, map GS back to d-space to get a $(d-1)$-sphere S.

Because o is a center point of $\pi(P)$ in $(d+1)$-space, the sphere S $\frac{d+1}{d+2}$-splits P. Moreover, as proved in the continuous separator theorem, with probability $1/2$, S has area

$$(\text{Total-Volume(f)})^{\frac{d-1}{d}}.$$

The step 2 above uses a random number generator which uniformly generates a real number from $[0,1]$ (Knuth [54]; Hammersley and Handscomb [41]). In Chapter 10, I will show that this requirement can be relaxed when applying to neighborhood systems.

The run time of the above algorithm crucially depends on the time needed to compute a center point in $(d+1)$-space. All other steps of the algorithm can be performed in $O(n)$ time, and in constant time using $O(n)$ processors.

Unfortunately, no linear time or even $n \cdot \text{polylog}(n)$ time algorithm is known to compute a center point in d-space ($d \geq 2$). However, I will show in Chapter 8 that a $(\frac{d}{d+1}+\epsilon)$-center point (with very small ϵ) can be computed in random constant time in parallel and deterministic linear time sequentially. Therefore,

Theorem 5.10 *Suppose f is a density function on $\mathbb{R}^d$ and P is a set of n points in $\mathbb{R}^d$. Then a $(\frac{d+1+\epsilon}{d+2})$-splitting sphere S of P of area*

$$\text{Area}_f(S) = O\left((\text{Total-Volume}(f))^{\frac{d-1}{d}}\right)$$

can be computed in deterministic linear time, and in random constant time using $O(n)$ processors, where $\frac{1}{n^{1/2d}} \leq \epsilon \leq 1$.

Chapter 6

Separator Theorems

Now that we have built our tool kit, we need to use it to prove that intersection graphs and overlap graphs have small separators.

6.1 Separating intersection graphs

We first start with the simpler class, the intersection graph, and hope to develop some background and intuition to handle the more general one. The following is the main theorem of this section.

Theorem 6.1 (Intersection graphs) *If G is the intersection graph of a d-dimensional neighborhood system with density μ, then G has an $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$-separator that $(\frac{d+1}{d+2})$-splits.*

Moreover, the bound in Theorem 6.1 is best possible in term of n, and μ (see section 6.3).

6.1.1 The initial idea

In order to use Theorem 5.9 to prove Theorem 6.1, we need to come up with a density function f and show that Total-volume$(f) = O(\mu^{\frac{1}{d-1}} n)$. Then by the continuous separator theorem, there is a sphere separator S with weighted surface area $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$ $\frac{d+1}{d+2}$-splitting the centers of the given neighborhood system. However, not all density functions have the property that a vertex separator can be deduced from a sphere separator, whose size is linear in the area of the sphere separator.

So we need to select a density function with special care. Not only should the total volume be linearly bounded by n, but also the density function should be *faithful* to the underlying graph in the sense that the area of a sphere separator is linearly related with size of a vertex separator. In other words, we have to construct a density function f which encodes the separator property of the underlying graph.

By Lemma 5.2, each sphere separator S induces a vertex separator whose size is bounded by the intersection number of S. In the following sections, we will construct a density function which approximates the intersection number, and prove that the total volume of this function is $O(\mu^{\frac{1}{d-1}} n)$.

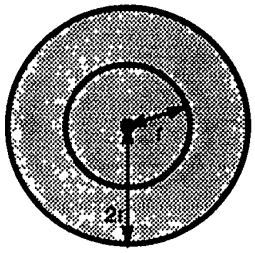

Figure 6.1: Influential area of a kingdom and its density

6.1.2 Local density functions

To visualize the intersection number of a sphere in 2-space, imagine we are riding a circle (1-sphere) in a hypothetical world which has n kingdoms. Geographically, each kingdom has the shape of a disk, and kingdoms may share land among themselves. Land which is not claimed by any kingdom is public and is free to drive on. The cost of such a drive is the total number of kingdoms we have visited.

If we view the density function as the inverse of the speed of driving, then the weighted length of a circle represents the total time it takes to drive along the circle. The goal is thus to assign a speed at each point so that the total time of a drive is linearly related with the number of kingdoms been visited.

One way to achieve this goal is to spend at least one hour but no more than two hours in a kingdom, so the total time of a drive is at least the number of kingdoms being visited, but no more than the twice of that number. In order not to spend any time in the public area, we will drive at infinity speed there, i.e., with zero density. In general, the larger the kingdom, the fast the speed in it – we will drive at a speed proportional to the radius of a kingdom.

But there is a serious problem with the above suggestion: Suppose a drive just touches kingdom or barely crosses one, how can we manage to spend a decent period of time there? (Shall we stop for coffee to kill time?). One solution is to make the density larger when closer to the boundary of a kingdom. However, this still does not put any charge on just touching a kingdom.

The solution we are now proposing it to think that each kingdom has an *influential area* which has radius twice as big as that of the kingdom itself. Using the notation of neighborhoods, $2 \cdot B_i$ is the influential area of B_i. Rather than assign a density within B_i itself, we assign a uniform density function within $2 \cdot B_i$ for each kingdom B_i, which is inversely propositional to the radius of B_i (See Figure 6.1).

To justify the intuition, we now show that if the circle passes a kingdom B_i, then the length of the curve within $2 \cdot B_i$ is bounded from below by a constant times the radius of B_i – it takes some constant fraction of an hour to cross the influential area of each kingdom. Formally,

Lemma 6.2 *Suppose B is a ball in d-space with radius r and S a $(d-1)$-sphere of radius at least $2r$. If S intersects B, then area of $S \cap (2 \cdot B)$ is at least $\left(\frac{\sqrt{7}}{2}r\right)^{d-1} V_{d-1}$, where V_{d-1} is the volume of a unit $(d-1)$-dimensional ball.*

Proof: First, consider the case when the center of B is in the exterior of S. Assume S is a sphere

that has smallest common area with $2 \cdot B$ satisfying the condition of the lemma. The optimality of S implies that the radius of S is equal to $2r$. For otherwise, the sphere S', which has radius $2r$ and the same center as S, would have smaller common area with $2 \cdot B$. Because the center of B is in the exterior of S, the optimality of S also implies that S touches B (see Figure 6.2),

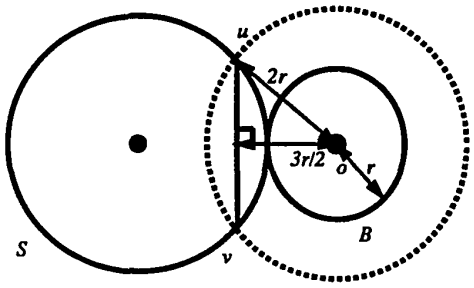

Figure 6.2: The common area of S and $2 \cdot B$.

The points common to S and the boundary of $2 \cdot B$ is a $(d-2)$-sphere, denoted by C. Notice that the $(d-2)$-dimensional volume of C is smaller than the area of $S \cap 2 \cdot B$ (see 6.2 for the 2-dimensional case where the length of uv is less than the length of the arc uv).

Let uv be a diameter of C, Figure 6.2 shows that the length of uv is $\sqrt{7}r$. So the volume of C is least $(\frac{\sqrt{7}}{2}r)^{d-1}V_{d-1}$.

Similarly, if the center of B is in the interior of S, then the common area of S with $2 \cdot B$ is at least $(\sqrt{3}r)^{d-1}V_{d-1}$. □

Now that we have solved the problem of barely crossing a kingdom, but may be at the cost of the problem when the circle marginally missing a kingdom. Did we overcharge ourselves?

We will show later, this is not quite a problem, thanks to the nice properties of neighborhood systems, and the Big-O notation.

To formally summarize the above discussion, let r_i be the radius of B_i and let $\gamma_i = 2r_i$, define

$$f_i(x) = \begin{cases} 1/\gamma_i & \text{if } \|x - p_i\| \leq \gamma_i \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, f_i sets up a cost on each $(d-1)$-sphere S such that the closer S is to B_i, the larger B_i contributes to the surface area of S. The function f_i is called the *local* density function of B_i.

6.1.3 Putting local density functions together

Now we need to construct a global density function from the local ones. The simplest way may be to take the sum, i.e., $f = \sum_i f_i$. But is the sum really a good choice? To gain more insight, let us calculate the total volume in the following extreme case when Ξ contains n identical balls. In this case, the intersection graph of Ξ is the complete graph with n vertices and the total volume of the sum function is $n^d V_d$, which is clearly too big for our applications because no sphere has intersection number more than n. So $f = \sum_i f_i$ is not an optimal choice, although the following lemma can be proved:

Lemma 6.3 *Let $\Xi = \{B_1, \ldots, B_n\}$ be a d-dimensional neighborhood system with density μ. If f_i is defined as above and $f = \sum_i f_i$, then* Total-Volume$(f) = O(\mu^d n)$.

Notice that when μ is a constant, the above lemma yields a weaker version of Theorem 6.1.

6.1.4 Using norms: Vavasis's idea

Suppose $a_1, \ldots, a_n$ are n real numbers. For each positive integer p, the L_p *norm* of $a_1, \ldots, a_n$, denoted as $L_p(a_1, \ldots, a_n)$, is given by

$$L_p(a_1, \ldots, a_n) = \left(\sum_{i=1}^{n} |a|^p \right)^{1/p}.$$

The following lemma states the relationship between different norms.

Lemma 6.4 *Let $a_1, \ldots, a_n$ be n reals. If $p \leq q$, then $L_p(a_1, \ldots, a_n) \geq L_q(a_1, \ldots, a_n)$.*

Proof: See Hardy, Littlewood and Pólya [43] (pp 26 and pp 144). □

The density function[1] of the intersection graph G is then defined to be the L_{d-1} norm of $f_1, \ldots, f_n$, i.e.,

$$f(x) = L_{d-1}(f_1, \ldots, f_n) = \left(\sum_{i=1}^{n} (f_i(x))^{d-1} \right)^{\frac{1}{d-1}}$$

Note that the L_d norm is not a good choice as a global function because its total volume is $V_d n$ for all neighborhood systems,.

6.1.5 Relating to the intersection number

Lemma 6.5 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ be a d-dimensional neighborhood system with density μ. Let $f_1, \ldots, f_n$ be local density functions defined above and let $f = L_{d-1}(f_1, \ldots, f_n)$. For each $(d-1)$ -sphere S,*

$$\iota_\Xi(S) \leq 3^d \mu + \left(\frac{4\sqrt{7}}{7} \right)^{d-1} \frac{1}{V_{d-1}} \text{Area}_f(S).$$

Proof: Let C be indices of neighborhoods in Ξ which intersect S. C can be partitioned into two subsets $C = C_1 \cup C_2$, where

$$\begin{aligned} C_1 &= \{i \in C : \gamma_i \geq \text{radius}(S)\} \\ C_2 &= C - C_1 \end{aligned}$$

By Lemma 3.6, $|C_1| \leq 3^d \mu$, and by the definition of area, we have

$$\text{Area}_f(S) = \int_{v \in S} \sum_{i=1}^{n} f_i(v)^{d-1} (dv)^{d-1} \geq \sum_{i \in C_2} \int_{v \in S} f_i(v))^{d-1} (dv)^{d-1}.$$

[1] In the construction of Miller and Vavasis for density graph, the cost function is defined to be the L_1 norm of the functions defined over each vertices.

By the definition of C_2, for each $i \in C_2$, S has a common point with B_i. Further, $radius(S) \geq \gamma_i$. So, by lemma 6.2, the area of $S \cap (2 \cdot B_i)$ is at least $\left(\frac{\sqrt{7}}{4}\gamma_i\right)^{d-1} V_{d-1}$.

Therefore,

$$\int_{v \in S} (f_i(v))^{d-1}(dv)^{d-1} \geq \text{Area}(A \cap (2 \cdot B_i)) \left(\frac{1}{\gamma_i}\right)^{d-1} \geq \left(\frac{\sqrt{7}}{4}\right)^{d-1} V_{d-1}.$$

Hence,

$$\text{Area}_f(S) \geq \sum_{i \in C_2} \int_{v \in S} f_i(v))^{d-1}(dv)^{d-1} \geq |C_2| \left(\frac{\sqrt{7}}{4}\right)^{d-1} V_{d-1},$$

which implies $|C_2| \leq \left(\frac{4\sqrt{7}}{7}\right)^{d-1} \frac{1}{V_{d-1}} Cost_f(S)$. □

6.1.6 Some basic lemmas

The following set of lemmas will be used in the next section. The proofs we present here are given by Vavasis [95].

Lemma 6.6 *Let $a_1, \ldots, a_n$ be nonnegative numbers, and suppose $p \geq 1$. Then*

$$\left(\sum_{i=1}^{n} a_i\right)^p \leq p \left[\sum_{i=1}^{n} a_i \left(\sum_{j=i}^{n} a_j\right)^{p-1}\right].$$

Proof: Define the function

$$\phi(x_1, \ldots, x_n) = \left(\sum_{i=1}^{n} x_i\right)^p.$$

We notice that

$$\frac{\partial \phi}{\partial x_i} = p \left(\sum_{i=1}^{n} x_i\right)^{p-1}.$$

Let (i) be the vector in $\mathbb{R}^n$ given by:

$$(i) = (0, \ldots, 0, a_i, a_{i+1}, \ldots, a_n).$$

Then

$$\begin{aligned}
\phi(a_1, \ldots, a_n) &= \phi((1)) - \phi((n+1)) \\
&= \sum_{i=1}^{n} [\phi((i)) - \phi((i+1))] \\
&= \sum_{i=1}^{n} \int_0^{a_i} \frac{\partial \phi}{\partial x_i}(0, \cdots, 0, t, a_{i+1}, \ldots, a_n)\, dt \\
&= p \sum_{i=1}^{n} \int_0^{a_i} \left(t + \sum_{j=i+1}^{n} a_j\right)^{p-1} dt
\end{aligned}$$

$$\leq p\sum_{i=1}^{n} a_i \cdot \left(a_i + \sum_{j=i+1}^{n} a_j\right)^{p-1}$$
$$\leq p\sum_{i=1}^{n} a_i \cdot \left(\sum_{j=i}^{n} a_j\right)^{p-1}.$$

□

Lemma 6.7 *Let $\ldots, m_{-1}, m_0, m_1, m_2, \ldots$ be a doubly infinite sequence of nonnegative numbers such that each m_i is bounded above by θ and such that at most a finite number of m_i's are nonzero. Let $d \geq 2$ be an integer. Then*

$$\left(\sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)}\right)^{\frac{d}{d-1}} \leq c_d \theta^{\frac{1}{d-1}} \sum_{k=-\infty}^{\infty} m_k 2^{-kd}$$

where c_d is a positive number depending on d.

Proof: Since at most a finite number of the m_k are nonzero, then we can apply the preceding lemma because the above sums are actually finite. Applying the lemma, we see that

$$\begin{aligned}
\left(\sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)}\right)^{\frac{d}{d-1}} &\leq \frac{d}{d-1} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot \left(\sum_{j=k}^{\infty} m_j 2^{-j(d-1)}\right)^{\frac{1}{d-1}} \\
&\leq \frac{d}{d-1} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot \left(\sum_{j=k}^{\infty} \theta \cdot 2^{-j(d-1)}\right)^{\frac{1}{d-1}} \\
&\leq \frac{d}{d-1} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot \left(\frac{\theta \cdot 2^{-k(d-1)}}{1 - 2^{-(d-1)}}\right)^{\frac{1}{d-1}} \\
&\leq c_d \theta^{\frac{1}{d-1}} \sum_{k=-\infty}^{\infty} m_k 2^{-k(d-1)} \cdot 2^{-k} \\
&\leq c_d \theta^{\frac{1}{d-1}} \sum_{k=-\infty}^{\infty} m_k 2^{-kd}.
\end{aligned}$$

□

6.1.7 Bound the total volume

Lemma 6.8 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ be a d-dimensional neighborhood system with density μ. If $f_1, \ldots, f_n$ and f are defined as above, then*

$$\text{Total-Cost}(f) = O(\mu^{\frac{1}{d-1}} n).$$

Proof: Let V_d be the volume of a unit ball in $\mathbb{R}^d$. Clearly, $\int_{x \in \mathbb{R}^d} (f_i(x))^d (dx)^d = V_d$.

Consequently, letting

$$g(x) = L_d(f_1, \ldots, f_n) = \left(\sum_{i=1}^{n} (f_i(x))^d\right)^{\frac{1}{d}},$$

we have

$$\int_{x \in \mathbb{R}^d} (g(x))^d (dx)^d = V_d n$$

Therefore, Lemma 6.8 follows immediately from the following lemma. □

Lemma 6.9 *For all* $x \in \mathbb{R}^d$, $(g(x))^d \leq (f(x))^d \leq c_d 2^d (6^d \mu)^{\frac{1}{d-1}} \cdot (g(x))^d$.

Proof: The first inequality follows immediately from the definitions of f and g, and Lemma 6.4.

For the second inequality, we focus on a particular point $p \in \mathbb{R}^d$. Notices that if $g(p) = 0$, then, $f(p) = 0$ as well. The inequality follows. Now, assume $g(p) > 0$ and define

$$M_l = \left\{ i \in \{1, ..., n\} : 2^{-l} \leq f_i(p) \leq 2^{-l+1} \right\}$$

for all $l : -\infty \leq l \leq \infty$.

Because that $\cup_{-\infty \leq l \leq \infty} M_l = \{i : f_i(p) \neq 0\}$ and M_l's are pairwise disjoint, each indices $i : f_i(p) \neq 0$ occurs in exactly one of M_l's. Let $m_l = |M_l|$. We claim $m_l \leq 6^d \mu$.

We now prove the claim. For each $i \in M_l$, by the definition of M_l and f_i, $2^{l-1} \leq \gamma_i \leq 2^l$, where $\gamma_i = r_i$. Let B be a ball centered at p with radius $2^l + 2^{l-1}$. Since $||p - p_i|| \leq \gamma_i$, it follows $B_i \subset B$. Because the neighborhood system has density μ, we have

$$\mu \cdot \text{vol}(B) \geq \sum_{j \in M_l} \text{vol}(B_j)$$

Let $V_d(r)$ be the volume of a ball in $\mathbb{R}_d$ of radius r. Because for all $j \in M_l$, $\text{vol}(B_j) \geq V_d(2^{l-2})$,

$$\mu \cdot V_d(2^l + 2^{l-1}) \geq |M_l| V_d(2^{l-2}),$$

which implies $|M_l| \leq 6^d \mu$, completing the proof of the claim.

Now, we have

$$\begin{aligned}
(f(p))^d &= \left(\sum_{l=-\infty}^{\infty} \sum_{i \in M_l} f_i(p)^{(d-1)}\right)^{\frac{d}{d-1}} \\
&\leq \left(\sum_{l=-\infty}^{\infty} m_l (2^{-l+1})^{d-1}\right)^{\frac{d}{d-1}} \\
&\leq 2^d \left(\sum_{l=-\infty}^{\infty} m_l (2^{-l})^{d-1}\right)^{\frac{d}{d-1}}
\end{aligned}$$

where $m_l \leq 6^d \mu$.

Setting $\theta = 6^d\mu$ and applying Lemma 6.7, we obtain

$$f(p)^d \leq c_d 2^d (6^d\mu)^{\frac{1}{d-1}} \sum_{l=-\infty}^{\infty} m_l 2^{-ld}$$

This summation is a lower bound on $g(p)^d$ because for each $i \in M_l$, $f_i(p)^d \geq 2^{-ld}$. This concludes the proof of the lemma. □

Consequently, there exists a $(\frac{d+1}{d+2})$-splitting sphere S of P with $\text{Area}_f(S) = O\left(6\xi\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right)$, where $\xi = 2^{d-1}V_d$. Theorem 6.1 follows from Lemmas 6.8, 6.5, and Theorem 5.9.

6.2 Separating overlap graphs

We now prove the following separator theorem of overlap graphs. The intuition and the method is basically the same as that for Theorem 6.1. The only things remaining are the technical details.

Theorem 6.10 (Overlap graphs) *If G is the α-overlap graph of a neighborhood system with density μ in d-space, then G has an $O(\alpha\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$-separator that $(\frac{d+1}{d+2})$-splits.*

6.2.1 A density function for overlap graphs

The first step to prove Theorem 6.10 is to construct a density function. Again, we first define a local density function for each neighborhood, and then use the L_{d-1} norm as the global density function.

Suppose $\Xi = \{B_1, \ldots, B_n\}$ be a neighborhood system of density μ. Let $P = \{p_1, \ldots, p_n\}$ be the centers of Ξ and r_i the radius of B_i. Let $\gamma_i = 2\alpha r_i$. Define f_i as

$$f_i(x) = \begin{cases} 1/\gamma_i & \text{if } x \in (2\alpha) \cdot B_i, \text{ i.e., } ||x - p_i|| \leq \gamma_i \\ 0 & \text{otherwise} \end{cases}$$

The global density function is then defined to be $f(x) = L_{(p-1)}(f_1, \ldots, f_n)$.

6.2.2 The total volume

Lemma 6.11 *For each neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ of density μ in $\mathbb{R}^d$, if $f_1, \ldots, f_n$ and f are defined as above, then*

$$\text{Total-Volume}(f) = O(\alpha^{\frac{d}{d-1}} \mu^{\frac{1}{d-1}} n).$$

Proof: Let V_d be the volume of a unit ball in $\mathbb{R}^d$. Clearly, $\int_{x \in \mathbb{R}^d} (f_i(x))^d (dx)^d = V_d$.

Consequently, letting

$$g(x) = L_d(f_1, \ldots, f_n) = \left(\sum_{i=1}^{n} (f_i(x))^d \right)^{\frac{1}{d}},$$

we have

$$\int_{x \in \mathbb{R}^d} (g(x))^d (dx)^d = V_d n$$

Therefore, Lemma 6.11 follows immediately from the following lemma. □

Lemma 6.12 *For all* $x \in \mathbb{R}^d$, $(g(x))^d \leq (f(x))^d \leq c_d 2^d (6^d \alpha^d \mu)^{\frac{1}{d-1}} \cdot (g(x))^d$.

Proof: The first inequality follows immediately from the definitions of f and g and Lemma 6.4.

For the second inequality, we focus on a particular point $p \in \mathbb{R}^d$. Notices that if $g(p) = 0$, then, $f(p) = 0$ as well. The inequality follows.

Now, assume $g(p) > 0$.

Define

$$M_l = \left\{ i \in \{1, ..., n\} : 2^{-l} \leq f_i(p) \leq 2^{-l+1} \right\}$$

for all $l : -\infty \leq l \leq \infty$.

Because that $\cup_{-\infty \leq l \leq \infty} M_l = \{i : f_i(p) \neq 0\}$ and M_l's are pairwise disjoint, each indices $i : f_i(p) \neq 0$ occurs in exactly one of M_l's. Let $m_l = |M_l|$. We claim $m_l \leq 6^d \alpha^d \mu$.

We now prove the claim. For each $i \in M_l$, by the definition of M_l and f_i, $2^{l-1} \leq \gamma_i \leq 2^l$. Let B be a ball centered at p with radius $2^l + 2^{l-1}$. Since $||p - p_i|| \leq \gamma_i$, it follows $B_i \subset B$. Because the density of the neighborhood system is μ,

$$\mu \cdot \text{vol}(B) \geq \sum_{j \in M_l} \text{vol}(B_j)$$

Let $V_d(r)$ be the volume of a ball in $\mathbb{R}_d$ of radius r. We have for all $j \in M_l$, $\text{vol}(B_j) \geq V_d(2^{(l-2)/\alpha})$. Consequently, $\mu \cdot V_d(2^l + 2^{l-1}) \geq |M_l| V_d(2^{(l-2)/\alpha})$, which implies $|M_l| \leq 6^d \alpha^d \mu$, completing the proof of the claim.

Now, we have

$$\begin{aligned}
(f(p))^d &= \left(\sum_{l=-\infty}^{\infty} \sum_{i \in M_l} f_i(p)^{(d-1)} \right)^{\frac{d}{d-1}} \\
&\leq \left(\sum_{l=-\infty}^{\infty} m_l (2^{-l+1})^{d-1} \right)^{\frac{d}{d-1}} \\
&\leq 2^d \left(\sum_{l=-\infty}^{\infty} m_l (2^{-l})^{d-1} \right)^{\frac{d}{d-1}}
\end{aligned}$$

where $m_l \leq 6^d \alpha^d \mu$.

Setting $\theta = 6^d \alpha^d \mu$ and applying Lemma 6.7, we obtain

$$f(p)^d \leq c_d 2^d (6^d \alpha^d \mu)^{\frac{1}{d-1}} \sum_{l=-\infty}^{\infty} m_l 2^{-ld}$$

This summation is a lower bound on $g(p)^d$ because for each $i \in M_l$, $f_i(p)^d \geq 2^{-ld}$. This concludes the proof of the lemma. □

Consequently, there exists a $(\frac{d+1}{d+2})$-splitting sphere S of P with $\text{Area}_f(S) = O\left(\xi \alpha \mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right)$. where $\xi = 2^{d-1} V_d$.

6.2.3 A vertex separator from a continuous one

This step is somewhat more involved for overlap graphs than for intersection graphs. Nevertheless, the intuition is the same. First notice that the set of balls intersecting S is not a vertex separator in general, because some neighborhoods are dilated; and the number of balls whose α-dilations intersect S can be as big as $\Omega(n)$ in the worst case (otherwise we would prove that all α-dilation graphs have small separators).

One way to construct a vertex separator is to remove one of the two balls of each edge cut by S. We say S *cuts* an edge (B_i, B_j) if the line segment between p_i and p_j, the centers of B_i and B_j, respectively, has a common point with S. Let $q_{i,j}$ denote this common point. For each edge (B_i, B_j) cut by S, without loss of generality, assume $r_i \leq r_j$. Notice that either $q_{i,j}$ is in B_j or in $\alpha \cdot B_i$. We call B_j an *overlap neighborhood* of S in the first case and B_i an *overlap neighborhood* in the second case. Immediately,

Lemma 6.13 *For all positive reals $\alpha \geq 1$ and $0 < \delta < 1$, let S be a sphere that δ-splits the centers of a neighborhood system Ξ, and D the set of all overlap neighborhoods of S. Then D δ-splits the α-overlap graph of Ξ.* □

However, in the above construction, one have to check the structure of an overlap graph to induce a vertex separator from a sphere separator, in contrast to the fact that a vertex separator of an intersection graph can be computed directly from the neighborhood system given a sphere separator. This is a serious drawback for various applications including mapping sparse structures to parallel architectures and computational geometry (see chapter 9). We now show how to induce a vertex separator from a sphere one more directly.

First notice that the set D above can be partitioned into two subsets D_1 and D_2, where

$$\begin{aligned} D_1 &= \{B_i \in D : B_i \cap S \neq \emptyset\} \\ D_2 &= D - D_1 \end{aligned}$$

Lemma 6.14 *For all $B_i \in D_2$, $r_i \leq r$.*

Proof: If $B_i \in D_2$, then $B_i \cap S = \emptyset$. There are two possible cases:

- **Case 1:** If $p_i \in \text{int}(S)$, then it simply follows from $B_i \cap S = \emptyset$, that $r_i \leq r$;
- **Case 2:** If $p_i \in \text{ext}(S)$, then from $B_i \in D_2$, it follows $\alpha \cdot B_i \cap S \neq \emptyset$, and there is a neighborhood system B_j such that (1) $p_j \in \text{int}(S)$; (2) $B_i \cap S = \emptyset$; (3) $r_i \leq r_j$; and (4) $\alpha \cdot B_i \cap B_j \neq \emptyset$. But $r_j \leq r$, and hence $r_j \leq r_j \leq r$. □

In the remaining of the thesis, a ball B_i is an *overlap neighborhood* of a sphere S if either $B_i \cap S \neq \emptyset$, or $\alpha \cdot B_i \cap S \neq \emptyset$ and $r_i \leq r$. The number of the overlap neighborhoods of S is called *overlap number* of S, denoted by $\vartheta_\Xi(S)$. By Lemma 6.14, Lemma 6.13 holds when the new definition of overlap neighborhoods is used. The set of overlap neighborhoods of a sphere can be computed in $O(n)$ time directly from the neighborhood system.

Lemma 6.15 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system with density μ in $\mathbb{R}^d$. Let $f_1, \ldots, f_n$ be local density function defined for the α-overlap graph of Ξ and let $f = L_{d-1}(f_1, \ldots, f_n)$. For each sphere S,*

$$\vartheta_\Xi(S) \leq O(\alpha^d)\mu + \left(\frac{4\sqrt{7}}{7}\right)^{d-1} \frac{1}{V_{d-1}} \text{Area}_f(S).$$

Proof: Let D be the set of overlap neighborhoods in Ξ which intersect S. D can be partitioned into two subsets $D = D_1 \cup D_2$, where

$$\begin{aligned} D_1 &= \{B_i \in D : \alpha r_i \leq r\} \\ D_2 &= \{B_i \in D : \alpha r_i \geq r\} \end{aligned}$$

By Lemma 3.6, $|D_2| \leq O(\alpha^d \mu)$. To bound the the size of D_1. first notice that

$$\text{Area}_f(S) = \int_{v \in S} \sum_{i=1}^{n} f_i(v)^{d-1}(dv)^{d-1} \geq \sum_{i \in D_1} \int_{v \in S} f_i(v))^{d-1}(dv)^{d-1}.$$

By the definition of D_1, for each $i \in D_1$, S has a common point with either B_i or αB_i. Further, $radius(S) \geq \gamma_i$. By lemma 6.2, the area of $S \cap (2 \cdot B_i)$ is at least $\left(\frac{\sqrt{7}}{4}\gamma_i\right)^{d-1} V_{d-1}$.

Therefore,

$$\int_{v \in S} (f_i(v))^{d-1}(dv)^{d-1} \geq \text{Area}(A \cap (2 \cdot B_i)) \left(\frac{1}{\gamma_i}\right)^{d-1} \geq \left(\frac{\sqrt{7}}{4}\right)^{d-1} V_{d-1}.$$

Hence

$$\text{Area}_f(S) \geq \sum_{i \in D_1} \int_{v \in S} f_i(v))^{d-1}(dv)^{d-1} \geq |D_1| \left(\frac{\sqrt{7}}{4}\right)^{d-1} V_{d-1}$$

Thus, $|D_1| \leq \left(\frac{4\sqrt{7}}{7}\right)^{d-1} \frac{1}{V_{d-1}} \text{Area}_f(S)$. □

Theorem 6.10 follows from Lemmas 6.11, 6.15, and Theorem 5.9.

6.3 Lower Bounds

Using a method described in Leighton [55], assuming that k is big with respect to d. we get a lower bound of $\Omega(k^{\frac{1}{d}} \alpha n^{\frac{d-1}{d}})$ on α-overlap graphs of k-neighborhood systems. An $\Omega\left(\alpha n^{\frac{d-1}{d}}\right)$ bound on α-density graph appeared in [94].

Let P be the set of all points of the $m \times m \times \cdots \times m$ regular grid in $\mathbb{R}^d$. We write each point $p \in P$ as a vector $(a_1, \ldots, a_d)$ of integers with $1 \leq a_i \leq m$.

Let l be the largest integer satisfying

$$V_d(l + \sqrt{d})^d \leq \frac{k}{2^d}, \tag{6.1}$$

where V_d denotes the volume of a unit ball in d-space given by (see [63, 40])

$$V_d = \frac{2\pi^{d/2}}{d\Gamma(d/2)}.$$

By (6.1), we have

$$l = \lfloor (\frac{k^{\frac{1}{d}}}{2(V_d)^{\frac{1}{d}}}) - \sqrt{d} \rfloor. \tag{6.2}$$

Consider the following graph G_k (see Figure 6.3 for examples), where $V(G_k) = P$ and

$$\begin{aligned} E(G_k) \;=\; & \{[(a_1,\ldots,a_d),(b_1,\ldots,b_d)]: \\ & \text{if there is } 1 \le i \le d \text{ such that } |a_i - b_i| \le l \text{ and } a_j = b_j \text{ for } j \ne i.\} \end{aligned}$$

Figure 6.3: G_2 in one dimension and G_4 in two dimensions

Lemma 6.16 *The graph G_k is a subgraph of the k-nearest neighborhood graph M_k of P.*

Proof: Let B be the ball of radius l centered at o and B' be the ball of $l+\sqrt{d}$ centered at o. The set of all integer points, those with all integer coordinates, induce an infinite set of unit cubes. Clearly, all cubes that intersect B must be contained in B'. Since the volume of B' is $V_d(l+\sqrt{d})^d$ the number of cubes intersect B is bounded from above by $V_d(1+\sqrt{d})^d$. Since each cube has 2^d integer points, it follows from the definition of l (6.1) that B contains no more than k integer points. By translating the center of B to each grid point in P, it follows that all grid points in B are among the k nearest neighbors of its center. Hence, G_k is a subgraph of the k-nearest neighborhood graph M_k on the points P. □

Clearly, the degree of G_k is $2dl$. We next show how, in the VLSI sense (i.e., mapping vertices onto vertices and edges to paths), to embed the complete graph K_n containing $n = m^d$ vertices onto G_k with edge conjession at most $\frac{m^{d+1}}{k^{\frac{1}{d}}}$, where *edge conjession* of an embedding is the maximum number of paths embedded on an edge.

We start with the case when the dimension is one. For each positive integer l, the *l-line graph* of n vertices (See Figure 6.4) has vertices $V = \{1,2,...,n\}$ and edges $E = \{(i,j) : \text{if } |i-j| \le l\}$. An edge (i,j) with $|i-j| = s$ $(1 \le s \le l)$ is called an *s-edge.*

Lemma 6.17 *The complete graph of n vertices, K_n, can be embedded in an l-line graph G of n vertices with conjession at most $\frac{n^2}{l^2}$.*

Proof: We prove the lemma by induction on l. When $l = 1$, we map edge (i,j) of the complete graph to the path $(i,i+1),(i+1,i+2),...,(j-1,j)$. Clearly, the edge in the 1-line graph with the largest conjession is the edge $(\lfloor \frac{n}{2} \rfloor, \lceil \frac{n}{2} \rceil)$, whose conjession is $(\lfloor \frac{n}{2} \rfloor \cdot \lceil \frac{n}{2} \rceil) \le n^2$.

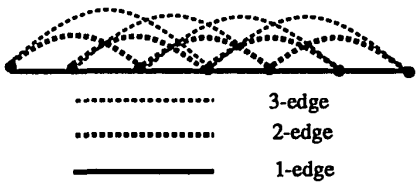

Figure 6.4: A 3-line graph

Now suppose the lemma is true for $l-1$. Let

$$n_1 = \lfloor \frac{(l-1)n}{l} \rfloor.$$

Let G' be the subgraph of G induced by vertices $\{1, 2, ..., n_1\}$. Applying the induction hypothesis on $l-1$ and n_1, we obtain an embedding of an n_1-clique on G', which does not use any l-edge of G'. To embed K_n on G, we take the above embedding on the first n_1 vertices, i.e., embedding the first n_1 vertices of K_n onto G'. For each remaining edge (i,j) with $i \leq j$, let j' be the largest integer such that $0 \leq j - j' \leq k$ and $(j'-i) \bmod k = 0$. To complete the embedding of the complete graph on to G, we embed each remaining edge (i,j) to the path $(i, i+k)(i+k, i+2k)\cdots(j'-k, j')(j', j)$. Clearly, each l-edge in G' has conjession at most $n_1(n-n_1)/l \leq \frac{n^2}{l^2}$ and each l-edge $(i, i+l)$ in $G - G'$ has conjession at most $i(n-i)/l$ which is also no more than $\frac{n^2}{l^2}$ because $i \geq n_1$. The edge (j', j) defined above has conjession at most j'/l which is also smaller than $\frac{n^2}{l^2}$. □

We now generalize Lemma 6.17 to higher dimensions.

Lemma 6.18 *There is a constant c depending only on d such that the complete graph of $n = m^d$ vertices, K_n, can be embedded in G_k with conjession at most $\frac{cm^{d+1}}{k^2}$.*

Proof: For simplicity, we view each vertex of K_n as a vector $(a_1, \ldots, a_d)$ of integers with $1 \leq a_i \leq m$. We first map K_n to G_1 by sending each pairs of vertices $(a_1, \ldots, a_d)$ and $(b_1, \ldots, b_d)$, to the natural path which goes through the vertices:

$$(a_1, \ldots, a_d), (b_1, a_2, \ldots, a_d), \cdots (b_1, \ldots, b_i, a_{i+1}, \cdots, a_d)$$

This embedding has conjession at most m^{d+1}. To see this, consider an edge

$$e = [(c_1, \ldots, c_i, \ldots, c_d), (c_1, \ldots, c_i + 1, \ldots, b_d)]$$

then the pairs of points (p, q) whose path contains e will be those of the following form:

$$[(a_1, \ldots, a_i, c_{i+1}, \ldots, c_d), (c_1, \ldots, c_{i-1}, b_i, \ldots, c_d)]$$

where $a_i > c_i \geq b_i$ or $b_i > c_i \geq a_i$. The number of such pairs is at most m^{d+1}.

Notice that each dimension of G_k is a collection of one dimensional G_l (see Figure 6.3), where l is given by (6.2). So, we can use Lemma 6.17 to embed edges along each dimension. Therefore, we reduce the conjession to at most $\frac{m^{k+1}}{l^2}$, which is $\frac{cm^{d+1}}{k^2}$ for an appropriate choice of constant c, because we have assumed that k is big with respect to d. □

Since each edge separator with constant splitting ratio of K_n must have size $\Omega(n^2)$, Lemma 6.18 implies a lower bound of $\Omega\left(\frac{n^2 k=frac2d}{m^{d+1}}\right) = \Omega\left(k^{\frac{2}{d}} m^{d-1}\right)$ on the edge separator of a d-dimensional G_k. Furthermore, because the degree of a d-dimensional G_k is $2dl$, the lower bound on the edge separator of G_k implies a lower bound of $\Omega(k^{\frac{1}{d}} m^{d-1})$ on the vertex separator of G_k. Therefore,

Corollary 6.19 *Any edge separator for G_k with constant splitting ratio must have size $\Omega(k^{\frac{2}{d}} m^{d-1})$ and any vertex separator for G_k with constant splitting ratio must have size $\Omega(k^{\frac{1}{d}} m^{d-1})$*

Consequently, because G_k is a subgraph of the k-nearest neighborhood graph of P, we get a lower bound of $\Omega(k^{\frac{1}{d}} n^{\frac{d-1}{d}})$ on the separator size of the intersection graphs of k-neighborhood systems. Moreover, for each k-neighborhood system Ξ with centers P, there is an $\alpha^d k$-neighborhood system Ξ' with centers P such that the α-overlap graph Ξ is a subgraph of the intersection graph of Ξ'. We thus get a lower bound of $\Omega(k^{\frac{1}{d}} \alpha n^{\frac{d-1}{d}})$ on the separator size of the α-overlap graphs of k-neighborhood systems. Therefore, the separator bounds in Theorem 6.1 and Theorem 6.10 are best possible in term of n, μ, and α. In summary,

Theorem 6.20 (Lower bounds) *For all positive integers d, k, and for each positive real $\alpha \geq 1$, there is a d-dimensional k-neighborhood system Ξ, letting G be the α-overlap graph of Ξ, such that each separator of G with a constant splitting ratio must have size $O\left(k^{\frac{1}{d}} \alpha n^{\frac{d-1}{d}}\right)$.*

Chapter 7

Graphs in Minkowski Spaces

Immediately, one may ask: does the separator theorem for k-nearest neighborhood graphs still hold if the L_1 norm or the L_∞ norm is used as metrics for the distance? Which metric space has the property that every α-overlap graph of a k-neighborhood system in it has a small separator?

7.1 Metric spaces

A *metric space* is a set $\mathcal{M}$ with a distance function $d_{\mathcal{M}}$ satisfying for all $x, y, z \in \mathcal{M}$,

1. $d_M(x, y) \geq 0$ and $d_M(x, y) = 0$ iff $x = y$;
2. $d_{\mathcal{M}}(x, y) = d_{\mathcal{M}}(y, x)$; (symmetry)
3. $d_{\mathcal{M}}(x, y) + d_{\mathcal{M}}(y, z) \geq d_{\mathcal{M}}(x, z)$. (triangle inequality)

Like an Euclidean neighborhood, each *$\mathcal{M}$-neighborhood* B has a center p and radius r, where B is the set of points in $\mathcal{M}$ whose distance from p is $\leq r$. A *neighborhood system* in $\mathcal{M}$ is a collection of neighborhoods. $\Xi = \{B_1, \ldots, B_n\}$ is a k-neighborhood system if the interior of each B_i contains no more than k points from P, where P is the set of centers of Ξ. The classes of intersection graphs, overlap graphs, and dilation graphs can be naturally extended to a metric space $\mathcal{M}$.

Clearly, not all metric space has a small separator theorem for its overlap graphs. An example is a *graph-based* metric space, $\mathcal{M}_G$, which contains all vertices of a graph G with the distance between each pair of vertices the length of the shortest path between the two vertices in G.

If G is a complete graph, then the intersection graph of a 1-neighborhood system is a complete graph, and hence it has no small separator. If G is a star graph with m vertices, however, the intersection graph of a 1-neighborhood system is either a graph which has no edges, or a star graph. Thus it has a small separator. While the intersection graph of a 2-neighborhood system is always a complete graph. For any integer k, let the *k-fuzz* of a graph G be the graph G_k obtained from G by adding an edge between vertices u and v if it is possible to go from u to v in at most k steps. It is not hard to see that if G has an $f(n)$-separator and the maximum degree is δ, then the k-fuzz of G has an $\delta^k f(n)$-separator. The small separator property of G can thus be extended to the metric space defined by G.

In this chapter, small separator theorems of Euclidean space are extended to normed linear spaces or Minkowski spaces. See below for the formal definitions. Consequently, for each positive

integer p, each k-nearest neighborhood graph in d-space with the L_p norm as the metric for the distance has an $O(k^{1/d}n^{(d-1)/d})$-separator that $\frac{d+1}{d+2}$-splits.

7.2 Overlap graphs in Minkowski spaces

Suppose $\{D, +, \cdot, 0\}$ be a linear space. $\mathcal{M}$ is a *normed linear space* of D provided there is associated with each element of D a nonnegative real number, called its *norm*, and denoted by $||x||$, such that

1. $||x|| = 0$ iff $x = 0$;
2. $\lambda \cdot x = |\lambda| \cdot ||x||$ for all $\lambda \in \mathbb{R}$;
3. $||x + y|| \leq ||x|| + ||y||$, for all pairs of $x, y \in D$.

The distance between any pair of elements x and y is defined to be $||x - y||$. Immediately from property (3), the distance yields a metric space.

The *dimension* of a linear space is the largest set of elements that are linearly independent. As shown in [14], each normed linear space with dimension $d < \infty$ is equivalent to a *Minkowski space* of dimension d.

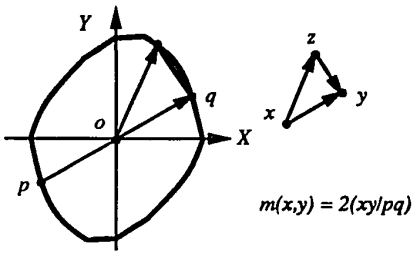

Figure 7.1: An example of Minkowski distance

Let Γ be the boundary of a bounded central symmetrical convex set in $\mathbb{R}^d$, whose center is at 0. The distance $m(x, y)$ of a pair of points x and $y \in \mathbb{R}^d$ is defined by $m(xy) = 2(xy/pq)$, where xy is the Euclidean distance of points x and y and pq is the Euclidean length of the chord of Γ contained in the line space through the center of Γ that is parallel to the line determined by x and y (See Figure 7.1). The corresponding space is denoted by Minkowski(Γ).

For example, the Euclidean space of dimension d is a Minkowski space with the choice of Γ to be a unit $(d-1)$-sphere; the L_∞ normed space is one with the choice of Γ the unit cube (See Figure 7.2), and in general the L_p normed space is the one with

$$\Gamma = \left\{ x : \sum_{i=1}^{d} (x_i)^d = 1 \right\}.$$

For each normed linear space $\mathcal{M}$ of dimension d, let Γ be the unit ball centered at 0. It can be shown that $\mathcal{M}$ is equivalent to Minkowski(Γ); and conversely, each Minkowski space in d dimensions, Minkowski(Γ), is a normed linear space of dimension d, whose unit ball is Γ (see Blumenthal and Menger [14] page 18–21 for a proof).

The following is the main result of this chapter.

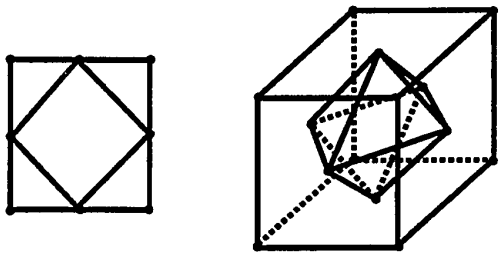

Figure 7.2: Cubes and co-cubes

Theorem 7.1 *Let $\mathcal{M}$ be a normed linear space with dimension d. For all positive integers μ and k, for each positive real $\alpha \geq 1$, for each d-dimensional neighborhood system Ξ with density μ, the α-overlap graph of Ξ has an $O\left(\mu^{1/d}\alpha n^{\frac{d-1}{d}}\right)$-separator that $\frac{d+1}{d+2}$-splits. Further, such a separator can be computed in random linear time.*

7.3 Some basic geometric properties

Define the *kissing number* of a normed linear space $\mathcal{M}$, $\tau_{\mathcal{M}}$, to be the maximum number of nonoverlapping unit $\mathcal{M}$-balls which can be arranged so that they all touch a central unit M-ball. Also, for each positive real δ, define $A_{\mathcal{M}}(\delta)$ to be the maximum number of points that can be arranged on a unit $\mathcal{M}$-sphere, such that the distance between each pair of points is at least δ. By a volume argument, $\tau_{\mathcal{M}} \leq 3^d - 1$ for each normed linear space $\mathcal{M}$ of dimension d.

Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system in a d-dimensional normed linear space $\mathcal{M}$. For each $p \in M$, define $\text{density}_{\Xi}(p)$ to be the number of neighborhoods in Ξ that contain p.

All proofs in Chapter 4 and Chapter 5 can be applied to normed linear space to obtain the following set of lemmas.

Lemma 7.2 (Density Lemma) *For each k-neighborhood system Ξ in a d-dimensional normed linear space $\mathcal{M}$, for each $p \in \mathbb{R}^d$, $\text{density}_{\Xi}(p) \leq \tau_{\mathcal{M}} k$.*

Lemma 7.3 (Ball Intersection) *Let $\Xi = \{B_1, \ldots, B_n\}$ be a k-neighborhood system in a d-dimensional normed linear space $\mathcal{M}$. Then for each ball $B \in \mathcal{M}$ (with center p and radius r), and for each α,*

$$|\{i : B_i \cap B \neq \emptyset \text{ and } p_i \in \mathcal{M} - \alpha \cdot B\}| \leq A_{\mathcal{M}}(\frac{\alpha - 1}{\alpha})k,$$

where p_i is the center of B_i $(1 \leq i \leq n)$.

Lemma 7.4 *The intersection graph of a k-neighborhood system in a normed linear space $\mathcal{M}$ is $2A_{\mathcal{M}}(1/2)k$-inductive.*

Lemma 7.5 *For each d, there is a constant c_d, which only depends on d, such that the α-overlap graph of a k-neighborhood system in a d-dimensional normed linear space $\mathcal{M}$ is $(c_d\alpha^d k)$-inductive.*

7.3.1 Linear transformation and the small separator property

We prove Theorem 7.1 in this section. First recall the steps to prove the counterpart of Theorem 7.1 in Euclidean spaces.

1. Dilate each neighborhood by a factor of 2α and assign a uniform local density within each dilated neighborhood;
2. Define the global density to be the L_{d-1} norm of the local ones;
3. Show that the global density is linearly related with the L_d norm of local ones at each points in $\mathbb{R}^d$, and hence bound the total volume of the global density;
4. Apply Theorem 5.9 to compute a sphere separator of low surface area;
5. Pull the sphere separator back to a vertex separator.

To prove Theorem 7.1, we follow the same outline. Let $\Xi = \{B_1, \ldots, B_n\}$ be a neighborhood system in $\mathcal{M}$ with density μ. Let r_i be the radius of B_i and let $\gamma_i = 2\alpha r_i$. Let f_i be uniform within $\gamma_i \cdot B_i$ so that the total volume of f_i is $V_{\mathcal{M}}$, the volume of a unit ball in $\mathcal{M}$.

Let $f(x) = L_{(d-1)}(f_1, \ldots, f_n)$ be the global density, and let $g = L_d(f_1, \ldots, f_n)$. Clearly, Total-volume$(g) = V_{\mathcal{M}} n$. Similar to Lemma 6.12 is the following lemma.

Lemma 7.6 *For all* $x \in \mathbb{R}^d$, $(g(x))^d \leq (f(x))^d \leq c_d 2^d (6^d \alpha^d \mu)^{1/(d-1)} \cdot (g(x))^d$.

Therefore, by Theorem 5.9, there is a $(d-1)$-dimensional **Euclidean** sphere of surface area $O(\mu^{1/d} \alpha n^{(d-1)/d})$ that $\frac{d+1}{d+2}$-splits the centers of Ξ.

To derive a small vertex separator from a sphere separator, we would like to use a similar construction of Section 6.2.3. But there is one problem – the unit ball in a normed linear space may have unbounded aspect-ratio, so that even a sphere separator intersects a neighborhood B, it may only pick up a very small amount of surface area from $2\alpha \cdot B$ (See Figure 7.3).

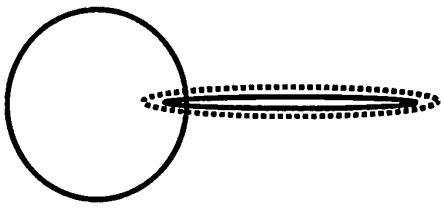

Figure 7.3: The problem with unbounded aspect-ratio

Fortunately, this problem can be solved using the observation that the structure of an α-overlap graph is preserved by a linear transformation.

Let $\mathcal{M}$ be a normed linear space with the norm $||x||$. Let L be a linear transformation, i.e., $L(x+y) = L(x)+L(y)$ for all $x, y \in \mathcal{M}$. Let $L(\mathcal{M})$ be the normed linear space transformed from $\mathcal{M}$ by L, with a norm denoted by $||x||_L$, where $||L(x)||_L = ||x||$. Clearly, L transforms a neighborhood B in $\mathcal{M}$ to one in $L(\mathcal{M})$, denoted by $L(B)$. In term of Minkowski space, Minkowski(Γ') is transformed from Minkowski(Γ) by a linear transformation L provided $\Gamma' = L(\Gamma)$.

Lemma 7.7 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system in $\mathcal{M}$ with density μ. Let L be a linear transformation of $\mathcal{M}$ and let $L(\Xi) = \{L(B_1), \cdots, L(B_n)\}$. Then $L(\Xi)$ is a neighborhood system in $L(\mathcal{M})$ with density μ; and the α-overlap graph Ξ is isomorphic to the α-overlap graph of $L(\Xi)$.*

Proof: The lemma follows straightforwardly from the claim that for each neighborhood B and a point $q \in \mathcal{M}$, $q \in \text{int}(B)$ iff $L(q) \in \text{int}(L(B))$.

Let p be the center of B and r the radius. The center of $L(B)$ is $L(p)$ and the radius is still r. Notice that $||L(q)||_L = ||q||$ and $||L(p) - L(q)||_L = ||L(p-q)||_L = ||p-q||$. The distance is preserved by L and hence the claim follows. □

Linear transformation can be used to relate all normed linear spaces.

Lemma 7.8 *For all bounded central symmetrical convex figure Γ, there is a linear transformation L such that*

$$CC \subseteq L(\Gamma) \subseteq C$$

where $CC = \{x : \sum_{i=1}^{d} |x_i| = 1\}$ is the unit co-cube *and $C = \{x : \max_{i=1}^{d} |x_i| = 1\}$, the unit cube (see Figure 7.4).*

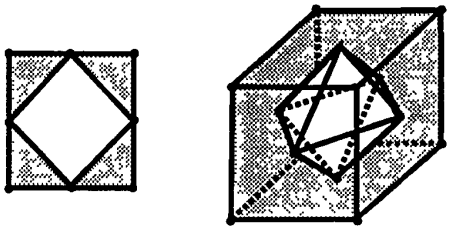

Figure 7.4: All symmetrical convex figures can be mapped between the cube and the co-cube

Proof: Here we give an informal description of a proof due to Lovasz [50].

Let $p_1, \ldots, p_d$ be d points on Γ such that the volume of the convex hull defined by $p_1, \ldots, p_d$ and o is maximum. Linearly transform p_i to e_i, the unit vector on the i^{th} dimension. Denoted the linear transformation by L. By the convexity and central symmetricity of Γ, we can conclude that $CC \subseteq L(\Gamma)$; by the maximality of the volume, we can deduce that $L(\Gamma) \subseteq C$. □

Therefore, without loss of generality, we can assume that all normed linear spaces under consideration have the property that their unit balls Γ satisfy the following inequalities.

$$CC \subseteq L(\Gamma) \subseteq C.$$

Lemma 7.9 *Let B be the co-cube in d dimensions with radius r and S a $(d-1)$-sphere of radius at least $2r$. If S intersects B, then the surface area of $S \cap (2 \cdot B)$ is at least $\left(\frac{\sqrt{7}-1}{2} r\right) V_{d-1}$, where V_{d-1} is the volume of a unit $(d-1)$-dimensional ball.*

Proof: For simplicity, we only prove the lemma when $d = 2$. The general case can be proved similarly. In the case when $d = 2$, the common area of S and B is just the length of the curve of S within B,

Assume S is a sphere with smallest common area with $2 \cdot B$ satisfying the condition of the lemma. The optimality of S implies that the radius of S is equal to $2r$. For otherwise, the sphere S', which has radius $2r$ and the same center as S, would have smaller common area with $2 \cdot B$ than that of S. The optimality also implies that the line passing the center of S and the center of B is parallel to x-coordinate axis and S touches B (See Figure 7.5).

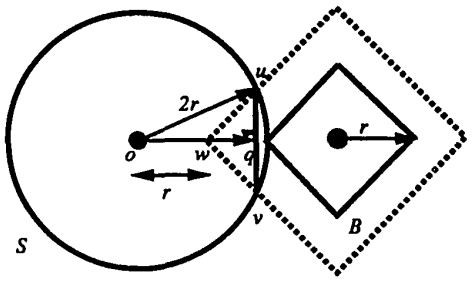

Figure 7.5: The common area of S and $2 \cdot B$.

Let u and v be the common points of S and the boundary $2 \cdot B$. Clearly, the length of the curve of S in B is longer than the length of the line segment between u and v. As shown in Figure 7.5, the line segment oq is perpendicular to the line segment uv and the length of the line segment ow is r. Since B is a 2-dimensional co-cube, the angle $\angle wuq = \pi/4$. Hence, the length of the line segment wq is equal to the length of the line segment uq. Let x be the length of uq, we have

$$(2r)^2 = x^2 + (r + x)^2,$$

which implies that $x = \frac{\sqrt{7}-1}{2}r$, and the length of the line segment of uv is $\sqrt{7} - 1$. □

Consequently,

Corollary 7.10 *Let Γ be a d-dimensional central symmetrical convex figure satisfying $CC \subseteq \Gamma \subseteq C$. Let B be a ball in* Minkowski(Γ) *with radius r and S a $(d-1)$-dimensional Euclidean sphere of radius at least $2r$. If S intersects B, then the surface area of $S \cap (2 \cdot B)$ is at least $\left(\frac{\sqrt{7}-1}{2}r\right) V_{d-1}$.*

Using a similar argument and construction as in the Euclidean space, we can deduce a vertex separator from a sphere separator, whose size is linearly bound by the surface area of its continuous counterpart, and thus Theorem 7.1 follows.

Chapter 8

Center Points and Point Divisions

The central algorithmic issue in computing sphere separators is the one of finding a center point of a given set of points (see chapter 5). Although many applications, such as computing sphere separators, are in fixed dimensions, some applications (for example in decision theory [53]) require efficiently computing a center point in variable dimensional spaces. By variable dimension I mean that the dimension is also specified as part of the input. For this reason, I will study this problem in both fixed and variable dimension.

I have proved some structural and complexity-theoretical properties of center points which are quite interesting. Although those results are somewhat peripheral to the main theorem of the thesis, I have decided to include them.

This chapter is self-contained. Results are given for computing, approximating, and testing center point and also for some related problems.

8.1 Introduction

Foe each positive real $0 < \alpha < 1$, a d-dimensional point c is an *α-center point* of a set of d-dimensional points P if every hyperplane containing c α-splits P. The set of all $\frac{d}{d+1}$-center points are *center points* of P, denoted by center(P).

Although the existence of center points follows simply from Helly's Theorem (Theorem 5.5), no efficient algorithm is known for computing center points. There is a straightforward method which requires solving a set of $\Theta(n^d)$ linear inequalities. The only nontrivial result, due to Cole, Sharir, and Yap [17], is that a center point in two dimension can be computed in $O(n \log^5 n)$ time, and in three dimensions $O(n^2 \log^7 n)$ time. No subquadratic algorithm is known that always returns an approximate center point.

A closely related problem is the one of testing whether a given point is a center point of a set of points. This problem will be referred as TESTING CENTER. The importance of this problem is illustrated in the following lemma.

Lemma 8.1 *The problem of computing a center point is polynomial time reducible to TESTING CENTER.*

Proof: The lemma follows from the general Ellipsoid method of Grötschel, Lovász, and Schrijver [40] that a point in a convex set can be found in polynomial time using the membership oracle. □

However, I will show that TESTING CENTER in variable dimensions is co-NP complete. As an interesting contrast, I will present a deterministic linear time algorithm for testing whether a point is an approximate center point in fixed dimensions.

The above co-NP completeness result motivates the following fundamental question: does there always exist a center point which has a succinct witness? By a succinct witness, I mean a space-efficient witness, with the help of it, we can efficiently determine whether the given point is a center point. If there is such a point, can it and its witness be computed in polynomial time?

This question is important for the following reasons.

1. Each point computed by a polynomial time center point algorithm has a short-proof of being a center point – such an algorithm itself serves as a proof. Therefore, the structure of center points with a short-proof can be used as a road-map for designing efficient center point algorithms.
2. Since the problem of testing center point is co-NP complete, it is especially important to design center point algorithms which also output a succinct proof of or a witness to the point it computes.

To prove a given point p is not a center point of P, it is enough to demonstrate a hyperplane whose splitting ratio is larger than $\frac{d}{d+1}$. Then, how to prove a point is a center point? One way to do this is to show that each hyperplane defined by p and $(d-1)$ other points of P has splitting ratio bounded by $d/(d+1)$. However, this requires checking $\Omega(n^{(d-1)})$ hyperplanes!

Suppose one have found a partition of P into $r = \lceil \frac{n}{d+1} \rceil$ disjoint subsets, $P_1, \ldots, P_r$ (magically), such that for all $i : 1 \leq i \leq r$, p is in the convex hull of P_i, $\text{conv}(P_i)$. Then he can claim that p is a center point of P. To see this, let H be a hyperplane containing p. For each $i : 1 \leq i \leq r$, since $p \in \text{conv}(P_i)$, there must be one point from P_i in each closed halfspace of H. This implies that there are at least r points in each closed half space of H. So, p is a center point. H. Tverberg [88] is the first who made this observation. The above discussion motivates the following definition.

Definition 8.2 (Point Division) *Suppose P is a set of n points in $\mathbb{R}^d$ and k be a positive integer. A point $p \in \mathbb{R}^d$ is* **k-divisible w.r.t** *P if P can be partitioned into k pairwise disjointed subsets $P_1, \ldots, P_k$, such that $p \in conv(P_i)$ for all $1 \leq i \leq k$. Such a k-partition $P_1, \ldots, P_k$ is said to be a* **k-division** *of P.*

Let $\text{tver}(P)$ be the set of all $\lceil \frac{n}{d+1} \rceil$-divisible points of P. It is known from Birch [10] that when $d = 2$, the convex hull of $\text{tver}(P)$ is equal to $\text{center}(P)$ and from a remarkable result of Tverberg [88] that $\text{tver}(P)$ is not empty (in all dimensions). However, it is not known whether in higher dimensions the convex hull of tverberg points is always equal to the set of center points.

Interestingly, the co-NP completeness result of TESTING CENTER implies that the equality does not always hold, unless NP = co-NP! This claim follows from the following lemma.

Lemma 8.3 *Each point in the convex hull of* $\text{tver}(P)$ *has a short proof of being a center point of P.*

Proof: It follows from the Caratheodory's theorem: when $P \subset \mathbb{R}^d$, each point of $\text{conv}(P)$ is a convex combination of $d+1$ (or fewer) points of P, that for each p in the convex hull of $\text{tver}(P)$,

there are $d+1$ points $p_1, \ldots, p_{d+1} \in \text{tver}(P)$ whose convex hull contains p. Therefore the $\frac{n}{d+1}$-divisions of p_i's together with the proof that p is in the convex hull of $p_1, \ldots, p_{d+1}$ yields a proof that p is a center point of P. □

In contrast, I will show that the problem of testing whether a point is a k-divisible point (or in tver(P)) is NP-complete in general. Together with the co-NP completeness TESTING CENTER, it implies an interesting structural picture of center points (see Figure 8.1).

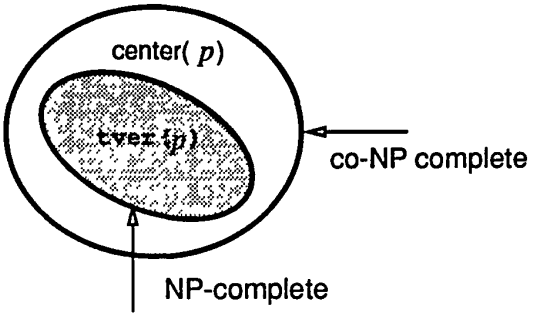

Figure 8.1: The Structure of Center Points

Geometrically, this means that if NP $\neq$ co-NP, then conv(tver(P)) is in general properly contained in center(P). Each point outside center(P) has a short proof of not being a center point of P, and each point in conv(tver(P)) has a short proof of being a center point. However, it is computational hard to prove whether a point is in center(P) or outside conv(tver(P)).

Define *DIVISIBILITY* to be the problem of testing whether a given point is k-divisible. When $k = 1$, this problem is known as the Caratheodory problem or the Convex Hull Problem. The Caratheodory problem is computationally equivalent to LINEAR PROGRAMMING [76] and hence can be solved in polynomial time. DIVISIBILITY with $k > 1$ is a natural generalization of the Caratheodory Problem. However, I will show that DIVISIBILITY is NP-complete even when $k = 2$.

This contrasts interestingly with the following observation: a 2-divisible point can be computed in polynomial time.

Let radon(P) be the set of all the 2-divisible points of P, called *Radon points.* It follows from Radon's proof of the Radon Theorem that a point in radon(P) can be computed in $O(d^3)$-time. But it is NP-complete to test whether a given point is a Radon point. This implies that the problem of computing a Radon point is polynomial time solvable but not polynomial time checkable (Blum and Kannan [13]).

In the second part of the chapter, efficient sequential and parallel algorithms for approximating center points and point-division in a fixed dimensions will be presented. Our algorithms can be either randomized or deterministic.

The randomized algorithm is based on the random sampling, and computes an approximate center point in random $O(\log n)$ time, using $n/\log n$ processors together with a division-type of proof of the point it computes.

Further, the concept of point division is generalized to the integer lattice. Lower bound and upper bound are given on the number of points for the existence of a k-divisible integer points.

8.2 The complexity of TESTING CENTER

Theorem 8.4 *In variable dimensions TESTING CENTER is co-NP complete.*

It is easy to see that TESTING CENTER is in co-NP. To prove it is complete for co-NP, the following NP-complete problem, the MAXIMUM 2-SATISFIABILITY, is shown to be polynomial time reducible to the problem of testing whether a point is not a center point. The later problem will be referred as TESTING NOT CENTER.

Problem 8.5 (MAXIMUM 2–SATISFIABILITY) *Given a positive integer N, a set U of m variables, and a collection C of s clauses over U such that each clause $c \in C$ has $|c| = 2$. Question: is there a truth assignment for U that simultaneously satisfies at least N clauses of C.*

It is known that MAXIMUM 2–SATISFIABILITY is NP-complete (See Garey, Johnson, and Stockmeyer [30]).

To reduce MAXIMUM 2–SATISFIABILITY to TESTING NOT CENTER, It will be shown that given an instance of first problem (N, U, C), a set of points P and a point p can be constructed in polynomial time such that p is a center point of P iff there is no truth assignment for U that simultaneously satisfies at least N clauses of C.

In the NP-completeness proof of the DENSEST HEMISPHERE PROBLEM (Johnson and Preparata [48]), Johnson and Preparata gave the following construction. Let $t = \lceil \log(ms + 1) \rceil$. The dimension d in the construction is $m + 1 + 3t$. For each $1 \leq i \leq m$, define

$$\begin{aligned} A_i &= (0)^{i-1}(1)(0)^{m-i}(1)\{1,-1\}^{3t} \\ \overline{A}_i &= (0)^{i-1}(-1)(0)^{m-i}(1)\{1,-1\}^{3t} \\ B_i &= (0)^{i-1}(4)(0)^{m-i}(-2)(0)^{2t}\{1,-1\}^{t} \\ \overline{B}_i &= (0)^{i-1}(-4)(0)^{m-i}(-2)(0)^{2t}\{1,-1\}^{t} \end{aligned}$$

In the above definition, the sets A_i, $\overline{A}_i$, B_i, and $\overline{B}_i$ are points corresponding to the i^{th} variable of U. Let $A = \bigcup_i (A_i \cup \overline{A}_i)$ and $B = \bigcup_i (B_i \cup \overline{B}_i)$.

Further, they define a set D which consists of one representative for each $c \in C$, as following. Let $x_j[1] = x_j$ and $x_j[-1] = \overline{x}_j$. Suppose $c = \{x_i[e_i], x_j[e_j]\}$, with $1 \leq i \leq j \leq m$ and $e_i, e_j \in \{1, -1\}$. The points in D corresponds c is

$$P_c = (0)^{i-1}(4e_i)(0)^{j-i}(4e_j)(0)^{m-j}(1)(0)^{3t},$$

Let $P' = A \cup B \cup D$ and $M = 2m \cdot 2^{3t} + m \cdot 2^t + N$. Let $o = (0)^d$ be the origin of the space. Johnson and Preparata proved the following lemma.

Lemma 8.6 (Johnson and Preparata) *There is a truth assignment for U that simultaneously satisfies at least N clauses of C iff there is a hyperplane passing o, such that one of its open halfspace contains more than M points of P'.*

Note that Johnson and Preparata's construction does not give a direct proof to Theorem 8.4, because $M >> \lceil \frac{d}{d+1} \rceil$. But some simple modification is sufficient.

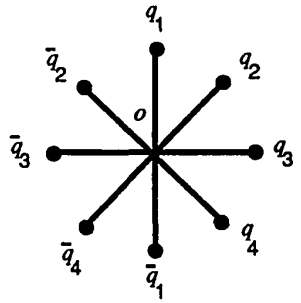

Figure 8.2: A 4-pencil

Let a *k-pencil* be $2k$ point $\{q_1, \ldots, q_k, \bar{q}_1, \ldots, \bar{q}_k\}$ such that for all $i : 1 \leq i \leq k$, $o \in \text{conv}(\{q_i, \bar{q}_i\})$. A k pencil has the property that each hyperplane passing o partitions the k–pencil evenly (see Figure 8.2).

Let k be an integer solution of the following equation.

$$2m \cdot 2^{3t} + m \cdot 2^t + N + k = \lceil (\frac{d}{d+1})(2m \cdot 2^{3t} + 2m \cdot 2^t + N + 2k) \rceil$$

Such solution always exists and is positive. Let Q be a k–pencil and let $P = P' \cup Q$. The following Lemma follows from Lemma 8.6 and property that each hyperplane passing o partitions the k–pencil evenly.

Lemma 8.7 *There is a truth assignment for U that simultaneously satisfies at least N clauses of C iff o is not a center point for P.*

Theorem 8.4 follows simply from the above lemma.

8.3 The complexity of DIVISIBILITY

Theorem 8.8 *DIVISIBILITY is NP-Complete even when $k = 2$ and $n = 2d$, where n is the number of points and d is the dimension.*

It is easy to see that DIVISIBILITY is in NP. To prove it is NP-complete, we need to show that there is an NP-complete problem which is polynomial time reducible to DIVISIBILITY. This time, the reduction is from the partition problem.

Suppose $A = \{a_1, \ldots, a_m\}$ be a set of m positive integers. A is *partitionable* if there is a subset $S \subset A$, such that

$$\sum_{a \in S} a = \sum_{b \in A - S} b.$$

The partition problem is formally defined as,

Problem 8.9 (PARTITION) *Given $A = \{a_1, \ldots, a_m\}$, a set of positive integers. Question: is A partitionable.*

It is knows from Karp ([52]) that PARTITION is NP-complete.

To prove PARTITION is polynomial time reducible to DIVISIBILITY, for each instance $A = \{A_1, \ldots, A_m\}$ of PARTITION, we will show that a set of $2m+1$ points $p_1, \ldots, p_{2m}$ and p in $(m+1)$-space can be constructed in polynomial time such that p is a radon point of $P = \{p_1, \ldots, p_{2m}\}$ iff A is partitionable.

Given an instance of the partition problem $A = \{a_1, \ldots, a_m\}$, let M be an $(m+1)$ by $2m$ matrix given below:.

$$M = \begin{pmatrix} 1_m & 0_m \\ I_m & I_m \end{pmatrix}$$

where $1_m = (\underbrace{1, 1, ..., 1}_{m})$, $0_m = (\underbrace{0, 0, ..., 0}_{m})$, and I_m is the m by m identity matrix.

Let p_i $(1 \leq i \leq 2m)$ be the i^{th} column of M – p_i can be viewed as a point in $(m+1)$-space.

Let $a = \sum_{i=1}^{m} a_i$, $b = (\frac{a}{2}, a_1, \ldots, a_m)$, and $p = \frac{1}{a}b = (\frac{1}{2}, \frac{a_1}{a}, ..., \frac{a_m}{a})$. Also p can be viewed as a point in $(m+1)$-space.

In the following subsection, I will prove the following lemma from which Theorem 8.8 follows.

Lemma 8.10 *A is partitionable iff p is a radon point of $p_1, \ldots, p_{2m}$.*

8.3.1 A connector: the transportation problem

To illustrate the relation between the partition problem and DIVISIBILITY, Let us look at the following classical combinatorial problem, the *transportation problem.*

The input to the transportation problem is a 2 by m complete bipartite graph G, with vertices $\{l_1, l_2\}$ on one side and $\{r_1, \ldots, r_m\}$ the other side. Also there are $\frac{1}{2}\sum_i^n a_i$ units of goods respectively at vertices l_1 and l_2, and there are a_i units of requirement from r_i. The transportation problem is to find how many units should be shipped on each edge (l_i, r_j), for all $1 \leq i \leq 2$ and $1 \leq j \leq m$ (see Figure 8.3).

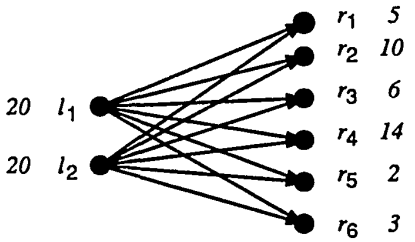

Figure 8.3: A 2 by m Transportation Problem

Computationally, the above transportation problem is not very hard. The greedy method works just fine. But what is useful and interesting here is that this problem serves as a bridge connecting the partition problem with DIVISIBILITY. Let $x_{i,j}$ denote the units shipped on edge (l_i, r_j) in a solution of the transportation problem. It is not hard to see that

$$X = (x_{1,1}, ..., x_{1,m}, x_{2,1}, ..., x_{2,m})$$

is a solution of the following linear system (8.1).

$$M \cdot X^T = \begin{pmatrix} 1_m & 0_m \\ I_m & I_m \end{pmatrix} \begin{pmatrix} x_{1,1} \\ \vdots \\ x_{1,m} \\ x_{2,1} \\ \vdots \\ x_{2,m} \end{pmatrix} = b = \begin{pmatrix} a/2 \\ a_1 \\ \vdots \\ a_m \end{pmatrix} \tag{8.1}$$

The solution of linear system (8.1) has the following properties.

Lemma 8.11

1. *Each solution of linear system (8.1) has at least m nonzero entries;*
2. *If there are exactly m entries in a solution of linear system (8.1) are nonzero, then no entry is negative.*

Proof: Statement (1) follows simply from the observation that one of $x_{1,i}$ and $x_{2,i}$ must be nonzero; We now prove statement (2). In a solution to linear system (8.1) which contains exactly m nonzero entries, it must be the case that exactly one of $x_{1,i}$ and $x_{2,i}$ are nonzero, $(1 \leq i \leq m)$. Moreover, the nonzero one must be equal to a_i which is positive. □

The following lemma show the relation between linear system (8.1) and the partition problem.

Lemma 8.12 *The linear system (8.1) has a solution which contains exactly m nonzero entries iff $A = \{a_1, \ldots, a_m\}$ is partitionable.*

Proof: ($\Rightarrow$) Suppose that the linear system (8.1) has a solution X which contains exactly m nonzero entries. Let $R_1 = \{i | x_{1,i} = a_i \neq 0\}$. Clearly,

$$\sum_{i \in R_1} a_i = \sum_{j=1}^{m} x_{1,j} = \frac{1}{2} \sum_{i=1} a_i$$

Therefore, A is partitionable.

($\Leftarrow$) Now assume that A is partitionable. Let $S \subset A$ be such a solution. Let $R_1 = \{i | a_i \in S\}$ and $R_2 = \{i | a_i \notin S\}$. It is not hard to see that

$$\begin{aligned} x_{1,i} &= \begin{cases} a_i & \text{if } i \in R_1 \\ 0 & \text{if } i \in R_2 \end{cases} \\ x_{2,i} &= \begin{cases} 0 & \text{if } i \in R_1 \\ a_i & \text{if } i \in R_2 \end{cases}, \end{aligned}$$

is a solution to the linear system (8.1) and contains exactly m nonzero entries. □

Therefore, Lemma 8.10 and Theorem 8.8 follows from the following lemma.

Lemma 8.13 *The linear system (8.1) has a solution which contains exactly m nonzero entries iff p is a radon point of $p_1, \ldots, p_{2m}$.*

Proof: ($\Rightarrow$) Suppose that p is a radon point of $p_1, \ldots, p_{2m}$. Then there is $U_1 \subset \{p_1, \ldots, p_{2n}\}$ and $U_2 = \{p_1, \ldots, p_{2m}\} - U_1$, such that $p \in \text{conv}(U_1)$ and $p \in \text{conv}(U_2)$. Hence, there are $0 \leq \alpha_1, \ldots, \alpha_{2m} \leq 1$ such that $\sum_{i:p_i \in U_1} \alpha_i = 1$, $\sum_{i:p_i \in U_2} \alpha_i = 1$, and

$$\begin{aligned} \sum_{i:p_i \in U_1} \alpha_i p_i &= p \\ \sum_{i:p_i \in U_2} \alpha_i p_i &= p \end{aligned}$$

Let

$$\begin{aligned} y_i &= \begin{cases} a\alpha_i & \text{if } p_i \in U_1 \\ 0 & \text{otherwise} \end{cases} \\ z_i &= \begin{cases} a\alpha_i & \text{if } p_i \in U_2 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Clearly, both $(y_1, \ldots, y_{2m})$ and $(z_1, \ldots, z_{2m})$ are solutions of the linear system (8.1). Since the total number of nonzero entries in $y's$ and $z's$ are $2m$, it follows from Lemma 8.11 that both solutions have exactly m nonzero entries.

($\Leftarrow$) Suppose that linear system (8.1) has a solution $x = (x_{1,1}, \ldots, x_{1,m}, x_{2,1} \ldots, x_{2,m})$ which has exactly m nonzero entries. Let $\overline{x} = (\overline{x}_{1,1}, \ldots, \overline{x}_{2,2n})$ where for each $1 \leq j \leq 2$ and $1 \leq i \leq m$,

$$\overline{x}_{j,i} = \begin{cases} 0 & \text{if } x_{j,i} = a_i \\ a_i & x_{j,i} = 0 \end{cases}$$

Clearly, $\overline{x}$ is also a solution to linear system (8.1).

For notation simplicity, $x_i = x_{k_i,l_i}$ and $\overline{x}_i = \overline{x}_{k_i,l_i}$, where $(k_i - 1)m + l_i = i$. Note that k_i $(1 \leq k_i \leq 2)$ and l_i $(1 \leq l_i \leq n)$ are uniquely determined by $i : 1 \leq i \leq 2n$.

Let $U_1 = \{p_i | x_i = a_{l_i}\}$ and $U_2 = \{p_i | \overline{x}_i = a_{l_i}\}$. It is not hard to see that $U_2 = \{p_1, \ldots, p_{2n}\} - U_1$, and the above discussion implies that

$$\begin{aligned} \sum_{i:p_i \in U_1} x_i &= a \\ \sum_{i:p_i \in U_2} \overline{x}_i &= a \\ \sum_{i:p_i \in U_1} x_i p_i &= ap \\ \sum_{i:p_i \in U_2} \overline{x}_i p_i &= ap \end{aligned}$$

We now show that $p \in \text{conv}(U_1)$ and $p \in \text{conv}(U_2)$. Let $\alpha_i = x_i/a$ and $\beta_i = \overline{x}_i/a$ $(1 \leq i \leq 2n)$. It follows from $x_i \geq 0$ and $\overline{x}_i \geq 0$, that $\alpha_i \geq 0$, $\beta_i \geq 0$, and

$$\begin{aligned} \sum_{i:p_i \in U_1} \alpha_i &= 1 \\ \sum_{i:p_i \in U_2} \beta_i &= 1 \end{aligned}$$

$$\sum_{i:p_i \in U_1} \alpha_i p_i = p$$
$$\sum_{i:p_i \in U_2} \beta_i p_i = p,$$

which implies that $p \in \text{conv}(U_1)$ and $p \in \text{conv}(U_2)$ and hence p is a radon point of $\{p_1, \ldots, p_{2m}\}$. □

This completes the proof of Theorem 8.8. Similarly,

Theorem 8.14 *The problem of testing whether a point p is a tverberg point of a given set of points $P = \{p_1, \ldots, p_n\}$ is NP-complete.* □

8.3.2 Computing a Radon point

The NP-completeness of DIVISIBILITY in the case when $k = 2$ contrasts interestingly with the fact that a radon point can be computed in $O(d^3)$ time, where d is the dimension of the space.

Theorem 8.15 (Radon) *For all $P \geq d+2$, a radon point can be computed in $O(d^3)$ time.*

Proof: Suppose $P = \{p_1, \ldots, p_n\}$ with $n \geq d+2$. Consider the system of $n+1$ homogeneous linear equations

$$\sum_{i=1}^{n} \alpha = 0 = \sum_{i=1}^{m} \alpha p_i^j \qquad (i \leq j \leq d),$$

where $p_i = (p_i^1, \ldots, p_i^d)$ in the usual coordinate system of $\mathbb{R}^d$. Since $n \geq d+2$, the system has a nontrivial solution $(\alpha_1, \ldots, \alpha_n)$. Let U be the set of all i for which $\alpha_i \geq 0$, and V the set for which $\alpha < 0$, and $c = \sum_{i \in U} \alpha_i > 0$. Then $\sum_{i \in V} \alpha_i = -c$ and $\sum_{i \in U} (\frac{\alpha_i}{c}) p_i = \sum_{i \in V} \frac{1 \alpha_i}{c} p_i$.

Note that to compute a radon point of P, it is sufficient to compute a 2-division for the first $d+2$ point. Hence, a radon point can be computed in $O(d^3)$ time. □

As a consequence of both Theorem 8.8 and 8.15,

Corollary 8.16 *The problem of computing a radon point is polynomial time solvable but not polynomial time checkable.*

8.4 Point division versus center points

The complexity results in the above sections imply that it is unlikely the convex hull of tverberg points is always equal to the set of center points. But is there a proof without using any computational assumption?

It is not hard to see that in 1-space, the convex hull of tverberg points is precisely the set of all center points. To see this, considering the following two cases: If P has odd number of points, then there is only one center point, the median of P, which is also the only tverberg point of P; If P consists of even number of points (say $P = \{p_1, \ldots, p_{2n}\}$ with $p_i \leq p_{i+1}$), then the set of center points is the closed interval between p_n and p_{n+1}. But both p_n and p_{n+1} are tverberg of P.

Then what is the smallest dimension that the convex hull of tverberg points is not always equal to the set of center points?

Suppose there are $3n$ points $P = \{p_1, \ldots, p_{3n}\}$ in plane (2-space) and p is a center point of P. Let θ_i be the angle between the ray pp_i and the x-coordinate. Without loss of generality assume $\theta_i \leq \theta_{i+1}$. It can be easily shown that if p is a center point of P, then p is in the convex hull of points p_i, p_{n+i}, p_{2n+i}, for all $1 \leq i \leq n$. This was first observed by Birch [10]. Therefore,

Lemma 8.17 (Birch) *For each P of $3n$ points in $\mathbb{R}^2$, $\text{tver}(P) = \text{center}(P)$.*

Now suppose there are $3n+1$ points $P = \{p_1, \ldots, p_{3n+1}\}$ in the plane. It is implied in the proof of Lemma 5.7 that center(P) is a convex set and for each vertex c of center(P) there are four points p_i, p_j, p_k, p_l from P such that c is a common point of line segments between p_i, p_j and p_k, p_l. So c is also a center point of points $P - \{p_i, p_j, p_k, p_l\}$. Therefore, by lemma 8.17, c is a tverberg of P. Similarly, the same assertion can be proved when $3n+2$ points. Thus,

Corollary 8.18 *For any set P of points in $\mathbb{R}^2$, the convex hull of* tver(P) *is equal to* center(P).

However, using the intuition behind the proof of Theorem 8.8, I have found the following example which implies that it is not always true the convex hull of tverberg points is equal to the set of center points in 3-space.

Example 8.19 *Let $P = \{p_1, p_2, p_3, p_4, p_5, p_6\}$, where*

$$\begin{aligned} p_1 &= (1,0,0) \\ p_2 &= (0,1,0) \\ p_3 &= (0,0,1) \\ p_4 &= (0,-1,-1) \\ p_5 &= (-1,0,-1) \\ p_6 &= (-1,-1,0). \end{aligned}$$

Then the convex hull of tver(P) $\neq$ center$P$.

8.5 Approximating center points

Now we have shown that the problem of testing whether a point is a center point is co-NP complete. One may wonder how much of a hint the co-NP complete result gives about the complexity of computing a center point? Can we then conclude that it is computational hard to compute a center point in variable dimension? Unfortunately, I do not know the answer but I do not quite think so – the NP-completeness of testing a radon point does not mean too much for computing a radon point.

Let us study the center point problem in fixed dimensions. After all, fixed dimensions are what we mainly care for applications of small separators.

8.5.1 Computing an approximate center point in fixed dimensions

In the remainder of this section, I will use the following notations.

Definition 8.20 (ϵ-center points) *A point p is an* ϵ-center point *of X if all hyperplanes containing p have splitting ratio $\leq d/(d+1)+\epsilon$, where d is the dimension.*

The randomized algorithm uses random sampling [16, 75, 44] - an important algorithmic technique for designing efficient geometric algorithms. To illustrate the idea, we show how to use random sampling to compute an approximate center point in 1-dimension. In this case, the input is a set of n integers $P = \{p_1, \ldots, p_{2n}\}$, without loss of generality assume $p_i < p_j$. Clearly, center$(P) = [p_n, p_{n+1}]$.

Now suppose we randomly select an element from P, say p. The probability that $p \in \{p_n, p_{n+1}\}$ is $\frac{1}{n}$, while the probability that $p_{\lceil \frac{n}{2} \rceil} \leq p \leq p_{\lceil \frac{3n}{2} \rceil}$ is 0.5. So, with probability 0.5, a randomly selected element from P is an $\epsilon = \frac{1}{4}$ center point.

We can improve ϵ using larger samples! Suppose l random elements $S = \{r_1, \ldots, r_l\}$ are selected and their median r is the output. Let $I(r)$ be the rank of r in P, it follows from a simple analyse that $E[I(r)] = n$, and $V[I(r)] = \frac{(2n+1)(2n-1-2l)}{8k+6}$. By Chebyshev's inequality,

$$P(|I(r) - n|) \geq t) \leq \frac{n^2}{2lt^2}$$

Thus, with probability at least 0.5, $|I(r) - n| \leq \frac{n}{\sqrt{l}}$, i.e., r is a $\frac{1}{2\sqrt{l}}$ center point of $|P|$. A $\frac{1}{2\sqrt{l}}$ center point of $|P|$ can be computed in $O(l)$ time.

The above algorithm can be generalized to higher dimensions. In d dimensions, the randomized ϵ-center point algorithm has the following form:

Algorithm A: Random ϵ–Center Point
Input $P \subset \mathbb{R}^d$

1. Randomly (uniformly) select l points from P, say $S \subseteq P$;
2. Compute a center point c_S of S;
3. Output c_S.

The feasibility of the above algorithm is specified in the following question: What is the probability that c_S computed above is an ϵ-center point?

We now introduce a notation which will be very useful in quantifying the quality of the c_S computed by the above sampling algorithm.

Definition 8.21 (ϵ-Good Sample) *For each $P \in \mathbb{R}^d$, $S \subseteq P$ is an ϵ-good sample if*

$$\sup_{h \in HS} \left| \frac{|h \cap S|}{|S|} - \frac{|h \cap P|}{|P|} \right| \leq \epsilon$$

Lemma 8.22 *For each $P \in \mathbb{R}^d$, if $S \subseteq P$ is an ϵ–good sample, then each center point of S is an ϵ–center point of P.*

Proof: Let c_S be a center point of S. Suppose c_S is not an ϵ center point of P, then there is a half-space $h \in HS$ passing c_S with

$$\frac{|h \cap P|}{|P|} > \frac{d}{d+1} + \epsilon.$$

On the other hand, c_S is a center point of of S, thus

$$\frac{|h \cap S|}{|S|} \leq \frac{d}{d+1}.$$

This implies that

$$\frac{|h \cap P|}{|P|} - \frac{|h \cap S|}{|S|} > \frac{d}{d+1} + \epsilon - \frac{d}{d+1} = \epsilon,$$

contradicting with the assumption that S is an ϵ-good sample. □

Now the question becomes: how often a set of l randomly chosen points forms an ϵ-good sample? This is not a trivial question, but was in fact answered many many years ago by Vapnik and Chervonenkis [93].

8.5.2 A theory of Vapnik and Chervonenkis

Let $\mathcal{X}$ be a set called *domain*, $\mathcal{S}$ a set of subsets over $\mathcal{X}$, and $\mathcal{P}_\mathcal{X}$ a fixed distribution over $\mathcal{X}$. Let $\mathcal{X}^{(l)}$ denote the space of samples in $\mathcal{X}$ of size l. On the space $\mathcal{X}^{(l)}$, a probability measure $\mathcal{P}$ is defined as, for each $(y_1, \ldots, y_l)$ in $\mathcal{X}^{(l)}$,

$$\mathcal{P}[(y_1, \ldots, y_l)] = \prod_i^l \mathcal{P}_\mathcal{X}(y_i)$$

For any $\{x_1, \ldots, x_l\} \in \mathcal{X}^{(l)}$, each $\mathcal{A} \in \mathcal{S}$ *induces* a subsample, $S_\mathcal{A} = \{x_i \in \mathcal{A}\}$. Let $\Delta^\mathcal{S}(x_1, \ldots, x_l)$ be the number of all different subsamples induced by the set of $\mathcal{S}$ with respect to the sample $\{x_1, \ldots, x_l\}$. The *growth function* $m^\mathcal{S}(l)$ is defined as

$$m^\mathcal{S}(l) = \max_{\mathcal{X}^{(l)}} \Delta^\mathcal{S}(x_1, \ldots, x_l)$$

The *VC dimension* of $\mathcal{S}$, denoted by $\mathcal{VC}(\mathcal{S})$, is defined to be the constant that

$$\mathcal{VC}(\mathcal{S}) = \min\{l : m^\mathcal{S}(l) \neq 2^l\}.$$

For each $\mathcal{A} \in \mathcal{S}$, let $\mathcal{P}_\mathcal{A} = \mathcal{P}(\{x | x \in \mathcal{A}\})$. Each sample $(x_1, \ldots, x_l)$ and $\mathcal{A} \in \mathcal{S}$ determines a relative frequency, $v_\mathcal{A}^{(l)}(x_1, \ldots, x_l)$, which is equal to

$$\frac{|\{i : x_i \in \mathcal{A}\}|}{l}.$$

Let $\pi^{(l)} = \sup_{\mathcal{A} \in \mathcal{S}} |v_\mathcal{A}^{(l)} - P_\mathcal{A}|$, the maximum difference over the class $\mathcal{S}$ between relative frequency and probability. The quantity $\pi^{(l)}$ is a point function in $\mathcal{X}^l$.

Note that if $\mathcal{S} = \{\mathcal{A}\}$, then the Bernoulli's Theorem states that $\lim_{l \to \infty} P[|v_\mathcal{A}^{(l)} - P_a|] = 0$. Vapnik and Chervonenkis generalized the Bernoulli's Theorem and related the convergence of $\pi^{(l)}$ to the VC-dimension of $\mathcal{S}$. They prove the following result.

Theorem 8.23 (Vapnik and Chervonenkis) *For each* $0 < \epsilon \leq 1$ *and* $l \geq \frac{2}{\epsilon^2}$,

$$P[\pi^{(l)} > \epsilon] \leq 4 m^\mathcal{S}(2l) e^{\frac{-\epsilon^2 l}{8}}$$

Corollary 8.24 *If the* $\mathcal{VC}(\mathcal{S}) = d$, *then* $P[\pi^{(l)} > \epsilon] \leq \eta$, *provided*

$$l \leq \frac{16}{\epsilon^2}\left(d \log \frac{16n}{\epsilon^2} + \log \frac{4}{\eta}\right) = O\left(d \frac{1}{\epsilon^2} \log \frac{n}{\epsilon} + \log \frac{1}{\eta}\right)$$

8.5.3 The probability of ϵ-good-samples

Let H be the set of all half-spaces in d space. For each set of points $P \subseteq \mathbb{R}^d$, let $H(P) = \{h \cap P | h \in H\}$. It is not hard to see that for each $P \subset \mathbb{R}^d$, the VC dimension of $H(P)$ is bounded above by $d+1$. For each $P \subset \mathbb{R}^d$, let $\mathcal{X} = P$, $\mathcal{S} = H(P)$, and $\mathcal{P}_{\mathcal{X}}$ be the uniform distribution over $\mathcal{X}$. Let $\pi^{(l)} = \sup_{A \in \mathcal{S}} |v_A^{(l)} - \mathcal{P}_A|$ be the random variable of maximum difference over class $\mathcal{S}$ between the relative frequency and probability. It follows from the definition of ϵ-good sample that

Lemma 8.25 *$P[S \subseteq P$, with $|S| = l$, is an ϵ-good sample$] = P[\pi^{(l)} \leq \epsilon]$.*

Consequently,

Theorem 8.26 (Algorithm A) *For all $P \in \mathbb{R}^d$, an ϵ-center point of P can be computed in*

$$O\left(\left[\frac{d}{\epsilon^2}\log\frac{d}{\epsilon} + \log\frac{1}{\eta}\right]^d\right)$$

time, with probability at least $1-\eta$.

Notice that all computation above can be efficiently implemented in parallel.

8.5.4 Testing ϵ-center points

The above random sampling algorithm has a major drawback. It does not always output an ϵ-center point. One way to overcome this problem is to design efficient algorithm to test whether a given point is an ϵ-center point. Note that the problem to testing whether a point is an ϵ-center point in variable dimensions is again co-NP complete. The proof is just a simple modification of the one for Theorem 8.4. When the dimension is fixed, there is a straightforward algorithm which runs in $O(n^{d-1})$ time, which is too slow for any practical application.

We now present an efficient algorithm which is based on point division and random sampling. The intuition behind the algorithm is expressed in the following lemma.

Lemma 8.27 *Suppose P is partitioned into r disjointed subsets $P_1, \ldots, P_r$. Then if p is an ϵ-center point of all $P_1, \ldots, P_r$, respectively, then p is an ϵ-center point of P.*

The above lemma motivates that following concept. Suppose P is a set of points in $\mathbb{R}^d$. A partition $(P_1, \ldots, P_r)$ is an *(ϵ, k)-partition of P*, if there is a point p such that $|P_i| \leq k$ and p is an ϵ-center point of P_i for all $1 \leq i \leq r$. Such a partition $(P_1, \ldots, P_r)$ is also called an (ϵ, k)-partition **w.r.t** p.

So if an (ϵ, k)-partition **w.r.t** p could be computed such that k is of manageable size, then the $O(k^{d-1})$-time algorithm can be used to prove that p is an ϵ-center point of each member of the partition and hence p is an ϵ-center point of P. However, no deterministic method known to generate an (ϵ, k)-partition efficiently. Fortunately, a random partition works well with high probability.

A partition $(P_1, \ldots, P_r)$ is a *regular k-partition* iff $k-1 \leq |P_i| \leq k$, for all $i : 1 \leq i \leq r$. Clearly, $r = \lfloor n/k \rfloor$. The following lemma estimates the probability that a random regular k-partition is an (ϵ, k)-partition.

Lemma 8.28 *For each $P \subset \mathbb{R}^d$, for each $1 \leq \epsilon \leq 1$, and for all $k \geq \frac{2}{\epsilon^2}$, the probability that a random k-partition is an (ϵ, k)-partition is at least $1 - 8n(2k)^d e^{\frac{-\epsilon^2 k}{8}}$.*

Proof: Note that if S is a random $\lceil n/2 \rceil$-sample of P, then $P - S$ is also a random $\lfloor n/2 \rfloor$–sample of P. Similarly, if $(P_1, \ldots, P_r)$ is a random regular k-partition of P, then each P_i is a random $|P_i|$ sample of P. Therefore, the probability that $(P_1, \ldots, P_r)$ is a regular (ϵ, k)-partition is equal to the probability that all P_i are ϵ-good-sample, which is at least r times the probability that a random k–sample is an ϵ-good-sample. Thus, the lemma follows from Theorem 8.23. □.

So we can efficiently generate an (ϵ, k)-partition of P. But how often it is also an (ϵ, k)-partition w.r.t an δ center point p? This question is partially answered in the following lemma which follows directly from the definition of ϵ-good-sample.

Lemma 8.29 *If S is an ϵ-good-sample, then each δ center point of P is an $(\epsilon + \delta)$-center point of S.*

Consequently,

Corollary 8.30 *Let $(P_1, \ldots, P_r)$ be an (ϵ, k)-partition of P and p a δ-center point of P. Then $(P_1, \ldots, P_r)$ is a $(\delta + \epsilon, k)$-partition* w.r.t *p.*

Simply from Corollary 8.30 and Lemma 8.28 and the fact that a center point of n points in d-space can be computed in $O(n^d)$ time is the following observation.

Observation 8.31 *For all $P \subset \mathbb{R}^d$, for each ϵ center point p of P, for each $\delta = O(\frac{1}{polylog(n)})$, and for all $\eta = O(\frac{1}{poly(n)})$, a proof that p is an $(\delta + \epsilon)$ center point of P can be computed in $O(n \cdot polylog(n))$ time sequentially, and $O(\log n)$ time, using in $O(n \cdot polylog(n))$ processors, with probability at least $1 - \eta$.*

Combining with the randomized algorithm of the last section, we can in random $O(\log n)$ time, compute an ϵ center point and an division type of proof, using $O(n \cdot polylog(n))$ processors.

8.5.5 An improved algorithm

In this subsection, we present a random linear time algorithm for computing an ϵ center point. This algorithm is always correct. It never outputs a point which in not an ϵ center point. The idea of this algorithm will used in the next section to design a deterministic linear time algorithm for computing an ϵ center point.

By Lemma 8.22, if S is an ϵ-good-sample of P, then the center point of S is an ϵ center point of P. So it is sufficient to design a random linear time algorithm which always outputs an ϵ-good-sample of size $O(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon})$.

One method is to test whether a randomly selected sample S is an ϵ-good sample of P. However, we know of no efficient way to do this. There is a straightforward $O(n^d)$-time algorithm which simply tests whether the condition of the ϵ-good-sample holds for each $(d-1)$-dimensional hyperplane defined by d points of P, which is too slow for any practical purpose. But we will apply this algorithm on sets of small number of points. The following lemma is used to reduce the size of the problem.

Lemma 8.32 (Union Lemma) *Suppose P is partitioned into r disjoint subsets $P_1, \ldots, P_r$ and S_i is an ϵ-good-sample of P_i $(1 \leq i \leq r)$. Then $(\cup_{i=1}^{r} S_i)$ is also an ϵ-good-sample of P.*

The lemma itself is very simple and follows simply from the definition of ϵ-good-sample. It suggests the following procedure for computing an ϵ-good-sample.

Procedure

1. Partition P into r disjointed subsets $P_1, \ldots, P_r$ of about equal size;
2. Randomly select a sample S_i for P_i and prove S_i is an ϵ-good-sample using the $O(n^d)$-time algorithm;
3. If S_i fails to be an ϵ-good-sample of P_i, then choose another random sample and prove again until succeed;
4. Output $S = \cup_{i=1}^{r} S_i$.

By Lemma 8.32, S is an ϵ-good-sample of P. If all subsets P_i's $(1 \leq i \leq r)$ are of constant size, then the above procedure runs in random linear time. But, the number of elements in S is fairly big. In fact it is a constant fraction of n. To further reduce the size of the sample, we can apply the above procedure on S. However, this might decrease the quality of the sample. Then how bad can it be? This question is answered in the following lemma which follows straightforwardly from the definition of good sample.

Lemma 8.33 (Composition Lemma) *Let P be a set of points in d-space, S_1 an ϵ-good-sample of P, and S_2 a δ-good-sample of S_1. Then S_2 is an $(\epsilon + \delta)$-good-sample of P.*

Now, to compute a small ϵ-good-sample, we just repeatedly apply above procedure. The key point to maintain the quality of the sample is the following observation: The problem size is reduced after each round and hence we can spend more time in finding better samples. To guarantee that the quality of the final sample is as good as that of the top level sample, the quality factor is improved at each round geometrically, i.e., if the quality factor at the i^{th} round is ϵ_i, then the quality factor of the $(i+1)^{st}$ round is $\epsilon_{i+1} = \epsilon_i/2$. By the Composition Lemma above, the set obtained after iterative applications of the above procedure is an $\sum_{i=1} \epsilon/2^i \leq 2\epsilon_1$ good sample of P.

Let us now work out the detail number. To compute an ϵ_i-good-sample randomly, we have to choose a sample of size $O(\frac{1}{\epsilon_i^2} \log \frac{1}{\epsilon_i})$. Let U be the constant of the above big-O notation. Let l_i denote the size of subsets in the i^{th} round; let s_i be the size of random sample for each subset in the i^{th} round; and let n_i be the total number of points at the beginning of the i^{th} round. The expected time of the i^{th} round is $T_i = \frac{n_i}{l_i} l_i^d = n_i l_i^{d-1}$. Let $\gamma = \gamma_i = l_i/s_i$. We have $n_i = n_{i-1}/\gamma_i$.

To maintain the quality of the sample, s_i is chosen as

$$s_i = \lceil U \frac{1}{\epsilon_i^2} \log \frac{1}{\epsilon_i} \rceil.$$

So $s_i \leq 5 s_{i-1}$. Let $\gamma = 2 \cdot 5^{d-1}$, and let $l_i = \gamma s_i$.

The expect time T_i of each round is bounded as below.

$$\begin{aligned} T_i &= n_i l_i^{d-1} \\ &= \frac{n_{i-1}}{\gamma_i} l_i^{d-1} \\ &\leq \frac{n_{i-1}}{\gamma_i} (5l_{i-1})^{l-1} \\ &\leq (n_{i-1}(l_{i-1})^{d-1})/2 \\ &\leq T_{i-1}/2 \end{aligned}$$

Therefore the total time is bounded by $\sum_i T_i \leq 2T_1$. However, The above approach has one serious problem. The sample size of the i^{th} round is five times as big as that of the $(i-1)^{st}$ round. But this problem can be fixed. The idea is decrease the sample size s_i after a predefined round. Of course, this is done at the cost of reducing the quality factor of the sample. But this will not cause any problem if it is done in a controlled fashion. More specifically, let j be the smallest integer such that $s_1 = n/\gamma^j$. After the $(j/2)^{th}$ round, the sample size is decreased, but in a backward fashion. For $i \geq j/2$, assign $\epsilon_i = 2\epsilon_{i-1}$, and $s_i = s_i = \lceil U \frac{1}{\epsilon_i^2} \log \frac{1}{\epsilon_i} \rceil$, and $l_i = l_{i-1}/\gamma$. (See Figure 8.4).

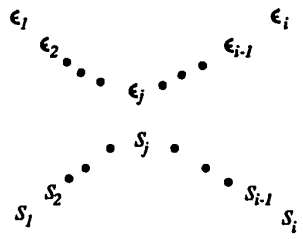

Figure 8.4: The change of ϵ_i and s_i

By the similar time analysis, $T_{i+1} \leq 2T_i$. The algorithm stops when $\epsilon_i = \epsilon_1$. Hence the total complexity is bounded by $4T_1$ which is linear in n if s_1 is a constant. The number of points left is approximately equal to s_1. Thus if the algorithm starts with $\epsilon_1 = \epsilon/4$, then it outputs an ϵ-good-simple of size $O(\frac{1}{\epsilon^2} \log \frac{1}{\epsilon})$.

Theorem 8.34 *There is a random linear time algorithm which always outputs an ϵ center point.*

Notice that the above algorithm can be easily parallelized. It can be implement on an CRCW PRAM in $O(\log n)$ time, using n processors.

8.5.6 A deterministic algorithm

In this section, we present a deterministic linear time algorithm for computing an ϵ-center point. This algorithm is basically a de-randomization version of the random one given in the last section.

I have to point out that Matoušek [62] recently gave an linear time deterministic algorithm for this problem independently. The basic idea of his algorithm is quite similar to ours. But his algorithm is more general – it can be applied to all abstract range spaces with bounded VC dimensions. Because his result has already been published in conference proceedings, I refer interesting readers to his work. Here I will only present the high level idea of our construction.

The result of the last section can be phrased as: if there is an $O(n^{poly(d)})$ time algorithm for computing an ϵ-good-sample of size $O(\frac{1}{\epsilon^2}\log\frac{1}{\epsilon})$, where n is the size of inputs and d the dimension, then there is an linear time algorithm for computing an ϵ-center point in d-space. Moreover, if the former algorithm is deterministic, then the later one is also deterministic.

We first notice that the problem of computing small ϵ-good-sample in fixed dimensions can be expressed by a integer linear programming with n variables and n^d inequalities.

Let $P = \{p_1, \ldots, p_n\}$ be the set of input points. Let x_i be a boolean variable which is 1 if p_i is chosen in the sample and 0 otherwise. Clearly, for each hyperplane, the condition of ϵ-good sample can be expressed by a linear inequality. Since, it is sufficient to check n^d hyperplanes, those which are defined by d points from P. To compute a small ϵ-good-sample, the sum $\sum_{i=1}^{n} x_i$ is minimized provided all inequality holds.

Using the general derandomization technique of Raghavan [74] and the fact that linear programming can be solved in polynomial time, an small ϵ-good-sample can be computed in $O(n^{poly(d)})$ time.

Theorem 8.35 *There is a deterministic linear time algorithm for computing an ϵ-center point in fixed dimensions.*

8.6 Point divisions on the integer lattice

In this section the point division problem is investigated over a natural convexity space $\mathcal{Z}^d$. On one hand, it will be shown that each set of $rd2^d$ points in $\mathcal{Z}^d$ has an r-division, i.e., it can be partitioned into r subsets, such that the convex hull of the r disjointed subsets has a common integer point. On the other hand, the argument of Shmuel Onn is generalized to show that for each r, there is a set of $(r-1)2^d+1$ points which does not have an r-division.

8.6.1 An upper bound

Definition 8.36 (α-Center Points) *A point $c \in \mathcal{Z}^d$ is an α-center point of a finite set of points $P \subset \mathcal{Z}^d$ provides every hyperplane containing c has a splitting ratio no more than α. The set of α-center points of P is denoted by α-center(P).*

Similar to the relation between center points and division point, the following basic relation holds between α-center point and r-division in integer lattice.

Lemma 8.37 *For each set $P \subset \mathcal{Z}^d$,*

1. *If $(P_1, \ldots, P_r)$ is an r division of P, and $c \in (\bigcap_{i=1}^{r} conv(P_i)) \bigcap \mathcal{Z}^d$, then c is a $(1 - \frac{r}{|P|})$ center point of P;*
2. *If $c \in \mathcal{Z}^d$ is an α-center point of P, then there is a $\lfloor \frac{(1-\alpha)|P|}{d} \rfloor$ division of P such that c is contained in the convex hull of all the subsets.*

Proof: Statement 1 follows simply from the definitions of α-center points and r-division. Statement 2 can be easily proved by applying the Caratheodory theorem. □

The upper bound result is stated in the following theorem.

Theorem 8.38 *Each set $P \in \mathcal{Z}^d$ of $rd2^d$ points has an r division.*

By Lemma 8.37, it is sufficient to prove the following lemma.

Lemma 8.39 *Each set $P \in \mathcal{Z}^d$ of $rd2^d$ points has an $\frac{2^d-1}{2^d}$-center point.*

Proof: Let H be the set of all half-spaces that contain more than $\frac{2^d-1}{2^d}$ points from P. We will show

$$\left(\bigcap_{h\in H} h\right) \bigcap Z^d \neq \emptyset \tag{8.2}$$

We use the following result due to Bell and Scarf [78].

Proposition 8.40 (Bell and Scarf) *Let F be a finite collection of closed half spaces in $\mathbb{R}^d$. If $(\cap F)\cap \mathcal{Z}^d = \emptyset$, then there exists a subfamily $L \subset F$ such that $|L| \leq 2^d$ and $(\cap L) \cap \mathcal{Z}^d = \emptyset$.*

Since P is finite, there is a finite subset $H' \subset H$ such that $\cap H' = \cap H$.

Assume relation (8.2) is not true, i.e., $(\cap H') \cap Z^d = \emptyset$. By Proposition 8.40, there is $L \subset H'$ such that $|L| \leq 2^d$ and $(\cap L) \cap \mathcal{Z}^d = \emptyset$. For each $h \in L$, let $h' = \mathbb{R}^d - h$. We have $|h' \cap X| < \frac{1}{2^d}|X|$. Thus $|(\cup_{h\in L} h') \cap X^d| < |X|$, which implies $|X| - (\cup_{h\in L} h') \neq \emptyset$. A contradiction to our assumption. Therefore the lemma holds. □

8.6.2 A lower bound

Let $\delta(r, d)$ denote the smallest integer such that each set of $\delta(r, d)$ points in $\mathcal{Z}^d$ has an r-division.

The result in this section is motivated by the result of Onn [70] who gave a lower on the radon number of the integer lattice. Since the radon partition is a 2-division. The Onn's lower bound can be specified as

Theorem 8.41 (Onn) *For all $d \geq 2$, $\delta(2, d) \geq 2^d + 2^{d-1} + 1$.*

We now generalize the Onn's lower bound to all r.

Theorem 8.42 (Division Number: A Lower Bound) *For all $d \geq 1$, for all r,*

$$\delta(r, d) \geq 2^d r - 2^d + 1.$$

Proof: The claimed lower bound follows from the following inequality.

$$\delta(r, d+1) \geq 2\delta(r, d) - 1 \tag{8.3}$$

Assume $\delta(r, d+1) \leq 2\delta(r, d) - 2$. Let $A \subset \mathcal{Z}^d$, $|A| = \delta(r, d) - 1$ be a set that has no r-division. Let

$$X_i = \{(a, i) : a \in A\} \qquad (i = 0, 1).$$

Let $X = X_0 \cup X_1 \in \mathcal{Z}^{d+1}$. Since $|X| = 2(\delta(r, d) - 1) \geq \delta(r, d+1)$, it follows from the assumption that X has an r-division, say $(Y_1, \ldots, Y_r)$. Let c be an integer division point of $(Y_1, \ldots, Y_r)$, i.e., $c \in \mathcal{Z}^{d+1}$ and $c \in \cap_{i=1}^r (\text{conv}(Y_i))$. So, $c \in (\mathcal{Z}^d \times \{0, 1\})$, i.e., if $c = (c_1, \ldots, c_{d+1})$, then $c_{d+1} \in \{0, 1\}$. Without loss of generality, assume $c_{d+1} = 0$. Then let $Z_i = Y_i \cap X_0$, it follows that $c \in \cap_{i=1}^r (\text{conv}(Y_i))$. This contradict with the assumption that A does not have an r division. Hence, the Inequality (8.3) holds. Note that for all r, $\delta(r, 1) = 2r - 1$. It follows by a simple induction on d that $\delta(r, d) \geq 2^d r - 2^d + 1$. □

8.6.3 On $4n$ points in the planar integer lattice

In 2-dimensional integer lattice, the upper bound of the division number can be tighten. We shall show that each set of $4n$ lattice points has an n-division. The proof is based on the idea of Birch [10].

Theorem 8.43 *Each set $Z \in \mathcal{Z}^2$ of $4n$ points has an n-division.*

Proof: It follows from Lemma 8.39 that each $Z \in \mathcal{Z}^2$ (of $4n$ points) has an $\frac{3}{4}$ center point. Let $c \in \mathcal{Z}^2$ be a $\frac{3}{4}$ center point of Z. Since division and center point are preserved by any affine transformation. Without loss of generality assume $c = (0,0)$. Each point $z = (x,y) \in Z$ determines a vector $\vec{zc}$. Let θ_z be the angle formed by $\vec{zc}$ and $\vec{(1,0)c}$, the x coordinate. Let $Z = \{z_1, \ldots, z_{4n}\}$ be a sorted list of Z by θ_i, the angle of $\vec{z_ic}$ and x ray. That is $\theta_i \leq \theta_{i+1}$.

Define an n-partition $(Z_0, ..., Z_{n-1})$ of Z by $Z_i = \{z_j | j \equiv i \bmod n\}$. We claim that $(Z_0, ..., Z_{n-1})$ forms an n-division of Z. More specifically, we shall prove that for all $0 \leq i \leq n-1$, $c \in \text{conv}(Z_i)$.

Let $Z_i = \{p_0, p_1, p_3, p_3\}$, where $p_j = z_{(i+jn)}$, $0 \leq j \leq 3$. Let h_j be the closed half-space (containing c) defined by the line passing p_j and $pj+1 \bmod 4$. Since $Z = \{z_1, \ldots, z_{4n}\}$ is a sorted list of Z by θ_i, the angle of $\vec{z_ic}$ and x ray. Hence,

$$h_j \cap Z \supseteq Z - \{z_{i+jn+1 \bmod 4n}, z_{i+jn+2 \bmod 4n}, \cdots, z_{i+(j+1)n-1 \bmod 4n}\},$$

which implies $|h_j| > \frac{3n}{4}$. It follows that c is an $\frac{3}{4}$ center point that $c \in h_j$, for all $0 \leq j \leq 3$, thus $c \in \text{conv}(Z_i)$. □

However, if $|Z| \neq 4n$, it is not true that every $\frac{3}{4}$-center point is a $\lceil |Z|/4 \rceil$-division point. Moreover, there is a set Z, $|Z| \not\equiv 4$, such that the convex hull of all $\lceil |Z|/4 \rceil$ division points does not contain all integer $\frac{3}{4}$-center points of Z. This contrast interestingly with the relation between center points and division points in $\mathbb{R}^2$, where the convex hull of $\lceil |Z|/3 \rceil$-division points is equal to the set of $\frac{2}{3}$-center points.

Chapter 9

Separator Based Divide and Conquer

In this chapter, a divide and conquer paradigm is presented based on the separator results proved in previous chapters. The usefulness of this paradigm in computational geometry is demonstrated. The new paradigm is compared with a commonly used paradigm for solving geometry problems – the multi-dimensional divide and conquer of Bentley [8]. It will be shown that this paradigm outperforms the multi-dimensional divide and conquer on various geometry problems. The new paradigm also provides a good support for designing efficient parallel algorithms for geometry problems in fixed dimensions.

In this chapter, we are more concerned with expressing the concept of the separator based divide and conquer paradigm than examining the detail of a particular algorithm. For this reason, we will omit some more sophisticated but asymptotically faster algorithms which are based on the same paradigm.

In Section 9.1, we present an efficient search structure for a geometric query problem. This search structure will be used in subsequent sections. In Section 9.2, we develop a separator based divide and conquer algorithm for constructing the intersection graph of a given neighborhood system. We also give an efficient parallel implementation of this algorithm. This algorithm will be modified in Section 9.3 to compute the overlap graph of a neighborhood system. In Section 9.4, we present a separator based divide and conquer algorithm for computing k-nearest neighborhood graphs in fixed dimensions. This algorithm can be implemented on parallel machines more efficiently than the known parallelization of Vaidya's algorithm [91].

9.1 The neighborhood query problem

The *neighborhood query problem* is defined as: given a neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ in d-space, preprocess the input to organize it into a search structure so that queries of the form "output all neighborhoods that contain a given point p" can be answered efficiently.

Like other geometry query problems, there are three costs associated with the neighborhood query problem: the *preprocessing* time $T(n,d)$ required to build the search structure, the *query* time $Q(n,d)$ required to answer a query, and the *space* $S(n,d)$ required to represent the search structure in memory.

If Ξ is an arbitrary neighborhood system, then there may exist some point p such that the density of p – the number of neighborhoods that cover p – is $\Omega(n)$. In this case, just to print the

output would require $\Omega(n)$ work. However, if Ξ is restricted to be a k-neighborhood system or a neighborhood system with density μ, then the number of balls in the output is bounded by $\tau_d k$ or μ, respectively. Can the neighborhood query problem be solved more efficiently?

To our knowledge, there are no prior result on this problem. Nevertheless, it is relatively straightforward that the multi-dimensional divide and conquer yields a search structure with the following properties.

$$\begin{aligned} T(n,d) &= O(n \log^{d-1} n) \\ Q(n,d) &= O(\mu + \log^d n) \\ S(n,d) &= O(n \log^{d-1} n) \end{aligned}$$

In contrast, using separator based divide and conquer, we are able to construct a search structure with the following properties.

$$\begin{aligned} T(n,d) &= \text{random } O(n \log n) \\ Q(n,d) &= O(\mu + \log n) \\ S(n,d) &= O(n) \end{aligned}$$

By saying an algorithm runs in random $t(n)$ time we mean that the algorithm never give a wrong output but may not terminate. The probability of success, namely, it produces a correct output in $t(n)$ steps, is at least $1 - \frac{1}{\log n}$. The probability of success of randomized algorithms in this chapter can be easily amplified to $1 - \frac{1}{n}$ at the cost an $O(\log \log n)$ factor increasing in the run time. More sophisticate analysis can be used to eliminate this extra $O(\log \log n)$ overhead (see Frieze, Miller, Teng [28]). To simplify the discussion, in the following sections, we assume that the density μ is a constant.

9.1.1 A separator based search structure

The main idea is to use a sphere separator which intersects substantial small number of balls to partition the neighborhoods into two subsets of roughly equal size, and then recursively build search structures for each subsets.

Given a neighborhood system Ξ with density μ, we will build a binary tree of height $O(\log n)$ to guide the search in answering a query. Associated with each leaf of the tree is a subset of neighborhoods in Ξ, and the search structure has the property that for all $p \in \mathbb{R}^d$, the set of neighborhoods that covers p can be found in one of the leaves.

In the following construction, any sublinear separator with constant splitting ratio can be used. The asymptotical complexity is the same. Hence, I will not refer to any particular constant. All that matter is that β, δ and c are constants with the property that $0 < \beta < 1$, $0 < \delta < 1$ and c is a positive real that only depends on the dimension d. If not further specified, it is assumed that $\beta = \frac{d-1}{d}$, $\delta = \frac{d+1}{d+2}$, and c is the constant term in Theorem 5.3.

Let S be a sphere separator with intersection number $c \cdot n^\beta$ that δ-splits centers of Ξ. Let Ξ_0 be the subset of neighborhoods which intersect either S or the interior of S; and Ξ_1 the subset of neighborhoods which intersect either S or the exterior of S. Clearly $|\Xi_0|, |\Xi_1| \leq \delta n + c \cdot n^\beta$, and $|\Xi_1| + |\Xi_2| \leq n + c \cdot n^\beta$. We store the information of S, its center and radius, in the root of the search tree, and recursively build binary search trees for Ξ_0 and Ξ_1, respectively. The roots of the

tree for Ξ_0 and Ξ_1 are respectively the left and right children of the node associated with S. The recursive construction stops when the subset has cardinality smaller than m_0, for a parameter m_0 satisfying $cm_0^\beta \leq \left(\frac{1-\delta}{2}\right) m_0$, to be specified later.

To answer a query when given a point $p \in \mathbb{R}^d$, we first check p against S, the sphere separator associated with the root of the search tree. There are three cases:

- **Case 1**: if p is in the interior of S then recursively search on the left subtree of S;
- **Case 2**: if p is in the exterior of S then recursively search the right subtree of S;
- **Case 3**: if p is on S then recursively search on the left subtree of S.

When reaching a leaf, we then check p against all neighborhoods associated with the leaf and print all those which cover p.

The correctness of the search structure and the above searching procedure is obvious and can be proved by induction: if p is in the interior (exterior) of S, then all neighborhoods which cover p must intersect either S or the interior (exterior) of S, and hence are in the left (right) subtree of S. The time complexity to answer a query is clearly bounded by $O(h(n) + m_0)$, where $h(n)$ is the height of the search tree, given by the following recurrence.

$$h(m) \leq \begin{cases} 1 & \text{if } m \leq m_0 \\ h(\delta m + cm^\beta) + 1 & \text{if } m \geq m_0. \end{cases} \tag{9.1}$$

Lemma 9.1 *Let h be a function defined above. Then $h(n) = O(\log n)$.*

Proof: By definition of m_0, for all $m \geq m_0$, $cm^\beta \leq \left(\frac{1-\delta}{2}\right) m$. Because $h(m)$ is a nondecreasing function in m, we have

$$h(m) \leq \begin{cases} 1 & \text{if } m \leq m_0 \\ h\left(\left(\frac{1+\delta}{2}\right)m\right) + 1 & \text{if } m > m_0. \end{cases}$$

Since $\delta < 1$ and hence $\frac{2}{1+\delta} > 1$, we can infer $h(n) = \lceil \log_{\frac{2}{1+\delta}} n \rceil = O(\log n)$. □

Consequently,

$$Q(n, d) = O(\log n + m_0).$$

We now analyze the space requirement of the search structure. First observe that each internal node requires a constant amount of space and each leaf requires $O(m_0)$ space. To bound the total space, it is sufficient to bound the total number of leaves in the tree.

Let $s(m)$ denote the number of leaves in the search tree for m neighborhoods. Clearly $s(m)$ is given by the following recurrence.

$$s(m) \leq \begin{cases} 1 & \text{if } m \leq m_0 \\ s(\delta_1 m + cm^\beta) + s((1-\delta_1)m) & \text{if } m > m_0, \end{cases} \tag{9.2}$$

where $1 - \delta_1 \leq \delta_1 \leq \delta$.

Lemma 9.2 *Let s be the function defined above. Then $s(n) = O(n)$ for a sufficient large constant m_0 that depends only on d, δ, β, and c.*

Proof: For any constant $\gamma : \beta < \gamma < 1$, we use induction to establish $s(n) \leq C(n - n^\gamma)$ for a sufficient large constant m_0 that depends only on d, δ, β, and c and an appropriate choice of the constant $C > 1$.

Because $s(m) = 1$ for $m \leq m_0$, we need to choose C, and m_0 such that $C(m_0 - m_0^\gamma) \geq 1$. Because $\gamma < 1$, this condition holds for a sufficient large m_0, establishing the base for the induction. Now assume that the lemma is true for all $m < n$, using the substitution method of [18], we have

$$\begin{aligned} s(n) &\leq s(\delta_1 n + cn^\beta) + s((1-\delta_1)n) \\ &\leq C(\delta_1 n + cn^\beta) + C((1-\delta_1)n) - C(\delta_1 n + cn^\beta)^\gamma - C((1-\delta_1)n)^\gamma \quad \text{(By induction)} \\ &= Cn - Cn^\gamma + Cn^\gamma + Ccn^\beta - C(\delta_1 n + cn^\beta)^\gamma - C((1-\delta_1)n)^\gamma \\ &\leq Cn - Cn^\gamma + Cn^\gamma + Ccn^\beta - C(\delta_1 n)^\gamma - C((1-\delta_1)n)^\gamma \\ &\leq C(n - n^\gamma), \end{aligned}$$

as long as we choose C and m_0 such that

$$C(\delta_1 n)^\gamma + C((1-\delta_1)n)^\gamma - Cn^\gamma - Ccn^\beta \geq 0. \tag{9.3}$$

By Taylor expansion of around the point 1, we have

$$(\delta_1)^\gamma + (1-\delta_1)^\gamma \geq (\delta_1)^\gamma + 1 - \gamma\delta_1.$$

Because $0 < \gamma < 1$ and $1 < \delta_1 < \delta < 1$, the above inequality (9.3) holds if

$$C[(\delta^\gamma - \gamma\delta)n^\gamma - cn^\beta] \geq 0. \tag{9.4}$$

Because $0 < \beta < \gamma < 1$, inequality 9.4 holds for all

$$n \geq \left(\frac{c}{\delta^\gamma - \gamma\delta}\right)^{\frac{1}{\gamma-\beta}}.$$

Therefore, the lemma hold for a sufficient large constant m_0 and C which only depend on d, δ, β and c. □

Because the number of nodes in a proper binary tree is no more than twice the number of leaves, we have for all sufficient large m_0 satisfying $cm_0^\beta \leq \left(\frac{1-\delta}{2}\right) m_0$ and the condition of Lemma 9.2,

$$s(n) = O(n/m_0).$$

Therefore, the total space requirement of the above search structure is bounded by

$$S(n, d) = O(n).$$

Now let us look at the time required in building such a search structure.

From the previous chapters, each d-dimensional neighborhood system of m balls and with density μ has a sphere separator with intersection number $O\left(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right)$ that $\frac{d+1}{d+2}$-splits. If we could compute such a sphere separator in $O(n)$ time, then the time required in computing such a search structure, $T(m)$, would be given by the following recurrence.

$$T(m) \leq \begin{cases} 1 & \text{if } m \leq m_0 \\ T(\delta_1 m + cm^\beta) + T((1-\delta_1)m) + O(m) & \text{if } m > m_0, \end{cases} \tag{9.5}$$

where $\delta_1 \leq \delta$.

Lemma 9.3 *Let T be the function defined above. Then $T(n) = O(n \log n)$ for a sufficient large constant m_0 that depends only on d, δ, β, and c.*

Proof: We use induction to establish $T(n) \leq Cn \log n$ for an appropriate choice of the constant $C > 1$.

Clearly, $T(m) = 1 \leq C$ for $m \leq m_0$. This is the base of the induction. Now, assume the lemma is true for all $m < n$, we have

$$\begin{aligned}
T(n) &\leq T(\delta_1 n + cn^\beta) + T((1-\delta_1)n) + c_2 n \\
&\leq C(\delta_1 n + cn^\beta)\log(\delta_1 n + cn^\beta) + C((1-\delta_1)n)\log((1-\delta_1)n) + c_2 n \qquad \text{(By induction)} \\
&\leq Cn\log(\delta n + cn^\beta) + Ccn^\beta \log(\delta + cn^\beta) + c_2 n \\
&\leq Cn\log\left(\frac{1+\delta}{2}n\right) + Ccn^\beta \log(\delta n + cn^\beta) + c_2 n \\
&= Cn\log n - Cn\log\left(\frac{2}{1+\delta}\right) + Ccn^\beta \log(\delta + cn^\beta) + c_2 n \\
&\leq Cn\log n,
\end{aligned}$$

as long as we choose C and m_0 such that

$$Cn\log\left(\frac{2}{1+\delta}\right) - Ccn^\beta \log(\delta + cn^\beta) - c_2 n \geq 0. \tag{9.6}$$

Because $\beta < 1$ and $\delta < 1$, inequality (9.6) holds for a sufficiently large m_0 and $C > c_2$ which only depend on d, δ, β and c. □

Consequently,

$$T(n, d) = O(n \log n).$$

However, we know of no deterministic linear time algorithm to compute a sphere separator with intersection number $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$ for a neighborhood system of density μ. Instead, as shown in the previous chapters, such a sphere separator can be computed in random linear time. From Chapter 5, if $\beta = \frac{d-1}{d} + \epsilon$ for some constant $\beta : 0 < \beta < \frac{1}{d}$, then the probability such a randomized algorithm outputs a sphere separator with intersection number $O(\mu^{\frac{1}{d}} n^\beta)$ is at least $1 - \frac{1}{n^\epsilon}$. Moreover, we can check whether the intersection number of a sphere separator is $O(\mu^{1/d} n^\beta)$ in linear time. If the construction above stops when $m_0 = O(\log n)$, then it follows from Lemma 9.3 that the search structure can be constructed in random $O(n \log n)$ time with a probability of success $1 - \frac{1}{\log n}$. To amplify the probability of success, if the subproblem is of size m, we run the randomized sphere separator algorithm $O(\frac{\log n}{\log m})$ times. It can be easily shown that with probability $\frac{1}{n}$ the above algorithm terminates in $O(n \log n \log\log n)$-time with a correct search structure. This extra $O(\log\log n)$ overhead can be eliminated using a more sophisticate analysis (Frieze, Miller, and Teng [28]).

9.1.2 A parallel construction

We now examine the parallel time required in computing a search structure of the neighborhood query problem. In this context, the time complexity has a new parameter p, the number of processors used in the algorithm. For the neighborhood query problem, let $T(n, d, p)$ represent the

preprocessing time required to build the search structure with p processors. We now show that the search structure presented in the above subsection has the following properties.

$$\begin{aligned} Q(n,d) &= O(\mu + \log n) \\ S(n,d) &= O(n) \\ T(n,d,n) &= \text{random } O(\log^2 n). \end{aligned}$$

The space complexity and the query time is the same as they used to be. The $O(\log^2 n)$ parallel time follows from the observation that with n processor, we can in random $O(\log n)$ time compute a sphere separator S with intersection number $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$ and decide which neighborhoods intersect with S, whose centers are in the interior or exterior. Since the height of the search tree is $O(\log n)$, the total time complexity is $O(\log^2 n)$ on a CRCW PRAM. However, if we use a SCAN CRCW PRAM, i.e., we assume prefix sum can be performed in constant time, [11], then we can check to see if the sphere S has a small intersection number in constant time and divide neighborhoods into two sets, those balls that intersect the interior of the sphere and those that intersect the exterior, in constant time. Therefore the algorithm runs in random $O(\log n)$ time with total work at most $O(n \log n)$.

Without assuming the scan primitive, we know how to improve the parallel time bound to random $O(\log n \log\log n)$ on a CRCW PRAMs. However, the solution is much more complex and I have decided not to present it here. Nevertheless, I believe the parallel time bound can be reduced to random $O(\log n)$.

In comparison, the search structure constructed by the multi-dimensional divide and conquer has the following properties.

$$\begin{aligned} T(n,d,p) &= O(\log^d n) \\ Q(n,d) &= O(\mu + \log^d n) \\ S(n,d) &= O(n \log^{d-1} n) \end{aligned}$$

9.2 Constructing intersection graphs

The problem of this section is to construct the intersection graph of a given neighborhood system. There is a simple solution for this problem: testing each pair of neighborhoods to decide whether they intersect. Since there are $O(n^2)$ pairs and the testing of each pair can be performed in constant time, the whole construction can be done in $O(n^2)$ time. If we require the algorithm to report all edges of the intersection graph, the above algorithm is optimal if we are working with general neighborhood systems. This is because that there could be as much as $\Omega(n^2)$ number of edges in some intersection graphs. However, if the neighborhood system is restricted to the one with density μ, then the intersection graph has at most $O(\mu n)$ edges. Is there a more efficient way to compute the intersection graph when μ is small and d is fixed?

We answer this question in the affirmative by presenting two algorithms. The first one, using the separator based divide and conquer, is a randomized algorithm with time complexity $O(\mu n \log n)$. The second one is deterministic and has the similar time complexity $O(kn \log n)$, although is only for k-neighborhood systems. The deterministic algorithm uses an idea of Vaidya [91]. However, the second algorithm is sequential in nature and we know of no efficient method to

implement it in parallel. In contrast the first algorithm can be efficiently parallelized in random $O(\log^2 n)$ time[1] using $O(\mu n)$ processors.

9.2.1 A separator based algorithm

To illustrate the idea, let us view the graph construction problem as a search problem: a problem of exploring the structure of an unknown graph with the help of some oracles. First, assume that n, the number of vertices, is known in advance and we have an oracle – the *edge oracle* – which answers the question of the form "is there is an edge between vertex u and v?" in constant time. It is not hard to see that even though the number of edges is known in advance, $\Omega(n^2)$ queries have to be asked in the worst case.

Now suppose that there is more information available: it is known in advance that the graph has a cn^β-separator that δ-splits and moreover each subgraph of $m > m_0$ vertices also has a cm^β-separator that δ-splits, for some constant c, m_0, and $0 < \beta < 1$. Can the number of queries be reduced? It is a nice question, but I have no answer.

However, suppose in addition, we have an oracle – a *separator oracle* – which when presented with a subset of m vertices, delivers three sets, A, B, and C, where C is an cm^β-separator that δ-splits the subgraph induced by those m vertices into A and B. Then it is sufficient to consult with the oracle $O(n)$ times to compute the structure of the unknown graph G

The strategy is divide and conquer. We first present the separator oracle the whole set of vertices and get back from the oracle three sets A, B, C, where C is a cn^β-separator that δ-splits G into A and B. We then recursively search the structure of subgraphs induced by $A \cup C$ and $B \cap C$ until the size of subproblems is below m_0. Finally, we use the edge oracle to finish off.

The total number of query $q(n)$ to the separator oracle is clearly given by the following recurrence.

$$q(n) \leq \begin{cases} 0 & \text{if } n \leq m_0 \\ q(\delta_1 n + cn^\beta) + q((1-\delta_1)n) + 1 & \text{if } n > m_0, \end{cases}$$

where $\delta_1 \leq \delta$.

By a similar argument as Lemma 9.2, it can be shown $q(n) = O(n)$.

Now suppose each query to the separator oracle costs $O(m)$ time, where m is the size of query. It is not hard to see that the total time $T(n)$ needed to search the structure of the graph is given by the following recurrence.

$$T(n) \leq \begin{cases} O(1) & \text{if } n \leq m_0 \\ T(\delta_1 n + cn^\beta) + T((1-\delta_1)n) + O(n) & \text{if } n > m_0, \end{cases}$$

By Lemma 9.3, we have $T(n) = O(n \log n)$.

The divide and conquer algorithm for constructing intersection graphs is based on the following interesting observation: we do not need to have the intersection graph in order to compute a small separator efficiently – all we need is the neighborhood system!

To see this, let us recall how we compute a small separator of an intersection graph. First we find a sphere separator S of low cost. This step involves computing an approximate center point

[1] Again, if we use Scan CRCW PRAM, the we can reduce the time to random $O(\log n)$. We also know a parallel algorithm with time complexity random $O(\log n \log\log n)$ on a CRCW PRAM.

and a conformal map. We then compute a vertex separator from the sphere separator S. The rule of choosing vertices is very simple: If ball B_i has a common point with S, then the vertex corresponding to B_i is placed in the separator. The time complexity of the above step is $O(n)$.

One important aspect of the above separator based divide and conquer algorithm is that it can be efficiently parallelized. In particular, the above algorithm can be implemented on an CRCW PRAM with μn processors in random $O(\log^2 n)$ time, as shown in the last section. In conclusion,

Theorem 9.4 *The intersection graph of a neighborhood system with density μ in a fixed dimensional space can be computed in random $O(\mu n \log n)$ time sequentially and in random $O(\log^2 n)$ time using μn processors on a CRCW PRAM. Moreover, the algorithm uses only $O(n)$-space.*

In contrast, we can only derive an $O(\mu n \log^{d-1} n)$ time, $O(n \log^{d-1} n)$-space algorithm for computing an intersection graph from the multi-dimensional divide and conquer paradigm.

9.2.2 A deterministic algorithm

In this section, we present an $O(kn \log n)$ time deterministic algorithm for computing the intersection graph of a given k-neighborhood system in a fixed dimensions. This algorithm is optimal up to a constant factor when k is a constant. The lower bound result of Ben-Or [7] implies that in the algebraic decision tree model of computation, there is a lower bound of $\Omega(n \log n)$ in the time complexity for this problem even when $d = 1$.

The basic idea of our algorithm is motivated by the method of Vaidya [91] for solving the k-nearest-neighbors problem. However our specific construction is quite different and the complexity argument is more sophisticated.

The algorithm constructs a sequence of graphs G_0, $G_1, \ldots, G_m$ where G_0 is a graph with a single node and G_m is the desired intersection graph. For each $1 \leq i \leq m$, G_i is a *refinement* of G_{i-1} and is obtained from G_{i-1} by a constant number of modifications.

More specifically, we start with the smallest d-dimensional cube b_0 that contains all centers P of the neighborhood system, and G_0 is the graph with vertex $\{b_0\}$. At each stage of refinement, from G_{i-1} to G_i, a cube b in G_{i-1} of the largest size is chosen and b is split into 2^d equal-size cubes $b_1, \ldots, b_{2^d}$ by d hyperplanes each passing through the center of b and being perpendicular to one of the coordinate axes. The splitting induces a partition of points in b into sets $b_i \cap P$. We discard those cubes whose interiors are empty and shrink each of the remaining boxes as much as possible to the extent that the resulting boxes contain the same set of points. A cube is *reduced* if it contains less than k points from P. Each reduced cube is further decomposed into cubes, and each contains exactly one of the points in the reduced cubes. Such cubes are called *singleton* cubes. The set of shrunk cubes and singleton cubes are called *children* of b, and b their *parent*. The set of cubes of G_i is then the set of all cubes in G_{i-1} but b plus the set of children of b.

Graph G_i has vertices, one for each cube in it. There is an edge between two nonreduced cubes c_i and c_j if $2 \cdot c_i$ and $2 \cdot c_j$ have a common point; between two singleton cubes containing points p_i and p_j, respectively if B_i and B_j have a common point; between a nonreduced cube c_i and a singleton cube containing p_j if $2 \cdot c_i$ and B_j have a common point. Such a graph is called a *pseudo 2-dilation graph* of cubes in G_i.

Note that the refinement from G_{i-1} to G_i is performed by locally changing the structure of G_{i-1}. To be more specific, let $N_{i-1}(b)$ denote the set of cubes connected with cube b in G_{i-1}.

Clearly, G_i can be obtained from G_{i-1} by inserting vertices one for each cube in children(b) and deleting the vertex corresponding to b and making proper connections between $N_{i-1}(b)$ and children(b).

We know that the number of children(b) is bounded by $k2^d$. The following lemma bounds the size of $N_{i-1}(b)$ for all nonreduced cubes b.

Lemma 9.5 *There is a constant c_d that depends only on d such that if b is nonreduced in G_{i-1}, then the number of cubes in $N_{i-1}(b)$ is bounded from above by c_d.*

Proof: First observe that at each stage, the set of nonreduced cubes forms a 1-neighborhood system in the L_∞ norm. For each b' in $N_{i-1}(b)$, we know the size of *parent*(b') is at least as large as that of b, because b' is split before b. There are two cases depending on whether *parent*(b') is an ancestor of b. Let Ξ_1 be the set of all *parent*(b')'s which are not ancestors of b, together with b. Clearly, all cubes in Ξ_1 are nonreduced and Ξ_1 is a 1-neighborhood system and b is the smallest cube in this neighborhood system. Therefore, $(2 \cdot b)$ can only intersect some constant number of such cubes (See Ball Intersection Lemma I and II). Let this constant be c'. Also it is not hard to see that $2 \cdot b$ can only intersect some constant number of the first type of parent cubes, say c''. To combine these, let $c''' = (c' + c'')2^d$, we know $(2 \cdot b)$ can not intersect more than c''' nonreduced cubes of G_{i-1}. The same argument when applied to cubes in Ξ_1 also implies that for each $b' \in \Xi_1$, $2 \cdot b'$ can not intersect more than c''' cubes from Ξ_1. Therefore, $\Xi_2 = \{2 \cdot b' : b' \in \Xi_1\}$ forms a c'''-neighborhood system, and $2 \cdot b$ is the smallest cube. By Ball Intersection Lemma I and II again, we know that there is a constant c'_d such that $2 \cdot b$ can not intersect more than c'_d number of cubes from Ξ_2. Letting $c_d = 2^d c'_d$, we conclude the degree of b in the pseudo 2-dilation graph of cubes in G_{i-1} is bounded from above by c_d. □

Consequently, G_i can be obtained from G_{i-1} by a constant number of modifications.

The remainder of the algorithm is very like the one of Vaidya [91] for computing nearest neighborhood graphs. We maintain a heap of cubes in G_i. This allows us to pick a cube that has the largest volume in constant time. Since the total size of the cube tree is bounded by $2n$, the total time for heap maintenance is thus $O(n \log n)$. To split the set of centers in a cube b efficiently, we, like Vaidya, use d ordered lists $List_i(b)$, $1 \leq i \leq k$, the i^{th} list containing the centers in $b \cap P$ ordered on the i^{th}-coordinate. When b is split, we can obtain similar lists for all cubes in *children*(b) from the ordered lists for b. We use exactly the same splitting procedure of Vaidya. By the Splitting Lemma 1 of Vaidya, we can bound the total splitting time by $O(kn \log n)$. For details, we refer the reader to Vaidya [91].

Therefore, we have shown

Theorem 9.6 *The intersection graph of a k-neighborhood system can be computed in $O(kn \log n)$ time deterministically.*

9.3 Constructing overlap graphs

In this section, we generalize the construction algorithm in the last section to overlap graphs. Here the problem is to compute the α-overlap graph of a neighborhood system for a positive real $\alpha > 1$.

Again, the divide and conquer algorithm uses an important fact of overlap graphs: we do not need to have the overlap graph in order to compute a small separator – we can directly compute one from the neighborhood system itself.

To see this, let us go back to the construction in Chapter 6 when proving the overlap graph has a small separator. In that construction, we first find a sphere separator S with low surface area. Then we compute a vertex separator from S. Here is the rule for choosing vertices in a separator: if ball B_i intersects S, then the vertex corresponding to B_i is chosen. Otherwise, if $\alpha \cdot B_i$ intersects with S and the radius of B_i is no more than the radius of S, the vertex of B_i is also chosen. Hence this step can be computed in $O(n)$ time sequentially and in constant time if we have n processors.

Therefore, we have a random linear time oracle for computing a small separator for an overlap graph. Incorporating this into the generic algorithm presented in the last section, we have

Theorem 9.7 *The α-overlap graph of a neighborhood system with density μ in a fixed d-dimensional space can be computed in random $O(\alpha^d \mu n \log n)$ time sequentially and in random $O(\log^2 n)$ time using $\alpha^d \mu n$ processor on a CRCW PRAM. Moreover, our algorithm uses $O(n)$ space.*

In contrast, we can only derive an $O(\alpha^d \mu n \log^{d-1} n)$ time, $O(n \log^{d-1} n)$ space algorithm for computing an overlap graph from the multi-dimensional divide and conquer paradigm.

Similarly, we can generalize the deterministic algorithm of the last section to overlap graphs. In particular, we can prove,

Theorem 9.8 *The α-overlap graph of a k-neighborhood system can be computed in deterministic $O(\alpha^d k n \log n)$ time.*

9.4 Constructing k-nearest neighborhood graphs

Now we come to a more interesting but harder problem – to compute the k-nearest neighborhood graph when given a set of n points $P = \{p_1, \ldots, p_n\}$ in a fixed dimension. The reason that we say this problem is harder is the following: the neighborhood of each point is not known in advance and thus it is harder to compute a small separator of the graph by only looking at the points.

The nearest neighbor problem has been intensively studied [8, 91, 75, 16]. The simplest algorithm for this problem, the one that compares each point with all other points, runs in $\Theta(n^2)$-time. Using the multi-dimensional divide and conquer, Bentley was the first to develop an $O(n(\log n)^{d-1})$-time algorithm in d-dimensions (Bentley [8]). Clarkson presented a randomized algorithm for the problem with expected run time $O(n \log n)$ (Clarkson [16]) and Vaidya gave an optimal $O(kn \log n)$ time algorithm for the k-nearest neighbor problem in any fixed dimensions (Vaidya [91]). However, there is no optimal parallel algorithm known for this problem in arbitrary fixed dimensions. Vaidya's algorithm is very sequential. A straightforward parallelization of Bentley's algorithm runs in $O(\log^d n)$ time using n processors, although the work of Cole and Goodreich pushes the time complexity down to $O(\log^{d-1} n)$. When randomization is allowed, Reif and San presented an $O(\log n)$ random time n processor algorithm for this problem in 2 dimensions by reducing it to the Voronoi problem. We know of no other parallel results in an arbitrary fixed dimensions.

We now present a random $O(\log^3 n)$ time, n processors parallel algorithm for constructing a k-nearest neighborhood graph in a fixed dimensions. To our knowledge, this is the first logarithmic time linear processors parallel algorithm whose exponent in the $\log n$ term (of the time bound)

is independent of d. This new algorithm uses the separator based divide and conquer paradigm. Recently, Frieze, Miller, and Teng [28] obtained an improved divide and conquer algorithm using a more sophisticate construction and a new probabilistic lemma.

We first present an random $O(n \log^2 n)$ time sequential algorithm. The algorithm works as follows. The first step chooses a hyperplane h perpendicular to the x_1-coordinate axis, dividing the points P into two subsets P_l and P_r, each containing $n/2$ points. The second step recursively computes the k-nearest neighbor graph of A and B, respectively. Let B_i be the ball associated with point p_i from the recursive construction. Assume $p_i \in P_l$. There are two cases:

- **Case 1:** B_i does not intersect h. In this case, B_i is the correct ball in the sense that B_i is the largest ball center at p_i with less than k points from P in its interior. Hence, the k-nearest neighbors of p_i is correctly computed.
- **Case 2:** B_i intersects h. In this case, B_i could be bigger than what it is supposed to be because there might be points from P_r in the interior of B_i. So we have to correct the radius of B_i.

The key observation is that the correction step for case 2 is just a neighborhood query problem which can be solved in random $O(n \log n)$ time. It is interesting to point out that the correction step is where Bentley used the multi-dimensional reduction. So, he can not simply stop the recurrence. Instead, he reduces the problem to one lower dimensions. The separator based divide and conquer removes the necessity of a multi-dimensional recursive call and hence improves the algorithm.

Let Ξ_l be the set of balls from P_l that intersect h. Clearly Ξ_l is a k-neighborhood system. So we can build a search structure for the neighborhood query problem of Ξ_l. Using this search structure, we can decide in $O(\log n)$ time for each $p \in P_r$ the set of balls in Ξ_l whose interior contains p and we correct B_i's in Ξ_l accordingly. Thus in random $O(n \log n)$ time, we can correct all balls in Ξ_l. Similarly, we can correct all balls in Ξ_r, the set of balls from P_r that intersect h.

The time complexity of the above algorithm is given by the following recurrence.

$$T(n) = T(\delta_1 n + cn^{\beta}) + T((1 - \delta_1)n) + O(n \log n).$$

Therefore, $T(n) =$ random $O(n \log^2 n)$.

The above algorithm can be implemented in parallel. The first step can be done in constant time using n processors if we presort the points according to their x_1-coordinates. The second step can be performed in random $O(\log^2 n)$ time using n processors. Hence the total time complexity is bounded from above by $O(\log^3 n)$. Using a more sophisticated construction, we can reduce the time complexity to random $O(\log n \log\log n)$ without increasing the number of processors (Frieze, Miller, Teng [28]).

9.5 Testing k-neighborhoodness

Finally, we give an efficient algorithm for answering the following question: whether a given neighborhood system is a k-neighborhood system. We refer this problem as *testing k-neighborhoodness.*

There is a straightforward $O(n^2)$-time algorithm for testing k-neighborhoodness. This algorithm tests for each pair of centers to see if one is in the interior of the neighborhood of another. However, a much more efficient algorithm exists which is based on the following observation.

Lemma 9.9 *Let $\Xi = \{B_1, \ldots, B_n\}$ be a neighborhood system with centers $P = \{p_1, \ldots, p_n\}$, and for all $1 \leq i \leq n$, let q_i be the point among p_i's $(k+1)$-nearest neighbors that is furthest away from p_i. Then Ξ is a k-neighborhood system iff for all i, q_i is not in the interior of B_i.*

Proof: If q_i is in the interior of B_i, then the interior of B_i contains at least $(k+1)$ points from P, and hence Ξ is not a k-neighborhood system. On the other hand, if q_i is not in the interior of B_i, then we know that no more than k points from P can be there, and the lemma follows. □

Immediately from the above lemma, we get the following algorithm for testing k-neighborhoodness.

Input: A neighborhood system $\Xi = \{B_1, \ldots, B_n\}$.

1. Let $P = \{p_1, \ldots, p_n\}$ be the set of centers of Ξ;
2. For all i, compute the set N_i of $k+1$-nearest neighbors of p_i;
3. for all i, compute for each p_i, the point q_i in N_i that is farthest away from p_i;
4. if there is an i, such that q_i is in the interior of B_i, then output Ξ *is not a k-neighborhood system* otherwise output Ξ *is a k-neighborhood system.*

Consequently,

Theorem 9.10 *The problem of testing k-neighborhoodness can be solved in $O(kn \log n)$ time sequentially and in random $O(\log^2 n)$ time using kn processors on a CRCW PRAM.*

Chapter 10

Final Remarks

Now that we have proved that each neighborhood system with a bounded density in a fixed dimensions has a sphere separator with a small cost (in term of the intersection or overlap number). Furthermore, we have presented a randomized algorithm that computes such a sphere separator in random linear time with high probability if given a random number generator which uniformly generates a real number from $[0, 1]$ (Knuth [54, 41]).

However, if the number of random bits is bounded, then it may be the case that all points in the sample space are "bad" in the sense that they do not correspond to any sphere separator with a small cost. How to use only a finite number of random bits to generate a small cost sphere separator is thus of a major concern.

In this chapter, a new method to derive the separator property of intersection graphs and overlap graphs is presented. This approach not only provides a mean to bound the number of random bits required in computing a small cost sphere separator but also generates a more direct and elementary proof of the main separator theorem (Theorem 6.1 and 6.10).

To simplify the discussion, I will present proofs only for bounding the intersection number. The argument can be easily generalized to the overlap number. It is worthwhile to mention that it is still not known how to bound the number of random bits for the general case as stated in Theorem 5.9. The result to be presented relies on the structural properties of neighborhood systems.

10.1 Great rings of the unit sphere

As shown in Chapter 5, a d-dimensional neighborhood system $\Xi = \{B_1, \ldots, B_n\}$ can be conformally mapped to a unit d-sphere U_d so that each d-dimensional hyperplane containing the center of U_d $\frac{d+1}{d+2}$-splits Ξ. Notice that each B_i is mapped to a patch D_i on U_d, whose boundary C_i has the shape of a (d-1)-sphere. The radius of D_i is defined to be the radius of C_i. Clearly, the density of $\{D_i : 1 \leq i \leq n\}$ is μ iff the density of Ξ is μ. Let A_d be the surface area of U_d. Since the number of patches with surface area greater than $A_d/2$ is bounded by $O(\mu)$, without loss of generality, we assume that each patch has a surface area at most $A_d/2$. Let r_i be the radius of C_i. The volume of the d-dimensional ball with boundary C_i is then equal to $V_d(r_i)^d$, which is a lowerbound on the surface area of the patch D_i. In the remaining of this chapter, I will identify B_i with D_i, and assume that $\Xi = \{B_1, \ldots, B_n\}$ is given on the unit d-sphere.

10.1.1 Great rings

Recall that a great sphere of U_d is the intersection of U_d with a hyperplane passing through the center of U_d. Each great sphere GS can be identified with the pair of points p_{GS} and q_{GS} on U that lay on the normal to GS (see Figure 10.1 for an example in 2 dimensions).

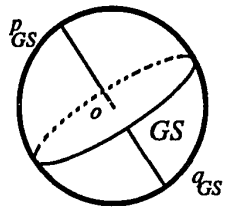

Figure 10.1: GC and its p_{GS}, q_{GS}

There is a dual relation between points on U_d and their great spheres:

Proposition 10.1 (Duality) *For each pair of great spheres GS and GS' of U_d, GS contains $p_{GS'}$ (and hence $q_{GS'}$ as well) if and only if GS' contains p_{GS} (and hence q_{GS}).*

Proof: If $p_{GS'} \in GS$, then the hyperplane associated with GS is perpendicular to the hyperplane associated with GS'. Therefore, p_{GS} is contained in the hyperplane of GS', and hence contained in GS'. □

Let a *great ring* be the set of points of U_d that lay between a pair of parallel hyperplanes symmetric to the center of U_d. The *width* of a great ring is then the distance of its two hyperplanes. Notice that a great sphere is a great ring with width 0. Clearly, the surface area of a great ring of width r is bounded by $V_{d-1}r$, where V_{d-1} is the volume of a unit $(d-1)$-ball.

It simply follows from the duality of points and their great spheres (Proposition 10.1) that

Lemma 10.2 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system on U_d. Then for each $1 \leq i \leq n$, there is a great ring R_i such that a great sphere GS intersects B_i iff p_{GS} and q_{GS} is contained in R_i. Moreover, the width of R_i is equal to $2r_i$, where r_i is the radius of B_i (see Figure 10.2).*

For each point x on U_d, let $\iota_\Xi(x)$ be the number of neighborhoods in Ξ that intersect the great sphere $GS(x)$ associated with x. Let

$$\psi(\Xi) = \frac{1}{A_d}\left(\int_{x \in U_d} \iota_\Xi(x)(dx)^d\right).$$

By the duality of points and their great spheres (Proposition 10.1), we have

Proposition 10.3 *The above defined $\psi(\Xi)$ is equal to the expect intersection number of a random great sphere of U_d.*

Lemma 10.4 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system on U_d. Let R_i be the great ring defined by B_i. Then*

$$\psi(\Xi) = \frac{1}{A_d}\left(\sum_{i=1}^{n} \text{Area}(R_i)\right)$$

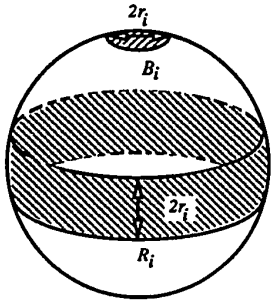

Figure 10.2: The great ring induced by B_i

Proof: For each $1 \leq i \leq n$, for each $x \in U_d$, Let $\Phi_i(x)$ be the function which takes value 1 if $GS(x)$ intersects B_i and 0 otherwise.

$$\begin{aligned}
\psi(\Xi) &= \frac{1}{A_d}\left(\int_{x\in U_d} \iota_\Xi(x)(dx)^d\right) \\
&= \frac{1}{A_d}\left(\int_{x\in U_d} [|\{B_i : GS(x)\cap B_i \neq \emptyset\}|](dx)^d\right) \\
&= \frac{1}{A_d}\left(\sum_{i=1}^{n}\left[\int_{x\in U_d} \Phi_i(x)(dx)^d\right]\right) \\
&= \frac{1}{A_d}\left(\sum_{i=1}^{n}[\text{Area}(R_i)]\right)
\end{aligned}$$

The first equality follows from the definition of $\iota_\Xi(x)$; the second equality is obvious; and the third equality follows from Lemma 10.2. □

10.1.2 The total area of great rings

The above relation (Lemma 10.4) enable us to obtain an upper bound on $\sum_{i=1}^{n}\text{Area}(R_i)$. This upper bound will be used in the next section to bound the number of random bits required in computing a small separator for the intersection graph.

Lemma 10.5 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system on U_d with density μ. Let R_i be the great ring defined by B_i. Then*

$$\sum_{i=1}^{n}\text{Area}(R_i) = O\left(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right).$$

Proof: Let π be the stereographic projection which maps $\mathbb{R}^d$ onto U_d. Let Ξ' be the pre-image of Ξ in $\mathbb{R}^d$, i.e., $\Xi' = \{B_1', \ldots, B_n'\}$ such that $\pi(B_i') = B_i$. Clearly, Ξ' is a neighborhood system in $\mathbb{R}^d$, whose density is also μ.

Let f'_i be the local density function defined on B'_i as in Chapter 6 and let $f' = L_{d-1}(f'_1, \ldots, f'_n)$. It follows from Lemma 6.8 that

$$\text{Total-Cost}(f') = O\left(\mu^{\frac{1}{d-1}} n\right).$$

Let f be the density function on U_d obtained from f' via the stereographic projection π. We have

$$\text{Total-Cost}(f) = \text{Total-Cost}(f') = O\left(\mu^{\frac{1}{d-1}} n\right).$$

It follows from Proposition 5.1, that $\text{avg}(f) = O\left(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right)$, recall that $\text{avg}(f)$ denotes the average area of great spheres of U_d. We now relate $\text{avg}(f)$ with $\psi(\Xi)$.

For each great sphere GS of U_d, let S be the pre-image of GS in $\mathbb{R}^d$, i.e., S is a $(d-1)$-sphere $\mathbb{R}^d$ such that $\pi(S) = GS$. By the definition of f and the fact that π is conformal, $\text{Area}_{f'}(S) = \text{Area}_f(GS)$. Furthermore, the intersection number of GS over Ξ is equal to the intersection number of S' over Ξ'. Let $c_1 = 3^d$ and

$$c_2 = \left(\frac{4\sqrt{7}}{7}\right)^{d-1} \frac{1}{V_{d-1}}$$

It follows from Lemma 6.5 that $\iota_\Xi(p_{GS}) \leq c_1\mu + c_2\text{Area}_f(GS)$.

In other words, for each point $x \in U_d$, $\iota_\Xi(x) \leq c_1\mu + c_2\text{Area}_f(GS(x))$. Therefore,

$$\begin{aligned}
\psi(\Xi) &= \frac{1}{A_d}\left(\int_{x\in U_d} \iota_\Xi(x)(dx)^d\right) \\
&\leq \frac{1}{A_d}\left(\int_{x\in U_d} [c_1\mu + c_2\text{Area}_f(GS(x))](dx)^d\right) \\
&= \frac{1}{A_d}(c_1 A_d \mu + c_2 A_d \text{avg}(f)) \\
&= c_1\mu + c_2\text{avg}(f) \\
&= O\left(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right)
\end{aligned}$$

The lemma then follows from Lemma 10.4. □

Remark 10.6 *In fact, the analysis of Chapter 6 can be adapted to prove the above Lemma directly. The direct proof leads to much smaller constant term in the Big-O notation. However, the proof is quite lengthy (as long as Chapter 6), I have decided not to include it in the thesis. In section 10.3, I will present a simpler and more elementary proof of a weaker version Lemma 10.5 as well as Theorem 6.1.*

10.2 The number of random bits and a deterministic algorithm

From the last section, we have $\sum_{i=1}^n \text{Area}(R_i) \leq O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$. Hence, there is a point on U_d which is contained in at most $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$ great rings from $\{R_1, \ldots, R_n\}$. For each given set of points $Q = \{q_1, \ldots, q_m\}$ on U_d, let

$$\phi(Q) = \frac{1}{m}\sum_{i=1}^{m} |\{R_j : q_i \in R_j\}|$$

Now the question becomes: can we find a finite (hopefully small) set Q of points on U so that $\phi(Q) = O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$?

An affirmative answer to this question has two consequences:

1. If the cardinality of this set is m, then we only need $O(\log m)$ random bits to do the sampling in computing a sphere separator with a low cost;
2. If m is polynomial in n, then we will have a deterministic polynomial time algorithm for computing a small (sphere or vertex) separator[1].

Lemma 10.7 *For each positive real $\epsilon < 1$, let $Q = \{q_1, \ldots, q_m\}$ be a set of points on the unit d-sphere U such that each great ring R of width $r \geq \epsilon$ satisfies $|R \cap Q| \leq c_d \text{Area}(R)m$, where c_d is a constant that depends only on d. Then there is a constant $c > c_d$ depending only on d such that for each set of great rings $\{R_1, \ldots, R_n\}$,*

$$\phi(Q) \leq cn\epsilon + c\sum_{i=1}^{n} \text{Area}(R_i)$$

Proof: Partition $\{R_1, \ldots, R_n\}$ into two sets, Δ_1 and Δ_2, where

$$\begin{aligned} \Delta_1 &= \{R_i : r_i < \epsilon\} \\ \Delta_2 &= \{R_1, \ldots, R_n\} - \Delta_1 \end{aligned}$$

Notice for each $R_i \in \Delta_1$, R_i is contained in a great ring with width ϵ, and there is a constant c' such that the area of a great ring of width ϵ is $c'\epsilon$. Hence there is a constant $c > c_d$ depending only on d such that $|R_i \cap Q| \leq c\epsilon m$. Since $|\Delta_1| \leq n$, we have

$$\sum_{R_i \in \Delta_1} |R_i \cap Q| \leq cn\epsilon m.$$

It follows from the condition of the lemma, we have

$$\sum_{R_i \in \Delta_2} |R_i \cap Q| \leq \sum_{R_i \in \Delta_2} c_d \text{Area}(R_i)m \leq \sum_{i=1}^{n} c\text{Area}(R_i)m.$$

Therefore,

$$\phi(Q) = \frac{1}{m}\left(\sum_{i=1}^{n} |R_i \cap Q|\right) \leq cn\epsilon + c\sum_{i=1}^{n} \text{Area}(R_i)$$

□

By Lemma 10.5, we have

$$\sum_{i=1}^{n} \text{Area}(R_i) = O\left(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right)$$

So if $\epsilon = \left(\frac{\mu}{n}\right)^{\frac{1}{d}}$, then the assumption of Lemma 10.7 implies that

$$\phi(Q) = O\left(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}}\right).$$

The following lemma must be known before. But I did not find any reference.

[1]Recall that an approximated center point can be computed in deterministic linear time (Chapter 8).

Lemma 10.8 *There is a constant c which depends only on d such that for each positive real $\epsilon \leq 1$, there is a set Q of $m = (2d+2)\left(\lceil \frac{1}{\epsilon} \rceil\right)^d$ points on the unit d-sphere U_d (whose center is o, the origin) such that each great ring R of width $r \geq \epsilon$ covers at most $c\text{Area}(R)m$ points of Q.*

Proof: Without loss of generality, assume $1/\epsilon$ is an odd integer. Let C be the unit $(d+1)$-cube centered at o and let C' be $\frac{1}{\sqrt{d}} \cdot C$, i.e., C' can be obtained from C by shrinking by a factor of $\frac{1}{\sqrt{d}}$.

Let Q_C be the set of $(d+1)$-dimensional grid points of resolution ϵ on the surface of the unit cube C. Algebraically,

$$\begin{aligned} Q_C \;=\; & \{(x_1, \ldots, x_{d+1}) : \\ & \exists i \text{ such that } x_i = \pm 1 \text{ and for all } j \neq i \; x_j = \pm l_j \epsilon \text{ for some integer } 0 \leq l_j < \tfrac{1}{2\epsilon}\} \end{aligned}$$

Clearly, the cardinality of Q_C is $m = (2d+2)\left(\frac{1}{\epsilon}\right)^d$.

For each point $q_C \in Q_C$, the line segment between o and q_C has a unique common point with U_d as well as with C', denoted by q and q'_C respectively. Let $Q = \{q : q_C \in Q_C\}$ and $Q'_C = \{q'_C : q_C \in Q_c\}$. Notice that Q'_C is the set of $(d+1)$-dimensional grid points of resolution $\frac{\epsilon}{\sqrt{d}}$ on the surface of C'.

We now show that the set Q satisfies the lemma. For each great ring R of width r, let h_1 and h_2 be the pair of parallel hyperplanes that define R. Let A be the set of points of Q'_C which lay between h_1 and h_2, let B be the set of points of Q which lay between h_1 and h_2, i.e., covered by R. By convexity, $B \subset \{q : q'_C \in A\}$. Therefore, the number of points of Q covered by R is at most the number of points from Q'_C that lay between h_1 and h_2. To bound the later number, we dilate the whole $(d+1)$-dimensional space by a factor of $\sqrt{d}$, i.e., C' is mapped to the unit cube centered at o. Notice that Q'_C is mapped to the set of $(d+1)$-dimensional grid points of resolution ϵ on the surface of the unit cube and h_1 and h_2 to a pair of parallel hyperplanes of distance $\sqrt{d}r$. Lemma 10.8 follows from Lemma 10.9 below. □

Lemma 10.9 *Let Q be the set of grid points of resolution ϵ on the surface of the unit $(d+1)$-dimensional cube centered at o. Let h_1 and h_2 be a pair of parallel hyperplanes that are symmetric to o and have distance r, with $r \geq \sqrt{d}\epsilon$. Then the number of points of Q which lies between h_1 and h_2 is bounded from above by $(2d+2)2^d(\sqrt{d}r + 2)\left(\frac{1}{\epsilon}\right)^d$.*

Proof: Each cube in $(d+1)$ dimensions has $2(d+1)$ d-dimensional faces. So, it is sufficient to bound the maximum possible number of grid points on each face that lay between h_1 and h_2.

Let F be a face of the unit cube which intersects h_1 or h_2. Let h'_1 and h'_2 be the pair of parallel d-dimension hyperplanes induced by h_1 and h_2, respectively, on F (or on the hyperplane defined by F). Since h_1 and h_2 are symmetric to the center of the cube, it follows that the distance between h'_1 and h'_2 is bounded by $\sqrt{d}r$ (because of the bounded angle). Notice that no d-dimensional hyperplane on F can intersect more than $\left(\frac{1}{\epsilon}\right)^d$ grids of F, and there are no more than $\sqrt{d}r\left(\frac{1}{\epsilon}\right)^d$ grids of F that lay between h'_1 and h'_2 (by a volume argument). Therefore, the number of grid points of F that are between h_1 and h_2 is at most $2^d(2+\sqrt{d}r)\left(\frac{1}{\epsilon}\right)^d$, because each grid contains 2^d grid points. A more careful analysis can be used to reduce this number. □

Corollary 10.10 *In the randomized algorithm for computing a small sphere separator, it is sufficient to use $O(\log n)$ random bits.*

It is worthwhile to mention that in the randomized algorithm for computing a small sphere separator, it is more suitable to choose grid points on C, the unit cube, as sample points rather than those from the unit sphere. This is because the grid points on C have rational coordinates.

Because an approximate center point in a fixed dimension can be computed in deterministic linear time (see chapter 8). It follows from the above lemmas (with $\epsilon = O\left((\frac{\mu}{n})^{\frac{1}{d}}\right)$),

Corollary 10.11 *Given a d-dimensional neighborhood system Ξ with density μ, a sphere S can be computed in deterministic $O(n^2)$ time, such that S has an intersection number $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$ and $\frac{d+1}{d+2}$-splits Ξ.*

10.3 An elementary proof of the separator results

In this section, we give a very simple proof of the following lemma.

Lemma 10.12 *Suppose $\Xi = \{B_1, \ldots, B_n\}$ is a neighborhood system on U_d with density μ. Then*

$$\psi(\Xi) = O\left(\mu^{\frac{1}{d+1}} n^{\frac{d}{d+1}}\right).$$

Although the bound in the above lemma is weaker than what have been proved in Theorem 6.1 and Lemma 10.5, I hope this proof will shed new light on the separator property of neighborhood system and lead to more elementary proof to Theorem 6.1 and Lemma 10.5.

Proof: It follows from Lemma 10.4 that

$$\psi(\Xi) = \frac{1}{A_d}\left(\sum_{i=1}^{n} \text{Area}(R_i)\right) \leq \frac{1}{A_d}\left(\sum_{i=1}^{n} 2V_{d-1} r_i\right),$$

where r_i is the radius of B_i. The second inequality follows from Lemma 10.2,

Let $\gamma = (\frac{\mu}{n})^{\frac{1}{d+1}}$. To establish the lemma, we partition Ξ into disjoint subsets Ξ_1 and Ξ_2, where $\Xi_1 = \{B_i : r_i \geq \gamma\}$ and $\Xi_2 = \{B_i : r_i < \gamma\}$, Clearly, $\psi(\Xi) = \psi(\Xi_1) + \psi(\Xi_2)$.

For each $B_i \in \Xi_1$, $\text{Area}(B_i) \geq V_d r_i^d$. Because the density of Ξ is μ,

$$\sum_{i=1}^{n} \text{Area}(B_i) \leq \mu A_d.$$

Therefore, $|\Xi_1| \leq O\left(\mu^{\frac{1}{d+1}} n^{\frac{d}{d+1}}\right)$, and hence $\psi(\Xi_1) = O\left(\mu^{\frac{1}{d+1}} n^{\frac{d}{d+1}}\right)$. By Lemma 10.4,

$$\psi(\Xi_2) = \frac{1}{A_d}\left(\sum_{B_i \in \Xi_1} \text{Area}(R_i)\right) \leq \frac{1}{A_d} 2V_{d-1} n\gamma = O\left(\mu^{\frac{1}{d+1}} n^{\frac{d}{d+1}}\right).$$

Therefore, $\psi(\Xi) = O\left(\mu^{\frac{1}{d+1}} n^{\frac{d}{d+1}}\right)$. □

10.4 Open questions

There are still some interesting questions unanswered. They may represent the future direction to extend the work of this thesis. The following is a set of open questions.

1. Can a disk packing representation of a planar graph be computed in polynomial time?

   The proof of Thurston [83] hinges on the convergence of an iterative procedure. But it is not clear that his iterative procedure converges in polynomial time. Notice that an efficient algorithm for the above problem together with the results of this thesis would provide a new method for computing a small separator of a planar graph.

   Concerning about disk packing representation of planar graphs, Brightwell and Scheinerman [15] recently generalized Andreev-Thurston's theorem and deduced from their new results the following theorem: each 3-connected planar graph G and its dual can be represented simultaneously in the plane with straight line edges so that edges of G cross the dual edges at right angles, answering an question first asked by Tutte [87].

2. We have showed that each neighborhood system with density μ has a sphere separator of intersection number $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$. Does it also have a cube or other symmetrical figure separator of intersection number $O(\mu^{\frac{1}{d}} n^{\frac{d-1}{d}})$?

3. Is there a geometrical characterization of graphs with bounded genus and bounded excluded minor?

   In general, we would like to know whether all graphs that have a family of $O(\sqrt{n})$ separators are k-embeddable in 2-space for some constant k.

4. Do other geometrical graphs, such as relative neighborhood graphs or the 1-skeleton of Voronoi diagrams, have a small separator?

5. Can the separator results of this thesis be applied to other problems such as the N-body problem?

6. What is the complexity for deciding whether a graph is (α, k)-embeddable? Can we compute a center point in fixed dimension in $O(n\text{polylog}(n))$ time? Can we compute a sphere separator with a small intersection number or overlap number in deterministic $O(n\text{polylog}(n))$ time, when given a neighborhood system with density μ?

Bibliography

[1] Aggarwal, A., R. J. Anderson, and M.-Y. Kao. "Parallel depth-first search in general directed graphs". In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, ACM*, 297–306, 1989.

[2] Aho, A., J. Hopcroft and J. Ullman. *The Design and Analysis of Computer Algorithms.* Addison-Wesley, 1974.

[3] Andreev, E. M. "On convex polyhedra in Lobacevskii space". *Math. USR Sbornik*, 10(3):413–440, 1970.

[4] Andreev, E. M. "On convex polyhedra of finite volume in Lobacevskii space". *Math. USR Sbornik*, 12(2):270–259, 1970.

[5] Alon, N., P. Seymour, and R. Thomas. "A separator theorem for graphs with an excluded minor and its applications". In *ACM STOC*, 293–299. 1990.

[6] Batcher, K. "Sorting networks and their applications". In *Proceedings AFIPS, SJCC* 32: 307–314, 1968.

[7] Ben-or, M. "Lower bounds for algebraic computation trees". In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, ACM*, 80–86, 1983.

[8] Bentley, J. L. "Multidimensional divide-and-conquer". CACM 23: 214–229, 1980.

[9] Bern, M., D. Eppstein and J. R. Gilbert. "Provable good mesh generation". In *31st Annual Symposium on Foundations of Computer Science, IEEE*, 231–241, 1990.

[10] Birch, B. J. "On $3N$ points in a plane". *Proc. Bambridge Philos. Soc.*, 55(4): 289–293, 1959.

[11] Blelloch, G. E. *Vector Models for Data-Parallel Computing.* MIT-Press, Cambridge MA, 1990.

[12] Blum, M., R. W. Floyd, V. R. Pratt, R. L. Rivest, and R. E. Tarjan. "Time bounds for selection". *JCSS* 7(4): 448-461, 1972.

[13] Blum, M. and S. Kannan. "Program correctness checking... and the design of programs that check their work". In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, ACM*, 86–97, 1989.

[14] Blumenthal, L. and K. Menger. *Studies in Geometry.* W. H. Freeman and Company, San Francisco, 1970.

[15] Brightwell, G. R. and E. R. Scheinerman. *Representations of Planar Graphs.* manuscript, 1991.

[16] Clarkson, K. "Fast algorithm for the All-nearest-neighbors problem". In *the 24th Annual Symposium on Foundations of Computer Science*, 226–232, 1983.

[17] Cole, R., M. Sharir and C. K. Yap. "On k-hulls and related problems". *SIAM J. Computing*, 61, 1987.

[18] Cormen, T. H., C. E. Leiserson, R. L. Rivest. *Introduction to Algorithms.* MIT-Press. Cambridge MA, 1990.

[19] Conway J. H., and N. J. A. Sloane. *Sphere Packings, Lattices and Groups.* Springer-Verlag, 1988.

[20] Danzer, L., J. Fonlupt and V. Klee. "Helly's theorem and its relatives". *Proceedings of Symposia in Pure Mathematics*, American Mathematical Society, 7: 101-180, 1963.

[21] Djidjev, H.N. "On the problem of partitioning planar graphs". *SIAM J. ALG. DISC. METH.*, 3(2):229–240, June 1982.

[22] Doyle, P. G., and J. L. Snell. *Random Walks and Electric Networks.* The Carus Mathematical Monographs, The Mathematical Association of America. 1984.

[23] Dubrovin, B. A., A. T. Fomenko, and S. P. Novikov. *Modern Geometry – Methods and Applications.* Springer–Verlag, 1984.

[24] Erdós, P., R. L. Graham, and E. Szemerédi. "On sparse graphs with dense long paths". *Comp. and Math. with Appl.* 1: 365–369, 1975.

[25] Fáry, I. "On straight line representing of planar graphs". *Acta. Sci. Math.* 24: 229–233, 1948.

[26] De Fraysseix, H., J. Pach, and R. Pollack. "Small sets supporting Fáry embeddings of planar graphs". In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, 426–433, 1988.

[27] Fredickson, G. N., and R. Janardan. "Separator-based strategies for efficient message routing". In *27st Annual Symposium on Foundation of Computation Science,IEEE*, 428–237, 1986.

[28] Frieze, A. M., G. L. Miller, and S.-H. Teng. *Separator Based Divide and Conquer in Computational Geometry.* manuscript, Carnegie Mellon University, 1991.

[29] Garey, M. R. and D. S. Johnson. *Computer and Intractability: a guide to the theory of NP-completeness.* Freeman, San Francisco, 1979.

[30] Garey, M. R., D. S. Johnson, and L. Stockmeyer. "Some simplified NP-complete graph problems." *Theor. Comput. Sci.* 1: 237–267, 1976.

[31] Gazit, H. *An Improved Algorithm for Separating a Planar Graph.* manuscript, USC, 1986.

[32] Gazit, H. "A Deterministic Parallel Algorithm for Planar Graphs Isomorphism". In *32nd Annual Symposium on Foundations of Computer Science, IEEE*, to appear, 1991.

[33] Gazit, H. and G. L. Miller. "A parallel algorithm for finding a separator in planar graphs". In *IEEE FOCS*, 238–248, Los Angeles, October 1987.

[34] Gazit, H., G. L. Miller and S.-H. Teng. "Optimal tree contraction in the EREW model". In *Current Computations* S. K. Tewsburg, B. W. Dickinson, and S, C, Schwartz editors. 139–156, 1988.

[35] George, J. A., "Nested dissection of a regular finite element mesh". *SIAM J. Numerical Analysis*, 10: 345–363, 1973.

[36] George, A and J. W. H. Liu. *Computer Solution of Large Sparse Positive Definite Systems.* Prentice-Hall, 1981.

[37] Gibbons A. and W. Rytter. *Efficient Parallel Algorithms.* Cambridge University Press, 1988.

[38] Gilbert, J. R., J. P. Hutchinson, and R. E. Tarjan. "A separation theorem for graphs of bounded genus". J. Algorithms, 5, 391–407, 1984.

[39] Gilbert, J. R. and R. E. Tarjan. "The analysis of a nested dissection algorithm". *Numerische Mathematik*, 50(4):377–404, 1987.

[40] Grötschel, M., L. Lovász and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization.* Springger-Verlag, 1988.

[41] Hammersley, J. M. and D. C. Handscomb. *Monte Carlo Methods* Chapman and Hall, 1964.

[42] Hammond, S. and R. Schreiber. Solving unstructured grid problems on massively parallel computers. Technical Report TR 90.22, Research Institute for Advanced Computer Science, 1990.

[43] Hardy, G., J. E. Littlewood and G. Pólya. *Inequalities* Second edition, Cambridge University Press, 1952.

[44] Haussler, D. and E. Welzl. "ϵ-net and simplex range queries". *Discrete & Computational Geometry*, 2: 127–151, 1987.

[45] He, X., and Y. Yesha. "A nearly optimal parallel algorithm for constructing depth first spanning trees in planar graphs". *SIAM Journal on Comput.* 17(3):486–491, June, 1988.

[46] Hutchinson, J. P. and G. L. Miller. "On deleting vertices to make a graph of positive genus planar". In *Japan-U.S.A. Joint Seminar on Discrete Algorithms and Complexity Theory*, Kyoto, Japan, June 1986.

[47] Hilbert, D. and S. Cohn-Vossen. *Geometry and the Imagination.* (Translated by P. Nemenyi), Chelsea Publishing Company, New York, 1952.

[48] Johnson, D. S., and F. P. Preparata. "The densest hemisphere problem". *Theoretical Computer Science* 6: 93–107, 1978.

[49] Kabatiansky, G. A. and V. I. Levenshtein. "Bounds for packings on a sphere and in space". PPI 14(1): 3–25, 1978.

[50] Kannan, R. Private communications, 1990 at Carnegie Mellon University.

[51] Kao, M.-Y. "All graphs have cycle separators and planar directed depth-first search is in DNC". In *Proceedings of the 3rd Aegean Workshop on Computing, VLSI Algorithms and Architectures*, J. H. Reif editor, 53–63, 1988.

[52] Karp, R. M. "Reducibility among combinatorial problems". In *Complexity of Computer Computation*, R. E. Miller and J. W. Thatcher editor, Plenum, New York, 85–103, 1972.

[53] Khachiyan, L. Private communications, 1990 at Carnegie Mellon University.

[54] Knuth, D. E. *The Art of Computer Programming.* vol 1, vol 2, vol 3. Addison-Wesley, 1973.

[55] Leighton, F. T. *Complexity Issues in VLSI.* Foundations of Computing. MIT Press, Cambridge, MA, 1983.

[56] Leiserson, C. E. *Area Efficient VLSI Computation.* Foundations of Computing. MIT Press, Cambridge, MA, 1983.

[57] Li, M. "Simulating two pushdown stores by on tape in $O(n^{1.5}\sqrt{\log n})$ time". In *26th Annual Symposium on Foundations of Computer Science, IEEE*, 56–64, 1985.

[58] Lovász, L. *An Algorithmic Theory of Numbers, Graphs and Convexity.* SIAM. CBMS-NSF Regional Conference Series in Applied Mathematics, 1986.

[59] Lipton, R. J., D. J. Rose, and R. E. Tarjan. "Generalized nested dissection". *SIAM J. on Numerical Analysis*, 16:346–358, 1979.

[60] Lipton, R. J. and R. E. Tarjan. "A separator theorem for planar graphs". *SIAM J. of Appl. Math.*, 36:177–189, April 1979.

[61] Lipton, R. J. and R. E. Tarjan. "Applications of planar separator theorem". *SIAM J. Comput*, 9(3): 615–627, Auguest 1981.

[62] Matousek, J. "Approximations and optimal geometric divide-and-conquer". In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing.* 512–522, 1991.

[63] Mcl. D. and Y. Sommerville. *An Introduction to the Geometry of n dimensions*, Dover, New York, 1958.

[64] Miller, G. L. "Finding small simple cycle separators for 2-connected planar graphs". *Journal of Computer and System Sciences*, 32(3):265–279, June 1986.

[65] Miller, G. L. and S.-H. Teng. "Systematic methods for tree based parallel algorithm development". In *Second International Conference on Supercomputing*, Santa Clara, 392–403, 1987.

[66] Miller, G. L. and S.-H. Teng. *Center Points and Point Divisions.* manuscript, School of Computer Science, Carnegie Mellon University, 1990.

[67] Miller, G. L. and W. Thurston. "Separators in two and three dimensions". In *ACM STOC*, pages 300–309, Baltimore, May 1990.

[68] Miller, G. L., S.-H. Teng, and S. A. Vavasis. "An unified geometric approach to graph separators". In *32nd Annual Symposium on Foundations of Computer Science, IEEE*, to appear, 1991.

[69] Miller, G. L. and S. A. Vavasis. "Density graphs and separators". In *Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 331–336, San Fransciso, January 1991. ACM-SIAM.

[70] Onn S. *On the Radon number of the integer lattice* Technical Report No. 870, October 1989.

[71] Paterson, M. S. "Tape bounds for time-bounded turing machines". *J. Comp. Syst. Sci.*, 6:116–124, 1972.

[72] Pan, V. and J. Reif. "Efficient parallel solution of linear systems". In *ACM STOC*, pages 143–152, Providence,RI, May 1985.

[73] Preparata, F. P. and J. E. Vuillemin. "The Cube-Connected-Cycles: a versatile network for parallel computation". *CACM* 8(5): 300–309, 1981.

[74] Raghavan, P. "Probabilistic construction of deterministic algorithms approximating packing integer programs". In *27th Annual Symposium on Foundations of Computer Science, IEEE*, 10–18, 1986.

[75] Reif, J. H. and S. Sen. "Polling: a new randomized sampling technique for computational geometry". In *Proceedings of the 21st annual ACM Symposium on Theory of Computing.* 394–404, 1989.

[76] Reiss, S. P. and D. P. Dobkin. "The complexity of linear programming". Technical Report–69, Department of Computer Science, Yale University. 1976.

[77] Schönhage, A. and V. Strassen. "Schnelle multiplikation grosser zahlen". *Computing*, 7: 281–292. 1971.

[78] Schrijver, A. *Theory of Linear and Integer Programming.* Wiley, 1986.

[79] Smith, J. R. "Parallel algorithms for depth first search I. planar graphs". *SIAM Journal on Comput*, 15(3):814–830, Auguest 1986.

[80] Strang, G., and G. J. Fix. *An Analysis of the Finite Element Mothod*, Prentice-Hall, 1973.

[81] Strassen, V. "Gaussian elimination is not optimal". *Numerische Mathematik*, 13: 254–256. 1969.

[82] Teng, S.-H. "Functional inversion and communication complexity". *CRYPTO, IACR*, to appear, 1991.

[83] Thurston, W. P. *The Geometry and Topology of 3-manifolds.* Princeton University Notes, 1988.

[84] Thomassen, C. "Planarity and duality of finite and infinite graphs". *Journal of Combinatorial Theory*, Series B, 29: 244–271, 1980.

[85] Thompson, C. D. *A Complexity Theory for VLSI.* PhD thesis, Carnegie-Mellon University, Department of Computer Science, 1980.

[86] Tutte, W. T. "Convex representations of graphs". *Proc. London Math. Soc.* 10(3): 304–320, 1960.

[87] Tutte, W. T. "How to draw a graph". *Proc. London Math. Soc.* 13(3): 743–768, 1963.

[88] Tverberg, H. "A generalization of Radon's theorem". *Journal London Math Society*, 123–128, 1966.

[89] Ullman, J. D. *Computational Aspects of VLSI.* Computer Science Press, Rockville MD, 1984.

[90] Ungar, P. "A theorem on planar graphs". *Journal London Math Soc.* 26: 256–262, 1951.

[91] Vaidya, P. M. "An $O(n \log n)$ algorithm for the All-nearest-neighbors problem". *Discrete & Computational Geometry*,4:101–115, 1989.

[92] Valiant, L. G. "Universality consideration in VLSI circuits". *IEEE Transaction on Computers*, 30(2): 135–140, February, 1981.

[93] Vapnik, V. N. and A. Ya. Chervonenkis. "On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16: 264-280, 1971.

[94] Vavasis, S. A. "Automatic domain partitioning in three dimensions". *SIAM J. Sci. Stat. Comp.*, July 1991. to appear.

[95] Vavasis, S. A. Private Communication, 1991.

[96] Wyner, A. D. "Capabilities of bounded discrepancy decoding". *BSTJ*, 44: 1061–1122, 1965.

[97] Yao, F. F. "A 3-space partition and its application." In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, ACM*, 258–263, 1983.