DISS. ETH No. 15901, 2005

Structural Parameters in Combinatorial Objects

A dissertation submitted to the
Swiss Federal Institute of Technology Zürich
for the degree of Doctor of Technical Sciences

presented by
Yoshio Okamoto
Master of Systems Science, the University of Tokyo, Japan
born July 22, 1976 in Hekinan, Japan
citizen of Japan

submitted to
Prof. Emo Welzl, ETH Zürich, examiner
Prof. Komei Fukuda, ETH Zürich, co-examiner

Acknowledgements

First of all, I would like to express my sincere gratitude to my supervisor Emo Welzl. It was a great opportunity for me to work at his teaching/research group in Zurich. His persistent encouragement and support were really helpful, and his way of looking at problems, way of presenting materials, and everything were very exciting to me. He was also a speaker of the Berlin/Zurich Joint Graduate Program "Combinatorics, Geometry, and Combinatorics" (CGC) financed by ETH Zurich and German Science Foundation. Without this program, I would not have imagined how I completed the study.

My sincere gratitude also goes to Komei Fukuda who kindly admitted to be a co-referee. Indeed, he was the first person who I made a contact when I tried to move to Zurich. At that time we did not know each other; Nevertheless he was so kind that he gave a lot of information and encouraged me to come to Zurich.

Some of the chapters in the thesis are based on joint papers with the following people: Vladimir Deĭneko, Michael Hoffmann, Kenji Kashiwabara, Masataka Nakamura, Takeaki Uno and Gerhard Woeginger. I would like to thank them for fruitful discussion.

It was my great pleasure to share the time in Zurich with current and former members of the research group: Udo Adamy, Christoph Ambühl, Alexx Below, Robert Berke, Péter Csorba, Kaspar Fischer, Bernd Gärtner, Jochen Giesen, Matthias John, Michael Hoffmann, Dieter Mitsche, Leo Rüst, Eva Schuberth, Ingo Schurr, Shakhar Smorodinsky, Simon Spalinger, Bettina Speckmann, Miloś Stojakovič, Tibor Szabó, Csaba Tóth, Falk Tschirschnitz, and Uli Wagner. I would like to thank all of them, and the same thanks also go to our long-term visitors: József Beck, Jiří Matoušek, Sonoko Moriyama, and Frans Wessendorp. Especially for this thesis, Michael helped me a lot around the TEXnical stuffs; Robert kindly did a thorough proof-reading; Dieter was also kind enough to translate the abstract into German.

I am also indebted to Floris Tschurr and Franziska Hefti-Widmer for a lot of administrative matters. Without their help, I could have not finished my Ph.D.

Several seminars and workshops held at here in ETH Zurich and/or organized by ETH people greatly inspired me. I would like to have an opportunity to thank Markus Bläser, Thomas Erlebach, Komei Fukuda, Dmitry Kozlov, Angelika Steger, Emo Welzl, Peter Widmayer and their group members for keeping a nice atmosphere there.

As mentioned already, my study in Zurich was greatly supported by the Berlin/Zurich Joint Graduate Program "Combinatorics, Geometry, and Combinatorics" (CGC) financed by ETH Zurich and German Science Foundation (DFG). I would like to acknowledge the faculty members and the colleagues of CGC, both in Berlin and Zurich, and a financial support from ETH and DFG. Special thanks go to Hiroyuki Miyazawa, who was the first CGC student from Japan and encouraged me to move to Zurich.

I am also grateful to the group at the University of Tokyo, where I spent my master study. The (current and former) members are Masataka Nakamura, Kenji Kashiwabara, Yoshimasa Hiraishi, Tadashi Sakuma, Takashi Takabatake, Masahiro Hachimori, Masaharu Kato, Kei Sakamoto, and Aiko Kurushima. Sometimes I visited them, and from discussion with them I could elaborate several ideas. I would also like to thank the KTYY seminar members at the University of Tokyo for their kindness. I am also grateful to many other researchers in Japan, especially the members of Combinatorial Mathematics Seminar and Combinatorial Optimization Seminar. They kindly (financially and/or spiritually) supported my visits in Japan.

Finally I would like to express my gratitude to my family for their warm support.

4th March 2005, Zurich
Yoshio Okamoto

Abstract

When we deal with combinatorial objects mathematically or algorithmically, we may observe that the intrinsic difficulty is governed by some structural parameters. The use of structural parameters is diverse: it often happens that finding a good structural parameter opens a door to proofs of mathematical statements. On the computational side, there are lots of algorithmic results stating that an algorithm is more efficient if a certain structural parameter is smaller. Also, successful structural parameters are further generalized or specialized for particular purposes so that the use can fit into broader context. Treewidths of graphs and VC-dimension used in discrete and computational geometry are such examples.

This thesis tries to identify some nice structural parameters for three combinatorial or geometric objects. It consists of three rather independent parts.

In Part I, we consider the clique complex of a graph, which is the family of all cliques in the graph and is a special independence system (i.e., closed under taking subsets). It is known that every independence system is the intersection of finitely many matroids, and a natural greedy algorithm gives a solution of value at most k times away from the optimal value for the maximum weight independent set problem when the independence system is the intersection of k matroids. Therefore, we regard this k as a nice structural parameter. Our main result is the characterization of the clique complexes which are the intersections of k matroids for each natural number k. The same question was asked by Fekete, Firla & Spille for matching complexes, but since a matching complex is a special clique complex, our result is more general than theirs. Several related results are provided.

In Part II, we study abstract convex geometries introduced by Edelman & Jamison in 1985. An abstract convex geometry is a combinatorial abstraction of convexity concepts appearing in a lot of objects such as point configurations, partially ordered sets, trees and rooted graphs, and it is defined in a purely combinatorial way. Our result states that actually each abstract convex geometry can be obtained from some point configuration. This result can be seen as an analogue of the topological representation theorem for oriented matroids by Folkman

& Lawrence. However, our theorem gives an affine-geometric representation of an abstract convex geometry. This suggests the intrinsic simplicity of abstract convex geometries, and that the minimum dimension of a representation can be considered as a good structural parameter for an abstract convex geometry. As an application of our representation theorem, we study open problems raised by Edelman & Reiner about local topology of the free complex of an abstract convex geometry. We settle their problem affirmatively when the realization is 2-dimensional and separable. This can be seen as a first step to the solution of their problems.

In Part III, we design fixed-parameter algorithms for some geometric optimization problems. Fixed-parameter tractability is a concept capturing hardness of the problem when some parameter associated to the problem is small. We consider the number of inner points as a parameter for geometric optimization problems on a 2-dimensional point set. Since many of such problems can be solved in polynomial time when the number of inner points is zero (i.e., the points are in convex position), this parameter should be a nice choice. To support this intuition, we consider two specific problems, namely the traveling salesman problem and the minimum weight triangulation problem. For both of them, we devise fixed-parameter algorithms, and show that they can be solved in polynomial time when the number of inner points is at most logarithmic in the number of input points.

Zusammenfassung

Bei der mathematischen oder algorithmischen Betrachtung kombinatorischer Objekte stellt man fest, dass die intrinsische Schwierigkeit von einigen strukturellen Parametern der Objekte beherrscht wird. Die Verwendung struktureller Parameter dient zu verschiedenen Zwecken: einerseits tritt häufig der Fall ein, dass eine gute Auswahl eines strukturellen Parameters eine Tür zu Beweisen mathematischer Aussagen öffnet. Andererseits existiert eine Vielzahl von algorithmischen Ergebnissen, die zeigen, dass ein Algorithmus effizienter ist, wenn ein bestimmter struktureller Parameter klein ist. Zudem werden erfolgreiche strukturelle Parameter oftmals weiter verallgemeinert oder für bestimmte Zwecke weiter spezialisiert, um so in einem breiteren Zusammenhang Verwendung zu finden. Beispiele für solche Parameter sind Baumweiten von Graphen und VC-Dimension in diskreter Geometrie.

In dieser Arbeit wird versucht, einige gute strukturelle Parameter für drei kombinatorische oder geometrische Objekte zu identifizieren. Die Arbeit besteht aus drei unabhängigen Teilen.

In Teil I betrachten wir den Clique-Komplex eines Graphen, unter dem die Familie aller Cliquen des Graphen verstanden wird und der ein spezielles Unabhängigkeitssystem darstellt. Es ist bekannt, dass jedes Unabhängigkeitssystem Durchschnitt von endlich vielen Matroiden ist, und dass ein natürlicher Greedy-Algorithmus für das Problem einer unabhängigen Menge maximalen Gewichts einen Wert liefert, der um höchstens einen Faktor k von der optimalen Lösung abweicht, wenn das Unabhängigkeitssystem Durchschnitt von k Matroiden ist. Daher betrachten wir dieses k als guten strukturellen Parameter. Das Hauptresultat dieses Abschnitts ist die Charakterisierung von Clique-Komplexen, die Durchschnitt von k Matroiden sind (für jede natürliche Zahl k). Die gleiche Frage wurde von Fekete, Firla und Spille für Matching-Komplexe gestellt. Da ein Matching-Komplex ein Spezialfall eines Clique-Komplexes ist, ist unser Resultat wesentlich allgemeiner. In diesem Teil werden darüber hinaus weitere Resultate in Zusammenhang mit diesem Problem präsentiert.

Teil II behandelt das Studium abstrakter konvexer Geometrien, die von Edelman und Jamison 1985 eingeführt wurden. Eine ab-

strakte konvexe Geometrie ist eine kombinatorische Abstraktion von Konzepten der Konvexität, die in vielen Objekten wie Konfigurationen von Punkten, partiell geordneten Mengen, Bäumen oder Wurzelgraphen eine wichtige Rolle spielen; eine abstrakte konvexe Geometrie wird allerdings auf rein kombinatorische Weise definiert. Wir zeigen in diesem Abschnitt, dass man jede abstrakte konvexe Geometrie von bestimmten Konfigurationen von Punkten erhalten kann. Dieses Ergebnis kann als Analogon zum Satz der topologischen Darstellung von orientierten Matroiden von Folkman und Lawrence betrachtet werden. Im Gegensatz zu Folkman und Lawrence geben wir jedoch eine affin-geometrische Darstellung von abstrakten konvexen Geometrien an. Dies suggeriert zum einen die intrinsische Einfachheit von abstrakten konvexen Geometrien, zum anderen suggeriert es auch die Wahl der minimalen Dimension einer Darstellung als guten strukturellen Parameter. Als eine Anwendung unseres Theorems betrachten wir offene Probleme von Edelman und Reiner über die lokale Topologie des freien Komplexes einer abstrakten konvexen Geometrie. Wir lösen dieses Problem im positiven Sinn für den Fall, dass die Realisierung zweidimensional und separabel ist. Dies kann als erster Schritt zur vollständigen Lösung der Probleme von Edelman und Reiner betrachtet werden.

In Teil III entwerfen wir parametrisierte Algorithmen für einige geometrische Optimierungsprobleme. Parametrisierte Komplexität ist ein Konzept, um die Schwierigkeit des Problems zu erfassen, wenn ein bestimmter Parameter des Problems klein ist. Wir betrachten die Anzahl innerer Punkte als einen Parameter für geometrische Optimierungsprobleme auf einer zweidimensionalen Punktmenge. Da viele Probleme dieser Art in polynomieller Zeit gelöst werden können, wenn die Anzahl der inneren Punkte null ist (i.e., die Punkte sind in konvexer Lage), sollte dieser Parameter eine gute Auswahl sein. Um die Intuition zu bekräftigen, betrachten wir zwei spezielle Probleme: das Rundreiseproblem sowie das Problem der Triangulierung minimalen Gewichts. Für beide Probleme entwerfen wir parametrisierte Algorithmen und zeigen, dass die Probleme in polynomieller Zeit gelöst werden können, wenn die Anzahl der inneren Punkte logarithmisch in der Anzahl der Eingabepunkte ist.

Contents

> The Phantom: Still hungry for an introduction!
>
> *The Phantom of the Opera* (1989)

Chapter 0

Introduction

0.1 Use of Structural Parameters

When we deal with combinatorial objects mathematically or algorithmically, we may observe that the intrinsic difficulty is governed by some structural parameters. The use of structural parameters is diverse. It often happens that finding a good structural parameter opens the door to proofs of mathematical statements. On the computational side, there are lots of algorithmic results stating that an algorithm is more efficient if a certain structural parameter is small. Also, successful structural parameters are further generalized and specialized for particular purposes so that they can fit into a broader context.

This thesis does not aim at defining what a "structural parameter" is. We think that it makes little sense to define it formally; it is not the main goal of this thesis to classify several parameters as structural ones or non-structural ones. However, we would like structural parameters to have some characteristics. First of all, since it is structural, the parameter value must only depend on the object itself. Second, it must reflect some "complexity" of the object; for example, if an object has a smaller parameter value, then the problem under investigation should be easier to solve.

Indeed, discrete mathematics and theoretical computer science have been investigating a lot of structural parameters. To obtain some intuition, let us look at two prominent examples: one from graph theory and one from discrete and computational geometry.

0.1.1 Example: the Graph Minor Theorem and Treewidth

The first example comes from the graph minor theorem, which has been known as Wagner's conjecture [Wag37] and has recently had a complete proof by the series of articles of Robertson & Seymour [RS83, RS86a, RS84, RS90a, RS86b, RS86c, RS88, RS90c, RS90b, RS91, RS94, RS95b, RS95c, RS95a, RS96, RS03a, RS99, RS03b, RS04a, RS04b]. (Diestel's book [Die00] contains a relatively short exposition of that proof.) The following is a statement of the graph minor theorem: every infinite family of finite graphs contains two graphs such that one is a minor of the other.

Before Robertson & Seymour, it was proven that the graph minor theorem is true for trees by Kruskal [Kru60], and they relied on this fact. One of the ideas which Robertson & Seymour had was to introduce a tree-decomposition and the treewidth of a graph. We are not going to define what a tree-decomposition and a treewidth are, but at least it is good to notice that the treewidth can be seen as a parameter representing how close a graph is to a tree. (For example, a graph has treewidth 1 if and only if it is a tree.) With these concepts accompanied by a lot of far-from-trivial lemmas, they are able to apply the proof strategy which Kruskal used in order to complete the proof of the graph minor theorem.

After the invention, tree-decompositions and treewidths have become important in graph algorithms and a deep theory has arisen. For example, a result by Courcelle [Cou91] states that any decision problem which can be described by the so-called monadic second order logic can be solved in linear time for graphs of bounded treewidth, and the result also gives such an algorithm assuming that a tree-decomposition is given. (Since Bodlaender [Bod96] provides a linear-time algorithm to compute a tree-decomposition when a given graph has bounded treewidth, this assumption is reasonable.) This is generalized to optimization problems and counting problems by Arnborg, Lagergren & Seese [ALS91]. There are even linear-time algorithms which do not require a tree-decomposition [ACPS93, BvAdF01]. Survey articles by Bodlaender [Bod93, Bod97, Bod05] nicely explain some of the algorithmic aspects on treewidths.

Recent development on fixed-parameter algorithms for planar graphs is also influenced by tree-decompositions. They are nicely explained in Niedermeier's habilitation thesis [Nie02]. Most recent results are given by Fomin & Thilikos [FT04]. This direction of research is extended to the so-called bidimensional graph problems, and there a local treewidth plays an important role. Demaine & Hajiaghayi [DH05a, DH05b] gave surveys on this topic.

Thanks to the algorithmic success of treewidths, similar notions have been introduced. One of the most important one is the cliquewidth, which was introduced by Courcelle, Engelfriet & Rozenberg [CER93] and was related to graph algorithms by Courcelle & Olariu [CO00]. Analogously to the treewidth, there is a result due to Courcelle, Makowsky & Rotics [CMR00] stating that every optimization problem which can be expressed by the monadic second order logic is linear-time solvable for graphs of bounded cliquewidth.

0.1.2 Example: Vapnik–Chervonenkis Dimension in Discrete and Computational Geometry

Vapnik & Chervonenkis [VC71] introduced a concept which is nowadays called the Vapnik–Chervonenkis dimension or shortly VC-dimension. (The name was coined by Haussler & Welzl [HW87].) Vapnik & Chervonenkis defined this concept in the context of statistical learning theory. The VC-dimension becomes one of the super-key concepts in the theory of empirical processes and computational learning theory [Vap98, SS01]. One of the theorems by Vapnik & Chervonenkis is the following. For a set system of VC-dimension d, there exists an ε-approximation of the system of size $O(\frac{d}{\varepsilon^2}\log\frac{d}{\varepsilon})$, and such an ε-approximation can be found by random sampling. Here, an ε-approximation of a set system is a subset of the ground set which approximates the size of every set in the family within an additive error ε. (We are not going to give a formal definition.)

A concept similar to an ε-approximation was introduced by Haussler & Welzl [HW87] while they studied the simplex range query in arbitrary dimension, which is a problem in computational geometry. They proved the so-called ε-net theorem, which states that for a set system of VC-dimension d, there exists an ε-net of the system of size $\mathrm{O}(\frac{d}{\varepsilon}\log\frac{d}{\varepsilon})$, and such an ε-net can be found by random sampling. Here, an ε-net of a set system is a subset of the ground set which intersects every set in the family that has at least ε-fraction many elements of the size of the ground set. With proper definitions, we can show that an ε-approximation is an ε-net, but not the other way around in general. However, the results above show that ε-nets are better in size than ε-approximations by a factor of $1/\varepsilon$.

Haussler & Welzl [HW87] applied the ε-net theorem to give a better upper bound for the simplex range query, and it turned out that the concept of VC-dimension is quite useful in discrete and computational geometry. For example, in the proof by Alon & Kleitman of the (p,q)-theorem (for convex sets), which was proposed by Hadwiger & Debrunner, the ε-net theorem plays a key role. Recently, Matoušek [Mat04] showed that any set system of bounded VC-dimension has a (p,q)-theorem. So, families of bounded VC-dimension behave nicely. Additionally, as an answer to a conjecture by Kavraki, Latombe, Motwani & Raghavan [KLMR98] on the art gallery problem, Kalai & Matoušek [KM97] proved that if X is a compact simply-connected set in the plane of Lebesgue measure 1, such that any point $x \in X$ sees a part of X of measure at least ε, then one can choose a set G of at most $\mathrm{O}(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon})$ points in X such that any point of X is seen by some point of G. To show that, they just needed to prove that a certain set system has bounded VC-dimension.

On the computational side, Matoušek [Mat95] designed an efficient deterministic algorithm to construct an ε-net for a set system of bounded VC-dimension. (A simplified presentation is given by Chazelle & Matoušek [CM96], for example.) This was used for derandomizing many geometric algorithms, for example, by Har-Peled [HP99], by Chazelle & Matoušek [CM96], and by Ramos [Ram01].

The ε-net theorem and the notions around it play a great role in discrepancy theory. Books by Matoušek [Mat99], Chazelle [Cha01] and Matoušek [Mat02] are nice guides for the ε-net theorem and related materials

With its success, VC-dimension is generalized in several ways. For example, Matoušek [Mat99, Exercise 5.2.10] discusses a generalization of VC-dimension to k-valued functions. (The usual VC-dimension corresponds to binary functions.) Raz [Raz00] introduced the concept of VC-dimension of a set of permutations.

0.1.3 Viewpoint of the Thesis

Treewidth and VC-dimension are just examples of structural parameters; structural parameters are all around in combinatorial objects. However, we hope that these two examples convince us so that structural parameters in combinatorial objects should be of interest.

Looking at the two examples above, we observe that the first key step was actually to find a nice structural parameter of objects. Then, the associated results were established and the theory has been enriched. In this thesis, we study three kinds of discrete objects with their structural parameters. Our emphasis is on the identification of new structural parameters which we expect to be useful for discrete mathematics and theoretical computer science. (So, we are still hungry for parameters.)

0.2 Summary and Organization

According to the objects we study, the rest of the thesis is decomposed into three parts. Table 1 shows the correspondence of the parts, the objects we look at, the structural parameters and the type of results we obtain. Each part is written in a self-contained way so that the reader can start anywhere. (As a consequence, some definitions are repeated.)

0.2.1 Part I

In the first part, we consider independence systems with their relation to matroids. An independence system is a set system such that any subset of a member of the system is also a member. Depending on

	Object	Structural Parameter	Type of Result
I	clique complex	the minimum number of matroids we need to represent a given clique complex as their intersection	characterization
II	abstract convex geometry	the minimum dimension of an affine realization	realization theorem and its application to topological combinatorics
III	planar point set	the number of inner points	fixed-parameter algorithms

Table 1: Correspondence of the concepts and the results

the context, an independence system is called an abstract simplicial complex (in combinatorial topology and topological combinatorics), a hereditary hypergraph (in hypergraph theory) and an order ideal (in poset theory). An "independence system" is a term mainly used in matroid theory, which was initiated by Whitney [Whi35]. A matroid is an independence system which additionally satisfies the so-called augmentation axiom. Good sources about matroids are a book by Oxley [Oxl92] on fundamental properties of matroids and representability questions, and the three-volume book by Schrijver [Sch03] on relations to combinatorial optimization. We often encounter matroids in combinatorics. For example, if we have a set of vectors in some vector space, the independent subsets of the vectors form a matroid. If we have a graph, the edge sets of its forests form a matroid. From lattice theory, it is known that a geometric lattice is an equivalent notion to a matroid. In theory of hyperplane arrangements, the enumerative aspects are well treated through matroids [Sta04]. Polyhedral combinatorics has been founded on the basis of matroid theory [Sch03]. Matroids have numerous applications in combinatorial design [Dez92], combinatorial optimization [Sch03] which leads to submodular-type optimization [Fuj91] and further to discrete convex analysis [Mur03], secret sharing schemes in cryptography [BD91, NW01], rigidity [Whi92, Rec89, GSS93], electric engineering [Rec89, Nar97], systems analysis [Rec89, Mur00], and so on.

It is folklore that every independence system is the intersection of finitely many matroids. Therefore, we can associate with each independence system the minimum number of matroids whose intersection is the independence system. We consider this number as a structural parameter of independence systems, and then we see that the larger the parameter value of an independence system is, the more complex the description of the system becomes. Another fact which supports this choice of a parameter is the following one due to Jenkyns [Jen76] and Korte & Hausmann [KH78]: the greedy algorithm approximates the optimal value within a factor of k for the maximum weight independent set problem in the intersection of k matroids. This means that the smaller the structural parameter of an independence system is, the better approximation the greedy algorithm gives.

As a foundation of this structural parameter, we try to characterize the independence systems which are the intersections of k matroids, for each $k \in \mathbb{N}$. Unfortunately we do not find an answer to this question, but for the special case of clique complexes we find an answer.

Given a graph, the clique complex of the graph is the collection of all cliques of the graph, where a clique of a graph is defined to be a vertex subset whose elements are pairwise adjacent. In the literature a clique complex is also called a flag complex. A clique complex is a natural independence system arising from a graph, and it has connections to many other questions in, for example, extremal combinatorics [Bol95], algebraic combinatorics [Ham90, CS04], topological combinatorics [CD95] and hypergraph theory [ABM03, Mes01, Mes03]. Furthermore, the class of clique complexes contains other important classes of independence systems such as the matching complexes of graphs and the order complexes of partially ordered sets.

The main theorem of this part is the following. The clique complex of a graph is the intersection of k matroids if and only if there exist k partitions of the vertex set of the graph into stable sets such that a pair of vertices is not an edge of the graph if and only if it is contained in some stable set from one of the partitions. This theorem gives a polynomial-time checkable certificate for a graph to have a clique complex which is the intersection of k matroids. Thus, the theorem implies that the corresponding decision problem belongs to NP.

As a corollary of the main theorem, we show a good characterization of graphs with clique complexes which are the intersections of two matroids. This enables us to determine in polynomial time whether the clique complex of a given graph is the intersection of two matroids. As well, we consider the following extremal problem: what is the maximum number of matroids we need for the representation of the clique complex of a graph with n vertices? We give a complete answer to this question: "$n-1$ matroids are always sufficient, and necessary in some case."

Recently the same problem was studied by Fekete, Firla & Spille [FFS03] for matching complexes. Since matching complexes are always clique complexes, we are able to derive most of their results as corollaries of our main theorem.

Part I is based on joint work with Kenji Kashiwabara & Takeaki Uno [KOU03].

0.2.2 Part II

Part II is devoted to an abstraction of the concept of convexity. Convexity is one of the central concepts in geometry. We study abstract convex geometries introduced by Edelman & Jamison [EJ85]. They are finite set systems fulfilling additional conditions, and obtained from diverse objects and processes. For example, a proper definition of a "convex set" in a finite point set gives rise to abstract convex geometries. This also applies to partially ordered sets, trees, chordal graphs, Ptolemaic graphs, and acyclic oriented matroids. Searching processes in a rooted graph give an abstract convex geometry as well.

Abstract convex geometries appear in a lot of seemingly unrelated contexts. Since it is an abstraction of a geometric concept, it naturally arises in discrete geometry [EJ85, ER00, ERW02]. In combinatorial optimization, abstract convex geometries are considered equivalent to antimatroids, where an antimatroid is a specialization of a greedoid and a greedoid is a generalization of a matroid [KLS91], so a certain bottleneck-type optimization problem can be solved by a greedy algorithm [BF90]. From lattice theory, we can see that finite lower semimodular lattices are equivalent to abstract convex geometries [EJ85, Ste99]. From the viewpoint of closure operators, abstract convex geometries can be seen as closure spaces with anti-exchange closure operators [EJ85, KLS91, And02]. Namely an abstract convex geometry is a nice counterpart of a matroid since a matroid can be seen as a closure space with exchange closure operator. In submodular-type optimization, it has turned out that abstract convex geometries are "essential" for the so-called dual greedy algorithm to work [KO03, Fuj04]. While pursuing a well-behaving structure in queuing theory, Glasserman & Yao [GY94] arrived at abstract convex geometries. In mathematical psychology, an abstract convex geometry arises in the study of knowledge spaces [FD88, Fal89, Kop98, DF99a]. Furthermore, a path-independence choice function in social choice theory is also an equivalent notion to an abstract convex geometry [JD01, Kos99, MR01, And02]. In scheduling, a generalization of usual precedence constraints to the so-called AND/OR precedence constraints gives rise to an abstract convex geometry [MSS04]. Therefore, the relationship with directed hypergraphs and Horn theory of boolean functions can also be seen.

An abstract convex geometry is introduced through an extraction of geometric convexity of finite point sets, and it is defined in a purely combinatorial manner. This is a direction from geometry to combinatorics. Our result exhibits a kind of the opposite direction. As a main result, we show that every abstract convex geometry can be defined via finite point sets. More precisely speaking, we define a generalized convex shelling, which is an abstract convex geometry specified by two finite point sets, and we prove that every abstract convex geometry is isomorphic to some generalized convex shelling. We call such a generalized convex shelling an affine realization of the abstract convex geometry.

The main theorem naturally provides a structural parameter for abstract convex geometries. Namely, we can take the minimum dimension of point sets of an affine realization of an abstract convex geometry. Therefore, our theorem gives a fresh view to the theory of abstract convex geometries.

The main theorem enables us to study abstract convex geometries in a geometric setting, which allows us to use geometric machineries. As an application of the main theorem we study an open problem posed by Edelman & Reiner [ER00]. The problem is concerned with local topology of a certain simplicial complex associated with an abstract convex geometry. We settle the problem affirmatively when abstract convex geometries have two-dimensional affine realizations.

The main theorem in this part is based on joint work with Kenji Kashiwabara & Masataka Nakamura [KNO05]. The result about topology is the author's individual work [Oka04].

0.2.3 Part III

Finally in Part III, we study the algorithmic aspects of a finite set of points in the Euclidean plane.

Several optimization problems on finite planar point sets are known to be solvable in polynomial time when the set forms the vertex set of some convex polygon. (In such a case the points are said to be in convex position.) However, the problems often become hard (or not known to be easy) when the points are arbitrarily placed. Having these two cases, we observe "inner points make the problems hard." Here we define an inner point of a point set as a point in the interior of the convex hull.

From this observation, we take the number of inner points as a structural parameter of a finite planar point set. Especially we look at some optimization problems which can be efficiently solved when this structural parameter is small. Then the next question would be "how large the parameter can get in order to assure a polynomial-time algorithm?"

To study this question, we adapt the viewpoint of parameterized computation. In parameterized computation we consider a parameterized problem, formally defined as a pair of a usual computational problem and a parameter. An algorithm for a parameterized problem is called a fixed-parameter algorithm (or an FPT algorithm) if it runs in time $\mathrm{O}(f(k)n^c)$ where n is the size of the input of the problem, k is the parameter, c is a constant independent of n and k, and f is a computable function. For example, an algorithm with running time $\mathrm{O}(3^k n^2)$ is allowed, but $\mathrm{O}(n^k)$ is not. If we consider the case where k is small, say $k = 10$, the first algorithm runs in $\mathrm{O}(n^2)$, while the second one runs in $\mathrm{O}(n^{10})$. Furthermore, the first algorithm runs in polynomial time in n even if $k = \mathrm{O}(\log n)$, while this is not true for the second one. So, the power of fixed-parameter algorithms is clear. Recently quite a few survey articles on parameterized computation have been written [Nie98, DFS99a, DFS99b, Fel01, Fel02, Dow03, Fel03a, Fel03b, Nie04, FG04, DM04] and books on this topic [DF99a, Nie02] are also available.

Niedermeier [Nie04] gives a state-of-the-art survey on parameterization in parameterized computation. He identifies several ways of parameterizing computational problems based on some previous results. The most common one is to parameterize with respect to outputs. For example, we take the size of an output as a parameter. Another view gives a parameterization with respect to inputs. The latter one is the parameterization that our result use, and indeed our study is one of the results which motivates Niedermeier [Nie04] to name it "distance-from-triviality" parameterization. Namely, in our case, the number of inner points measures the distance from the trivially solvable (or polynomially solvable) case in which the points are in convex position. This way of parameterization has only been studied for graph problems by Cai [Cai03], Marx [Mar04], and Guo, Hüffner & Niedermeier [GHN04], for the satisfiability problem by Szeider [Sze04], and for a string problem by Guo, Hüffner & Niedermeier [GHN04]. Therefore, our results are the first and unique fixed-parameter contribution to the distance-from-triviality approach for geometric problems.

As concrete examples, we study two geometric problems. The first one is the traveling salesman problem. In this problem, we are given n points in the Euclidean plane and we want to find a shortest tour (i.e., a way to visit all points and go back to where we started). This problem is in general NP-hard, as proven by Garey, Graham & Johnson [GGJ76] and independently by Papadimitriou [Pap77]. However, the problem becomes easy when the points are in convex position. Therefore, parameterization by the number of inner points makes sense. We design two fixed-parameter algorithms for this problem. The first one runs in $\mathrm{O}(k!kn)$ time and $\mathrm{O}(k)$ space, and the second one runs in $\mathrm{O}(2^k k^2 n)$ time and $\mathrm{O}(2^k kn)$ space. Here k represents the number of inner points among given n points. We also study some variants of the traveling salesman problem such as the prize-collecting traveling salesman problem and the partial traveling salesman problem, and give fixed-parameter algorithms for them as well.

The second problem we investigate is the minimum weight triangulation problem. In this problem, we are again given n points in the Euclidean plane, and we want to find a triangulation of the point set (i.e., a subdivision of the convex hull of the set into triangles such that each edge connects points from the set and each triangle contains no point from the set in its interior) of minimum weight. Here, the weight of a triangulation is measured by the sum of the lengths of its edges. In general, the minimum weight triangulation problem is not known to be solvable in polynomial time nor to be NP-hard. This is one of the open problems listed in the book by Garey & Johnson [GJ79] which are still unsolved. For the parameterized version of this problem, we give a fixed-parameter algorithm running in $O(6^k n^5 \log n)$ time. Here again, k represents the number of inner points among given n points.

The basic technique common to these algorithms is dynamic programming. As well, a well-established enumeration technique plays an important role.

The results on the traveling salesman problem in this part is based on joint work with Vladimir Deĭneko, Michael Hoffmann & Gerhard Woeginger [DHOW04], and the results on the minimum weight triangulation problem is based on joint work with Michael Hoffmann [HO04].

Part I

Graphs and Matroids

Chapter 1

> And it seems that overall, we are short of examples. The methods for coming up with useful examples (or counterexamples for commonly believed conjectures) are even less clear than the methods for proving.
>
> Gil Kalai (2000)

Matroid Representation of Clique Complexes

1.1 Introduction

An independence system is a family of subsets of a non-empty finite set such that all subsets of a member of the family are also members of the family. A lot of combinatorial optimization problems can be seen as optimization problems on the corresponding independence systems. For example, in the minimum cost spanning tree problem, we want to find a maximal set with minimum total weight in the collection of all forests of a given graph, and this collection is an independence system. In the maximum weight matching problem we consider the collection of all matchings of a given graph. This is also an independence system. More examples are provided by Korte & Vygen [KV02]. In this chapter, we study independence systems arising from the maximum weight clique problem.

A clique in a graph is a subset of the vertex set which induces a complete graph. In the maximum weight clique problem, we are given a graph and a weight function on the vertex set, and we want to find a clique which maximizes the total weight of its vertices. As is well known, the maximum weight clique problem is NP-hard even if the weight function is constant [GJ79]. This means that there exists no polynomial-time algorithm for this problem unless P = NP. Moreover, Håstad [Hås99] proved that there exists no polynomial-time algorithm for this problem which approximates the optimal value within a factor $n^{1-\varepsilon}$ for any $\varepsilon > 0$ unless NP = ZPP. (Here, n stands for the number of vertices in a given graph.) Therefore, the maximum clique problem is deeply inapproximable. Thus, one wants to determine classes of graphs on which they can perform well. To do that, we adapt the viewpoint of independence systems and matroids. For the maximum weight clique problem, we look at the family of all cliques of a graph as an independence system. Such an independence system is called a clique complex.

It is known that every independence system can be represented as the intersection of a finite number of matroids. Jenkyns [Jen76] and Korte & Hausmann [KH78] showed that, for the maximum weight base problem on an independence system which can be represented as the intersection of k matroids, a natural greedy algorithm approximates the optimal value within a factor k. (Their result can be seen as a generalization of the validity of the greedy algorithm for matroids, shown by Rado [Rad57] and Edmonds [Edm71], although their results showed that the validity of the greedy algorithm even characterizes matroids.) Thus, this number k is a measure of "how complex an independence system is with respect to the corresponding optimization problem."

Here, we want to notice the importance of clique complexes in fields other than combinatorial optimization. In extremal combinatorics, the f-vector of a clique complex (namely, the sequence of the numbers of cliques of all sizes in a graph) is studied in connection with Turán's problem. (See Bollobás [Bol95].) Related to that, in algebraic combinatorics, problems on the roots of the f-polynomial of a clique complex are studied. For example, Hamidoune [Ham90] asked whether the f-polynomial of the clique complex of a graph whose complement is claw-free has only real roots. (It was only recently that the problem has been solved by Chudnovsky & Seymour [CS04].) Also, Charney & Davis [CD95] made a conjecture on clique complexes which triangulate a homology sphere of odd dimension. For this topic, see Stanley's survey article [Sta00]. Finally, in topological combinatorics, when we refer to the topology of a graph, it usually means the topology of the clique complex of the graph. The topology of clique complexes plays an important role when one investigates Hall-type theorems in hypergraphs [ABM03, Mes01, Mes03]. Similarly, when we refer to the topology of a partially ordered set, it usually means the topology of the order complex of the partially ordered set, which turns out to be a clique complex.

In this chapter, we investigate how many matroids we need for the representation of the clique complex of a graph as their intersection. We show that the clique complex of a given graph G is the intersection of k matroids if and only if there exists a family of k stable-set partitions of G such that every edge of $\overline{G}$ (the complement of G) is contained in a stable set of some stable-set partition in the family. This theorem implies that the following decision problem belongs to NP: given a graph G and a natural number $k > 0$, determine whether the clique complex of G has a representation by k matroids or not. This is not a trivial fact since in general the size of an independence system can be exponential. As another consequence, we show that the class of clique complexes is the same as the class of the intersections of partition matroids. This may open a new direction of research to attack some open problems on clique complexes.

Formerly, Fekete, Firla & Spille [FFS03] investigated the same problem for matching complexes, and they characterized a graph whose matching complex is the intersection of k matroids, for every natural number k. Since the matching complexes form a subclass of the class of clique complexes, we observe that some of their results can be derived from our theorems as corollaries.

Further, we consider an extremal problem related to our theorem. Namely, we determine how many matroids are necessary and sufficient for the representation of every graph with n vertices. This number turns out to be $n-1$. We also investigate the case of two matroids more thoroughly. This case is especially important since the maximum weight base problem can be solved exactly in polynomial time for the intersection of two matroids by Frank's algorithm [Fra81]. (Namely, in this case, the maximum weight clique problem can be solved in polynomial time for any non-negative weight vector.) There, we will see that an algorithm by Protti & Szwarcfiter [PS02] checks whether a given clique complex has a representation by two matroids or not in polynomial time. Additionally, we show that the clique complex of a graph G is the intersection of k matroids if and only if G itself is the intersection of k matroids. (Here, we regard graphs themselves as independence systems of rank 2.) Thus, this reveals the intimate relationship between a graph and its clique complex in terms of matroid intersection.

The organization of this chapter is as follows. In Section 1.2, we introduce some terminology on independence systems. The proof of the main theorem is given in Section 1.3. Some immediate consequences of the main theorem are also given there. In Section 1.4, we consider an extremal problem related to our theorem. In Section 1.5, we investigate the case of two matroids. In Section 1.6, we study a graph itself as an independence system and relate it to our theorem. In Section 1.7, we deduce some results by Fekete, Firla & Spille [FFS03] from our theorems. We conclude with Section 1.8.

1.2 Preliminaries

1.2.1 Graphs

A *graph* is a pair $G = (V, E)$ of a finite set V, called the vertex set of G, and a family $E \subseteq \binom{V}{2}$ of two-element subsets of V, called the edge set of G. An element of V is called a *vertex* of V, and an element of E is called an *edge* of V. The vertex set of G is denoted by $V(G)$ and the edge set of G is denoted by $E(G)$. When we make a picture of a graph for illustration, we draw vertices as points and edges as arcs connecting two corresponding vertices.

A *subgraph* of a graph $G = (V, E)$ is a graph H such that $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. For a vertex subset $W \subseteq V(G)$, the subgraph *induced by* W is a subgraph H defined as $V(H) := W$ and $E(H) := E(G) \cap \binom{W}{2}$. The subgraph induced by W is denoted by $G[W]$. The *complement* of G is a graph H defined as $V(H) := V$ and $E(H) := \binom{V}{2} \setminus E$, and denoted by $\overline{G}$. A *complete graph* is a graph in which every two vertices form an edge. A *clique* of a graph $G = (V, E)$ is a subset $K \subseteq V$ such that the induced subgraph $G[K]$ is complete. A *stable set* of a graph $G = (V, E)$ is a subset $S \subseteq V$ such that the induced subgraph $G[S]$ contains no edge.

For a graph $G = (V, E)$, the *degree* of a vertex $v \in V$ is the number of edges containing v. The *maximum degree* of G is the maximum of the degrees over all vertices, and denoted by $\Delta(G)$. A *proper k-coloring* of G is a map from V to $\{1, \ldots, k\}$ (regarded as the set of "colors") such that every two vertices forming an edge are mapped to different colors. A *proper coloring* of G is a proper k-coloring of G for some $k \in \mathbb{N}$. The preimage of each color is called a *color class*. The *chromatic number* of G is the minimum $k \in \mathbb{N}$ such that a proper k-coloring of G exists. We denote the chromatic number of G by $\chi(G)$. A graph G is *k-colorable* if $\chi(G) \leq k$. Similarly, we may define a *proper k-edge-coloring* as a map from E to $\{1, \ldots, k\}$ such that every two edges sharing a vertex are mapped to different colors. A proper edge-coloring, a color class, the *edge-chromatic number* and the *k-edge-colorability* are defined in an analogous way. We denote the edge-chromatic number of G by $\chi'(G)$.

1.2.2 Independence Systems and Matroids

Now we introduce the notions of independence systems and matroids. For details, see Oxley's book [Oxl92]. Given a non-empty finite set V, an *independence system* on V is a non-empty family $\mathcal{I}$ of subsets of V such that $X \in \mathcal{I}$ implies $Y \in \mathcal{I}$ for all $Y \subseteq X \subseteq V$. The set V is called the *ground set* of the independence system. In the literature, an independence system is also called an *abstract simplicial complex*. A *matroid* is an independence system $\mathcal{I}$ additionally satisfying the following *augmentation axiom*: for $X, Y \in \mathcal{I}$ with $|X| > |Y|$ there exists $z \in X \setminus Y$ such that $Y \cup \{z\} \in \mathcal{I}$. For an independence system $\mathcal{I}$, a set X is called *independent* if $X \in \mathcal{I}$, and X is called *dependent* otherwise. A *base* of an independence system is a maximal independent set, and a *circuit* of an independence system is a minimal dependent set. (Notice that, in this chapter, we use the word "circuit" only for independence systems, not for graphs.) We denote the family of bases of an independence system $\mathcal{I}$ and the family of circuits of $\mathcal{I}$ by $\mathcal{B}(\mathcal{I})$ and $\mathcal{C}(\mathcal{I})$, respectively. Note that we can reconstruct an independence system $\mathcal{I}$ from $\mathcal{B}(\mathcal{I})$ as

$$\mathcal{I} = \{X \subseteq V \mid X \subseteq B \text{ for some } B \in \mathcal{B}(\mathcal{I})\}, \tag{1.1}$$

and from $\mathcal{C}(\mathcal{I})$ as

$$\mathcal{I} = \{X \subseteq V \mid C \not\subseteq X \text{ for all } C \in \mathcal{C}(\mathcal{I})\}. \tag{1.2}$$

It can be shown that $\mathcal{B}(\mathcal{I}_1) = \mathcal{B}(\mathcal{I}_2)$ if and only if $\mathcal{I}_1 = \mathcal{I}_2$; similarly $\mathcal{C}(\mathcal{I}_1) = \mathcal{C}(\mathcal{I}_2)$ if and only if $\mathcal{I}_1 = \mathcal{I}_2$. We can see that all the bases of a matroid have the same size from the augmentation axiom, but this is not necessarily the case for an independence system in general.

Let $\mathcal{I}$ be a matroid on V. We say that $x, y \in V$ are *parallel* in $\mathcal{I}$ if $\{x, y\}$ is a circuit of the matroid $\mathcal{I}$ or $x = y$. The next is a well known fact.

Lemma 1.1. *For a matroid, the relation that "x is parallel to y" is an equivalence relation on its ground set.*

Proof. Let $\mathcal{I}$ be a matroid on V. Choose three distinct elements $x, y, z \in V$ such that $\{x, y\}$ and $\{y, z\}$ are circuits of $\mathcal{I}$. We claim that $\{x, z\}$ is a circut of $\mathcal{I}$ as well. Since $\{x, y\}$ and $\{y, z\}$ are circuits, it holds that $\{x\} \in \mathcal{I}$ and $\{z\} \in \mathcal{I}$. Therefore, it suffices to show that $\{x, z\} \notin \mathcal{I}$.

For the sake of contradiction, suppose that $\{x, z\} \in \mathcal{I}$. Since $\{x, y\}$ is a circuit of $\mathcal{I}$, it holds that $\{y\} \in \mathcal{I}$. By the augmentation axiom for matroids, we have that $\{x, y\} \in \mathcal{I}$ or $\{y, z\} \in \mathcal{I}$. However, this contradicts the assumption that $\{x, y\}$ is a circuit of $\mathcal{I}$ (implying $\{x, y\} \notin \mathcal{I}$) and $\{y, z\}$ is a circuit of $\mathcal{I}$ (implying $\{y, z\} \notin \mathcal{I}$). This is a contradiction. $\square$

Let $\mathcal{I}_1$ and $\mathcal{I}_2$ be independence systems on the same ground set V. The *intersection* of $\mathcal{I}_1$ and $\mathcal{I}_2$ is just $\mathcal{I}_1 \cap \mathcal{I}_2$. The intersection of three or more independence systems is defined in a similar way. Note that the intersection of independence systems is an independence system as well. In addition, we have the following lemma. For a set system $\mathcal{F}$, we denote by $\mathrm{MIN}(\mathcal{F})$ the family of minimal sets in $\mathcal{F}$, namely,

$$\mathrm{MIN}(\mathcal{F}) := \{X \in \mathcal{F} \mid Y \not\subseteq X \text{ for any } Y \in \mathcal{F} \setminus \{X\}\}.$$

Lemma 1.2. *Let $\mathcal{I}_1, \ldots, \mathcal{I}_m$ be independence systems on the same ground set. Then, the family of circuits of $\bigcap_{i=1}^m \mathcal{I}_i$ is the family of the minimal sets in $\bigcup_{i=1}^m \mathcal{C}(\mathcal{I}_i)$, i.e.,*

$$\mathcal{C}\left(\bigcap_{i=1}^{m} \mathcal{I}_i\right) = \mathrm{MIN}\left(\bigcup_{i=1}^{m} \mathcal{C}(\mathcal{I}_i)\right).$$

Proof. We prove the case $m = 2$. The general case can be proven in the same way.

Since any two sets in $\mathrm{MIN}(\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2))$ have no inclusion relationship, there exists an independence system $\mathcal{I}$ which has $\mathrm{MIN}(\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2))$ as its family of circuits. We claim that $\mathcal{I}_1 \cap \mathcal{I}_2 = \mathcal{I}$.

First let us look at the outline of the proof with logic flavor.

$$
\begin{aligned}
& X \in \mathcal{I}_1 \cap \mathcal{I}_2 \\
& \quad \Leftrightarrow (X \in \mathcal{I}_1) \wedge (X \in \mathcal{I}_2) \\
& \quad \Leftrightarrow \neg((X \notin \mathcal{I}_1) \vee (X \notin \mathcal{I}_2)) \\
& \quad \Leftrightarrow \neg((\exists C_1 \in \mathcal{C}(\mathcal{I}_1) : C_1 \subseteq X) \vee (\exists C_2 \in \mathcal{C}(\mathcal{I}_2) : C_2 \subseteq X)) \\
& \quad \Leftrightarrow \neg(\exists C \in \mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2) : C \subseteq X) \\
& \quad \Leftrightarrow \forall C \in \mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2) : C \not\subseteq X \\
& \quad \Leftrightarrow \forall C \in \mathrm{MIN}(\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2)) : C \not\subseteq X \\
& \quad \Leftrightarrow X \in \mathcal{I}.
\end{aligned}
$$

The first equivalence is clear. The second one is de Morgan's law. The third equivalence is due to Equality (1.2). The fourth equivalence is clear. The fifth is again de Morgan's law in predicate calculus. The sixth needs an argument. We will discuss it later. The seventh one is again due to Equality (1.2).

Now, we claim that the sixth equivalence is true. Since $\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2) \supseteq \mathrm{MIN}(\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2))$, the direction from left to right (or from top to bottom) is true. How about the opposite direction? Assume that $C \not\subseteq X$ for any $C \in \mathrm{MIN}(\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2))$. We want to show that $D \not\subseteq X$ for any $D \in \mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2)$. Fix $D \in \mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2)$ arbitrarily. Then, there exists a (not necessarily proper) subset D' of D such that $D' \in \mathrm{MIN}(\mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2))$. If X contains D, then X should also contain D'. However, by our assumption it holds that $D' \not\subseteq X$. This is a contradiction. Hence, X does not contain D, i.e., $D \not\subseteq X$. Thus, we have shown that $D \not\subseteq X$ for any $D \in \mathcal{C}(\mathcal{I}_1) \cup \mathcal{C}(\mathcal{I}_2)$. The proof is completed. □

The following well-known observation is crucial for this chapter.

Lemma 1.3. *Every independence system can be represented as the intersection of a finite number of matroids on the same ground set.*

Proof. Let $\mathcal{I}$ be an independence system, and denote the circuits of an independence system $\mathcal{I}$ by $C^{(1)}, \ldots, C^{(m)}$ (i.e., $\mathcal{C}(\mathcal{I}) = \{C^{(1)},$

$\ldots, C^{(m)}\})$. Consider the independence system $\mathcal{I}_i$ with a unique circuit $\mathcal{C}(\mathcal{I}_i) = \{C^{(i)}\}$ for each $i \in \{1, \ldots, m\}$. Then, it follows that

$$\mathcal{C}\left(\bigcap_{i=1}^{m} \mathcal{I}_i\right) = \mathrm{MIN}(\{C^{(1)}, \ldots, C^{(m)}\}) = \{C^{(1)}, \ldots, C^{(m)}\} = \mathcal{C}(\mathcal{I}),$$

where the first identity is by Lemma 1.2 and the second one is due to the fact that $\{C^{(1)}, \ldots, C^{(m)}\}$ is the family of circuits of $\mathcal{I}$ and no two circuits have inclusion relationship. The third identity is from the definition. Since the family of circuits determines an independence system uniquely, it follows that $\mathcal{I} = \bigcap_{i=1}^{m} \mathcal{I}_i$.

What remains to show is that the independence system $\mathcal{I}_i$ is a matroid for each $i \in \{1, \ldots, m\}$. To prove this claim, we have to check the augmentation axiom. Fix an arbitrary $i \in \{1, \ldots, m\}$ and take two sets $X, Y \in \mathcal{I}_i$ such that $|X| > |Y|$. Now, for the sake of contradiction, suppose that $C^{(i)} \subseteq Y \cup \{z\}$, namely $Y \cup \{z\}$ is dependent in $\mathcal{I}_i$, for all $z \in X \setminus Y$. Since $Y \in \mathcal{I}_i$, we know that Y does not contain $C^{(i)}$. Therefore, to have $C^{(i)} \subseteq Y \cup \{z\}$, it must hold that Y is a proper subset of $C^{(i)}$. This implies that $C^{(i)} = Y \cup \{z\}$. Since this holds for all $z \in X \setminus Y$, we can see that $X \setminus Y = \{z\}$. However, this implies that $X = Y \cup \{z\} = C^{(i)}$, which means that X is not independent in $\mathcal{I}_i$. This is a contradiction. $\square$

Note that the matroids $\mathcal{I}_1, \ldots, \mathcal{I}_m$ in the proof are actually graphic matroids. (A graphic matroid is an independence system isomorphic to the family of forests in some multigraph.) Therefore, Lemma 1.3 itself can be strengthened to "every independence system can be represented as the intersection of a finite number of graphic matroids on the same ground set," although it is not important for the discussion in the rest of the chapter.

Due to Lemma 1.3, we are interested in the representation of an independence system as the intersection of matroids. Following the construction in the proof of Lemma 1.3, we see that at most $|\mathcal{C}(\mathcal{I})|$ matroids are enough to represent $\mathcal{I}$ as their intersection. However, we might do better. In the rest of this chapter, we study clique complexes.

1.3 Clique Complexes and the Main Theorem

A graph gives rise to various independence systems. Among them, we study clique complexes.

The *clique complex* of a graph $G = (V, E)$ is the collection of all cliques of G. We denote the clique complex of G by $\mathfrak{C}(G)$. Note that the empty set is a clique and $\{v\}$ is also a clique for each $v \in V$. So we see that the clique complex is actually an independence system on V. We also say that an independence system is a clique complex if it is isomorphic to the clique complex of some graph. Notice that a clique complex is also called a *flag complex* in the literature.

Here, we give some examples of clique complexes. (We omit necessary definitions.) (1) The family of stable sets of a graph G is nothing but the clique complex of $\overline{G}$. (2) The family of matchings of a graph G is the clique complex of the complement of the line graph of G, which is called the *matching complex* of G. (3) The family of chains of a partially ordered set P is the clique complex of the comparability graph of P, which is called the *order complex* of P. (4) The family of antichains of a partially ordered set P is the clique complex of the cocomparability graph (i.e., the complement of the comparability graph) of P.

The next lemma may be folklore.

Lemma 1.4. *Let $\mathcal{I}$ be an independence system on a finite set V. Then, $\mathcal{I}$ is a clique complex if and only if the size of every circuit in $\mathcal{I}$ is two. In particular, the circuits of the clique complex of G are the edges of $\overline{G}$ (i.e., $\mathcal{C}(\mathfrak{C}(G)) = E(\overline{G})$).*

Proof. Let $\mathcal{I}$ be the clique complex of $G = (V, E)$. Since a single vertex $v \in V$ forms a clique, the size of each circuit in $\mathcal{I}$ is greater than one. Each dependent set of size two in $\mathcal{I}$ is an edge of the complement of G. Observe that they are minimal dependent sets since the size of each dependent set in $\mathcal{I}$ is greater than one. In order to show that they are the only minimal dependent sets, suppose that there exists a circuit C of size more than two in $\mathcal{I}$. Then each two elements in C form an edge of G because of the minimality of C. Hence C is a clique in G. However, this is a contradiction to the assumption that C is dependent in $\mathcal{I}$ (i.e., not a clique in G).

Conversely, let $\mathcal{I}$ be an independence system on V and assume that the size of every circuit of $\mathcal{I}$ is two. Now construct a graph $G' = (V, E')$ with $E' = \{\{u,v\} \in \binom{V}{2} \mid \{u,v\} \notin \mathcal{C}(\mathcal{I})\}$, and consider the clique complex $\mathfrak{C}(G')$. By the opposite direction which we have just shown above, we can see that the circuits of $\mathfrak{C}(G')$ are the edges of $\overline{G'}$, and they are the circuits of $\mathcal{I}$. Therefore we have that $\mathcal{C}(\mathfrak{C}(G')) = \mathcal{C}(\mathcal{I})$. Since the family of circuits uniquely determines an independence system, this concludes that $\mathcal{I}$ is the clique complex of G'. $\square$

Now, we start studying the number of matroids which we need for the representation of a clique complex as their intersection. For a graph G, denote by $\mu(G)$ the minimum number of matroids such that the clique complex $\mathfrak{C}(G)$ is the intersection of them. Namely,

$$\mu(G) := \min \left\{ k \mid \mathfrak{C}(G) = \bigcap_{i=1}^{k} \mathcal{I}_i \text{ where } \mathcal{I}_1, \ldots, \mathcal{I}_k \text{ are matroids} \right\}.$$

First, we characterize the graphs G satisfying $\mu(G) = 1$ (namely the graphs whose clique complexes are indeed matroids). To do this, we define a partition matroid. A *partition matroid* is a matroid $\mathcal{I}(\mathcal{P})$ associated with a partition $\mathcal{P} = \{P_1, P_2, \ldots, P_r\}$ of V (that is, $V = \bigcup_{i=1}^{r} P_i$ and $P_i \cap P_j = \emptyset$ for all $i \neq j$), defined as

$$\mathcal{I}(\mathcal{P}) := \{I \subseteq V \mid |I \cap P_i| \leq 1 \text{ for all } i \in \{1, \ldots, r\}\}.$$

To justify the name "partition matroid" we need to show the following.

Lemma 1.5. *A partition matroid is a matroid.*

Proof. Let $\mathcal{I}(\mathcal{P})$ be the partition matroid associated with a partition $\mathcal{P} = \{P_1, \ldots, P_r\}$ of a ground set V. First we show that $\mathcal{I}(\mathcal{P})$ is an independence system. Let $I \in \mathcal{I}(\mathcal{P})$ and $J \subseteq I$. By definition, for each $i \in \{1, \ldots, r\}$ we have $|I \cap P_i| \leq 1$. Since $J \subseteq I$, we have $J \cap P_i \subseteq I \cap P_i$. Therefore, it follows that $|J \cap P_i| \leq |I \cap P_i| \leq 1$ for every $i \in \{1, \ldots, r\}$. This proved that $\mathcal{I}(\mathcal{P})$ is an independence system.

Now, we check that $\mathcal{I}(\mathcal{P})$ satisfies the augmentation axiom. Let $X, Y \in \mathcal{I}(\mathcal{P})$ such that $|X| > |Y|$. Then, there must exist $i \in \{1, \ldots, r\}$

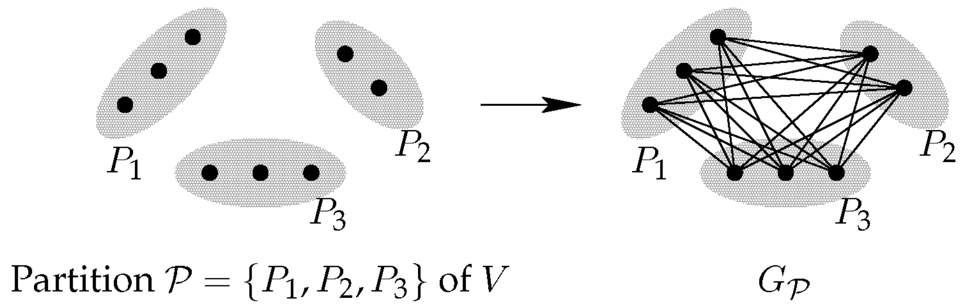

Figure 1.1: The correspondence of a partition matroid and a complete multipartite graph.

such that $|X \cap P_i| = 1$ and $|Y \cap P_i| = 0$. Let us denote $\{z\} := X \cap P_i$. We claim that $Y \cup \{z\} \in \mathcal{I}(\mathcal{P})$. Indeed, for $j \in \{1, \dots, r\} \setminus \{i\}$, we have

$$\begin{aligned} |(Y \cup \{z\}) \cap P_j| &= |(Y \cap P_j) \cup (\{z\} \cap P_j)| \\ &= |Y \cap P_j| + |\{z\} \cap P_j| \leq 1 + 0 = 1, \end{aligned}$$

and for i we have

$$|(Y \cup \{z\}) \cap P_i| = |Y \cap P_i| + |\{z\} \cap P_i| = 0 + 1 = 1.$$

The proof is completed. □

Next we observe that $\mathcal{I}(\mathcal{P})$ is a clique complex. Indeed we can see that $\mathcal{I}(\mathcal{P}) = \mathfrak{C}(G_{\mathcal{P}})$ if we construct the following graph $G_{\mathcal{P}} = (V, E)$ from $\mathcal{P}$: two vertices $u, v \in V$ are adjacent in $G_{\mathcal{P}}$ if and only if u and v are elements of distinct partition classes in $\mathcal{P}$. See Figure 1.1 for an illustration.

An alternative argument is to observe that

$$\mathcal{C}(\mathcal{I}(\mathcal{P})) = \left\{ \{u, v\} \in \binom{V}{2} \mid \{u, v\} \subseteq P_i \text{ for some } i \in \{1, \dots, r\} \right\}.$$

Then, we find out that $\mathcal{I}(\mathcal{P})$ satisfies the condition in Lemma 1.4, and this shows that $\mathcal{I}(\mathcal{P})$ is a clique complex. Note that $G_{\mathcal{P}}$ constructed above is a complete r-partite graph with the partition $\mathcal{P}$. (In Figure 1.1, $G_{\mathcal{P}}$ is a complete tripartite graph.) In particular, this means that, if G is a complete multipartite graph, then $\mu(G) = 1$. In the following characterization of a matroidal clique complex, we prove that the converse also holds.

Lemma 1.6. *Let $G = (V, E)$ be a graph. Then the following are equivalent.*

(1) The clique complex of G is a matroid.

(2) The clique complex of G is a partition matroid.

(3) G is complete r-partite for some r.

Note that the equivalence of (1) and (3) in the lemma is noticed by Okamoto [Oka03].

Proof. Since a partition matroid is a matroid, "(2) $\Rightarrow$ (1)" is clear. From the discussion above, "(3) $\Rightarrow$ (2)" is immediate. So we only have to show "(1) $\Rightarrow$ (3)."

Assume that the clique complex $\mathfrak{C}(G)$ is a matroid. By Lemma 1.4, every circuit of $\mathfrak{C}(G)$ is of size two, which corresponds to an edge of $\overline{G}$. Therefore, the elements of each circuit are parallel in $\mathfrak{C}(G)$. By Lemma 1.1, the parallel elements induce an equivalence relation on V, which yields a partition $\mathcal{P} = \{P_1, \dots, P_r\}$ of V for some r. By the construction, this equivalence relation is the same as "x and y are equivalent if and only if there is no edge between x and y in G." Thus, we can see that G is a complete r-partite graph with the vertex partition $\mathcal{P}$. $\square$

For the case of two or more matroids, we use a stable-set partition. A *stable-set partition* of a graph $G = (V, E)$ is a partition $\mathcal{P} = \{P_1, \dots, P_r\}$ of V such that for each $i \in \{1, \dots, r\}$, the set P_i is a stable set of G. (Note that a stable-set partition is nothing but a proper coloring of a graph. However, here we are not interested in how many colors we need (i.e., the size of $\mathcal{P}$) as we do not study the proper coloring problem here.) The following theorem is the main result of this chapter. It tells us how many matroids we need to represent a given clique complex as their intersection.

Theorem 1.7. *Let $G = (V, E)$ be a graph. Then, the following are equivalent.*

(1) The clique complex $\mathfrak{C}(G)$ can be represented as the intersection of k matroids (i.e., $\mu(G) \le k$).

(2) There exist k stable-set partitions $\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(k)}$ of G which fulfill the following condition.

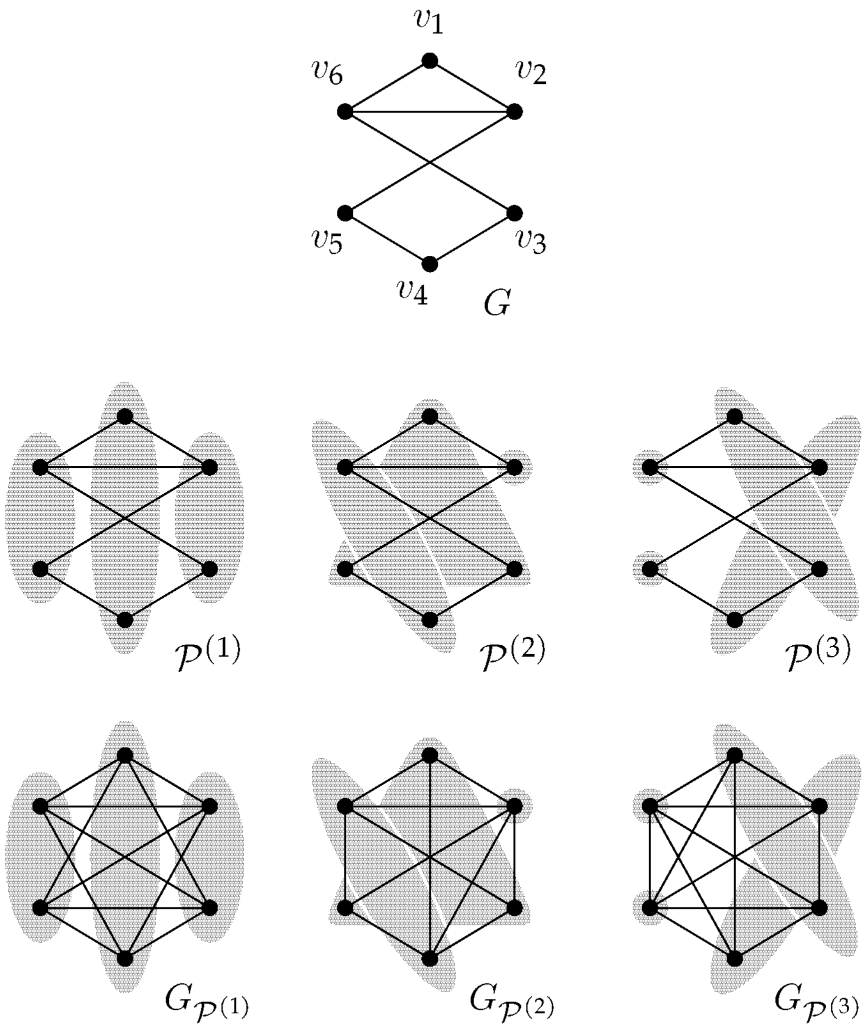

Figure 1.2: An example for Theorem 1.7.

> *Condition P:*
>
> $\{u,v\} \in \binom{V}{2}$ *is an edge of* $\overline{G}$ *if and only if* $\{u,v\} \subseteq S$ *for some* $S \in \bigcup_{i=1}^{k} \mathcal{P}^{(i)}$.

In particular, when Condition P is fulfilled, it holds that

$$\mathfrak{C}(G) = \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)}).$$

Before proving Theorem 1.7, we illustrate the theorem by a pictorial example. Look at Figure 1.2. In the graph $G = (\{v_1, \ldots, v_6\}, E)$, there are seven edges, and

$$\mathcal{P}^{(1)} = \{\{v_1,v_4\},\{v_2,v_3\},\{v_5,v_6\}\},$$
$$\mathcal{P}^{(2)} = \{\{v_1,v_3,v_5\},\{v_2\},\{v_4,v_6\}\},$$
$$\mathcal{P}^{(3)} = \{\{v_1,v_3\},\{v_2,v_4\},\{v_5\},\{v_6\}\}$$

are stable-set partitions of G. We can see that these stable-set partitions meet Condition P, that is, for each $\{u,v\} \in E(\overline{G})$, there exists a stable set $S \in \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)} \cup \mathcal{P}^{(3)}$ such that $\{u,v\} \subseteq S$. For example, look at $\{v_1,v_5\} \in E(\overline{G})$. Then we have a stable set $\{v_1,v_3,v_5\} \in \mathcal{P}^{(2)}$ such that $\{v_1,v_5\} \subseteq \{v_1,v_3,v_5\}$. Indeed, the clique complex $\mathfrak{C}(G)$ can be written as the intersection $\mathcal{I}(\mathcal{P}^{(1)}) \cap \mathcal{I}(\mathcal{P}^{(2)}) \cap \mathcal{I}(\mathcal{P}^{(3)})$ of three partition matroids, or in other words, the intersection $\mathfrak{C}(G_{\mathcal{P}^{(1)}}) \cap \mathfrak{C}(G_{\mathcal{P}^{(2)}}) \cap \mathfrak{C}(G_{\mathcal{P}^{(3)}})$ of the clique complexes of complete multipartite graphs, which are partition matroids (Lemma 1.6).

The intuition behind Condition P in Theorem 1.7 is as follows. Suppose that we consider the clique complex $\mathfrak{C}(G)$ of a given graph G, and we want to gather some complete multipartite graphs $G_1,\ldots,G_k$ so that we can ensure that $\mathfrak{C}(G) = \bigcap_{i=1}^k \mathfrak{C}(G_i)$. Then by Lemma 1.2 it holds that $\mathcal{C}(\mathfrak{C}(G)) = \mathrm{MIN}(\bigcup_{i=1}^k \mathcal{C}(\mathfrak{C}(G)))$. Let us assume that this is equal to $\bigcup_{i=1}^k \mathcal{C}(\mathfrak{C}(G))$, just for the sake of an intuitive discussion. Then it follows that $E(\overline{G}) = \bigcup_{i=1}^k E(\overline{G_i})$ by Lemma 1.4. Therefore, for every edge e of $\overline{G}$ there must be $i \in \{1,\ldots,k\}$ such that e is not an edge of G_i. Actually, Condition P in Theorem 1.7 makes it sure that this requirement is satisfied.

To prove Theorem 1.7, we use the following lemmas.

Lemma 1.8. *Let $G = (V,E)$ be a graph. If the clique complex $\mathfrak{C}(G)$ can be represented as the intersection of k matroids (i.e., $\mu(G) \leq k$), then there exist k stable-set partitions $\mathcal{P}^{(1)},\ldots,\mathcal{P}^{(k)}$ such that $\mathfrak{C}(G) = \bigcap_{i=1}^k \mathcal{I}(\mathcal{P}^{(i)})$.*

Proof. Assume that $\mathfrak{C}(G)$ is represented as the intersection of k matroids $\mathcal{I}_1, \ldots, \mathcal{I}_k$. Choose $j \in \{1,\ldots,k\}$ arbitrarily, and look at $\mathcal{I}_j$.

By Lemma 1.1, the parallel elements of $\mathcal{I}_j$ induce an equivalence relation on V. Let $\mathcal{P}^{(j)}$ be the partition of V arising from this equivalence relation. Then, we can see that the two-element circuits of $\mathcal{I}_j$

are the circuits of the partition matroid $\mathcal{I}(\mathcal{P}^{(j)})$, namely, it holds that $\{C \in \mathcal{C}(\mathcal{I}_j) \mid |C| = 2\} = \mathcal{C}(\mathcal{I}(\mathcal{P}^{(j)}))$. Furthermore, by Lemmas 1.2 and 1.4, it holds that

$$\begin{aligned}
\mathcal{C}(\mathfrak{C}(G)) &= \mathrm{MIN}\left(\bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}_i)\right) \\
&= \mathrm{MIN}\left(\bigcup_{i=1}^{k} \{C \in \mathcal{C}(\mathcal{I}_i) \mid |C| = 2\}\right) \\
&= \mathrm{MIN}\left(\bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))\right).
\end{aligned}$$

Here, the first identity is due to Lemma 1.2. The second one is due to Lemma 1.4 and the assumption that $\mathfrak{C}(G) = \bigcap_{i=1}^{k} \mathcal{C}(\mathcal{I}_i)$. (If we do not have this assumption, there is no assurance for this identity to be true.) The last one is what we observed just above. Now finally, Lemma 1.2 concludes that $\mathfrak{C}(G) = \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)})$. □

Here is another lemma.

Lemma 1.9. *Let $G = (V, E)$ be a graph and $\mathcal{P}$ be a partition of V. Then $\mathfrak{C}(G) \subseteq \mathcal{I}(\mathcal{P})$ if and only if $\mathcal{P}$ is a stable-set partition of G.*

Proof. Assume that $\mathcal{P}$ is a stable-set partition of G. Choose $I \in \mathfrak{C}(G)$ arbitrarily. Then we have that $|I \cap P| \leq 1$ for each $P \in \mathcal{P}$ by the definitions of a clique and a stable set. Hence it follows that $I \in \mathcal{I}(\mathcal{P})$. Thus we have that $\mathfrak{C}(G) \subseteq \mathcal{I}(\mathcal{P})$.

Conversely, assume that $\mathfrak{C}(G) \subseteq \mathcal{I}(\mathcal{P})$ for a partition $\mathcal{P}$ of $V(G)$. Choose $P \in \mathcal{P}$ and a clique $K \in \mathfrak{C}(G)$ of G arbitrarily. From our assumption, we have that $K \in \mathcal{I}(\mathcal{P})$. Therefore, it holds that $|K \cap P| \leq 1$ from the definition of a partition matroid. This means that P is a stable set of G. Hence, $\mathcal{P}$ is a stable-set partition of G. □

Now we are ready for the proof of Theorem 1.7.

Proof of Theorem 1.7. Assume that the clique complex $\mathfrak{C}(G)$ of a given graph $G = (V,E)$ is the intersection of k matroids $\mathcal{I}_1,\ldots,\mathcal{I}_k$. From Lemma 1.8, $\mathfrak{C}(G)$ can be represented as the intersection of k matroids associated with some stable-set partitions $\mathcal{P}^{(1)},\ldots,\mathcal{P}^{(k)}$ of G. We show that these partitions $\mathcal{P}^{(1)},\ldots,\mathcal{P}^{(k)}$ fulfill Condition P. By Lemma 1.4, $\{u,v\}$ is an edge of $\overline{G}$ if and only if $\{u,v\}$ is a circuit of the clique complex $\mathfrak{C}(G)$. Then, it follows that

$$\{u,v\} \in \mathcal{C}(\mathfrak{C}(G)) = \mathrm{MIN}\left(\bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))\right) = \bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)})).$$

(The first identity is due to Lemma 1.2, and the last identity relies on the fact that the size of each circuit of a partition matroid is two.) This implies that there exists at least one index $i \in \{1,\ldots,k\}$ such that $\{u,v\} \in \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))$. Since we have

$$\mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)})) = \left\{\{u,v\} \in \binom{V}{2} \mid \{u,v\} \subseteq S \text{ for some } S \in \mathcal{P}^{(i)}\right\},$$

we can conclude that $\{u,v\} \subseteq S$ for some $S \in \mathcal{P}^{(i)}$ if and only if $\{u,v\}$ is an edge of $\overline{G}$. One direction of the theorem is finished.

Conversely, assume that we are given k stable-set partitions $\mathcal{P}^{(1)},\ldots,\mathcal{P}^{(k)}$ of V satisfying Condition P. We show that

$$\mathfrak{C}(G) = \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)}).$$

By Lemma 1.9, we can see that $\mathfrak{C}(G) \subseteq \mathcal{I}(\mathcal{P}^{(i)})$ for each $i \in \{1,\ldots,k\}$. This implies that $\mathfrak{C}(G) \subseteq \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)})$. What remains to prove is $\mathfrak{C}(G) \supseteq \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)})$. However, in order to show that, we only have to prove that $\mathcal{C}(\mathfrak{C}(G)) \subseteq \bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))$. Why is it true? Assume that $\mathcal{C}(\mathfrak{C}(G)) \subseteq \bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))$, and take $X \in \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)})$ arbitrarily.

Then, we have the following chain of implications.

$$
\begin{aligned}
X &\in \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)}) \\
&\Leftrightarrow \forall i \in \{1,\ldots,k\} : X \in \mathcal{I}(\mathcal{P}^{(i)}) \\
&\Leftrightarrow \forall i \in \{1,\ldots,k\} \forall C \in \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)})) : C \nsubseteq X \\
&\Leftrightarrow \forall C \in \bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)})) : C \nsubseteq X \\
&\Rightarrow \forall C \in \mathcal{C}(\mathfrak{C}(G)) : C \nsubseteq X \\
&\Leftrightarrow X \in \mathfrak{C}(G).
\end{aligned}
$$

(Here, the second and fifth equivalences are due to Equality (1.2); the fourth implication is by our assumption.)

Now, we know that we have to show $\mathcal{C}(\mathfrak{C}(G)) \subseteq \bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))$. Pick $C \in \mathcal{C}(\mathfrak{C}(G))$ arbitrarily. By Lemma 1.4 we can see that C is an edge of $\overline{G}$. Set $\{u,v\} := C \in E(\overline{G})$. From Condition P, there exists some $S \in \bigcup_{i=1}^{k} \mathcal{P}^{(i)}$ such that $\{u,v\} \subseteq S$. This means that $\{u,v\} \in \bigcup_{i=1}^{k} \mathcal{C}(\mathcal{I}(\mathcal{P}^{(i)}))$. Thus the proof is completed. □

Next, let us look at some consequences of the discussion in this section. First of all, Theorem 1.7 implies that the clique complex $\mathfrak{C}(G)$ of a graph G can be represented as the intersection of k matroids if and only if $\mathfrak{C}(G)$ can be represented as the intersection of k *partition matroids arising from stable-set partitions of* G. Therefore, if you want to find $\mu(G)$, you only have to search within the partition matroids arising from stable-set partitions of G. This considerably reduces the time/cost of the search.

In Lemma 1.8, we showed that, for a given graph G on the vertex set V whose clique complex $\mathfrak{C}(G)$ is the intersection of k matroids, we can find k partition matroids whose intersection is $\mathfrak{C}(G)$. Moreover, we can show the following "converse" statement.

Corollary 1.10. *For any collection of k partitions $\mathcal{P}^{(1)},\mathcal{P}^{(2)},\ldots,\mathcal{P}^{(k)}$ of a finite set V, there exists a graph G on V such that $\mathfrak{C}(G)$ is the intersection of the partition matroids $\mathcal{I}(\mathcal{P}^{(1)}),\mathcal{I}(\mathcal{P}^{(2)}),\ldots,\mathcal{I}(\mathcal{P}^{(k)})$.*

Proof. From a given collection of partitions $\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(k)}$ of V, we construct a graph G as follows. The vertex set of G is V. Two vertices u and v are connected by an edge in G if and only if they do not lie in a common class of $\mathcal{P}^{(i)}$ for any $i \in \{1, \dots, k\}$ (i.e., there exists no $S \in \mathcal{P}^{(i)}$ such that $\{u, v\} \subseteq S$ for any $i \in \{1, \dots, k\}$). Then we can see that $\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(k)}$ are stable-set partitions of G. Moreover, they satisfy Condition P in the statement of Theorem 1.7. Therefore, by Theorem 1.7, we can conclude that $\mathfrak{C}(G) = \bigcap_{i=1}^{k} \mathcal{I}(\mathcal{P}^{(i)})$. □

This leads to the following important consequence, which characterizes the clique complexes as the intersections of partition matroids.

Corollary 1.11. *For every $k > 0$, the class of clique complexes which are the intersections of k matroids is the same as the class of the intersections of k partition matroids; in particular, the class of clique complexes is the same as the class of the intersections of partition matroids.*

Proof. Combine Lemma 1.8 and Corollary 1.10. □

At the end of this section, we would like to notice that Theorem 1.7 implies that the following decision problem belongs to NP.

Problem:	CLIQUE COMPLEX k-MATROID REPRESENTATION
Instance:	a graph G and a positive integer k
Question:	Is $\mu(G) \leq k$?

Let us state this fact as a corollary.

Corollary 1.12. CLIQUE COMPLEX k-MATROID REPRESENTATION *belongs to* NP.

Note that this corollary is not trivial since a matroid itself can have an exponential number of independent sets.

Proof. By Theorem 1.7, k stable-set partitions satisfying Condition P is a certificate for the positive answer to the decision problem above. Since the size of stable-set partition is polynomial in the size of a graph

G and k is at most the number of vertices in G, these k stable-set partitions constitute a polynomial-size certificate. Furthermore, Condition P can be checked in polynomial time for a given graph and given k stable-set partitions of the graph. This concludes that the decision problem CLIQUE COMPLEX k-MATROID REPRESENTATION belongs to NP. □

However, we do not know that CLIQUE COMPLEX k-MATROID REPRESENTATION belongs to P, or even to coNP. It could be NP-complete. When k is fixed, the situation is somehow changed. For $k = 1$, due to Lemma 1.6 the problem can be solved in polynomial time. The case $k = 2$ is discussed in Section 1.5, and we prove that in this case the problem can also be solved in polynomial time.

1.4 An Extremal Problem for Clique Complexes

Remember that $\mu(G)$ is the minimum number of matroids needed for the representation of the clique complex of G as their intersection. Furthermore, let $\mu(n)$ be the maximum of $\mu(G)$ over all graphs G with n vertices. Namely,

$$\mu(n) := \max\{\mu(G) \mid G \text{ has } n \text{ vertices}\}.$$

In this section, we determine $\mu(n)$ exactly. It is straightforward to observe that $\mu(1) = 1$. For the case of $n \geq 2$, we can immediately obtain $\mu(n) \leq \binom{n}{2}$ from Lemmas 1.3 and 1.4. However, the following theorem tells us that the truth is much better.

Theorem 1.13. *For every $n \geq 2$, it holds that $\mu(n) = n - 1$.*

First, we prove that $\mu(n) \geq n - 1$. Consider the graph $K_1 \cup K_{n-1}$. (Figure 1.3 shows $K_1 \cup K_5$.)

Lemma 1.14. *For $n \geq 2$, it holds that $\mu(K_1 \cup K_{n-1}) = n - 1$. Particularly it follows that $\mu(n) \geq n - 1$.*

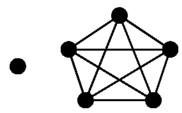

$K_1 \cup K_5$

Figure 1.3: The graph $K_1 \cup K_5$.

Proof. First, observe that $\overline{K_1 \cup K_{n-1}}$ has $n-1$ edges. Therefore, Lemma 1.4 implies that the number of the circuits of $\mathfrak{C}(K_1 \cup K_{n-1})$ is $n-1$. Then, by the argument below the proof of Lemma 1.3, it follows that $\mu(K_1 \cup K_{n-1}) \leq |\mathcal{C}(\mathfrak{C}(K_1 \cup K_{n-1}))| = n-1$.

Now, suppose that $\mu(K_1 \cup K_{n-1}) \leq n-2$. By Theorem 1.7, there exist at most $n-2$ stable-set partitions $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(n-2)}$ of $K_1 \cup K_{n-1}$ satisfying Condition P, namely, each edge e of $\overline{K_1 \cup K_{n-1}}$ is contained in some set $S \in \bigcup_{i=1}^{n-2} \mathcal{P}^{(i)}$. Then, the pigeon hole principle tells us that there exists an index $i^* \in \{1, \ldots, n-2\}$ such that at least two edges of $\overline{K_1 \cup K_{n-1}}$ are contained in sets of $\mathcal{P}^{(i^*)}$. Let e, e' be such (distinct) edges of $\overline{K_1 \cup K_{n-1}}$ and $P_e, P_{e'} \in \mathcal{P}^{(i^*)}$ be unique sets such that $e \subseteq P_e$ and $e' \subseteq P_{e'}$. (The uniqueness follows from the fact that $\mathcal{P}^{(i^*)}$ is a partition.) Remember that e and e' share a vertex (since e, e' are edges of $\overline{K_1 \cup K_{n-1}}$). This implies that $P_e \cap P_{e'} \neq \emptyset$. Therefore, it holds that $P_e = P_{e'}$ since $\mathcal{P}^{(i^*)}$ is a partition. Set $e = \{u, v\}$ and $e' = \{u, v'\}$. We can find $\{v, v'\}$ is also contained in P_e. However, $\{v, v'\}$ is an edge of $K_1 \cup K_{n-1}$. This contradicts the fact that $\mathcal{P}^{(i^*)}$ is a stable-set partition (i.e., P_e is a stable set of $K_1 \cup K_{n-1}$). Thus, it follows that $\mu(K_1 \cup K_{n-1}) = n-1$.

For the second part, we just follow the definition of $\mu(n)$. Then we can find that $\mu(n) \geq \mu(K_1 \cup K_{n-1}) = n-1$. ☐

Next we prove that $\mu(n) \leq n-1$. To do that, first we look at the relation of $\mu(G)$ with the edge-chromatic number $\chi'(\overline{G})$ of the complement.

Lemma 1.15. *It holds that $\mu(G) \leq \chi'(\overline{G})$ for every graph G. Particularly, if the number n of vertices of G is even then we have that $\mu(G) \leq n-1$, and if n is odd then we have that $\mu(G) \leq n$. Moreover, if $\mu(G) = n$ then n is odd and the maximum degree of $\overline{G}$ is $n-1$ (i.e., G has an isolated vertex).*

Proof. Consider a minimum proper edge-coloring of $\overline{G}$, and let $k :=$

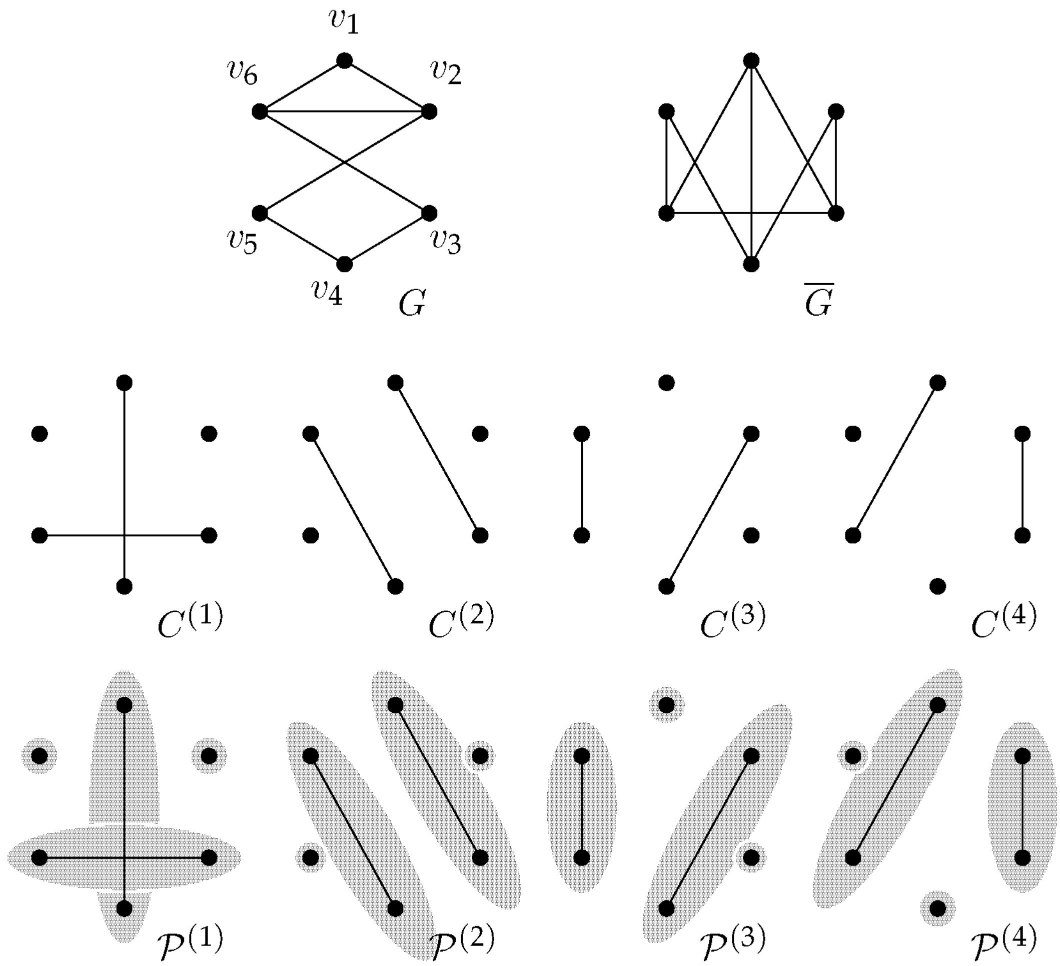

Figure 1.4: The construction in the proof of Lemma 1.15.

$\chi'(\overline{G})$. We construct k stable-set partitions of a graph G with n vertices from this edge-coloring.

We have the color classes $C^{(1)}, \ldots, C^{(k)}$ of the edges from the minimum proper edge-coloring. (Figure 1.4 is an illustration. In the example of the figure, we have $\chi'(\overline{G}) = 4$. The first row shows a given graph G and its complement $\overline{G}$. In the second row, we can find a minimum proper edge-coloring of $\overline{G}$, and each $C^{(i)}$ depicts a color class of this coloring.)

Take a color class $C^{(i)} = \{e_1^{(i)}, \ldots, e_{\ell_i}^{(i)}\}$ ($i \in \{1, \ldots, k\}$) and construct a stable-set partition $\mathcal{P}^{(i)}$ of G from $C^{(i)}$ as follows: S is a member of $\mathcal{P}^{(i)}$ if and only if either

(1) S is a two-element set belonging to $C^{(i)}$ (i.e., $S = e_j^{(i)}$ for some $j \in \{1, \ldots, \ell_i\}$) or

(2) S is a one-element set $\{v\}$ which is not used in $C^{(i)}$ (i.e., $v \notin e_j^{(i)}$ for any $j \in \{1, \ldots, \ell_i\}$).

Notice that $\mathcal{P}^{(i)}$ is actually a stable-set partition of G. Then we collect all the stable-set partitions $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(k)}$ constructed by the procedure above. Moreover, these stable-set partitions satisfy Condition P in Theorem 1.7 since each edge of $\overline{G}$ appears in exactly one of the $C^{(i)}$'s. Hence, by Theorem 1.7, it follows that $\mu(G) \leq k = \chi'(\overline{G})$. In Figure 1.4, the constructed stable-set partitions are put in the third row. Thus we have shown that $\mu(G) \leq \chi'(\overline{G})$.

Notice that $\chi'(\overline{G}) \leq \chi'(K_n)$ for any graph G with n vertices. Therefore, if n is even, then we can conclude that $\mu(G) \leq n-1$ since $\chi'(K_n) = n-1$. Similarly, if n is odd, then we can conclude that $\mu(G) \leq n$ since $\chi'(K_n)$ is n.

For the last part of the lemma, assume that $\mu(G) = n$. From the discussion above, n should be odd. Now, let us remind Vizing's theorem [Viz64, Viz65], which states that for a (simple) graph H with maximum degree $\Delta(H)$ we have that $\chi'(H) = \Delta(H)$ or $\Delta(H) + 1$. Since $\Delta(\overline{G}) \leq n-1$, it follows that

$$n = \mu(G) \leq \chi'(\overline{G}) \leq \Delta(\overline{G}) + 1 \leq n.$$

Therefore, $\mu(G) = n$ holds only if $\Delta(\overline{G}) + 1 = n$. $\square$

Finally, we show that if a graph G with n vertices, n odd, has an isolated vertex then $\mu(G) \leq n-1$. This completes the proof of Theorem 1.13.

Lemma 1.16. *Let n be odd and G be a graph with n vertices which has an isolated vertex. Then it holds that $\mu(G) \leq n-1$.*

Proof. Let v' be an isolated vertex of G. Consider the subgraph of G induced by $V(G) \setminus \{v'\}$. Denote this induced subgraph by G' (i.e., $G' = G[V(G) \setminus \{v'\}]$). Since G' has $n-1$ vertices and this is even, we have $\mu(G') \leq n-2$ from Lemma 1.15.

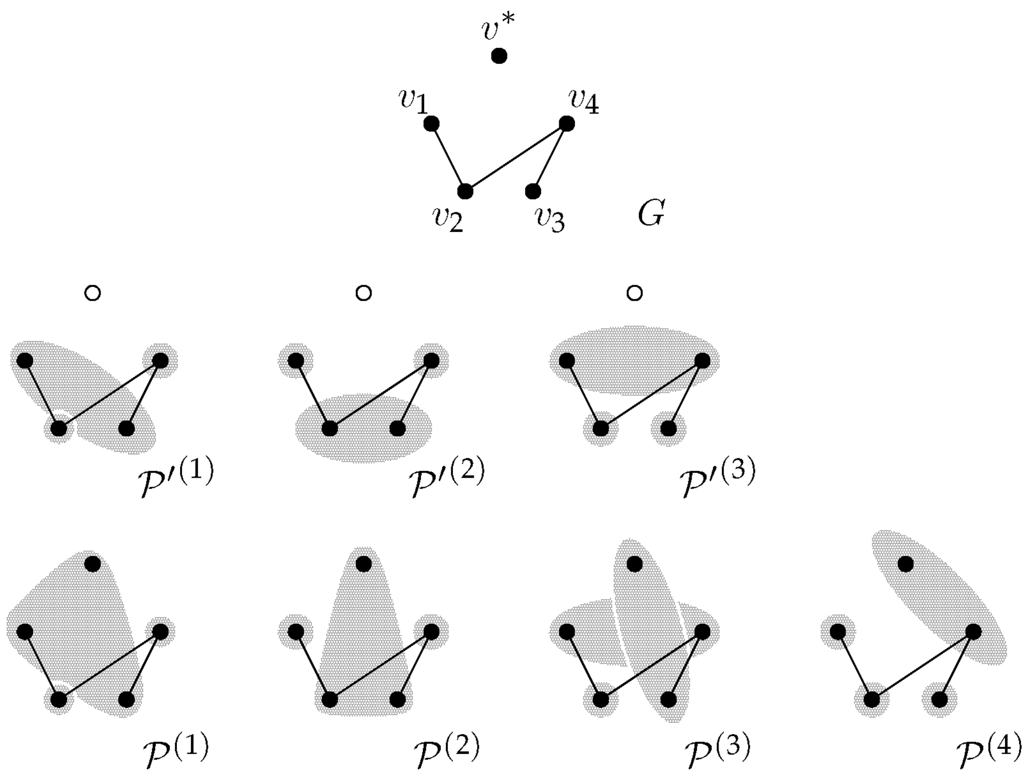

Figure 1.5: The construction of stable-set partitions from an edge-coloring.

Now we construct $n-1$ stable-set partitions of G satisfying Condition P from $n-2$ stable-set partitions of G' satisfying Condition P as well. Denote the vertices of G' by $v_1, \ldots, v_{n-1}$, and stable-set partitions of G' satisfying Condition P by $\mathcal{P}'^{(1)}, \ldots, \mathcal{P}'^{(n-2)}$ (where some of them may be identical in case $\mu(G') < n-2$). Then construct stable-set partitions $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(n-2)}, \mathcal{P}^{(n-1)}$ of G as follows. For each $i \in \{1, \ldots, n-2\}$, put $P \in \mathcal{P}^{(i)}$ if and only if either

(1) $P \in \mathcal{P}'^{(i)}$ and $v_i \notin P$ or

(2) $v' \in P$, $P \setminus \{v'\} \in \mathcal{P}'^{(i)}$ and $v_i \in P$.

Furthermore, put $P \in \mathcal{P}^{(n-1)}$ if and only if either

(1) $P = \{v_i\}$ $(i \in \{1, \ldots, n-2\})$ or

(2) $P = \{v', v_{n-1}\}$.

Figure 1.5 illustrates the construction of $\mathcal{P}^{(i)}$ $(i \in \{1, \ldots, n-1\})$. The first row shows a given graph G where the topmost vertex v' is isolated. In the second row, we can find three stable-set partitions of $G' = G[\{v_1, v_2, v_3, v_4\}]$ satisfying Condition P. In this row, the symbol $\circ$ is used to indicate the neglected vertex v'. In the third row (lowest), the constructed stable-set partitions of G are shown according to the considered vertices.

For conclusion, it is enough to check that the stable-set partitions $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(n-1)}$ constructed above satisfy Condition P. Choose any edge e of $\overline{G}$. If e is also an edge of $\overline{G'}$, then we can find a set $S' \in \bigcup_{i=1}^{n-2} \mathcal{P}'^{(i)}$ such that $e \subseteq S'$ since $\mathcal{P}'^{(1)}, \ldots, \mathcal{P}'^{(n-2)}$ satisfy Condition P. From the construction of $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(n-2)}$, we observe that for each $i \in \{1, \ldots, n-2\}$ and each $P' \in \mathcal{P}'^{(i)}$ there exists a set $P \in \mathcal{P}^{(i)}$ such that $P' \subseteq P$. Therefore, for S' above, we also have $S \in \bigcup_{i=1}^{n-2} \mathcal{P}^{(i)}$ such that $S' \subseteq S$, which implies that $e \subseteq S$. If e is not an edge of $\overline{G'}$, then e is expressed as $e = \{v', v_i\}$ for some $i \in \{1, \ldots, n-1\}$. Then it turns out that e is contained in a member of $\mathcal{P}^{(i)}$ which was put in $\mathcal{P}^{(i)}$ by the second alternative. In this way, we have verified that $\mathcal{P}^{(1)}, \ldots, \mathcal{P}^{(n-1)}$ satisfy Condition P. $\square$

1.5 Characterizations for Two Matroids

In this section, we look more closely at clique complexes that can be represented as the intersections of two matroids.

To do this, we invoke another concept. The *stable-set graph* of a graph $G = (V, E)$ is a graph whose vertices are the maximal stable sets of G and two vertices of which are adjacent if the corresponding two maximal stable sets of G share a vertex in G. We denote the stable-set graph of a graph G by $\mathcal{S}(G)$. Figure 1.6 shows an example of a stable-set graph.

The next lemma establishes the relationship between $\mu(G)$ and the chromatic number $\chi(\mathcal{S}(G))$ of the stable-set graph.

Lemma 1.17. *Let G be a graph and k be a natural number. Then the clique complex $\mathfrak{C}(G)$ can be represented as the intersection of k matroids if the stable-set graph $\mathcal{S}(G)$ is k-colorable. In other words, it holds that $\mu(G) \leq \chi(\mathcal{S}(G))$.*

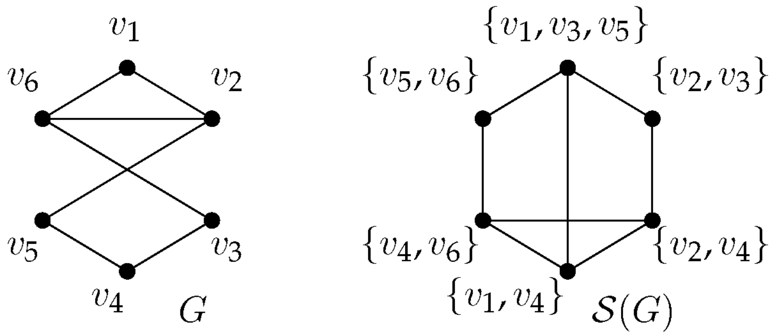

Figure 1.6: A stable-set graph.

Proof. Assume that we are given a proper k-coloring c of $\mathcal{S}(G)$, i.e., a map $c : V(\mathcal{S}(G)) \to \{1,\dots,k\}$ where $c(S) \neq c(T)$ whenever $S \cap T \neq \emptyset$. Then gather the maximal stable sets of G which have the same color with respect to the coloring c, that is, put $C_i := \{S \in V(\mathcal{S}(G)) \mid c(S) = i\}$ for each $i \in \{1,\dots,k\}$. By construction, the members of C_i are disjoint maximal stable sets of G for each $i \in \{1,\dots,k\}$.

Now we construct a graph G_i from C_i as follows. The vertex set of G_i is the same as that of G, and two vertices of G_i are adjacent if and only if either

(1) one belongs to a maximal stable set in C_i and the other belongs to another maximal stable set in C_i, or

(2) one belongs to a maximal stable set in C_i and the other belongs to no maximal stable set in C_i.

Figure 1.7 explains the construction of G_i by the example in Figure 1.6. In that figure, three colors of $\mathcal{S}(G)$ are depicted by $\bullet$, $\blacksquare$ and $\circ$, and in the second row, the shaded groups show maximal cliques corresponding to the vertices in $\mathcal{S}(G)$ colored by identical colors.

Note that G_i is complete r-partite, where r is equal to $|C_i|$ plus the number of the vertices which do not belong to any maximal stable set in C_i. (This holds in general, not just in the picture above.) Then consider $\mathfrak{C}(G_i)$, the clique complex of G_i. By Lemma 1.6, we can see that $\mathfrak{C}(G_i)$ is actually a matroid. Since an edge of G is also an edge of G_i (or by Lemma 1.9), we have that $\mathfrak{C}(G) \subseteq \mathfrak{C}(G_i)$.

Now we consider the intersection $\mathcal{I} = \bigcap_{i=1}^{k} \mathfrak{C}(G_i)$. Since $\mathfrak{C}(G) \subseteq \mathfrak{C}(G_i)$ for every $i \in \{1,\dots,k\}$, we have $\mathfrak{C}(G) \subseteq \mathcal{I}$. Since each circuit of

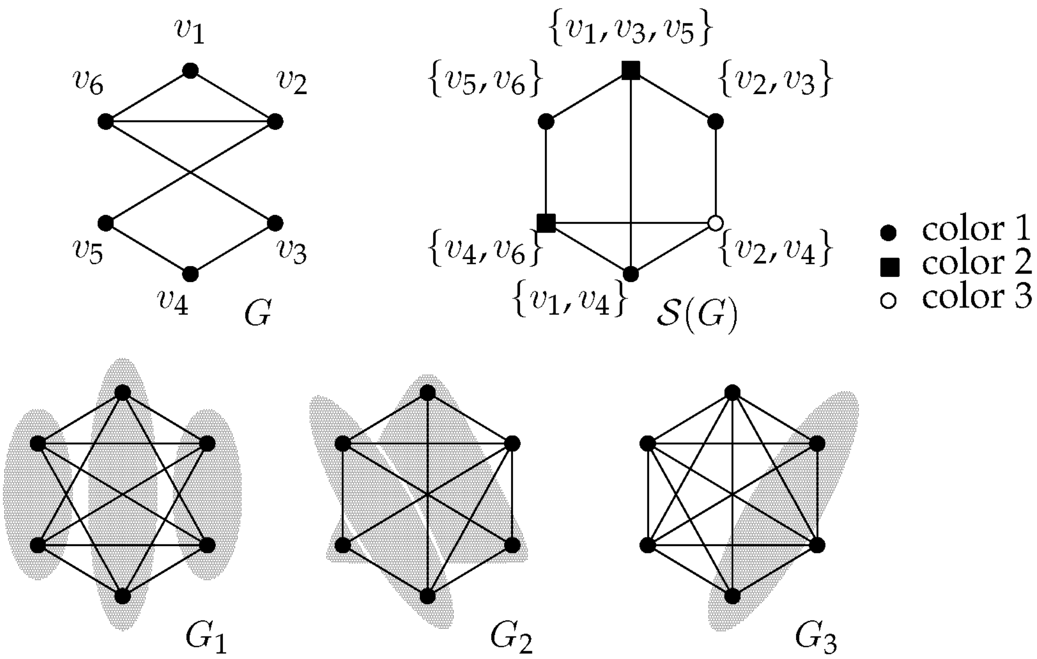

Figure 1.7: The construction of G_i in the proof of Lemma 1.17.

$\mathfrak{C}(G)$ is also a circuit of $\mathfrak{C}(G_i)$ for some $i \in \{1, \ldots, k\}$ (recall Lemma 1.4), we also have $\mathcal{C}(\mathfrak{C}(G)) \subseteq \mathcal{C}(\mathcal{I})$, which implies $\mathfrak{C}(G) \supseteq \mathcal{I}$. Thus we have $\mathfrak{C}(G) = \mathcal{I}$. □

Let us note that the inequality $\mu(G) \leq \chi(\mathcal{S}(G))$ is tight. Indeed, a tight example has already appeared in Figure 1.6. In that example, it holds $\chi(\mathcal{S}(G)) = 3$. On the other hand, by Theorem 1.18 below, we can see that $\mu(G) \geq 3$ since $\mathcal{S}(G)$ is not 2-colorable. Figure 1.7 shows three stable-set partitions satisfying Condition P. So, by Theorem 1.7 we conclude that $\mu(G) = 3$.

Furthermore, note that the converse of Lemma 1.17 does not hold in general even if $k = 3$. A counterexample is the graph $G = (V, E)$ defined as

$$V = \{v_1, v_2, v_3, v_4, v_5, v_6\} \text{ and}$$
$$E = \{\{v_1, v_2\}, \{v_3, v_4\}, \{v_5, v_6\}\}.$$

See Figure 1.8. In the graph shown in Figure 1.8, consider the following

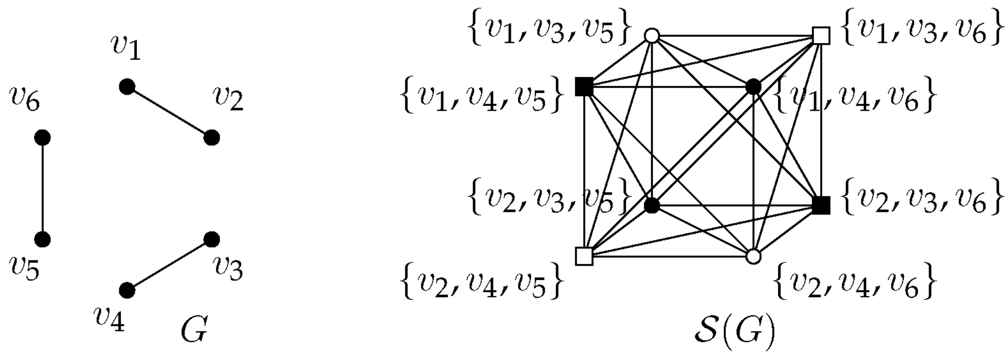

Figure 1.8: A counterexample for the converse of Lemma 1.17.

stable-set partitions of G:

$$\mathcal{P}^{(1)} = \{\{v_1, v_3, v_5\}, \{v_2, v_4, v_6\}\},$$
$$\mathcal{P}^{(2)} = \{\{v_1, v_3, v_6\}, \{v_2, v_4, v_5\}\},$$
$$\mathcal{P}^{(3)} = \{\{v_1, v_4, v_5\}, \{v_2, v_3, v_6\}\}.$$

We can check that these stable-set partitions fulfill Condition P in Theorem 1.7. Therefore, by Theorem 1.7, we can see that $\mathfrak{C}(G)$ is the intersection of three partition matroids $\mathcal{I}(\mathcal{P}^{(1)}), \mathcal{I}(\mathcal{P}^{(2)})$ and $\mathcal{I}(\mathcal{P}^{(3)})$. However, $\mathcal{S}(G)$ is not 3-colorable but 4-colorable. (Since $\mathcal{S}(G)$ has a clique of size four, it is not 3-colorable. On the other hand, we can color $\mathcal{S}(G)$ with four colors as indicated in Figure 1.8.)

By a similar argument, we can also see that, if we consider a graph consisting of only $n/2$ independent edges (i.e. a graph which is a matching), then the difference between $\mu(G)$ and $\chi(\mathcal{S}(G))$ can be arbitrarily large.

However, the converse holds if $k = 2$.

Theorem 1.18. *Let G be a graph. The clique complex $\mathfrak{C}(G)$ can be represented as the intersection of two matroids if and only if the stable-set graph $\mathcal{S}(G)$ is 2-colorable (i.e., bipartite).*

Proof. The if-part is straightforward from Lemma 1.17. Now we prove the only-if-part. Assume that $\mathfrak{C}(G)$ is represented as the intersection of two matroids. Due to Theorem 1.7, we may assume that these two matroids are associated with stable-set partitions $\mathcal{P}^{(1)}, \mathcal{P}^{(2)}$ of G satisfying Condition P.

Let S be an arbitrary maximal stable set of G. Now we claim the following.

Claim 1.18.1. *It holds that $S \in \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)}$. Namely, every maximal stable set of G is contained in $\mathcal{P}^{(1)}$ or $\mathcal{P}^{(2)}$.*

Proof of Claim 1.18.1. To prove this claim, from the maximality of S, we only have to show that $S \subseteq P$ for some $P \in \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)}$. (Then, the maximality of S tells us that $S = P$.) Since $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$ are partitions of $V(G)$, this claim clearly holds if $|S| = 1$. If $|S| = 2$, the claim holds from Condition P.

Assume that $|S| \geq 3$. Then consider the following independence system on S:

$$\mathcal{I} := \{I \subseteq S \mid I \subseteq P \text{ for some } P \in \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)}\}.$$

Choose a base B (i.e., a maximal independent set) of $\mathcal{I}$ arbitrarily. If $B = S$ holds, then we are done (since $B \subseteq P$ for some $P \in \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)}$). Since $B \subseteq S$, it suffices to show that $B \supseteq S$.

Now, suppose that $S \setminus B \neq \emptyset$ for a contradiction. Without loss of generality, we may assume that B is contained in a set P of $\mathcal{P}^{(1)}$. If there exists an element $u \in (S \setminus B) \cap P$, then it follows that $\{u\} \cup B \subseteq P$. This contradicts the maximality of B in $\mathcal{I}$. Therefore, $(S \setminus B) \cap P = \emptyset$; in other words $B = S \cap P$.

Since $B \subseteq S$ holds and S is a stable set of G, every two-element subset of S is a circuit of $\mathfrak{C}(G)$ (i.e., an edge of $\overline{G}$). Fix $u \in S \setminus B$ arbitrarily and also choose $v \in B$ arbitrarily. Then we have $\{u, v\} \subseteq S$ is an edge of $\overline{G}$. Therefore, by Condition P there must exist $Q \in \mathcal{P}^{(1)} \cup \mathcal{P}^{(2)}$ containing $\{u, v\}$. Since $B \subseteq P$, we have $v \in P$. If Q would be a member of $\mathcal{P}^{(1)}$, then it holds that $v \in P \cap Q$, which contradicts the fact that $\mathcal{P}^{(1)}$ is a partition. Therefore, Q must be a member of $\mathcal{P}^{(2)}$. Since Q contains u and $\mathcal{P}^{(2)}$ is a partition of $V(G)$, we can see that Q must be a (unique) set of $\mathcal{P}^{(2)}$ which contains u. Since the choice of v was arbitrary, this set Q must contain *all* $v \in B$, which implies that $\{u\} \cup B$ is contained in Q. However, we may conclude that $\{u\} \cup B \in \mathcal{I}$ by the definition of $\mathcal{I}$ and the fact that $Q_v \in \mathcal{P}^{(2)}$. This contradicts the maximality of B. The claim has been verified. $\square$

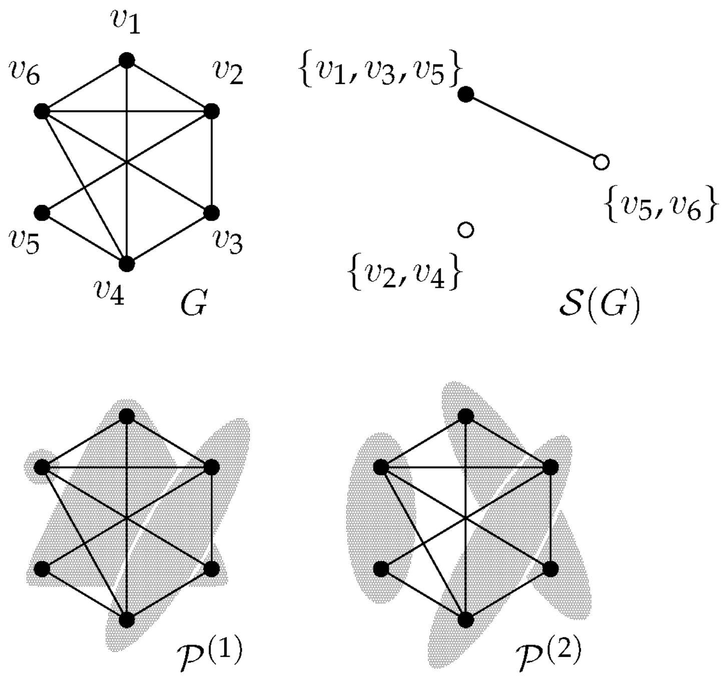

Figure 1.9: An illustration of the proof of Theorem 1.18.

Coming back to the proof of Theorem 1.18, we now color the vertices of $\mathcal{S}(G)$, i.e., the maximal stable sets of G, according to $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$. If a maximal stable set S belongs to $\mathcal{P}^{(1)}$, then S is colored by 1. Similarly, if S belongs to $\mathcal{P}^{(2)}$, then S is colored by 2. (If S belongs to both, then S can be colored by either 1 or 2 arbitrarily.) By the claim above, this procedure can color all vertices of $\mathcal{S}(G)$, and furthermore this coloring is certainly a proper 2-coloring of $\mathcal{S}(G)$ since $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$ are partitions of $V(G)$. $\square$

Figure 1.9 is an illustration of what we saw in the proof. The graph G in Figure 1.9 has three maximal stable sets, and they form the vertex set of the stable-set graph $\mathcal{S}(G)$. In the second row, we can see two stable-set partitions satisfying Condition P. According to these stable-set partitions, we can color the vertices in $\mathcal{S}(G)$. In this example, $\{v_1, v_3, v_5\}$ is colored by $\bullet$ (color 1) since $\{v_1, v_3, v_5\}$ appears in $\mathcal{P}^{(1)}$, and $\{v_5, v_6\}$ is colored by $\circ$ (color 2) since $\{v_5, v_6\}$ appears in $\mathcal{P}^{(2)}$. Then, $\{v_2, v_4\}$ appears in both of $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$. Therefore we can color it by either $\bullet$ or $\circ$ arbitrarily. In the picture above, we just chose $\circ$ by chance.

Some researchers already noticed that the bipartiteness of $\mathcal{S}(G)$ is

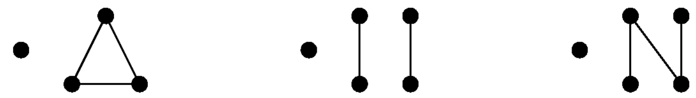

$K_1 \cup K_3$ $\quad$ $K_1 \cup K_2 \cup K_2$ $\quad$ $K_1 \cup P_3$

Figure 1.10: The forbidden induced subgraphs for Proposition 1.19.

characterized by other properties. We gather them in the following proposition. Here, the *line graph* of a multigraph G is a graph $L(G)$ such that the vertex set of $L(G)$ is the edge set of G and two vertices in $L(G)$ are adjacent through an edge if and only if the corresponding two edges in G share a vertex in G.

Proposition 1.19. *Let G be a graph. Then the following are equivalent.*

(1) The stable-set graph $\mathcal{S}(G)$ is bipartite.

(2) G is the complement of the line graph of a bipartite multigraph.

(3) G has no induced subgraph isomorphic to

$$K_1 \cup K_3, K_1 \cup K_2 \cup K_2, K_1 \cup P_3 \text{ or } \overline{C_{2k+3}} \ (k = 1, 2, \ldots),$$

where C_{2k+3} denotes a cycle of length $2k + 3$. See Figure 1.10.

Proof. "(1) $\Leftrightarrow$ (2)" is immediate from a result by Cai, Corneil & Proskurowski [CCP96]. Also, "(1) $\Leftrightarrow$ (3)" is immediate from a result by Protti & Szwarcfiter [PS02]. ☐

Notice that we can decide whether the stable-set graph of a graph is bipartite or not in polynomial time using the algorithm described by Protti & Szwarcfiter [PS02]. Here, we briefly describe their algorithm. To sketch their algorithm, first we have to observe that if $\mathcal{S}(G)$ is bipartite then G contains at most $2n$ maximal stable sets. (This is not trivial. For a proof, see the original paper [PS02].) Using this observation, they provided the following algorithm. At the first step, we list the maximal stable sets of G using an algorithm with polynomial delay by Tsukiyama, Ide, Ariyoshi & Shirakawa [TIAS77], for example. If the algorithm starts to generate more than $2n$ maximal stable sets then we stop the whole execution and answer "No" (since $\mathcal{S}(G)$ cannot be bipartite from the observation above). If it generates at most $2n$ maximal stable sets and halts, then we proceed to the second step. At the second step, we explicitly construct $\mathcal{S}(G)$, which can be done in polynomial time since the number of vertices of $\mathcal{S}(G)$ is at most $2n$. Then, as the third step, we check that $\mathcal{S}(G)$ is bipartite or not, which can also be done in polynomial time. If it is bipartite then answer "Yes," otherwise "No." In total, this procedure runs in polynomial time.

For the class of graphs satisfying a condition in Proposition 1.19 we can solve the maximum weight clique problem exactly in polynomial time by Frank's algorithm [Fra81] for the maximum weight base problem in the intersection of two matroids. Notice that in Frank's algorithm we need to know what the two matroids are. (More precisely speaking, we are required to have a polynomial-time algorithm to decide whether a given set is an independent set of each of the two matroids.) However, the above algorithm by Protti & Szwarcfiter [PS02] explicitly gives a proper 2-coloring of the stable-set graph if the answer is "Yes." Hence, from the argument in the proof of Theorem 1.18 we can find the corresponding stable-set partitions of the graph, which are sufficient for Frank's algorithm.

You may wonder about the intersection of three matroids. As for the recognition problem, we do not know so far that the problem to decide whether the clique complex of a given graph is the intersection of three matroids or not can be solved in polynomial time, or is NP-complete. We leave this question as an open problem. As for optimization, the problem to find a maximum weight clique in a graph whose clique complex is the intersection of three matroids turns out to be NP-hard, even for the unweighted case. Here, we want to describe the reason briefly. In Corollary 1.11, we mentioned that the class of clique complexes which are the intersections of three matroids is the same as the class of the intersections of three partition matroids. Therefore, our problem is exactly to find a maximum weight base in the intersection of three partition matroids. However this problem contains the maximum 3-dimensional matching problem as a special case, which is known to be NP-hard [GJ79] (and even MAX-SNP-hard [Kan91]). This implies that our problem is intractable for three matroids.

1.6 Graphs as Independence Systems and the Intersection of Matroids

We can regard a graph as an independence system such that a subset of the vertex set is independent if and only if it is either

(1) the empty set,

(2) a vertex of the graph or

(3) an edge of the graph.

In this section we consider how many matroids we need to represent a graph (viewed as an independence system) by their intersection. First, we establish a lemma on the matroidal case.

Lemma 1.20. *Let G be a graph. Then the following are equivalent.*

(1) G is a matroid.

(2) $\mathfrak{C}(G)$ is a matroid.

(3) G is complete r-partite for some r.

For the proof, we need the concept of truncation. Let $\mathcal{I}$ be an independence system on V. For $k \geq 0$, the *k-th truncation* of $\mathcal{I}$ is the subfamily $\mathcal{I}^{\leq k}$ of $\mathcal{I}$ defined as

$$\mathcal{I}^{\leq k} := \{X \in \mathcal{I} \mid |X| \leq k\}.$$

We can see that the truncation of an independence system $\mathcal{I}$ is also an independence system, and if $\mathcal{I}$ is a matroid then $\mathcal{I}^{\leq k}$ is also a matroid for every $k \geq 0$. Note that the k-th truncation is also called the $(k-1)$-*dimensional skeleton* especially when we study "simplicial complexes" instead of "independence systems."

Proof of Lemma 1.20. "(2) ⇔ (3)" is nothing else Lemma 1.6. "(2) ⇒ (1)" is immediate from the two facts that G is the 2-truncation of $\mathfrak{C}(G)$ and that the truncation of a matroid is also a matroid. Now we prove "(1) ⇒ (3)." Suppose that G is not complete r-partite for any r. Then, G has three vertices u, v, w such that $\{u, v\}$ is an edge but neither $\{u, w\}$ nor $\{v, w\}$ is an edge of G. However, this contradicts the augmentation axiom of a matroid. □

The following theorem states that the minimum number of matroids for a graph equals that for the clique complex of this graph.

Theorem 1.21. *Let G be a graph and k be a natural number. Then G can be represented as the intersection of k matroids if and only if the clique complex $\mathfrak{C}(G)$ can be represented as the intersection of k matroids.*

Proof. First, we show the if-part. Let $\mathfrak{C}(G)$ be represented as the intersection of k matroids $\mathcal{I}_1,\dots,\mathcal{I}_k$, i.e., $\mathfrak{C}(G) = \bigcap_{i=1}^k \mathcal{I}_i$. Due to Theorem 1.7, without loss of generality, we may assume that $\mathcal{I}_i$ is a partition matroid for each $i \in \{1,\dots,k\}$. Then consider the truncations $\mathcal{I}_1^{\leq 2},\dots,\mathcal{I}_k^{\leq 2}$, and observe that $\bigcap_{i=1}^k \mathcal{I}_i^{\leq 2} = (\bigcap_{i=1}^k \mathcal{I}_i)^{\leq 2}$. On the other hand, it follows that $G = \mathfrak{C}(G)^{\leq 2} = (\bigcap_{i=1}^k \mathcal{I}_i)^{\leq 2}$ by our assumption. Thus we conclude that $G = (\bigcap_{i=1}^k \mathcal{I}_i)^{\leq 2}$, namely G is the intersection of k matroids.

Next we show the only-if-part. Let G be represented as the intersection of k matroids $\mathcal{J}_1,\dots,\mathcal{J}_k$, namely $G = \bigcap_{i=1}^k \mathcal{J}_i$. Without loss of generality, we may assume that the size of every base of $\mathcal{J}_i$ is at most two for each $i \in \{1,\dots,k\}$. (If not, then consider the truncation $\mathcal{J}_i^{\leq 2}$, which does not change the intersection that we are considering since the size of every base in G is at most two.) Then we can regard $\mathcal{J}_i$ as a graph for every $i \in \{1,\dots,k\}$. Let us denote this graph by G'_i. From Lemma 1.20, the clique complex of G'_i is a matroid (since G'_i is a matroid). Now we have that $G = \bigcap_{i=1}^k G'_i$. Therefore, it holds that $\mathfrak{C}(G) = \mathfrak{C}(\bigcap_{i=1}^k G'_i) = \bigcap_{i=1}^k \mathfrak{C}(G'_i)$. Since we have just observed that $\mathfrak{C}(G'_i)$ is a matroid for each $i \in \{1,\dots,k\}$, this completes the proof. □

1.7 Matching Complexes

In this section, we apply our theorems to matching complexes of graphs, and observe that some results by Fekete, Firla & Spille [FFS03] can be deduced from our more general theorems.

A *matching* of a graph $G = (V,E)$ is a subset $M \subseteq E$ of the edge set in which the edges are pairwise disjoint, that is, $e \cap e' = \emptyset$ for each $e, e' \in M$. The *matching complex* of a graph G is the family of matchings of G, and denoted by $\mathfrak{M}(G)$. We can see that the matching complex

$\mathfrak{M}(G)$ is indeed an independence system on E. Note that the matching complex $\mathfrak{M}(G)$ is identical to the clique complex of the complement of the line graph of G, i.e., $\mathfrak{M}(G) = \mathfrak{C}(\overline{L(G)})$. Recall that the *line graph* of a graph G is a graph $L(G)$ such that the vertex set of $L(G)$ is the edge set of G and two vertices in $L(G)$ are adjacent through an edge if and only if the corresponding two edges in G share a vertex in G. We also call a graph G a *line graph* if there exists some graph whose line graph is G. For a line graph G, a graph H is called a *root graph* of the line graph G if $G = L(H)$. Note that a root graph of a line graph is not unique in general. For example, K_3 is the line graph of K_3 and also of $K_{1,3}$, i.e., both K_3 and $K_{1,3}$ are root graphs of K_3. Also, note that not every graph is a line graph; for example, $K_{1,3}$ is not a line graph.

First, let us deduce the characterization of matroidal matching complexes from Lemma 1.6.

Corollary 1.22. *Let G be a graph. The matching complex $\mathfrak{M}(G)$ is a matroid if and only if G is the disjoint union of stars and triangles.*

Proof. Assume that $\mathfrak{M}(G)$ is a matroid. Since $\mathfrak{M}(G) = \mathfrak{C}(\overline{L(G)})$ holds, it holds that $\overline{L(G)}$ is a complete r-partite graph for some r by Lemma 1.6. This means that $L(G)$ is a disjoint union of complete graphs. Let K be a connected component of $L(G)$, which is a complete graph. Now, we want to find the root graphs of K. Then we observe that a root graph of K_1 is $K_2 (= K_{1,1})$, and this is a unique root graph of K_1; a root graph of K_2 is $K_{1,2}$, and this is a unique root graph of K_2; root graphs of K_3 are K_3 and $K_{1,3}$, and they are the only root graphs of K_3; a root graph of K_n ($n \geq 4$) is $K_{1,n}$, and this is a unique root graph of K_n. (Note that our graph is always simple, i.e., without a loop or a multiple edge.) Therefore, G is the disjoint union of stars and triangles.

Let us show the converse. Assume that G is the disjoint union of stars and triangles. Then we can see that $\overline{L(G)}$ is a complete multipartite graph. From Lemma 1.6, it follows that $\mathfrak{M}(G) = \mathfrak{C}(\overline{L(G)})$ is a matroid. The proof is completed. □

Fekete, Firla & Spille [FFS03] studied the matching complex in the same spirit as we did in this chapter. They proved the following statement for the intersection of two matroids. In this chapter, we derive this result as a corollary from our theorem.

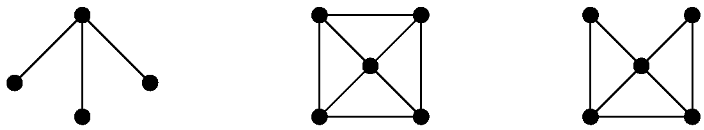

Figure 1.11: Graphs appearing in the proof of Corollary 1.23.

Corollary 1.23 ([FFS03]). *Let G be a graph. The matching complex $\mathfrak{M}(G)$ is the intersection of two matroids if and only if G contains no subgraph (not necessarily induced) isomorphic to C_{2k+3} ($k = 1,2,\ldots$), and each triangle in G has at most one vertex of degree more than two.*

To prove Corollary 1.23, we use the following fact on a line graph.

Lemma 1.24. *Let G be a graph, H be a line graph, and $R_1,\ldots,R_k$ be the root graphs of H. Then $L(G)$ contains no induced subgraph isomorphic to H if and only if G contains no subgraph (not necessarily induced) isomorphic to any of $R_1,\ldots,R_k$.*

Proof. Straightforward from the definitions of a line graph and a root graph. ☐

With use of Lemma 1.24, we can prove Corollary 1.23.

Proof of Corollary 1.23. Assume that there exist two matroids $\mathcal{I}_1, \mathcal{I}_2$ on $E(G)$ such that $\mathfrak{M}(G) = \mathcal{I}_1 \cap \mathcal{I}_2$. From the observation above, this is equivalent to $\mathfrak{C}(\overline{L(G)}) = \mathcal{I}_1 \cap \mathcal{I}_2$. By Theorem 1.18, this is also equivalent to $\overline{L(G)}$ contains no induced subgraph isomorphic to $K_1 \cup K_3$, $K_1 \cup K_2 \cup K_2$, $K_1 \cup P_3$ or $\overline{C_{2k+3}}$ ($k = 1,2,\ldots$). Therefore, by Lemma 1.24, we can see that this is also equivalent to $L(G)$ contains no subgraph (not necessarily induced) isomorphic to $K_{1,3} = \overline{K_1 \cup K_3}$, $W_4 = \overline{K_1 \cup K_2 \cup K_2}$, $W_4^- = \overline{K_1 \cup P_3}$ or C_{2k+3} ($k = 1,2,\ldots$). See Figure 1.11 for the shapes of these graphs.

About the root graphs of $K_{1,3}$, W_4, W_4^-, and C_{2k+3} ($k = 1,2,\ldots$), we observe the following.

(1) There is no root graph of $K_{1,3}$ (i.e., $K_{1,3}$ is not a line graph).

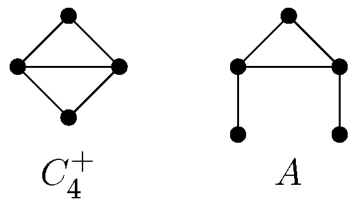

Figure 1.12: The root graphs appearing in the proof of Corollary 1.23.

(2) A root graph of W_4 is C_4^+ (in Figure 1.12) and this is a unique root graph of W_4.

(3) A root graph of W_4^- is A (in Figure 1.12) and this is a unique root graph of W_4^-.

(4) For each $k = 1, 2, \ldots$, a root graph of C_{2k+3} is C_{2k+3} and this is a unique root graph of C_{2k+3}.

Thus, we can see that Lemma 1.24 implies that the matching complex $\mathfrak{M}(G)$ is the intersection of two matroids if and only if G contains no subgraph isomorphic to C_4^+, A or C_{2k+3} ($k = 1, 2, \ldots$). Hence, for the proof of the corollary, it is enough to observe that G contains no subgraph isomorphic to C_4^+ or A if and only if each triangle in G has at most one vertex of degree more than two.

To observe that, first assume that G contains no subgraph isomorphic to C_4^+ or A and also suppose that there exists a triangle in G which has at least two vertices of degree more than two. Let u and v be such vertices in the triangle ($u \neq v$). Then the above assumption means that there exist edges $\{u, x\}$ and $\{v, y\}$ in G. In case $x = y$, we see that G contains C_4^+ as a subgraph. In case $x \neq y$, we see that G contains A as a subgraph. Therefore, in both cases this is a contradiction.

Conversely, assume that each triangle in G has at most one vertex of degree more than two. Pick a triangle T in G arbitrarily. Then we see that T cannot be contained in a subgraph isomorphic to C_4^+ or A in G since C_4^+ and A have two vertices of degree more than two. This means that G contains no subgraph isomorphic to C_4^+ or A. This concludes the proof. □

Fekete, Firla & Spille [FFS03] also gave a characterization of the matching complex which can be represented as the intersection of k matroids for a general k. Their characterization involves an integer programming formulation of the problem to find the right k. We observe that their characterization is also a corollary of our theorem. We need to introduce their formulation.

First, we introduce the variables. Since the circuits of $\mathfrak{M}(G)$ corresponds to the paths of length two (this is an immediate consequence from Lemma 1.4 and the fact that $\mathfrak{M}(G) = \mathfrak{C}(\overline{L(G)})$), it makes sense that we use a variable $x \in \{0,1\}^{\{1,\dots,k\}\times\mathcal{P}(G)}$ where $\mathcal{P}(G)$ denotes the family of all paths of length 2 in G. We denote a path of length 2 in G by (u,v,w) when v is the midpoint of the path and u, w are the endpoints of the path. Note that the path (w,v,u) is identified with the path (u,v,w), so exactly one of them belongs to $\mathcal{P}(G)$. The interpretation of the variable x is as follows. Assume that $\mathfrak{M}(G)$ is the intersection of k matroids $\mathcal{I}_1,\dots,\mathcal{I}_k$. For $i \in \{1,\dots,k\}$ and $(u,v,w) \in \mathcal{P}(G)$, $x[i,(u,v,w)] = 1$ if (u,v,w) is a circuit of $\mathcal{I}_i$; otherwise $x[i,(u,v,w)] = 0$. Fekete, Firla & Spille [FFS03] considered the following set of constraints.

Cover condition: $\sum_{i=1}^{k} x[i,(u,v,w)] \geq 1$ for all $(u,v,w) \in \mathcal{P}(G)$,

Claw condition: $x[i,(u,v,w)] + x[i,(u,v,t)] + x[i,(w,v,t)] \neq 2$ for all $i \in \{1,\dots,k\}$ and $(u,v,w),(u,v,t),(w,v,t) \in \mathcal{P}(G)$,

Triangle condition: $x[i,(u,v,w)] + x[i,(v,w,u)] + x[i,(w,u,v)] \neq 2$ for all $i \in \{1,\dots,k\}$ and $(u,v,w),(v,w,u),(w,u,v) \in \mathcal{P}(G)$,

Matching condition: $x[i,(u,v,w)] + x[i,(v,w,t)] \leq 1$ for all $i \in \{1,\dots,k\}$ and (u,v,w), $(v,w,t) \in \mathcal{P}(G)$.

(See Fekete, Firla & Spille [FFS03] for the detail of these constraints.) Note that the Claw condition and the Triangle condition can be written as linear inequality constraints as well.

Corollary 1.25 ([FFS03]). *Let G be a graph. Then $\mathfrak{M}(G)$ is the intersection of k matroids if and only if there exists a vector $x \in \{0,1\}^{\{1,\dots,k\}\times\mathcal{P}(G)}$ which satisfies all of the four conditions above (namely, the Cover condition, the Claw condition, the Triangle condition and the Matching condition).*

Proof. Let $G = (V,E)$ be a given graph. First, let us assume that $\mathfrak{M}(G) = \mathfrak{C}(\overline{L(G)})$ is the intersection of k matroids. Then, by Theorem 1.7, there exist k stable-set partitions $\mathcal{P}^{(1)},\dots,\mathcal{P}^{(k)}$ of $\overline{L(G)}$ which satisfy the following condition: $\{e,f\} \in \binom{E}{2}$ is an edge of $L(G)$ if and only if $\{e,f\} \subseteq S$ for some $S \in \bigcup_{i=1}^{k} \mathcal{P}^{(i)}$. Put $e = \{u,v\}$ and $f = \{w,t\}$ for some $u,v,w,t \in V$. Then, we can see that this condition is equivalent to that $\{e,f\} \in \binom{E}{2}$ forms a path $(u,v=t,w)$ of length 2 in G if and only if $\{e,f\} \subseteq S$ for some $S \in \bigcup_{i=1}^{k} \mathcal{P}^{(i)}$. In the sequel, we write "$(u,v,w) \subseteq S$" instead of "$\{e,f\} \subseteq S$" when $e = \{u,v\}$ and $f = \{w,t\}$ form the path $(u,v=t,w)$ of length 2. Let us summarize this condition as follows and call it Condition P′.

> Condition P′:
>
> $\{e,f\} \in \binom{E}{2}$ forms a path $(u,v=t,w)$ of length 2 in G (where $e=\{u,v\}$ and $f=\{w,t\}$) if and only if $(u,v,w) \subseteq S$ for some $S \in \bigcup_{i=1}^{k} \mathcal{P}^{(i)}$.

Now, we construct $\overline{x} \in \{0,1\}^{\{1,\dots,k\}\times\mathcal{P}(G)}$ from our stable-set partitions. For $i \in \{1,\dots,k\}$ and $(u,v,w) \in \mathcal{P}(G)$, set $\overline{x}[i,(u,v,w)] = 1$ if $(u,v,w) \subseteq S$ for some $S \in \mathcal{P}^{(i)}$; set $\overline{x}[i,(u,v,w)] = 0$ otherwise. We show that the vector $\overline{x}$ constructed above satisfies the four conditions.

First, check the Cover condition. Fix a path (u,v,w) of length 2 in G arbitrarily. Then, from Condition P′, there exists at least one index i^* such that $(u,v,w) \subseteq S$ for some $S \in \mathcal{P}^{(i^*)}$. Our construction implies that $\overline{x}[i^*,(u,v,w)] = 1$. Therefore, we have that $\sum_{i=1}^{k} \overline{x}[i,(u,v,w)] \geq 1$. Since the choice of the path (u,v,w) was arbitrary, this inequality holds for all paths of length 2 in G. Hence, $\overline{x}$ satisfies the Cover condition.

Secondly, we check the Claw condition. Suppose that the Claw condition is violated, namely there exist an index $i \in \{1,\dots,k\}$ and paths $(u,v,w),(u,v,t),(w,v,t) \in \mathcal{P}(G)$ such that $\overline{x}[i,(u,v,w)] + \overline{x}[i,(u,v,t)] + \overline{x}[i,(w,v,t)] = 2$. By the symmetry of $(u,v,w),(u,v,t),(w,v,t)$, we may assume that $\overline{x}[i,(u,v,w)] = 1$, $\overline{x}[i,(u,v,t)] = 1$ and $\overline{x}[i,(w,v,t)] = 0$ without loss of generality. The construction of $\overline{x}$ implies that there exist $S_{uw}, S_{ut} \in \mathcal{P}^{(i)}$ such that $(u,v,w) \subseteq S_{uw}$ and $(u,v,t) \subseteq S_{ut}$. Therefore, $\{u,v\} \in S_{uw}$ and $\{u,v\} \in S_{ut}$. This implies that $S_{uv} \cap S_{ut} \neq \emptyset$. On the other hand, since $\mathcal{P}^{(i)}$ is a partition of E and $S_{uw}, S_{ut} \in \mathcal{P}^{(i)}$, it holds that $S_{uw} \cap S_{ut} = \emptyset$. This is a contradiction. Thus, we have shown that $\overline{x}$ satisfies the Claw condition.

Next, we check the Triangle condition. Suppose that the Triangle condition is violated, i.e., there exist an index $i \in \{1,\dots,k\}$

and paths $(u,v,w),(v,w,u),(w,u,v) \in \mathcal{P}(G)$ such that $\overline{x}[i,(u,v,w)] + \overline{x}[i,(v,w,u)] + \overline{x}[i,(w,u,v)] = 2$. By the symmetry of $(u,v,w),(v,w,u)$, (w,u,v), we may assume that $\overline{x}[i,(u,v,w)] = 1$, $\overline{x}[i,(v,w,u)] = 1$ and $\overline{x}[i,(w,u,v)] = 0$, without loss of generality. Then, our construction implies that there exist $S_u, S_v \in \mathcal{P}^{(i)}$ such that $(u,v,w) \subseteq S_u$ and $(v,w,u) \subseteq S_v$. Therefore, we can see that $\{v,w\} \in S_u$ and $\{v,w\} \in S_v$. This means that $S_u \cap S_v \neq \emptyset$. On the other hand, since $\mathcal{P}^{(i)}$ is a partition of E, and $S_u, S_v \in \mathcal{P}^{(i)}$, it holds that $S_u \cap S_v = \emptyset$. So, they contradict each other. Thus, we have shown that $\overline{x}$ satisfies the Triangle condition.

Finally, we check the Matching condition. Suppose that the Matching condition is violated, i.e., there exist $i \in \{1,\dots,k\}$ and $(u,v,w),(v,w,t) \in \mathcal{P}(G)$ such that $\overline{x}[i,(u,v,w)] + \overline{x}[i,(v,w,t)] > 1$. Since $\overline{x}$ is a $\{0,1\}$-vector, we have that $\overline{x}[i,(u,v,w)] = 1$ and $\overline{x}[i,(v,w,t)] = 1$. Because of our construction, there exist $S_u, S_v \in \mathcal{P}^{(i)}$ such that $(u,v,w) \subseteq S_u$ and $(v,w,t) \subseteq S_v$. Therefore, we can see that $\{v,w\} \in S_u$ and $\{v,w\} \in S_v$. Then, by the same reason as in the case of the Triangle condition, we obtain a contradiction. Thus, we have checked that $\overline{x}$ meets the Matching condition.

The discussion above concludes the only-if-part of the corollary. So it remains to show the if-part.

Assume that there exists a vector $x \in \{0,1\}^{\{1,\dots,k\}\times\mathcal{P}(G)}$ which satisfies the Cover condition, the Claw condition, the Triangle condition and the Matching condition. From this vector, we construct k stable-set partitions $\mathcal{Q}^{(1)},\dots,\mathcal{Q}^{(k)}$ of $\overline{L(G)}$ which satisfy Condition P$'$ above. Since Condition P$'$ is equivalent to Condition P in Theorem 1.7, this concludes the proof.

Fix $i \in \{1,\dots,k\}$. Then we put $\{\{u,v\}\} \in \mathcal{Q}^{(i)}$ if there exists no $(u,v,w) \in \mathcal{P}(G)$ such that $x[i,(u,v,w)] = 1$ and also there exists no $(v,u,t) \in \mathcal{P}(G)$ such that $x[i,(v,u,t)] = 1$. Furthermore, we put $\{\{u,v\},\{v,w\}\} \in \mathcal{Q}^{(i)}$ if $x[i,(u,v,w)] = 1$.

We now must check that $\mathcal{Q}^{(i)}$ is indeed a stable-set partition of $\overline{L(G)}$ for each $i \in \{1,\dots,k\}$ as desired. Fix $i \in \{1,\dots,k\}$ arbitrarily. First, let us check that $\mathcal{Q}^{(i)}$ is a partition of $V(\overline{L(G)})$, i.e., a partition of $E(G)$. Clearly $E(G) = \bigcup \mathcal{Q}^{(i)}$ for each $i \in \{1,\dots,k\}$. Suppose, for contradiction, that there exist two distinct sets $S,T \in \mathcal{Q}^{(i)}$ such that $S \cap T \neq \emptyset$.

Since each set in $\mathcal{Q}^{(i)}$ is of size 1 or 2, we have the following two cases. As the first case, assume that $|S| = 1$ and $|T| = 2$, say $S = \{\{u,v\}\}$ and $T = \{\{u,v\},\{v,w\}\}$. However, this contradicts our construction of $\mathcal{Q}^{(i)}$. The second case is where $|S| = |T| = 2$. We have two subcases. Assume that $S = \{\{u,v\},\{v,w\}\}$ and $T = \{\{u,v\},\{v,t\}\}$ where $w \neq t$. Then from our construction we have that $x[i,(u,v,w)] = 1$ and $x[i,(u,v,t)] = 1$. By the Claw condition, we should have $x[i,(t,v,w)] = 1$. However, the Matching condition requires $x[i,(u,v,t)] + x[i,(t,v,w)] \leq 1$. This is a contradiction. Next, assume that, say, $S = \{\{u,v\},\{v,w\}\}$ and $T = \{\{v,u\},\{u,t\}\}$. In this case, again from the construction we have that $x[i,(u,v,w)] = 1$ and $x[i,(v,u,t)] = 1$. If $w \neq t$, then this contradicts the Matching condition. If $w = t$, then from the Triangle condition we should have that $x[i,(u,w,v)] = 1$. However, this again contradicts the Matching condition. Thus, $\mathcal{Q}^{(i)}$ partitions $E(G)$.

Secondly, we check that each set $S \in \mathcal{Q}^{(i)}$ is a stable set of $\overline{L(G)}$. If $|S| = 1$, then clearly S is stable. Assume that $|S| = 2$, say, $S = \{\{u,v\},\{v,w\}\}$. Since (u,v,w) is a path of length 2 in G, $\{u,v\}$ and $\{v,w\}$ are adjacent in $L(G)$. This means that they are not adjacent in $\overline{L(G)}$. Therefore $\{\{u,v\},\{v,w\}\}$ is stable in $\overline{L(G)}$. Thus, we proved that $\mathcal{Q}^{(i)}$ is a stable-set partition of $\overline{L(G)}$ for each $i \in \{1,\ldots,k\}$.

Now, we check the constructed stable-set partitions $\mathcal{Q}^{(1)},\ldots,\mathcal{Q}^{(k)}$ satisfy Condition P′ above. However, this can be easily checked with the Cover condition. This concludes the whole proof. ☐

1.8 Concluding Remarks

In this chapter, motivated by a quality of a natural greedy algorithm for the maximum weight clique problem, we characterized the clique complex of a graph which can be represented as the intersection of k matroids (Theorem 1.7). This implies that the problem to determine the clique complex of a given graph has a representation by k matroids or not belongs to NP (Corollary 1.12). Furthermore, in Section 1.5 we observed that the corresponding problem for two matroids can be solvable in polynomial time. However, the problem for three or more matroids is not known to be solved in polynomial time. We leave the further issue on computational complexity of this problem as an open problem. In addition, we showed that $n-1$ matroids are necessary and sufficient for the representation of the clique complexes of all graphs with n vertices (Theorem 1.13), and looked at the relationship between the clique complex of a graph and the graph itself as independence systems (Theorem 1.21).

In Corollary 1.11, we proved that the class of clique complexes is the same as the class of the intersections of partition matroids. This result sheds more light on the structure of clique complexes, and may give a new research direction to attack some open problems on them.

Formerly, Fekete, Firla & Spille [FFS03] studied matching complexes from the viewpoint of matroid intersections. In Section 1.7, we have observed that some of their results can be derived from our more general theorems.

Finally, we would like to mention open problems arising from the study. As mentioned at the end of Section 1.5, we are not aware of a polynomial-time algorithm to decide whether the clique complex of a given graph is the intersection of three matroids or not. This is open. As another open question, we want to mention the following. In Theorem 1.13, we showed that $\mu(n) = n-1$ for $n \geq 2$. There, the graph showing $\mu(n) \geq n-1$ was disconnected. Therefore, we can ask what we can say if we are restricted to graphs with a certain connectivity requirement. For example, what is the maximum possible $\mu(G)$ when G is connected, or 2-connected, ... and so on. This problem remains open.

Part II

Abstract Convex Geometries

Chapter 2

An attentive reader might have noticed that while here we call the function $e^{\sqrt{n}}$ subexponential, the title implicitly calls $e^{n^{1/3}}$ exponential. We believe that this is excusable: what one calls a mountain depends very much on whether one lives in Holland or in Switzerland, for example.

Jirka Matoušek (2004)

The Affine Representation Theorem for Abstract Convex Geometries

2.1 Introduction

Abstract models of geometric concepts are useful. For example, a matroid is considered as the abstraction of linear and affine dependence [Oxl92], and plays an important role in finite geometry and coding theory, and also in systems analysis [Mur00] and combinatorial optimization [Sch03], etc. Another example is an oriented matroid, also considered as the abstraction of linear and affine dependence, and it captures essences of convex polytopes, point configurations, and hyperplane arrangements [BLS$^+$99, Zie98]. Oriented matroids play an important role in the theory of convex polytopes, discrete geometry, computational geometry and linear programming, and they are known to be quite powerful models.

One of the most important theorems in oriented matroid theory is the "topological representation theorem" by Folkman & Lawrence [FL78]. The topological representation theorem states that every simple oriented matroid can be represented as a "pseudohyperplane arrangement." Therefore, in principle, when we study an oriented matroid, we only have to look at the corresponding pseudohyperplane arrangement. (Another proof of the topological representation theorem for oriented matroids has recently been found by Bokowski, King, Mock & Streinu [BKMS], built on the earlier work by Bokowski, Mock & Streinu [BMS01].) A recent study by Swartz [Swa03] revealed the topological representation of matroids, stating that every simple matroid can be represented as the arrangement of homotopy spheres.

In this chapter, we study yet another example of combinatorial abstraction of geometric concepts, namely abstract convex geometries. Abstract convex geometries (or convex geometries for short) were introduced by Edelman & Jamison [EJ85] as an abstraction of convexity, and they can be seen as a "dual" (or a "polar" or a "complement") of antimatroids [Die89]. (Therefore, we sometimes use the word "antimatroid" instead of "convex geometry" to express the same object.) Convex geometries and antimatroids appear not only in papers on discrete geometry [EJ85, ER00, ERW02] but also in some other areas like social choice theory [JD01, Kos99, MR01], knowledge spaces in mathematical psychology [DF99a], the discrete-event system [GY94], and semimodular lattices [Ste99]. Furthermore, convex geometries form a greedily solvable special case of a certain optimization problem [BF90], and a recent development has uncovered the relationship of convex geometries with submodular-type optimization [Fuj04, KO03]. From the opposite side of view, convex geometries form a special subclass of closure spaces, and antimatroids form a subclass of greedoids [BZ92, KLS91].

In this chapter, we prove the representation theorem for convex geometries. The theorem states that every convex geometry can be represented as a "generalized convex shelling." Since a generalized convex shelling is defined in a purely affine-geometric manner, this theorem gives an affine-geometric representation of a convex geometry. Since an affine-geometric representation theorem does exist neither for matroids nor for oriented matroids, our affine-geometric representation theorem for convex geometries indicates the intrinsic simplicity of convex geometries. Just as the topological representation theorem for oriented matroids plays a significant role in the theory of oriented matroids, we hope that our theorem plays a similar role in the theory of convex geometries.

Organization In Section 2.2, we give a definition of convex geometries and state our theorem precisely. The proof of the theorem is constructive. In Section 2.3, we give a construction for the proof. In Section 2.4, we collect facts on convex geometries which will be used for showing the validity of the construction. In Section 2.5, we conclude the proof. Section 2.6 summarizes the chapter and gives some recent progresses to which our result has opened the direction.

Geometric Preliminaries and Notation The set of non-negative real numbers and the set of positive real numbers are denoted by $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_{>0}$, respectively.

We call points $x_1, x_2, \ldots, x_n \in \mathbb{R}^d$ *affinely independent* if for any real numbers $\lambda_1, \lambda_2, \ldots, \lambda_n \in \mathbb{R}$,

$$\sum_{i=1}^{n} \lambda_i x_i = \mathbf{0} \qquad \text{and} \qquad \sum_{i=1}^{n} \lambda_i = 0$$

imply $\lambda_1 = \cdots = \lambda_n = 0$. An *affine basis* of $\mathbb{R}^d$ is a maximal set of points in $\mathbb{R}^d$ which are affinely independent. From the definition, we can see that the cardinality of an affine basis is always $d+1$, and if $\{x_1, x_2, \ldots, x_{d+1}\}$ is an affine basis, then for every point $x \in \mathbb{R}^d$ there exist real numbers $\lambda_1, \ldots, \lambda_{d+1}$ summing up to "1" such that $x = \sum_{i=1}^{d+1} \lambda_i x_i$.

A *convex combination* of points $x_1, x_2, \ldots, x_n \in \mathbb{R}^d$ is an expression $\sum_{i=1}^{n} \lambda_i x_i$ where $\lambda_1, \ldots, \lambda_n \in \mathbb{R}_{\geq 0}$ are non-negative real numbers summing up to 1. A set $S \subseteq \mathbb{R}^d$ is *convex* if all convex combinations of points in S also lie in S. For a set $S \subseteq \mathbb{R}^d$, the *convex hull* of S is a unique minimal convex set containing S, and is denoted by $\mathsf{conv}(S)$.

2.2 Convex Geometries and the Representation Theorem

In this section, we give the definition of a convex geometry, which was introduced by Edelman & Jamison [EJ85], and we state our theorem precisely.

Let E be a non-empty finite set. A family $\mathcal{L}$ of subsets of E is called a *convex geometry* on E if $\mathcal{L}$ satisfies the following three axioms:

(L1) $\emptyset \in \mathcal{L}$ and $E \in \mathcal{L}$;

(L2) if $X, Y \in \mathcal{L}$, then $X \cap Y \in \mathcal{L}$;

(L3) if $X \in \mathcal{L} \setminus \{E\}$, then there exists some element $e \in E \setminus X$ such that $X \cup \{e\} \in \mathcal{L}$.

A member of a convex geometry $\mathcal{L}$ is called a *convex set*. Two convex geometries $\mathcal{L}_1$ on E_1 and $\mathcal{L}_2$ on E_2 are *isomorphic* if there exists a bijection $\psi : E_1 \to E_2$ such that $\psi(X) \in \mathcal{L}_2$ if and only if $X \in \mathcal{L}_1$.

Let us look at some examples of convex geometries.

Example 2.1 (convex shelling). Let P be a finite set of distinct points in $\mathbb{R}^d$, and define

$$\mathcal{L} := \{X \subseteq P \mid \mathsf{conv}(X) \cap P = X\}.$$

Then, we can see that $\mathcal{L}$ is a convex geometry on P, and we call this kind of convex geometries a *convex shelling* on P. A convex geometry isomorphic to the convex shelling on some finite point set P is also called a convex shelling.

Example 2.2 (poset shelling). Let E be a partially ordered set endowed with a partial order $\preceq$, and define

$$\mathcal{L} := \{X \subseteq E \mid e \in X \text{ and } f \preceq e \text{ imply } f \in X\}.$$

(Namely, $\mathcal{L}$ is the family of order ideals of E.) Then we can see that $\mathcal{L}$ is a convex geometry on E, and we call this kind of convex geometries a *poset shelling* on E.

Example 2.3 (tree shelling). Let V be the vertex set of a (graph-theoretic) tree T, and define

$$\mathcal{L} := \{X \subseteq V \mid G[X] \text{ is connected}\},$$

where $G[X]$ denotes the subgraph of G induced by X. Then we can see that $\mathcal{L}$ is a convex geometry on V, and we call this kind of convex geometries a *tree shelling*.

Example 2.4 (graph search). Let $G = (V, E)$ be a connected graph with root $r \in V$, and define

$$\mathcal{L} := \{X \subseteq V \setminus \{r\} \mid G[V \setminus X] \text{ is connected}\},$$

where $G[V \setminus X]$ again denotes the subgraph of G induced by $V \setminus X$. Then we can see that $\mathcal{L}$ is a convex geometry on $V \setminus \{r\}$, and we call this kind of convex geometries a *graph search*.

In the literature, we can find more examples of convex geometries arising from various objects [EJ85, KLS91].

Now we introduce yet another example of convex geometries, which was so far not mentioned explicitly.

Example 2.5 (generalized convex shelling). Let P and Q be finite point sets in $\mathbb{R}^d$ such that $P \cap \mathsf{conv}(Q) = \emptyset$. (In particular, $P \cap Q = \emptyset$.) Then define

$$\mathcal{L} := \{X \subseteq P \mid \mathsf{conv}(X \cup Q) \cap P = X\}.$$

We call $\mathcal{L}$ the *generalized convex shelling on P with respect to Q*. If $Q = \emptyset$, this just gives a convex shelling on P. So, as the name indicates, a generalized convex shelling is a generalization of a convex shelling. While at first sight it is not obvious that a generalized convex shelling is indeed a convex geometry, later we will prove it as Lemma 2.2.

A generalized convex shelling is related to a minor of a convex geometry. Let $\mathcal{L}$ be a convex geometry and $A, B \in \mathcal{L}$ be convex sets of $\mathcal{L}$ such that $A \subseteq B$. Then, define

$$\mathcal{L}[A, B] := \{X \subseteq B \setminus A \mid X \cup A \in \mathcal{L}\}.$$

As in the following lemma, it is known that $\mathcal{L}[A,B]$ is a convex geometry on $B \setminus A$ and it is called a *minor* of $\mathcal{L}$. (Remark that the definition of a minor is different from that in a paper of Edelman & Jamison [EJ85]. Rather, our definition obeys that in the book by Korte, Lovász & Schrader [KLS91].)

Lemma 2.1. *Let $\mathcal{L}$ be a convex geometry on E and $A, B \in \mathcal{L}$ satisfy $A \subseteq B \subseteq E$. Then $\mathcal{L}[A,B]$ is a convex geometry on $B \setminus A$.*

Proof. We only have to check that $\mathcal{L}[A,B]$ satisfies (L1), (L2) and (L3). Let us check (L1) first. Since $A \in \mathcal{L}$, we have $\emptyset \cup A = A \in \mathcal{L}$. Hence $\emptyset \in \mathcal{L}[A,B]$. Similarly, since $B \in \mathcal{L}$, we have $(B \setminus A) \cup A = B \in \mathcal{L}$. Hence $B \setminus A \in \mathcal{L}[A,B]$.

Secondly, we check (L2). Choose $X, Y \in \mathcal{L}[A,B]$ arbitrarily. Then, it follows that $X \cup A, Y \cup A \in \mathcal{L}$. Using (L2) for $\mathcal{L}$, we get $(X \cup A) \cap (Y \cup A) \in \mathcal{L}$, namely $(X \cap Y) \cup A \in \mathcal{L}$. Therefore, it holds that $X \cap Y \in \mathcal{L}[A,B]$.

Finally, we check (L3). Choose $X \in \mathcal{L}[A,B] \setminus \{B \setminus A\}$ arbitrarily. Then we have $X \cup A \in \mathcal{L}$, $X \cap A = \emptyset$ and $X \cup A \subsetneq B$. Applying (L3) to $X \cup A$ many times, we can find a sequence $e_1, e_2, \ldots, e_k \in E \setminus (X \cup A)$ of elements such that $(X \cup A) \cup \{e_1, \ldots, e_i\} \in \mathcal{L}$ for all $i \in \{1, \ldots, k\}$ and $(X \cup A) \cup \{e_1, \ldots, e_k\} = E$. Let i^* be the minimal index in $\{1, \ldots, k\}$ such that $e_{i^*} \in B \setminus (X \cup A)$. Then we can see that $((X \cup A) \cup \{e_1, \ldots, e_{i^*}\}) \cap B = (X \cup A) \cup \{e_{i^*}\}$ and from (L2) we can also see that this belongs to $\mathcal{L}$. Thus we have found $e_{i^*} \in B \setminus (X \cup A)$ such that $(X \cup A) \cup \{e_{i^*}\} \in \mathcal{L}$, namely $X \cup \{e_{i^*}\} \in \mathcal{L}[A,B]$. $\square$

In this proof, we have used the "chain argument," which is useful in the theory of convex geometries, and will be used again in the rest of this chapter.

The next lemma shows that a generalized convex shelling is a minor of some convex shelling. Together with Lemma 2.1, this implies that a generalized convex shelling is a convex geometry.

Lemma 2.2. *Let P and Q be finite point sets in $\mathbb{R}^d$ such that $P \cap \mathrm{conv}(Q) = \emptyset$. Furthermore, let $\mathcal{L}$ be the generalized convex shelling on P with respect to Q, and $\widetilde{\mathcal{L}}$ be the convex shelling on $P \cup Q$. Then it holds that $\mathcal{L} = \widetilde{\mathcal{L}}[Q, P \cup Q]$.*

Proof. First, because of the condition that $P \cap \mathsf{conv}(Q) = \emptyset$, it follows that $Q \in \widetilde{\mathcal{L}}$. So, $\widetilde{\mathcal{L}}[Q, P \cup Q]$ is well-defined. Since

$$\widetilde{\mathcal{L}} = \{X \subseteq P \cup Q \mid \mathsf{conv}(X) \cap (P \cup Q) = X\},$$

it follows that

$$\begin{aligned}
\widetilde{\mathcal{L}}[Q, P \cup Q] &= \{X \subseteq (P \cup Q) \setminus Q \mid X \cup Q \in \widetilde{\mathcal{L}}\} \\
&= \{X \subseteq P \mid X \cup Q \in \widetilde{\mathcal{L}}\} \\
&= \{X \subseteq P \mid \mathsf{conv}(X \cup Q) \cap (P \cup Q) = X \cup Q\} \\
&= \{X \subseteq P \mid \mathsf{conv}(X \cup Q) \cap P = X\} \\
&= \mathcal{L}.
\end{aligned}$$

Notice that the derivations of the second and the fourth identities use the assumption that $P \cap \mathsf{conv}(Q) = \emptyset$, in particular $P \cap Q = \emptyset$. This concludes the proof. $\square$

We are ready to state our main theorem. This states that the class of convex geometries coincides with the class of generalized convex shellings defined geometrically, although convex geometries arise from diverse objects as we have seen.

Theorem 2.3. *Every convex geometry is isomorphic to a generalized convex shelling.*

The main concern of this chapter is the proof of Theorem 2.3. For the proof of Theorem 2.3, in the next section we construct finite sets P_0 and Q_0 of points from a given convex geometry $\mathcal{L}$ with the following property: $\mathcal{L}$ is isomorphic to the generalized convex shelling on P_0 with respect to Q_0. In Section 2.4, we prepare more concepts from convex geometries which are needed in the proof. Section 2.5 completes the proof of the validity of the construction.

2.3 Construction of Point Sets

In our construction, we use rooted circuits of a convex geometry. So, at the beginning of this section, we introduce rooted circuits. A rooted circuit of a convex geometry was originally defined by Korte & Lovász [KL84].

In order to define a rooted circuit, we need other technical terms. For a convex geometry $\mathcal{L}$ on E and $A \subseteq E$, the *trace* of $\mathcal{L}$ on A is defined as

$$\mathsf{Tr}(\mathcal{L}, A) := \{X \cap A \mid X \in \mathcal{L}\}.$$

A *rooted set* is a pair (X, r) of a set X and an element r of X. A *rooted subset* of E is a rooted set (X, r) such that $X \subseteq E$.

Here comes the definition of a rooted circuit. Let $\mathcal{L}$ be a convex geometry on E. A rooted subset (C, r) of E is called a *rooted circuit* of $\mathcal{L}$ if $\mathsf{Tr}(\mathcal{L}, C) = 2^C \setminus \{C \setminus \{r\}\}$. We denote by $\mathcal{C}(\mathcal{L})$ the family of rooted circuits of a convex geometry $\mathcal{L}$.

Now we are ready for our construction. We construct point sets P_0 and Q_0 from a given convex geometry $\mathcal{L}$ on E so that $\mathcal{L}$ can be isomorphic to the generalized convex shelling on P_0 with respect to Q_0.

Let $n := |E|$. We use an $(n-1)$-dimensional space $\mathbb{R}^{n-1}$. For each element $e \in E$, we take a point $p(e) \in \mathbb{R}^{n-1}$ such that the points $\{p(e) \in \mathbb{R}^{n-1} \mid e \in E\}$ form an affine basis of $\mathbb{R}^{n-1}$. Then, for each rooted circuit $(C, r) \in \mathcal{C}(\mathcal{L})$ of $\mathcal{L}$ we choose a point $q(C, r) \in \mathbb{R}^{n-1}$ determined as

$$q(C, r) := |C|p(r) - \sum_{e \in C \setminus \{r\}} p(e). \tag{2.1}$$

Note that $p(r)$ is a convex combination of the points in $\{p(e) \mid e \in C \setminus \{r\}\} \cup \{q(C, r)\}$ with positive coefficients for every rooted circuit $(C, r) \in \mathcal{C}(\mathcal{L})$. Actually, this property is all that is needed in the construction. The definition of $q(C, r)$ above is just one of such choices, but it eases the later calculation. Thus, we have set up $|E| + |\mathcal{C}(\mathcal{L})|$ points in $\mathbb{R}^{n-1}$.

Let $P_0 = \{p(e) \mid e \in E\}$ and $Q_0 = \{q(C, r) \mid (C, r) \in \mathcal{C}(\mathcal{L})\}$. Then it holds that $P_0 \cap Q_0 = \emptyset$. Now our claim is as follows.

Lemma 2.4. *For P_0 and Q_0 constructed as above, the generalized convex shelling on P_0 with respect to Q_0 is isomorphic to $\mathcal{L}$.*

This lemma proves Theorem 2.3. The proof of Lemma 2.4 will be done in Section 2.5.

To illustrate the construction, let us look at examples for $n = 3$. Below we enumerate all of the six non-isomorphic convex geometries on $E = \{1,2,3\}$ together with their rooted circuits.

$$\begin{aligned}
&\mathcal{L}_1 = 2^{\{1,2,3\}}, && \mathcal{C}(\mathcal{L}_1) = \emptyset,\\
&\mathcal{L}_2 = \mathcal{L}_1 \setminus \{\{1,3\}\}, && \mathcal{C}(\mathcal{L}_2) = \{(\{1,2,3\},2)\},\\
&\mathcal{L}_3 = \mathcal{L}_2 \setminus \{\{3\}\}, && \mathcal{C}(\mathcal{L}_3) = \{(\{2,3\},2)\},\\
&\mathcal{L}_4 = \mathcal{L}_3 \setminus \{\{2,3\}\}, && \mathcal{C}(\mathcal{L}_4) = \{(\{1,3\},1),(\{2,3\},2)\},\\
&\mathcal{L}_5 = \mathcal{L}_3 \setminus \{\{1\}\}, && \mathcal{C}(\mathcal{L}_5) = \{(\{1,2\},2),(\{2,3\},2)\},\\
&\mathcal{L}_6 = \mathcal{L}_4 \setminus \{\{2\}\}, && \mathcal{C}(\mathcal{L}_6) = \{(\{1,2\},1),(\{1,3\},1),(\{2,3\},2)\}.
\end{aligned}$$

Figure 2.1 depicts the construction of the point sets for these examples.

2.4 More Properties of Convex Geometries

In this section, we introduce more concepts from the theory of convex geometries, which will be needed in the proof of Lemma 2.4 (i.e., Theorem 2.3). The reader is encouraged to interpret these concepts and lemmas with the examples in Section 2.2.

Let $\mathcal{L}$ be a convex geometry on E. Then the *closure operator* of $\mathcal{L}$ is a map $\tau_{\mathcal{L}} : 2^E \to 2^E$ defined as

$$\tau_{\mathcal{L}}(A) := \bigcap\{X \in \mathcal{L} \mid A \subseteq X\}$$

for $A \subseteq E$. By (L2) in the definition of a convex geometry, we can see that $\tau_{\mathcal{L}}(A) \in \mathcal{L}$ for every $A \subseteq E$. Furthermore, from the definition of a closure operator, we can prove the following facts.

Lemma 2.5. *Let $\mathcal{L}$ be a convex geometry on E, and $\tau_{\mathcal{L}}$ the closure operator of $\mathcal{L}$.*

(T1) *(Characterization of convex sets.) For $X \subseteq E$, it holds that $X \in \mathcal{L}$ if and only if $\tau_{\mathcal{L}}(X) = X$.*

(T2) *(Extensionality.) $A \subseteq \tau_{\mathcal{L}}(A)$ for $A \subseteq E$.*

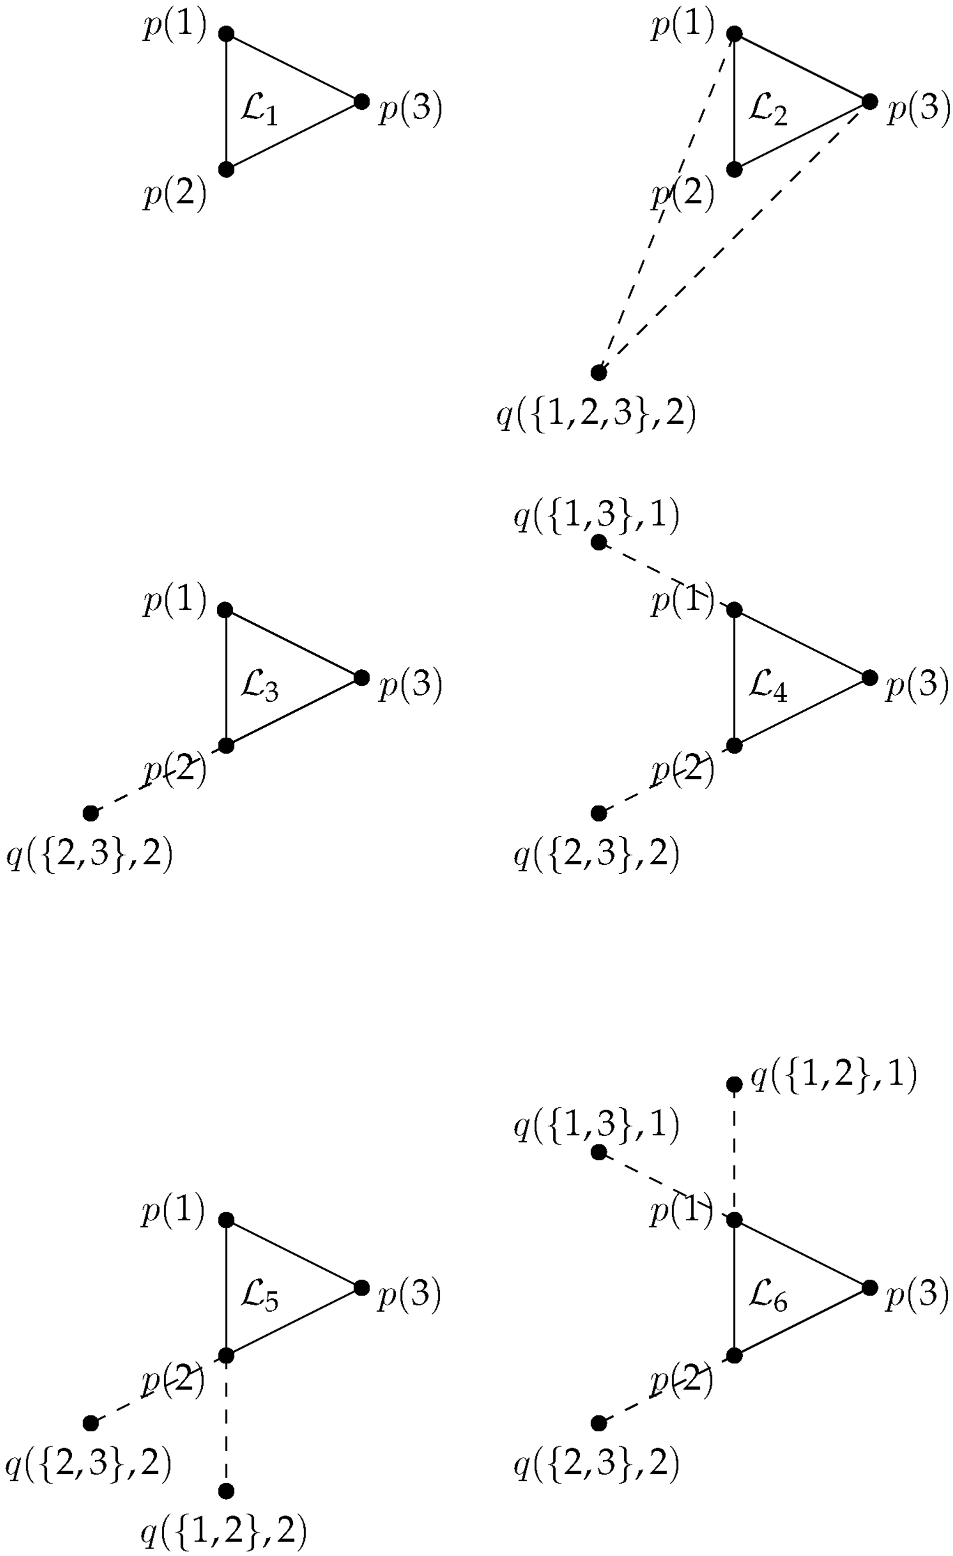

Figure 2.1: Construction of the point sets for $n = 3$.

(T3) *(Idempotence.)* $\tau_{\mathcal{L}}(\tau_{\mathcal{L}}(A)) = \tau_{\mathcal{L}}(A)$ *for* $A \subseteq E$.

(T4) *(Monotonicity.)* $A \subseteq B$ *implies* $\tau_{\mathcal{L}}(A) \subseteq \tau_{\mathcal{L}}(B)$.

(T5) *(Anti-exchange property.) Let* $A \subseteq E$ *and* $e, f \in E$ *such that* $e \neq f$ *and* $e, f \notin \tau_{\mathcal{L}}(A)$*. If* $f \in \tau_{\mathcal{L}}(A \cup \{e\})$ *then* $e \notin \tau_{\mathcal{L}}(A \cup \{f\})$*.*

Proof. The properties (T1)–(T4) are immediate from the definitions. The proof of the antiexchange property (T5) goes as follows.

Let A, e and f be as in the description of (T5). Further, let X be a set such that $X \subseteq E \setminus \{e\}$, $X \in \mathcal{L}$ and X is maximal (in the sense of set-inclusion) with respect to these two properties. Since $\tau_{\mathcal{L}}(A) \subseteq E \setminus \{e\}$ and $\tau_{\mathcal{L}}(A) \in \mathcal{L}$, such a set X always exists, and we have $A \subseteq \tau_{\mathcal{L}}(A) \subseteq X$. By (L3) in the definition of a convex geometry, there exists some element $e' \in E \setminus X$ such that $X \cup \{e'\} \in \mathcal{L}$. If $e' \neq e$, then $X \cup \{e'\} \subseteq E \setminus \{e\}$. This contradicts the maximality of X. Therefore $e' = e$ follows. This implies that $X \cup \{e\} \in \mathcal{L}$.

Assume that $f \in \tau_{\mathcal{L}}(A \cup \{e\})$. Since $A \cup \{e\} \subseteq X \cup \{e\}$ and $X \cup \{e\} \in \mathcal{L}$, it holds that $f \in \tau_{\mathcal{L}}(A \cup \{e\}) \subseteq \tau_{\mathcal{L}}(X \cup \{e\}) = X \cup \{e\}$. (Here, we have used (T4) and (T1).) Since $e \neq f$, we have $f \in X$. This means that $X \cup \{f\} = X$. Therefore, it follows that $\tau_{\mathcal{L}}(X \cup \{f\}) = \tau_{\mathcal{L}}(X) = X \not\ni e$. (Here again we have used (T1).) By the monotonicity (T4) we have that $\tau_{\mathcal{L}}(A \cup \{f\}) \subseteq \tau_{\mathcal{L}}(X \cup \{f\})$. Hence it holds that $e \notin \tau_{\mathcal{L}}(A \cup \{f\})$. $\square$

Note that the properties (T1)–(T4) of Lemma 2.5 even hold for more general "closure spaces" [EJ85, KLS91]. So the anti-exchange property (T5) is a characteristic feature of convex geometries. Actually, (T5) characterizes a convex geometry in the following sense: a map $\tau : 2^E \to 2^E$ satisfying extensionality, idempotence, monotonicity and also $\tau(\emptyset) = \emptyset$ is the closure operator of some convex geometry if and only if τ additionally satisfies the anti-exchange property [EJ85, KLS91].

In the following lemma, we can see that a trace of a convex geometry is again a convex geometry and that the closure operator of a trace is nicely combined with the closure operator of the original convex geometry.

Lemma 2.6. *Let* $\mathcal{L}$ *be a convex geometry on* E*, and* $\tau_{\mathcal{L}}$ *the closure operator of* $\mathcal{L}$*. Then,* $\mathsf{Tr}(\mathcal{L}, A)$ *is a convex geometry on* A *for every* $A \subseteq E$*. Moreover, the closure operator* $\tau_{\mathsf{Tr}(\mathcal{L},A)} : 2^A \to 2^A$ *of* $\mathsf{Tr}(\mathcal{L}, A)$ *is derived as*

$$\tau_{\mathsf{Tr}(\mathcal{L},A)}(B) = \tau_{\mathcal{L}}(B) \cap A \text{ for } B \subseteq A.$$

Proof. Fix $A \subseteq E$ and check that $\mathsf{Tr}(\mathcal{L}, A)$ satisfies (L1), (L2) and (L3). First we check (L1). Since $\emptyset \cap A = \emptyset$, we have $\emptyset \in \mathsf{Tr}(\mathcal{L}, A)$. In addition, since $E \cap A = A$, we have $A \in \mathsf{Tr}(\mathcal{L}, A)$.

Next we check (L2). Choose $X \cap A, Y \cap A \in \mathsf{Tr}(\mathcal{L}, A)$ where $X, Y \in \mathcal{L}$. Since $X \cap Y \in \mathcal{L}$ by (L2), we have $(X \cap A) \cap (Y \cap A) = (X \cap Y) \cap A \in \mathsf{Tr}(\mathcal{L}, A)$.

Finally we check (L3). Choose $X \cap A \in \mathsf{Tr}(\mathcal{L}, A)$ where $X \in \mathcal{L}$. By (L3) there exists $e \in E \setminus X$ such that $X \cup \{e\} \in \mathcal{L}$. If $e \in A \setminus X$, we are done. If not, applying (L3) to X many times, we get a sequence $e_1, \ldots, e_k \in E \setminus X$ such that $X \cup \{e_1, \ldots, e_i\} \in \mathcal{L}$ for all $i = 1, \ldots, k$ and $X \cup \{e_1, \ldots, e_k\} = E$. Let i^* be the minimum index such that $e_{i^*} \in A$. Then we have $(X \cup \{e_1, \ldots, e_{i^*}\}) \cap A = (X \cup \{e_{i^*}\}) \cap A \in \mathsf{Tr}(\mathcal{L}, A)$. Thus we have found $e_{i^*} \in A \setminus (X \cap A)$ such that $(X \cup \{e_{i^*}\}) \cap A \in \mathsf{Tr}(\mathcal{L}, A)$.

For the second part, we just calculate as follows. For any $B \subseteq A$,

$$\begin{aligned}
\tau_{\mathsf{Tr}(\mathcal{L},A)}(B) &= \bigcap \{X \in \mathsf{Tr}(\mathcal{L}, A) \mid B \subseteq X\} \\
&= \bigcap \{X \cap A \mid X \in \mathcal{L}, B \subseteq X\} \\
&= (\bigcap \{X \in \mathcal{L} \mid B \subseteq X\}) \cap A \\
&= \tau_{\mathcal{L}}(B) \cap A.
\end{aligned}$$

Here, the first and the last identities are due to the definition of the closure operator. The second one comes from the definition of the trace. □

Now, we look at how the closure operator reveals properties of rooted circuits.

Lemma 2.7. *Let $\mathcal{L}$ be a convex geometry on E. If (C, r) is a rooted circuit of $\mathcal{L}$, then $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$.*

Proof. Assume that $(C, r) \in \mathcal{C}(\mathcal{L})$. This means that $\mathsf{Tr}(\mathcal{L}, C) = 2^C \setminus \{C \setminus \{r\}\}$. Since $\tau_{\mathcal{L}}(C \setminus \{r\}) = \bigcap \{X \in \mathcal{L} \mid C \setminus \{r\} \subseteq X\}$ by definition, in order to show that $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$ we only have to check that $r \in X$ for all $X \in \mathcal{L}$ such that $C \setminus \{r\} \subseteq X$. Take such a set X arbitrarily. Then it holds that

$$X \cap C = \begin{cases} C & (r \in X), \\ C \setminus \{r\} & (r \notin X), \end{cases}$$

since $C \setminus \{r\} \subseteq X$. However, if $X \cap C = C \setminus \{r\}$, one would conclude that $C \setminus \{r\} \in \mathsf{Tr}(\mathcal{L}, C)$. (Recall the definition of the trace: $\mathsf{Tr}(\mathcal{L}, C) = \{X \cap C \mid X \in \mathcal{L}\}$.) This contradicts our assumption. So it should hold that $X \cap C = C$, which means $r \in X$. □

Here is another lemma.

Lemma 2.8. *Let $\mathcal{L}$ be a convex geometry on E, and $r \notin X \subseteq E$. Then $r \in \tau_{\mathcal{L}}(X) \setminus X$ if and only if there exists $C \subseteq X \cup \{r\}$ such that (C, r) is a rooted circuit of $\mathcal{L}$.*

Proof. First we prove the if-part. Assume that there exists $C \subseteq X \cup \{r\}$ such that $(C, r) \in \mathcal{C}(\mathcal{L})$. Then, from Lemma 2.7, we can see that $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$. Combining this with $\tau_{\mathcal{L}}(C \setminus \{r\}) \subseteq \tau_{\mathcal{L}}(X)$ (following by the monotonicity (T4)) and $r \notin X$, we obtain $r \in \tau_{\mathcal{L}}(X) \setminus X$.

To prove the converse, assume that $r \in \tau_{\mathcal{L}}(X) \setminus X$. Let $D \subseteq X$ be a minimal subset of X satisfying $r \in \tau_{\mathcal{L}}(D)$. Note that such a set D always exists because X itself satisfies $r \in \tau_{\mathcal{L}}(X)$, and D is not empty because $\tau_{\mathcal{L}}(\emptyset) = \emptyset$ by (L1) and (T1). Now we claim that $(D \cup \{r\}, r)$ is a rooted circuit of $\mathcal{L}$. This will finish the whole proof.

Since D is not empty, we may choose an arbitrary element $e \in D$. By the minimality of D, it holds that $r \notin \tau_{\mathcal{L}}(D \setminus \{e\})$. Then, we claim the following.

Claim 2.8.1. *It holds that $e \notin \tau_{\mathcal{L}}(D \setminus \{e\})$ for every $e \in D$.*

Proof of 2.8.1. Suppose the contrary; namely, $e \in \tau_{\mathcal{L}}(D \setminus \{e\})$. Then, using monotonicity (T4), we obtain

$$D = (D \setminus \{e\}) \cup \{e\} \subseteq \tau_{\mathcal{L}}(D \setminus \{e\}) \cup \{e\} = \tau_{\mathcal{L}}(D \setminus \{e\}).$$

By monotonicity (T4) and idempotence (T3), we can see that $\tau_{\mathcal{L}}(D) \subseteq \tau_{\mathcal{L}}(\tau_{\mathcal{L}}(D \setminus \{e\})) = \tau_{\mathcal{L}}(D \setminus \{e\})$. On the other hand, it holds that $\tau_{\mathcal{L}}(D \setminus \{e\}) \subseteq \tau_{\mathcal{L}}(D)$ again by the monotonicity (T4). Therefore, it holds that

$\tau_{\mathcal{L}}(D) = \tau_{\mathcal{L}}(D \setminus \{e\})$. Since $r \in \tau_{\mathcal{L}}(D)$ by the choice of D, this implies that $r \in \tau_{\mathcal{L}}(D \setminus \{e\})$. However, this contradicts the observation that $r \notin \tau_{\mathcal{L}}(D \setminus \{e\})$ as seen just above the statement of this claim. Thus, the claim is proven. ☐

By Claim 2.8.1, the observation above that $r \notin \tau_{\mathcal{L}}(D \setminus \{e\})$ and the monotonicity (T4), it follows that

$$D \setminus \{e\} = (D \cup \{r\}) \cap \tau_{\mathcal{L}}(D \setminus \{e\}) \in \mathsf{Tr}(\mathcal{L}, D \cup \{r\}). \tag{2.2}$$

(Remember that $\tau_{\mathcal{L}}(A) \in \mathcal{L}$ for all $A \subseteq E$.) Furthermore, since $r \in \tau_{\mathcal{L}}(D)$, it holds that $(D \setminus \{e\}) \cup \{r\} \subseteq \tau_{\mathcal{L}}(D)$ by the monotonicity (T4). Hence, we obtain

$$(D \setminus \{e\}) \cup \{r\} = ((D \setminus \{e\}) \cup \{r\}) \cap \tau_{\mathcal{L}}(D) \in \mathsf{Tr}(\mathcal{L}, D \cup \{r\}). \tag{2.3}$$

Since the expressions (2.2) and (2.3) hold for all $e \in D$, by using (L2) we can see that $\mathsf{Tr}(\mathcal{L}, D \cup \{r\}) = 2^{D \cup \{r\}} \setminus \{D\}$. ☐

The following lemma due to Korte & Lovász [KL84] says that the family of rooted circuits of a convex geometry determines it uniquely.

Lemma 2.9. *Let $\mathcal{C}(\mathcal{L})$ be the family of rooted circuits of a convex geometry $\mathcal{L}$ on E. Then we have*

$$\mathcal{L} = \{X \subseteq E \mid (E \setminus X) \cap C \neq \{r\} \textit{ for all } (C, r) \in \mathcal{C}(\mathcal{L})\}.$$

Proof. First we show that

$$\mathcal{L} \subseteq \{X \subseteq E \mid (E \setminus X) \cap C \neq \{r\} \text{ for all } (C, r) \in \mathcal{C}(\mathcal{L})\}.$$

Choose $X \in \mathcal{L}$ arbitrarily, and suppose that there exists some rooted circuit $(C, r) \in \mathcal{C}(\mathcal{L})$ such that $(E \setminus X) \cap C = \{r\}$. Then we have $X \cap C = C \setminus \{r\}$. However, this means that $C \setminus \{r\} \in \mathsf{Tr}(\mathcal{L}, C)$, which is a contradiction to the definition of a rooted circuit. So it should hold that $(E \setminus X) \cap C \neq \{r\}$ for all $(C, r) \in \mathcal{C}(\mathcal{L})$.

Let us show the other direction. Choose $X \notin \mathcal{L}$ arbitrarily. This means that $X \subsetneq \tau_{\mathcal{L}}(X)$ by (T1) and (T2). Therefore, there exists an element $r \in \tau_{\mathcal{L}}(X) \setminus X$. By Lemma 2.8, we have a set $C \subseteq X \cup \{r\}$ such that (C, r) is a rooted circuit of $\mathcal{L}$. So it follows that $(E \setminus X) \cap C = \{r\}$, concluding the lemma. ☐

The next lemma shows that rooted circuits are minimal in a certain sense.

Lemma 2.10. *Let $\mathcal{L}$ be a convex geometry on E, and (C, r) a rooted circuit of $\mathcal{L}$. Then $\mathsf{Tr}(\mathcal{L}, D) = 2^D$ for any proper subset $D \subsetneq C$.*

Proof. First of all, observe that

$$\begin{aligned}\mathsf{Tr}(\mathcal{L}, D) &= \{X \cap D \mid X \in \mathcal{L}\} \\ &= \{(X \cap C) \cap D \mid X \in \mathcal{L}\} \\ &= \{Y \cap D \mid Y \in \mathsf{Tr}(\mathcal{L}, C)\} \\ &= \{Y \cap D \mid Y \in 2^C \setminus \{C \setminus \{r\}\}\}.\end{aligned}$$

Here, the first and the third identities are due to the definition of a trace. The second one comes from the assumption that $D \subsetneq C$, and the last one from the definition of a rooted circuit.

Now, we have two cases. First consider the case in which $D \neq C \setminus \{r\}$. Then, all subsets of D belong to $2^C \setminus \{C \setminus \{r\}\}$. Therefore, $\mathsf{Tr}(\mathcal{L}, D) = 2^D$ holds. Next consider the case in which $D = C \setminus \{r\}$. Then, $C \cap D = C \setminus \{r\}$ and every proper subset of D belongs to $2^C \setminus \{C \setminus \{r\}\}$. Therefore, we also have $\mathsf{Tr}(\mathcal{L}, D) = 2^D$. □

Here are more properties of rooted circuits.

Lemma 2.11. *Let $\mathcal{L}$ be a convex geometry on E, and $\mathcal{C}$ be the family of rooted circuits of $\mathcal{L}$. Then the following properties hold.*

(C1) *If $(C_1, r), (C_2, r) \in \mathcal{C}$ and $C_1 \subseteq C_2$, then $C_1 = C_2$.*

(C2) *If $(C_1, r_1), (C_2, r_2) \in \mathcal{C}$ and $r_1 \in C_2 \setminus \{r_2\}$, then there exists $(C_3, r_2) \in \mathcal{C}$ such that $C_3 \subseteq C_1 \cup C_2 \setminus \{r_1\}$.*

Proof. Let us first prove (C1). Suppose $C_1 \subsetneq C_2$. Then, using Lemma 2.10, we can see that $\mathsf{Tr}(\mathcal{L}, C_1) = 2^{C_1}$. This is a contradiction to the assumption that (C_1, r) is a rooted circuit. Hence (C1) follows.

Next we prove (C2). Let $X := (C_1 \cup C_2) \setminus \{r_1, r_2\}$. Since

$$\begin{aligned}C_2 \setminus \{r_2\} &\subseteq (C_1 \cup C_2) \setminus \{r_2\} \\ &= ((C_1 \cup C_2) \setminus \{r_1, r_2\}) \cup \{r_1\} \\ &= X \cup \{r_1\},\end{aligned}$$

we have $r_2 \in \tau_{\mathcal{L}}(C_2 \setminus \{r_2\}) \subseteq \tau_{\mathcal{L}}(X \cup \{r_1\})$ by Lemma 2.7 and the monotonicity (T4) of $\tau_{\mathcal{L}}$. Similarly, we have $r_1 \in \tau_{\mathcal{L}}(C_1 \setminus \{r_1\}) \subseteq \tau_{\mathcal{L}}(X \cup \{r_2\})$. Therefore by the anti-exchange property (T5), we obtain $r_1 \in \tau_{\mathcal{L}}(X)$ or $r_2 \in \tau_{\mathcal{L}}(X)$.

If $r_1 \in \tau_{\mathcal{L}}(X)$, then it follows that

$$C_2 \setminus \{r_2\} \subseteq X \cup \{r_1\} \subseteq \tau_{\mathcal{L}}(X) \cup \{r_1\} = \tau_{\mathcal{L}}(X),$$

where the second inclusion is due to the monotonicity (T4), and in the last identity we use the assumption $r_1 \in \tau_{\mathcal{L}}(X)$. Therefore, it should hold that

$$r_2 \in \tau_{\mathcal{L}}(C_2 \setminus \{r_2\}) \subseteq \tau_{\mathcal{L}}(\tau_{\mathcal{L}}(X)) = \tau_{\mathcal{L}}(X)$$

by Lemma 2.7, the monotonicity (T4) again and the idempotence (T3). Hence in either case we have $r_2 \in \tau_{\mathcal{L}}(X) \setminus X$. Then, by Lemma 2.8, there exists $C_3 \subseteq X \cup \{r_2\}$ such that (C_3, r_2) is a rooted circuit of $\mathcal{L}$. □

Let us note that (C1) and (C2) in Lemma 2.11 actually characterize the family of rooted circuits of a convex geometry among families of rooted subsets. That is, a given family $\mathcal{C}$ of rooted subsets of E satisfies (C1) and (C2) if and only if $\mathcal{C}$ is the family of rooted circuits of some convex geometry on E. This characterization is due to Dietrich [Die87, Die89].

Here, we observe the relation of a rooted circuit and the closure operator.

Lemma 2.12. *Let $\mathcal{L}$ be a convex geometry on E. Then $(C, r) \in \mathcal{C}(\mathcal{L})$ if and only if $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$ and $r \notin \tau_{\mathcal{L}}(D \setminus \{r\})$ for every proper subset $D \subsetneq C$.*

Proof. Assume that $(C, r) \in \mathcal{C}(\mathcal{L})$. From Lemma 2.7 it follows that $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$. Now we show that $r \notin \tau_{\mathcal{L}}(D \setminus \{r\})$ for every proper subset $D \subsetneq C$. Choose a proper subset $D \subsetneq C$ arbitrarily. Then Lemma 2.6 tells us $\tau_{\mathsf{Tr}(\mathcal{L},C)}(D \setminus \{r\}) = \tau_{\mathcal{L}}(D \setminus \{r\}) \cap C$. Since (C, r) is a rooted circuit of $\mathcal{L}$, we have $D \setminus \{r\} \in \mathsf{Tr}(\mathcal{L}, C)$, which implies $\tau_{\mathsf{Tr}(\mathcal{L},C)}(D \setminus \{r\}) = D \setminus \{r\}$ by (T1) in Lemma 2.5. Therefore, we have $\tau_{\mathcal{L}}(D \setminus \{r\}) \cap C = D \setminus \{r\}$. Since $r \notin D \setminus \{r\}$ and $r \in C$, it follows that $r \notin \tau_{\mathcal{L}}(D \setminus \{r\})$. The only-if-part has been proven.

Next, we prove that if $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$ and $r \notin \tau_{\mathcal{L}}(D \setminus \{r\})$ for any proper subset $D \subsetneq C$, then $(C, r) \in \mathcal{C}(\mathcal{L})$. Since $r \in \tau_{\mathcal{L}}(C \setminus \{r\})$ (by the assumption) and $r \notin C \setminus \{r\}$, we have $r \in \tau_{\mathcal{L}}(C \setminus \{r\}) \setminus (C \setminus \{r\})$. Therefore, by Lemma 2.8, there exists $C' \subseteq (C \setminus \{r\}) \cup \{r\} = C$ such that $(C', r) \in \mathcal{C}(\mathcal{L})$. By Lemma 2.7, we have $r \in \tau_{\mathcal{L}}(C' \setminus \{r\})$. Since we have assumed that $r \notin \tau_{\mathcal{L}}(D \setminus \{r\})$ for any proper subset $D \subsetneq C$, it should hold that $C' = C$. This implies that $(C, r) \in \mathcal{C}(\mathcal{L})$. $\square$

Now, we determine the closure operator of a generalized convex shelling.

Lemma 2.13. *Let P and Q be finite point sets in $\mathbb{R}^d$ such that $P \cap \mathsf{conv}(Q) = \emptyset$, and $\mathcal{L}$ be the generalized convex shelling on P with respect to Q. Then*

$$\tau_{\mathcal{L}}(A) = \mathsf{conv}(A \cup Q) \cap P$$

for $A \subseteq P$.

To prove Lemma 2.13, we use the following lemma.

Lemma 2.14. *Let $\mathcal{L}$ be a convex geometry on E, and $S \subseteq E$. Consider the minor $\mathcal{L}' := \mathcal{L}[S, E]$. Then, it holds that $\tau_{\mathcal{L}'}(T) = \tau_{\mathcal{L}}(T \cup S) \setminus S$ for each $T \subseteq E \setminus S$.*

Proof. From the definitions of the closure operator and a minor, it holds that

$$\begin{aligned}
\tau_{\mathcal{L}'}(T) &= \bigcap \{X \in \mathcal{L}' \mid T \subseteq X\} \\
&= \bigcap \{X \subseteq E \mid X \cup S \in \mathcal{L}, T \subseteq X\} \\
&= \bigcap \{Y \setminus S \mid Y \in \mathcal{L}, T \cup S \subseteq Y\} \\
&= \left(\bigcap \{Y \in \mathcal{L} \mid T \cup S \subseteq Y\}\right) \setminus S \\
&= \tau_{\mathcal{L}}(T \cup S) \setminus S.
\end{aligned}$$

In the third identity, we replaced $X \cup S$ by Y. $\square$

Proof of Lemma 2.13. First observe that the closure operator $\tau_{\mathcal{L}^*}$ of the convex shelling $\mathcal{L}^*$ on $P \cup Q$ is given by $\tau_{\mathcal{L}^*}(B) = \mathsf{conv}(B) \cap (P \cup Q)$

for each $B \subseteq P \cup Q$. From Lemma 2.2, the generalized convex shelling $\mathcal{L}$ on P with respect to Q is the same as $\mathcal{L}^*[Q, P \cup Q]$. Therefore, from Lemma 2.14, we obtain

$$\begin{aligned}\tau_{\mathcal{L}}(A) &= \tau_{\mathcal{L}^*}(A \cup Q) \setminus Q \\ &= (\mathsf{conv}(A \cup Q) \cap (P \cup Q)) \setminus Q \\ &= \mathsf{conv}(A \cup Q) \cap P.\end{aligned}$$

This concludes the proof. □

Combining Lemmas 2.12 and 2.13, we can obtain a characterization of the family of rooted circuits of a generalized convex shelling.

Lemma 2.15. *Let $\mathcal{L}$ denote the generalized convex shelling on P with respect to Q, and let $C \subseteq P$ and $r \in C$. Then $(C, r) \in \mathcal{C}(\mathcal{L})$ if and only if $r \in \mathsf{conv}((C \setminus \{r\}) \cup Q)$ and $r \notin \mathsf{conv}((D \setminus \{r\}) \cup Q)$ for any proper subset $D \subsetneq C$.*

Proof. This is a direct consequence of Lemmas 2.12 and 2.13. □

2.5 Proof of the Main Theorem

As explained in Section 2.3, for a given convex geometry $\mathcal{L}$ on E, we construct point sets P_0 and Q_0. We denote by $\mathcal{L}'$ the generalized convex shelling on P_0 with respect to Q_0.

First we have to check that P_0 and Q_0 satisfy the precondition of a generalized convex shelling, namely $P_0 \cap \mathsf{conv}(Q_0) = \emptyset$.

Lemma 2.16. *For P_0 and Q_0 constructed in Section 2.3, it holds that $\mathsf{conv}(P_0) \cap \mathsf{conv}(Q_0) = \emptyset$; in particular, $P_o \cap \mathsf{conv}(Q_0) = \emptyset$.*

To show Lemma 2.16, the next fact is useful, which will be used later again and again.

Lemma 2.17. *Let V be a set of affinely independent points in $\mathbb{R}^d$ and $V_1, V_2 \subseteq V$. If there exist sets $\{\alpha_v \in \mathbb{R}_{>0} \mid v \in V_1\}$ and $\{\beta_v \in \mathbb{R}_{>0} \mid v \in V_2\}$ of positive* *numbers such that*

$$\sum_{v \in V_1} \alpha_v = \sum_{v \in V_2} \beta_v \quad \text{and} \quad \sum_{v \in V_1} \alpha_v v = \sum_{v \in V_2} \beta_v v,$$

then it holds that $V_1 = V_2$.

Proof. Compute as

$$\begin{aligned} \mathbf{0} &= \sum_{v \in V_1} \alpha_v v - \sum_{v \in V_2} \beta_v v \\ &= \sum_{v \in V_1 \cap V_2} (\alpha_v - \beta_v) v + \sum_{v \in V_1 \setminus V_2} \alpha_v v - \sum_{v \in V_2 \setminus V_1} \beta_v v. \end{aligned}$$

Since V consists of affinely independent points, so does $V_1 \cup V_2$. The affine independence of the points in $V_1 \cup V_2$ and our assumption that $\sum\{\alpha_v \mid v \in V_1\} = \sum\{\beta_v \mid v \in V_2\}$ imply that

- $\alpha_v - \beta_v = 0$ for $v \in V_1 \cap V_2$,
- $\alpha_v = 0$ for $v \in V_1 \setminus V_2$,
- $\beta_v = 0$ for $v \in V_2 \setminus V_1$.

Since $\alpha_v > 0$ for $v \in V_1$ and $\beta_v > 0$ for $v \in V_2$, this is possible only if $V_1 = V_2$. □

Now we are ready to show Lemma 2.16 with Lemma 2.17.

Proof of Lemma 2.16. Let $E' \subseteq E$ and $\{(C_1, r_1), \ldots, (C_k, r_k)\} \subseteq \mathcal{C}(\mathcal{L})$ be minimal sets such that the convex hull of the points from $\{p(e) \mid e \in E'\}$ intersects $\mathsf{conv}(\{q(C_i, r_i) \mid i = 1, \ldots, k\})$. By the minimality, there exist some positive numbers $\lambda_e \in \mathbb{R}_{>0}$ for every $e \in E'$ and $\mu_1, \ldots, \mu_k \in \mathbb{R}_{>0}$ such that

$$\sum_{e \in E'} \lambda_e = 1, \quad \sum_{i=1}^{k} \mu_i = 1 \quad \text{and} \quad \sum_{e \in E'} \lambda_e p(e) = \sum_{i=1}^{k} \mu_i q(C_i, r_i).$$

By the construction of Q_0, we have

$$\sum_{e \in E'} \lambda_e p(e) + \sum_{i=1}^{k} \mu_i \left(\sum_{e \in C_i \setminus \{r_i\}} p(e) \right) = \sum_{i=1}^{k} \mu_i |C_i| p(r_i).$$

Since the points in $\{p(e) \mid e \in E\}$ are affinely independent, using Lemma 2.17 we obtain

$$E' \cup \bigcup_{i=1}^{k} (C_i \setminus \{r_i\}) = \{r_i \mid i \in \{1,\dots,k\}\}.$$

Let us denote $R = \{r_i \mid i \in \{1,\dots,k\}\}$. Then the identity above implies that

$$R = R \cup \bigcup_{i=1}^{k} (C_i \setminus \{r_i\}) = \bigcup_{i=1}^{k} C_i. \tag{2.4}$$

By the conditions (L1) and (L3) in the definition of a convex geometry, there exists a subfamily $\{X_j \mid j = 0,1,\dots,n\} \subseteq \mathcal{L}$ such that $X_0 \subsetneq X_1 \subsetneq \cdots \subsetneq X_n$ and $|X_j| = j$ for each $j \in \{0,1,\dots,n\}$. Especially, $X_0 = \emptyset$ and $X_n = E$. Fix such a subfamily $\{X_j \mid j \in \{0,\dots,n\}\}$. Then there exists an index j^* such that $|(E \setminus X_{j^*}) \cap R| = 1$. Let $\{r\} := (E \setminus X_{j^*}) \cap R$. From the identity (2.4), there exists a rooted circuit $(C,r) \in \mathcal{C}(\mathcal{L})$ such that $C \subseteq R$ since $r \in R$. Then it follows that $(E \setminus X_{j^*}) \cap C = \{r\}$. However this implies that $X_{j^*} \notin \mathcal{L}$ by Lemma 2.9, which is a contradiction. $\square$

Lemma 2.16 tells us that $\mathcal{L}_0$ is well-defined. In order to prove Lemma 2.4, we only have to show that $\mathcal{C}(\mathcal{L})$ is isomorphic to $\mathcal{C}(\mathcal{L}_0)$ thanks to Lemma 2.9. Namely, we want a bijection $\psi : E \to P_0$ such that $(\psi(C),\psi(r)) \in \mathcal{C}(\mathcal{L}_0)$ if and only if $(C,r) \in \mathcal{C}(\mathcal{L})$. In our case, the natural bijection $\psi : E \to P_0$ is as follows: $\psi(e) = p(e)$ for $e \in E$. Thus we only have to show the next lemma.

Lemma 2.18. *In the setting above, it holds that*

$$\mathcal{C}(\mathcal{L}_0) = \{(\psi(C),\psi(r)) \mid (C,r) \in \mathcal{C}(\mathcal{L})\}.$$

This lemma follows from the following two lemmas (Lemmas 2.19 and 2.20) and (C1) in Lemma 2.11.

Lemma 2.19. *In the setting above, for every rooted circuit $(C,r) \in \mathcal{C}(\mathcal{L})$, there exists $(C_0,r_0) \in \mathcal{C}(\mathcal{L}_0)$ such that $C_0 \subseteq \psi(C)$ and $r_0 = \psi(r)$.*

Lemma 2.20. *In the setting above, for every rooted circuit $(C_0,r_0) \in \mathcal{C}(\mathcal{L}_0)$, there exists $(C,r) \in \mathcal{C}(\mathcal{L})$ such that $C \subseteq \psi^{-1}(C_0)$ and $r = \psi^{-1}(r_0)$.*

Before proving Lemmas 2.19 and 2.20, let us show how Lemma 2.18 can be derived from them.

Proof of Lemma 2.18. First we prove that if $(C, r) \in \mathcal{C}(\mathcal{L})$ then $(\psi(C), \psi(r)) \in \mathcal{C}(\mathcal{L}_0)$. Take an arbitrary rooted circuit $(C, r) \in \mathcal{C}(\mathcal{L})$. Then from Lemma 2.19, there exists some $(C_0, r_0) \in \mathcal{C}(\mathcal{L}_0)$ such that $C_0 \subseteq \psi(C)$ and $r_0 = \psi(r)$. Note that $r = \psi^{-1}(r_0)$ since ψ is a bijection. Then from Lemma 2.20, there exists some $(\tilde{C}, \tilde{r}) \in \mathcal{C}(\mathcal{L})$ such that $\tilde{C} \subseteq \psi^{-1}(C_0)$ and $\tilde{r} = \psi^{-1}(r_0)$. So we have $r = \psi^{-1}(r_0) = \tilde{r}$.

Now using (C1) in Lemma 2.11, we have

$$(C, r) = (\psi^{-1}(C_0), \psi^{-1}(r_0)) = (\tilde{C}, \tilde{r}).$$

Since ψ is a bijection, we also have

$$(\psi(C), \psi(r)) = (C_0, r_0) = (\psi(\tilde{C}), \psi(\tilde{r})).$$

Therefore, we have $(\psi(C), \psi(r)) \in \mathcal{C}(\mathcal{L}_0)$ since $(C_0, r_0) \in \mathcal{C}(\mathcal{L}_0)$.

Similarly, we can show that if $(C_0, r_0) \in \mathcal{C}(\mathcal{L}_0)$ then it holds that $(\psi^{-1}(C_0), \psi^{-1}(r_0)) \in \mathcal{C}(\mathcal{L})$. □

To prove Lemma 2.19, we use Lemma 2.15.

Proof of Lemma 2.19. Take an arbitrary rooted circuit $(C, r) \in \mathcal{C}(\mathcal{L})$. From our construction, we have

$$p(r) \in \mathsf{conv}(\{p(e) \mid e \in C \setminus \{r\}\} \cup \{q(C, r)\}),$$

which implies

$$\psi(r) \in \mathsf{conv}(\psi(C \setminus \{r\}) \cup Q_0).$$

Take a subset $C_0 \subseteq \psi(C)$ such that $\psi(r) \in \mathsf{conv}((C_0 \setminus \{\psi(r)\}) \cup Q_0)$ and $\psi(r) \notin \mathsf{conv}((D_0 \setminus \{\psi(r)\}) \cup Q_0)$ for any proper subset $D_0 \subsetneq C_0$. (Note that such a set C_0 exists because if $\psi(r) \in A \subseteq B$ and $\psi(r) \in \mathsf{conv}((A \setminus \{\psi(r)\}) \cup Q_0)$ then $\psi(r) \in \mathsf{conv}((B \setminus \{\psi(r)\}) \cup Q_0)$.) From Lemma 2.15, it follows that $(C_0, \psi(r)) \in \mathcal{C}(\mathcal{L}_0)$. □

In order to prove Lemma 2.20, we prepare another lemma.

Lemma 2.21. *In the setting above, let $\overline{e} \in E$ be an element which satisfies the following condition: there exist a subset $F \subseteq E \setminus \{\overline{e}\}$ and some $(C_1,r_1),\ldots,(C_k,r_k) \in \mathcal{C}(\mathcal{L})$ such that $p(\overline{e})$ is a convex combination of the points from $\{p(f) \mid f \in F\} \cup \{q(C_i,r_i) \mid i \in \{1,\ldots,k\}\}$ with all positive coefficients.*

(1) It holds that

$$F \cup \{r_i \mid i \in \{1,\ldots,k\}\} = \{\overline{e}\} \cup \bigcup_{i=1}^{k} (C_i \setminus \{r_i\}).$$

(2) It holds that $\overline{e} \in \tau_{\mathcal{L}}(F)$.

Proof. Let us first prove (1). From the assumption, there exist some $\{\lambda_f \in \mathbb{R}_{>0} \mid f \in F\}$ and $\{\mu_i \in \mathbb{R}_{>0} \mid i \in \{1,\ldots,k\}\}$ such that

$$\sum_{f\in F} \lambda_f + \sum_{i=1}^{k} \mu_i = 1 \quad \text{and} \quad \sum_{f \in F} \lambda_f p(f) + \sum_{i=1}^{k} \mu_i q(C_i,r_i) = p(\overline{e}).$$

From the construction of Q_0, it follows that

$$p(\overline{e}) = \sum_{f\in F} \lambda_f p(f) + \sum_{i=1}^{k} \mu_i \left(|C_i| p(r_i) - \sum_{e \in C_i \setminus \{r_i\}} p(e) \right),$$

meaning that

$$\sum_{f\in F} \lambda_f p(f) + \sum_{i=1}^{k} \mu_i |C_i| p(r_i) = p(\overline{e}) + \sum_{i=1}^{k} \sum_{e \in C_i \setminus \{r_i\}} p(e).$$

By Lemma 2.17, it holds that

$$F \cup \{r_i \mid i = 1,\ldots,k\} = \{\overline{e}\} \cup \bigcup_{i=1}^{k} (C_i \setminus \{r_i\}).$$

Thus, part (1) is proven.

For part (2), set $R := \{r_i \mid i \in \{1,\ldots,k\}\}$ and

$$F^* := \left(\left(\bigcup_{i=1}^{k} (C_i \setminus \{r_i\}) \right) \cup \{\overline{e}\} \right) \setminus R.$$

By part (1) of this lemma, we have $F^* \subseteq F$. Moreover, by part (1) again, it follows that $\overline{e} \in R$. Therefore, F^* can be represented as

$$F^* = \left(\bigcup_{i=1}^{k} C_i \right) \setminus R.$$

We claim that for every $X \in \mathcal{L}$ satisfying $F^* \subseteq X$ it holds that $\overline{e} \in X$. To show that by a contradiction, we suppose that there exists $X^* \in \mathcal{L}$ such that $F^* \subseteq X^*$ and $\overline{e} \notin X^*$. Since $\overline{e} \in R$ and $\overline{e} \notin X^*$, we have $\overline{e} \in (E \setminus X^*) \cap R$. This implies that $|(E \setminus X^*) \cap R| \geq 1$. So by the chain argument, there exists $Z \in \mathcal{L}$ such that $|(E \setminus Z) \cap R| = 1$ and $Z \supseteq X^*$. Without loss of generality, let us say $\{r_1\} = (E \setminus Z) \cap R$. Since $F^* \subseteq X^* \subseteq Z$ we have $(E \setminus Z) \cap F^* = \emptyset$. Therefore, it follows that

$$\begin{aligned}
(E \setminus Z) \cap \left(\bigcup \{C_i \mid i = 1, \ldots, k\} \right) &= (E \setminus Z) \cap (F^* \cup R) \\
&= ((E \setminus Z) \cap F^*) \cup ((E \setminus Z) \cap R) \\
&= \emptyset \cup \{r_1\} \\
&= \{r_1\}.
\end{aligned}$$

Then we obtain $(E \setminus Z) \cap C_1 = \{r_1\}$. However this implies that $Z \notin \mathcal{L}$, together with Lemma 2.9. This is a contradiction.

Let us consider $\tau_{\mathcal{L}}(F^*)$. Since $F^* \subseteq \tau_{\mathcal{L}}(F^*) \in \mathcal{L}$ (the extensionality of $\tau_{\mathcal{L}}$), it holds that $\overline{e} \in \tau_{\mathcal{L}}(F^*)$. (Here, we have used the claim above.) By the monotonicity (T4) of $\tau_{\mathcal{L}}$ we obtain $\tau_{\mathcal{L}}(F^*) \subseteq \tau_{\mathcal{L}}(F)$. From this we conclude that $\overline{e} \in \tau_{\mathcal{L}}(F)$. □

Now we are ready to prove Lemma 2.20.

Proof of Lemma 2.20. Let $(C_0, r_0) \in \mathcal{C}(\mathcal{L}_0)$. From Lemma 2.15, we can see that $r_0 \in \mathsf{conv}((C_0 \setminus \{r_0\}) \cup Q_0)$ and $r_0 \notin \mathsf{conv}((D_0 \setminus \{r_0\}) \cup Q_0)$ for any proper subset $D_0 \subsetneq C_0$. Let us observe the following.

Claim 2.21.1. *There exists some subset $Q_1 \subseteq Q_0$ such that r_0 is a convex combination of the points from $\mathsf{conv}((C_0 \setminus \{r_0\}) \cup Q_1)$ with all positive coefficients.*

Proof of Claim 2.21.1. To see this, suppose contrarily that there exists no such set. Namely, for any subset $Q_1 \subseteq Q_0$, every convex combination of the points from $(C_0 \setminus \{r_0\}) \cup Q_1$ representing r_0 has a term with a zero coefficient. Take $Q_1 = \emptyset$. Then this particularly implies that if we write r_0 as

$$r_0 = \sum_{p \in C_0 \setminus \{r_0\}} \lambda_p p$$

for some non-negative real numbers λ_p, $p \in C_0 \setminus \{r_0\}$, and let

$$F := \{p \in C_0 \setminus \{r_0\} \mid \lambda_p > 0\},$$

then F is a proper subset of $C_0 \setminus \{r_0\}$. It holds that $r_0 \in \mathsf{conv}(F) \subseteq \mathsf{conv}(F \cup Q_0)$. However, this contradicts the assumption that $r_0 \notin \mathsf{conv}((D_0 \setminus \{r_0\}) \cup Q_0)$ for any proper subset $D_0 \subsetneq C_0$. The claim is proven. $\square$

Using Claim 2.21.1 together with Lemma 2.21(2), we obtain $\psi^{-1}(r_0) \in \tau_{\mathcal{L}}(\psi^{-1}(C_0 \setminus \{r_0\}))$. Choose $C \subseteq \psi^{-1}(C_0)$ such that $\psi^{-1}(r_0) \in \tau_{\mathcal{L}}(C \setminus \{\psi^{-1}(r_0)\})$ and $\psi^{-1}(r_0) \notin \tau_{\mathcal{L}}(D \setminus \{\psi^{-1}(r_0)\})$ for any proper subset $D \subsetneq C$. (Note that such a set C exists because of the same reason as in the proof of Lemma 2.19.) By Lemma 2.12, it follows that $(C, \psi^{-1}(r_0)) \in \mathcal{C}(\mathcal{L})$. $\square$

This completes the whole proof. Q.E.D.

2.6 Conclusion

In this chapter, we have provided the affine representation theorem for (abstract) convex geometries. This should be as useful as the representation theorem for oriented matroids by Folkman & Lawrence [FL78]. Actually, the theorem has opened several new directions of research. We indicate some of them here.

1. Our theorem makes it possible to talk about the dimension of the space in which a given convex geometry can be realized. Hachimori & Nakamura [HN04] studied stem clutters of a convex geometry which can be realized in the 2-dimensional space. They

gave a characterization of a stem clutter in dimension 2 with the max-flow min-cut property.

2. The next chapter studies an open problem posed by Edelman & Reiner [ER00] from the viewpoint of our theorem. Especially, we solve the question affirmatively for 2-dimensional separable generalized convex shellings. (Here, "separable" means that $\mathsf{conv}(P_0) \cap \mathsf{conv}(Q_0) = \emptyset$. Because of Lemma 2.16, our theorem can be strengthened to: every convex geometry is isomorphic to some separable generalized convex shelling.)

We hope that our theorem gives a fruitful tool in the theory of convex geometries and related fields.

A good proof should be easy to verify — and it should be easy to falsify, in case it is wrong.

Emo Welzl (2003, in an explicitly written form)

Chapter 3

Local Topology of the Free Complex of a Two-Dimensional Generalized Convex Shelling

3.1 Introduction

An Euler-Poincaré type formula for the number of interior points in a d-dimensional point configuration was proved by Ahrens, Gordon & McMahon [AGM99] for $d = 2$, and proved by Edelman & Reiner [ER00] and Klain [Kla99] independently for arbitrary d. The approach by Klain [Kla99] used a more general theorem on valuation, while that by Edelman & Reiner [ER00] was topological. (Edelman, Reiner & Welker [ERW02] gave another proof based on oriented matroids. Pinchasi, Radoičić & Sharir [PRS04] gave yet another proof using elementary geometric arguments for point configurations in general position.) In the paper by Edelman & Reiner [ER00], they studied the topology of deletions of the free complex of a convex shelling (arising from a point configuration), and also mentioned a possible generalization to a convex geometry. More precisely speaking, their open problems are as follows. (The necessary definitions will be given in the following section.)

Open Problem 3.1 (Edelman & Reiner [ER00]). *Let $\mathcal{L}$ be a convex geometry on E and denote the free complex of $\mathcal{L}$ by* $\mathrm{Free}(\mathcal{L})$.

1. *Is the deletion* $\mathrm{del}_{\mathrm{Free}(\mathcal{L})}(x)$ *of an element $x \in E$ contractible if* $\mathrm{Dep}_{\mathcal{L}}(x) \neq E$?
2. *Is* $\mathrm{del}_{\mathrm{Free}(\mathcal{L})}(x)$ *homotopy equivalent to a bouquet of spheres if* $\mathrm{Dep}_{\mathcal{L}}(x) = E$,

Edelman & Reiner [ER00] answered Open Problem 3.1 affirmatively for convex shellings to derive the forementioned Euler-Poincaré type formula, and further for poset double shellings and simplicial shellings of chordal graphs as well. Subsequently, Edelman, Reiner & Welker [ERW02] answered it affirmatively for convex shellings of acyclic oriented matroids.

With the view from the previous chapter, we resolve Open Problem 3.1 for 2-dimensional separable generalized convex shellings.

Theorem 3.2. *Let P and Q be non-empty finite point sets in $\mathbb{R}^2$ such that* $\mathsf{conv}(P) \cap \mathsf{conv}(Q) = \emptyset$. *In addition, let $\mathcal{L}$ be the generalized convex shelling on P with respect to Q. Consider the free complex* $\mathrm{Free}(\mathcal{L})$ *of $\mathcal{L}$. Then the following holds.*

1. *If* $\mathrm{Dep}_{\mathcal{L}}(x) \neq P$, *then the deletion* $\mathrm{del}_{\mathrm{Free}(\mathcal{L})}(x)$ *of an element $x \in P$ is contractible (i.e., homotopy equivalent to a single point).*
2. *If* $\mathrm{Dep}_{\mathcal{L}}(x) = P$, *then* $\mathrm{del}_{\mathrm{Free}(\mathcal{L})}(x)$ *is contractible or homotopy equivalent to a 0-dimensional sphere (i.e., two distinct points).*

Theorem 3.2 settles the problem for the special case of 2-dimensional separable generalized convex shellings. However, our case is not just a special case. Thanks to Theorem 2.3 of the previous chapter, every convex geometry is isomorphic to some separable generalized convex shelling in some dimension. Therefore, our result is a step toward the solution of Open Problem 3.1.

The organization of this chapter is as follows. In the next section we introduce the necessary terminology about simplicial complexes and convex geometries. Section 3 sketches the proof of our theorem. We conclude the chapter in Section 4 with some examples.

3.2 Preliminaries

In this chapter, we assume a moderate familiarity with graphs and convex geometries. You can consult Chapter 1 for graphs, and Chapter 2 for convex geometries.

3.2.1 Simplicial Complexes

Let E be a finite set. An *abstract simplicial complex* on E is a non-empty family Δ of subsets of E satisfying that: if $X \in \Delta$ and $Y \subseteq X$ then $Y \in \Delta$. Often an abstract simplicial complex is just called a simplicial complex and this is nothing but an independence system as it appeared in Chapter 1. However, since we are talking about topology we use the term "abstract simplicial complex" which is more usual in that field. For a simplicial complex Δ on E, a subset of E is called a *face* of the simplicial complex Δ if it belongs to Δ; if not it is called a *non-face*. (Namely, in the terminology of independence systems, a face is an independent set and a non-face is a dependent set.)

For a simplicial complex Δ on E and an element $x \in E$, the *deletion* of x in Δ is defined by $\mathrm{del}_\Delta(x) := \{X \in \Delta \mid x \notin X\}$. Note that the deletion is a simplicial complex on $E \setminus \{x\}$.

When we talk about topology of a simplicial complex, we refer to a geometric realization of the simplicial complex. A *d-dimensional simplex* is the convex hull of $d{+}1$ affinely independent points. Conventionally, the empty set is considered a (-1)-dimensional simplex. Let Δ be a simplicial complex on E. A *geometric realization* of Δ is a collection C of simplices satisfying the following condition: there exists a mapping $\psi : E \to \mathrm{I\!R}^d$ for some natural number d such that the convex hull of the image $\psi(X)$ is an $(\ell{-}1)$-dimensional simplex of C for each $X \in \Delta$ of size ℓ, where ℓ is a natural number, and every two simplices in C

do not intersect in their relative interiors. It is known that for every simplicial complex Δ such that the faces have at most k elements there exists a geometric realization of Δ in $\mathbb{R}^{2k-1}$, and this bound is tight. Details can be found in Matoušek's book [Mat03] for example.

Our topological investigation is restricted to the Euclidean case. So we just define some terms within the Euclidean space. Let X and Y be sets in $\mathbb{R}^d$. Two continuous maps $f_0, f_1 : X \to Y$ are *homotopic* if there exists a continuous map $F : X \times [0,1] \to Y$ such that $F(x,0) = f_0(x)$ and $F(x,1) = f_1(x)$ for all $x \in X$. Two sets $X, Y \subseteq \mathbb{R}^d$ are *homotopy equivalent* if there exist two continuous maps $f : X \to Y$ and $g : Y \to X$ such that the composition $f \circ g : Y \to Y$ and the identity map $\mathrm{id}_Y : Y \to Y$ are homotopic and also the composition $g \circ f : X \to X$ and the identity map $\mathrm{id}_X : X \to X$ are homotopic. Intuitively speaking, X and Y are homotopy equivalent if one of them can be deformed continuously to the other.

3.2.2 Convex Geometries

Here we introduce terminology for convex geometries which have not appeared in Chapter 2.

Let E be a non-empty finite set and $\mathcal{L}$ be a convex geometry on E. For a set $A \subseteq E$, an element $e \in A$ is called an *extreme point* if $e \notin \tau_{\mathcal{L}}(A \setminus \{e\})$. We denote the set of extreme points of A by $\mathrm{ex}_{\mathcal{L}}(A)$. Namely, define the operator $\mathrm{ex}_{\mathcal{L}} : 2^E \to 2^E$ as

$$\mathrm{ex}_{\mathcal{L}}(A) := \{e \in A \mid e \text{ is an extreme point of } A \text{ in } \mathcal{L}\}.$$

We call $\mathrm{ex}_{\mathcal{L}}$ the *extreme point operator* of $\mathcal{L}$. Notice that the extreme point operator $\mathrm{ex}_{\mathcal{L}}$ of a convex geometry $\mathcal{L}$ on E satisfies the following properties:

Intensionality: $\mathrm{ex}_{\mathcal{L}}(A) \subseteq A$ for all $A \subseteq E$,

which is clear from the definition. Ando [And02] gives a detailed treatment on closure operators and extreme point operators in a more general setting.

A set $A \subseteq E$ is called *independent* if $\mathrm{ex}_{\mathcal{L}}(A) = A$. We say that e *depends on* f if there exists an independent set A such that $f \in A$,

$e \in \tau_{\mathcal{L}}(A)$ and $e \notin \tau_{\mathcal{L}}(A \setminus \{f\})$. We denote the set of all elements on which e depends by $\mathrm{Dep}_{\mathcal{L}}(e)$. A set $X \subseteq E$ is called *free* if $X \in \mathcal{L}$ and $\mathrm{ex}_{\mathcal{L}}(X) = X$. We denote the family of free sets of a convex geometry $\mathcal{L}$ by $\mathrm{Free}(\mathcal{L})$. Note that $\mathrm{Free}(\mathcal{L})$ forms a simplicial complex for any convex geometry $\mathcal{L}$, as shown in the next lemma.

Lemma 3.3. *Let $\mathcal{L}$ be a convex geometry on E. Then* $\mathrm{Free}(\mathcal{L})$ *is a simplicial complex on E.*

Proof. Let $X \in \mathrm{Free}(\mathcal{L})$ and $Y \subseteq X$. We want to show that

$$\mathrm{ex}_{\mathcal{L}}(Y) = Y \text{ and} \tag{3.1}$$
$$\tau_{\mathcal{L}}(Y) = Y. \tag{3.2}$$

First we prove Equality (3.1). By the intensionality of $\mathrm{ex}_{\mathcal{L}}(Y)$, it holds that $\mathrm{ex}_{\mathcal{L}}(Y) \subseteq Y$. So, we only have to prove that $Y \subseteq \mathrm{ex}_{\mathcal{L}}(Y)$. Choose $e \in Y$ arbitrarily. Since $X = \mathrm{ex}_{\mathcal{L}}(X)$ and e also belongs to X, e is an extreme point of X. Namely, we have $e \notin \tau_{\mathcal{L}}(X \setminus \{e\})$. By the monotonicity (T4) of $\tau_{\mathcal{L}}$, it follows that $\tau_{\mathcal{L}}(Y \setminus \{e\}) \subseteq \tau_{\mathcal{L}}(X \setminus \{e\})$. Therefore, it concludes that $e \notin \tau_{\mathcal{L}}(Y \setminus \{e\})$. This shows that $Y \subseteq \mathrm{ex}_{\mathcal{L}}(Y)$.

Next, we prove Equality (3.2). To prove it, we only have to show that $\tau_{\mathcal{L}}(X \setminus \{e\}) = X \setminus \{e\}$ for all $e \in X$. That is because, from Condition (2) in the definition of convex geometries, we have $Y = \bigcap_{e \in X \setminus Y} X \setminus \{e\} = \bigcap_{e \in X \setminus Y} \tau_{\mathcal{L}}(X \setminus \{e\}) \in \mathcal{L}$, which means $\tau_{\mathcal{L}}(Y) = Y$.

Fix an arbitrary element $e \in X$. By the extensionality (T2), we have $X \setminus \{e\} \subseteq \tau_{\mathcal{L}}(X \setminus \{e\})$. By the monotonicity (T4) and the assumption that $\tau_{\mathcal{L}}(X) = X$, we also have $\tau_{\mathcal{L}}(X \setminus \{e\}) \subseteq \tau_{\mathcal{L}}(X) = X$. Therefore, $\tau_{\mathcal{L}}(X \setminus \{e\})$ is either $X \setminus \{e\}$ or X. However, since e is an extreme point of X, we have $e \notin \tau_{\mathcal{L}}(X \setminus \{e\})$. This concludes that $\tau_{\mathcal{L}}(X \setminus \{e\}) = X \setminus \{e\}$. Hence, we have $\tau_{\mathcal{L}}(X \setminus \{e\}) = X \setminus \{e\}$. □

Thus, it is natural that we call $\mathrm{Free}(\mathcal{L})$ the *free complex* of a convex geometry $\mathcal{L}$. Note that in general there might exist an element $x \in E$ such that $\{x\} \notin \mathrm{Free}(\mathcal{L})$.

Let us remind the definition of a generalized convex shelling. Let P and Q be finite point sets in $\mathbb{R}^d$ (where d is a positive integer) such that $P \cap \mathsf{conv}(Q) = \emptyset$. Then the *generalized convex shelling* on P with respect to Q is a convex geometry $\mathcal{L}$ defined as follows: $\mathcal{L} := \{X \subseteq P \mid P \cap \mathsf{conv}(X \cup Q) = X\}$. We also call a convex geometry $\mathcal{L}$ a *d-dimensional generalized convex shelling* if there exist finite point sets P and Q in $\mathbb{R}^d$ such that $P \cap \mathsf{conv}(Q) = \emptyset$ and $\mathcal{L}$ is isomorphic to the generalized convex shelling on P with respect to Q. The next lemma tells us the closure operator and the extreme point operator of a generalized convex shelling.

Lemma 3.4. *Let $\mathcal{L}$ be a generalized convex shelling on P with respect to Q. Then, we have*

$$\begin{aligned} \tau_{\mathcal{L}}(X) &= P \cap \mathsf{conv}(X \cup Q), \\ \mathsf{ex}_{\mathcal{L}}(X) &= \{x \in X \mid x \textit{ is an extreme point of } \mathsf{conv}(X \cup Q)\} \end{aligned}$$

for each set $X \subseteq P$.[1] *In particular, $X \subseteq P$ is free if and only if $P \cap \mathsf{conv}(X \cup Q) = X$ and every element of X is an extreme point of $\mathsf{conv}(X \cup Q)$.*

Proof. The statement for the closure operator has already been proved as Lemma 2.13. Here, we prove that the extreme point operator is as claimed. The proof is based on the following chain of equivalences.

$$\begin{aligned} p \in \mathsf{ex}_{\mathcal{L}}(X) &\Leftrightarrow p \notin \tau_{\mathcal{L}}(X \setminus \{p\}) \\ &\qquad \text{(by the definition of } \mathsf{ex}_{\mathcal{L}}) \\ &\Leftrightarrow p \notin P \cap \mathsf{conv}((X \setminus \{p\}) \cup Q) \\ &\qquad \text{(from the first part of this lemma)} \\ &\Leftrightarrow p \notin \mathsf{conv}((X \setminus \{p\}) \cup Q) \\ &\qquad \text{(since } p \in P) \\ &\Leftrightarrow p \notin \mathsf{conv}((X \cup Q) \setminus \{p\}) \\ &\Leftrightarrow p \text{ is an extreme point of } \mathsf{conv}(X \cup Q) \\ &\qquad \text{(by the definition of an extreme point).} \end{aligned}$$

The second part is immediate from the first two parts of this lemma and the definition of a free set. ☐

In this chapter, we study the free complex of a 2-dimensional separable generalized convex shelling. Since we already know that Open

[1] Here, you would notice that we are using the phrase "extreme point" in two different meanings. One for an extreme point in a convex geometry, one for an extreme point of the convex hull. But they should be clear from the context.

Problem 3.1 has been solved when $Q = \emptyset$ [ER00], we may make the following assumption, which is important in this chapter.

Assumption 3.5. *When we talk about the generalized convex shelling on P with respect to Q in the rest of this chapter, Q is always non-empty unless stated otherwise.*

3.3 Proof of Theorem 3.2

3.3.1 Basic Properties and the Outline

Now we concentrate on 2-dimensional separable generalized convex shellings. Let P and Q be two non-empty finite point sets in $\mathbb{R}^2$ such that $\mathsf{conv}(P) \cap \mathsf{conv}(Q) = \emptyset$. Denote by $\mathcal{L}$ the generalized convex shelling on P with respect to Q. Since $\mathsf{conv}(P) \cap \mathsf{conv}(Q) = \emptyset$, there exists a line which strictly separates $\mathsf{conv}(P)$ and $\mathsf{conv}(Q)$. Fix such a line, and call it ℓ. In the rest of the chapter, we visualize ℓ as a vertical line, and P is put left to ℓ and Q right to ℓ.

To prove Theorem 3.2, we use the following fact.

Lemma 3.6 (Hachimori & Nakamura [HN04]). *A minimal non-face of the free complex* Free$(\mathcal{L})$ *of a d-dimensional generalized convex shelling is of size at most d.*

Lemma 1.4 from Chapter 1 shows that a simplicial complex whose minimal non-faces are of size 2 is a clique complex of some graph. (Let us remind the definition of a clique complex: the *clique complex* of G is the family of cliques of G.) Therefore, the free complex of a 2-dimensional generalized convex shelling $\mathcal{L}$ is the clique complex of some graph, and this graph is actually the 1-dimensional skeleton of Free$(\mathcal{L})$. (The *d-dimensional skeleton* of a simplicial complex Δ is a collection $\{X \in \Delta \mid |X| \leq d{+}1\}$. Note that a 1-dimensional skeleton can be regarded as a graph.) Denote by $G(\mathcal{L})$ the 1-dimensional skeleton of Free$(\mathcal{L})$ regarded as a graph. The following lemma tells what $G(\mathcal{L})$ is.

Lemma 3.7. *A point $x \in P$ is a vertex of $G(\mathcal{L})$ if and only if $P \cap \mathsf{conv}(\{x\} \cup Q) = \{x\}$ holds, i.e., $\mathsf{conv}(\{x\} \cup Q)$ contains no point of P except for x. Two points $x, y \in P$ form an edge of $G(\mathcal{L})$ if and only if they are vertices of $G(\mathcal{L})$*

and $P \cap \mathsf{conv}(\{x,y\} \cup Q) = \{x,y\}$ holds, i.e., $\mathsf{conv}(\{x,y\} \cup Q)$ contains no point of P except for x,y.

Proof. First of all, notice that $x \in P$ is a vertex of $G(\mathcal{L})$ if and only if $\{x\} \in \mathrm{Free}(\mathcal{L})$, and that $\{x,y\} \subseteq P$ is an edge of $G(\mathcal{L})$ if and only if $\{x,y\} \in \mathrm{Free}(\mathcal{L})$.

Assume that $x \in P$ satisfies $\{x\} \in \mathrm{Free}(\mathcal{L})$. Then, from Lemma 3.4, this is equivalent to saying that $P \cap \mathsf{conv}(\{x\} \cup Q) = \{x\}$ and x is an extreme point of $\mathsf{conv}(\{x\} \cup Q)$. However, x is always an extreme point of $\mathsf{conv}(\{x\} \cup Q)$ since we have the assumption that $P \cap \mathsf{conv}(Q) = \emptyset$. Thus, we have shown that $x \in P$ is a vertex of $G(\mathcal{L})$ if and only if $P \cap \mathsf{conv}(\{x\} \cup Q) = \{x\}$.

For the second part, first choose arbitrary two vertices x,y of $G(\mathcal{L})$. Namely, x and y satisfy the condition in the first part. Now we show that $\{x,y\}$ is an edge of $G(\mathcal{L})$ if and only if $P \cap \mathsf{conv}(\{x,y\} \cup Q) = \{x,y\}$. Assume that $\{x,y\}$ is an edge of $G(\mathcal{L})$. Again, from Lemma 3.4, this is equivalent to saying that $P \cap \mathsf{conv}(\{x,y\} \cup Q) = \{x,y\}$ and x and y are extreme points of $\mathsf{conv}(\{x,y\} \cup Q)$. However, the property that x and y are extreme points of $\mathsf{conv}(\{x,y\} \cup Q)$ can be derived from our assumption that x and y are vertices of $G(\mathcal{L})$. To verify this, suppose that x is not an extreme point of $\mathsf{conv}(\{x,y\} \cup Q)$. This means that $x \in \mathsf{conv}(\{y\} \cup Q)$. However, this implies that y violates the condition that $P \cap \mathsf{conv}(\{y\} \cup Q) = \{y\}$. So this is a contradiction to the first part of this lemma. Thus, we have shown the second part. □

Thanks to Lemma 3.7, we can regard $G(\mathcal{L})$ as a geometric graph. Namely, we can geometrically construct $G(\mathcal{L})$ in the following way. First, we remove a point $x \in P$ if and only if the condition that $P \cap \mathsf{conv}(\{x\} \cup Q) = \{x\}$ is violated. The remaining points from P are the vertices of $G(\mathcal{L})$ (by Lemma 3.7). Among these remaining points, we connect two points $x,y \in P$ by a line segment if and only if $P \cap \mathsf{conv}(\{x,y\} \cup Q) = \{x,y\}$ holds. This process gives the edges of $G(\mathcal{L})$. Figure 3.1 is an example of $G(\mathcal{L})$, where P consists of eight points $1,\ldots,8$ and Q of two points q_1 and q_2. The right one is the resulting geometric graph $G(\mathcal{L})$. The point 2 does not remain in $G(\mathcal{L})$ as a vertex since $P \cap \mathsf{conv}(\{2\} \cup Q) = \{2,5,6\}$.

The rest of the proof is organized in the following way.

1. We prove that $G(\mathcal{L})$ is connected (Lemma 3.8).
2. We prove that $G(\mathcal{L})$ is chordal (Lemma 3.9).
3. We observe that the clique complex of a connected chordal graph is contractible (Lemma 3.10).
4. We show the relation of a cut-vertex of $G(\mathcal{L})$ and a dependency set (Lemmas 3.13 and 3.14).

The rest of the section is divided according to the proof scheme above.

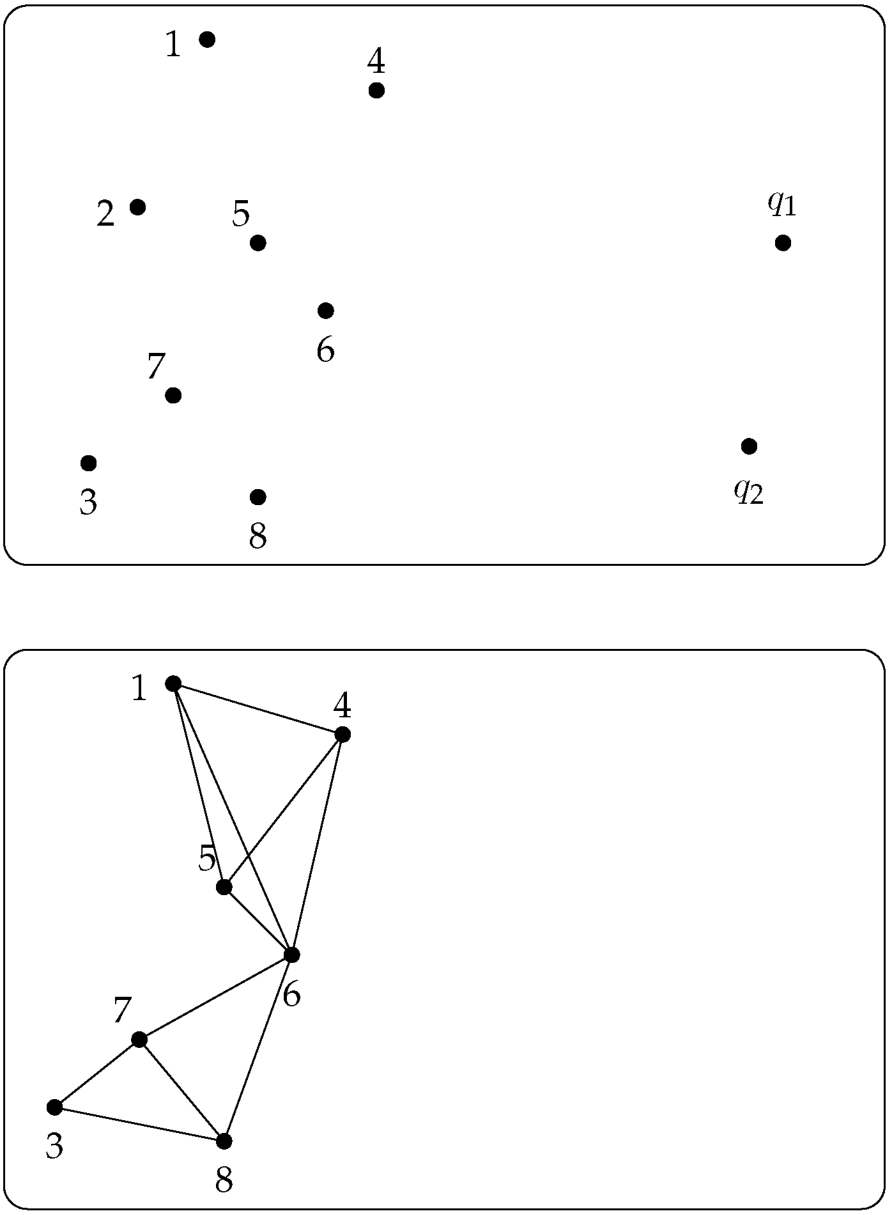

Figure 3.1: (Top) given sets of points. (Bottom) the resulting geometric graph $G(\mathcal{L})$.

3.3.2 Connectedness of the Graph

First, we show the connectedness of $G(\mathcal{L})$.

Lemma 3.8. *In the setup above, $G(\mathcal{L})$ is connected.*

Proof. The proof is done by induction on the number of points in P. When $|P| = 1$, $G(\mathcal{L})$ consists of only one vertex. So $G(\mathcal{L})$ is connected. Fine.

Assume that $|P| > 1$. Let us choose a point v of P which is the furthest from $\mathsf{conv}(Q)$.

Let $P' = P \setminus \{v\}$ and $\mathcal{L}'$ be the generalized convex shelling on P' with respect to Q. We have two cases.

Case 1: v is not a vertex of $G(\mathcal{L})$. In this case, we claim that $G(\mathcal{L}') = G(\mathcal{L})$. First we show that the vertex sets are the same. To show that, suppose not. If $G(\mathcal{L}')$ owns a vertex u which is not a vertex of $G(\mathcal{L})$, then it must hold that $v \in \mathsf{conv}(\{u\} \cup Q)$. However, this means that v is closer to $\mathsf{conv}(Q)$ than u. This contradicts the choice of v. On the other hand, if $G(\mathcal{L})$ owns a vertex w which is not a vertex of $G(\mathcal{L}')$, then there must exist a point $x \in P' \setminus P$ such that $x \in \mathsf{conv}(\{w\} \cup Q)$. However, this is impossible because $P' \subseteq P$, consequently $P' \setminus P = \emptyset$. Thus, the vertex sets of $G(\mathcal{L})$ and $G(\mathcal{L}')$ are the same.

Secondly we show that the edge sets are the same. This can be done in a similar way to the vertex sets. Thus, the claim follows.

By induction hypothesis, $G(\mathcal{L}')$ is connected. Then from the claim above, we conclude that $G(\mathcal{L})$ is connected.

Case 2: v is a vertex of $G(\mathcal{L})$. In this case, we introduce further symbols. Let ℓ be a line supporting $\mathsf{conv}(Q)$ and perpendicular to the line spanned by v and the point in $\mathsf{conv}(Q)$ closest to v. Further, let ℓ_v be a line parallel to ℓ and passing through v. Denote by $\ell_\top$ and $\ell_\perp$ the lines supporting $\mathsf{conv}(\{v\} \cup Q)$ and passing through v. These lines ℓ, ℓ_v, $\ell_\top$ and $\ell_\perp$ are well-defined since $\mathsf{conv}(P) \cap \mathsf{conv}(Q) = \emptyset$. See Figure 3.2. Note that $\ell_\top$ and $\ell_\perp$ coincide when $|Q| = 1$. By an argument similar to

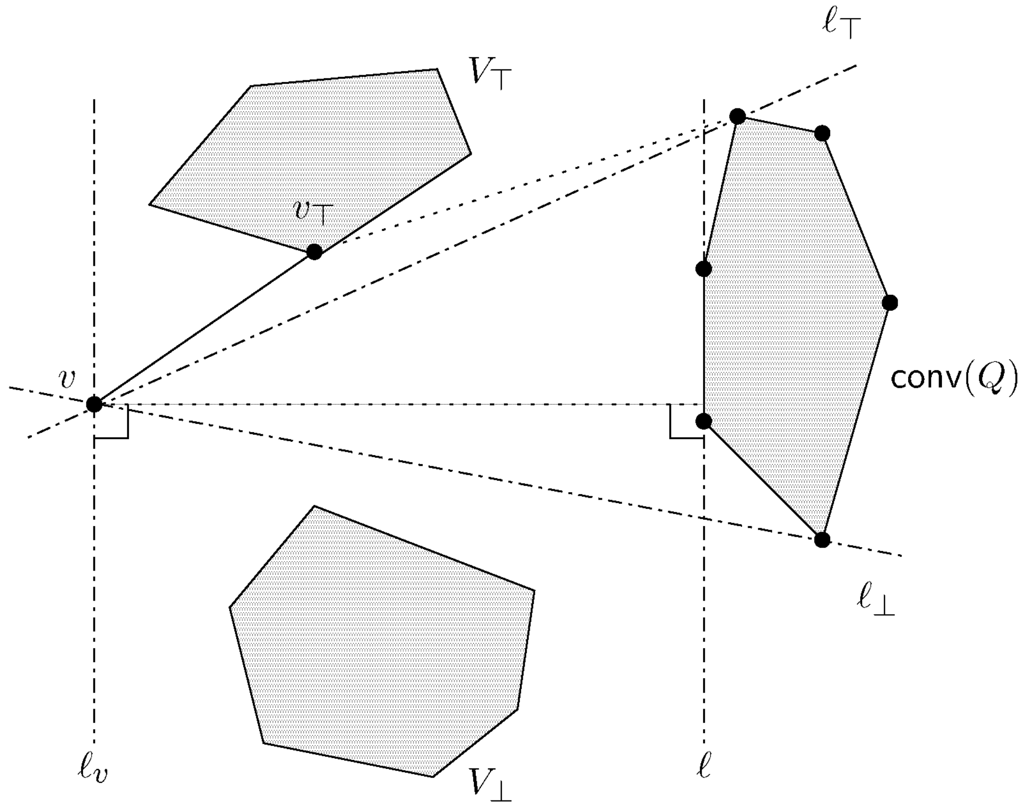

Figure 3.2: v is not an isolated vertex.

the first case, we can observe that $G(\mathcal{L}') = G(\mathcal{L}) - v$.

Now, by the induction hypothesis, $G(\mathcal{L}')$ is connected. Therefore, it suffices to show that v is not an isolated vertex of $G(\mathcal{L})$.

From our choices, the vertices of $G(\mathcal{L})$ other than v should lie either in the space bounded by ℓ_v and $\ell_\top$ or in the space bounded by ℓ_v and $\ell_\perp$. Let $V_\top$ (and $V_\perp$) be the set of vertices of $G(\mathcal{L})$ lying in the former (and latter, respectively) space, as in Figure 3.2. Note that at least one of the two is non-empty since the number of vertices of $G(\mathcal{L})$ is more than one. Assume that $V_\top$ is non-empty, without loss of generality. Then choose a vertex in $V_\top$ which is closest to $\ell_\top$ and name it $v_\top$. We can see that $P \cap \mathsf{conv}(\{v, v_\top\} \cup Q) = \{v, v_\top\}$ because of our choices. This means that $\{v, v_\top\}$ forms an edge in $G(\mathcal{L})$, thus v is not an isolated vertex of $G(\mathcal{L})$. It concludes the whole proof. $\square$

3.3.3 Chordality of the Graph

Next, we show the chordality of $G(\mathcal{L})$. A graph is *chordal* if it has no induced cycle of length more than three.

Lemma 3.9. *In the setup above, $G(\mathcal{L})$ is chordal.*

Proof. Suppose, for the contradiction, that $G(\mathcal{L})$ has an induced cycle of length more than 3. Choose such an induced cycle C arbitrarily, and denote by V_C the set of vertices of C.

The convex hull of V_C and the convex hull of Q have two outer common tangents ℓ_1 and ℓ_2.[2] Choose $v_1 \in V_C \cap \ell_1$ and $v_2 \in V_C \cap \ell_2$ arbitrarily.

We observe that $v_1 \neq v_2$. To show that, suppose not. Then, since ℓ_1 and ℓ_2 are outer common tangents of $\mathsf{conv}(V_C)$ and $\mathsf{conv}(Q)$, all points of V_C must be contained in $\{v_1\} \cup \mathsf{conv}(Q)$. However, this is a contradiction to the fact that v_1 is a vertex of $G(\mathcal{L})$ (remember Lemma 3.7). Therefore, v_1 is distinct from v_2.

Now, we have two cases.

Case 1: $\{v_1, v_2\}$ is an edge of C. In the cycle C, two vertices v_1 and v_2 are joined by two distinct paths. By our assumption, one of them is v_1v_2, namely a path of length one. Let $v_1u_1 \cdots u_kv_2$ be the other path. (Here, a path is denoted by the sequence of consecutive vertices on it.) Since the length of C is more than three, it holds that $k \geq 2$.

Since $\{v_1, v_2\}$ is an edge of $G(\mathcal{L})$, by Lemma 3.7 it follows that $\mathsf{conv}(\{v_1, v_2\} \cup Q)$ contains no point of $P \setminus \{v_1, v_2\}$, in particular none of $\{u_1, \ldots, u_k\}$. Since we chose v_1 and v_2 via the outer common tangents of $\mathsf{conv}(V_C)$ and $\mathsf{conv}(Q)$, this implies that all points of $\{u_1, \ldots, u_k\}$ lie in the region bounded by ℓ_1, ℓ_2 and the line spanned by v_1, v_2. Take a point $u_i \in \{u_1, \ldots, u_k\}$ which is closest to the line segment $\overline{v_1v_2}$. Since $k \geq 2$, at least one of $\{v_1, u_i\}$ and $\{v_2, u_i\}$ is not an edge of $G(\mathcal{L})$. Without loss of generality, assume that $\{v_1, u_i\}$ is not an edge. Since all points of $\{u_1, \ldots, u_k\}$ lie in the region bounded

[2] Here, an *outer common tangent* of two convex sets A and B is a line ℓ which touches A, B and determines a halfplane containing both of A and B.

by ℓ_1, ℓ_2 and the line spanned by v_1, v_2, we have $\mathsf{conv}(\{v_1, v_2\} \cup Q) \subseteq \mathsf{conv}(\{v_1, v_2, u_i\} \cup Q)$. Since $\{v_1, u_i\}$ is not an edge of $G(\mathcal{L})$, by Lemma 3.7 there must exist a point $p \in \mathsf{conv}(\{v_1, u_i\} \cup Q)$. However, $\{v_1, u_1\}, \{u_1, u_2\}, \ldots, \{u_{i-1}, u_i\}$ are edges of $G(\mathcal{L})$ and we have $\mathsf{conv}(\{v_1, u_i\} \cup Q) \subseteq \bigcup_{j=0}^{i-1} \mathsf{conv}(\{u_j, u_{j+1}\} \cup Q)$ by our choices, where u_0 is set to v_1. This means that there exists some index $j \in \{0, \ldots, i-1\}$ such that the set $\mathsf{conv}(\{u_j, u_{j+1}\} \cup Q)$ contains p. Lemma 3.7 implies that $\{u_j, u_{j+1}\}$ is not an edge of $G(\mathcal{L})$. This is a contradiction.

Case 2: $\{v_1, v_2\}$ is not an edge of C. By Lemma 3.7, there must exist a point of $P \setminus \{v_1, v_2\}$ belonging to $\mathsf{conv}(\{v_1, v_2\} \cup Q)$. Let p be the furthest point from the line spanned by v_1 and v_2 among all such points in $P \setminus \{v_1, v_2\}$. Consider a path in C joining v_1 and v_2, and denote it by $v_1 u_1 \cdots u_k v_2$. Since $\{v_1, v_2\}$ is not an edge, we have $k \geq 1$.

Now we claim that this path has p as a vertex. To show that, denote by ℓ the line spanned by v_1 and v_2 and further denote by ℓ_p the line parallel to ℓ which passes the point p. Because of our choice, the points $u_1, \ldots, u_k$ must lie in the region bounded by ℓ, ℓ_p, ℓ_1 and ℓ_2. Then, we can see that $\bigcup_{j=0}^{k} \mathsf{conv}(\{u_j, u_{j+1}\} \cup Q)$ contains p, where u_0 and u_{k+1} are set to v_1 and v_2 respectively. This implies the existence of some index $j \in \{0, \ldots, k\}$ such that $\mathsf{conv}(\{u_j, u_{j+1}\} \cup Q)$ contains p. This contradicts the fact that $\{u_j, u_{j+1}\}$ is an edge of $G(\mathcal{L})$. Thus the claim is proved.

Now, we know that a path in C joining v_1 and v_2 passes p. However, we have two such paths in C. Since they must not share a vertex other than v_1 and v_2, this is a contradiction. $\square$

Then we observe the next lemma.

Lemma 3.10. *The clique complex of a connected chordal graph is homotopy equivalent to a single point.*

Proof (Sketch). We prove it by induction on the number of vertices. If a graph has only one vertex, it is always connected and chordal, and the clique complex consists of a single point. So the statement is true.

Assume that a connected chordal graph G has at least two vertices. Then we use a useful property of chordal graphs due to Dirac [Dir61]:

every chordal graph has a vertex whose neighbors form a clique. Let us take such a vertex and name it v. Then v and its neighbors form a clique in G. Since G is connected, the neighborhood of v is not empty. Now remove v from G to obtain a smaller graph $G' := G - v$. Since G' is also connected and chordal, the clique complex of G' is homotopy equivalent to a single point by the induction hypothesis. Then we put v back to G. This corresponds to gluing the clique complex of G' and a simplex by a facet of the simplex. So the result is also homotopy equivalent to a single point.

For a complete proof following the definition of homotopy equivalence, we have to give two continuous functions. This can be done along the line of the arguments above. □

Therefore, from Lemmas 3.9 and 3.10, we immediately obtain the following.

Corollary 3.11. *The free complex* $\mathrm{Free}(\mathcal{L})$ *of a* 2*-dimensional generalized convex shelling is homotopy equivalent to a single point.*

Note that Corollary 3.11 holds for all d-dimensional generalized convex shellings even if $Q = \emptyset$. A proof of Corollary 3.11 has already been given by Edelman & Reiner [ER00] (based on a theorem in Edelman & Jamison [EJ85]). However, our approach is discrete-geometric while they used tools from topological combinatorics.

Since an induced subgraph of a chordal graph is also chordal, we can immediately see the following.

Lemma 3.12. *Let x be a vertex of $G(\mathcal{L})$ and c_x be the number of connected components of $G(\mathcal{L}) - x$. Then* $\mathrm{del}_{\mathrm{Free}(\mathcal{L})}(x)$ *is homotopy equivalent to c_x distinct points.*

Therefore, in order to prove Theorem 3.2, we only have to show the following two lemmas.

3.3.4 Relationship of a Cut-Vertex and a Dependency Set

Lemma 3.13. *Let x be a cut-vertex of $G(\mathcal{L})$. Then $G(\mathcal{L}){-}x$ has exactly two connected components.*

Proof. Since x is a vertex of $G(\mathcal{L})$, we have $P \cap \mathsf{conv}(\{x\} \cup Q) = \{x\}$. Consider two connected components C_1 and C_2 of $G(\mathcal{L}){-}x$. Choose $u \in V(C_1)$ and $v \in V(C_2)$ such that $\{x,u\}$ and $\{x,v\}$ are edges of $G(\mathcal{L})$. Since $\{u,v\}$ is not an edge of $G(\mathcal{L})$, it should hold that $P \cap \mathsf{conv}(\{u,v\} \cup Q) \neq \{u,v\}$. Let $P' := (P \cap \mathsf{conv}(\{u,v\} \cup Q)) \setminus \{u,v\}$. From the observation above, $P' \neq \emptyset$. We claim that $x \in P'$. To show that, suppose that $x \notin P'$ for the sake of contradiction. Let P'' be the set of vertices of $G(\mathcal{L})$ which also belong to P, namely $P'' := \{y \in P' \mid P \cap \mathsf{conv}(\{y\} \cup Q) = \{y\}\}$. (Note that $P'' \neq \emptyset$.) Then each $y \in P''$ lies in either

(1) $\mathsf{conv}(\{u\} \cup Q)$,

(2) $\mathsf{conv}(\{v\} \cup Q)$, or

(3) $\mathsf{conv}(\{u,v\} \cup Q) \setminus (\mathsf{conv}(\{u\} \cup Q) \cup \mathsf{conv}(\{v\} \cup Q))$.

When (1) or (2) happens, u or v cannot be a vertex of $G(\mathcal{L})$ by Lemma 3.7, respectively. This is a contradiction. Therefore, it holds that $P'' \subseteq \mathsf{conv}(\{u,v\} \cup Q) \setminus (\mathsf{conv}(\{u\} \cup Q) \cup \mathsf{conv}(\{v\} \cup Q))$. Now, let us take the convex hull of $P'' \cup \{u,v\}$, and it has two chains of edges connecting u and v. By our assumption, one is the edge $\{u,v\}$ and the other consists of at least two edges. Consider the latter one. (In Figure 3.3, the gray region is the convex hull of $P'' \cup \{u,v\}$.) Then this chain corresponds to a path from u to v in $G(\mathcal{L})$. However, this means that C_1 and C_2 are not distinct connected components of $G(\mathcal{L}){-}x$. A contradiction. Thus, we have $x \in P'$.

Now, suppose that $G(\mathcal{L}){-}x$ has at least three connected components, say C_1, C_2, C_3. As before, choose $u \in V(C_1), v \in V(C_2), w \in V(C_3)$ such that $\{x,u\}$, $\{x,v\}$ and $\{x,w\}$ are edges of $G(\mathcal{L})$. Consider two outer common tangents ℓ_1, ℓ_2 of $\mathsf{conv}(\{u,v,w\})$ and $\mathsf{conv}(Q)$. Without loss of generality, let u be the intersection of ℓ_1 and $\mathsf{conv}(\{u,v,w\})$, and v be the intersection of ℓ_2 and $\mathsf{conv}(\{u,v,w\})$. Note that these intersection points must be distinct by the same reason as in the proof of Lemma 3.9. Let ℓ be the line spanned by u and v. We have two cases.

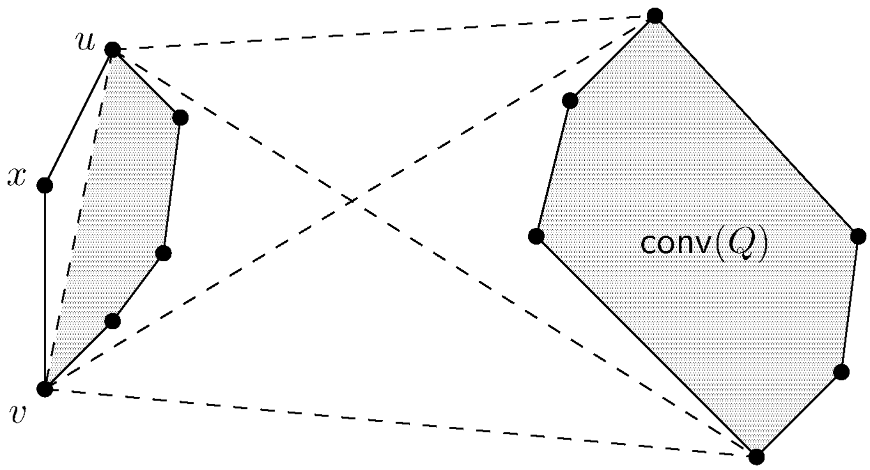

Figure 3.3: Where does x lie?

Case 1: w and Q lie on the same side of ℓ. In this case, we can see that $\mathsf{conv}(\{w\} \cup Q)$ is identical to the intersection of $\mathsf{conv}(\{u, v\} \cup Q)$, $\mathsf{conv}(\{v, w\} \cup Q)$ and $\mathsf{conv}(\{u, w\} \cup Q)$. By the claim above, x belongs to all of these three sets. Therefore, x belongs to $\mathsf{conv}(\{w\} \cup Q)$. However, since w is a vertex of $G(\mathcal{L})$, this contradicts Lemma 3.7.

Case 2: w and Q lie on the different sides of ℓ. By an argument similar to Case 1, we can observe that x belongs to $\mathsf{conv}(\{w\} \cup Q)$, which is again a contradiction. $\square$

Lemma 3.14. *Let x be a vertex of $G(\mathcal{L})$. If x is a cut-vertex of $G(\mathcal{L})$, then* $\mathrm{Dep}_{\mathcal{L}}(x) = P$.

Proof. Assume that x is a cut-vertex of $G(\mathcal{L})$. We have to show that $\mathrm{Dep}_{\mathcal{L}}(x) = P$, namely, for every $y \in P$ there exists a set $A \subseteq P$ such that

(1) $\mathrm{ex}_{\mathcal{L}}(A) = A$,

(2) $y \in A$,

(3) $x \in \tau_{\mathcal{L}}(A)$, and

(4) $x \notin \tau_{\mathcal{L}}(A \setminus \{y\})$.

Fix $y \in P$ arbitrarily. According to the position of y, we have several cases. Let $\ell_\top$ and $\ell_\perp$ be lines supporting $\mathsf{conv}(\{x\} \cup Q)$ which pass through x. (In case $|Q| = 1$, they coincide.) Denote by $\ell_\top^{\supseteq}$ the closed

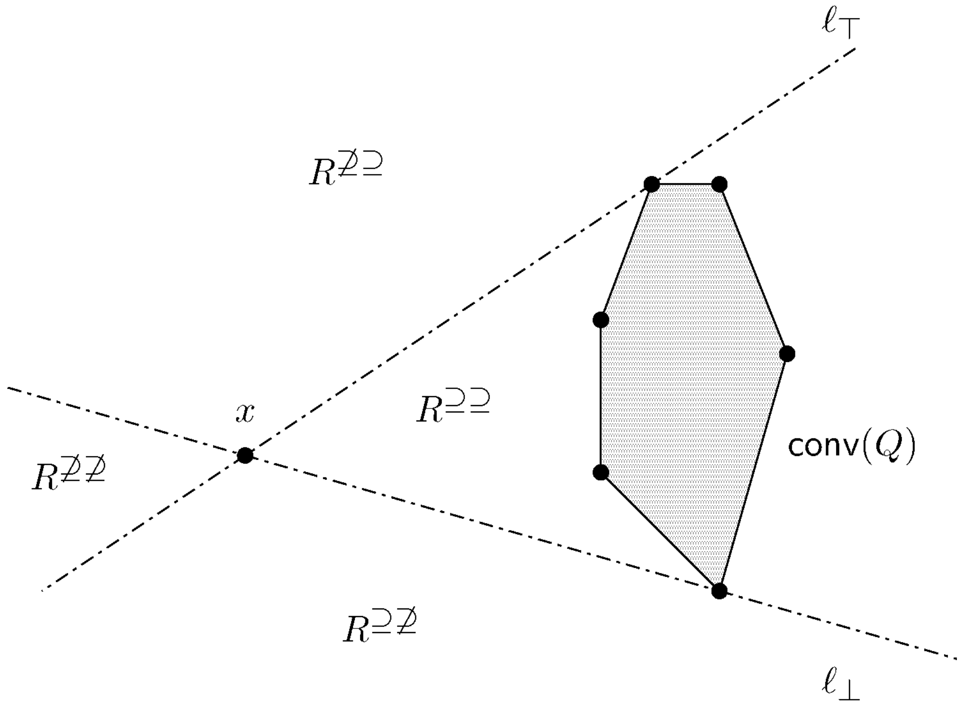

Figure 3.4: The whole plane is divided into four parts.

halfplane determined by $\ell_\top$ which contains Q, and by $\ell_\top^{\not\supseteq}$ the closed halfplane determined by $\ell_\top$ which does not contain Q. We define $\ell_\bot^{\supseteq}$ and $\ell_\bot^{\not\supseteq}$ analogously. Then, the whole plane is decomposed into four parts:

$$\begin{aligned} R^{\supseteq\supseteq} &:= \ell_\top^{\supseteq} \cap \ell_\bot^{\supseteq}, \\ R^{\supseteq\not\supseteq} &:= \ell_\top^{\supseteq} \cap \ell_\bot^{\not\supseteq}, \\ R^{\not\supseteq\supseteq} &:= \ell_\top^{\not\supseteq} \cap \ell_\bot^{\supseteq}, \\ R^{\not\supseteq\not\supseteq} &:= \ell_\top^{\not\supseteq} \cap \ell_\bot^{\not\supseteq}. \end{aligned}$$

Figure 3.4 illustrates this decomposition.

First, let us observe that $R^{\supseteq\supseteq}$ contains no point from $P \setminus \{x\}$. To show that, suppose that it contains a point $p \in P \setminus \{x\}$. If it lies in "front" of conv(Q) (i.e., the bounded region determined by $\ell_\top$, $\ell_\bot$ and conv(Q)), then it holds that $p \in \mathsf{conv}(\{x\} \cup Q)$. However, this means that x is not a vertex of $G(\mathcal{L})$ by Lemma 3.7. A contradiction. Otherwise, the line segment connecting p and x intersects conv(Q). However, this implies that $\mathsf{conv}(P) \cap \mathsf{conv}(Q)$ is not empty. A contradiction to our assumption. Thus, $R^{\supseteq\supseteq}$ contains no point from $P \setminus \{x\}$.

Hence we obtain three cases to consider about the position of y. However, the cases of $R^{\supseteq\not\supseteq}$ and $R^{\not\supseteq\supseteq}$ are symmetric. So the essential cases are the following two.

Case 1: y lies in $R^{\not\supseteq\not\supseteq}$. In this case, we can choose $\{y\}$ as A. We claim that this A satisfies conditions (1)–(4) above. Since y is an extreme point of $\mathsf{conv}(\{y\} \cup Q)$, by Lemma 3.4 condition (1) is fulfilled. The second condition is true by definition. The third and fourth conditions can be verified via Lemma 3.4. This case is done.

Case 2: y lies in $R^{\not\supseteq\supseteq}$. From the argument in the proof of Lemma 3.13, we can see that one component $G_\top$ of $G(\mathcal{L})-x$ lies in $R^{\not\supseteq\supseteq}$ and the other component $G_\perp$ of $G(\mathcal{L})-x$ is contained in $R^{\supseteq\not\supseteq}$. Both of them are non-empty. Now, let A be the set of points of P which are moreover the extreme points of $\mathsf{conv}(\{y\} \cup V(G_\perp) \cup Q)$. We claim that this A satisfies conditions (1)–(4) above.

By Lemma 3.4, condition (1) is clear. Since y lies on the different side of $\ell_\top$ than Q and $V(G_\perp)$, we can see that y is an extreme point of $\mathsf{conv}(\{y\} \cup V(G_\perp) \cup Q)$. Hence, condition (2) is fulfilled. Since Q and $V(G_\perp)$ lie on the different sides of $\ell_\perp$, and no vertex of $G_\perp$ lies on ℓ (because of Lemma 3.7), we see that $x \notin \mathsf{conv}(V(G_\perp) \cup Q)$, which means that condition (4) is satisfied.

To verify condition (3), take any vertex v of $G_\perp$. By the anti-exchange property (T5) of the closure operator (Lemma 2.5) and Lemma 3.4, we can find a point $z \in \mathsf{conv}(\{z\} \cup Q)$ which is a vertex of $G_\top$. Since x is a cut-vertex of $G(\mathcal{L})$, $\{z, v\}$ is not an edge of $G(\mathcal{L})$. Then, by Lemma 3.7 and the fact that x is a cut-vertex, we can see that $\mathsf{conv}(\{z, v\} \cup Q)$ contains x. Namely, we have

$$\begin{aligned} x \in \mathsf{conv}(\{z, v\} \cup Q) &\subseteq \mathsf{conv}(\{y, v\} \cup Q) \\ &\subseteq \mathsf{conv}(\{y\} \cup V(G_\perp) \cup Q), \end{aligned}$$

which implies that condition (3) holds by Lemma 3.4. In this way, the whole proof is completed. □

Thus, we are able to conclude the proof of Theorem 3.2. Q.E.D.

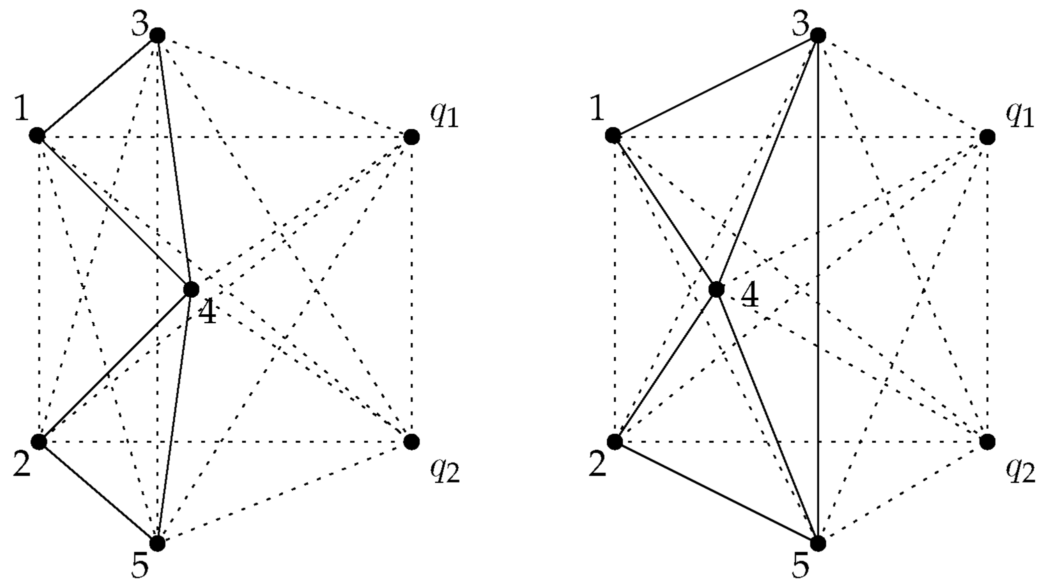

Figure 3.5: Examples.

3.4 Examples

In this section, we show that both cases in Theorem 3.2.2 can really occur by exhibiting such examples.

Look at Figure 3.5. In both of the examples, $P = \{1,2,3,4,5\}$ and $Q = \{q_1, q_2\}$. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be the generalized convex shellings on P with respect to Q of the left-hand side and the right-hand side of Figure 3.5, respectively. The solid lines show the edges of $G(\mathcal{L})$, and the dotted lines are just used for the clarification of the placement of points.

In both cases, we can observe that $\mathrm{Dep}_{\mathcal{L}_1}(4) = P$ and $\mathrm{Dep}_{\mathcal{L}_2}(4) = P$. In the left case, the deletion of 4 from $G(\mathcal{L}_1)$ results in a disconnected graph, therefore $\mathrm{del}_{\mathrm{Free}(\mathcal{L}_1)}(4)$ is homotopy equivalent to two distinct points. However, in the right case, the deletion of 4 from $G(\mathcal{L}_2)$ results in a connected graph, therefore $\mathrm{del}_{\mathrm{Free}(\mathcal{L}_2)}(4)$ is contractible.

The right example is especially interesting because before this work we did not have an example of a convex geometry $\mathcal{L}$ on a finite set E which has an element $e \in E$ such that $\mathrm{Dep}_{\mathcal{L}}(e) = E$ and $\mathrm{del}_{\mathrm{Free}(\mathcal{L})}(e)$ is contractible. Namely, this is the first example of such a kind.

Note added after the examination Recently, Hachimori & Kashiwabara [HK04] completely solved Open Problem 3.1. According to their solution, the first problem is affirmative while the second one is negative in general.

Part III

Geometric Optimization with Few Inner Points

Chapter 4

There is something in common between Hungarian, and Japanese, and Hispanic mathematics. We can start a paper by saying "Let P be a set of n points in the plane..." and no one will ask us, "Why? Why?"

János Pach (2004)

The Traveling Salesman Problem with Few Inner Points

4.1 Introduction

A lot of NP-hard optimization problems on graphs can be solved in polynomial time when the input is restricted to partial k-trees, that is, graphs with treewidth at most k, where k is fixed. In this sense, the treewidth is regarded as a natural parameter to measure the complexity of graphs. This is based on the observation that "some NP-hard optimization problems on graphs are easy when the class is restricted to trees."

We try to address the following question: what is a natural parameter that could play a similar role for geometric problems as the treewidth does for graph problems? One basic observation is that "some NP-hard optimization problems on a point set in the Euclidean plane are easy when the points are in convex position," namely, they are the vertices of a convex polygon. Therefore, the number of inner points can be regarded as a natural parameter for the complexity of geometric problems. Here, an inner point is a point in the interior of the convex hull of the given point set. Intuitively, we might say that "fewer inner points make the problem easier to solve."

In this chapter, we concentrate on one specific problem: the traveling salesman problem. The traveling salesman problem (TSP) is one of the most famous optimization problems, which comes along many kinds of applications such as logistics, scheduling, VLSI manufacturing, etc. In many practical applications, we have to solve TSP instances arising from the two-dimensional Euclidean plane, which we call the 2DTSP. Also most of the benchmark instances for TSP belong to this class. Theoretically speaking, the general 2DTSP is strongly NP-hard [GGJ76, Pap77]. On the other hand, the problem is trivial if the points are in convex position. Therefore, the following natural question is asked: what is the influence of the number of inner points on the complexity of the problem? Here, an inner point of a finite point set P is a point from P which lies in the interior of the convex hull of P.

We provide simple algorithms based on the dynamic programming paradigm. The first one runs in $\mathrm{O}(k!kn)$ time and $\mathrm{O}(k)$ space, and the second runs in $\mathrm{O}(2^k k^2 n)$ time and $\mathrm{O}(2^k kn)$ space, where n is the total number of input points and k is the number of inner points. Observe that the second algorithm gives a polynomial-time solution to the problem when $k = \mathrm{O}(\log n)$. Although the first algorithm is inferior to the second one in terms of time complexity, the first one has a benefit in its space complexity and also it is easy to parallelize.

From the viewpoint of parameterized computation [DF99b, Nie02], these algorithms are fixed-parameter algorithms when the number of inner points is taken as a parameter, hence the problem is fixed-parameter tractable (FPT). Here, a *fixed-parameter algorithm* has running time $\mathrm{O}(f(k)n^c)$, where n is the input size, k is a parameter, c is a constant independent of n and k, and $f : \mathbb{N} \to \mathbb{N}$ is an arbitrary computable function. For example, an algorithm with running time $\mathrm{O}(440^k n)$ is a fixed-parameter algorithm whereas one with $\mathrm{O}(n^k)$ is not.

We also study two variants of the traveling salesman problem: the prize-collecting traveling salesman problem, introduced by Balas [Bal89], and the partial traveling salesman problem. Both problems are also strongly NP-hard. We show that these problems in the Euclidean plane are FPT as well.

Computational Model Here, let us notice the computational model we use in this and the next chapters. As usual for computational problems, we use a random-access machine (RAM) but we allow it to deal with real numbers. Namely, a real number of any precision can be stored at a single place. Also we allow some operations on real numbers at unit cost. They include addition, subtraction, multiplication, division, comparison, and taking the square root. The square root is important for our algorithms since we have to look at Euclidean distances. This model is called the *real RAM* model (see the book of Preparata & Shamos [PS85] for example).

Related Work Since the literature on the TSP and its variants is vast, we only point out studies on the TSP itself which are closely related to our result. To the author's knowledge, only few papers studied the parameterized complexity of the 2DTSP. Probably the most closely related one is a paper by Deĭneko, van Dal & Rote [DvDR94]. They studied the 2DTSP where the inner points lie on a line. The problem is called the *convex-hull-and-line TSP*. They gave an algorithm running in $O(kn)$ time, where k is the number of inner points. Deĭneko & Woeginger [DW96] studied a slightly more general problem called the *convex-hull-and-ℓ-line TSP*, and gave an algorithm running in $O(k^\ell n^2)$ time. Compared to these results, our algorithms deal with the most general situation, and are fixed-parameter algorithms with respect to k. As for approximation algorithms, Arora [Aro98] and Mitchell [Mit99] found polynomial-time approximation schemes (PTAS) for the 2DTSP. Rao & Smith [RS98] gave a PTAS with better running time $O(n \log n + 2^{\mathrm{poly}(1/\varepsilon)} n)$. As for exact algorithms, Held & Karp [HK62] and independently Bellman [Bel62] provided a dynamic programming algorithm to solve the TSP optimally in $O(2^n n^2)$ time and $O(2^n n)$ space. For geometric problems, Hwang, Chang & Lee [HCL93] gave an algorithm to solve the 2DTSP in $n^{O(\sqrt{n})}$ time based on the so-called separator theorem.

Organization The next section introduces the problem formally, and gives fundamental lemmas. Sections 4.3 and 4.4 describe algorithms for the 2DTSP. Variations are discussed in Section 4.5. We conclude with an open problem in the final section.

4.2 Traveling Salesman Problem with Few Inner Points

Let $P \subseteq \mathbb{R}^2$ be a set of n points in the Euclidean plane. The *convex hull* of P is the smallest convex set containing P. A point $p \in P$ is called an *inner* point if p lies in the interior of the convex hull of P. We denote by $\mathsf{Inn}(P)$ the set of inner points of P. A point $p \in P$ is called an *outer* point if it is not an inner point, i.e., it is on the boundary of the convex hull of P. We denote by $\mathsf{Out}(P)$ the set of outer points of P. Note that $P = \mathsf{Inn}(P) \cup \mathsf{Out}(P)$ and $\mathsf{Inn}(P) \cap \mathsf{Out}(P) = \emptyset$. Let $n := |P|$ and $k := |\mathsf{Inn}(P)|$. (So, we have $|\mathsf{Out}(P)| = n-k$.)

A *tour* on P is a linear order $(x_1, x_2, \ldots, x_n)$ of the points in P. We say that this tour *starts at* x_1, and we assume that all indices are taken modulo n in the tour so that x_{n+1} can be identified with x_1. We often identify the tour $(x_1, \ldots, x_n)$ on P with a closed polygonal curve consisting of the line segments $\overline{x_1x_2}$, $\overline{x_2x_3}$, $\ldots$, $\overline{x_{n-1}x_n}$, $\overline{x_nx_1}$. The *length* of the tour is the Euclidean length of this polygonal curve, i.e., $\sum_{i=1}^{n} d(x_i, x_{i+1})$, where $d(x_i, x_{i+1})$ stands for the Euclidean distance from x_i to x_{i+1}. The objective of the *traveling salesman problem* (TSP) is to find a shortest tour. The following lemma was probably first noted by Flood [Flo56] and nowadays it is folklore.

Lemma 4.1 (Flood [Flo56]). *Every shortest tour has no crossing.*

Proof. Let P be a given set of n points in the plane, and $(x_1, \ldots, x_n)$ be a shortest tour on P which has a crossing on $\overline{x_ix_{i+1}}$ and $\overline{x_jx_{j+1}}$ ($i < j$). See the left part of Figure 4.1 where $n = 8$, $i = 3$ and $j = 6$. Then, we remove $\overline{x_ix_{i+1}}$ and $\overline{x_jx_{j+1}}$, and add $\overline{x_ix_j}$ and $\overline{x_{i+1}x_{j+1}}$ to form another tour (the right part of Figure 4.1). Then by the triangle inequality we obtain $d(x_i, x_j) + d(x_{i+1}, x_{j+1}) < d(x_i, x_{i+1}) + d(x_j, x_{j+1})$. Therefore the second tour is shorter than the first tour. This is a contradiction to the assumption that the first one is shortest. $\square$

This lemma immediately implies the following lemma, which plays a fundamental role in our algorithm. We call a linear order on $\mathsf{Out}(P)$ *cyclic* if every two consecutive points in the order are also consecutive on the boundary of the convex hull of P.

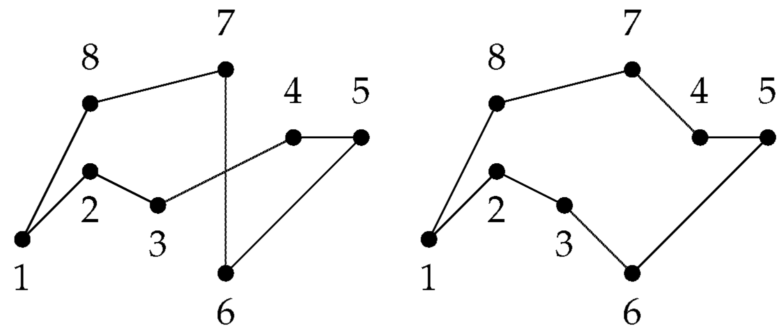

Figure 4.1: Proof of Lemma 4.1.

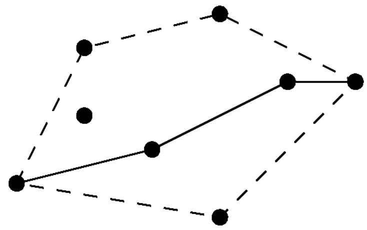

Figure 4.2: Proof of Lemma 4.2.

Lemma 4.2. *In every shortest tour on P, the points of* $\mathsf{Out}(P)$ *appear in a cyclic order.*

Proof. For the sake of contradiction, suppose that there exists a shortest tour τ which does not respect a cyclic order on $\mathsf{Out}(P)$. This means that when $(x^o_1, x^o_2, \dots, x^o_{n-k})$ is a cyclic order on $\mathsf{Out}(P)$, there exists an index $i \in \{1, \dots, n-k\}$ such that x^o_i and x^o_{i+1} are not consecutive on τ (where $i+1$ is considered modulo $n-k$). Let x^o_j be the next point of x^o_i in τ. By the assumption, j is not $i+1$ modulo m. Therefore, a polygonal chain from x^o_i to x^o_j divides the convex hull of P into two parts and both parts contain an outer point which does not participate in this polygonal chain. See Figure 4.2. The tour τ continues itself to one of the partitioned side, and when it tries to enter the other side, we obtain a crossing. This contradicts Lemma 4.1 since τ is a shortest tour. □

With Lemma 4.2, we can establish the following naive algorithm: take an arbitrary cyclic order on $\mathsf{Out}(P)$, then look through all tours (i.e., the linear orders) π on P which respect[1] this cyclic order; com-

[1]For a set S and a subset $S' \subseteq S$, we say that a linear order π on S *respects* a linear

pute the length of each tour and output the best one among them. The number of such tours is $O(k!n^k)$. Since the computation of the length of a tour takes $O(n)$ time, in total the running time of this algorithm is $O(k!n^{k+1})$. So, if k is constant, this algorithm runs in polynomial time. However, it is not a fixed-parameter algorithm with respect to k since k appears in the exponent of n.

4.3 First Fixed-Parameter Algorithm

First, let us notice that later on we always assume that, when P is given to an algorithm as input, the algorithm already knows $\mathsf{Out}(P)$ together with a cyclic order $\gamma = (p_1, \ldots, p_{n-k})$ on $\mathsf{Out}(P)$. Also, note that the space complexity in the algorithms below do not count the input size, as usual in theoretical computer science.

Our first algorithm adapts the following idea. We look through all linear orders on $\mathsf{Inn}(P)$. Let π be a linear order π on $\mathsf{Inn}(P)$. We will try to find a shortest tour on P which respects both the cyclic order γ on $\mathsf{Out}(P)$ and the linear order π on $\mathsf{Inn}(P)$. Then, we exhaust this procedure for all linear orders on $\mathsf{Inn}(P)$, and output a minimum one. Later we will show that we can compute such a tour in $O(kn)$ time and $O(k)$ space. Then, since the number of linear orders on $\mathsf{Inn}(P)$ is $k!$ and they can be enumerated in amortized $O(1)$ time per one linear order with $O(k)$ space [Sed77], overall the algorithm runs in $O(k!kn)$ time and $O(k)$ space.

Now, given a cyclic order γ on $\mathsf{Out}(P)$ and a linear order π on $\mathsf{Inn}(P)$, we describe how to compute a shortest tour among those which respect γ and π by dynamic programming. For dynamic programming in algorithmics, see the textbook by Cormen, Leiserson, Rivest & Stein [CLRS01], for example.

We consider a three-dimensional array $F_1[i, j, m]$, where $i \in \{1, \ldots, n-k\}$, $j \in \{0, 1, \ldots, k\}$, and $m \in \{\mathsf{Inn}, \mathsf{Out}\}$. The first index i represents the point p_i in $\mathsf{Out}(P)$, the second index j represents the point q_j in $\mathsf{Inn}(P)$, and the third index m represents the position.

order π' on S' if the restriction of π onto S' is π'.

The value $F_1[i,j,m]$ represents the length of a shortest "path" on $\{p_1,\dots,p_i\} \cup \{q_1,\dots,q_j\}$ that satisfies the following conditions.

- It starts at $p_1 \in \mathsf{Out}(P)$.
- It visits exactly the points in $\{p_1,\dots,p_i\} \cup \{q_1,\dots,q_j\}$. (If $j = 0$, set $\{q_1,\dots,q_j\} = \emptyset$.)
- It respects the orders γ and π.
- If $m = \mathsf{Out}$, then it ends at p_i (an outer point). If $m = \mathsf{Inn}$, then it ends at q_j (an inner point).

Then, the length of a shortest tour respecting π and γ can be computed as

$$\min\{F_1[n{-}k,k,\mathsf{Out}] + d[p_{n-k},p_1], F_1[n{-}k,k,\mathsf{Inn}] + d(q_k,p_1)\}.$$

Therefore, it suffices to know the values $F_1[i,j,m]$ for all possible i,j,m.

To do that, we establish a recurrence. First let us look at the boundary cases.

- Since we start at p_1, set $F_1[1,0,\mathsf{Out}] = 0$.
- There are some impossible states for which we set the values to ∞. Namely, for every $j \in \{1,\dots,k\}$ set $F_1[1,j,\mathsf{Out}] = \infty$, and for every $i \in \{1,\dots,n{-}k\}$ set $F_1[i,0,\mathsf{Inn}] = \infty$.

Now, assume we want to visit the points of $\{p_1,\dots,p_i\} \cup \{q_1,\dots,q_j\}$ while respecting the orders γ and π and arrive at q_j. How can we get to this state? Since we respect the orders, either (1) first we have to visit the points of $\{p_1,\dots,p_i\} \cup \{q_1,\dots,q_{j-1}\}$ to arrive at p_i then move to q_j, or (2) first we have to visit the points of $\{p_1,\dots,p_i\} \cup \{q_1,\dots,q_{j-1}\}$ to arrive at q_{j-1} then move to q_j. Therefore, we have

$$F_1[i,j,\mathsf{Inn}] = \min\{F_1[i,j{-}1,\mathsf{Out}] + d(p_i,q_j), F_1[i,j{-}1,\mathsf{Inn}] + d(q_{j-1},q_j)\} \tag{4.1}$$

for $(i,j) \in \{1,\dots,n{-}k\} \times \{1,\dots,k\}$. Similarly, we have

$$F_1[i,j,\mathsf{Out}] = \min\{F_1[i{-}1,j,\mathsf{Out}] + d(p_{i-1},p_i), F_1[i{-}1,j,\mathsf{Inn}] + d(q_j,p_i)\} \tag{4.2}$$

for $(i,j) \in \{2,\ldots,n-k\} \times \{0,\ldots,k\}$, where $d(q_0,p_i)$ is considered ∞ for convenience. Since what is referred to in the right-hand sides of Equalities (4.1) and (4.2) has smaller indices, we can solve this recursion in a bottom-up way by dynamic programming. This completes the dynamic-programming formulation for the computation of $F_1[i,j,m]$.

The size of the array is $(n-k) \times (k+1) \times 2 = \mathrm{O}(kn)$, and the computation of each entry requires to look up at most two other entries of the array. Therefore, we can fill up the array in $\mathrm{O}(kn)$ time and $\mathrm{O}(kn)$ space.

Now, we describe how to reduce the space requirement to $\mathrm{O}(k)$. For each $(i,j) \in \{1,\ldots,n-k\} \times \{1,\ldots,k\}$, consider when $F_1[i,j,\mathsf{Inn}]$ is looked up throughout the computation. It is looked up only when we compute $F_1[i,j+1,\mathsf{Inn}]$ and $F_1[i+1,j,\mathsf{Out}]$. So the effect of $F_1[i,j,\mathsf{Inn}]$ is local. Similarly, the value $F_1[i,j,\mathsf{Out}]$ is looked up only when we compute $F_1[i,j+1,\mathsf{Inn}]$ and $F_1[i+1,j,\mathsf{Out}]$. We utilize this locality in the computation.

We divide the computation process into some phases. For every $i \in \{1,\ldots,n-k\}$, in the i-th phase, we compute $F_1[i,j,\mathsf{Inn}]$ and $F_1[i,j,\mathsf{Out}]$ for all $j \in \{0,\ldots,k\}$. Within the i-th phase, the computation starts with $F_1[i,1,\mathsf{Out}]$ and proceeds along $F_1[i,2,\mathsf{Out}], F_1[i,3,\mathsf{Out}],\ldots$, until we get $F_1[i,k,\mathsf{Out}]$. Then, we start calculating $F_1[i,1,\mathsf{Inn}]$ and proceed along $F_1[i,2,\mathsf{Inn}], F_1[i,3,\mathsf{Inn}],\ldots$, until we get $F_1[i,k,\mathsf{Inn}]$. From the observation above, all the computation in the i-th phase only needs the outcome from the $(i-1)$-st phase and the i-th phase itself. Therefore, we only have to store the results from the $(i-1)$-st phase for each i. This requires only $\mathrm{O}(k)$ storage.

In this way, we obtain the following theorem. Let us remind that $\log(k!) = \Theta(k \log k)$.

Theorem 4.3. *The 2DTSP on n points including k inner points can be solved in $\mathrm{O}(k!kn)$ time and $\mathrm{O}(k)$ space. In particular, it can be solved in polynomial time if $k = \mathrm{O}(\log n / \log\log n)$.* □

4.4 Second Fixed-Parameter Algorithm with Better Running Time

To obtain a faster algorithm, we make use of the trade-off between the time complexity and the space complexity. Compared to the first algorithm, the second algorithm has a better running time $O(2^k k^2 n)$ but needs more space $O(2^k kn)$. The idea of trade-off is also taken by the dynamic programming algorithm for the general traveling salesman problem due to Bellman [Bel62] and Held & Karp [HK62], and our second algorithm is essentially a generalization of their algorithm. (For a nice exposition of this "dynamic programming across the subsets" technique together with other methods for exact computation, see Woeginger's survey article [Woe03].)

In the second algorithm, we first fix a cyclic order γ on $\mathsf{Out}(P)$. Then, we immediately start the dynamic programming. This time, we consider the following three-dimensional array $F_2[i,S,r]$, where $i \in \{1,\dots,n-k\}$, $S \subseteq \mathsf{Inn}(P)$, and $r \in S \cup \{p_i\}$. We interpret $F_2[i,S,r]$ as the length of a shortest path on $\{p_1,\dots,p_i\} \cup S$ that satisfies the following conditions.

- It starts at $p_1 \in \mathsf{Out}(P)$.
- It visits exactly the points in $\{p_1,\dots,p_i\} \cup S$.
- It respects the order γ.
- It ends at r.

Then, the length of a shortest tour can be computed as

$$\min\{F_2[n-k,\mathsf{Inn}(P),r] + d(r,p_1) \mid r \in \mathsf{Inn}(P) \cup \{p_{n-k}\}\}.$$

Therefore, it suffices to know the values $F_2[i,S,r]$ for all possible triples (i,S,r).

To do that, we establish a recurrence. The boundary cases are as follows.

- Since we start at p_1, set $F_2[1,\emptyset,p_1] = 0$.
- Set $F_2[1,S,p_1] = \infty$ when $S \neq \emptyset$, since this is an unreachable situation.

Let $i \in \{2,\ldots,n-k\}$ and $S \subseteq \mathsf{Inn}(P)$. We want to visit the points of $\{p_1,\ldots,p_i\} \cup S$ while respecting the order γ and arrive at p_i. How can we get to this state? Since we respect the order γ, we first have to visit the points of $\{p_1,\ldots,p_{i-1}\} \cup S$ to arrive at a point in $S \cup \{p_{i-1}\}$ and then move to p_i. Therefore, we have

$$F_2[i,S,p_i] = \min\{F_2[i-1,S,t] + d(t,p_i) \mid t \in S \cup \{p_{i-1}\}\} \tag{4.3}$$

for $i \in \{2,\ldots,n-k\}$ and $S \subseteq \mathsf{Inn}(P)$. Similarly, we have

$$F_2[i,S,r] = \min\{F_2[i,S \setminus \{r\},t] + d(t,r) \mid t \in (S \setminus \{r\}) \cup \{p_i\}\} \tag{4.4}$$

for $i \in \{2,\ldots,n-k\}$, $S \subseteq \mathsf{Inn}(P)$, $S \neq \emptyset$ and $r \in S$. This completes the dynamic-programming formulation for the computation of $F_2[i,S,r]$.

The size of the array in this algorithm is $(n-k)\sum_{s=0}^{k}\binom{k}{s}s = \mathrm{O}(2^k kn)$, and the computation of each entry requires to look up $\mathrm{O}(k)$ other entries. Therefore, we can fill up the array in $\mathrm{O}(2^k k^2 n)$ time and in $\mathrm{O}(2^k kn)$ space. Thus, we obtain the following theorem.

Theorem 4.4. *The 2DTSP on n points including k inner points can be solved in $\mathrm{O}(2^k k^2 n)$ time and $\mathrm{O}(2^k kn)$ space. In particular, it can be solved in polynomial time if $k = \mathrm{O}(\log n)$.* □

4.5 Variants of the Traveling Salesman Problem

Since our approach to the TSP in the previous section is based on the general dynamic programming paradigm, it is also applicable to other variants of the TSP. In this section, we discuss two of them.

4.5.1 Prize-Collecting Traveling Salesman Problem

In the *prize-collecting TSP*, we are given an n-point set $P \subseteq \mathbb{R}^2$ with a distinguished point $h \in P$ called the *home*, and a non-negative number $c(p) \in \mathbb{R}$ for each point $p \in P$ which we call the *penalty* of p. The goal is to find a subset $P' \subseteq P \setminus \{h\}$ and a tour on $P' \cup \{h\}$ starting at h

which minimizes the length of the tour minus the penalties over all $p \in P' \cup \{h\}$. In this section, the value of a tour (or a path) refers to the value of this objective.

For this problem, we basically follow the same procedure as the TSP, but a little attention has to be paid because in this case we have to *select* some of the points from P. In addition, we have to consider two cases: $h \in \mathsf{Inn}(P)$ or $h \in \mathsf{Out}(P)$.

First Algorithm

First, let us consider the case $h \in \mathsf{Out}(P)$. In this case, we fix a cyclic order γ on $\mathsf{Out}(P)$, which starts at h, and we look through all linear orders on $\mathsf{Inn}(P)$. Let $\gamma = (p_1, p_2, \ldots, p_{n-k})$, where $p_1 = h$, and fix one linear order $\pi = (q_1, q_2, \ldots, q_k)$ on $\mathsf{Inn}(P)$. Then, we consider a three-dimensional array $F_1[i,j,m]$, where $i \in \{1,\ldots,n-k\}$, $j \in \{0,1,\ldots,k\}$ and $m \in \{\mathsf{Inn}, \mathsf{Out}\}$. The value $F_1[i,j,m]$ is interpreted as the value of an optimal path on $\{p_1,\ldots,p_i\} \cup \{q_1,\ldots,q_j\}$ which satisfies the following conditions.

- It starts at $p_1 \in \mathsf{Out}(P)$.
- It visits *some* points from $\{p_1,\ldots,p_i\} \cup \{q_1,\ldots,q_j\}$, and not more. (If $j = 0$, set $\{q_1,\ldots,q_j\} = \emptyset$.)
- It respects the orders γ and π.
- If $m = \mathsf{Out}$, then it ends at p_i. If $m = \mathsf{Inn}$, then it ends at q_j.

We want to compute the values $F_1[i,j,m]$ for all possible triples (i,j,m).

The boundary cases are:

- $F_1[1,j,\mathsf{Out}] = -c(p_1)$ for every $j \in \{1,\ldots,k\}$; and
- $F_1[i,0,\mathsf{Inn}] = \infty$ for every $i \in \{1,\ldots,n-k\}$,

and the main part of the recurrence is:

$$\begin{aligned} F_1[i,j,\mathsf{Inn}] = \min\{ & \min_{i' \in \{1,\ldots,i\}} \{F_1[i',j-1,\mathsf{Out}] + d(p_{i'},q_j) - c(q_j)\}, \\ & \min_{j' \in \{0,\ldots,j-1\}} \{F_1[i,j',\mathsf{Inn}] + d(q_{j'},q_j) - c(q_j)\}\} \end{aligned}$$

for $(i,j) \in \{1,\ldots,n-k\} \times \{1,\ldots,k\}$, and

$$F_1[i,j,\mathsf{Out}] = \min\{\min_{i' \in \{1,\ldots,i-1\}} \{F_1[i',j,\mathsf{Out}] + d(p_{i'},p_i) - c(p_i)\}, \min_{j' \in \{0,\ldots,j\}} \{F_1[i-1,j',\mathsf{Inn}] + d(q_{j'},p_i) - c(p_i)\}\}$$

for $(i,j) \in \{2,\ldots,n-k\} \times \{0,\ldots,k\}$. For convenience, $d(q_0,p_i)$ is considered to be ∞.

Then, the value of an optimal prize-collecting tour respecting π and γ can be computed as

$$\min\{\min_{i \in \{1,\ldots,n-k\}} \{F_1[i,k,\mathsf{Out}] + d(p_i,p_1)\}, \min_{j \in \{0,\ldots,k\}} \{F_1[n-k,j,\mathsf{Inn}] + d(q_j,p_1)\}\}.$$

Since the size of the array is $\mathrm{O}(kn)$ and each entry can be filled by looking up $\mathrm{O}(n)$ other entries, the running time is $\mathrm{O}(kn^2)$. Therefore, looking through all linear orders on $\mathsf{Inn}(P)$, the overall running time of the algorithm is $\mathrm{O}(k!kn^2)$.

Next, let us consider the case $h \in \mathsf{Inn}(P)$. In this case, we look through all linear orders on $\mathsf{Inn}(P)$ staring at h, and also all cyclic orders on $\mathsf{Out}(P)$. Fix one linear order $\pi = (q_1,q_2,\ldots,q_k)$ on $\mathsf{Inn}(P)$, where $q_1 = h$, and one cyclic order $\gamma = (p_1,p_2,\ldots,p_{n-k})$ on $\mathsf{Out}(P)$. Then, we consider a three-dimensional array $F_1[i,j,m]$, where $i \in \{0,1,\ldots,n-k\}$, $j \in \{1,\ldots,k\}$ and $m \in \{\mathsf{Inn},\mathsf{Out}\}$. The interpretation and the obtained recurrence is similar to the first case, hence omitted. However, in this case, the number of orders we look through is $\mathrm{O}(k!n)$. Therefore, the overall running time of the algorithm in this case is $\mathrm{O}(k!kn^3)$. Thus, we obtain the following theorem.

Theorem 4.5. *The prize-collecting TSP in the Euclidean plane can be solved in $\mathrm{O}(k!kn^3)$ time and $\mathrm{O}(kn)$ space, when n is the total number of input points and k is the number of inner points. In particular, it can be solved in polynomial time if $k = \mathrm{O}(\log n/ \log\log n)$.* □

Second Algorithm

Now we adapt the second algorithm for the 2DTSP to the prize-collecting TSP. Let us consider the case $h \in \mathsf{Out}(P)$. (The case

$h \in \mathsf{Inn}(P)$ can be handled in the same way.) For a cyclic order $\gamma = (p_1, \ldots, p_{n-k})$ on $\mathsf{Out}(P)$ with $p_1 = h$, we define a three-dimensional array $F_2[i, S, r]$ where $i \in \{1, \ldots, n-k\}$, $S \subseteq \mathsf{Inn}(P)$ and $r \in S \cup \{p_i\}$. We interpret $F_2[i, S, r]$ as the value of an optimal path on $\{p_1, \ldots, p_i\} \cup S$ that satisfies the following conditions.

- It starts at $p_1 \in \mathsf{Out}(P)$.
- It visits some points of $\{p_1, \ldots, p_i\} \cup S$.
- It respects the order γ.
- It ends at r.

Then, the value of an optimal tour can be computed as

$$\min\{F_2[n-k, \mathsf{Inn}(P), r] + d(r, p_1) \mid r \in P\}.$$

The boundary cases are:

- $F_2[1, \emptyset, p_1] = -c(p_1)$;
- $F_2[1, S, r] = \infty$ when $S \neq \emptyset$.

The main part of the recurrence is

$$F_2[i, S, p_i] = \min_{t \in S \cup \{p_{i-1}\}} \{F_2[i-1, S, t] + d(t, p_i) - c(p_i)\}$$

for $i \in \{2, \ldots, n-k\}$ and $S \subseteq \mathsf{Inn}(P)$; and

$$F_2[i, S, r] = \min_{t \in (S \setminus \{r\}) \cup \{p_i\}} \{F_2[i, S \setminus \{r\}, t] + d(t, r) - c(r)\}$$

for $i \in \{2, \ldots, n-k\}$, $S \subseteq \mathsf{Inn}(P)$ and $r \in S$. Then, we see that the computation can be done in $\mathrm{O}(2^k k^2 n)$ time and $\mathrm{O}(2^k kn)$ space.

Theorem 4.6. *The prize-collecting TSP in the Euclidean plane can be solved in $\mathrm{O}(2^k k^2 n)$ time and $\mathrm{O}(2^k kn)$ space, when n is the total number of input points and k is the number of inner points. In particular, it can be solved in polynomial time if $k = \mathrm{O}(\log n)$.* $\square$

4.5.2 Partial Traveling Salesman Problem

In the *ℓ-partial TSP*[2], we are given an n-point set $P \subseteq \mathbb{R}^2$ with a distinguished point $h \in P$ called the home, and a positive integer $\ell \leq n$. We are asked to find a shortest tour starting at h and consisting of ℓ points from P.

We do not give an adaptation of the first algorithm for the TSP, although it is possible but too tedious to elaborate. So, we just describe a variation of the second algorithm.

Second Algorithm

Similarly to the prize-collecting TSP, we have to consider two cases: $h \in \mathsf{Inn}(P)$ or $h \in \mathsf{Out}(P)$. Here we only consider the case $h \in \mathsf{Out}(P)$. (The case $h \in \mathsf{Inn}(P)$ is similar.) Fix a cyclic order $\gamma = (p_1, \ldots, p_{n-k})$ on $\mathsf{Out}(P)$, where $p_1 = h$. We consider a four-dimensional array $F_2[i, S, r, m]$, where $i \in \{1, \ldots, n-k\}$, $S \subseteq \mathsf{Inn}(P)$, $r \in S \cup \{p_i\}$, and $m \in \{1, \ldots, \ell\}$. Then, $F_2[i, S, r, m]$ is interpreted as the length of a shortest path that satisfies the following conditions.

- It starts at $p_1 \in \mathsf{Out}(P)$.
- It visits exactly m points of $\{p_1, \ldots, p_i\} \cup S$.
- It respects the order γ.
- It ends at r.

Note that the fourth index m represents the number of points which have already been visited. Then, the length of a shortest tour through ℓ points is

$$\min\{F_2[i, \mathsf{Inn}(P), r, \ell] + d(r, p_1) \mid i \in \{1, \ldots, n-k\}, r \in \mathsf{Inn}(P) \cup \{p_i\}\}.$$

Therefore, it suffices to compute the values $F_2[i, S, r, m]$ for all possible i, S, r, m.

The boundary cases are:

[2] Usually the problem is called the k-partial TSP. However, since k is reserved for the number of inner points in the current work, we will use ℓ instead of k.

- $F_2[i, S, r, 1] = 0$ if $i = 1$ and $r = p_1$; Otherwise $F_2[i, S, r, 1] = \infty$;
- $F_2[1, S, p_1, m] = \infty$ for $m > 1$.

The main part of the recurrence is:

$$F_2[i, S, p_i, m] = \min_{t \in S \cup \{p_{i-1}\}} \{F_2[i-1, S, t, m-1] + d(t, p_i)\}$$

for $i \in \{2, \ldots, n-k\}$, $S \subseteq \mathsf{Inn}(P)$ and $m \in \{2, \ldots, \ell\}$;

$$F_2[i, S, r, m] = \min_{t \in (S \setminus \{r\}) \cup \{p_i\}} \{F_2[i, S \setminus \{r\}, t, m-1] + d(t, r)\}$$

for $i \in \{1, \ldots, n-k\}$, $S \subseteq \mathsf{Inn}(P)$, $r \in S$ and $m \in \{2, \ldots, \ell\}$.

Although the size of the array is $O(2^k k \ell n)$, we can reduce the space requirement to $O(2^k kn)$ because of the locality with respect to the fourth index m. In this way, we obtain the following theorem.

Theorem 4.7. *The ℓ-partial TSP in the Euclidean plane can be solved in $O(2^k k^2 \ell n)$ time and $O(2^k kn)$ space, where n is the total number of input points and k is the number of inner points. In particular, it can be solved in polynomial time if $k = O(\log n)$.*

4.6 Concluding Remarks

We have investigated the influence of the number of inner points in the two-dimensional Euclidean traveling salesman problem. Our results support the intuition "fewer inner points make the problem easier to solve," and nicely "interpolate" triviality when we have no inner point and intractability for the general case. This interpolation has been explored from the viewpoint of parameterized computation. Let us note that the results in this chapter can also be applied to the two-dimensional Manhattan traveling salesman problem, where the distance is measured by the ℓ_1-norm. That is because Lemmas 4.1 and 4.2 are also true for that case. More generally, our algorithms solve any TSP instance (not necessarily geometric) for which $n-k$ points have to be visited in a specified order.

The major open question is to improve the time complexity $O(2^k k^2 n)$. For example, is there a polynomial-time algorithm for the 2DTSP when $k = O(\log^2 n)$?

Franz Sanchez: Problem solver.
James Bond: More of a problem eliminator.

License to Kill (1989)

Chapter 5

The Minimum Weight Triangulation Problem with Few Inner Points

5.1 Introduction

Following the line of the previous chapter, we continue the study of geometric optimization problems with few inner points. This chapter is devoted to the minimum weight triangulation problem, which is notorious as one of the problems not known to be NP-hard nor solvable in polynomial time for a long time [GJ79]. However, when the points are in convex position, the problem can be solved in polynomial time by dynamic programming. The main result in this chapter is an exact algorithm to compute a minimum weight triangulation in $O(6^k n^5 \log n)$ time, where n is the total number of input points and k is the number of inner points. From the viewpoint of parameterized complexity [DF99b, Nie02] this is a fixed-parameter algorithm if k is taken as a parameter. Furthermore, the algorithm implies that the problem can be solved in polynomial time if $k = O(\log n)$.

Actually, our algorithm also works for simple polygons with inner points. Or, rather we should say that the algorithm is designed for such objects, and as a special case, we can compute a minimum weight triangulation of a point set. This digression to simple polygons is essential because our strategy is based on recursion and in the recursion we encounter simple polygons.

Related work Since the literature on the minimum weight triangulation problem is vast, we just mention some articles that are closely related to ours. As already mentioned, finding a minimum weight triangulation of a finite point set is not known to be NP-hard nor solvable in polynomial time [GJ79]. For an n-vertex convex polygon, the problem can be solved in $O(n^3)$ using dynamic programming. For an n-vertex simple polygon, Gilbert [Gil79] and Klincsek [Kli80] independently gave a dynamic-programming algorithm running in $O(n^3)$ time. But with inner points the problem seems more difficult. Another polynomial-time solvable case was discussed by Anagnostou & Corneil [AC93]: they considered the case where a given point set lies on a constant number of nested convex hulls. As for exact algorithms for the general case, Kyoda, Imai, Takeuchi & Tajima [KITT97] took an integer programming approach and devised a branch-and-cut algorithm. Aichholzer [Aic99] introduced the concept of a "path of a triangulation," which can be used to solve any kinds of "decomposable" problems (in particular the minimum weight triangulation problem) by recursion. These algorithms were not analyzed in terms of worst-case time complexity. As for approximation of minimum weight triangulations, Levcopoulos & Krznaric [LK98] gave a constant-factor polynomial-time approximation algorithm, but with a huge constant.

5.2 Preliminaries and Description of the Result

We start our discussion by introducing some notation and definitions used in this chapter. Then we state our result in a precise manner. From now on, we assume that input points are in general position, that is, no three points are on a single line and no two points have the same x-coordinate. When a point p has a larger x-coordinate than a point q, we say p is *right* of q; otherwise p is *left* of q.

The line segment connecting two points $p, q \in \mathbb{R}^2$ is denoted by $\overline{pq}$. The length of a line segment $\overline{pq}$ is denoted by $d(p, q)$, which is measured by the Euclidean distance. A *polygonal chain* is a planar shape described as $\gamma = \bigcup_{i=0}^{\ell-1} \overline{p_i p_{i+1}}$ where $p_0, \ldots, p_\ell \in \mathbb{R}^2$ are distinct points except that p_0 and p_ℓ can be identical (in such a case, the chain is *closed*). For a closed polygonal chain we assume in the following that all indices are taken modulo ℓ. The length of γ is the sum of the lengths of the line segments, that is, $\mathsf{length}(\gamma) = \sum_{i=0}^{\ell-1} d(p_i, p_{i+1})$. We say γ is *selfintersecting* if there exist two indices $i, j \in \{0, \ldots, \ell-1\}$, $i \neq j$, such that $(\overline{p_i p_{i+1}} \cap \overline{p_j p_{j+1}}) \setminus \{p_i, p_{i+1}, p_j, p_{j+1}\} \neq \emptyset$. Otherwise, we say γ is *non-selfintersecting*. The points $p_0, \ldots, p_\ell$ are the *vertices* of γ. When γ is not closed, p_0 and p_ℓ are called the *endpoints* of γ. In this case, we say γ *starts from* p_0 (or p_ℓ).

A *simple polygon* P is a simply connected compact region in the plane bounded by a closed non-selfintersecting polygonal chain. A *vertex* of P is a vertex of the polygonal chain which is the boundary of P. We denote the set of vertices of P by $\mathsf{Vert}(P)$. A *neighbor* of a vertex $p \in \mathsf{Vert}(P)$ is a vertex $q \in \mathsf{Vert}(P)$ such that the line segment $\overline{pq}$ lies on the boundary of P.

Following Aichholzer, Rote, Speckmann & Streinu [ARSS03], we call a pair $\Pi = (S, P)$ a *pointgon* when P is a simple polygon and S is a finite point set in the interior of P. We call S the set of *inner points* of Π. The *vertex set* of Π is $\mathsf{Vert}(P) \cup S$, and denoted by $\mathsf{Vert}(\Pi)$. Figure 5.1 shows an example of a pointgon.

Let $\Pi = (S, P)$ be a pointgon. A *triangulation* $\mathcal{T}$ of a pointgon $\Pi = (S, P)$ is a subdivision of P into triangles whose edges are straight line segments connecting two points from $\mathsf{Vert}(\Pi)$ and which have no point from $\mathsf{Vert}(\Pi)$ in their interior. The *weight* of $\mathcal{T}$ is the sum of the edge lengths used in $\mathcal{T}$. (Especially, all segments on the boundary of P are used in any triangulation and counted in the weight.) A *minimum weight triangulation* of a pointgon Π is a triangulation of Π which has minimum weight among all triangulations.

In this chapter, we study the problem of computing a minimum weight triangulation of a given pointgon $\Pi = (S, P)$. The input size is proportional to $|\mathsf{Vert}(\Pi)|$. In the sequel, for a given pointgon $\Pi =$

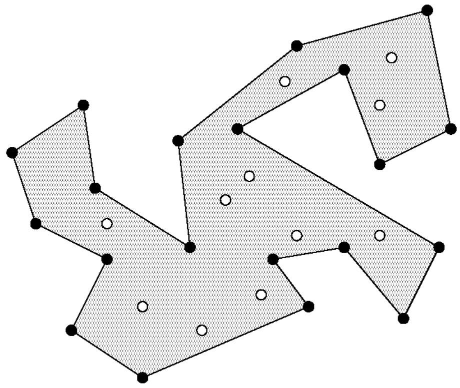

Figure 5.1: A pointgon $\Pi = (S, P)$. In this chapter, the points of S are drawn by empty circles and the points of $\mathsf{Vert}(P)$ are drawn by solid circles.

(S, P), we set $n := |\mathsf{Vert}(\Pi)|$ and $k := |S|$. Our goal is to find an exact algorithm for a pointgon $\Pi = (S, P)$ when $|S|$ is small. The main theorem of this chapter is the following.

Theorem 5.1. *Let $\Pi = (S, P)$ be a pointgon. Let $n := |\mathsf{Vert}(\Pi)|$ and $k := |S|$. Then we can find a minimum weight triangulation of Π in $O(6^k n^5 \log n)$ time. In particular, if $k = O(\log n)$, then a minimum weight triangulation can be found in polynomial time.*

This theorem shows that, in the terminology of parameterized computation, the problem is fixed-parameter tractable, when the size of S is taken as a parameter.

In the next section, we prove this theorem by providing an algorithm.

5.3 A Fixed-Parameter Algorithm for Minimum Weight Triangulations

First, we describe a basic strategy for our algorithm. The details are then discussed in Sections 5.3.2 and 5.3.3.

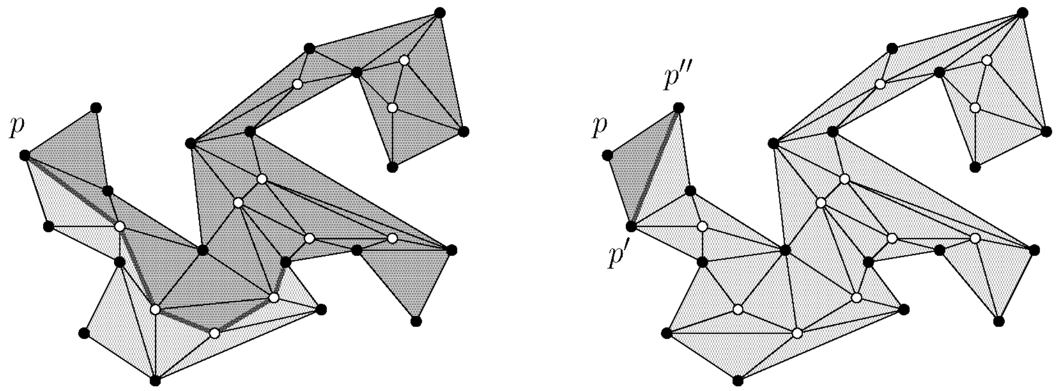

(a) Case (1).

(b) Case (2).

Figure 5.2: Situations in Observation 5.2.

5.3.1 Basic Strategy

An *inner path* of a pointgon $\Pi = (S, P)$ is a polygonal chain $\gamma = \bigcup_{i=0}^{\ell-1} \overline{p_i p_{i+1}}$ such that $p_0, \ldots, p_\ell$ are all distinct, $p_i \in S$ for each $i \in \{1, \ldots, \ell-1\}$, $p_0, p_\ell \in \mathsf{Vert}(P)$, and $\gamma \setminus \{p_0, p_\ell\}$ is contained in the interior of P. An inner path $\bigcup_{i=0}^{\ell-1} \overline{p_i p_{i+1}}$ is called *x-monotone* if the x-coordinates of $p_0, \ldots, p_\ell$ are either increasing or decreasing.

The basic fact we are going to use is the following.

Observation 5.2. *Let $\Pi = (S, P)$ be a pointgon and p be a vertex of Π with the smallest x-coordinate. Denote by p', p'' the neighbors of p in P. Then, for every triangulation $\mathcal{T}$ of Π, either*

(1) there exists a non-selfintersecting x-monotone inner path starting from p and consisting of edges of $\mathcal{T}$, or

(2) the three points p, p', p'' form a triangle of $\mathcal{T}$.

The situation in Observation 5.2 is illustrated in Figure 5.2.

We would like to invoke Observation 5.2 for our algorithm.

Let $\Pi = (S, P)$ be a pointgon, and $p \in \mathsf{Vert}(P)$ the vertex with the smallest x-coordinate. A non-selfintersecting x-monotone inner path divides a pointgon into two smaller pointgons. (See Figure 5.2(a) and recall the general position assumption.) Hence, by looking at all non-selfintersecting x-monotone inner paths starting from P, we can recursively solve the minimum weight triangulation problem. To establish an appropriate recursive formula, denote by $\mathcal{D}(p)$ the set that consists of the line segment $\overline{p'p''}$ and of all non-selfintersecting x-monotone inner paths starting from p. Each non-selfintersecting inner path $\gamma \in \mathcal{D}(p)$ divides our pointgon Π into two smaller pointgons, say Π'_γ and Π''_γ. Then, we can see that

$$\mathsf{mwt}(\Pi) = \min_{\gamma \in \mathcal{D}(p)} \{\mathsf{mwt}(\Pi'_\gamma) + \mathsf{mwt}(\Pi''_\gamma) - \mathsf{length}(\gamma)\}. \tag{5.1}$$

To see that Equality (5.1) is really true, the following observation should be explicitly mentioned, although the proof is straightforward and thus omitted.

Observation 5.3. *Let $\Pi = (S, P)$ be a pointgon and $\mathcal{T}$ be a minimum weight triangulation of Π. Choose an inner path γ which uses edges of $\mathcal{T}$ only, and let Π' and Π'' be two smaller pointgons obtained by subdividing Π with respect to γ. Then, the restriction of $\mathcal{T}$ to Π' is a minimum weight triangulation of Π'. The same holds for Π'' as well.*

Therefore, by solving Recursion (5.1) with an appropriate boundary (or initial) condition, we can obtain a minimum weight triangulation of Π. Note that even if Π is a convex pointgon, the pointgons Π'_γ and Π''_γ encountered in the recursion might not be convex. Thus, our digression to simple polygons is essential also for the minimum weight triangulation problem for a finite point set, i.e., a convex pointgon.

5.3.2 Outline of the Algorithm

Now, we describe how to solve Recursion (5.1) with the dynamic-programming technique.

First, let us label the elements of $\mathsf{Vert}(P)$ in a cyclic order, i.e., the order following the appearance along the boundary of P. According to this order, let us denote $\mathsf{Vert}(P) = \{p_0, p_1, \ldots, p_{n-k-1}\}$. Then, pick a vertex $p_i \in \mathsf{Vert}(P)$, and consider a non-selfintersecting x-monotone inner path γ starting from p_i. Let $p_j \in \mathsf{Vert}(P)$ be the other endpoint of γ. Note that $\mathsf{Vert}(\gamma) \setminus \{p_i, p_j\}$ consists of inner points of Π only. Therefore, such a path can be uniquely specified by a subset $T \subseteq S$.

That is, we associate a triple (i,j,T) with an x-monotone inner path $\overline{p_i q_1} \cup \overline{q_1 q_2} \cup \cdots \cup \overline{q_{t-1} q_t} \cup \overline{q_t p_j}$ where $T = \{q_1, q_2, \ldots, q_t\}$. For the sake of brevity we write $\gamma(T)$ to denote the inner path associated with T when the endpoints p_i, p_j are clear from the context.

For two vertices $p_i, p_j \in \mathsf{Vert}(P)$ on the boundary of Π, and a set $T \subseteq S$ of inner points, let $\Pi(i,j,T)$ be the pointgon obtained from Π as follows: the boundary polygon is the union of the polygonal chains $\bigcup_{\ell=i}^{j-1} \overline{p_\ell p_{\ell+1}}$ and $\gamma(T)$. (Note that we only consider the case where $\gamma(T)$ is well-defined, that is, it does not intersect the boundary polygon.) The inner points of $\Pi(i,j,T)$ consist of the inner points of Π contained in the boundary polygon specified above. Furthermore, denote by $\mathsf{mwt}(i,j,T)$ the weight of a minimum weight triangulation of the pointgon $\Pi(i,j,T)$. Then, Equality (5.1) can be rewritten in the following way if we take p_0 for the role of p:

$$\begin{aligned}
&\mathsf{mwt}(\Pi) = \\
&\min\Big\{ \min_{1 \le i < n-k,\, T \subseteq S} \{\mathsf{mwt}(0,i,T) + \mathsf{mwt}(i,0,T) - \mathsf{length}(\gamma(T))\}, \\
&\mathsf{mwt}(1, n-k-1, \emptyset) + d(p_0, p_1) + d(p_0, p_{n-k-1}) \Big\} . \qquad (5.2)
\end{aligned}$$

The number of values considered in the right hand side of Equality (5.2) is $O((n-k)2^k) = O(2^k n)$. Hence, for the computation of $\mathsf{mwt}(\Pi)$ it is sufficient to know $\mathsf{mwt}(i,j,T)$ for every triple (i,j,T) of two indices $i,j \in \{0, \ldots, n-k-1\}$ and a subset $T \subseteq S$. Since the number of such triples is $O(2^k n^2)$, the efficient computation of each value results in fixed-parameter tractability of the minimum weight triangulation problem.

Nevertheless, to compute these values, we have to generalize the class of pointgons under consideration. That is because pointgons we encounter in the recursion might not be of the form $\Pi(i,j,T)$. Therefore we introduce two additional types of pointgons.

First of all, we call a pointgon which can be defined as the form of $\Pi(i,j,T)$ a *type-1 pointgon* in the following. See Figure 5.3(a) for illustration.

Another class of pointgons is defined for two indices $i,j \in \{0,...,n-k-1\}$, two disjoint subsets $T_1, T_2 \subseteq S$, and a vertex $r \in \mathsf{Vert}(\Pi)$. Then, $\Pi(i,j,T_1,T_2,r)$ is a pointgon bounded by the x-monotone path

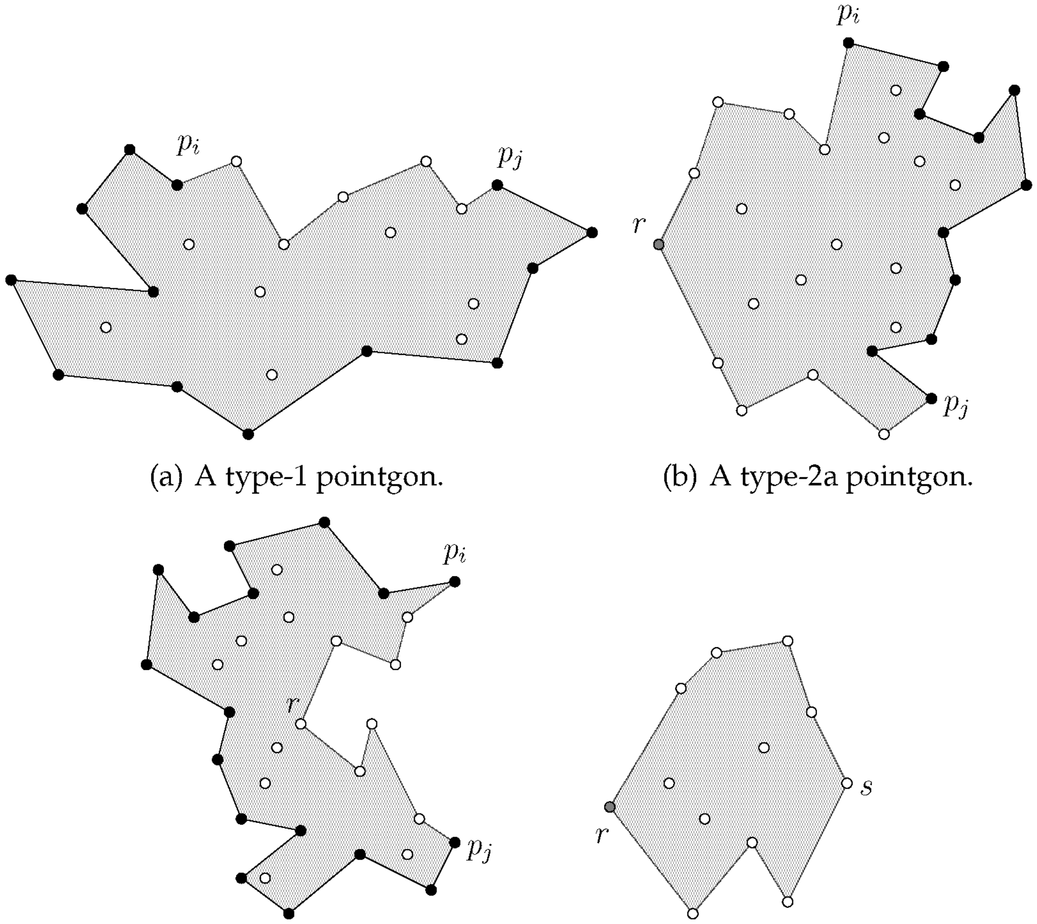

(a) A type-1 pointgon. (b) A type-2a pointgon.

(c) A type-2b pointgon. (d) A type-3 pointgon.

Figure 5.3: The three types of subpointgons of Π. The vertex r belongs either to $\mathsf{Vert}(P)$ or S.

connecting p_i and r through T_1, the x-monotone path connecting p_j and r through T_2, and $\bigcup_{\ell=i}^{j-1} \overline{p_\ell p_{\ell+1}}$. (Again we only consider those tuples which are well-defined, that is, where the paths described above are indeed x-monotone and do not cross each other.) We call such a pointgon a *type-2 pointgon* of Π, and divide them into two subclasses according to whether r is a convex (type-2a) or reflex (type-2b) vertex of the pointgon. Figures 5.3(b) & 5.3(c) illustrate the definition.

The last kind of pointgons uses at most one vertex of P. For a vertex $r \in \mathsf{Vert}(\Pi)$ and two subsets $T_1, T_2 \subseteq S$ with $T_1 \cap T_2 = \{s\}$, we define the pointgon $\Pi(T_1, T_2, r)$ as one which is bounded by two x-monotone paths connecting r and s through T_1 and through T_2, respectively. We call such a pointgon a *type-3 pointgon* of Π. See Figure 5.3(d) for an example.

Let us count the number of these pointgons. The number of type-1 pointgons is $O(2^k n^2)$; the number of type-2 pointgons is $O(3^k n^3)$; the number of type-3 pointgons is $O(3^k n)$. Therefore, the total number of these pointgons is $O(3^k n^3)$. Our goal in the following is to compute the weights of minimum weight triangulations of these pointgons efficiently. Denote by $\mathsf{mwt}(i,j,T)$ the weight of a minimum weight triangulation of a type-1 pointgon $\Pi(i,j,T)$. Similarly, we define $\mathsf{mwt}(i,j,T_1,T_2,r)$ and $\mathsf{mwt}(T_1,T_2,r)$ for type-2 and type-3 pointgons, respectively.

Before describing the algorithm, let us discuss why we encounter these three types of pointgons only in the recursion. For this, we have to be careful which vertex to choose as p in the recursion step. Recall that in any step of Recursion (5.1) there are two cases: either p is cut off by joining its neighbors by an edge, or the pointgon is subdivided by an x-monotone inner path starting from p. Also recall that in Observation 5.2 we required p to be the leftmost point of the pointgon. If we apply the same argument as in Observation 5.2 to an arbitrary vertex of the pointgon, in the first case there appears an inner path starting from p that is *almost x-monotone*, i.e., x-monotone except for the first edge incident to p.

Initially we have a given pointgon $\Pi = (S, P)$ and choose the leftmost vertex as p. If p is cut off (Figure 5.4(a)) the result is a type-1 pointgon where $T = \emptyset$. Otherwise, any x-monotone inner path starting from p divides the pointgon into two type-1 pointgons (Figure 5.4(b)).

When we apply Recursion (5.1) to a type-1 pointgon $\Pi(i,j,T)$, we choose as p the leftmost vertex of the inner path $\gamma(T)$ (which might just consist of a single edge joining p_i and p_j). If p is cut off, the result is either again a type-1 pointgon (Figure 5.5(a)) or a type-2a pointgon (Figure 5.5(b)). Otherwise, consider the vertex q on $\gamma(T)$ next to p. In every triangulation, the edge $\overline{pq}$ must belong to some triangle. To make such a triangle we need another vertex, say z. Let us choose z to be such that $\overline{pz}$ is the first edge of an almost x-monotone inner path γ' starting from p. If $z \in \mathsf{Vert}(P)$, then we get a type-1 pointgon, the triangle pqz and a type-2a pointgon when z is right of p (Figure 5.6(a)), or a type-1 pointgon, the triangle pqz and a type-1 pointgon when z is left of p (Figure 5.6(b)). If $z \in S$, then we have four subcases. When z

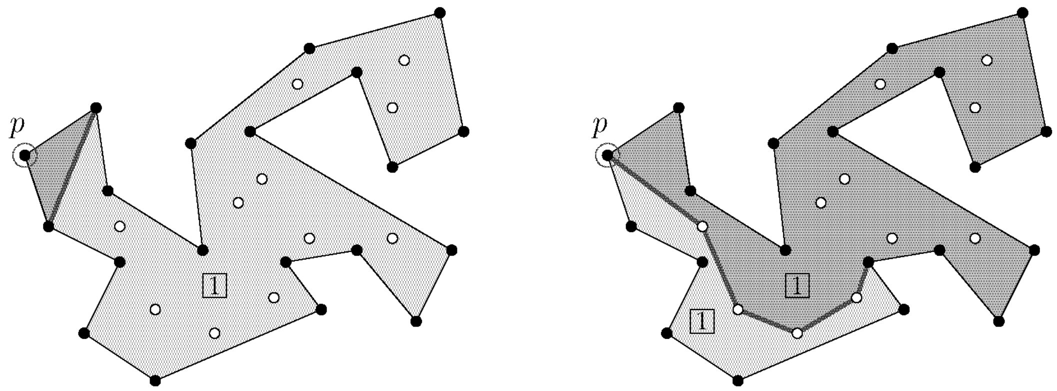

(a) The vertex p is cut off. (b) An inner path from p.

Figure 5.4: Subdivisions obtained from Π. From now on in the pictures, the vertex p is indicated by a larger circle, and the numbers enclosed by squares represent the types of subpointgons obtained by the corresponding subdivisions.

is right of p and γ' ends at a vertex of $\gamma(T)$, we get a type-1 pointgon, the triangle pqz and a type-3 pointgon (Figure 5.7(a)). When z is right of p and γ' ends at a vertex of P, we get a type-1 pointgon, the triangle pqz and a type-2a pointgon (Figure 5.7(b)). When z is left of p and γ' ends at a vertex of $\gamma(T)$, we get a type-2b pointgon, the triangle pqz and a type-3 pointgon (Figure 5.8(a)). When z is left of p and γ' ends at a vertex of P, we get a type-2b pointgon, the triangle pqz and a type-2a pointgon (Figure 5.8(b)).

We choose r as p when we apply Recursion (5.1) to a type-2a pointgon $\Pi(i, j, T_1, T_2, r)$. If p is cut off, then the result is either again a type-2a pointgon or a type-1 pointgon (Figure 5.9(a)). Otherwise, consider an x-monotone inner path starting from p. If the path ends at a vertex of P, we get two type-2a pointgons (Figure 5.9(b)). If, on the other hand, the inner path ends at a vertex in S, then it subdivides the pointgon into a type-2a and a type-3 pointgons (Figure 5.9(c)).

When we apply Recursion (5.1) to a type-2b pointgon, we choose as p the leftmost vertex of the inner path. Since p is a reflex vertex, p cannot be cut off. So, every x-monotone inner path starting from p subdivides the pointgon into two type-1 pointgons (Figure 5.9(d)).

When we apply Recursion (5.1) to a type-3 pointgon $\Pi(T_1, T_2, r)$, we choose r as p. Then, no matter how we divide the pointgon by the op-

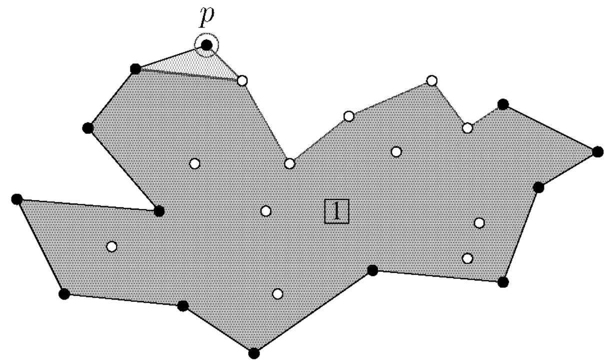

(a) The vertex p is cut off and a neighbor of p is left of p.

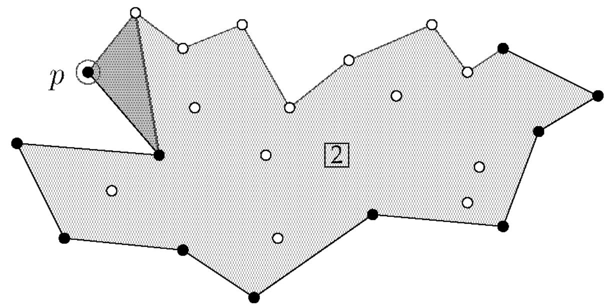

(b) The vertex p is cut off and the neighbors of p are right of p.

Figure 5.5: Subdivisions obtained from a type-1 pointgon.

erations in the recursion, the result again consists of type-3 pointgons (Figure 5.10).

So much for preparation, and now we are ready to give the outline of our algorithm.

Step 1: enumerate all possible type-1 pointgons $\Pi(i,j,T)$, type-2 pointgons $\Pi(i,j,T_1,T_2,r)$, and type-3 pointgons $\Pi(T_1,T_2,r)$.

Step 2: compute the values $\mathsf{mwt}(i,j,T)$, $\mathsf{mwt}(i,j,T_1,T_2,r)$, and $\mathsf{mwt}(T_1,T_2,r)$ for some of them, which are sufficient for Step 3, by dynamic programming.

Step 3: compute $\mathsf{mwt}(\Pi)$ according to Equality (5.2).

We already argued that Step 3 can be done in $O(2^k n)$ time. In the next section we will show that Steps 1 & 2 can be done in $O(6^k n^5 \log n)$ time, which dominates the overall running time.

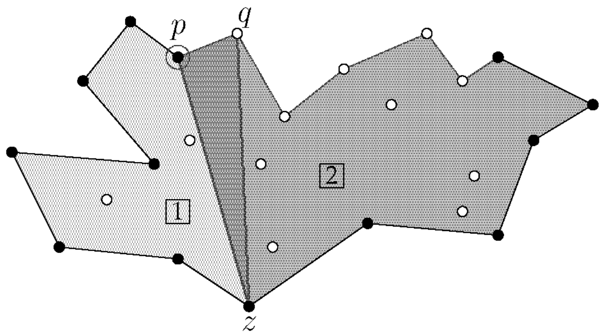

(a) The vertex z belongs to $\mathsf{Vert}(P)$ and z is right of p.

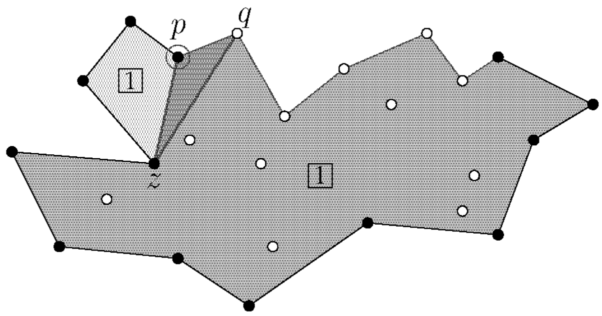

(b) The vertex z belongs to $\mathsf{Vert}(P)$ and z is left of p.

Figure 5.6: Subdivisions obtained from a type-1 pointgon (continued).

5.3.3 Dynamic Programming

Now, we are going to explain how to compute the values of $\mathsf{mwt}(i,j,T)$, $\mathsf{mwt}(i,j,T_1,T_2,r)$, and $\mathsf{mwt}(T_1,T_2,r)$ for all possible choices of i,j,T_1,T_2,r.

First we enumerate all possibilities of i,j,T_1,T_2,r. Each of them can be enumerated in O(1) time, and each of them can be identified as a well-defined pointgon or not (i.e., the inner paths do not intersect each other nor the boundary) in $O(n \log n)$ time. (Apply the standard line segment intersection algorithm [SH76].) Therefore, they can be enumerated in $O(3^k n^3 \cdot 1 \cdot n \log n) = O(3^k n^4 \log n)$ time. This completes Step 1 of our algorithm.

Then, we perform the dynamic programming. For each pointgon enumerated in Step 1, determine the vertex p and consider all possible subdivisions with respect to p as described in the previous section.

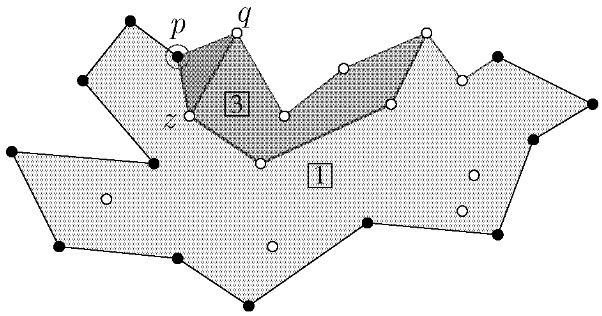

(a) The vertex z is right of p and the path γ' ends at a vertex in S.

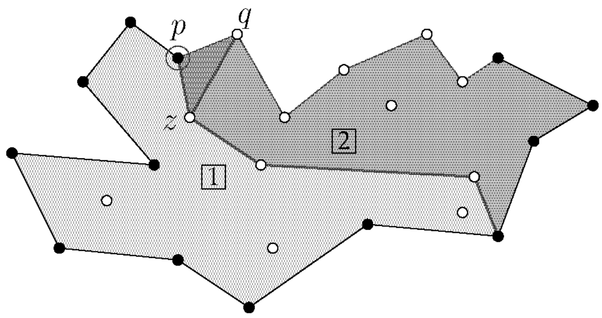

(b) The vertex z is right of p and the path γ' ends at a vertex of P.

Figure 5.7: Subdivisions obtained from a type-1 pointgon (continued). The vertex z belongs to S.

Each subdivision replaces Π by two smaller pointgons. Then, as we have argued in the previous section, these two pointgons can be found among those enumerated in Step 1.

We can associate a parent-child relation between two pointgons Π_1, Π_2 in our enumeration: Π_1 is a parent of Π_2 if Π_2 is obtained as a smaller pointgon when we subdivide Π_1 by a path starting from p (which is fixed as in the previous section) or through the edge cutting off p. It can also be thought as defining a directed graph on the enumerated pointgons: namely, draw a directed edge from Π_1 to Π_2 if the same condition as above is satisfied.

Observe that if Π_1 is a parent of Π_2, then the number of inner points in Π_2 is less than that in Π_1 or $|T_1| + |T_2|$ is smaller in Π_2 than in Π_1. Therefore, the parent-child relation is well-defined (i.e., there is no directed cycle in the directed-graph formulation).

Now, to do the bottom-up computation, we first look at the lowest

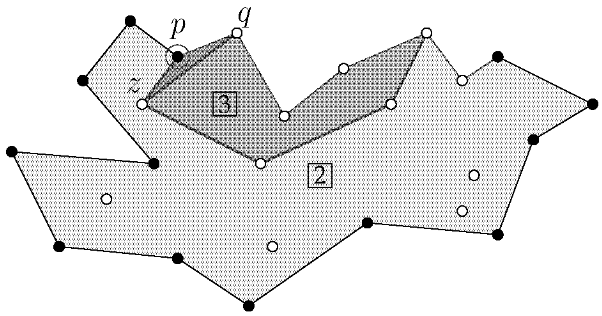

(a) The vertex z is left of p and the path γ' ends at a vertex in S.

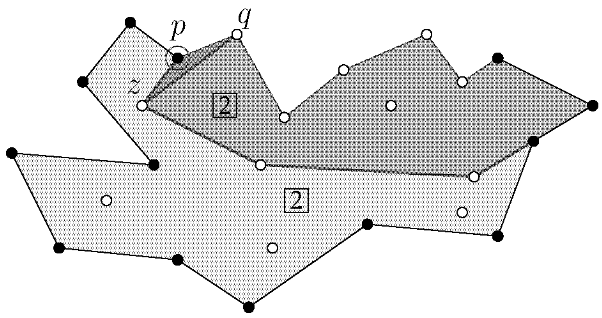

(b) The vertex z is left of p and the path γ' ends at a vertex of P.

Figure 5.8: Subdivisions obtained from a type-1 pointgon (continued). The vertex z belongs to S.

descendants (or the sinks in the directed-graph formulation). They are triangles. So, the weights can be easily computed in constant time. Then, we proceed to their parents. For each parent, we look up the values of its children. In this way, we go up to the highest ancestor, which is Π. Thus, we can compute $\mathsf{mwt}(\Pi)$.

What is the time complexity of the computation? First, let us estimate the time for the construction of the parent-child relation. The number of enumerated pointgons is $O(3^k n^3)$. For each of them, the number of possible choices of non-selfintersecting x-monotone paths is $O(2^k n)$. For each of the paths, we can decide whether it really defines a non-selfintersecting path in $O(n \log n)$ time. Therefore, the overall running time for the construction is $O(3^k n^3 \cdot 2^k n \cdot n \log n) = O(6^k n^5 \log n)$.

In the bottom-up computation, for each pointgon we look up at most $O(2^k n)$ entries and compute the value according to Equality (5.1). Therefore, this can be done in $O(3^k n^3 \cdot 2^k n) = O(6^k n^4)$.

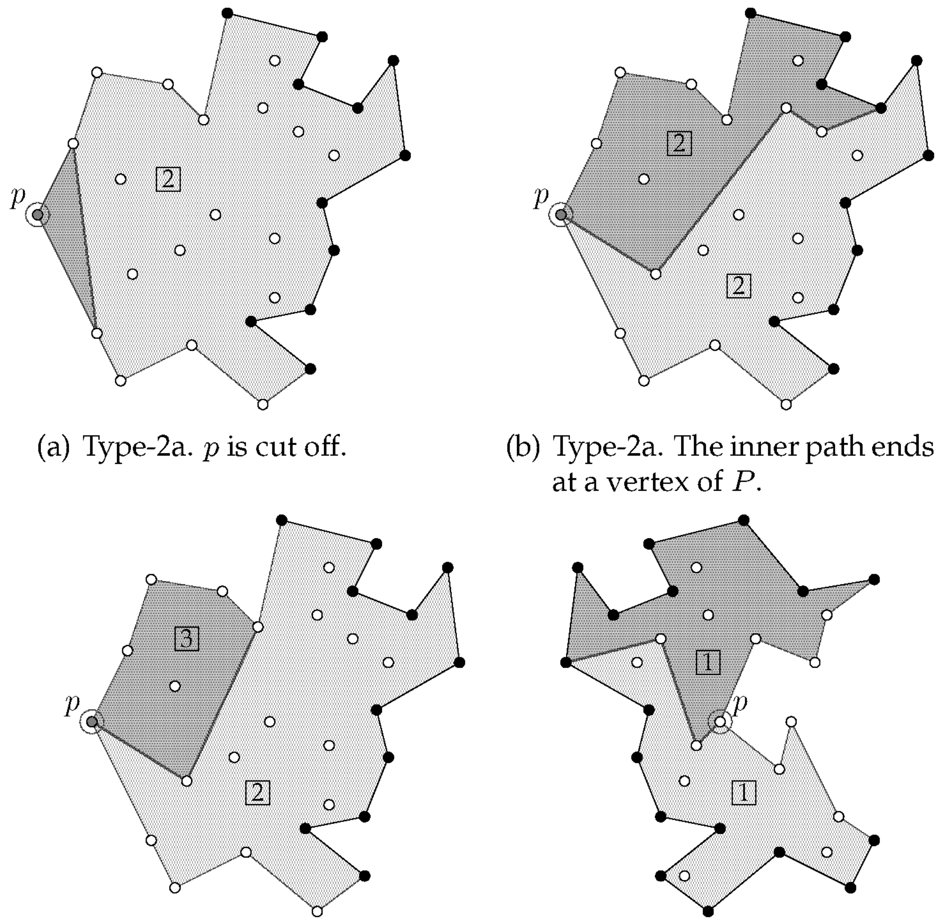

(a) Type-2a. p is cut off.

(b) Type-2a. The inner path ends at a vertex of P.

(c) Type-2a. The inner path ends at a vertex in S.

(d) Type-2b.

Figure 5.9: Subdivisions obtained from a type-2 pointgon.

Hence, the overall running time of the algorithm is $\mathrm{O}(3^k n^4 \log n + 6^k n^5 \log n + 6^k n^4) = \mathrm{O}(6^k n^5 \log n)$. This completes the proof of Theorem 5.1.

5.4 Conclusion

In this chapter, we studied the minimum weight triangulation problem from the viewpoint of parameterized computation. We established an algorithm to solve this problem for a simple polygon with some inner points. The running time is $\mathrm{O}(6^k n^6 \log n)$ when n is the total number of input points and k is the number of inner points. Therefore, the problem is fixed-parameter tractable with respect to the number

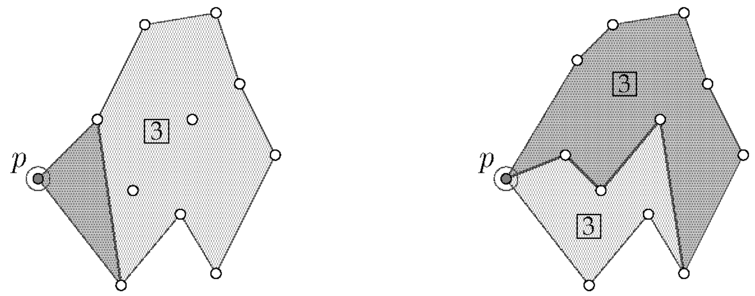

(a) The vertex p is cut off.

(b) An inner path from p.

Figure 5.10: Subdivisions obtained from a type-3 pointgon.

of inner points. We believe the number of inner points in geometric optimization problems plays a role similar to the treewidth in graph optimization problems.

Since our algorithm is based on a simple idea, it can be extended in several ways. For example, we can also compute a maximum weight triangulation in the same time complexity. (It seems quite recent that attention has been paid to maximum weight triangulations [Hu03, WCY99].) To do that, we just replace "min" in Equalities (5.1) and (5.2) by "max." By a similar idea, we can also compute a triangulation which minimizes the length of a longest edge, which maximizes the length of a shortest edge, which minimizes the area of a largest triangle, which maximizes the area of a smallest triangle (studied by Keil & Vassilev [KV03]), which minimizes the largest angle, and so on. Another direction of extension is to incorporate some heuristics. For example, there are some known pairs of vertices which appear as edges in all minimum weight triangulations, e.g. the β-skeleton for some β and the LMT-skeleton; see [BDE02, CX01, WY01] and the references therein. Because of the flexibility of our algorithm, we can insert these pairs at the beginning of the execution as edges, and proceed in the same way except that we can use the information from these prescribed edges.

The obvious open problem is to improve the time complexity of our algorithm. For example, is it possible to provide a polynomial-time algorithm for the minimum weight triangulation problem when $k = \mathrm{O}(\log^2 n)$?

Bibliography

The references are sorted alphabetically by the abbreviations (rather than by the authors' names).

[ABM03] Ron Aharoni, Eli Berger, and Roy Meshulam. Eigenvalues and homology of flag complexes and vector representations of graphs. Available at `http://www.arxiv.org/abs/math.CO/0312482`, December 2003. (Cited on pages 7, 17)

[AC93] Efthymios Anagnostou and Derek Corneil. Polynomial-time instances of the minimum weight triangulation problem. *Computational Geometry: Theory and Applications*, 3(5):247–259, 1993. (Cited on page 124)

[ACPS93] Stefan Arnborg, Bruno Courcelle, Andrej Proskurowski, and Detlef Seese. An algebraic theory of graph reduction. *Journal of the Association for Computing Machinery*, 40(5):1134–1164, 1993. (Cited on page 3)

[AGM99] Christina Ahrens, Gary Gordon, and Elizabeth W. McMahon. Convexity and the beta invariant. *Discrete & Computational Geometry*, 22(3):411–424, 1999. (Cited on page 85)

[Aic99] Oswin Aichholzer. The path of a triangulation. In *Proceedings of 15th ACM Annual Symposium on Computational Geometry (SCG '99)*, pages 14–23, New York, 1999. Association for Computing Machinery. (Cited on page 124)

[ALS91] Stefan Arnborg, Jens Lagergren, and Detlef Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2):308–340, 1991. (Cited on page 3)

[And02] Kazutoshi Ando. Extreme point axioms for closure spaces. Discussion Paper 969, Institute of Policy and Planning Sciences, University of Tsukuba, 2002. (Cited on pages 9, 88)

[Aro98] Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998. (Cited on page 109)

[ARSS03] Oswin Aichholzer, Günter Rote, Bettina Speckmann, and Ileana Streinu. The zigzag path of a pseudo-triangulation. In Frank Dehne, Jörg-Rüdiger Sack, and Michiel Smid, editors, *Proceedings of 8th International Workshop on Algorithms and Data Structures (WADS 2003)*, volume 2748 of *Lecture Notes in Computer Science*, pages 377–388, New York, 2003. Springer-Verlag. (Cited on page 125)

[Bal89] Egon Balas. The prize collecting traveling salesman problem. *Networks*, 19(6):621–636, 1989. (Cited on page 108)

[BD91] Ernest F. Brickell and Daniel M. Davenport. On the classification of ideal secret sharing schemes. *Journal of Cryptology*, 4(2):123–134, 1991. (Cited on page 6)

[BDE02] Prosenjit Bose, Luc Devroye, and William Evans. Diamonds are not a minimum weight triangulation's best friend. *International Journal of Computational Geometry & Applications*, 12(6):445–453, 2002. (Cited on page 138)

[Bel62] Richard Bellman. Dynamic programming treatment of the travelling salesman problem. *Journal of the Association for Computing Machinery*, 9(1):61–63, 1962. (Cited on pages 109, 115)

[BF90] E. Andrew Boyd and Ulrich Faigle. An algorithmic characterization of antimatroids. *Discrete Applied Mathematics*, 28(3):197–205, 1990. (Cited on pages 8, 60)

[BKMS] Jürgen Bokowski, Simon King, Susanne Mock, and Ileana Streinu. A topological representation theorem for oriented matroids. *Discrete & Computational Geometry*. To appear. Available at `http://www.arxiv.org/abs/math.CO/0209364`. (Cited on page 60)

[BLS+99] Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler. *Oriented Matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2nd edition, 1999. (Cited on page 59)

[BMS01] Jürgen Bokowski, Susanne Mock, and Ileana Streinu. On the Folkman-Lawrence topological representation theorem for oriented matroids of rank 3. *European Journal of Combinatorics*, 22(5):601–615, 2001. (Cited on page 60)

[Bod93] Hans L. Bodlaender. A tourist guide through treewidth. *Acta Cybernetica*, 11(1–2):1–22, 1993. (Cited on page 3)

[Bod96] Hans L. Bodlaender. A linear time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996. (Cited on page 2)

[Bod97] Hans L. Bodlaender. Treewidth: Algorithmic techniques and results. In Igor Prívara and Peter Ruzicka, editors, *Proceedings of 22nd Mathematical Foundations of Computer Science (MFCS'97)*, volume 1295 of *Lecture Notes in Computer Science*, pages 19–36. Springer-Verlag, 1997. (Cited on page 3)

[Bod05] Hans L. Bodlaender. Discovering treewidth. In Peter Vojtás, Mária Bieliková, Bernadette Charron-Bost, and Ondrey Sýkora, editors, *Proceedings of 31st Annual Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM 2005)*, volume 3381 of *Lecture Notes in Computer Science*, pages 1–16, Berlin Heidelberg, 2005. Springer-Verlag. (Cited on page 3)

[Bol95] Béla Bollobás. Extremal graph theory. In Ronald L. Graham, Martin Grötschel, and László Lovász, editors, *Handbook of Combinatorics*, chapter 23, pages 1231–1292. Elsevier, Amsterdam, 1995. (Cited on pages 7, 16)

[BvAdF01] Hans L. Bodlaender and Babette van Antwerpen-de Fluiter. Reduction algorithms for graphs of small treewidth. *Information and Computation*, 167(2):86–119, 2001. (Cited on page 3)

[BZ92] Anders Björner and Günter M. Ziegler. Introduction to greedoids. In Neil White, editor, *Matroid Applications*, volume 40, chapter 8, pages 284–357. Cambridge University Press, Cambridge, 1992. (Cited on page 60)

[Cai03] Leizhen Cai. Parameterized complexity of vertex colouring. *Discrete Applied Mathematics*, 127(3):415–429, 2003. (Cited on page 11)

[CCP96] Leizhen Cai, Derek Corneil, and Andrzej Proskurowski. A generalization of line graphs: (X,Y)-intersection graphs. *Journal of Graph Theory*, 21(3):267–287, 1996. (Cited on page 44)

[CD95] Ruth Charney and Michael Davis. The Euler characteristic of a nonpositively curved, piecewise linear Euclidean manifold. *Pacific Journal of Mathematics*, 171(1):117–137, 1995. (Cited on pages 7, 16)

[CER93] Bruno Courcelle, Joost Engelfriet, and Grzegorz Rozenberg. Handle-rewriting hypergraph grammars. *Journal of Computer and System Sciences*, 46(2):218–270, 1993. (Cited on page 3)

[Cha01] Bernard Chazelle. *The Discrepancy Method: Randomness and Complexity*. Cambridge University Press, Cambridge, paperback edition, 2001. (Cited on page 4)

[CLRS01] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to algorithms*. The MIT Press, Cambridge, 2nd edition, 2001. (Cited on page 112)

[CM96] Bernard Chazelle and Jiří Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *Journal of Algorithms*, 21(3):579–597, 1996. (Cited on page 4)

[CMR00] Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, 2000. (Cited on page 3)

[CO00] Bruno Courcelle and Stephan Olariu. Upper bounds to the clique width of graphs. *Discrete Applied Mathematics*, 101(1–3):77–114, 2000. (Cited on page 3)

[Cou91] Bruno Courcelle. The monadic second-order logic of graphs I: Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1991. (Cited on page 2)

[CS04] Maria Chudnovsky and Paul Seymour. The roots of the stable set polynomial of a clawfree graph. Submitted for publication, 2004. (Cited on pages 7, 16)

[CX01] Siu-Wing Cheng and Yin-Feng Xu. On β-skeleton as a subgraph of the minimum weight triangulation. *Theoretical Computer Science*, 262(1–2):459–471, 2001. (Cited on page 138)

[Dez92] Michel Deza. Perfect matroid designs. In Neil White, editor, *Matroid Applications*, volume 40, chapter 2, pages 284–357. Cambridge University Press, Cambridge, 1992. (Cited on page 6)

[DF99a] Jean-Paul Doignon and Jean-Claude Falmagne. *Knowledge Spaces*. Springer-Verlag, Berlin, 1999. (Cited on pages 9, 11, 60)

[DF99b] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer-Verlag, New York, 1999. (Cited on pages 108, 123)

[DFS99a] Rodney G. Downey, Michael R. Fellows, and Ulrike Stege. Computational tractability: The view from Mars. *Bulletin of the European Association for Theoretical Computer Science*, 69:73–97, 1999. (Cited on page 11)

[DFS99b] Rodney G. Downey, Michael R. Fellows, and Ulrike Stege. Parameterized complexity: A framework for systematically confronting computational intractability. In Ronald L. Graham, Jan

Kratochvíl, Jaroslav Nešetřil, and Fred S. Roberts, editors, *Contemporary Trends in Discrete Mathematics*, volume 49 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 49–99. American Mathematical Society, Providence, 1999. (Cited on page 11)

[DH05a] Erik D. Demaine and MohammadTaghi Hajiaghayi. Bidimensionality: New connections between FPT algorithms and PTASs. In *Proceedings of 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 590–601, New York, 2005. Association for Computing Machinery. (Cited on page 3)

[DH05b] Erik D. Demaine and MohammadTaghi Hajiaghayi. Fast algorithms for hard graph problems: Bidimensionality, minors, and local treewidth. In János Pach, editor, *Proceedings of 12th International Symposium on Graph Drawing (GD 2004)*, volume 3383 of *Lecture Notes in Computer Science*, pages 517–533, Berlin, 2005. Springer-Verlag. (Cited on page 3)

[DHOW04] Vladimir G. Deĭneko, Michael Hoffmann, Yoshio Okamoto, and Gerhard J. Woeginger. The traveling salesman problem with few inner points. In Kyung-Yong Chwa and J. Ian Munro, editors, *Proceedings of 10th International Computing and Combinatorics Conference (COCOON 2004)*, volume 3106 of *Lecture Notes in Computer Science*, pages 268–277, Berlin, 2004. Springer-Verlag. A journal version will appear in Operations Research Letters. (Cited on page 12)

[Die87] Brenda L. Dietrich. A circuit set characterization of antimatroids. *Journal of Combinatorial Theory Series B*, 43(3):314–321, 1987. (Cited on page 74)

[Die89] Brenda L. Dietrich. Matroids and antimatroids — a survey. *Discrete Mathematics*, 78(3):223–237, 1989. (Cited on pages 60, 74)

[Die00] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 2000. (Cited on page 2)

[Dir61] Gabriel A. Dirac. On rigid circuit graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 25:71–76, 1961. (Cited on page 97)

[DM04] Rodney G. Downey and Catherine McCartin. Some new directions and questions in parameterized complexity. In Cristian S. Calude, Elena Calude, and Michael J. Dinneen, editors, *Proceedings of 8th International Conference on Developments in Language Theory (DLT 2004)*, volume 3340 of *Lecture Notes in Computer Science*, pages 12–26, Berlin Heidelberg, 2004. Springer-Verlag. (Cited on page 11)

[Dow03] Rodney G. Downey. Parameterized complexity for the skeptic. In *Proceedings of 18th IEEE Conference on Computational Complexity (CCC 2003)*, pages 147–168, Los Alamitos, 2003. IEEE Computer Society Press. (Cited on page 11)

[DvDR94] Vladimir Deĭneko, René van Dal, and Günter Rote. The convex-hull-and-line traveling salesman problem: A solvable case. *Information Processing Letters*, 51(3):141–148, 1994. (Cited on page 109)

[DW96] Vladimir G. Deĭneko and Gerhard J. Woeginger. The convex-hull-and-k-line traveling salesman problem. *Information Processing Letters*, 59(6):295–301, 1996. (Cited on page 109)

[Edm71] Jack Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1:127–136, 1971. (Cited on page 16)

[EJ85] Paul H. Edelman and Robert E. Jamison. The theory of convex geometries. *Geometriae Dedicata*, 19(3):247–270, 1985. (Cited on pages 8, 9, 60, 62, 63, 64, 69, 98)

[ER00] Paul H. Edelman and Victor Reiner. Counting the interior points of a point configuration. *Discrete & Computational Geometry*, 23(1):1–13, 2000. (Cited on pages 8, 9, 60, 83, 85, 86, 91, 98)

[Eri04] Jeff Erickson. JCDCG highlights. Ernie's 3D Pancakes, October 12 2004. `http://3dpancakes.typepad.com/ernie/2004/10/jcdcg_highlight.html`. (Cited on page 161)

[ERW02] Paul H. Edelman, Victor Reiner, and Volkmar Welker. Convex, acyclic, and free sets of an oriented matroid. *Discrete & Computational Geometry*, 27(1):96–116, 2002. (Cited on pages 8, 60, 85, 86)

[Fal89] Jean-Claude Falmagne. A latent trait theory via a stochastic learning theory for a knowledge space. *Psychometrika*, 54(2):283–303, 1989. (Cited on page 9)

[FD88] Jean-Claude Falmagne and Jean-Paul Doignon. A Markovian procedure for assessing the state of a system. *Journal of Mathematical Psychology*, 32(3):232–258, 1988. (Cited on page 9)

[Fel01] Michael R. Fellows. Parameterized complexity: The main idea and some research frontiers. In Peter Eades and Tadao Takaoka, editors, *Proceedings of 12th International Symposium on Algorithms and Computation (ISAAC 2001)*, volume 2223 of *Lecture Notes in Computer Science*, pages 291–307, Berlin Heidelberg, 2001. Springer-Verlag. (Cited on page 11)

[Fel02] Michael R. Fellows. Parameterized complexity: The main ideas and connections to practical computing. In *Proceedings of Computing: The Australasian Theory Symposium (CATS) 2002*, volume 61 of *Electronic Notes in Theoretical Computer Science*, pages 1–19, 2002. (Cited on page 11)

[Fel03a] Michael R. Fellows. Blow-ups, win/win's, and crown rules: Some new directions in FPT. In Hans L. Bodlaender, editor, *Proceedings of 29th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2003)*, volume 2880 of *Lecture Notes in Computer Science*, pages 1–12, Berlin Heidelberg, 2003. Springer-Verlag. (Cited on page 11)

[Fel03b] Michael R. Fellows. New directions and new challenges in algorithm design and complexity, parameterized. In Frank Dehne, Jörg-Rüdiger Sack, and Michiel Smid, editors, *Proceedings of 8th International Workshop on Algorithms and Data Structures (WADS 2003)*, volume 2748 of *Lecture Notes in Computer Science*, pages 505–519, Berlin Heidelberg, 2003. Springer-Verlag. (Cited on page 11)

[FFS03] Sándor P. Fekete, Robert T. Firla, and Bianca Spille. Characterizing matchings as the intersection of matroids. *Mathematical Methods of Operations Research*, 58(2):319–329, 2003. (Cited on pages 8, 17, 18, 47, 48, 49, 50, 51, 55)

[FG04] Jörg Flum and Martin Grohe. Parameterized complexity and subexponential time. *Bulletin of the European Association for Theoretical Computer Science*, 84, 2004. To appear. (Cited on page 11)

[FL78] Jon Folkman and Jim Lawrence. Oriented matroids. *Journal of Combinatorial Theory Series B*, 25(2):199–235, 1978. (Cited on pages 59, 82)

[Flo56] Merrill M. Flood. Traveling-salesman problem. *Operations Research*, 4:61–75, 1956. (Cited on page 110)

[Fra81] András Frank. A weighted matroid intersection algorithm. *Journal of Algorithms*, 2(4):328–336, 1981. (Cited on pages 17, 45)

[FT04] Fedor V. Fomin and Dimitrios M. Thilikos. A simple and fast approach for solving problems on planar graphs. In Volker Diekert and Michel Habib, editors, *Proceedings of 21st Annual Symposium on Theoretical Aspects of Computer Science (STACS 2004)*, volume 2996 of *Lecture Notes in Computer Science*, pages 56–67, Berlin, 2004. Springer-Verlag. (Cited on page 3)

[Fuj91] Satoru Fujishige. *Submodular functions and optimization*, volume 47 of *Annals of Discrete Mathematics*. North-Holland, Amsterdam, 1991. (Cited on page 6)

[Fuj04] Satoru Fujishige. Dual greedy polyhedra, choice functions, and abstract convex geometries. *Discrete Optimization*, 1(1):41–49, 2004. (Cited on pages 9, 60)

[GGJ76] Michael R. Garey, Ronald L. Graham, and David S. Johnson. Some NP-complete geometric problems. In *Proceedings of 8th Annual ACM Symposium on Theory of Computing (STOC '76)*, pages 10–22, New York, 1976. Association for Computing Machinery. (Cited on pages 11, 108)

[GHN04] Jiong Guo, Falk Hüffner, and Rolf Niedermeier. A structural view on parameterizing problems: Distance from triviality. In Rod Downey, Michael Fellows, and Frank Dehne, editors, *Proceedings of 1st International Workshop on Parameterized and Exact Computation (IWPEC 2004)*, volume 3162 of *Lecture Notes in Computer Science*, pages 162–173, Berlin, 2004. Springer-Verlag. (Cited on page 11)

[Gil79] P.D. Gilbert. New results in planar triangulations. Master's thesis, University of Illinois Urbana, 1979. (Cited on page 124)

[GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. A Series of Books in the Mathematical Sciences. W.H. Freeman and Co., San Francisco, 1979. (Cited on pages 12, 16, 45, 123, 124)

[GSS93] Jack Graver, Brigitte Servatius, and Herman Servatius. *Combinatorial Rigidity*, volume 2 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 1993. (Cited on page 6)

[GY94] Paul Glasserman and David D. Yao. *Monotone Structure in Discrete-Event Systems*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, New York, 1994. (Cited on pages 9, 60)

[Ham90] Yahya Ould Hamidoune. On the number of independent k-sets in a claw free graph. *Journal of Combinatorial Theory Series B*, 50(2):241–244, 1990. (Cited on pages 7, 16)

[Hås99] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999. (Cited on page 16)

[HCL93] R.Z. Hwang, Ruei-Chuan Chang, and Richard Chia-Tung Lee. The searching over separators strategy to solve some NP-hard problems in subexponential time. *Algorithmica*, 9(4):398–423, 1993. (Cited on page 109)

[HK62] Michael Held and Richard M. Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10:196–210, 1962. (Cited on pages 109, 115)

[HK04] Masahiro Hachimori and Kenji Kashiwabara. On the topology of the free complexes of convex geometries. Submitted for publication, 2004. (Cited on page 104)

[HN04] Masahiro Hachimori and Masataka Nakamura. Rooted circuits of closed-set systems and the max-flow min-cut property of stem clutters of affine convex geometries. Submitted for publication, 2004. (Cited on pages 82, 91)

[HO04] Michael Hoffmann and Yoshio Okamoto. The minimum triangulation problem with few inner points. In Rod Downey, Michael Fellows, and Frank Dehne, editors, *Proceedings of 1st International Workshop on Exact and Parameterized Computation (IWPEC 2004)*, volume 3162 of *Lecture Notes in Computer Science*, pages 200–212, Berlin, 2004. Springer-Verlag. (Cited on page 12)

[HP99] Sariel Har-Peled. Constructing approximate shortest path maps in three dimensions. *SIAM Journal on Computing*, 28(4):1182–1197, 1999. (Cited on page 4)

[Hu03] Shiyan Hu. A constant-factor approximation for maximum weight triangulation. In *Proceedings of 15th Canadian Conference on Computational Geometry (CCCG 2003)*, pages 150–154, 2003. (Cited on page 138)

[HW87] David Haussler and Emo Welzl. ε-nets and simplex range queries. *Discrete & Computational Geometry*, 2(2):127–151, 1987. (Cited on pages 3, 4)

[JD01] Mark R. Johnson and Richard A. Dean. Locally complete path independent choice functions and their lattices. *Mathematical Social Sciences*, 42(1):53–87, 2001. (Cited on pages 9, 60)

[Jen76] Thomas A. Jenkyns. The efficacy of the "greedy" algorithm. In F. Hoffmann, L. Lesniak, R.C. Mullin, K.B. Reid, and R.G. Stanton, editors, *Proceedings of the 7th Southeastern Conference on Combinatorics, Graph Theory, and Computing*, volume 17 of *Congressus Numerantium*, pages 341–350. Utilitas Mathematica Publishing Inc. Winnipeg, 1976. (Cited on pages 7, 16)

[Kal00] Gil Kalai. Combinatorics with a geometric flavor. *Geometric and Functional Analysis*, Special Volume II:742–791, 2000. (Cited on page 161)

[Kan91] Viggo Kann. Maximum bounded 3-dimensional matching is MAX SNP-complete. *Information Processing Letters*, 37(1):27–35, 1991. (Cited on page 45)

[KH78] Bernhard Korte and Dirk Hausmann. An analysis of the greedy algorithm for independence systems. In Brian Alspach, Pavol Hell, and Donald J. Miller, editors, *Algorithmic Aspects of Combinatorics*, volume 2 of *Annals of Discrete Mathematics*, pages 65–74. North-Holland, Amsterdam, 1978. (Cited on pages 7, 16)

[KITT97] Yoshiaki Kyoda, Keiko Imai, Fumihiko Takeuchi, and Akira Tajima. A branch-and-cut approach for minimum weight triangulation. In Hon Wai Leong, Hiroshi Imai, and Sanjay Jain, editors, *Proceedings of 8th International Symposium on Algorithms and Computation (ISAAC '97)*, volume 1350 of *Lecture Notes in Computer Science*, pages 384–393, Berlin, 1997. Springer-Verlag. (Cited on page 124)

[KL84] Bernhard Korte and László Lovász. Shelling structures, convexity and a happy end. In Béla Bollóbas, editor, *Graph Theory and Combinatorics: Proceedings of the Cambridge Combinatorial Conference in Honour of Paul Erdős*, pages 219–232. Academic Press, London, 1984. (Cited on pages 66, 72)

[Kla99] Daniel A. Klain. An Euler relation for valuations on polytopes. *Advances in Mathematics*, 147(1):1–34, 1999. (Cited on page 85)

[Kli80] G.T. Klincsek. Minimal triangulations of polygonal domains. In *Combinatorics 79, Part II*, volume 9 of *Annals of Discrete Mathematics*, pages 121–123, 1980. (Cited on page 124)

[KLMR98] Lydia E. Kavraki, Jean-Claude Latombe, Rajeev Motwani, and Prabhakar Raghavan. Randomized query processing in robot path planning. *Journal of Computer and System Sciences*, 57(1):50–60, 1998. (Cited on page 4)

[KLS91] Bernhard Korte, László Lovász, and Rainer Schrader. *Greedoids*, volume 4 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1991. (Cited on pages 8, 9, 60, 63, 64, 69)

[KM97] Gil Kalai and Jiří Matoušek. Guarding galleries where every point sees a large area. *Israel Journal of Mathematics*, 101:125–140, 1997. (Cited on page 4)

[KNO05] Kenji Kashiwabara, Masataka Nakamura, and Yoshio Okamoto. The affine representation theorem for abstract convex geometries. *Computational Geometry: Theory and Applications*, 30(2):129–144, 2005. (Cited on page 10)

[KO03] Kenji Kashiwabara and Yoshio Okamoto. A greedy algorithm for convex geometries. *Discrete Applied Mathematics*, 131(2):449–465, 2003. (Cited on pages 9, 60)

[Kop98] Mathieu Koppen. On alternative representations for knowledge spaces. *Mathematical Social Sciences*, 36(2):127–143, 1998. (Cited on page 9)

[Kos99] Gleb A. Koshevoy. Choice functions and abstract convex geometries. *Mathematical Social Sciences*, 38(1):35–44, 1999. (Cited on pages 9, 60)

[KOU03] Kenji Kashiwabara, Yoshio Okamoto, and Takeaki Uno. Matroid representation of clique complexes. In Tandy Warnow and Binhai Zhu, editors, *Proceedings of 9th International Computing and Combinatorics Conference (COCOON 2003)*, volume 2697 of *Lecture Notes in Computer Science*, pages 192–201, Berlin, 2003. Springer-Verlag. (Cited on page 8)

[Kru60] Joseph B. Kruskal. Well-quasi-ordering, the tree theorem, and Vazsonyi's conjecture. *Transactions of the American Mathematical Society*, 95(2):210–225, 1960. (Cited on page 2)

[KV02] Bernhard Korte and Jens Vygen. *Combinatorial Optimization: Theory and Algorithms*, volume 21 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2nd edition, 2002. (Cited on page 15)

[KV03] J. Mark Keil and Tzvetalin Vassilev. An algorithm for the MaxMin area triangulation of a convex polygon. In *Proceedings of 15th Canadian Conference on Computational Geometry (CCCG 2003)*, pages 145–149, 2003. (Cited on page 138)

[LK98] Christos Levcopoulos and Drago Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. *Journal of Algorithms*, 27(2):303–338, 1998. (Cited on page 124)

[Mar04] Dániel Marx. Parameterized coloring problems on chordal graphs. In Rod Downey, Michael Fellows, and Frank Dehne, editors, *Proceedings of 1st International Workshop on Parameterized and Exact Computation (IWPEC 2004)*, volume 3162 of *Lecture Notes in Computer Science*, pages 83–95, Berlin, 2004. Springer-Verlag. (Cited on page 11)

[Mat95] Jiří Matoušek. Approximations and optimal geometric divide-and-conquer. *Journal of Computer and System Sciences*, 50(2):203–208, 1995. (Cited on page 4)

[Mat99] Jiří Matoušek. *Geometric Discrepancy: An Illustrated Guide*, volume 18 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin Heidelberg, 1999. (Cited on pages 4, 5)

[Mat02] Jiří Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002. (Cited on page 4)

[Mat03] Jiří Matoušek. *Using the Borsuk–Ulam Theorem: Lectures on Topological Methods in Combinatorics and Geometry*. Universitext. Springer-Verlag, Berlin, 2003. Written in cooperation with Anders Björner and Günter M. Ziegler. (Cited on page 88)

[Mat04] Jiří Matoušek. Bounded VC-dimension implies a fractional Helly theorem. *Discrete & Computational Geometry*, 31(2):251–255, 2004. (Cited on page 4)

[Mes01] Roy Meshulam. The clique complex and hypergraph matching. *Combinatorica*, 21(1):89–94, 2001. (Cited on pages 7, 17)

[Mes03] Roy Meshulam. Domination numbers and homology. *Journal of Combinatorial Theory Series A*, 102(2):321–330, 2003. (Cited on pages 7, 17)

[Mit99] Joseph S.B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999. (Cited on page 109)

[MR01] Bernard Monjardet and Vololonirina Raderanirina. The duality between the anti-exchange closure operators and the path independent choice operators on a finite set. *Mathematical Social Sciences*, 41(2):131–150, 2001. (Cited on pages 9, 60)

[MS04] Jiří Matoušek and Tibor Szabó. RANDOM EDGE can be exponential on abstract cubes. In *Proceedings of 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004)*, pages 92–100, Los Alamitos, 2004. IEEE Computer Society Press. (Cited on page 161)

[MSS04] Rolf H. Möhring, Martin Skutella, and Frederik Stork. Scheduling with AND/OR precedence constraints. *SIAM Journal on Computing*, 33(2):393–415, 2004. (Cited on page 9)

[Mur00] Kazuo Murota. *Matrices and Matroids for Systems Analysis*, volume 20 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2000. (Cited on pages 6, 59)

[Mur03] Kazuo Murota. *Discrete Convex Analysis*. SIAM Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics, Philadelphia, 2003. (Cited on page 6)

[Nar97] H. Narayanan. *Submodular Functions and Electrical Networks*, volume 54 of *Annals of Discrete Mathematics*. North-Holland, Amsterdam, 1997. (Cited on page 6)

[Nie98] Rolf Niedermeier. Some prospects for efficient fixed parameter algorithms. In Branislav Rovan, editor, *Proceedings of 25th Conference on Current Trends in Theory and Practice of Informatics (SOFSEM'98)*, volume 1521 of *Lecture Notes in Computer Science*, pages 168–185, Berlin Heidelberg, 1998. Springer-Verlag. (Cited on page 11)

[Nie02] Rolf Niedermeier. Invitation to fixed-parameter algorithms. Habilitation thesis, Universität Tübingen, 2002. (Cited on pages 3, 11, 108, 123)

[Nie04] Rolf Niedermeier. Ubiquitous parameterization — invitation to fixed-parameter algorithms. In Jiří Fiala, Václav Koubek, and Jan Kratochvíl, editors, *Proceedings of 29th International Symposium on Mathematical Foundations of Computer Science (MFCS 2004)*, volume 3153 of *Lecture Notes in Computer Science*, pages 84–103, Berlin, 2004. Springer-Verlag. (Cited on page 11)

[NW01] Siaw-Lynn Ng and Michael Walker. On the composition of matroids and ideal secret sharing schemes. *Designs, Codes and Cryptography*, 24(1):49–67, 2001. (Cited on page 6)

[Oka03] Yoshio Okamoto. Submodularity of some classes of the combinatorial optimization games. *Mathematical Methods of Operations Research*, 58(1):131–139, 2003. (Cited on page 26)

[Oka04] Yoshio Okamoto. Local topology of the free complex of a two-dimensional generalized convex shelling. Submitted for publication, 2004. (Cited on page 10)

[Oxl92] James G. Oxley. *Matroid Theory*. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1992. (Cited on pages 6, 19, 59)

[Pap77] Christos H. Papadimitriou. The Euclidean travelling salesman problem is NP-complete. *Theoretical Computer Science*, 4(3):237–244, 1977. (Cited on pages 11, 108)

[PRS04] Rom Pinchasi, Radoš Radoičić, and Micha Sharir. On empty convex polygons in a planar point set. In *Proceedings of 20th ACM Annual Symposium on Computational Geometry (SCG 2004)*, pages 391–400, New York, 2004. Association for Computing Machinery. (Cited on page 85)

[PS85] Franco P. Preparata and Michael Ian Shamos. *Computational Geometry: An Introduction*. Texts and Monographs in Computer Science. Springer-Verlag, New York, 1985. 2nd printing, 1988. (Cited on page 109)

[PS02] Fábio Protti and Jayme L. Szwarcfiter. Clique-inverse graphs of bipartite graphs. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 40:193–203, 2002. (Cited on pages 17, 44, 45)

[Rad57] Richard Rado. Note on independence functions. *Proceedings of the London Mathematical Society Third Series*, 7:300–320, 1957. (Cited on page 16)

[Ram01] Edger A. Ramos. An optimal deterministic algorithm for computing the diameter of a three-dimensional point set. *Discrete & Computational Geometry*, 26(2):233–244, 2001. (Cited on page 4)

[Raz00] Ran Raz. VC-dimension of sets of permutations. *Combinatorica*, 20(2):241–255, 2000. (Cited on page 5)

[Rec89] András Recski. *Matroid Theory and its Applications in Electric Network Theory and in Statics*, volume 6 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1989. (Cited on page 6)

[RS83] Neil Robertson and Paul D. Seymour. Graph minor I. Excluding a forest. *Journal of Combinatorial Theory Series B*, 35(1):39–61, 1983. (Cited on page 2)

[RS84] Neil Robertson and Paul D. Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory Series B*, 36(1):49–64, 1984. (Cited on page 2)

[RS86a] Neil Robertson and Paul D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 2(3):309–322, 1986. (Cited on page 2)

[RS86b] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory Series B*, 41(1):92–114, 1986. (Cited on page 2)

[RS86c] Neil Robertson and Paul D. Seymour. Graph minors. VI. Disjoint paths across a disc. *Journal of Combinatorial Theory Series B*, 41(1):115–138, 1986. (Cited on page 2)

[RS88] Neil Robertson and Paul D. Seymour. Graph minors. VII. Disjoint paths on a surface. *Journal of Combinatorial Theory Series B*, 45(2):212–254, 1988. (Cited on page 2)

[RS90a] Neil Robertson and Paul D. Seymour. Graph minors. IV. Tree-width and well-quasi-ordering. *Journal of Combinatorial Theory Series B*, 48(2):227–254, 1990. (Cited on page 2)

[RS90b] Neil Robertson and Paul D. Seymour. Graph minors. IX. Disjoint crossed paths. *Journal of Combinatorial Theory Series B*, 49(1):40–77, 1990. (Cited on page 2)

[RS90c] Neil Robertson and Paul D. Seymour. Graph minors. VIII. A Kuratowski theorem for general surfaces. *Journal of Combinatorial Theory Series B*, 48(2):255–288, 1990. (Cited on page 2)

[RS91] Neil Robertson and Paul D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory Series B*, 52(2):153–190, 1991. (Cited on page 2)

[RS94] Neil Robertson and Paul D. Seymour. Graph minors. XI. Circuits on a surface. *Journal of Combinatorial Theory Series B*, 60(1):72–106, 1994. (Cited on page 2)

[RS95a] Neil Robertson and Paul D. Seymour. Graph minor. XIV. Extending an embedding. *Journal of Combinatorial Theory Series B*, 65(1):23–50, 1995. (Cited on page 2)

[RS95b] Neil Robertson and Paul D. Seymour. Graph minors. XII. Distance on a surface. *Journal of Combinatorial Theory Series B*, 64(2):240–272, 1995. (Cited on page 2)

[RS95c] Neil Robertson and Paul D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory Series B*, 63(1):65–110, 1995. (Cited on page 2)

[RS96] Neil Robertson and Paul D. Seymour. Graph minors. XV. Giant steps. *Journal of Combinatorial Theory Series B*, 68(1):112–148, 1996. (Cited on page 2)

[RS98] Satish B. Rao and Warren D. Smith. Approximating geometric graphs via "spanners" and "banyans". In *Proceedings of 30th Annual ACM Symposium on Theory of Computing (STOC '98)*, pages 540–550, New York, 1998. Association for Computing Machinery. (Cited on page 109)

[RS99] Neil Robertson and Paul D. Seymour. Graph minors. XVII. Taming a vortex. *Journal of Combinatorial Theory Series B*, 77(1):162–210, 1999. (Cited on page 2)

[RS03a] Neil Robertson and Paul D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory Series B*, 89(1):43–76, 2003. (Cited on page 2)

[RS03b] Neil Robertson and Paul D. Seymour. Graph minors. XVIII. Tree-decompositions and well-quasi-ordering. *Journal of Combinatorial Theory Series B*, 89(1):77–108, 2003. (Cited on page 2)

[RS04a] Neil Robertson and Paul D. Seymour. Graph minors. XIX. Well-quasi-ordering on a surface. *Journal of Combinatorial Theory Series B*, 90(2):325–385, 2004. (Cited on page 2)

[RS04b] Neil Robertson and Paul D. Seymour. Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory Series B*, 92(2):325–357, 2004. (Cited on page 2)

[Sch03] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. (Cited on pages 6, 59)

[Sed77] Robert Sedgewick. Permutation generation methods. *ACM Computing Surveys*, 9(2):137–164, 1977. (Cited on page 112)

[SH76] Michael Ian Shamos and Dan Hoey. Geometric intersection problems. In *Proceedings of 17th Annual IEEE Symposium on*

Foundations of Computer Science (FOCS '76), pages 208–215, Long Beach, 1976. IEEE Computer Society. (Cited on page 134)

[SS01] Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels. Support Vector Machines, Regularizations, Optimization, and Beyond*. The MIT Press, Cambridge, 2001. (Cited on page 3)

[Sta00] Richard P. Stanley. Positivity problems and conjectures in algebraic combinatorics. In Vladimir I. Arnold, Michael F. Atiyah, Peter D. Lax, and Barry C. Mazur, editors, *Mathematics: Frontiers and Perspectives*, pages 295–319. American Mathematical Society, Providence, 2000. Updates available at `http://www-math.mit.edu/~rstan/papers/problems_update.html`. (Cited on page 16)

[Sta04] Richard P. Stanley. An introduction to hyperplane arrangements. Lecture Notes at IAS/Park City Mathematics Institute Graduate Summer School 2004 "Geometric Combinatorics", 2004. Available at `http://www-math.mit.edu/~rstan/arrangements/`. (Cited on page 6)

[Ste99] Manfred Stern. *Semimodular Lattices: Theory and Applications*, volume 73 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1999. (Cited on pages 8, 60)

[Swa03] Ed Swartz. Topological representations of matroids. *Journal of the American Mathematical Society*, 16(2):427–442, 2003. (Cited on page 60)

[Sze04] Stefan Szeider. Minimal unsatisfiable formulas with bounded clause-variable difference are fixed-parameter tractable. *Journal of Computer and System Sciences*, 69(4):656–674, 2004. (Cited on page 11)

[TIAS77] Shuji Tsukiyama, Mikio Ide, Hiromu Ariyoshi, and Isao Shirakawa. A new algorithm for generating all the maximal independent sets. *SIAM Journal on Computing*, 6(3):505–517, 1977. (Cited on page 44)

[Vap98] Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley Series on Adapative and Learning Systems for Signal Processing, Communications, and Control. John Wiley & Sons, New York, 1998. (Cited on page 3)

[VC71] Vladimir N. Vapnik and Alexey Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971. (Cited on page 3)

[Viz64] V.G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskretnyĭ Analiz*, 3:25–30, 1964. (Cited on page 36)

[Viz65] V.G. Vizing. Critical graphs with a given chromatic class. *Diskretnyĭ Analiz*, 5:9–17, 1965. (Cited on page 36)

[Wag37] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114:570–590, 1937. (Cited on page 2)

[WCY99] Cao An Wang, Francis Y. Chin, and Bo Ting Yang. Maximum weight triangulation and graph drawing. *Information Processing Letters*, 70(1):17–22, 1999. (Cited on page 138)

[Wel03] Emo Welzl. Boolean Satisfiability — Combinatorics and Algorithms. Lecture Notes, ETH Zurich, 2003+. (Cited on page 161)

[Whi35] Hassler Whitney. On the abstract properties of linear dependence. *American Journal of Mathematics*, 57(4):509–533, 1935. (Cited on page 6)

[Whi92] Walter Whiteley. Matroids and rigid structures. In Neil White, editor, *Matroid Applications*, volume 40 of *Encyclopedia of Mathematics and its Applications*, chapter 1, pages 1–53. Cambridge University Press, Cambridge, 1992. (Cited on page 6)

[Woe03] Gerhard J. Woeginger. Exact algorithms for NP-hard problems: A survey. In Michael Jünger, Gerhard Reinelt, and Giovanni Rinaldi, editors, *Combinatorial Optimization — Eureka, You Shrink!*, volume 2570 of *Lecture Notes in Computer Science*, pages 185–208, Berlin, 2003. Springer-Verlag. (Cited on page 115)

[WY01] Cao An Wang and Boting Yang. A lower bound for β-skeleton belonging to minimum weight triangulations. *Computational Geometry: Theory and Applications*, 19(1):35–46, 2001. (Cited on page 138)

[Zie98] Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, revised first edition, 1998. (Cited on page 59)

In total, 181 items are listed.

Index

Postscript

Each chapter started with a quotation.

The quotation from *The Phantom of the Opera* in Chapter 0 and the quotation from *License to Kill* in Chapter 5 have been found at the Internet Movie Database (IMDb) `http://www.imdb.com/`.

The quotation from Gil Kalai in Chapter 1 can be found in his paper [Kal00]. The paper not only provides nice short surveys about five topics, but also gives some thoughts on a role and a future direction of mathematics. This is communicated to me by Günter Ziegler at IAS/Park City Mathematics Institute Graduate Summer School 2004 "Geometric Combinatorics" where he gave a series of lectures on convex polytopes.

The quotation from Jirka Matoušek in Chapter 2 can be found in his paper coauthored with Tibor Szabó [MS04]. When I looked at a preprint version of their paper, I found a problem with $e^{n^{1/2}}$ and $e^{n^{1/3}}$. Then, he put this remark in the paper as a footnote. I am happy with his amazing response.

The quotation from Emo Welzl in Chapter 3 can be found in his lecture notes on satisfiability [Wel03] where he starts the discussion on NP-completeness. This sentence is one of his mottos and he repeated (and will repeat) again and again, but as for the citation the lecture notes mentioned above seems the first explicit appearance. I hope that the thesis follows his statement.

The quotation from János Pach in Chapter 4 was told by Jeff Erickson via his weblog [Eri04]. It seems that János Pach addressed the words at the banquet of Japanese Conference on Discrete and Computational Geometry 2004 held in Tokyo. However, I am not sure that what he said is true or not.

Curriculum Vitae

Yoshio Okamoto

born on July 22, 1976 in Hekinan, Japan

1992–1995 **high school in Okazaki, Japan**

1995–1999 **Studies at the University of Tokyo, Japan**
Systems Science
Degree: Bachelor of Systems Science

1999–2001 **Master study at the University of Tokyo, Japan**
Systems Science
Degree: Master of Systems Science

2001–2002 **Pre-doc study at Swiss Federal Institute of Technology, Switzerland**
Theoretical Computer Science

2002–2005 **PhD student at Swiss Federal Institute of Technology, Switzerland**
Theoretical Computer Science