Jesús A. De Loera
Jörg Rambau
Francisco Santos

ALGORITHMS AND COMPUTATION IN MATHEMATICS 25

Triangulations

Structures for Algorithms and Applications

[ACM]

Springer

Triangulations:

Structures for Algorithms and Applications

Jesús A. De Loera

Dept. of Mathematics
University of California
Davis, California, USA
deloera@math.ucdavis.edu

Jörg Rambau

Institut für Mathematik
Universität Bayreuth
Bayreuth, Germany
joerg.rambau@uni-bayreuth.de

Francisco Santos

Depto. de Matemáticas,
Estadística y Computación
Universidad de Cantabria
Santander, Spain
francisco.santos@unican.es

Preface

Triangulations appear in many different parts of mathematics and computer science since they are the natural way to decompose a region of space into smaller, easy-to-handle pieces. From volume computations and meshing to algebra and topology, there are many natural situations in which one has a fixed set of points that can be used as vertices for the triangulation. Typically one wants to find an optimal triangulation of those points or to explore the set of their all triangulations. The given points may represent the “sites” for a Delaunay triangulation computation, the test points for a surface reconstruction, or a set of monomials, represented as lattice points in $\mathbb{Z}^d$, in an algebraic-geometric meaning.

A central theme of this book is to use the rich geometric structure of the space of triangulations of a given set of points to solve computational problems (e.g., counting the number of triangulations or finding optimal triangulations with respect to various criteria), and for setting up connections to novel applications in algebra, computer science, combinatorics, and optimization. Thus at the heart of the book is a comprehensive treatment of the theory of regular subdivisions, secondary polytopes, flips, chambers, and their interactions. Again, we firmly believe that understanding the fundaments of geometry and combinatorics pays up for algorithms and applications.

The book is designed to serve as a textbook or for self-guided study. It was written with graduate students or advanced undergraduates as the target audience (in fact, several groups of students were kind enough to let us test the book with them). Beyond good knowledge of linear algebra, all that is required to use this book is maturity to read and understand proofs. With many examples and exercises, and with almost five hundred illustrations, we aim to gently introduce beginners to the properties of the spaces of triangulations of “highly-structured” (e.g., cubes, cyclic polytopes, lattice polytopes, etc.) and “pathological” situations (e.g., disconnected spaces of triangulations, NP-hardness constructions, etc.). We do this in arbitrary dimension, while using only elementary geometric principles. We are excited to present many open questions. Some are new, but many have been open for some time. Also, the book contains many new results appearing here for the first time, besides corrections and simpler proofs of well-known theorems.

Chapter 1 describes several instances where triangulations of a point set naturally arise in combinatorics, optimization, and algebra, as motivation for the rest. A reader may select which parts he or she is most interested in and skip the rest. None of the material is a prerequisite for later chapters, but we hope to communicate some of the exciting and diverse applications possible and to show that triangulations are worth studying by outsiders.

Chapters 2 and 4 lay out the formal language, notation, and basic constructions. Concerning the language, we have decided to distinguish between the points (or vectors) of a configuration and the labels used to denote them. This may look awkward to the beginner at first sight but it has many advantages in the long run.

Chapter 3, is an “interlude” devoted specifically to what happens in two dimensions and a quick glance at dimension three. This chapter is almost independent from the rest and we hope it will help the reader to build intuition and to motivate, in a visual way, the challenges to come in arbitrary dimension (e.g., the notion of flip, enumeration, optimization, etc). Because Chapter 3 lies in between two more technical chapters we included more examples and applications that helped balance the presentation. It is a fun detour through a very active area of computational geometry.

Chapter 5 contains what is probably the central theorem of the book: Gelfand, Kapranov and Zelevinsky’s ground-breaking construction of a polytope with face lattice equal to the poset of *regular subdivisions* for any given configuration. This theorem is the central tool for flipping algorithms such us the incremental randomized construction of Delaunay triangulations, customary in Computational Geometry.

The next two chapters are devoted to the study of important examples of configurations (Chapter 6) and triangulations (Chapter 7). In the first one, nice combinatorial structures allow for a deeper study of the properties of these examples, while in the latter the focus is on ingenious constructions of pathological triangulations (disconnected flip-graphs, and triangulations with very few or no flips).

Chapter 8 focuses on computation and algorithms. We start with a discussion of data structures and discuss methods for enumeration and optimization in triangulations of arbitrary dimension. Here we prove that the structural understanding helps with the design algorithms, software, and the analysis of computational complexity.

Finally, Chapter 9 explores generalizations or different ways of looking at some of the structures in the rest of the book. Some of these "further topics" are so rich they could be themselves central topics of books. Interesting directions discussed include fiber polytopes, Cayley's construction, Gröbner bases, connections to lattice points and Ehrhart functions, and the combinatorics of simplicial spheres.

If you are a teacher planning to give a one semester course based on this book, the core of it should be Chapters 2, 4 and 5, ending with Section 5.3. These chapters develop the structure on which most of the rest is based. Some parts can be omitted if you need to go to the essentials. These include Sections 2.6 and 4.5. The former relates triangulations with classical topics in polytope theory and the latter is meant as a comprehensive reference list of different ways in which triangulations can be characterized.

Despite our very best intentions there surely remain some errors or typos in the text. We take full responsability for that and plan to post a list of errors and typos at our web sites. Please do let us know via email (our web addresses are listed below) if you find any. (And feel free to write to us with your triangulation stories too!)

This book took too long to write and required the help of many friends and colleagues who taught us and inspired us with their mathematics and wisdom. We are truly grateful to the following people for their ideas, corrections, comments, questions, suggestions, or simply because they patiently kept asking about our seemingly never-ending book project:

Oswin Aichholzer, Nina Amenta, Christos Athanasiadis, David Avis, Miguel Azaola, Eric Babson, Imre Bárány, David Barnette, Alexander Barvinok, Margaret Bayer, Matthias Beck, Alexander Below, Andrew Berget, Marshall Bern, Louis J. Billera, Anders Björner, Jürgen Bokowski, Prosejit Bose, Lewis Bowen, Javier Bracho, Igor Chaplygin, David Bremner, Winfried Bruns, Robert Connelly, Mike Develin, Herbert Edelsbrunner, Richard Ehrenborg, David Eppstein, Komei Fukuda, Isidoro Gitler, David Gove, Peter Gritzmann, Joseph Gubeladze, Christian Haase, David Haws, Raymond Hemmecke, Martin Henk, Sven Herrmann, Takayuki Hibi, Serkan Hoşten, Katherine Jones, Michael Joswig, Gil Kalai, Jean-Michel Kantor, Victor Klee, Steve Klee, Edward D. Kim, Matthias Köppe, Francisco Larrión, James Lawrence, Carl W. Lee, Fu Liu, Monika Ludwig, Frank Lutz, Diane Maclagan, Peter Malkin, Jirka Matoušek, Tyrrell McAllister, Luis Montejano, Walter Morris, Oleg Musin, Victor Neumann-Lara, Eran Nevo, Isabella Novik, Mohamed Omar, David Orden, Shmuel Onn, Joseph O'Rourke, Edwin O'Shea, Lior Pachter, Igor Pak, Julian Pfeifle, Vincent Pilaud, Richard Pollack, Alexander Postnikov, Pedro Ramos, Tomás Recio, Victor Reiner, Jürgen Richter-Gebert, Leni Rostock, Günter Rote, Raman Sanyal, Achill Schürmann, Jozsef Solymosi, Richard P. Stanley, Mike Stillman, Ileana Streinu, Bernd Sturmfels, Francis Su, Thorsten Theobald, Rekha Thomas, Jorge Urrutia, Frank Vallentin, Michèle Vergne, Rafael Villarreal, Robert Weismantel, Roger J-B Wets, Emo Welzl, Jörg Wills, Andrei Zelevinksy, Günter M. Ziegler, and Chuanming Zong. Last, but not least, we thank the anonymous referees who gave us excellent comments and suggestions for the book.

Several institutions supported us during the writing journey and hosted the other visiting co-authors several times. Of course, our home institutions, University of California, Davis, Universität Bayreuth, and Universidad de Cantabria, were key hosts of our book activities, but earlier versions of the book were also discussed, taught, or written during our various stays at the Geometry Center of the University of Minnesota, ZIB-Berlin, ETH-Zürich, MSRI-Berkeley, and Universität Magdeburg. The Universidad de Cantabria and the Mathematical Sciences Research Institute at UC Berkeley gave us the unique opportunity to deliver Summer courses in 2001 and 2003 on the topic of the book, which very much shaped its current contents and format. The Mathematisches Forschunginstitut Oberwolfach allowed us to meet in their "Research in Pairs" program for two weeks in the Summer of 2005. This book benefited from the financial support received through grants from the National Science Foundation (DMS-0914107, DMS-0608785, DMS-0309694, and DMS-0073815), the Alexander von Humboldt Foundation, and the Spanish Ministry of Science and Innovation. We are grateful for the support we received.

Above all we are truly grateful to our wives Ingrid, Nicole, and Mónica, and our three pairs of children Antonio and Andrés, Magdalena and Anna, Jon Ander and Íñigo, for their love, encouragement and support. Whenever we were meeting for the book, either in Davis, Bayreuth, or Santander, we enjoyed the warm hospitality of the respective host family, their friendship, and their appreciation for our work on the book. Most of all we thank them for being with us in the first place. In a way this book belongs to them!

Jesús A. De Loera, Jörg Rambau, and Francisco Santos April 2010
deloera@math.ucdavis.edu, Joerg.Rambau@uni-bayreuth.de, francisco.santos@unican.es

Contents

Triangulations in Mathematics

1

When solving a difficult problem it is a natural idea to decompose complicated objects into smaller, easy-to-handle pieces. In this book we study such decompositions: *triangulations of point configurations*. We will look at triangulations from many different points of view. We will explore their combinatorial and geometric properties, as well as some algorithmic issues arising along the way.

This first chapter is designed to informally introduce the fundamental notions to come in later chapters. We provide motivating examples to convince the reader that triangulations are useful and that they appear in many areas of mathematics. The reader can skip most of this chapter safely, or read the sections in an order different than the one presented. The examples also provide an entry door for non-discrete-geometers (e.g., algebraic geometers, computer scientists, linear programming enthusiasts, etc.) that wish to learn about triangulations for their research, connecting our book to their topic. Without further delay we begin.

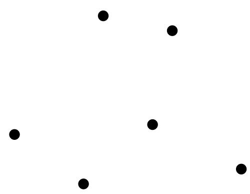

Figure 1.1: A point configuration.

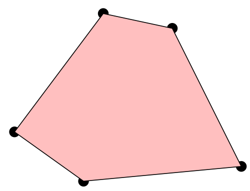

Figure 1.2: Its convex hull.

A *point configuration*[1] is a finite collection of points $\mathbf{A} = \{\mathbf{a}_1, \ldots, \mathbf{a}_n\}$ in Euclidean space $\mathbb{R}^d$.

The *convex hull* of $\mathbf{A}$ is by definition the intersection of all convex sets containing the points in $\mathbf{A}$. We denote it by $\operatorname{conv}(\mathbf{A})$.

A *k-simplex* is the convex hull of $k+1$ affinely independent points in $\mathbb{R}^d$ (clearly $d \geq k$). Simplices are the simplest of polyhedra: points, segments, triangles, tetrahedra, etc. A j-face of a k-simplex is the convex hull of $j+1$ of its vertices and thus in particular is a j-simplex itself. We say that the empty set is a (-1)-dimensional face common to all simplices, so that every k-simplex has exactly $\binom{k+1}{j+1}$ j-faces for $j = -1, 0, 1, \ldots, k$. A simplex of $\mathbf{A}$ is a simplex whose vertices are taken from $\mathbf{A}$.

Here is the main actor in this play:

Definition 1.0.1. A *triangulation* of a point configuration $\mathbf{A} \in \mathbb{R}^d$ is a collection $\mathscr{T}$ of d-simplices all of whose vertices are points in $\mathbf{A}$ that satisfies the following two properties:

1. The union of all these simplices equals $\operatorname{conv}(\mathbf{A})$. (*Union Property*)
2. Any pair of these simplices intersects in a common face (possibly empty). (*Intersection Property*)

[1]The word *configuration* is used to distinguish this from a *set* of points: in a subset of $\mathbb{R}^d$ there can be no multiple points, whereas in a configuration we, in principle, are allowed to have $\mathbf{a}_i = \mathbf{a}_j$ for some $i \neq j$ and still consider $\mathbf{a}_i$ and $\mathbf{a}_j$ two different elements of $\mathbf{A}$. See Chapter 2 for explanations and justifications of this.

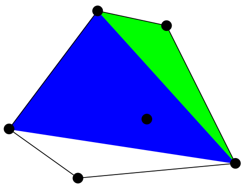

Figure 1.3: The union of simplices is not the whole convex hull (union property fails).

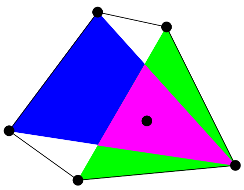

Figure 1.4: The intersection of simplices is not proper (intersection property fails).

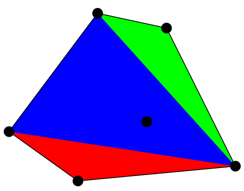

Figure 1.5: This example satisfies both the union and intersection properties, thus it is a triangulation.

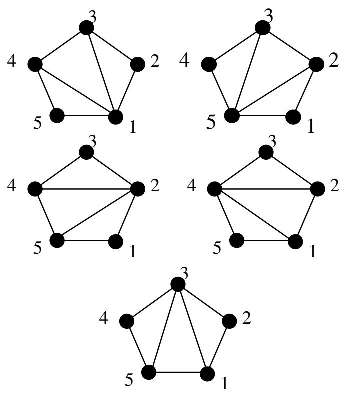

Figure 1.6: The five triangulations of a pentagon. In this case all triangulations are full.

As a particular case, by a triangulation of a convex polytope **P** we mean a triangulation of the point configuration given by the vertices of **P**.

In Figures 1.3 and 1.4 we show examples of the possible pathologies that may prevent a triangulation.

Let us emphasize two features that distinguish our definition from other definitions of the word triangulation that the reader may have seen before:

1. With very few exceptions, we will fix in advance the set **A** of points that can be used as vertices, and it is a finite set. In particular, the number of triangulations of a point configuration is always finite. This does not happen in classic combinatorial topology or in some applications, where one is free to use arbitrary additional points.

2. We do not insist that all points of **A** are used as vertices in a triangulation. For example, if our point configuration consists of points in $\mathbb{R}$, then there is one triangulation with only one simplex (the whole segment $\operatorname{conv}(\mathbf{A})$) and two vertices (the two convex hull extremes of the line segment $\operatorname{conv}(\mathbf{A})$), regardless of how many points we may have in **A**. This differs from the standard use of triangulations in Computational Geometry, where one usually requires all the points to be used. Thus we make the following definition:

Definition 1.0.2. Let $\mathbf{A} \subset \mathbb{R}^d$ be a point configuration. We call a triangulation of **A** *full* if all the points of **A** are vertices of it.

The first of these two features gives our setting a strong combinatorial flavor. Actually, to describe a particular triangulation we will normally number the points of A from 1 to n and give the list of vertex sets of the d-simplices in the triangulation. For example, the five triangulations of a pentagon would be written as:

$$\{\,\{1,2,3\},\{1,3,4\},\{1,4,5\}\,\}, \qquad \{\,\{1,2,5\},\{2,3,5\},\{3,4,5\}\,\},$$

$$\{\,\{1,2,5\},\{2,4,5\},\{2,3,4\}\,\}, \qquad \{\,\{1,4,5\},\{1,2,4\},\{2,3,4\}\,\},$$

$$\text{and} \qquad \{\,\{1,2,3\},\{1,3,5\},\{3,4,5\}\,\}.$$

We will even abbreviate $\{1,2,3\}$ as 123 and so on, whenever this creates no confusion.

Why should anyone care about studying triangulations of point configurations? It is our intention to illustrate, with some examples, how several of the fundamental defining properties of triangulations draw themselves into topics that, at first sight, seem far apart from the geometry of triangulations.

1.1 Combinatorics and triangulations

It is well-known that polyhedra can be quite useful when dealing with combinatorial problems. In this section we show two examples of combinatorial identities that have interpretations in terms of triangulations. Let us start with what is possibly the simplest example of the structures studied in this book: the set of triangulations of a convex polygon. Let $\mathbf{C}_n$ be a convex polygon with n vertices, numbered from 1 to n in clockwise order. The first observation is that the number of triangulations does not depend on the coordinates of the vertices. Indeed, a triangulation will be given by any $n-3$ diagonals not crossing one another, and two diagonals cross if and only if they involve four vertices in an alternating way. That is to say, if $1 \leq i < j < k < l \leq n$, then the only two diagonals involving these four points and crossing each other are ik and jl.

In particular, the number of triangulations of a convex n-gon is a number depending only on n and that we will denote by t_n. The first instances are easy to compute: $t_3 = 1$, $t_4 = 2$, $t_5 = 5$ (see Figure 1.6), and $t_6 = 14$ (see Figure 1.15).

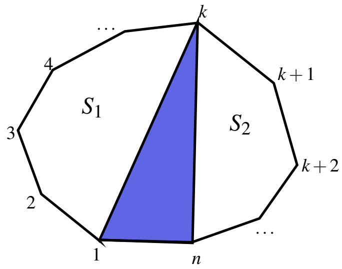

Figure 1.7: Setting up a recursion for $R(n)$.

Proposition 1.1.1. *Setting $t_2 = 1$ by convention, the sequence of numbers $t_2, t_3, t_4, \ldots$ satisfies the following recurrence relation:*

$$t_n = t_2 t_{n-1} + t_3 t_{n-2} + \cdots + t_{n-1} t_2.$$

Proof. In every triangulation of $\mathbf{C}_n$ the edge $\{1,n\}$ is a side of exactly one of the triangles, say $\{1,k,n\}$. The total number of triangulations, then, is the sum of the triangulations using the triangle $\{1,k,n\}$ for k ranging from 2 to $n-1$.

For a fixed k, the complement of the triangle $\{1,k,n\}$ consists of two polygons S_1 and S_2 with k and $n-k+1$ vertices, respectively. Since the polygons S_1 and S_2 can be triangulated independently, the number of triangulations of $\mathbf{C}_n$ using the triangle $\{1,k,n\}$ equals $t_k t_{n-k+1}$. (Of course, we admit S_1 or S_2 being a single edge, or a "2-gon", if $k=2$ or $k=n-1$. We take $t_2 = 1$ because this makes $t_2 t_{n-1}$ be equal to t_{n-1}).

The recurrence formula is obtained by adding this for $k = 2, \ldots, n-1$. □

The recurrence formula in the statement allows us to easily compute further terms in the sequence, for example $t_7 = 1 \cdot 14 + 1 \cdot 5 + 2 \cdot 2 + 5 \cdot 1 + 14 \cdot 1 = 42$ and $t_8 = 1 \cdot 42 + 1 \cdot 14 + 2 \cdot 5 + 5 \cdot 2 + 14 \cdot 1 + 42 \cdot 1 = 132$. A closed formula for t_i can be deduced from the recurrence with the method of *generating functions* (see Exercise 1.5), but here we use a more direct approach to find it.

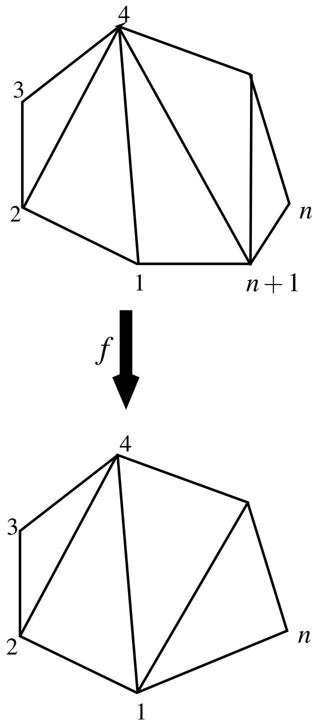

Figure 1.8: The contracting map.

Theorem 1.1.2. *The number t_n of triangulations of a convex n-gon equals*

$$\frac{1}{n-1}\binom{2n-4}{n-2}.$$

Proof. As before, we assume the vertices of the n-gon to be labeled from 1 to n in clockwise order. Denote by $\Delta(\mathbf{C}_n)$ the set of all triangulations of an n-gon, and their number by t_n. We are going to set up a simple surjective map f from $\Delta(\mathbf{C}_{n+1})$ onto $\Delta(\mathbf{C}_n)$. A triangulation in $\Delta(\mathbf{C}_{n+1})$ is mapped to the triangulation in $\Delta(\mathbf{C}_n)$ obtained by contracting the boundary edge $\{1,n+1\}$ (see Figure 1.8).

Our crucial observation is that the number of triangulations in $\Delta(\mathbf{C}_{n+1})$ mapped to a certain triangulation $\mathscr{T}$ in $\Delta(\mathbf{C}_n)$ equals the number of edges incident to vertex 1 in $\mathscr{T}$ (the *degree* of vertex 1 in $\mathscr{T}$). This is true because, to reverse the map f, we must choose one edge incident to 1 and "double" it to obtain a triangle incident to the edge $\{1, n+1\}$ (For example, in Figure 1.8 one has to double the edge $\{1,4\}$). This implies that

$$t_{n+1} = \sum_{\mathscr{T} \in \Delta(\mathbf{C}_n)} \deg_{\mathscr{T}}(1).$$

By cyclic symmetry of the n-gon, this same formula must hold for any other vertex of it. Hence:

$$nt_{n+1} = \sum_{i=1}^{n} \sum_{\mathscr{T} \in \Delta(\mathbf{C}_n)} \deg_{\mathscr{T}}(i) = \sum_{\mathscr{T} \in \Delta(\mathbf{C}_n)} \sum_{i=1}^{n} \deg_{\mathscr{T}}(i).$$

But it turns out that the sum $\sum_{i=1}^{n} \deg_{\mathscr{T}}(i)$ is independent of $\mathscr{T}$; it equals twice the number of edges of $\mathscr{T}$, that is, $2(2n-3)$. Hence:

$$t_{n+1} = \frac{2(2n-3)}{n} t_n, \quad \text{or} \quad t_n = \frac{2(2n-5)}{n-1} t_{n-1}.$$

From this we conclude that:

$$\begin{aligned} t_n &= \frac{2^{n-2}(2n-5)(2n-7)\cdots 3 \cdot 1}{(n-1)(n-2)\cdots 3 \cdot 2} && (1.1) \\ &= \frac{(2n-4)!}{(n-1)!(n-2)!} = \frac{1}{n-1}\binom{2n-4}{n-2}. && (1.2) \end{aligned}$$

□

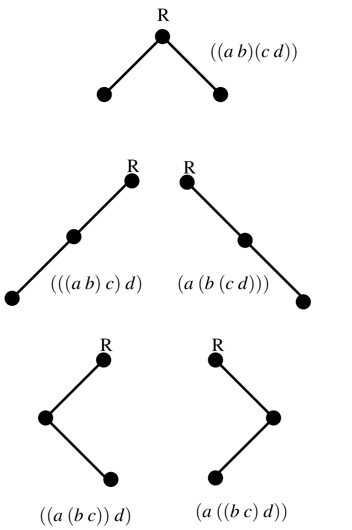

Figure 1.9: The five binary trees in 3 nodes, with their associated parenthesizations.

The sequence of numbers we have just found is known as the *Catalan numbers* (see Definition 1.1.4 below) and is one of the most important number sequences in combinatorics, perhaps comparable to the well-known Fibonacci sequence. We remark that asymptotically these numbers grow (up to a constant factor) like $4^n n^{-3/2}$. This may be seen by using Stirling's approximation of the factorial. As an example of the ubiquity of the Catalan numbers, the following statement lists four other combinatorial structures whose cardinality is given by the Catalan sequence, and Exercise 6.19 in [302, p. 219] contains 61 additional such examples.

Theorem 1.1.3. *There are as many triangulations of a convex polygon with $n+2$ vertices as:*

(i) Binary trees with n nodes (and hence $n-1$ edges).

(ii) Parenthesizations of the product of $n+1$ factors, that is to say, ways of placing n pairs of parentheses in order to perform the product.

(iii) Sequences of length $2n$ consisting of n plus signs and n minus signs, with the property that in every initial segment of the sequence there are at least as many pluses as minuses.

(iv) *Monotone paths in the integer grid, going from* $(0,0)$ *to* (n,n) *by steps of length 1 in the positive directions of the axes, and never going above the diagonal.*

One interesting feature of the equivalence to sign sequences is that it immediately shows that the number of triangulations of an n-gon is bounded above by 2^{2n-4}. Of course, that is also clear from Theorem 1.1.2, but its proof needed some work. Moreover, the equivalence explicitly tells us how to write a given triangulation as a binary number of length $2n-4$.

Before proving Theorem 1.1.3, let us define binary trees, which the reader may not be familiar with. A *tree* is a connected simple graph with no cycles [61]. We are interested in rooted trees, i.e., trees with a special distinguished node, called the *root*. In rooted trees, one can direct the edges naturally along the unique paths from each node to the root node. This establishes a hierarchy among the nodes: node v_1 becomes a *child of node* v_2 if they are adjacent and the edge joining them is directed from v_2 to v_1 (v_2 is the *parent of* v_1). Rooted trees are normally drawn with the root on top and with parents above their children.

A *binary tree* is a rooted tree in which each edge is marked as a *right* or *left* edge of its parent and each node has at most one right child and at most one left child (in particular, each node has either $0,1$ or 2 children). In Figure 1.9 we show the five different binary trees on three nodes. As it is customary, a left child is drawn toward the left-down direction and a right child is drawn toward the right-down direction. Binary trees are very useful combinatorial structures due to applications in data structures and the design of algorithms (see for example [191]).

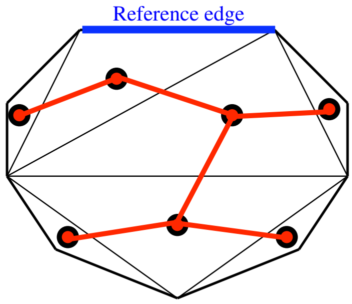

Figure 1.10: A binary tree dual of a triangulation.

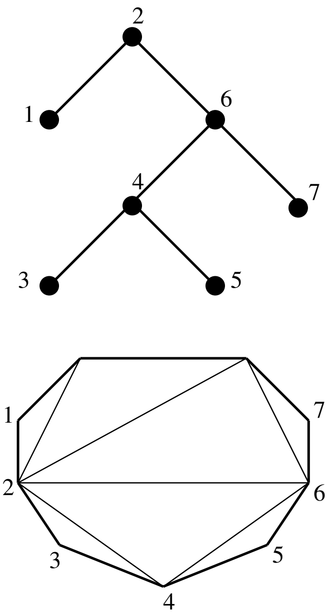

$+ - + + - + + - - + + - - -$
1 1 2 3 3 4 5 5 4 6 7 7 6 2

Figure 1.11: A binary tree with its symmetric order traversal, its associated triangulation, and its associated sequence of signs.

Proof of Theorem 1.1.3.

1. From triangulations to binary trees: Let us see how to build up a binary tree from a given triangulation of the $(n+2)$-gon. As usual, we assume vertices of the polygon labeled from 1 to $n+2$. We call the edge $\{1,n+2\}$ the *reference edge* of the polygon. We draw a node of the binary tree inside each of the n triangles of the triangulation, and join nodes of adjacent triangles by an edge. We declare the root of the tree to be the node of the unique triangle that contains the reference edge. The three sides of each triangle can be clearly identified as a "parent edge" (the one towards the root), a "right edge" (the next one in the clockwise direction) and a "left edge" (the third one). In particular, every node in the tree has a parent (unless it is the root node) and its children are labeled as right or left depending on whether the corresponding edge in the triangle is the right or the left one. See Figure 1.10.

To show that this construction is indeed a bijection, it suffices to show that it can be reversed: Starting with a binary tree, draw a triangle for the root and call its edges "parent", "right" and "left" appropriately. Then glue triangles to its right and left edge for the right and left children of the root, if they exist. Recursively continue with grand children and all the other descendants (great-grand-children, etc.) and, after you have finished, number the vertices of the $(n+2)$-gon starting and ending with the end-points of the parent edge of the root triangle.

There is actually a nice correspondence between the n nodes in the binary tree and the n vertices of the $(n+2)$-gon out of the reference edge, exhibited in Figure 1.11. It is given by the *symmetric order traversal* of the nodes of a binary tree: Start with the root node. If a node has at least one child, then process recursively first its left subtree, then insert the node itself into your order, then process recursively its right subtree, and finally return to its parent. We give numbers 1 to n to the n nodes as we visit them during the traversal. If a node has no children, then simply visit it by assigning to it the next available number from 1 to n, and return to its parent. We will use the symmetric order traversal a bit later.

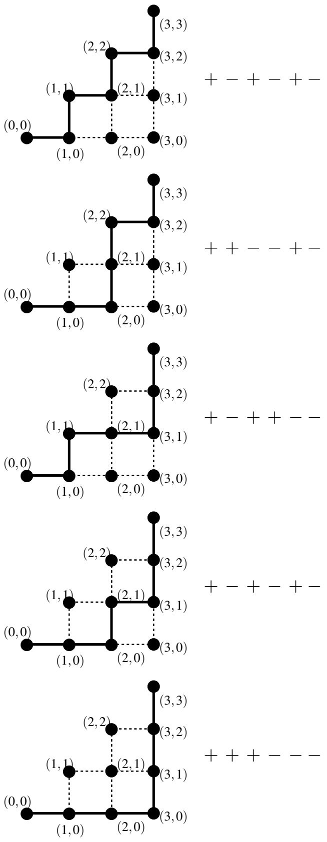

Figure 1.12: The five monotone paths/sign sequences of length six.

2. From binary trees to parenthesizations: The bijection between binary trees with n nodes and parenthesizations of products of $n+1$ factors was displayed in Figure 1.9. We place a pair of parentheses for each node of the binary tree, starting with the root parentheses which enclose the whole product and inserting inner parentheses for children with the following rule: if the right/left child of a given parent node has k descendants, the corresponding parentheses will enclose the $k+1$ leftmost/rightmost factors within the ones enclosed by the parent parentheses. Alternatively, one can start drawing parentheses for the leafs of the tree (leaving place holders for the two variables they contain, which we cannot still identify) and add greater and greater parentheses for their parents, inserting right or left factors (placeholders) depending on the type of edge leading to the parent. We leave it to the reader to convince him or herself that this is indeed a bijection.

3. From binary trees to sign sequences: Clearly, there is going to be a plus and a minus sign corresponding to each node in the tree, and the plus sign will appear before the minus to guarantee that every initial segment has at least as many pluses as minuses. The way to construct the sequence is: Go along the tree in the symmetric order traversal presented above. When visiting a node, first process its left subtree, then place the plus sign for this node, then visit the right subtree, then place the minus sign. Figure 1.11 shows an example where, to make things clear, each plus or minus is labeled by its corresponding vertex of the tree.

In the exercises you will see how to construct the sequences of signs directly from the triangulation.

4. From sign sequences to monotone paths: Figure 1.12 shows the monotone paths under consideration and, at the same time, their bijection to sign sequences. Essentially, plus signs correspond to steps to the right and minus signs to steps upwards. The condition that the monotone paths do not cross the diagonal is exactly equivalent to saying that every initial segment has at least as many plus signs as minus signs.

□

Definition 1.1.4. The n-th Catalan number, where $n = 0, 1, 2, \ldots,$ is the number C_n defined by the following recurrence formula:

$$C_0 = 1, \qquad C_n = \sum_{k=0}^{n-1} C_k C_{n-k-1}, \quad \forall n > 0. \tag{1.3}$$

Equivalently, it is the number of triangulations of the convex $(n+2)$-gon, which equals

$$C_n = \frac{1}{n+1}\binom{2n}{n}. \tag{1.4}$$

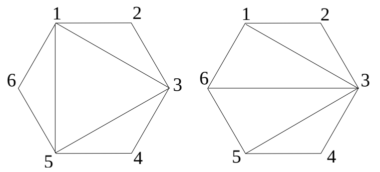

Figure 1.13: A diagonal flip in the quadrilateral 1356.

Theorem 1.1.3 can be read as saying that the five combinatorial structures described there are just different formulations of one and the same structure, that we can call the “Catalan structure”. Having the different formulations, besides its mathematical appeal, has practical consequences: properties which are obvious in one formulation may be invisible in others, and the many appearances of the structure provide additional insight and more tools to attack Catalan-type problems with.

As an example, our proof of Theorem 1.1.2 is heavily based on the cyclic symmetry of the convex n-gon, while none of the four other structures of Theorem 1.1.3 have a cyclic symmetry at all.

Once we know the cardinality of the set $\Delta(\mathbf{C}_n)$ of triangulations of the convex n-gon, let us see that it is “more than a set”, that there are natural relations between pairs of triangulations. Every internal edge in a triangulation is a diagonal of a convex quadrilateral formed by two adjacent triangles. One can change this diagonal to the opposite one and get a triangulation which is as similar as possible to the initial one. This operation is called a *diagonal flip* or simply a *flip* for short. Figure 1.13 shows a flip between two triangulations of a hexagon.

We can thus consider the set of triangulations of the n-gon as the vertices of a graph whose edges are diagonal flips. This graph is called the *graph of flips* of the triangulations of the n-gon. Some straightforward properties of the graph are:

1. It is regular of degree $n-3$ (that is to say, every triangulation has exactly $n-3$ flips). This is so because there is one flip associated to each internal diagonal.

2. It is connected. To prove this, let us pick any particular vertex, say the i-th one, and consider the unique triangulation in which all the triangles are incident to i. We call this the i-th *standard triangulation* of the n-gon. An example is in the right part of Figure 1.13. In any triangulation other than this one there is always at least one flip which increases the degree of vertex i: just flip the diagonal jk for any triangle ijk with j and k not consecutive vertices of the n-gon. This shows that every triangulation can be transformed into the standard one by a sequence of at most $n-3$ flips.

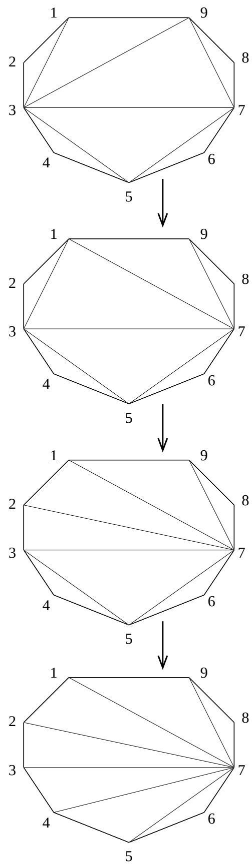
Figure 1.14: Flipping towards the 7th standard triangulation.

Other not-so-easy properties of the graph of flips are that it is Hamiltonian [217] and that it is the graph of a convex and simple polytope of dimension $n-3$ called the *associahedron* [202] (see also [334, Chapter 0]). The associahedron is a particular case of the *secondary polytope* or “polytope of triangulations and flips”, which exists for any finite point configuration in any dimension. Flips and secondary polytopes are introduced in Chapters 2 and 5, and are among the central topics in this book. See Figure 1.15 for a picture of the graph of flips on triangulations of a hexagon (well, we have forgotten to draw one edge. Can you find it?). You should try to verify in the picture all the properties of the graph mentioned so far.

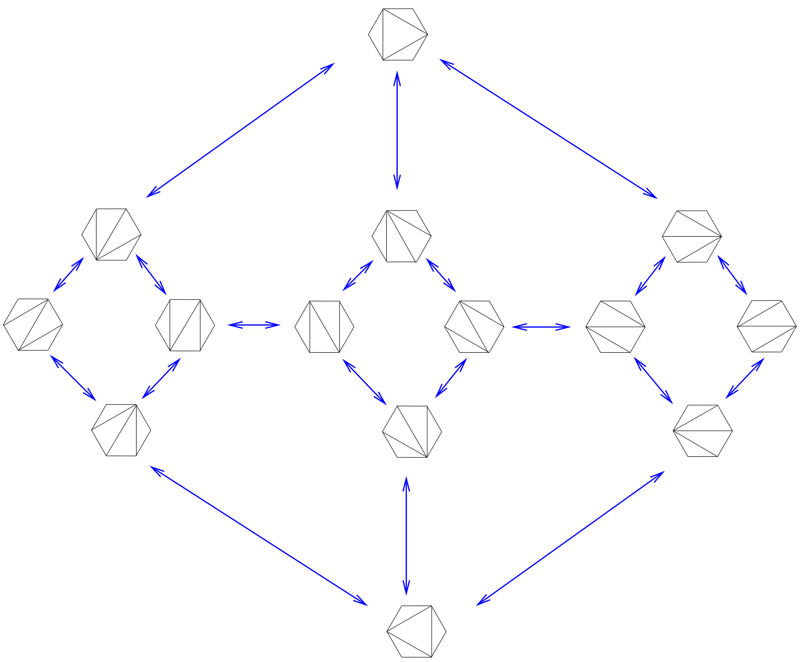

Figure 1.15: The Graph of flips for a hexagon, with one edge missing.

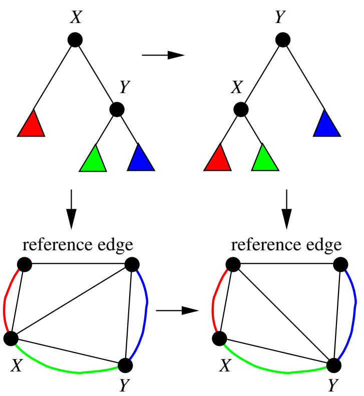

Figure 1.16: The diagonal flip corresponding to a rotation.

You may be wondering whether the graph of flips is meaningful in the other formulations of the Catalan structure that we have mentioned. The answer is yes and no. For example, the monotone path formulation possesses its own natural notion of flip (move the path along a single square of the grid), but these flips are certainly not equivalent to the flips in triangulations: In Figure 1.12 you can see paths with one, two, and three flips.

In the context of binary trees, however, diagonal flips can be described easily: They arise as the so-called "rotation of an edge". If an edge connects a node X to its right child Y, let P, A, B, and C denote the parent subtree of X, left subtree of X, left subtree of Y, and right subtree of Y, respectively. The rotation changes this to the binary tree in which X is a left child of Y and P, A, B, and C are, respectively, the parent subtree of Y, left subtree of X, right subtree of X, and right subtree of Y. A rotation and its correspondence to a flip in triangulations are depicted in Figure 1.16. In the context of parenthesizations a flip is given by a single application of the associative law $A(BC) \mapsto (AB)C$.

But how many flips does it take to move from one triangulation to another? Remember that the distance between two nodes in a connected graph is the minimal number of edges needed to go from one node to the other, and that the diameter of the graph is the maximum distance between nodes. It is interesting to say something about the diameter of the graph of flips:

Proposition 1.1.5. *Let $D(\mathbf{C}_n)$ be the diameter of the graph of flips between triangulations of the convex n-gon. Then:*

(i) $D(\mathbf{C}_n) \leq 2n - 10 + 12/n$ *for every n (in particular, it is bounded by* $2n - 10$ *for every* $n \geq 12$*).*

(ii) $D(\mathbf{C}_n) + 1 \leq D(\mathbf{C}_{n+1}) \leq D(\mathbf{C}_n) + 3$ *for every n.*

Proof. Part (i) can be proved by slightly refining the argument that proved connectedness. Let $\mathscr{T}$ and $\mathscr{T}'$ be two triangulations, and let d_j and d'_j denote the degrees of the vertex j in $\mathscr{T}$ and $\mathscr{T}'$ respectively, for each $j = 1, \ldots, n$. What we have shown is that for every $i = 1, \ldots, n$ there is a path from $\mathscr{T}$ to $\mathscr{T}'$ consisting of $2n - 2 - d_i - d'_i$ flips: just start flipping from $\mathscr{T}$ and $\mathscr{T}'$ in a way that always increases the degree of the i-th vertex. Now we wonder what is the minimum length of these n paths we constructed. This is a difficult question so we look instead at the *average* length, which is:

$$\begin{aligned} \frac{1}{n}\sum_{i=1}^{n}(2n-2-d_i-d'_i) &= 2n-2-\frac{1}{n}\left(\sum_{i=1}^{n} d_i + \sum_{i=1}^{n} d'_i\right) && (1.5)\\ &= 2n-2-\frac{8n-12}{n} && (1.6)\\ &= 2n-10+\frac{12}{n}. && (1.7) \end{aligned}$$

In Equation (1.6) we have used that $\sum d_i$ equals twice the number $2n - 3$ of edges, a property already used in the proof of Theorem 1.1.2.

Part (ii) is left as an exercise. For the left inequality, use the contraction map of Theorem 1.1.2. For the other one, use the arguments of Part (i), but flip through an "anti-standard" triangulation, that is to say, a triangulation with no internal edge at the given vertex i. □

Part (ii) says that the bound in Part (i) is not too bad, but the following statement says more; it gives the exact diameter for almost all values of n:

Theorem 1.1.6 (Sleator, Tarjan, Thurston). *The diameter of the graph of flips of an n-gon is* $2n - 10$ *for all sufficiently large values of n.*

No purely combinatorial proof of the lower bound implicit in this theorem is known. The proof contained in [295] is far from elementary and we will avoid all the details, but we will sketch the main idea. We wish to give a lower estimate on how many flips are necessary to move from one triangulation $\mathscr{T}$ to another $\mathscr{T}'$. For this, we associate to every sequence of flips going from $\mathscr{T}$ to $\mathscr{T}'$ an abstract simplicial complex in the following fashion:

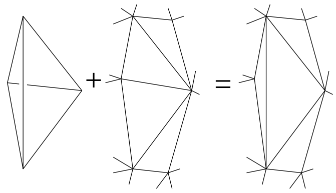

Figure 1.17: A diagonal flip viewed as glueing a tetrahedron to a surface.

Start with one of the triangulations, $\mathscr{T}$, and consider it as an abstract 2-dimensional simplicial complex. Then attach to this complex one tetrahedron for each flip in the sequence, in the same order. More precisely, if

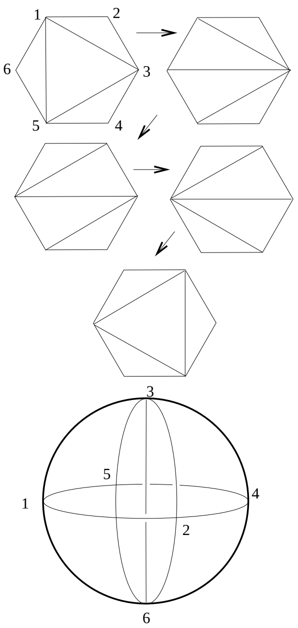

Figure 1.18: Four flips in a hexagon and the corresponding triangulated 3-ball.

a flip exchanges the diagonals pq and rs, add the tetrahedron $pqrs$ to the already constructed simplicial complex. This tetrahedron has two triangles in common to the complex we had before the flip, and it creates two new boundary triangles (see Figure 1.17). This interpretation of flips is also very popular in the context of Delaunay triangulations [127], as we will see in Theorem 3.4.2 of Chapter 3.

After we do this for all the flips in the sequence we get a simplicial complex that is topologically a 3-dimensional ball $\mathscr{B}$ and whose boundary, a 2-dimensional sphere, is made by glueing $\mathscr{T}$ and $\mathscr{T}'$ along their boundaries. (This is unless $\mathscr{T}$ and $\mathscr{T}'$ have interior edges in common; one step in the proof is to show that there is no loss of generality in assuming this.) Figure 1.18 shows an example where four flips in a hexagon give a triangulated ball with the four tetrahedra $\{1,3,5,6\}$, $\{1,2,3,6\}$, $\{3,4,5,6\}$, $\{2,3,4,6\}$. With this we get the following interesting result:

Lemma 1.1.7. *If the triangulated 2-sphere obtained by glueing $\mathscr{T}$ to $\mathscr{T}'$ along their boundaries cannot be extended to a triangulation of the 3-ball with less than k tetrahedra, then every sequence of flips going from $\mathscr{T}$ to $\mathscr{T}'$ has at least length k.*

In other words, the problem of finding paths in the flip-graph of a polygon is related to the problem of finding combinatorial simplicial 3-polytopes that do not admit small triangulations. This gives an easy way to remember where the number $2n-10$ comes from: It follows easily from Euler's formula that a triangulated 2-sphere with n vertices has exactly $2n-4$ triangles and, if $n \geq 13$, a vertex of degree at least six. Using this we can easily triangulate the 3-ball into at most $2n-10$ tetrahedra: just cone the highest degree vertex to all the triangles not incident to it.

But to obtain *lower bounds* for the flip distance we need to find 3-polytopes without small triangulations and, what is more difficult, prove that they do not have small triangulations. How can one do that? The key idea in [295] is to use hyperbolic polytopes. A useful fact about hyperbolic 3-space is that the volume of all hyperbolic tetrahedra is bounded by a certain constant σ, so that a hyperbolic 3-polytope of volume V clearly needs at least V/σ many tetrahedra to be triangulated. What the paper [295] proves is that for sufficiently big n, there are n-vertex hyperbolic 3-polyhedra with volume $(2n-10)\sigma$. We must remark, however, that this paper does not say how big n needs to be for the lower bound to be exact. The conjecture is that $n = 13$ is enough, and a related conjecture is the following:

Conjecture 1.1.8. *For every $n \geq 13$ there is a simplicial 3-polytope with n vertices whose interior cannot (even combinatorially) be triangulated with less than $2n-10$ tetrahedra.*

Let us also mention that the same idea of using hyperbolic volumes, now in arbitrary dimension, was used by W. Smith to give the best lower bound known for the size of the smallest triangulations of combinatorial n-dimensional cubes [296].

Triangulating and computing volumes are intimately related activities. Since the volume formula of a simplex in Euclidean space is just a determinant, an easy way to compute the Euclidean volume of a polytope is to add the volumes of simplices of any triangulation of it. Of course, for this to be a general algorithm we need the fact that every convex polytope can be triangulated (if you do not see why, read Proposition 2.2.4 in the next chapter).

Volume computations are useful throughout mathematics. For example, the calculation of volumes of hyperbolic convex polytopes has become of interest in topology. The reason is that every hyperbolic manifold can be obtained by identifying the faces of a convex polytope in hyperbolic space and its volume is a topological invariant. The volume has been used in the classification of hyperbolic manifolds (see [265] for references). (It should be said, however, that the calculation of the volume of a simplex in hyperbolic space is much more complicated than in Euclidean space).

The computation of volumes of polytopes in Euclidean space is also used in algebra [44, 138, 307]. But more important for us are the fascinating connections to combinatorics [301]. Here we show how the computation of volume is equivalent to counting linear extensions of posets, and that the linear extensions are simplices on a triangulation! This was first observed by R. Stanley in [300]:

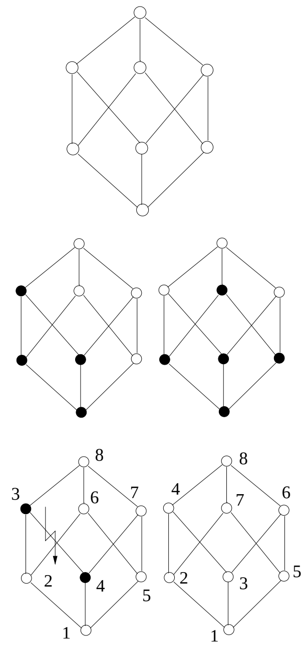

Figure 1.19: A poset P with two order ideals and two bijections $P \to \{1,\ldots,8\}$. Only one of them is a linear extension.

Definition 1.1.9. We define:

(i) A partially ordered set (or *poset*) is a finite set P with an ordering $<$ that is reflexive, antisymmetric, and transitive.

(ii) A *linear extension* of a poset on n vertices is a bijection λ from the set of vertices of P to $\{1,\ldots,n\}$ such that $\lambda(a) < \lambda(b)$ whenever $a < b$ in P.

(iii) An *order ideal* of a poset is a subset of the poset P such that if $a \in I$ and $b < a$, then $b \in I$.

Usually, a poset is represented by a graph, its *Hasse diagram*. We recommend the reader Chapter 3 of [303] for a thorough discussion of posets. Here we simply show in Figure 1.19 the Hasse diagram for the poset of subsets of the set $\{1,2,3\}$ ordered by containment, as well as two of its order ideals (black dots) and two linear orderings, only one of them extending the partial ordering. Given a poset P with elements $a_1,\ldots,a_n$, one can define the *order polytope* $\mathbf{O}(P)$ in $\mathbb{R}^n$ (see [300]) by the following linear constraints:

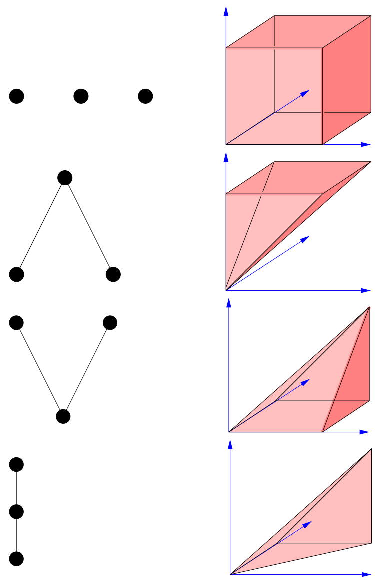
Figure 1.20: Some posets with three elements and their order polytopes.

$$\mathbf{O}(P) = \left\{ (x_1,\ldots,x_n) \in \mathbb{R}^n \,:\, 0 \le x_i \le 1, \text{ and } x_i \ge x_j \text{ if } a_i > a_j \text{ in } P \right\}.$$

Theorem 1.1.10. *The following properties hold for the order polytope* $\mathbf{O}(P)$ *of a poset* P*:*

(i) The vertices of the order polytope $\mathbf{O}(P)$ *are vectors with* $0-1$ *entries. They are in bijection with the order ideals of the poset* P*.*

(ii) The number of distinct linear extensions of the poset P *equals the number of simplices in a maximal size triangulation of the order polytope* $\mathbf{O}(P)$*.*