Dos problemas de combinatoria geométrica: Triangulaciones eficientes del hipercubo; Grafos planos y rigidez.

(Two problems in geometric combinatorics: Efficient triangulations of the hypercube; Planar graphs and rigidity).

David Orden Martín
Universidad de Cantabria

Tesis doctoral dirigida por el Profesor Francisco Santos Leal.

Santander, Marzo de 2003.

Índice

Agradecimientos **iii**

Preámbulo **v**

Preamble **vii**

1 Introducción y Preliminares **1**
1.1 Triangulaciones y subdivisiones de politopos 1
1.1.1 Preliminares sobre Teoría de Politopos 2
1.1.2 Triangulaciones de cubos 6
1.2 Teoría de Grafos 11
1.2.1 Preliminares y relaciones con politopos 11
1.2.2 Inmersiones planas de grafos: Teorema de Tutte 15
1.3 Teoría de Rigidez 20
1.3.1 La matriz de rigidez 20
1.3.2 Tensiones y grafos recíprocos 23
1.3.3 Rigidez infinitesimal y grafos isostáticos 26
1.4 Pseudo-triangulaciones de cuerpos convexos y puntos 29
1.4.1 Pseudo-triangulaciones de cuerpos convexos 29
1.4.2 Pseudo-triangulaciones de puntos 31
1.5 Resultados y Problemas abiertos 35
1.5.1 Triangulaciones eficientes de cubos 35
1.5.2 Grafos sin cruces, rigidez y pseudo-triangulaciones 36

1 Introduction and Preliminaries **41**
1.1 Triangulations and subdivisions of polytopes 41
1.1.1 Preliminaries about Polytope Theory 42
1.1.2 Triangulations of cubes 46
1.2 Graph Theory 50
1.2.1 Preliminaries and relations with polytopes 51
1.2.2 Plane embeddings of graphs: Tutte's Theorem 54
1.3 Rigidity Theory 59
1.3.1 The rigidity matrix 59
1.3.2 Stresses and reciprocal graphs 62

1.3.3 Infinitesimal rigidity and isostatic graphs 65
1.4 Pseudo-triangulations of convex bodies and points 67
1.4.1 Pseudo-triangulations of convex bodies 68
1.4.2 Pseudo-triangulations of points 69
1.5 Results and Open problems 73
1.5.1 Efficient triangulations of cubes 73
1.5.2 Non-crossing graphs, rigidity and pseudo-triangulations . . . 74

2 Asymptotically efficient triangulations of the d-cube 79
2.1 Introduction 79
2.2 Overview of the method and results 82
2.3 Polyhedral subdivision of $P \times \Delta^{k_1+\cdots+k_m-1}$ induced by one of $P \times \Delta^{m-1}$ 84
2.4 Triangulation of $P \times \Delta^{k_1+\cdots+k_m-1}$ induced by one of $P \times \Delta^{m-1}$ . . . 86
2.5 A triangulation of $P \times Q$ 88
2.6 Interpretation of our method via the Cayley Trick 92
2.7 Computation of $\rho_{l,m}$ 95
2.7.1 Smallest weighted efficiency of triangulations of $I^2 \times \Delta^{m-1}$. 95
2.7.2 Smallest weighted efficiency of triangulations of $I^3 \times \Delta^{m-1}$, for $m = 2, 3$ 97

3 The polytope of non-crossing graphs on a planar point set 101
3.1 Introduction 101
3.2 The graph of all pseudo-triangulations of $\mathcal{A}$ 105
3.2.1 Marked non-crossing graphs on $\mathcal{A}$. 109
3.3 The polyhedron of marked non-crossing graphs on $\mathcal{A}$ 110
3.4 Valid choices of f 116
3.5 Points in special position 120
3.5.1 The graph of all pseudo-triangulations of $\mathcal{A}$ 120
3.5.2 The case with only interior collinearities 123
3.5.3 Boundary collinearities 128

4 Planar minimally rigid graphs and pseudo-triangulations 131
4.1 Introduction 131
4.2 Embedding combinatorial pseudo-triangulations 132
4.3 3-connectedness of G 137
4.4 Laman graphs admit c.p.p.t. labelings. 140
4.5 Infinitesimally rigid graphs admit c.p.t. labelings. 145

5 Rigidity circuits, reciprocal diagrams and pseudo-triangulations 149
5.1 Introduction 149
5.2 Non-crossing spherical frameworks with non-crossing reciprocals . . . 152
5.3 Proof of Theorems 5.2.1 and 5.2.3 154
5.4 Rigidity circuits. The case of pseudo-triangulations 156

Agradecimientos

Quiero expresar aquí mi gratitud a todas las personas que de una u otra forma me han apoyado durante la realización de este trabajo. Muy especialmente a mis padres, pues entre lo mucho que les debo está el haber puesto siempre la mejor de sus voluntades para allanar el camino de la vida académica que ahora termino. Tanto ellos como mi hermano aportaron palabras de aliento al día a día de una investigación que muchas veces les era lejana. Además de su apoyo, a Encarna le agradezco también de una forma especial su compañía durante este tiempo.

Mi reconocimiento es también para aquellos amigos y compañeros que, cada uno en su medida, colaboraron y pusieron en tantas ocasiones un necesario paréntesis en mi trabajo. Asimismo, quiero hacer constar mi gratitud por el apoyo, a menudo incomparable en lo profesional y personal, de Paco Santos, sin cuya guía este trabajo habría resultado a buen seguro mucho menos llevadero. Otros miembros del departamento, desde la entrada en esta facultad hasta hoy, facilitaron también mi andadura y me aportaron en ocasiones sus opiniones y consejos, por lo que les estoy agradecido.

Finalmente, quiero mencionar a: Günter Ziegler, en cuyo grupo de la Technische Universität Berlin realicé una estancia en otoño de 2000. Emo Welzl y Jürgen Richter-Gebert, que permitieron mi asistencia a sus cursos en el ETH de Zürich. Komei Fukuda y Marc Noy, quienes accedieron a recomendar mi investigación. Jesús de Loera, por su calurosa acogida durante mi estancia en la University of California at Davis en el otoño de 2001, y por ayudarme a intentar mejorar algunos cálculos de esta tesis e invitarme a impartir una conferencia en Agosto de 2002. Ferrán Hurtado, a cuya iniciativa debo tanto mi inclusión en una acción integrada de la Universitat Politècnica de Catalunya como la solicitud de otras ayudas. Franz Aurenhammer, por aceptar mi estancia en la Technische Universität Graz en el próximo otoño. Günter Rote, que colaboró en parte de los resultados de esta tesis. Ruth Haas, Brigitte Servatius, Hermann Servatius, Diane Souvaine, Ileana Streinu y Walter Whiteley también aportaron conversaciones útiles a algunos de estos resultados.

Preámbulo

Esta tesis está dividida en dos partes independientes, aunque ambas proporcionan construcciones en combinatoria geométrica. En la primera parte, se aborda el estudio de métodos para construir triangulaciones "sencillas" de hipercubos de dimensión alta. Se ha demostrado que la obtención de triangulaciones eficientes del d-cubo regular $I^d = [0,1]^d$ tiene aplicaciones de interés como el cálculo de puntos fijos o la resolución de ecuaciones diferenciales por métodos de elementos finitos. En particular, se ha dedicado especial atención a tratar de determinar la triangulación más pequeña (en número de símplices [1]) o, al menos, el comportamiento asintótico de ese número mínimo de símplices. Los resultados previos se deben a Sallee (1984) y Haiman (1991). En el Capítulo 2 desarrollamos una construcción para obtener triangulaciones eficientes del producto de dos politopos generales. Aplicado al caso particular del producto de dos cubos, demostramos que este método mejora las cotas obtenidas por Haiman para el tamaño asintótico de una triangulación del d-cubo. Nuestra construcción comienza con una triangulación "eficiente" del producto entre un cubo y un símplice de dimensiones pequeñas. Este trabajo tiene una parte puramente teórica (cómo usar las triangulaciones de dimensión baja en objetos de dimensión alta) y una parte computacional (el cálculo de la triangulación óptima, para nuestros propósitos, de los objetos de dimensión baja).

La segunda parte de la tesis trata sobre las relaciones entre grafos planos, rigidez y pseudo-triangulaciones [2] de un conjunto de puntos en el plano. Éstas tienen ricas propiedades combinatorias y han sido aplicadas con un amplia variedad de propósitos (planificación de movimientos, visibilidad, estructuras de datos cinéticas) durante la última década. En el Capítulo 3 consideramos un conjunto $\mathcal{A}$ de puntos en el plano y, utilizando propiedades que relacionan las pseudo-triangulaciones con la teoría clásica de la rigidez (algunos de cuyos precursores son Euler y Cauchy), construimos un politopo simple cuyos vértices son todas las posibles pseudo-triangulaciones de $\mathcal{A}$. Al contrario que las construcciones similares propuestas previamente para triangulaciones, podemos describir completamente las ecuaciones de las facetas de nuestro politopo, lo que implica que se puede utilizar programación lineal para optimizar funcionales sobre él. Por otro lado, la estructura de caras del politopo es esencialmente la estructura de todos los grafos sin cruces que es posible dibujar en $\mathcal{A}$. La

[1] Un *d-símplice* es la envolvente convexa de $d+1$ puntos afínmente independientes.

[2] Un *pseudo-triángulo* es un polígono simple con exactamente tres vértices convexos.

construcción del politopo sirve incluso para un conjunto de puntos en posición no general, lo que conduce a una definición natural de pseudo-triangulación para ese caso.

En cuanto a sus relaciones con rigidez, Streinu (2000) demostró que las pseudo-triangulaciones puntiagudas [3] son grafos planos minimalmente rígidos. Demostramos que el inverso es también cierto: todo grafo plano minimalmente rígido admite una inmersión como pseudo-triangulación puntiaguda en $\mathbb{R}^2$, incluso bajo ciertas restricciones combinatorias como la imposición a priori de los tres ángulos convexos de cada cara. Este tipo de problemas de inmersión de grafos han sido ampliamente estudiados en la literatura, y nuestro resultado proporciona una caracterización de los grafos planos minimalmente rígidos que se añade a la debida a Laman. Yendo más lejos, obtenemos varios resultados que nos conducen a conjeturar que se puede extender la anterior caracterización para relacionar grafos planos rígidos y pseudo-triangulaciones generales.

Finalmente, el Capítulo 5 tiene que ver con la relación de Maxwell (1864) entre diagramas de fuerzas y figuras recíprocas. Obtenemos las condiciones para que un diagrama de fuerzas plano con un único equilibrio tenga un recíproco sin cruces. En particular, mostramos que si el diagrama es una pseudo-triangulación con exactamente un vértice no puntiagudo, entonces el recíproco es también una pseudo-triangulación del mismo tipo.

[3] En las que todo vértice es *puntiagudo*, e.d. tiene un ángulo incidente mayor de 180 grados.

Preamble

This thesis is divided into two independent parts, although both of them provide constructions in geometric combinatorics. The first part deals with the study of methods to construct "simple" triangulations of high dimensional hypercubes. Obtaining efficient triangulations of the regular d-cube $I^d = [0,1]^d$ has been proved to have interesting applications such as calculating fixed points or solving differential equations by finite element methods. In particular, it has brought special attention the problem of finding the smallest size (in number of simplices [4]) of a triangulation or, at least, the asymptotic behavior of this smallest number of simplices. Previous results are by Sallee (1984) and Haiman (1991). In Chapter 2 we develop a construction to obtain efficient triangulations of the product of two general polytopes. Applied to the particular case of the product of two cubes, we show this method to improve the bounds obtained by Haiman for the asymptotic size of a triangulation of the d-cube. Our construction starts from an "efficient" triangulation of the product of a cube and a simplex of small dimensions. This work has a purely theoretic part (how to use the triangulations of small dimensional objects in high dimensional ones) and a computational part (the computation of the optimal triangulation, for our purposes, of the small dimensional objects).

The second part of the thesis deals with the relations between planar graphs, rigidity and pseudo-triangulations [5] of a planar point set. These have rich combinatorial properties and have been applied with a wide range of purposes in computational geometry (motion planning, visibility, kinetic data structures) during the last decade. In Chapter 3 we consider a planar point set $\mathcal{A}$ and, using properties which relate pseudo-triangulations to the classical rigidity theory (somme of whose precursors are Euler and Cauchy), we construct a simple polytope whose vertices are all the possible pseudo-triangulations of $\mathcal{A}$. Contrary to previously proposed similar constructions for triangulations, we can completely describe the equations for the facets of our polytope, what implies that linear programming can be used to optimize functionals on it. In addition, the face poset of the polytope is essentially the poset of all non-crossing graphs which can be drawn on $\mathcal{A}$. The construction of the polytope works even for a point set in non-general position, leading to a natural definition of pseudo-triangulation for that case.

[4] A *d-simplex* is the convex hull of $d+1$ affinely independent points.

[5] A *pseudo-triangle* is a simple polygon with exactly three convex vertices.

Concerning their relations with rigidity, Streinu (2000) proved that pointed [6] pseudo-triangulations are minimally rigid planar graphs. We prove the converse to be true as well in Chapter 4: Every minimally rigid planar graph admits an embedding in $\mathbb{R}^2$ as a pointed pseudo-triangulation, even under certain combinatorial restrictions such as fixing a priori the three convex angles of each face. This type of embeddability problems have been thoroughly studied in the literature, and our result provides a characterization of minimally rigid planar graphs which adds to the well-known one due to Laman. Going further, we obtain several results leading us to conjecture that the above characterization can be extended to relate rigid planar graphs and general pseudo-triangulations.

Finally, Chapter 5 is related to Maxwell's relation (1864) between diagrams of forces and reciprocal figures. We obtain the conditions of a planar diagram of forces with a unique equilibrium to have a non-crossing reciprocal. In particular, we show that if the diagram is a pseudo-triangulation with exactly one non-pointed vertex, then the reciprocal is again such a pseudo-triangulation.

[6] In which every vertex is *pointed*, i.e. has an incident angle greater than 180 degrees.

Capítulo 1

Introducción y Preliminares

En este capítulo nuestra intención es proporcionar la base para una comprensión más sencilla del trabajo presentado en los capítulos siguientes. A pesar de que éstos han sido escritos con el propósito de ser lo más auto-contenidos posible, hay algunas nociones y resultados que pueden ser o no familiares para el lector. Comenzamos aquí por los más básicos, de manera que se puedan introducir lectores de distintos campos. Además, los ejemplos y construcciones han sido elegidos de modo que jugarán un papel más adelante en la obtención y discusión de los resultados.

La mayor parte de los enunciados en este capítulo introductorio aparecen sin demostración. No obstante, siempre damos referencias en las que pueden encontrarse y donde un lector interesado puede buscar información adicional. Sí incluimos aquellas demostraciones que van a resultar cruciales en el resto de los capítulos, bien por la técnica utilizada, o bien porque vayamos a hacer uso de ellas más adelante.

La estructura del capítulo es como sigue: en la primera sección introducimos algo de Teoría de Politopos, para centrarnos a continuación en triangulaciones de cubos. La segunda sección está dedicada a Teoría de Grafos, comenzando por un sumario de vocabulario y resultados básicos. Mostramos después dos conexiones con politopos, antes de dar una respuesta parcial al problema de inmersión de grafos, que va a resultar un resultado clave en el Capítulo 4. Posteriormente, trasladamos nuestra atención a la rigidez de grafos inmersos en el plano y enunciamos una serie de nociones y resultados que terminan con la definición de rigidez en la que estamos interesados. La cuarta sección está dedicada a presentar las pseudo-triangulaciones, tanto de cuerpos convexos como de puntos. Finalmente, en la última sección especificamos los resultados obtenidos en esta tesis, así como algunos problemas que quedan abiertos y pueden resultar objeto de un futuro trabajo.

1.1 Triangulaciones y subdivisiones de politopos

Las raíces históricas del estudio de politopos se remontan a los matemáticos de la antigua Grecia. De hecho, los polígonos y poliedros parecen haber estado entre los más tempranos objetos de una investigación matemática sistematizada, cuyo

primer resultado reseñable es la enumeración de los famosos Sólidos Platónicos, que son la solución al problema de encontrar todos los cuerpos convexos *regulares* de dimensión 3; aquéllos cuyas caras son copias de un mismo polígono convexo y para los que el número de caras que inciden en cada vértice es siempre el mismo. Éste puede ser considerado el punto de partida del estudio de las conexiones entre geometría y combinatoria, lo que desde entonces hasta nuestros días, pasando por el *Libro XII* de los *Elementos* de Euclides, ha sido una de las más reseñables características de los politopos.

1.1.1 Preliminares sobre Teoría de Politopos

Entre los objetos básicos en geometría se cuentan *puntos*, *rectas*, *planos* o *hiperplanos*, subespacios afines de dimensiones $0, 1, 2$ y $d-1$ respectivamente, en un espacio afín $\mathbb{R}^d$. Concentremos nuestra atención en conjuntos finitos de puntos etiquetados y que generen $\mathbb{R}^d$. Los conjuntos de puntos que consideraremos a lo largo de esta tesis estarán habitualmente en *posición general*, lo que quiere decir que todo subconjunto de $n+1$ puntos genera $\mathbb{R}^n$. Pero las definiciones y resultados en esta subsección sirven igualmente para el caso en que aparezcan *degeneraciones*.

Como de costumbre, los elementos de $\mathbb{R}^d$ se pueden ver bien como puntos o bien como vectores. Es decir, se puede considerar $\mathbb{R}^d$ como un espacio afín o vectorial. La siguiente definición establece, desde ambos puntos de vista, la que probablemente sea la noción crucial en Teoría de Politopos:

Definición 1.1.1

- Un conjunto de puntos $K \subseteq \mathbb{R}^d$ es *convexo* si para cualquier par de puntos $x, y \in K$, el segmento recto $[x, y] = \{\lambda x + (1-\lambda)y : 0 \leq \lambda \leq 1\}$ está contenido en K.
- Un conjunto no vacío de vectores $Y \subseteq \mathbb{R}^d$ es un *cono* si para cada subconjunto finito de vectores $\{y_1, \ldots, y_k\} \subseteq Y$ todas sus combinaciones lineales con coeficientes positivos están contenidas en Y. Por convenio, todo cono contiene el vector cero (obtenido como combinación del conjunto vacío).

Está claro que toda intersección de conjuntos convexos es a su vez un convexo, por lo que la *envolvente convexa* de K, definida como la intersección de todos los conjuntos convexos en $\mathbb{R}^d$ que contienen a K, es el menor convexo conteniendo a K:

$$\mathrm{conv}(K) := \bigcap \{K' \subseteq \mathbb{R}^d : K' \supseteq K, K' \text{ convexo}\}$$

Para el caso de un conjunto finito de puntos $K = \{x_1, \ldots, x_n\} \subseteq \mathbb{R}^d$, la envolvente convexa coincide con el conjunto de todas sus *combinaciones convexas*

$$\mathrm{conv}(K) = \{\lambda_1 x_1 + \cdots + \lambda_n x_n : \lambda_i \geq 0, \sum_{i=1}^{n} \lambda_i = 1\}$$

Análogamente, para un conjunto finito de vectores $Y = \{v_1, \ldots, v_n\} \subseteq \mathbb{R}^d$, la *envolvente cónica* de Y, definida como la intersección de todos los conos en $\mathbb{R}^d$ que contienen a Y, se puede expresar como:

$$\text{cone}(Y) = \{\mu_1 y_1 + \cdots + \mu_n y_n : \mu_i \geq 0\}$$

Antes de estar preparados para introducir dos versiones diferentes del principal objeto de estudio en esta sección, los *poliedros* y *politopos convexos*, necesitamos definir la *suma de Minkowski* de dos conjuntos $P, Q \subseteq \mathbb{R}^d$ como

$$P + Q := \{p + q : p \in P, q \in Q\}$$

A lo largo de esta tesis no vamos a considerar poliedros o politopos no convexos. Es por esto por lo que en adelante eliminaremos la palabra "convexo" cuando nos refiramos a ellos:

Definición 1.1.2 En un espacio afín $\mathbb{R}^d$:

- Un *$\mathcal{V}$-poliedro* es la suma de Minkowski de los dos siguientes conjuntos; la envolvente convexa de un conjunto finito de puntos, y la envolvente cónica de un conjunto finito de vectores:

$$\text{conv}(K) + \text{cone}(Y)$$

- Un *$\mathcal{H}$-poliedro* es una intersección de un número finito de semiespacios cerrados.

- Un *$\mathcal{V}$-politopo* es un $\mathcal{V}$-poliedro acotado, esto es, la envolvente convexa de un conjunto finito de puntos $\text{conv}(K)$.

- Un *$\mathcal{H}$-politopo* es un $\mathcal{H}$-poliedro acotado, esto es, una intersección de un número finito de semiespacios cerrados que no contiene ningún *rayo* $\{x + \mu y : \mu \geq 0\}$ para cualquier $y \neq 0$.

La equivalencia de la $\mathcal{V}$-definición y la $\mathcal{H}$-definición parece "geométricamente clara". Sin embargo, no es un hecho trivial de probar. Referimos al lector a [84] para una demostración del siguiente enunciado, del cual se sigue como consecuencia la equivalencia de las $\mathcal{V}$ y $\mathcal{H}$-definiciones para politopos. Mencionemos únicamente que en el fondo de la demostración está el Lema de Farkas o, en otras palabras, la dualidad de la programación lineal.

Teorema 1.1.3 (Teorema fundamental para poliedros) *Un subconjunto $P \subseteq \mathbb{R}^d$ es un $\mathcal{V}$-poliedro si, y sólo si, es un $\mathcal{H}$-poliedro.*

La *dimensión* $\dim(P)$ de un politopo o poliedro es la de su *envolvente afín*, la intersección de todos los subespacios afines que contienen a P. Un *d-politopo* es un politopo de dimensión d en algún $\mathbb{R}^e$ ($e \geq d$). El primer ejemplo de d-politopo que vamos a considerar es, en el sentido combinatorio, el "más pequeño":

Ejemplo 1.1.4 *Un* d-símplice τ *es la envolvente convexa de* $d+1$ *puntos afínmente independientes cualesquiera en algún* $\mathbb{R}^e$. *Por ejemplo, los triángulos y los tetraedros son los* 2-*símplices y* 3-*símplices, respectivamente. Es fácil ver que todos los* d-*símplices son afínmente isomorfos al* d-símplice estándar Δ^d *definido como*

$$\Delta^d := conv\{e_1, \ldots, e_{d+1}\}, \qquad \text{donde } e_i \text{ es el } i\text{-ésimo vector unidad en } \mathbb{R}^{d+1}.$$

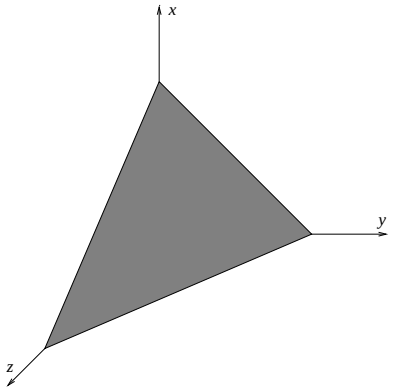

Figura 1.1: El 2-símplice estándar Δ^2.

Intuitivamente está claro lo que son las caras de un politopo; su caracterización matemática se basa en funcionales lineales:

Definición 1.1.5 Sea $P \subseteq \mathbb{R}^d$ un politopo convexo. Una *cara* de P es un conjunto de la forma

$$F = P \cap \{x \in \mathbb{R}^d : \langle c, x\rangle = c_0\},$$

donde $c \in \mathbb{R}^d, c_0 \in \mathbb{R}$ y la desigualdad $\langle c, x\rangle \leq c_0$ se satisface para todos los puntos $x \in P$. La *dimensión* de una cara es la de su envolvente afín.

Obsérvese que tanto el propio P como $\emptyset$ son caras del politopo P, definidas respectivamente por las desigualdades $\langle 0, x\rangle \leq 0$ y $\langle 0, x\rangle \leq 1$. Se les llama cara *impropia* y cara *trivial*. Los *vértices*, *aristas* y *facetas* de P son respectivamente las caras de dimensiones $0, 1$ y $\dim(P) - 1$; en particular, los vértices de P, denotados $\mathrm{vert}(P)$, son las caras no vacías minimales, y las facetas son las caras propias maximales. El k-*esqueleto* de un politopo es la unión de sus caras k-dimensionales. La *estrella* de una cara F, denotada $\mathrm{star}(F)$, se define como el conjunto de caras contenidas en la unión de todas las facetas que contienen a F.

Los dos siguientes tipos especiales de politopos nos introducen en el estudio de la combinatoria de estas estructuras geométricas:

Definición 1.1.6 Dado un politopo P de dimensión d,

- Se dice que P es *simple* si todo vértice está exactamente en d facetas.

- A P se le llama *simplicial* si todas sus facetas son símplices.

La información combinatoria de un politopo P va a estar codificada por su estructura de caras. Para estudiar ésta, necesitamos en primer lugar algunas nociones sobre *conjuntos parcialmente ordenados*, o "posets" (abreviatura de *partially ordered sets*). El lector interesado en el tema puede encontrar un estudio detallado en [16, Capítulo 11].

Definición 1.1.7 Un *poset* $(S, \leq)$ es un conjunto finito S junto con una relación "$\leq$" reflexiva, antisimétrica y transitiva.

- Un poset es *acotado* si tiene un único elemento minimal y un único elemento maximal.
- Un poset *Booleano* es uno isomorfo al poset $B_k = (2^{[k]}, \subseteq)$ de todos los subconjuntos de un conjunto de k elementos.
- Una *cadena* en un poset es un conjunto totalmente ordenado; la *longitud* de una cadena es su número de elementos menos 1.
- Un poset S es *graduado* si es acotado y todas las cadenas maximales tienen la misma longitud, que es llamada *longitud* o *rango* de S.
- Un *retículo* es un poset acotado tal que todo par de elementos $x, y \in S$ tienen una única cota superior minimal y una única cota inferior maximal en S.
- El *ideal inferior* de un elemento x en $(S, \leq)$ es el conjunto $\{y \in S : y \leq x\}$. El *filtro superior* es $\{y \in S : y \geq x\}$.

El siguiente resultado, bien conocido, describe la estructura de poset de las caras de un politopo, que es habitualmente referida como *poset de caras*:

Teorema 1.1.8 *Para todo politopo P, el poset de todas sus caras parcialmente ordenadas por la inclusión es un retículo graduado de longitud* $\dim(P) + 1$. *Por eso, se le llama también* retículo de caras *de P y se denota $L(P)$.*

Ejemplo 1.1.9 *En la Figura 1.2 mostramos el poset de caras de un pentágono convexo (izquierda), que es un retículo de longitud* 3*, y el de un* 3*-símplice (derecha), que es un retículo Booleano de longitud* 4*.*

Como avanzamos arriba, el retículo de caras es la noción apropiada para definir equivalencia combinatoria de politopos: dos politopos P y Q se dicen *combinatoriamente equivalentes* si sus retículos de caras son isomorfos, $L(P) \cong L(Q)$.

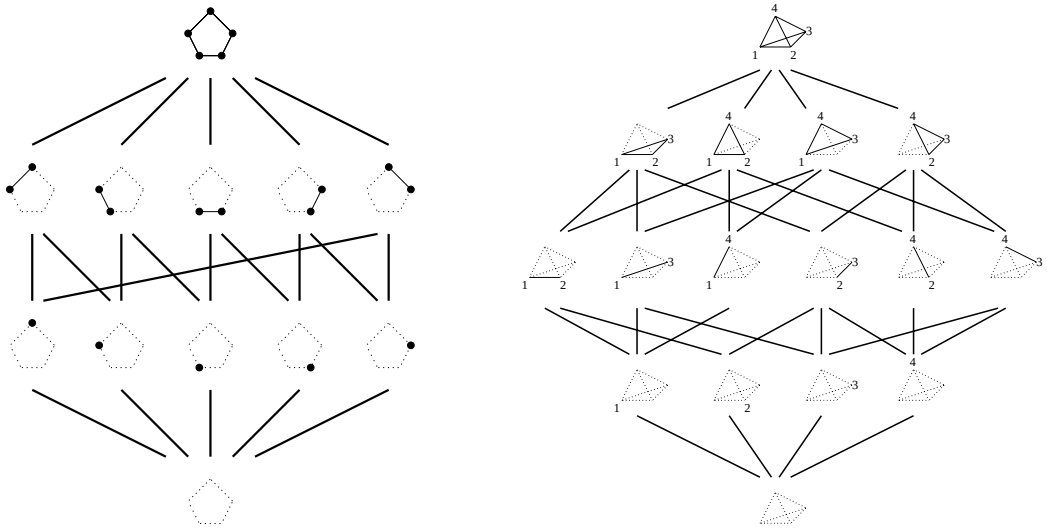

Figura 1.2: Retículos de caras de un pentágono convexo y un 3-símplice.

1.1.2 Triangulaciones de cubos

Cuando se tiene un objeto geométrico complicado (por ejemplo, un politopo "grande"), a veces resulta conveniente descomponerlo en partes más pequeñas y sencillas para simplificar su estudio:

Definición 1.1.10 Un *complejo politopal* $\mathcal{C}$ se define como una colección finita de politopos en $\mathbb{R}^d$ tal que:

(i) El politopo vacío está en $\mathcal{C}$,

(ii) Si $P \in \mathcal{C}$, entonces todas las caras de P están también en $\mathcal{C}$, y

(iii) La intersección $P \cap Q$ de dos politopos $P, Q \in \mathcal{C}$ es una cara de ambos.

La *dimensión* $\dim(\mathcal{C})$ es la mayor dimensión de un politopo en $\mathcal{C}$. El *conjunto soporte* de $\mathcal{C}$ es el conjunto de puntos $|\mathcal{C}| := \cup_{P \in \mathcal{C}} P$.

Como antes, la estructura combinatoria de un complejo politopal $\mathcal{C}$ es capturada por su *poset de caras* $(C, \subseteq)$, dado por el conjunto de politopos en $\mathcal{C}$ ordenados por inclusión. A continuación aparecen las definiciones que dan nombre a esta sección:

Definición 1.1.11 Dado un d-politopo P:

- Una *subdivisión politopal*, o simplemente *subdivisión*, de P es un complejo politopal $\mathcal{C}$ tal que su conjunto soporte es $|\mathcal{C}| = P$.
- Una *triangulación* de P es una subdivisión en la que todos los politopos en $\mathcal{C}$ son símplices.

- Dada una subdivisión $\mathcal{C}$ de P, un *refinamiento* es otra subdivisión $\mathcal{C}'$ tal que todo politopo de $\mathcal{C}$ es una unión de politopos de $\mathcal{C}'$.

Una descomposición de P en d-símplices que no se solapan y que no cumplen la condición (iii) de la Definición 1.1.10 se llama *disección* o, en caso de que todos los politopos sean símplices, *disección simplicial*. En particular, a lo largo de esta tesis estamos interesados en subdivisiones, triangulaciones y disecciones cuyo conjunto de vértices coincide con el de P, esto es, que no añaden vértices nuevos y tales que todo vértice es utilizado.

Observación 1.1.12 La misma definición sirve para *subdivisiones*, *triangulaciones*, *refinamientos* y *disecciones* de un conjunto finito de puntos $\mathcal{A}$, considerando el politopo definido por su envolvente convexa, excepto porque estamos asumiendo implícitamente que sólo se pueden utilizar como vértices puntos de $\mathcal{A}$. Además, en general supondremos que se utilizan todos.

Ejemplo 1.1.13 *La Figura 1.3 representa una triangulación (izquierda) y una disección simplicial (derecha) del cubo* 3*-dimensional. En el primer caso, la triangulación se compone de* 5 *símplices: los cuatro tetraedros obtenidos cortando esquinas no adyacentes y el tetraedro resultante, que usa diagonales opuestas de las caras superior e inferior. En el segundo caso, el cubo está partido en dos prismas triangulares, para una mejor visualización: cada prisma está triangulado en* 3 *tetraedros, pero las triangulaciones inducidas en la cara común a ambos prismas usan diagonales opuestas. Por lo tanto esos* 6 *tetraedros no forman una triangulación, puesto que se viola la condición (iii) de complejos politopales.*

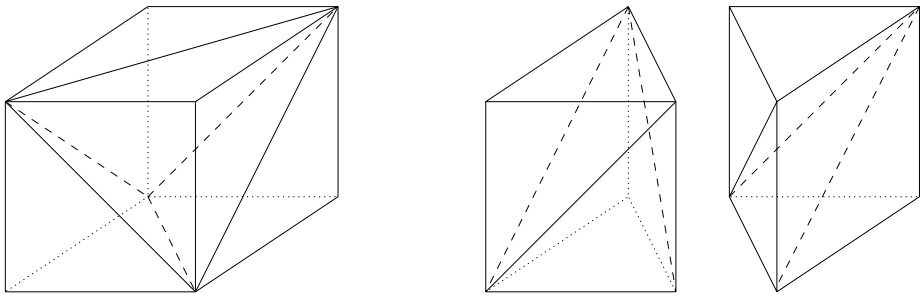

Figura 1.3: Izquierda: Una triangulación del 3-cubo. Derecha: Una disección del 3-cubo.

Señalemos aquí que toda subdivisión de un politopo P se puede refinar a una triangulación. Para ello, se necesitan triangulaciones T_i de cada uno de los politopos P_i en la subdivisión tales que, para todo par de índices i, j, T_i y T_j induzcan la misma triangulación en $T_i \cap T_j$ cuando ésta se considera como cara de T_i y de T_j. Existen construcciones bien conocidas (véase por ejemplo [76]) con esta característica.

El estudio de triangulaciones y sus propiedades ha atraído un gran interés durante las últimas décadas, especialmente estimulado por los avances en la computación, que han conducido a un buen número de aplicaciones en Geometría Computacional, entre las cuales podemos mencionar las relativas a gráficos por ordenador o simulaciones para ingeniería.

En particular, las triangulaciones de cubos han sido objeto de estudio por más de treinta años, y se tratan en el segundo capítulo de esta tesis, donde construimos triangulaciones "pequeñas" de cubos.

Definamos el *producto* de dos politopos generales como $P \times Q := \{(p, q) : p \in P, q \in Q\}$. Antes de considerar triangulaciones de cubos, vamos a concentrarnos en una manera de triangular el producto de dos símplices $\Delta^k \times \Delta^l$. La triangulación que mostramos aquí para un producto de este tipo es interesante por sí misma, debido a su simplicidad y fácil representación. Además, este interés quedará completamente justificado en el Capítulo 2, donde la generalizamos al producto de un politopo y un símplice. Un resultado clásico es que todas las triangulaciones de un producto $\Delta^k \times \Delta^l$ tienen el mismo número de símplices:

Lema 1.1.14 ([13]) *Toda triangulación de $\Delta^k \times \Delta^l$ tiene exactamente $\binom{k+l}{k}$ símplices.*

Nuestro propósito aquí es dar al lector una triangulación destacada de $\Delta^k \times \Delta^l$, que es conocida y utilizada desde hace tiempo en topología algebraica [30], y cuya fácil visualización la hace un buen ejemplo de lo útil que puede ser la combinatoria para entender estructuras geométricas en dimensiones altas. Esta utilidad se encuentra en el trasfondo de toda la tesis, y puede ser incluso más explícita en los Capítulos 2 y 3.

Asumamos que Δ^k y Δ^l tienen vértices $\{v_0, \ldots, v_k\}$ y $\{w_0, \ldots, w_l\}$ respectivamente. Entonces el producto $\Delta^k \times \Delta^l$ tiene por conjunto de vértices

$$\{(v_i, w_j) : 0 \leq i \leq k, 0 \leq j \leq l\},$$

que se puede representar como las casillas de una cuadrícula de tamaño $(k+1)\times(l+1)$ cuyas filas corresponden a vértices de Δ^k y cuyas columnas representan vértices de Δ^l. Véase el Ejemplo 1.1.15.

Consideremos ahora caminos del vértice (v_0, w_0) al vértice (v_k, w_l) cuyos pasos aumentan en una unidad bien el subíndice de v o bien el de w. En la representación, si etiquetamos filas y columnas de la cuadrícula de tal manera que (v_0, w_0) y (v_k, w_l) sean respectivamente las casillas inferior izquierda y superior derecha, estos caminos corresponden a escaleras monótonas, en las que cada paso es un movimiento desde una casilla bien a la siguiente por la derecha o bien a la inmediatamente superior. Cada uno de estos caminos selecciona un subconjunto de $k+l+1$ vértices de $\Delta^k \times \Delta^l$ y no es difícil ver que determina un símplice $(k + l)$-dimensional.

La colección de símplices dada por todos estos caminos, representada por la colección de todas las escaleras monótonas en la cuadrícula, forma una triangulación de $\Delta^k \times \Delta^l$, que se conoce como *triangulación escalera*. Nótese que el número de posibles caminos o escaleras es precisamente $\binom{k+l}{k}$.

Ejemplo 1.1.15 *La siguiente figura muestra la triangulación escalera del prisma triangular* $\Delta^1 \times \Delta^2$. *En la fila superior aparece la representación de su conjunto de vértices* $\{(v_i, w_j) : 0 \leq i \leq 1, 0 \leq j \leq 2\}$ *como las casillas de una cuadrícula de tamaño* 2×3. *Debajo mostramos las tres posibles escaleras en esta cuadrícula y los correspondientes tetraedros, que triangulan* $\Delta^1 \times \Delta^2$.

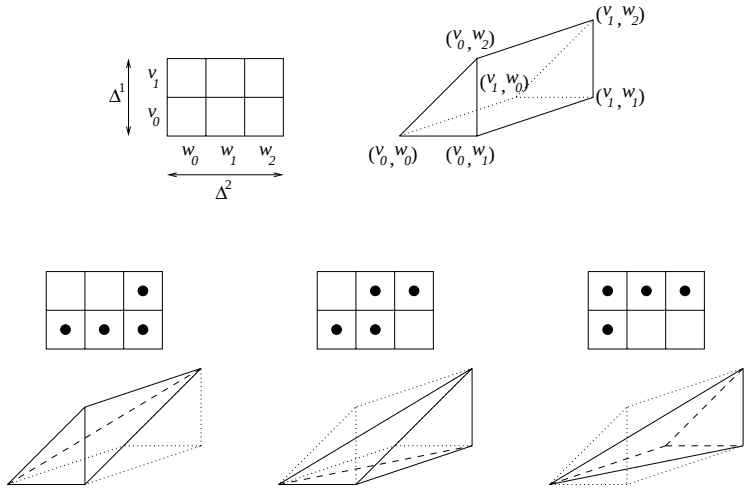

Figura 1.4: La triangulación escalera de $\Delta^1 \times \Delta^2$.

Trasladamos ahora nuestra atención a la familia de *d-cubos regulares* $I^d := [0,1]^d, d \geq 0$. Las triangulaciones "sencillas" del d-cubo regular tienen varias aplicaciones, como la resolución de ecuaciones diferenciales por métodos de elementos finitos o el cálculo de puntos fijos (veáse, por ejemplo, [72]). En particular, ha atraído una especial atención, tanto desde un punto de vista teórico como desde el práctico, la determinación del menor tamaño de una triangulación de I^d (ver [51, Sección 14.5.2] para un estudio reciente). Puntualicemos que el problema general de calcular la triangulación más pequeña de un politopo arbitrario es NP-completo incluso restringiéndose a dimensión 3, veáse [10].

Definición 1.1.16 El *tamaño* de una triangulación (o disección simplicial) T del d-cubo es su número de símplices $|T|$. El menor tamaño de triangulaciones de I^d se denota habitualmente por ϕ_d.

Definición 1.1.17 Dado un d-símplice $\tau = \text{conv}(p_1, \ldots, p_{d+1})$, su *volumen* se define como

$$\text{vol}(\tau) := \text{abs}\left(\begin{vmatrix} p_1 & \cdots & p_{d+1} \\ 1 & \ldots & 1 \end{vmatrix}\right)/d!$$

El mínimo volumen de un d-símplice que utilice los vértices de I^d es por tanto $1/d!$, y puesto que el del d-cubo es 1, el tamaño máximo de una triangulación de I^d

resulta ser $d!$ En el estudio antes mencionado se pueden encontrar un buen número de triangulaciones de tamaño máximo y de fácil construcción.

Si se nos dan dos triangulaciones de un d-cubo para d fija, es fácil decidir cuál de las dos es "mejor" (o "más sencilla"); la de menor tamaño. Sin embargo, en caso de que tengamos triangulaciones de d-cubos para distintas dimensiones d, necesitamos una noción que permita discernir cuál de ellas es la más conveniente. Ésta fue enunciada por primera vez por Todd en [72]:

Definición 1.1.18 La *eficiencia* de una triangulación T de un d-cubo se define como el número $(|T|/d!)^{1/d}$. La menor eficiencia $(\phi_d/d!)^{1/d}$ de triangulaciones de I^d se denota ρ_d.

Observación 1.1.19 Por la anterior observación sobre el tamaño máximo de una triangulación del d-cubo, la eficiencia de una triangulación de I^d es como mucho 1 y, cuanto menor sea, más eficiente es la triangulación. En particular, $\rho_d \leq 1, \forall d$.

En el segundo capítulo de esta tesis, presentamos una nueva construcción para triangular I^d, que conduce a buenas eficiencias asintóticas. Dicha construcción está basada en el siguiente método, debido a Haiman [40]:

Sean T_k y T_l triangulaciones de cubos regulares I^k e I^l respectivamente. El producto $T_k \times T_l$ de las dos triangulaciones (entendido como el producto por pares de sus símplices) da una descomposición del cubo $I^k \times I^l = I^{k+l}$ en $|T_k| \cdot |T_l|$ subpolitopos, cada uno de ellos isomorfo al producto de símplices $\Delta^k \times \Delta^l$.

Por el Lema 1.1.14, toda triangulación de $\Delta^k \times \Delta^l$ tiene tamaño $\binom{k+l}{k}$. Por tanto, refinando la subdivisión $T_k \times T_l$ de manera arbitraria se obtiene una triangulación de I^{k+l} de tamaño $|T_k| \cdot |T_l| \cdot \binom{k+l}{k}$, lo que lleva a:

Teorema 1.1.20 (Haiman) *Para todos k y l, ${\rho_{k+l}}^{k+l} \leq {\rho_k}^k {\rho_l}^l$.*

Esto implica que comenzando con una triangulación con una cierta eficiencia ρ de un cubo I^k, se puede construir una sucesión de triangulaciones de I^{nk} para $n \in \mathbb{N}$ con esa misma eficiencia. Lo que sugiere inmediatamente el siguiente resultado:

Corolario 1.1.21 *La sucesión $(\rho_n)_{n\in\mathbb{N}}$ es convergente, y*

$$\lim_{n\to\infty} \rho_n \leq \rho_d \qquad \forall d \in \mathbb{N}.$$

Demostración: Fijemos $d \in \mathbb{N}$ y $k \in \{1, \ldots, d\}$. El Teorema de Haiman implica que, para cada $i \in \mathbb{N}$,

$$\rho_{k+id} \leq {\rho_k}^{k/(k+id)} {\rho_d}^{id/(k+id)}.$$

Puesto que la expresión de la derecha converge a ρ_d cuando i crece, las d subsucesiones de $(\rho_n)_{n\in\mathbb{N}}$ consistentes en índices iguales módulo d, y por tanto la propia sucesión, tienen límite superior acotado por ρ_d. Esto se cumple para cada $d \in \mathbb{N}$. Un límite superior acotado por todos los términos de la sucesión debe coincidir con

el límite inferior. □

Denotamos por ρ_∞ y llamamos *eficiencia asintótica* de triangulaciones del cubo al límite de $(\rho_n)_{n\in\mathbb{N}}$. El corolario dice que la eficiencia asintótica de triangulaciones del cubo regular es a lo sumo la eficiencia mínima ρ_d para cada d. En particular, la mejor cota superior para ρ_∞ antes de nuestro trabajo era $\rho_\infty \leq \rho_7 = 0.840$. En el Capítulo 2 probamos $\rho_\infty \leq 0.8159$ con un método completamente nuevo.

1.2 Teoría de Grafos

Los investigadores en la materia coinciden en considerar que el artículo [31] de Euler (en el que se discutía el famoso problema de los puentes de Königsberg) marca el nacimiento de la Teoría de Grafos. Desde entonces, un buen número de matemáticos han aportado sus contribuciones a este campo, pero ha sido durante el último siglo cuando éste ha sido objeto de un gran interés, incrementado durante las últimas décadas por la sencillez de la representación de un grafo en ordenadores y el gran número de problemas que se pueden modelizar utilizando grafos.

En esta sección presentamos en primer lugar algunas relaciones entre grafos y politopos; la primera de ellas es el politopo de subdivisiones de un polígono convexo y está relacionada con los resultados del Capítulo 3. Las otras son interesantes propiedades de los grafos definidos por vértices y aristas de politopos. Finalmente, enunciamos un resultado de Tutte sobre inmersión de un grafo abstracto como uno geométrico, con puntos de $\mathbb{R}^2$ como vértices y segmentos rectos como aristas.

1.2.1 Preliminares y relaciones con politopos

Las nociones básicas que presentamos aquí le serán familiares al lector. Como referencia general sobre grafos recomendamos [15] ó [75], y para cada enunciado de esta sección especificamos dónde se puede encontrar una demostración.

Definición 1.2.1 Un *grafo* es un par de conjuntos $G = (V, E)$ donde $V = \{1, \ldots, n\}$ es un conjunto finito de *vértices* y $E \subseteq \{\{i, j\} : i, j \in V, i \neq j\}$ es un conjunto de *aristas* entre los vértices. Normalmente denotaremos las aristas por ij en lugar de $\{i, j\}$.

Por definición, los grafos que consideramos son *simples*; no tienen ni lazos (aristas que unen un elemento consigo mismo) ni aristas múltiples. En la siguiente definición resumimos parte del vocabulario básico. Denotamos por $\binom{V}{2}$ el conjunto de todos los pares, no ordenados, de vértices en V.

Definición 1.2.2 Sea $G = (V, E)$ un grafo.

- El número de aristas en E que contienen un vértice $v \in V$ se denomina *grado* de ese vértice.

- G es *regular de grado* k si todo vértice tiene grado k.
- G es *completo* si $E = \binom{V}{2}$. El grafo completo en n vértices se denota por K_n.
- El grafo obtenido a partir de G mediante el *borrado* de un conjunto de vértices $W \subseteq V$ se define como la restricción de G al conjunto de vértices $V \setminus W$, esto es; $G \setminus W := (V \setminus W, E \cap \binom{V \setminus W}{2})$.
- El conjunto de *vecinos* de un vértice v se define como $N(v) := \{w \in V : \{v, w\} \in E\}$.
- Un grafo $G_0 = (V_0, E_0)$ es un *subgrafo* de G, denotado $G_0 \subseteq G$, si $V_0 \subseteq V$ y $E_0 \subseteq E$.

Consideramos tres tipos distinguidos de subgrafos:

- Un *camino* en G entre dos vértices $v_1, v_k \in V$ (nótese que $1, k$ son subíndices y no necesariamente $v_i = i$) es un subgrafo $G_{v_1}^{v_k} \subseteq G$ con conjunto de vértices $V_{v_1}^{v_k} = \{v_1, \ldots, v_k\}$ y conjunto de aristas $E_{v_1}^{v_k} = \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots, \{v_{k-1}, v_k\}\}$.
- Un *k-ciclo* en G es un subgrafo $G_{v_1}^{v_1} \subseteq G$ con $V_{v_1}^{v_1} = \{v_1, \ldots, v_k\}$ y $E_{v_1}^{v_1} = \{\{v_1, v_2\}, \{v_2, v_3\}, \ldots, \{v_{k-1}, v_k\}, \{v_k, v_1\}\}$.
- Un *Y-grafo* en G de extremos $v_1, v_2, v_3 \in V$ es un subgrafo compuesto por un vértice $w \in V \setminus \{v_1, v_2, v_3\}$ y caminos $G_{v_1}^w, G_{v_2}^w, G_{v_3}^w$ que son disjuntos excepto por el punto común w.

Así como un tipo especial de grafos que aparecerán más adelante:

- G es un *grafo bipartito* si V es la unión disjunta de dos conjuntos V_1, V_2 y cada arista une un vértice de V_1 con otro de V_2. Se dice *completo* si sus aristas son todas las posibles entre V_1 y V_2. El grafo bipartito completo se denota por $K_{m,n}$, para m y n los cardinales de los conjuntos.

Los caminos y los ciclos se denotan a menudo por la secuencia de aparición de sus vértices, esto es, $(v_1, v_2, \ldots, v_k)$ o $(v_1, v_2, \ldots, v_k, v_1)$. Otra noción que nos interesa es la de conexión:

- G es *conexo* si para cada par de vértices distintos $v, w \in V$ hay un camino G_v^w que los une.
- G es *k-conexo* para $k > 1$ si para todo $v \in V$ el grafo $G \setminus v$ es $(k-1)$-conexo.

Así, un grafo abstracto $G = (V, E)$ está dado como un conjunto de vértices junto con un conjunto de pares no ordenados de vértices, y no se involucra ninguna geometría. Por otro lado, una *inmersión topológica* $G \hookrightarrow \mathbb{S}^2$ es una aplicación que envía los vértices V a puntos en la 2-esfera y las aristas $ij \in E$ a arcos. Se le llama inmersión *plana* o *sin cruces* si ningún par de arcos $p_i p_j$ y $p_k p_l$ correspondientes a aristas no adyacentes $ij, kl \in E, i, j \notin \{k, l\}$ tiene un punto en común.

Una *inmersión geométrica*, o un *grafo geométrico* $G(S) = (V, E, S)$ se define como un grafo $G = (V, E)$ junto con una aplicación $i \mapsto p_i \in S$ de los vértices V a un conjunto finito de puntos $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ en el plano euclídeo y de las aristas $ij \in E$ a segmentos rectos $p_i p_j$. Análogamente, es *plana*, o *sin cruces*, si ningún par de segmentos correspondientes a aristas no adyacentes se interseca. El Teorema de Fary [33] establece que se puede "enderezar" toda inmersión topológica plana y obtener una inmersión geométrica plana. Por otro lado, toda inmersión geométrica plana induce de manera trivial una inmersión topológica plana (simplemente proyectando a $\mathbb{S}^2$). Se dice que un grafo es *plano* si tiene una inmersión, topológica o geométrica, plana.

Las *celdas*, o *caras*, de una inmersión geométrica o topológica se definen como las componentes conexas del complementario (en $\mathbb{R}^2$ o $\mathbb{S}^2$) de la unión de aristas y vértices. Una inmersión geométrica plana tiene exactamente una cara no acotada, comúnmente llamada cara *exterior*.

En esta tesis sólo vamos a considerar inmersiones de grafos en $\mathbb{R}^2$ o $\mathbb{S}^2$. Usaremos las notaciones $G(S)$ y $G \hookrightarrow \mathbb{S}^2$ para inmersiones geométricas y topológicas, respectivamente. En ocasiones quizá abusemos de la notación y utilicemos el conjunto imagen S para referirnos a la inmersión geométrica $G(S)$. El siguiente resultado recoge algunos enunciados interesantes; sus demostraciones se pueden encontrar, por ejemplo, en [77]:

Teorema 1.2.3 *Sea G un grafo.*

(i) **(Teorema de Menger)** *G es k-conexo si, y sólo si, entre cada par de vértices hay k caminos $G_1, \ldots, G_k$ en G que son disjuntos excepto por los extremos.*

(ii) **(Teorema de Kuratowski)** *G es plano si, y sólo si, no tiene ningún subgrafo homeomorfo al grafo completo K_5 ni al grafo bipartito completo $K_{3,3}$.*

(iii) **(Teorema de Whitney)** *Si G es plano y 3-conexo, entonces todas sus inmersiones topológicas son equivalentes (tienen las mismas celdas).*

(iv) **(Teorema de Euler)** *Si G es conexo y plano, para cada inmersión plana $G(S)$ (o $G \hookrightarrow \mathbb{S}^2$) los números c de celdas, v de vértices y e de aristas están relacionados por*

$$v - e + c = 2$$

Consideremos ahora la clase de grafos geométricos *sin cruces* en un conjunto dado de vértices $\mathcal{A}$ en el plano $\mathbb{R}^2$. Este tipo de grafos es de interés en Geometría Computacional, Combinatoria Geométrica y áreas relacionadas. En particular, se ha dedicado mucho esfuerzo a la enumeración, recuento y optimización en el conjunto de los maximales de entre estos grafos, es decir, las triangulaciones. Nótese que aquí éstas se consideran con conjunto de vértices igual a $\mathcal{A}$, como se avanzó en la Observación 1.1.12. Puesto que cualquier triangulación de un conjunto de puntos con cardinal n tiene $n_i + 2n - 3$ aristas, para n_i el número de puntos en el interior de $\mathrm{conv}(\mathcal{A})$, es trivial una cota inferior de $2^{n_i+2n-3} = \Omega(4^n)$ grafos sin cruces en $\mathcal{A}$.

Una cota superior de tipo $O(c^n)$ para una constante c se demostró por primera vez en [5]. El mejor valor hasta la fecha para c es $59 \cdot 8 = 472$ [62].

Está claro que el conjunto de grafos geométricos sin cruces en un conjunto de puntos en el plano forma un poset con respecto a la inclusión de aristas. Sin embargo, hasta el desarrollo de esta tesis su estructura de poset solamente se comprendía bien cuando los puntos estaban en posición convexa. El artículo [34] contiene varios resultados enumerativos sobre grafos geométricos con n vértices en posición convexa. En particular, demuestra que hay $\Theta((6+4\sqrt{2})^n n^{-3/2})$ grafos sin cruces en total y da fórmulas explícitas para cada cardinal fijado.

Para puntos en posición convexa, los grafos sin cruces que contienen todas las aristas de la envolvente convexa coinciden con las subdivisiones poligonales del n-gono convexo. Es bien conocido que el poset que forman es el opuesto del poset de caras (menos la cara vacía) del $(n-3)$-*asociaedro*, un politopo simple de dimensión $n-3$ con las siguientes propiedades:

- Sus vértices se corresponden con todas las $\frac{1}{n-1}\binom{2n-4}{n-1}$ posibles maneras de poner paréntesis en una cadena de longitud $n-1$ o, equivalentemente, de multiplicar una expresión $a_1 a_2 \dots a_{n-1}$ cuando la multiplicación no es asociativa.
- Dos vértices son adyacentes si corresponden a una sola aplicación de la ley asociativa.

Así, los vértices del asociaedro se corresponden con las subdivisiones poligonales que tienen el máximo número de aristas, es decir, con triangulaciones.

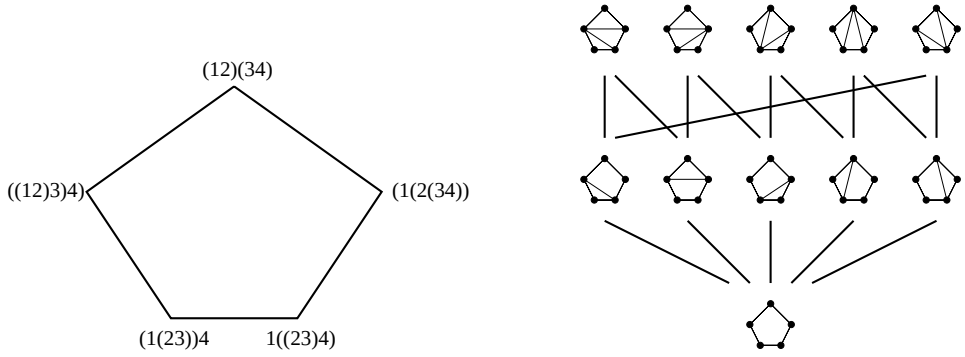

Figura 1.5: Izquierda: El 2-asociaedro. Derecha: Poset de subdivisiones poligonales del 5-gono convexo.

Ejemplo 1.2.4 *La Figura 1.5 representa el* 2*-asociaedro (izquierda), cuyos vértices corresponden a las* 5 *maneras de poner paréntesis en la cadena* 1234*. Por tanto es un pentágono, cuyo retículo de caras hemos encontrado ya en la Figura 1.2. Ignorando la cara vacía, se comprueba que éste coincide con el opuesto del poset de subdivisiones poligonales del* 5*-gono convexo, que mostramos en el dibujo de la derecha.*

El resultado principal en el tercer capítulo de esta tesis es una generalización del asociaedro; damos una descripción del poset de grafos sin cruces para cualquier conjunto de puntos, incluso en posición no general, como un subconjunto del poset de caras de cierto politopo que construimos.

Para una relación más directa entre grafos y politopos, obsérvese que los vértices y aristas de un politopo P (el 1-esqueleto) forman un grafo abstracto, al que se llama el *grafo de* P. Es natural la pregunta inversa sobre cuándo un grafo abstracto es el grafo de un politopo; para 3-politopos la respuesta la da el siguiente resultado clásico, cuya demostración original aparece en [66] y una muy comprensible se puede encontrar en [84]:

Teorema 1.2.5 (Steinitz) *G es el grafo de un politopo 3-dimensional si, y sólo si, es simple, plano y 3-conexo.*

La parte del "si" es la más complicada de demostrar. Para la parte del "y sólo si", puesto que los grafos de politopos son trivialmente simples y su planaridad puede comprobarse utilizando proyección radial sobre una esfera desde un punto interior, o una proyección lineal sobre un plano desde un punto más allá de una faceta, únicamente resta aplicar otro resultado bien conocido:

Teorema 1.2.6 (Balinski [8]) *El grafo de un d-politopo es d-conexo.*

1.2.2 Inmersiones planas de grafos: Teorema de Tutte

Por último, pero no por ello menos importante, introducimos el problema de inmersión de grafos. La construcción que presentamos para terminar esta sección da una respuesta a la siguiente pregunta: dado un grafo plano y 3-conexo G, ¿cómo podemos encontrar una inmersión geométrica en la que, además, las celdas sean polígonos convexos?

El Dibujo de Grafos (Graph Drawing) es un campo con una historia distinguida, y las inmersiones planas han recibido una atención sustancial en la literatura ([33], [73], [74], [28], [25], [63], [9]). También se han considerado recientemente (ver por ejemplo [55]) extensiones de inmersiones de grafos con segmentos rectos a pseudo-rectas. Es natural preguntarse cuando éste tipo de inmersiones se pueden "enderezar", esto es, cuándo se pueden dibujar con segmentos rectos manteniendo a su vez alguna subestructura combinatoria. De hecho, el resultado primordial sobre inmersión de grafos planos, el anteriormente mencionado Teorema de Fary [33], es simplemente un ejemplo de dar respuesta a una cuestión de este tipo.

Una explicación del transfondo de la siguiente construcción se puede dar fácilmente mediante este hecho físico: asumamos que las aristas de G son gomas elásticas. Tomemos los vértices correspondientes a una celda particular F_0 cuyo borde tenga m vértices. A continuación, pinchemos estos vértices en el plano $\mathbb{R}^2$ de modo que F_0 se convierta en un m-gono convexo y que todas las gomas elásticas interiores estén

en tensión. Voilà; la figura que se obtiene es una inmersión plana de G en la que todas las demás celdas vienen dadas como polígonos convexos que no se solapan.

Para modelizar las gomas elásticas se introduce la noción de *peso* $w_{ij} \in \mathbb{R}$ de una arista ij. Nótese que, puesto que estamos considerando aristas no dirigidas, se debe requerir la condición de simetría $w_{ij} = w_{ji}$.

Definición 1.2.7 Sean $G = (V, E)$ un grafo y $w : E \to \mathbb{R}$ una asignación de pesos a sus aristas. Consideremos también una asignación de posiciones $i \mapsto p_i \in \mathbb{R}^2$ a los vértices de G. Se dice que un vértice $i \in V$ está en *equilibrio* si

$$\sum_{ij \in E} w_{ij}(p_i - p_j) = 0.$$

Para un ejemplo, véase la Figura 1.6.

El siguiente importante resultado da una respuesta parcial a esta pregunta: ¿Cuándo un grafo se puede dibujar en el plano con aristas rectas, de tal manera que las celdas sean convexas y no se solapen? (Esto último equivale a "sin cruces")

Teorema 1.2.8 (Tutte) *Sea $G = (\{1, \dots n\}, E)$ un grafo plano y 3-conexo que tiene una celda $(k+1, \dots, n)$ para algún $k < n$. Sean $p_{k+1}, \dots, p_n$ (en este orden) los vértices de un $(n-k)$-gono convexo. Finalmente, sea $w : E' \to \mathbb{R}^+$ una asignación de pesos positivos a las aristas interiores. Entonces:*

(i) Existen unas únicas posiciones $p_1, \dots, p_k \in \mathbb{R}^2$ para los vértices interiores de modo que todos ellos están en equilibrio.

(ii) Todas las celdas de G son polígonos convexos que no se solapan.

Pese a que dar una demostración completa de este resultado está fuera de los propósitos de este capítulo introductorio, en el resto de esta sección damos un resumen de los pasos de la demostración que aparece en [60], la cual modificamos ligeramente en el Capítulo 4 para demostrar una versión dirigida del Teorema de Tutte. Incluímos aquí las nociones, resultados intermedios y demostraciones a las que nos referiremos en ese capítulo.
Demostración de (i): Asumamos dadas las posiciones de los puntos $p_{k+1}, \dots, p_n$ de la frontera. Sin pérdida de generalidad, podemos suponer $p_n = (0, 0)$. Considérese la función

$$E(x_1, \dots, x_k, y_1, \dots, y_k) = \frac{1}{2} \sum_{ij \in E'} w_{ij}((x_i - x_j)^2 + (y_i - y_j)^2) = \frac{1}{2} \sum_{ij \in E'} w_{ij} \|p_i - p_j\|^2,$$

que es cuadrática y no negativa en todo su dominio. Asumamos que el punto $z = (x_1, \dots, x_k, y_1, \dots, y_k)$ tiene al menos una entrada, digamos x_l, con un valor absoluto grande. Esto implica que p_l está lejos del punto $p_n = (0, 0)$ de la frontera. Puesto que el grafo G es conexo, existe un camino que une p_l con p_n; para al menos una arista ij en este camino, la distancia $\|p_i - p_j\|$ tiene que ser grande.

Así, para un $\alpha > 0$ suficientemente grande, la condición $|z| > \alpha$ implica $E(z) > E(0)$. Como E es cuadrática, esto implica que es estrictamente convexa (no degenerada con Hessiano definido positivo) y por tanto alcanza su único mínimo en $\{z : |Z| < \alpha\}$. La afirmación se sigue puesto que la condición $\nabla E = 0$ para un punto crítico de E:

$$\frac{\partial E}{\partial x_l} = \sum_{ij \in E'} w_{ij}(x_i - x_j) = 0 \quad \text{y} \quad \frac{\partial E}{\partial y_l} = \sum_{ij \in E'} w_{ij}(y_i - y_j) = 0, \quad \forall l \in \{1, \ldots, k\}$$

es exactamente la condición de equilibrio para los vértices interiores. □

Esquema de la demostración de (ii): De nuevo, asumamos que las posiciones de los puntos $p_{k+1}, \ldots, p_n$ del borde están dadas; por (i) las posiciones de los vértices interiores están a su vez unívocamente determinadas. Una configuración de puntos así se denomina *representación en equilibrio* de G.

Ejemplo 1.2.9 *En la Figura 1.6 se muestra una representación en equilibrio de un grafo G, junto con los pesos de cada arista interior. Damos una cuadrícula con el propósito de que el lector pueda comprobar fácilmente que los tres vértices interiores están en equilibrio.*

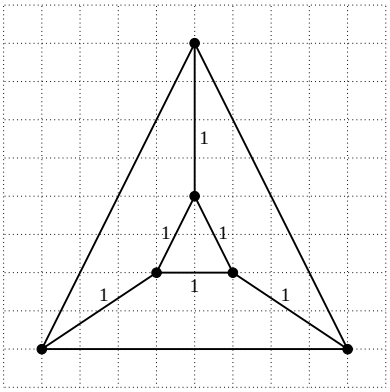

Figura 1.6: Un grafo en equilibrio.

El *interior relativo* de (la envolvente convexa de) una colección de puntos $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ se define como

$$\text{relint}(S) := \{\sum_{i=1}^{n} \lambda_i p_i : \sum_{i=1}^{n} \lambda_i = 1 \text{ and } \lambda_i > 0, \ \forall i = 1, \ldots, n\}.$$

El siguiente lema, cuya demostración es inmediata utilizando la condición de equilibrio, da una propiedad que resulta ser fundamental en el Teorema de Tutte. El enunciado puede comprobarse de nuevo en la Figura 1.6:

Lema 1.2.10 *En una representación en equilibrio $S := \{p_1, \ldots, p_n\} \in \mathbb{R}^{2n}$ de los vértices de G, todo vértice interior p está en el interior relativo de sus vecinos; $p \in relint(N(p))$.*

Definición 1.2.11 Sea G un grafo plano y 3-conexo con n vértices tales que los etiquetados como $k+1, \ldots, n$ forman, en este orden, una celda de G. Una configuración de puntos $S := \{p_1, \ldots, p_n\} \in \mathbb{R}^{2n}$ se dice que es una *buena representación* de G si:

(i) $p_{k+1}, \ldots, p_n$, en este orden, determinan un $(n-k)$-gono convexo, y

(ii) Para $i = 1, \ldots, k$ se tiene $p_i \in \text{relint}(N(p_i))$.

El Lema 1.2.10 establece que las representaciones en equilibrio son buenas representaciones. El Teorema de Tutte se sigue entonces de este hecho y la siguiente afirmación:

Teorema 1.2.12 ([60]) *Una buena representación de un grafo plano y 3-conexo G es una inmersión geométrica plana de G en la que todas las celdas interiores determinan polígonos convexos que no se solapan.*

La demostración se descompone a su vez en varias afirmaciones y tiene dos partes diferenciadas. En la primera, el siguiente lema demuestra que no pueden aparecer situaciones degeneradas, donde un vértice p_i se dice *degenerado* si el conjunto de sus vecinos $N(p_i)$ no genera afínmente $\mathbb{R}^2$. Su demostración puede evitarse en una primera lectura y, como apuntábamos antes, aparece aquí para una mejor comprensión de las diferencias con la demostración del Teorema de Tutte dirigido que damos en el Capítulo 4:

Lema 1.2.13 ([60]) *Una buena representación S de un grafo plano y 3-conexo G no tiene vértices degenerados.*

Demostración: Antes de comenzar la demostración, dada una recta $\ell = \{x : \phi(x) = \alpha\}$ para un funcional lineal ϕ, definamos un vértice a como *ℓ-activo* si sus vecinos no están todos en ℓ.

Es inmediato que los puntos del borde son no degenerados, puesto que son los vértices de un polígono estrictamente convexo. Considérese entonces un vértice interior degenerado $p \in S$, para el que debe haber una recta $\ell = \{x : \phi(x) = \alpha\}$ tal que $p \in \ell$ y $N(p) \subset \ell$.

Tomemos un punto $q \in S$ que no esté en ℓ. Puesto que G es 3-conexo, el Teorema de Menger 1.2.3.(i) garantiza la existencia de tres caminos A, B, C de p a q con aristas y vértices disjuntos, excepto por los extremos p, q. Como cada uno de ellos debe abandonar la recta ℓ en algún punto, contienen al menos un vértice ℓ-activo cada uno. Denotemos por a el primer punto ℓ-activo que aparezca después de p en el camino A. Entonces, el segmento inicial de este camino es de la forma

$A^0 = (p, a_1, \ldots, a_l, a)$ y todos sus puntos están en ℓ. Denotemos por b y c los primeros puntos ℓ-activos que aparezcan después de p en B y C respectivamente. Del mismo modo, sean B^0 y C^0 los correspondientes segmentos iniciales. Obsérvese que, todas juntas, las aristas de A^0, B^0 y C^0 forman un Y-grafo con extremos a, b, c, que denotamos Y^0 y que está contenido en ℓ.

El siguiente paso es demostrar que existen también Y-grafos Y^+ e Y^- con los mismos extremos a, b, c y cuyas aristas están todas respectivamente por encima y por debajo de la recta ℓ. Sólo mostramos la existencia de Y^+, pues la de Y^- se demuestra análogamente:

A lo sumo dos de los vértices del borde pueden estar en ℓ, por lo tanto entre los puntos a, b, c hay a lo más dos puntos del borde. Asumamos que a es interior. Para cada punto interior $q \in S$ la definición de buena representación implica $q \in \text{relint}(N(q))$, lo que a su vez implica la siguiente propiedad:

$$\begin{gathered}\text{Si hay un } q_- \in N(q) \text{ con } \phi(q) > \phi(q_-),\\ \text{entonces también hay un } q_+ \in N(q) \text{ con } \phi(q) < \phi(q_+)\end{gathered}$$

que se cumple también para puntos del borde que no sean maximales con respecto a ϕ. Como a es un punto ℓ-activo, hay o bien un punto $a_+ \in N(a)$ con $\phi(a) < \phi(a_+)$ o bien un punto $a_- \in N(a)$ con $\phi(a) > \phi(a_-)$, en cuyo caso la anterior propiedad implica la existencia también de un $a_+ \in N(a)$ con $\phi(a) < \phi(a_+)$.

Iterando, se puede generar un camino $A^+ = (a_0^+, \ldots, a_l^+)$ que conecte $a = a_0^+$ con un punto $a_l^+ = a^+$ del borde y maximal con respecto a ϕ. Análogamente, se construyen caminos B^+ y C^+ desde b y c a puntos del borde b^+ y c^+ maximales con respecto a ϕ.

Puesto que a lo sumo hay dos puntos del borde ϕ-maximales, o bien $a^+ = b^+ = c^+$ o bien corresponden a dos puntos conectados por una arista. Consideremos ahora el subgrafo G^+ de G que contiene los caminos A^+, B^+, C^+, el cual es conexo. Además, los tres vértices ℓ-activos a, b, c tienen grado uno en G^+. Así, la componente G^+ contiene un Y-grafo Y^+ de extremos a, b, c cuyas aristas están todas por encima de ℓ. De manera similar, existe un Y-grafo Y^- con los mismos extremos y todas las aristas por debajo de ℓ.

Por tanto, hay tres Y-grafos Y^0, Y^+ y Y^- de aristas disjuntas y con extremos a, b y c: forman un subgrafo homeomorfo al grafo bipartito completo $K_{3,3}$, lo que contradice la planaridad de G según el Teorema de Kuratowski 1.2.3.(ii). □

Para la segunda parte de la demostración del Teorema 1.2.12 se utiliza un argumento de consistencia global, dado por una serie afirmaciones que no incluímos aquí (referimos al lector a [60]), para probar que todas las celdas son de hecho convexas y no se solapan. □

1.3 Teoría de Rigidez

Fue de nuevo Euler quien en 1766 conjeturó "Una figura espacial cerrada no permite cambios, salvo que se la rompa" [32], estableciendo el punto de partida de la Teoría de la Rigidez. Sin embargo, el primer resultado importante publicado [18] se le atribuye a Cauchy en 1813: el grafo de todo 3-politopo simplicial es rígido. Esto constituyó un primer paso hacia el estudio de la Conjetura de Euler, que finalmente fue refutada en 1978 por R. Connelly [21]. Otra notoria contribución a la historia de la rigidez son los trabajos de J.C. Maxwell [52], quien estudió tensiones y fuerzas en grafos y sus relaciones con figuras recíprocas.

A pesar de que los resultados en Teoría de Rigidez han sido algo dispersos desde el establecimiento de la conjetura de Euler hasta su resolución, a partir de los años 70 un gran número de investigadores han dirigido sus esfuerzos a este campo, lo que ha conducido a una buena cantidad de aplicaciones a la ingeniería, véase por ejemplo [42]. Técnicas de Teoría de Rigidez se han aplicado también recientemente para resolver problemas en robótica o modelización molecular ([22], [69], [12]) abiertos durante bastante tiempo. Actualmente se trabaja en entender ciertas propiedades geométricas fundamentales de configuraciones moleculares por medio de este tipo de técnicas [47], con potenciales aplicaciones a plegamientos de moléculas [82], [71].

En este capítulo damos en primer lugar una serie de nociones y hechos, antes de dar una relación entre fuerzas y figuras recíprocas que aparecerá más adelante en nuestros resultados. A continuación damos la definición de rigidez infinitesimal genérica, que es la noción fundamental para nosotros, y finalmente presentamos los grafos con la propiedad de ser minimales entre los infinitesimalmente rígidos y esbozamos cómo se pueden construir. Consideramos principalmente inmersiones geométricas, pero también aparecerén inmersiones topológicas en la Subsección 1.3.2.

1.3.1 La matriz de rigidez

Una inmersión geométrica $G(S) = (V, E, S)$ de un grafo $G = (V, E)$ se suele llamar *armazón* (*framework*) cuando uno se refiere a sus propiedades de rigidez. El *armazón completo* para un conjunto de puntos V es el que tiene por aristas el conjunto completo $K = \binom{V}{2}$. Como observamos anteriormente, a lo largo de esta tesis consideraremos inmersiones geométricas sólo en $\mathbb{R}^2$, a pesar de que la mayor parte de esta sección sirve igualmente para inmersiones en un $\mathbb{R}^m$ general. En el plano afín (y en el espacio 3-dimensional) uno puede pensar en un armazón como un modelo matemático de una estructura física en la que cada vértice i corresponde a una articulación esférica colocada en el punto p_i y cada arista corresponde a una barra rígida que conecta dos articulaciones.

Obviamente, esta representación se puede utilizar por igual para describir estructuras rígidas (como puentes) o móviles (como máquinas o moléculas orgánicas). El objetivo de la Teoría de Rigidez es precisamente tratar la distinción entre armazones cuya representación sea rígida y los que admiten movimientos. En un primer

paso, se puede esperar que el que un armazón sea rígido o no dependa tanto del grafo $G = (V, E)$ como de la inmersión $G(S)$, lo que marca la diferencia entre la rigidez combinatoria y la geométrica. En el resto de esta sección nuestro objetivo es minimizar la importancia de la inmersión particular y concentrarnos principalmente en las propiedades del grafo. Con esta intención, introducimos las nociones básicas que se necesitan para entender los resultados sobre rigidez que presentamos en los capítulos siguientes. Salvo que se diga otra cosa, para los resultados de esta sección referimos al lector a [38].

Definición 1.3.1 Sea $G(S)$ una inmersión de un grafo G en un conjunto finito de puntos $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$. El conjunto S está en *posición general* y $G(S)$ es una *inmersión general* si ningún trío de elementos de S están alineados.

Si identificamos la inmersión $G(S)$ en el conjunto $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ con un punto $S = (p_1, \ldots, p_n) \in \mathbb{R}^{2n}$, podemos medir distancias de aristas mediante la *función de rigidez*

$$\begin{array}{rcl} r : \mathbb{R}^{2n} & \to & \mathbb{R}^{\binom{n}{2}} \\ S & \mapsto & r(S)_{ij} := \|p_i - p_j\|^2 \end{array}$$

que es diferenciable de manera continua. Entonces, la *matriz de rigidez* $R(S)$ de la inmersión $G(S)$ está definida como $r'(S) = 2R(S)$. Es una matriz de tamaño $\binom{n}{2} \times 2n$ cuya forma abreviada es

$$R(S) = \begin{pmatrix} p_1 - p_2 & p_2 - p_1 & 0 & 0 & \ldots & 0 & 0 \\ p_1 - p_3 & 0 & p_3 - p_1 & 0 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \ldots & p_{n-1} - p_n & p_n - p_{n-1} \end{pmatrix}$$

Será en términos de las filas de la matriz de rigidez como introduzcamos una noción crucial. Recuérdese que éstas corresponden a pares de vértices y por tanto a aristas del grafo completo en V:

Definición 1.3.2 Se dice que un conjunto de aristas $E \subseteq \binom{V}{2}$ es *independiente* con respecto a una inmersión S si las filas correspondientes de la matriz de rigidez son independientes.

Denotemos por $\delta(S)$ el determinante de un menor de tamaño $|E| \times |E|$ de la submatriz definida por E. Si hacemos variar S sobre todos los puntos de $\mathbb{R}^{2n}$, el conjunto solución de $\delta(S) = 0$ es, o bien todo $\mathbb{R}^{2n}$, o una hipersuperficie algebraica. Éste es el caso si hay una inmersión S con respecto a la cual E es independiente. Entonces, el conjunto de todas las inmersiones S para las que las filas correspondientes a E son dependientes está en la intersección $\mathcal{X}_E$ de las hipersuperficies algebraicas correspondientes a todos los menores de tamaño $|E| \times |E|$ de la submatriz. Si E es independiente con respecto a alguna inmersión, cada $\mathcal{X}_E$ es un conjunto algebraico cerrado de medida cero.

Definición 1.3.3 Una inmersión S se dice *genérica* si $S \notin \mathcal{X}$, donde

$$\mathcal{X} = \cup\{\mathcal{X}_E : \ E \text{ independiente para alguna inmersión}\}$$

Obsérvese que $\mathcal{X}$ es un conjunto cerrado de medida cero en $\mathbb{R}^{2n}$.

Definición 1.3.4 Un conjunto de aristas $E \subseteq \binom{V}{2}$ es *genéricamente independiente* si es independiente con respecto a todas las inmersiones genéricas de V.

Esta definición, junto con la de inmersión genérica, hace que se siga directamente este resultado:

Lema 1.3.5 *Sea V un conjunto finito de vértices. Si un conjunto de aristas $E \subseteq \binom{V}{2}$ es independiente con respecto a alguna inmersión de V, entonces es genéricamente independiente.*

Observemos que para una inmersión S en particular y un conjunto E de aristas, considerar los menores del conjunto de filas de $R(S)$ no es una manera muy práctica de comprobar independencia. En su lugar, normalmente se miran las relaciones de dependencia entre filas de la matriz de rigidez:

$$\sum_{ij \in E} w_{ij} r_{ij} = 0,$$

para r_{ij} la fila correspondiente a la arista ij. Un conjunto de aristas E es independiente con respecto a la inmersión S si, y sólo si, la única relación de dependencia que satisface el correspondiente conjunto de filas es la trivial (la que tiene $w_{ij} = 0$ para todo ij).

Definición 1.3.6 Una *tensión* (*stress*) en un armazón $G(S)$ es una asignación de escalares w_{ij}, no todos cero, a las aristas E de G de tal manera que, para cada vértice $i \in V$,

$$\sum_{ij \in E} w_{ij}(p_i - p_j) = 0.$$

Para una interpretación física, la asignación de estos escalares puede corresponder a la sustitución de cada arista ij por un muelle cuya compresión o extensión está indicada por el valor w_{ij}. Nótese que la última expresión es simplemente la anterior considerada de columna en columna. Entonces, con respecto a la inmersión S, el que un conjunto de aristas sea independiente, esto es, el que sus filas en la matriz de rigidez tengan sólo la dependencia trivial, es equivalente a que esas aristas tengan sólo la tensión trivial.

Se dice que un grafo tiene una *tensión* si sus aristas la tienen en cualquier inmersión genérica. Análogamente, un grafo es un *circuito de rigidez genérica* si para cualquier inmersión genérica sus aristas tienen una tensión no nula pero ningún subconjunto propio de aristas la tiene. En otras palabras, un *circuito* es un conjunto dependiente minimal de filas en la matriz de rigidez de cualquier inmersión genérica. Por tanto, un armazón $G(S)$ es un *circuito de rigidez* precisamente si tiene una tensión y ésta es distinta de cero en todas las aristas.

Ejemplo 1.3.7 *El lector se habrá percatado de que, dada una inmersión $G(S)$ de un grafo, tener una tensión en las aristas es equivalente a que todos los vértices (incluso los del borde) estén en equilibrio con sus vecinos. La diferencia con la situación en el Teorema de Tutte es que ahora permitimos pesos negativos en las aristas.*

En la siguiente figura añadimos los pesos de las aristas del borde de la Figura 1.6. Se puede comprobar fácilmente que los tres puntos del borde están también en equilibrio.

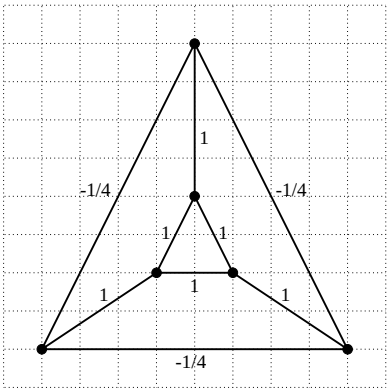

Figura 1.7: Una tensión en el grafo de la Figura 1.6.

1.3.2 Tensiones y grafos recíprocos

Hablando a grandes rasgos, en esta subsección mostramos que para aquellos armazones con tensiones tales que $w_{ij} \neq 0$ para toda arista ij, se puede dibujar siempre un dual con una interesante propiedad geométrica. Para ello, necesitamos en primer lugar algunas nociones:

Definición 1.3.8 Un *armazón esférico* se define como un par $(G \hookrightarrow \mathbb{S}^2, G(S))$ compuesto por una inmersión topológica plana $G \hookrightarrow \mathbb{S}^2$ y una inmersión geométrica (quizá no plana) $G(S)$, de un grafo G.

Un *armazón esférico sin cruces* es una inmersión geométrica plana $G(S)$ junto con la inmersión topológica plana $G \hookrightarrow \mathbb{S}^2$ que le corresponde trivialmente (la dada por $\mathbb{R}^2 \to \mathbb{S}^2$).

Definición 1.3.9 Dada una inmersión topológica plana $G \hookrightarrow \mathbb{S}^2$ de un grafo $G = (V, E)$ 2-conexo, el conjunto $\underline{E}$ de sus *parches de arista* (*edge patches*) se define como el que tiene por elementos las 4-tuplas ordenadas $\underline{e} = (i, j; h, k)$ tales que la arista que va del vértice i al vértice j tiene la cara h a la derecha y la cara k a la izquierda.

Por supuesto, la definición sirve también para inmersiones geométricas planas; de hecho, referimos al lector a la Figura 1.8 (izquierda) para una comprobación de

la noción de parche de arista. En ella, las aristas están denotadas por números y las caras por letras. Sin embargo, vamos a utilizar esta definición en el contexto de armazones esféricos, en el que la planaridad sólo está garantizada para la inmersión topológica.

Definición 1.3.10

- El *dual* $G' \hookrightarrow \mathbb{S}^2$ de una inmersión topológica plana $G \hookrightarrow \mathbb{S}^2$ fue introducido por Poincarè y es aquella descomposición de la esfera $\mathbb{S}^2$ que tiene un vértice por cada celda de G y en la que dos vértices están unidos por una arista si las celdas correspondientes son adyacentes en G. Obsérvese que el dual de un parche de arista $e = (i, j; h, k)$ es $e' = (h, k; j, i)$.

- Un armazón esférico $(G' \hookrightarrow \mathbb{S}^2, G'(S'))$ es un *recíproco* de otro armazón esférico $(G \hookrightarrow \mathbb{S}^2, G(S))$ si $G' \hookrightarrow \mathbb{S}^2$ es el dual de $G \hookrightarrow \mathbb{S}^2$ y para cada parche de arista $(i, j; h, k) \in \underline{E}$,

$$\langle p_i - p_j, p'_h - p'_k \rangle = 0.$$

  Es decir, las aristas de las inmersiones geométricas $G(S)$ y $G'(S')$ son perpendiculares.

La Figura 1.8 muestra las inmersiones geométricas $G(S)$ (izquierda) y $G'(S')$ (derecha) de un armazón esférico sin cruces y su recíproco. Los puntos $a, \ldots, d$ en el original representan caras.

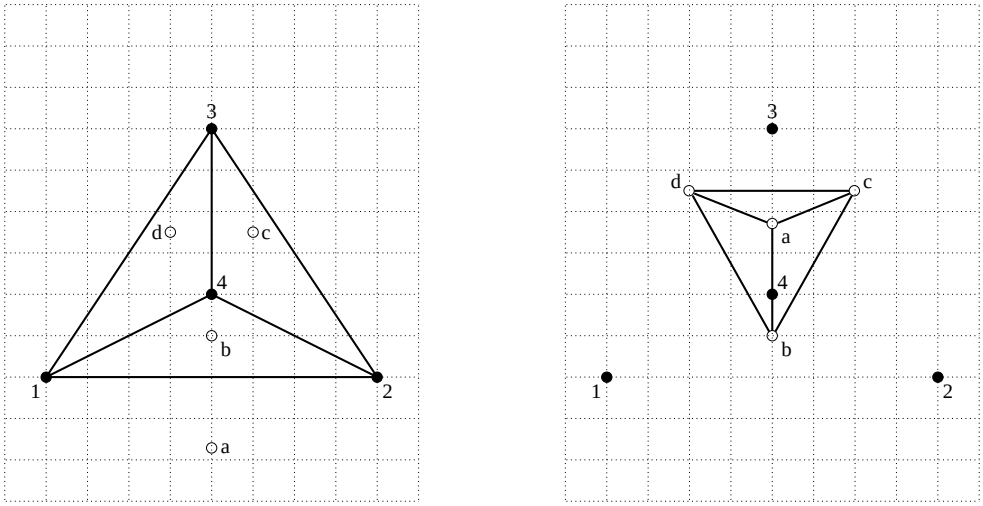

Figura 1.8: Inmersiones geométricas recíprocas.

Nótese que el dual de una inmersión topológica plana existe siempre. Por ello, dado un armazón esférico, la parte no trivial en la definición de recíproco es la existencia de $G'(S')$. Es por esto por lo que en el siguiente resultado abusamos de la notación y omitimos la parte topológica de los armazones esféricos:

Teorema 1.3.11 (Maxwell-Cremona) *Para un armazón esférico $G(S)$, son equivalentes:*

(i) $G(S)$ tiene un recíproco $G'(S')$,

(ii) $G(S)$ tiene una tensión $\sum_{ij\in E} w_{ij}(p_i - p_j) = 0$ tal que $w_{ij} \neq 0, \forall ij \in E$ (no nula en ninguna arista).

(iii) Hay un levantamiento lineal a trozos (quizá no convexo y con auto-intersecciones) de $G(S)$ a 3 dimensiones, lineal en cada cara.

Además, si las condiciones se cumplen para $G(S)$:

- *Las aristas con tensión positiva/negativa se corresponden con valles/crestas en el levantamiento de (iii). (En el sentido de que la pendiente aumenta/disminuye al cruzar la arista).*
- *Las condiciones se cumplen también para el recíproco $G'(S')$.*

Demostración: Nos referiremos a la equivalencia (i)⇔(ii), aparecida en [52], como Teorema de Maxwell. Damos aquí su demostración, que será crucial en el Capítulo 5. Para una demostración de la equivalencia con la tercera condición, veáse [24], [43] ó [78]. Nótese que los dos lados de un valle (resp. de una cresta) no van necesariamente "hacia arriba" (resp. "hacia abajo").

(ii)⇒(i) Asumamos que hay una tensión así. Elíjase un punto inicial p'_0 para el recíproco, correspondiendo a una celda F_0 del original, y defínase el punto correspondiente a otra celda F_c como $p'_c = p'_0 + (\sum_{e'} w_e d_e)^\perp$, donde la suma es sobre los parches de arista $e' = (h, k; j, i)$ de cualquier camino entre F_0 y F_c en el grafo dual, y donde $d_e = p_j - p_i$. Observar que éste es un proceso inductivo para el cual los puntos intermedios entre p'_0 y p'_k han sido previamente calculados. La primera tarea es demostrar que, dados dos caminos así entre F_0 y F_c, ambos conducen al mismo p'_c, lo que es cierto puesto que la suma a lo largo de un camino y de vuelta por el otro sigue una suma de ciclos de caras, y cada ciclo de caras es cero por la condición de equilibrio en el vértice recíproco. Puesto que ningún w_{ij} es cero, no hay ninguna arista degenerada en el recíproco.

(i)⇒(ii) Dado un recíproco, la construcción anterior se puede dar la vuelta, resolviendo el sistema $(w_e d_e)^\perp = p'_k - p'_h$ para recuperar la única tensión correspondiente en el armazón original. Puesto que la suma vectorial de aristas en cualquier ciclo de caras es cero, los vértices del original están en equilibrio. □

Observación 1.3.12 Sea $G'(S')$ el recíproco de $G(S)$ obtenido a partir de una cierta tensión con pesos w_{ij}. Sea $e = (i, j; h, k)$ un parche de arista en el original, entonces:

- $p'_k - p'_h = w_{ij}(p_j - p_i)^\perp$. Esto es; el segmento recto en el recíproco es perpendicular al correspondiente en el original, escalado por el valor absoluto de la tensión y en el sentido dado por el signo.

- $G'(S')$ tiene una tensión no nula en todas las aristas definida por $w'_{hk} := 1/w_{ij}$.

Ejemplo 1.3.13 *En la Figura 1.8 las posiciones de $a, \ldots, d$ en el recíproco (derecha) han sido obtenidas de acuerdo a la demostración anterior, una vez definida en el original la tensión $w_{14} = w_{24} = w_{34} = 1$ para las aristas interiores y $w_{12} = w_{23} = w_{13} = -1/3$ para las exteriores. Obsérvese que en el recíproco el punto a correspondiente a la cara exterior resulta estar por encima de b, lo que se corresponde con que la arista $\overline{12}$ tenga tensión negativa.*

Observación 1.3.14 Incluso para el caso de armazones esféricos sin cruces, el Teorema de Maxwell 1.3.11 afirma únicamente la existencia de un armazón esférico recíproco, que no se garantiza que sea sin cruces. En el Capítulo 5 caracterizamos las condiciones para que armazones esféricos sin cruces tengan recíproco también sin cruces.

1.3.3 Rigidez infinitesimal y grafos isostáticos

En esta subsección presentamos la noción de rigidez que vamos a utilizar en esta tesis, así como un tipo destacado de grafos, que son considerados los objetos fundamentales de la Teoría de Rigidez en dimensión 2. Volvemos a centrarnos en inmersiones geométricas.

Definición 1.3.15 Sea $G(S) = (\{1, \ldots, n\}, E, S)$ un armazón y considérese una asignación de un vector 2-dimensional v_i para cada vértice i. Tal asignación $v \in \mathbb{R}^{2n}$ se dice que es un *movimiento infinitesimal* del armazón si para toda arista $ij \in E$ se cumple la siguiente ecuación homogénea:

$$\langle v_i - v_j, p_i - p_j \rangle = 0.$$

El conjunto de movimientos infinitesimales de un armazón se denota $\mathcal{V}(E)$ y es claramente un subespacio lineal de $\mathbb{R}^{2n}$.

Cada v_i se puede interpretar como un vector de velocidad para el punto p_i en el plano. Los movimientos infinitesimales también son llamados a veces *flexiones* en la literatura sobre rigidez. Como interpretación, considérese la asignación de vectores v como las velocidades iniciales de un movimiento de los puntos p_i, tales que ni comprimen ni estiran las barras del armazón. Es por esto por lo que un vector v tal que

$$\langle v_i - v_j, p_i - p_j \rangle \geq 0,$$

esto es, tal que aumenta o mantiene la longitud de las barras, se dice que es un *movimiento infinitesimal expansivo*.

En este punto aparece la definición precisa de rigidez en la que estamos interesados:

Definición 1.3.16 Se dice que un armazón $G(S) = (\{1, \ldots, n\}, E, S)$ es *infinitesimalmente rígido* si $\mathcal{V}(E)$ es igual al subespacio $\mathcal{V}(K)$ de movimientos infinitesimales del armazón completo $G(K)$, que son las traslaciones o rotaciones de todo el grafo.

Un grafo G es *genéricamente infinitesimalmente rígido* si es infinitesimalmente rígido para cada inmersión genérica S.

Lema 1.3.17 ([38]) *Si $G(S)$ es infinitesimalmente rígido para alguna inmersión general S de V, entonces el grafo G es genéricamente infinitesimalmente rígido.*

En lo sucesivo nos centraremos en rigidez infinitesimal en el sentido genérico y así, cuando nos refiramos a un grafo como "infinitesimalmente rígido" querremos decir "genéricamente infinitesimalmente rígido". Un grafo $G = (V, E)$ se dice *isostático* o *minimalmente (infinitesimalmente) rígido* si es (infinitesimalmente) rígido y su conjunto E de aristas es independiente. La rigidez infinitesimal es entonces álgebra lineal, con grafos infinitesimalmente rígidos, tensiones y grafos isostáticos correspondiendo a conjuntos generadores, dependencias lineales y bases del espacio de filas de la matriz de rigidez del grafo completo.

Observemos que, para cada arista ij, la condición $\langle v_i - v_j, p_i - p_j \rangle = 0$ determina un hiperplano en $\mathbb{R}^{2n}$; el de los vectores v que mantienen la longitud de la barra $p_i p_j$. Por lo tanto, cada armazón tiene asociado un arreglo de hiperplanos por medio de sus aristas. Orientar los hiperplanos de acuerdo a si las barras resultan ser estiradas o comprimidas determina la llamada *matroide de rigidez*.

Para terminar la presente sección, consideramos el siguiente tipo de grafos, que dan una caracterización combinatoria de la rigidez infinitesimal:

Definición 1.3.18 Un grafo $G = (V, E)$ se dice *de Laman* si cumple las siguientes condiciones:

(i) $|E| = 2|V| - 3$, y

(ii) Para todo subgrafo (V', E') con $|V'| \geq 2$ vértices, $|E'| \leq 2|V'| - 3$.

La condición (ii), que a veces de denomina *propiedad de Laman hereditaria*, se puede reformular equivalentemente como:

(ii') Todo subconjunto V'' de vértices con $|V''| \leq |V| - 2$ es incidente a al menos $2|V''|$ aristas.

Hay toda una variedad de maneras de caracterizar los grafos de Laman. Aquélla en la que estamos más interesados para los propósitos de esta tesis, puesto que en el Capítulo 4 daremos una versión combinatoria, es la *construcción de Henneberg* [41]. Funciona inductivamente del siguiente modo: para un grafo de Laman (abstracto) en n vértices, se comienza con un triángulo para $n = 3$. En cada paso, se añade un nuevo vértice de una de las dos siguientes formas:

- **Henneberg I** (inserción de un vértice): El nuevo vértice se conecta por dos aristas a dos de los antiguos vértices.
- **Henneberg II** (partición de una arista): Se añade un nuevo vértice en alguna arista, que es partida en dos nuevas aristas y después conectada a un tercer vértice. Equivalentemente, se borra una arista y después se conecta un nuevo vértice con sus extremos y con algún otro vértice.

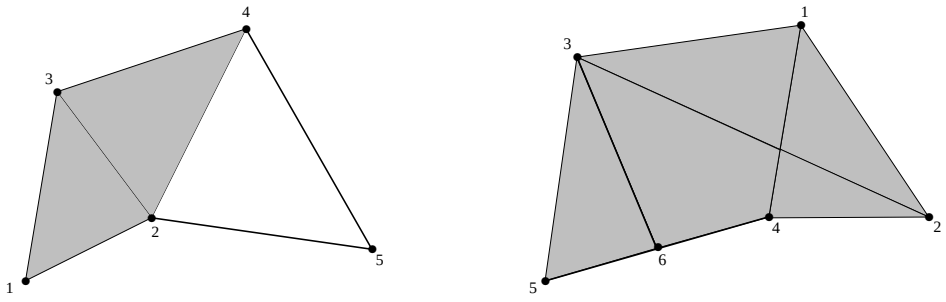

Figura 1.9: Pasos de la construcción de Henneberg.

Ejemplo 1.3.19 *Véase la Figura 1.9 para un ejemplo: los puntos han sido añadidos de acuerdo a sus etiquetas y las caras obtenidas anteriormente al presente paso están sombreadas. En la figura de la derecha hemos elegido un dibujo con aristas que se cruzan, para enfatizar que la construcción de Henneberg sirve para grafos* generales *y no sólo para los planos.*

El siguiente resultado muestra que los grafos de Laman son los conjuntos independientes maximales, y justifica por qué son fundamentales en la Teoría de Rigidez. Véase [50], [38], [80]:

Lema 1.3.20 ([81]) *Sea $G = (V, E)$ un grafo con $|V| \geq 2$. Entonces, las siguientes condiciones son equivalentes:*

(i) G es isostático,

(ii) G es infinitesimalmente rígido con $|E| = 2|V| - 3$,

(iii) G es independiente con $|E| = 2|V| - 3$,

(iv) G es un grafo de Laman (Teorema de Laman, [50]),

(v) Existe una construcción de Henneberg para G (Teorema de Henneberg, [41]).

Para concluir esta sección sobre rigidez, incluimos aquí el siguiente resultado que relaciona rigidez con politopos a través de los grafos de éstos. Para $d = 3$ fue probado por Dehn [29] en 1916, y la generalización a $d > 3$ es debida a Whiteley [79]:

Teorema 1.3.21 *El grafo de un d-politopo simplicial P es infinitesimalmente rígido.*

Mencionemos que el caso $d = 3$ es una consecuencia inmediata del teorema enunciado por Cauchy [18] en 1813 que citamos al comienzo de la sección, cuya demostración resultó tener varios errores menores y fue completada posteriormente (ver [67]).

1.4 Pseudo-triangulaciones de cuerpos convexos y puntos

Las pseudo-triangulaciones son objetos relativamente nuevos, pero que han sido utilizados ya en muchas aplicaciones en Geometría Computacional, entre ellas visibilidad [56, 58, 57, 65], "ray shooting" [37], estructuras de datos cinéticas [1, 49] y planificación de movimientos para brazos robóticos [69]. En este último artículo Streinu introdujo las *pseudo-triangulaciones mínimas* o *puntiagudas* y las utilizó para demostrar el Teorema de la Regla del Carpintero (del que una primera demostración fue dada un poco antes por Connelly et al. [22]).

Tienen ricas propiedades combinatorias, poliedrales y de teoría de rigidez, la mayor parte de las cuales no se han comenzado a investigar hasta recientemente, véanse también [61], [59], [48], [3], [11], [2], [17]. Encontrar su análogo 3-dimensional, lo que es probablemente el principal problema abierto sobre pseudo-triangulaciones y movimientos expansivos, podría conducir a la obtención de algoritmos eficientes de planificación de movimientos para ciertas clases de mecanismos 3-dimensionales, con potencial impacto en la comprensión de procesos de plegado de proteinas. Introducimos aquí las pseudo-triangulaciones en el contexto de visibilidad entre cuerpos convexos, en el que aparecieron originariamente, antes de trasladar nuestra atención a pseudo-triangulaciones de puntos, aquéllas en las que nos vamos a centrar en esta tesis.

1.4.1 Pseudo-triangulaciones de cuerpos convexos

A pesar de que en este trabajo no tratamos este tipo de pseudo-triangulaciones, su importancia deriva del hecho de que fue precisamente en este contexto en el que Pocchiola y Vegter comenzaron el estudio de pseudo-triangulaciones en los artículos pioneros [56] y [57], donde pueden encontrarse los resultados contenidos en esta subsección. Un precedente aparece en las *triangulaciones geodésicas* de polígonos introducidas en [19].

Pocchiola y Vegter consideran una colección $\mathcal{O} = \{B_1, \ldots, B_n\}$ de cuerpos convexos en el plano disjuntos dos a dos, que asumen *estrictamente convexos* (lo que significa que el segmento abierto que une dos puntos cualesquiera está contenido en el interior del objeto), con frontera *suave* C^1-diferenciable y en *posición general* en el sentido de que ningún trío de objetos comparte una recta tangente común. La siguiente definición recoge la mayoría de las nociones que se necesitan en esta subsección, y puede ser comprobada en la Figura 1.10:

Definición 1.4.1 En la situación anterior,

- Una *bitangente* para un par de cuerpos convexos es un segmento recto que es tangente a los dos objetos en sus extremos, en el sentido de que la recta que contiene al segmento es localmente una recta soporte. Hay cuatro de estas bitangentes; dos que separan los objetos y dos que no.
- El conjunto de *bitangentes* de una colección $\mathcal{O} = \{B_1, \ldots, B_n\}$ de objetos disjuntos dos a dos está compuesto por aquellos segmentos rectos que son bitangentes de algún par de objetos y que no intersecan a los demás. Se le llama *grafo de visibilidad tangente*.
- Una *pseudo-triangulación* de $\mathcal{O} = \{B_1, \ldots, B_n\}$ es la subdivisión del plano inducida por una familia maximal (con respecto a la inclusión) de bitangentes de $\mathcal{O}$ que no se cortan entre sí.

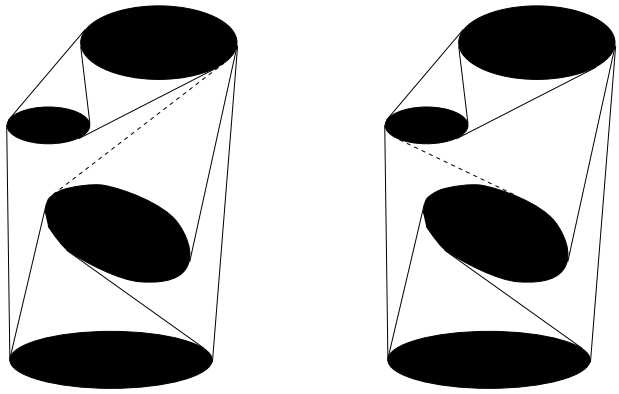

Figura 1.10: Flip entre dos pseudo-triangulaciones de una familia de cuerpos convexos.

Lema 1.4.2 *Toda pseudo-triangulación de n objetos convexos en posición general subdivide* $\mathrm{conv}(\cup_i B_i) \setminus \cup_i B_i$ *en $2n-2$ celdas, y el número de bitangentes es $3n-3$.*

El objetivo de Pocchiola y Vegter era el estudio del grafo de visibilidad tangente de una colección $\mathcal{O}$ de n cuerpos convexos. En particular, utilizaron pseudo-triangulaciones para obtener el primer algoritmo que construye el grafo en tiempo óptimo $O(n\log(n)+k)$ (para k el tamaño del grafo), utilizando sólo espacio lineal $O(n)$. En su algoritmo resultó ser crucial el análogo para pseudo-triangulaciones de la conocida noción de *flip diagonal* en triangulaciones (véase, por ejemplo, [46]):

Lema 1.4.3 *Dos celdas disjuntas T y T' tienen exactamente una recta tangente común. Es más, el borrado de una bitangente interior da una celda que se puede dividir de exactamente otra manera añadiendo otra bitangente. Esta operación se denomina* flip.

Se puede encontrar un ejemplo de flip entre las aristas de trazo discontinuo en la Figura 1.10.

Corolario 1.4.4 *De los dos lemas anteriores se deduce que toda pseudo-triangulación de $\mathcal{O}$ tiene exactamente $3n-3-n_b$ flips, donde n_b es el número de bitangentes en la frontera de* $\mathrm{conv}(\cup_i B_i)$.

Así, el *grafo de pseudo-triangulaciones* de $\mathcal{O}$, que tiene pseudo-triangulaciones por vértices y flips por aristas, es regular de grado $3n-3-n_b$. Pocchiola y Vegter demuestran que, además, es conexo.

1.4.2 Pseudo-triangulaciones de puntos

Consideremos un conjunto $\mathcal{A}$ de n puntos en posición general en el plano (una de las consecuencias de los resultados en el Capítulo 4 es una definición natural de pseudo-triangulaciones también en presencia de colinearidades). Una *triangulación* de $\mathcal{A}$ es una subdivisión del polígono $\mathrm{conv}(\mathcal{A})$ en triángulos, en el sentido de la Observación 1.1.12, asumiendo que *todos* los puntos de $\mathcal{A}$ son utilizados como vértices.

Definición 1.4.5 Llamemos *ángulos convexos* y *cóncavos* a aquellos que son respectivamente menores o mayores que π y observemos que la asunción de posición general implica la ausencia de ángulos iguales a π.

Un *pseudo-triángulo* es un polígono simple con únicamente tres vértices convexos (desde el interior), llamados *esquinas* y unidos por tres cadenas poligonales convexas hacia dentro, llamadas *pseudo-lados* del pseudo-triángulo. Véase la Figura 1.11.(a).

En general, un *pseudo-cuadrilátero*, *pseudo-pentágono*, etc. es un polígono simple con exactamente cuatro, cinco, etc. vértices convexos. Para ellos una *diagonal* se define como un segmento contenido en el pseudo-polígono y tangente a su frontera en los extremos. Nótese que todo k-gono convexo es un pseudo-k-gono. En particular, todo triángulo es un pseudo-triángulo. Notar también que todo polígono (quizá no convexo) es un pseudo-k-gono para algún k.

Una *pseudo-triangulación de* $\mathcal{A}$ es un grafo geométrico sin cruces con conjunto de vértices $\mathcal{A}$ y que subdivide $\mathrm{conv}(\mathcal{A})$ en pseudo-triángulos. Toda triangulación es una pseudo-triangulación.

Puesto que las triangulaciones de $\mathcal{A}$ coinciden con los grafos maximales sin cruces en $\mathcal{A}$, éstas son también las pseudo-triangulaciones maximales. Como es bien conocido, todas ellas tienen $2n_v+3n_i-3$ aristas, donde n_v y n_i denotan el número de puntos de $\mathcal{A}$ en la frontera y el interior de $\mathrm{conv}(\mathcal{A})$, respectivamente. También las pseudo-triangulaciones con el mínimo número posible de aristas resultan ser muy interesantes desde distintos puntos de vista:

Definición 1.4.6 Un vértice de un grafo geométrico se llama *puntiagudo* si sus aristas incidentes generan un ángulo menor de 180 grados desde ese vértice. El propio grafo se llama *puntiagudo* si todos sus vértices son puntiagudos.

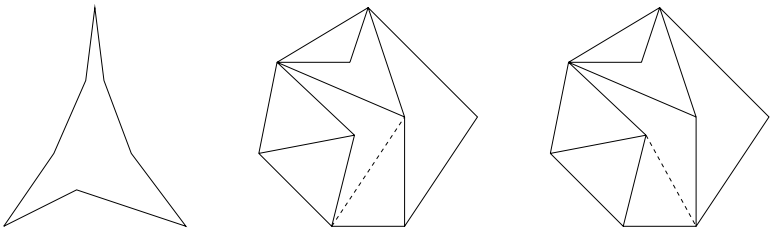

Figura 1.11: (a) Un pseudo-triángulo. (b) Una pseudo-triangulación puntiaguda. (c) Un flip de la arista discontinua en (b), que da otra pseudo-triangulación puntiaguda.

Una *pseudo-triangulación puntiaguda* es aquélla en la que todo vértice es puntiagudo.

El siguiente enunciado aparece originalmente en [69]. Damos aquí un esquema de su demostración, que se puede encontrar también en [61].

Proposición 1.4.7 (Streinu) *Sea $\mathcal{A}$ un conjunto de puntos en el plano como antes. Entonces:*

(i) Toda pseudo-triangulación de $\mathcal{A}$ con n_γ vértices no puntiagudos y n_ϵ vértices puntiagudos tiene $2n-3+n_\gamma = 3n-3-n_\epsilon$ aristas y $n-2+n_\gamma = 2n-2-n_\epsilon$ pseudo-triángulos.

(ii) Todo grafo en $\mathcal{A}$ puntiagudo y sin cruces tiene a lo sumo $2n-3$ aristas, y está contenido en alguna pseudo-triangulación puntiaguda de $\mathcal{A}$.

Demostración: La primera parte se sigue de la fórmula de Euler la cual, si llamamos t y e a los números de pseudo-triángulos y aristas, respectivamente, da $n-e+t=1$. Vamos a obtener una fórmula adicional por doble conteo del número de ángulos convexos: por un lado, cada uno de los pseudo-triángulos produce exactamente tres de éstos ángulos. Por el otro, para cada vértice no puntiagudo hay tantos ángulos convexos como aristas incidentes, y para cada vértice puntiagudo el número de ángulos convexos es uno menos que el anterior. Así, $2e - n_\epsilon = 3t$ y la fórmula anterior da los valores buscados de e y t.

Para la segunda parte, observamos en primer lugar que se puede asumir que el grafo es conexo, pues en caso contrario sus componentes se pueden analizar por separado. La condición de ser puntiagudo significa que el número de ángulos cóncavos es n. El número de ángulos convexos es al menos $3t$, con igualdad si, y sólo si, el grafo es una pseudo-triangulación. Puesto que el número total de ángulos suma $2e$, se obtiene $2e \geq n + 3t$ lo que junto con la fórmula de Euler implica $e \leq 2n-3$ (y $t \leq n-2$). La última frase en el enunciado se sigue entonces de dos hechos: por un lado, el que añadiendo *caminos geodésicos* (aquellos caminos sin cruces y de longitud mínima entre todos los caminos suficientemente próximos a ellos) entre

vértices convexos de un polígono el grafo se mantiene puntiagudo y sin cruces. Por otro, el hecho de que siempre haya un camino geodésico a través del interior de un polígono si éste no es un pseudo-triángulo. □

Observación 1.4.8 Obsérvese que todo vértice en la envolvente convexa de $\mathcal{A}$ es trivialmente puntiagudo. El número n_ϵ de vértices puntiagudos en una pseudo-triangulación va entonces desde n_v, en cuyo caso todos los pseudo-triángulos son de hecho triángulos y la pseudo-triangulación es una triangulación, hasta n para una pseudo-triangulación puntiaguda.

La parte (i) implica que, entre todas las pseudo-triangulaciones de $\mathcal{A}$, las puntiagudas tienen el mínimo número posible de aristas. Por esta razón se les llama también *pseudo-triangulaciones mínimas*.

La parte (ii) establece que las pseudo-triangulaciones puntiagudas coinciden con los grafos maximales puntiagudos y sin cruces.

Dada una configuración de puntos $\mathcal{A}$, es fácil construir una pseudo-triangulación puntiaguda de ella como sigue: en primer lugar, ordenar los puntos de $\mathcal{A}$ de acuerdo a su coordenada x. A continuación, considerar el triángulo definido por los tres primeros puntos en ese orden. Finalmente, añadir iterativamente las dos tangentes que unen cada nuevo punto con la pseudo-triangulación anteriormente construida. Véase la Figura 1.12. El Corolario 1.4.10 debajo mostrará que no es una coincidencia el que este procedimiento sea un caso particular de una construcción de Henneberg.

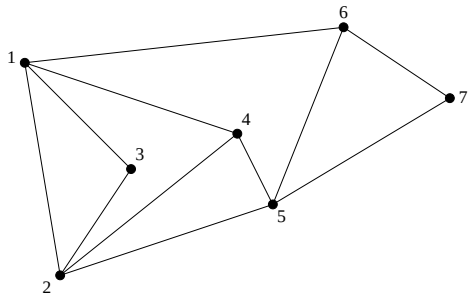

Figura 1.12: Construcción de una pseudo-triangulación puntiaguda.

Notemos que una pseudo-triangulación mínima (es decir, puntiaguda) es siempre "minimal" en el sentido de que el borrado de cualquier arista implica que ya no se tiene una pseudo-triangulación. Sin embargo, las pseudo-triangulaciones minimales no son necesariamente puntiagudas. Como ejemplo, considerar la Figura 1.13; la pseudo-triangulación de la izquierda es mínima (por tanto minimal), la del medio es minimal pero no mínima y la de la derecha no es minimal y por tanto tampoco mínima.

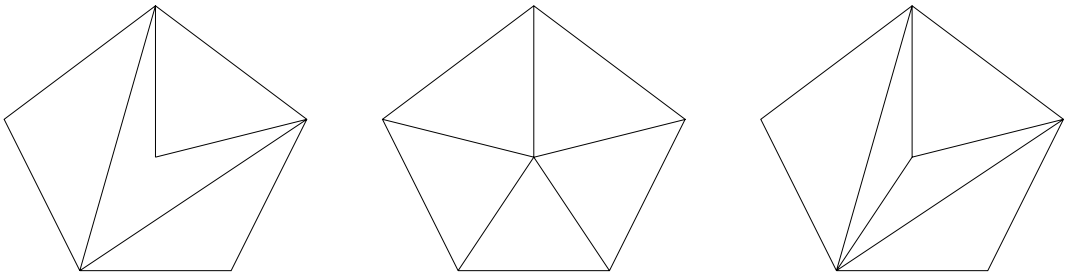

Figura 1.13: Pseudo-triangulaciones mínima, minimal y no minimal.

El siguiente resultado, tomado de [69], es un sumario de las anteriores caracterizaciones de pseudo-triangulaciones puntiagudas e introduce una nueva, que las relaciona con rigidez:

Teorema 1.4.9 (Caracterización de pseudo-triangulaciones puntiagudas)

Sea $G = (V, E)$ un grafo inmerso en el conjunto de puntos $\mathcal{A} = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^2$. Entonces, las siguientes propiedades son equivalentes:

(i) G es una pseudo-triangulación puntiaguda,

(ii) G es una pseudo-triangulación y $|E| = 2n - 3$ (equivalentemente, el número de pseudo-triángulos es $n - 2$),

(iii) G es una pseudo-triangulación mínima, esto es, tiene el mínimo número de aristas (equivalentemente, de pseudo-triángulos) entre todas las pseudo-triangulaciones de $\mathcal{A}$.

Es más, si alguna de las condiciones anteriores se satisface, entonces G cumple la propiedad de Laman hereditaria; para cada subgrafo (V', E') con $|V'| \geq 2$ vértices, $|E'| \leq 2|V'| - 3$.

Proof: (i)⇔(ii)⇔(iii) se sigue de la parte (i) de la Proposición 1.4.7. La propiedad de Laman se sigue de la parte (ii) de ése mismo resultado. □

Corolario 1.4.10 (Streinu) *Las pseudo-triangulaciones puntiagudas son grafos de Laman, por lo tanto isostáticos.*

En particular, el hecho de que además de ser grafos de Laman, cuando a una pseudo-triangulación puntiaguda se le elimina una arista de la envolvente convexa se convierta en un mecanismo expansivo [69], se ha demostrado crucial en el diseño de algoritmos eficientes para planificación de movimientos de brazos robóticos planos.

1.5 Resultados y Problemas abiertos

De acuerdo con los problemas estudiados, esta tesis se puede dividir en dos partes: La primera trata el problema de buscar triangulaciones eficientes del hipercubo y se lleva a cabo en el Capítulo 2. El resto de esta tesis, los Capítulos 3 a 5, está dedicada al estudio de grafos planos, pseudo-triangulaciones y sus relaciones con rigidez. Puesto que las técnicas utilizadas son diferentes, ambas partes pueden verse separadamente, a pesar de que las dos se enmarcan en el contexto de la Combinatoria Geométrica.

1.5.1 Triangulaciones eficientes de cubos

Una vez conocida una triangulación de eficiencia ρ de un cierto cubo, la construcción de Haiman da triangulaciones de tamaño $\rho^n n!$ de cubos de dimensión alta. Esto traslada el problema de decidir cómo de eficientes pueden ser las triangulaciones de cubos al estudio de triangulaciones de cubos de dimensiones pequeñas. Sin embargo, incluso este estudio ha resultado ser un problema difícil: hasta la fecha sólo se conocen triangulaciones de tamaño mínimo ϕ_d, y por tanto de eficiencia mínima ρ_d, hasta dimensión $d = 7$. Y no se han encontrado valores mejores prescindiendo de la condición de mínimas. En particular, el mejor tamaño asintótico de triangulaciones de cubos hasta el desarrollo de esta tesis era $0.840^n n!$, obtenido de ρ_7.

En el Capítulo 2 (colaboración con Francisco Santos [54]) presentamos una construcción para triangular productos de dos politopos generales. Como un caso particular, aplicamos nuestra construcción a productos de cubos y obtenemos triangulaciones de cubos de alta dimensión cuya eficiencia no es la de una triangulación de un cubo de dimensión pequeña, sino la "eficiencia" de una triangulación de un producto de un cubo y un símplice, ambos de dimensiones pequeñas.

En primer lugar, tenemos que definir la noción apropiada de eficiencia para una triangulación de un producto de este tipo, a la que llamamos *eficiencia ponderada* (Definición 2.2.1). El resultado principal del capítulo es:

(Teorema 2.2.3) Si hay una triangulación de $I^l \times \Delta^{m-1}$ con eficiencia ponderada ρ_0, entonces

$$\lim_{n\to\infty} \rho_n \leq \rho_0.$$

Como antes, esto nos conduce al estudio de la menor eficiencia ponderada de triangulaciones de $I^l \times \Delta^{m-1}$, que denotamos por $\rho_{l,m}$. Este número se puede calcular resolviendo un problema de programación lineal; de este modo hemos obtenido varios valores, de los cuales el menor es $\rho_{3,3} = 0.8159$. De aquí concluimos que:

- Para n suficientemente grande, el cubo I^n se puede triangular con $0.8159^n n!$ símplices.

El objetivo natural de un futuro trabajo en esta parte de la tesis debería ser el cálculo de mejores valores de $\rho_{l,m}$. En lo concerniente a este punto, tenemos que decir que no somos demasiado entusiastas sobre el uso de un método "de fuerza bruta", al menos por el momento: hemos tratado de obtener $\rho_{4,2}$ resolviendo el problema de programación lineal en un ordenador bastante potente de la Universidad de California en Davis sin ningún resultado después de semanas de cálculos. En su lugar, una aproximación interesante podría ser tratar de adaptar los métodos utilizados por Anderson y Hughes [6] para el cálculo de ϕ_6 y ϕ_7, que involucran maneras *ad hoc* de descomponer el sistema en subsistemas más pequeños. Pero están llenos de detalles técnicos y no estamos seguros sobre si mejorarían significativamente nuestros valores, puesto que la eficiencia resulta ser un parámetro muy poco sensible.

1.5.2 Grafos sin cruces, rigidez y pseudo-triangulaciones

La segunda parte de la tesis comienza en el Capítulo 3 (colaboración con Francisco Santos [53]). En él, consideramos un conjunto de puntos en el plano $\mathcal{A}$ y en primer lugar estudiamos sus pseudo-triangulaciones. Denotemos por n_i y n_v el número de puntos interiores y de vértices de $\text{conv}(\mathcal{A})$ respectivamente. Nótese que en posición no general pueden aparecer otros puntos en la frontera de $\text{conv}(\mathcal{A})$.

Con anterioridad al trabajo presentado en esta tesis, las pseudo-triangulaciones se habían considerado sólo para conjuntos de puntos en posición general. Había una noción de *flip* bien conocida para pseudo-triangulaciones puntiagudas, tal que el *grafo de pseudo-triangulaciones puntiagudas* que tiene a éstas por vértices y a los flips por aristas, es conexo y regular de grado $2n_i + n_v - 3$. Rote, Santos y Streinu [61] han demostrado que este grafo es el 1-esqueleto de un politopo simple $X_f(\mathcal{A})$ de dimensión $2n_i + n_v - 3$.

Nuestra primera observación fue que, todavía bajo la asunción de posición general, se puede extender la noción de flip a pseudo-triangulaciones generales, no necesariamente puntiagudas (Definición 3.2.6). Para el *grafo de pseudo-triangulaciones* resultante obtenemos:

(Proposición 3.2.7) El grafo de pseudo-triangulaciones de $\mathcal{A}$ es conexo y regular de grado $3n_i + n_v - 3 = 3n - 2n_v - 3$.

Como sugieren este resultado y el precedente anterior, damos una respuesta positiva a la pregunta de si este grafo es el 1-esqueleto de un politopo simple de dimensión $3n_i + n_v - 3$. Nuestros resultados principales son:

(Teorema 3.1.2) Sea $\mathcal{A}$ un conjunto finito de n puntos en el plano, no todos ellos contenidos en una recta. Entonces existe un politopo simple $Y_f(\mathcal{A})$ de dimensión $2n_i + n - 3$ tal que su conjunto de vértices está en biyección con el conjunto de todas las pseudo-triangulaciones de $\mathcal{A}$. El 1-esqueleto de $Y_f(\mathcal{A})$ es el grafo de flips entre ellas.

(Teorema 3.1.1) Hay una cara F de $Y_f(\mathcal{A})$, de dimensión $2n_i + n_v - 3$, tal que el complementario de la estrella de F en el poset de caras de $Y_f(\mathcal{A})$ es el opuesto del poset de grafos sin cruces en $\mathcal{A}$ que utilizan todas las aristas de la envolvente convexa.

En particular, es trivial obtener de éste el poset de *todos* los grafos sin cruces en $\mathcal{A}$, puesto que las aristas de la envolvente convexa no afectan a la planaridad. Además, la cara F coincide con el politopo $X_f(\mathcal{A})$ de [61].

El Teorema 3.1.2 da una respuesta a una cuestión abierta planteada en [69]. Además, probamos que nuestros resultados y la construcción del politopo $Y_f(\mathcal{A})$ funcionan incluso si $\mathcal{A}$ está en posición no general (Sección 3.5). Esto nos lleva a definir también grafos sin cruces, vértices puntiagudos, pseudo-triángulos, pseudo-triangulaciónes y flips en presencia de colinearidades (Definiciones 3.5.1, 3.5.2 y 3.5.5).

A pesar de que es en los Capítulos 4 y 5 en los que la rigidez viene a un primer plano, también juega un papel en las técnicas utilizadas para construir $Y_f(\mathcal{A})$, y en una interesante consecuencia de ellas: Streinu [69] demostró que las pseudo-triangulaciones puntiagudas son grafos minimalmente rígidos. Yendo más allá, por un lado probamos:

(Teorema 3.1.3) Las pseudo-triangulaciones son grafos rígidos.

Por otro lado, dedicamos el Capítulo 4 (colaboración con Günter Rote y Francisco Santos) a demostrar el inverso del resultado de Streinu:

(Teorema 4.1.1) Todo grafo plano minimalmente rígido admite una inmersión en $\mathbb{R}^2$ como pseudo-triangulación puntiaguda.

Por tanto, damos una caracterización combinatoria simple y elegante de todos los grafos que admiten inmersiones geométricas como pseudo-triangulaciones puntiagudas, dando respuesta a un problema abierto propuesto en [68]:

(Corolario 4.1.2) Dado un grafo G, las siguientes condiciones son equivalentes:

(i) G es plano y minimalmente rígido (isostático),

(ii) G es un grafo plano de Laman,

(iii) G admite una inmersión geométrica como pseudo-triangulación puntiaguda.

El problema abierto más famoso en Teoría de Rigidez (la *Conjetura de Rigidez* [38]) es encontrar el análogo 3-dimensional de los grafos planos de Laman, que hemos caracterizado como pseudo-triangulaciones puntiagudas. Sin embargo, encontrar una definición apropiada de pseudo-triangulaciones en 3 dimensiones no es

tarea fácil. Sería interesante generalizar la construcción anterior de $Y_f(\mathcal{A})$ utilizando rigidez 3-dimensional, para obtener un politopo cuyos vértices pudieran tomarse como las pseudo-triangulaciones de un conjunto $\mathcal{A}$ de puntos en $\mathbb{R}^3$.

También es natural preguntarse sobre la relación entre grafos rígidos en general (no necesariamente isostáticos) y pseudo-triangulaciones generales (no necesariamente puntiagudas). En particular, si alguna o todas las equivalencias en la caracterización anterior se pueden extender. Para esto tenemos una respuesta parcial: introducimos la noción de *pseudo-triangulación combinatoria (p.t.c.)*, definida como un grafo junto con un etiquetado de sus ángulos que induce la estructura combinatoria de una pseudo-triangulación (Definición 4.2.1). Además para estas p.t.c. definimos una *propiedad de Laman generalizada* tal que conjeturamos:

(Conjetura 4.5.3) Dado un grafo G, las siguientes condiciones son equivalentes:

(i) G es plano y rígido,

(ii) G tiene un etiquetado como p.t.c. con la propiedad de Laman generalizada,

(iii) G admite una inmersión geométrica como pseudo-triangulación.

Nuestra respuesta parcial incluye (iii)⇒(i) y (iii)⇒(ii), que se demuestran en el Teorema 3.1.3 anteriormente mencionado. Por otro lado, probamos (ii)⇒(iii) (Teorema 4.2.4). Sin embargo, dejamos abierto el demostrar (i)⇒(ii) o (i)⇒(iii), cualquiera de las cuales daría una caracterización combinatoria completa de los grafos planos rígidos.

Finalmente, en el Capítulo 5 (colaboración con varios de los co-autores de [39]) trasladamos nuestra atención a las condiciones para que un armazón esférico sin cruces $G(S)$ tenga un recíproco sin cruces. Es bien conocido que aquéllos con una tensión no nula en cada arista tienen un recíproco $G'(S')$, pero no se conocía nada sobre cuándo éste es también sin cruces. Damos una descripción completa de las condiciones en términos de la forma de las caras de $G(S)$ y los patrones de signos de la tensión a lo largo de sus aristas del borde.

(Teoremas 5.2.1 y 5.2.3) Sea G la inmersión geométrica de un armazón esférico sin cruces, con una tensión no nula en todas las aristas. Entonces el recíproco de G producido por esa tensión es también sin cruces si, y sólo si, se cumple:

1. El complementario de la cara exterior de G es convexo y todas sus aristas tienen el mismo signo.

2. Todas las caras interiores de G son o bien pseudo-triángulos o bien pseudo-cuadriláteros, con uno de los siguientes patrones de signos en sus aristas:

   2.*a* Un pseudo-triángulo en el que todas las aristas de un mismo pseudo-lado tienen el mismo signo y un pseudo-lado tiene signo opuesto a los otros dos. En otras palabras, hay dos cambios de signo a lo largo de la frontera del pseudo-triángulo y ambos ocurren en esquinas.

2.*b* Un pseudo-triángulo con cuatro cambios de signo a lo largo de su frontera, tres de los cuales ocurren en las tres esquinas.

2.*c* Un pseudo-cuadrilátero en el que todas las aristas de un mismo pseudo-lado tienen el mismo signo y pseudo-lados consecutivos tienen signos opuestos.

Además, en el grafo recíproco estas cuatro situaciones se corresponden respectivamente con:

1′. Un vértice no puntiagudo para el que todas las aristas incidentes tienen el mismo signo.

2.*a*′ Un vértice puntiagudo con dos cambios de signo a su alrededor, ninguno de los cuales ocurre en el ángulo grande.

2.*b*′ Un vértice puntiagudo con cuatro cambios de signo a su alrededor, uno de los cuales ocurre en el ángulo grande.

2.*c*′ Un vértice no puntiagudo con cuatro cambios de signo a su alrededor.

Para el caso particular en que sólo aparecen pseudo-triángulos como caras, y si la tensión es única, G es una pseudo-triangulación con un único vértice no puntiagudo. Demostramos que su recíproco es de nuevo una pseudo-triangulación del mismo tipo:

(Teorema 5.1.3) El recíproco de una pseudo-triangulación que sea circuito para la rigidez es de nuevo una pseudo-triangulación y circuito para la rigidez.

Chapter 1

Introduction and Preliminaries

In this chapter we intend to provide the background for an easier understanding of the work presented in the subsequent chapters. Although these have been written with the aim to be as self-contained as possible, there is a number of notions and results which the reader may or may not be familiar with. We start from the basics so that people from various fields can get introduced. In addition, the examples and constructions have been selected so that they will later play a role in the obtention and discussion of the results.

Most of the statements in this introductory chapter come without proof. Nevertheless, we always provide references where they can be found and where an interested reader could look for additional information. We do include those proofs which will become crucial in the rest of the chapters, either because of the technique applied or because we will make use of them later.

The structure of the chapter is as follows: In the first section, we introduce some Polytope Theory, to focus then on triangulations of cubes. The second section is devoted to Graph Theory, starting from a summary of basic vocabulary and results. A couple of connections with polytopes are shown before giving a partial answer to the embeddability problem for graphs which will become a key result in Chapter 4. Then, we turn our interest to rigidity of graphs embedded in the plane and state a series of notions and results ending with the definition of rigidity we are interested in. The fourth section is devoted to present pseudo-triangulations both for convex bodies and for points. In the last section, we introduce the results we obtain in this thesis, as well as some problems which remain open and could become object of a further work.

1.1 Triangulations and subdivisions of polytopes

The historical roots of the study of polytopes reach back to the Greek mathematicians. Actually, polygons and polyhedra appear to have been among the very earliest objects of systematic mathematical investigation, with the first remarkable result being the enumeration of the famous Platonic Solids, which are the solution

to the problem of finding all *regular* 3-dimensional convex bodies; those for which faces are copies of the same convex polygon and the number of faces meeting at each vertex is the same. This can be considered the starting point of the study of connections between geometry and combinatorics which, since then until nowadays, passing through *Book XII* of Euclid's *Elements*, has been one of the most remarkable features of polytopes.

1.1.1 Preliminaries about Polytope Theory

Among basic objects in geometry are *points*, *lines*, *planes* or *hyperplanes*, affine subspaces of dimensions $0, 1, 2$ and $d-1$, respectively, in an affine space $\mathbb{R}^d$. Let us focus our attention on finite sets of labeled points spanning $\mathbb{R}^d$. The point sets we will consider throughout this thesis are usually in *general position*, meaning that every subset of $n+1$ points spans $\mathbb{R}^n$. But the definitions and results in this subsection hold as well for the case in which *degeneracies* appear.

As usual, we can regard elements in $\mathbb{R}^d$ either as points or vectors. That is to say, we can consider $\mathbb{R}^d$ as an affine or a vector space. The following definition states, from both points of view, what is probably the crucial notion in Polytope Theory:

Definition 1.1.1

- A point set $K \subseteq \mathbb{R}^d$ is *convex* if for any pair of points $x, y \in K$, the straight line segment $[x, y] = \{\lambda x + (1-\lambda)y : 0 \leq \lambda \leq 1\}$ is contained in K.
- A nonempty set of vectors $Y \subseteq \mathbb{R}^d$ is a *cone* if for any finite subset of vectors $\{y_1, \ldots, y_k\} \subseteq Y$ all their linear combinations with positive coefficients are contained in Y. By convention, every cone contains the zero vector (obtained as combination of the empty set).

It is clear that every intersection of convex sets is convex as well, so the *convex hull* of K, defined as the intersection of all convex sets in $\mathbb{R}^d$ containing K, is the smallest convex set containing K:

$$\operatorname{conv}(K) := \bigcap \{K' \subseteq \mathbb{R}^d : K' \supseteq K, K' \text{ convex}\}$$

For the case of a finite set of points $K = \{x_1, \ldots, x_n\} \subseteq \mathbb{R}^d$, the convex hull coincides with the set of all their *convex combinations*

$$\operatorname{conv}(K) = \{\lambda_1 x_1 + \cdots + \lambda_n x_n : \lambda_i \geq 0, \sum_{i=1}^{n} \lambda_i = 1\}$$

Analogously, for a finite set of vectors $Y = \{v_1, \ldots, v_n\} \subseteq \mathbb{R}^d$, the *conical hull* or *positive span* of Y, defined as the intersection of all cones in $\mathbb{R}^d$ containing Y, can be expressed as:

$$\operatorname{cone}(Y) = \{\mu_1 y_1 + \cdots + \mu_n y_n : \mu_i \geq 0\}$$

Before being ready to introduce two different deffinitions of the main object of study in this section, *convex polyhedra* and *convex polytopes*, we need to define the *Minkowski sum* of two sets $P, Q \subseteq \mathbb{R}^d$ to be

$$P + Q := \{p + q : p \in P, q \in Q\}$$

Throughout this thesis we are not considering non-convex polyhedra or polytopes. That is why in the sequel we will drop the word "convex" when referring to them:

Definition 1.1.2 In an affine space $\mathbb{R}^d$:

- A *$\mathcal{V}$-polyhedron* is the Minkowski sum of the following two sets; the convex hull of a finite set of points, and the conical hull of a finite set of vectors:

$$\text{conv}(K) + \text{cone}(Y)$$

- An *$\mathcal{H}$-polyhedron* is an intersection of finitely many closed halfspaces.

- A *$\mathcal{V}$-polytope* is a bounded $\mathcal{V}$-polyhedron, that is, the convex hull of a finite set of points, $\text{conv}(K)$.

- An *$\mathcal{H}$-polytope* is a bounded $\mathcal{H}$-polyhedron, that is, an intersection of finitely many closed halfspaces containing no *ray* $\{x + \mu y : \mu \geq 0\}$ for any $y \neq 0$.

The equivalence of the $\mathcal{V}$-definition and the $\mathcal{H}$-definition seems to be "geometrically clear". However, it is not a trivial fact to prove. We refer the reader to [84] for a proof of the next statement, from which the equivalence of $\mathcal{V}$ and $\mathcal{H}$-definitions for polytopes follows as a consequence. Let us just mention that the heart of the proof is Farkas' Lemma or, in other words, linear programming duality.

Theorem 1.1.3 (Main Theorem for polyhedra) *A subset $P \subseteq \mathbb{R}^d$ is a $\mathcal{V}$-polyhedron if, and only if, it is an $\mathcal{H}$-polyhedron.*

The *dimension* $\dim(P)$ of a polytope or polyhedron is that of its *affine hull*, the intersection of all affine subspaces containing P. A *d-polytope* is a polytope of dimension d in some $\mathbb{R}^e$ ($e \geq d$). The first example of d-polytope we are going to consider is, in the combinatorial sense, the "smallest" one:

Example 1.1.4 *A* d-simplex τ *is the convex hull of any* $d+1$ *affinely independent points in some* $\mathbb{R}^e$*. For example, triangles and tetrahedra are the* 2*-simplices and* 3*-simplices, respectively. It is easy to see that all* d*-simplices are affinely isomorphic to the* standard d-simplex Δ^d *defined as*

$$\Delta^d := \textit{conv}\{e_1, \ldots, e_{d+1}\}, \quad \textit{where } e_i \textit{ is the } i\textit{-th unit vector in } \mathbb{R}^{d+1}.$$

It is intuitively clear what the faces of a polytope are; their mathematical characterization is based on linear functionals:

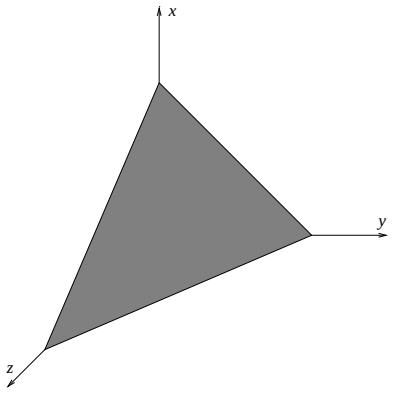

Figure 1.1: The standard 2-simplex Δ^2.

Definition 1.1.5 Let $P \subseteq \mathbb{R}^d$ be a convex polytope. A *face* of P is a set of the form

$$F = P \cap \{x \in \mathbb{R}^d : \langle c, x \rangle = c_0\},$$

where $c \in \mathbb{R}^d, c_0 \in \mathbb{R}$ and the inequality $\langle c, x \rangle \leq c_0$ is satisfied by all points $x \in P$. The *dimension* of a face is that of its affine hull.

Observe that P itself and $\emptyset$ are faces of the polytope P, defined respectively by the inequalities $\langle 0, x \rangle \leq 0$ and $\langle 0, x \rangle \leq 1$. They are called the *improper* and the *trivial* face. The *vertices*, *edges* and *facets* of P are respectively the faces of dimensions $0, 1$ and $\dim(P) - 1$; in particular, the vertices of P, denoted $\mathrm{vert}(P)$, are the minimal non-empty faces, and the facets are the maximal proper faces. The *k-skeleton* of a polytope is the union of its k-dimensional faces. The *star* of a face F, denoted $\mathrm{star}(F)$, is defined to be the set of faces contained in the union of all the facets containing F.

The following two special types of polytopes introduce us to the study of the combinatorics of these geometric structures:

Definition 1.1.6 Given a d-dimensional polytope P,

- P is said to be *simple* if every vertex is exactly in d facets.
- P is called *simplicial* if all its facets are simplices.

The combinatorial information about a polytope P will be encoded by its face structure. In order to study this, we first need some notions about *partially ordered sets*, or "posets" for short. A reader interested in the field could check [16, Chapter 11] for a detailed survey.

Definition 1.1.7 A *poset* $(S, \leq)$ is a finite set S together with a reflexive, antisymmetric and transitive relation "$\leq$".

- A poset is *bounded* if it has a unique minimal and a unique maximal elements.
- A *Boolean* poset is one isomorphic to the poset $B_k = (2^{[k]}, \subseteq)$ of all subsets of a k-element set.
- A *chain* in a poset is a totally ordered subset; the *length* of a chain is its number of elements minus 1.
- A poset S is *graded* if it is bounded and every maximal chain has the same length, which is said to be the *length* or *rank* of S.
- A *lattice* is a bounded poset such that every two elements $x, y \in S$ have a unique minimal upper bound and a unique maximal lower bound in S.
- The *lower ideal* of an element x in $(S, \leq)$ is the down-set $\{y \in S : y \leq x\}$. The *upper filter* is $\{y \in S : y \geq x\}$.

The following well-known result describes the poset structure of the faces of a polytope, to which one usually refers as the *face poset*:

Theorem 1.1.8 *For every polytope P, the poset of all its faces partially ordered by inclusion is a graded lattice of length* $\dim(P) + 1$. *Hence, it is also called the* face lattice *of P and denoted $L(P)$.*

Example 1.1.9 *In Figure 1.2 we show the face poset of a convex pentagon (left), which is a lattice of length* 3, *and that of a* 3*-simplex (right), which is a Boolean lattice of length* 4.

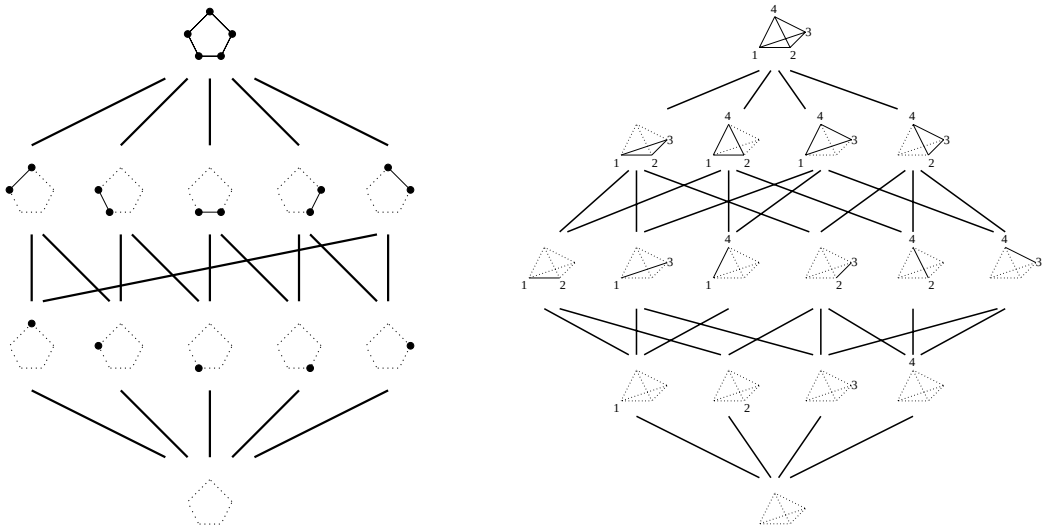

Figure 1.2: Face lattices of a convex pentagon and a 3-simplex.

As advanced above, the face lattice is the suitable notion to define combinatorial equivalence of polytopes: Two polytopes P and Q are said to be *combinatorially equivalent* if their face lattices are isomorphic, $L(P) \cong L(Q)$.

1.1.2 Triangulations of cubes

When one has a complicated geometric object (for example, a "big" polytope), it is sometimes convenient to decompose it into smaller and simple parts to simplify its study:

Definition 1.1.10 A *polytopal complex* $\mathcal{C}$ is defined to be a finite collection of polytopes in $\mathbb{R}^d$ such that:

(i) The empty polytope is in $\mathcal{C}$,

(ii) If $P \in \mathcal{C}$, then all the faces of P are also in $\mathcal{C}$, and

(iii) The intersection $P \cap Q$ of two polytopes $P, Q \in \mathcal{C}$ is a face of both of them.

The *dimension* $\dim(\mathcal{C})$ is the largest dimension of a polytope in $\mathcal{C}$. The *underlying set* of $\mathcal{C}$ is the point set $|\mathcal{C}| := \cup_{P \in \mathcal{C}} P$.

As above, the combinatorial structure of a polytopal complex $\mathcal{C}$ is captured by its *face poset* $(C, \subseteq)$ given by the set of polytopes in $\mathcal{C}$ ordered by inclusion. Here come the definitions giving name to this section:

Definition 1.1.11 Given a d-polytope P:

- A *polytopal subdivision*, or just *subdivision*, of P is a polytopal complex $\mathcal{C}$ such that its underlying set is $|\mathcal{C}| = P$.
- A *triangulation* of P is a subdivision in which all the polytopes in $\mathcal{C}$ are simplices.
- Given a subdivision $\mathcal{C}$ of P, a *refinement* is another subdivision $\mathcal{C}'$ such that every polytope of $\mathcal{C}$ is a union of polytopes of $\mathcal{C}'$.

A decomposition of P into non-overlapping d-simplices that do not fulfill condition (iii) of Definition 1.1.10 is called a *dissection* or, in case all the polytopes are simplices, a *simplicial dissection*. In particular, throughout this thesis we are interested in subdivisions, triangulations and dissections whose set of vertices coincides with that of P, that is, adding no new vertices and such that every vertex is used.

Remark 1.1.12 The same definition holds for *subdivisions*, *triangulations*, *refinements* and *dissections* of a finite point set $\mathcal{A}$, considering the polytope defined by its convex hull, except we implicitly are assuming that only points of $\mathcal{A}$ can be used as vertices. Furthermore, in general we will suppose that all of them are used.

Example 1.1.13 *Figure 1.3 depicts a triangulation (left) and a simplicial dissection (right) of the* 3*-dimensional cube. In the first case, the triangulation is composed of* 5 *simplices: the four tetrahedra obtained by cutting non-adjacent corners and the resultant tetrahedron, which uses opposite diagonals of the top and bottom faces. In the second case, the cube is split into two triangular prisms for a better visualization: each prism is triangulated into* 3 *tetrahedra, but the triangulations induced in the common face of the two prisms use opposite diagonals. Therefore those* 6 *simplices do not form a triangulation, since condition (iii) of polytopal complexes is violated.*

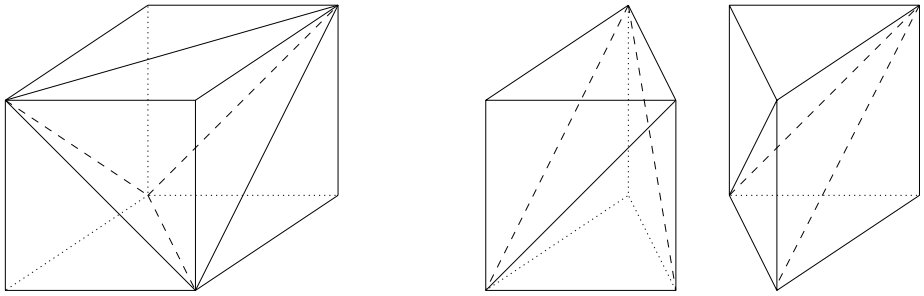

Figure 1.3: Left: A triangulation of the 3-cube. Right: A dissection of the 3-cube.

Let us point up that every subdivision of a polytope P can be refined to a triangulation. In order to do so, triangulations T_i of each of the polytopes P_i in the subdivision are required to be such that for every pair of indices i, j, T_i and T_j induce the same triangulation on $T_i \cap T_j$ when this is considered as a face of both T_i and T_j. There are well-known constructions (see for instance [76]) achieving this.

The study of triangulations and their properties has attracted a great interest during the last decades, specially stimulated by the advances on computations, which have led to a number of applications in Computational Geometry, among which can be mentioned those to computer graphics or simulations for engineering. In particular, triangulations of cubes have been object of study for more than thirty years and are treated in the second chapter of this thesis, where we construct "small" triangulations of cubes.

Let us define the *product* of two general polytopes as $P \times Q := \{(p, q) : p \in P, q \in Q\}$. Before considering triangulations of cubes, we are going to focus on a way to triangulate products $\Delta^k \times \Delta^l$ of two simplices. The triangulation we show here for such a product is interesting on its own, due to its simplicity and easy representation. In addition, this interest will be fully justified in Chapter 2 where we generalize it to the product of a polytope and a simplex. It is a classical result that all triangulations of a product $\Delta^k \times \Delta^l$ have the same number of simplices:

Lemma 1.1.14 ([13]) *Every triangulation of* $\Delta^k \times \Delta^l$ *has exactly* $\binom{k+l}{k}$ *simplices.*

Our aim here is to provide the reader with a remarkable triangulation of $\Delta^k \times \Delta^l$, known and used in algebraic topology [30] for a long time, and whose easy visualization makes it a good example of how helpful combinatorics can be in order to understand geometric structures in high dimensions. This utility is at the background of the whole thesis, and may be even more explicit in Chapters 2 and 3.

Assume Δ^k and Δ^l having vertices $\{v_0, \ldots, v_k\}$ and $\{w_0, \ldots, w_l\}$ respectively. Then the product $\Delta^k \times \Delta^l$ has vertex set

$$\{(v_i, w_j) : 0 \leq i \leq k, 0 \leq j \leq l\},$$

which can be represented as the squares of a $(k+1) \times (l+1)$ grid whose rows correspond to vertices of Δ^k and whose columns represent vertices of Δ^l. See Example 1.1.15.

Now, consider paths from the vertex (v_0, w_0) to the vertex (v_k, w_l) whose steps increase by one either the index of v or that of w. In the representation, if we label rows and columns of the grid in such a way that (v_0, w_0) and (v_k, w_l) are the lower-left and upper-right corner squares respectively, those paths correspond to monotone staircases, in which each step is a movement from a square either towards the next to the right or to the one inmediately above it. Each such path selects a subset of $k + l + 1$ vertices of $\Delta^k \times \Delta^l$ and it is not difficult to see that determines a $(k+l)$-dimensional simplex.

The collection of simplices given by all those paths, represented by the collection of all monotone staircases in the grid, forms a triangulation of $\Delta^k \times \Delta^l$, which is known as the *staircase triangulation*. Note that the number of possible such paths or staircases is precisely $\binom{k+l}{k}$.

Example 1.1.15 *The next figure shows the staircase triangulation of the triangular prism $\Delta^1 \times \Delta^2$. In the top row the representation of its vertex set $\{(v_i, w_j) : 0 \leq i \leq 1, 0 \leq j \leq 2\}$ as the squares of a 2×3 grid is depicted. Below we show the three possible staircases in this grid and the corresponding tetrahedra, which triangulate $\Delta^1 \times \Delta^2$.*

We turn now our attention to the family of *regular d-cubes* $I^d := [0,1]^d, d \geq 0$: "Simple" triangulations of the regular d-cube have several applications, such as solving differential equations by finite element methods or calculating fixed points (see, for example, [72]). In particular, it has brought special attention both from a theoretical point of view and from an applied one to determine the smallest size of a triangulation of I^d (see [51, Section 14.5.2] for a recent survey). Let us point up that the general problem of computing the smallest triangulation of an arbitrary polytope is NP-complete even when restricted to dimension 3, see [10].

Definition 1.1.16 The *size* of a triangulation (or simplicial dissection) T of the d-cube is its number of d-simplices $|T|$. The smallest size of triangulations of I^d is usually denoted ϕ_d.

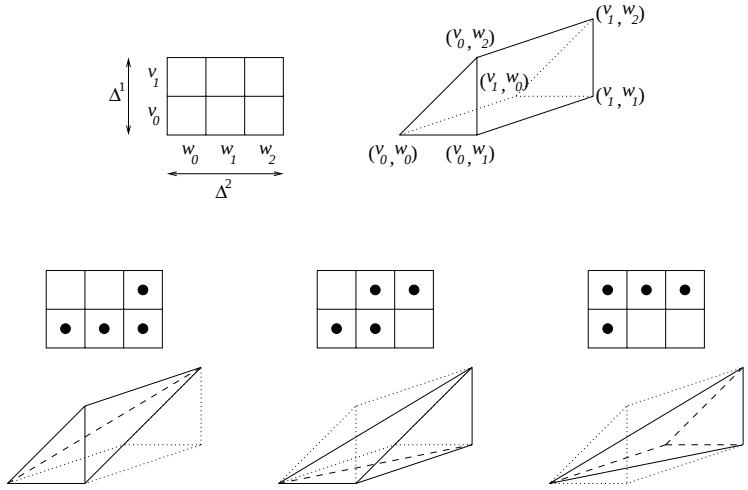

Figure 1.4: The staircase triangulation of $\Delta^1 \times \Delta^2$.

Definition 1.1.17 Given a d-simplex $\tau = \text{conv}(p_1, \ldots, p_{d+1})$, its *volume* is defined to be

$$\text{vol}(\tau) := \text{abs}\left(\begin{vmatrix} p_1 & \cdots & p_{d+1} \\ 1 & \ldots & 1 \end{vmatrix}\right)/d!$$

The minimum volume of a d-simplex using the vertices of I^d is therefore $1/d!$, and since that of the d-cube equals 1, the maximum size of a triangulation of I^d turns out to be $d!$ A number of easily constructed triangulations achieving the maximum size can be found in the above mentioned survey.

If we are given two triangulations of a d-cube for a fixed d, it is easy to decide which of them is "better" (or "simpler"); the one with smallest size. However, in case we have triangulations of d-cubes for different dimensions d, we need a notion to discern which of them is the most convenient. This was first stated by Todd in [72]:

Definition 1.1.18 The *efficiency* of a triangulation T of a d-cube is defined to be the number $(|T|/d!)^{1/d}$. The smallest efficiency $(\phi_d/d!)^{1/d}$ of triangulations of I^d is denoted ρ_d.

Remark 1.1.19 By the above observation about the maximum size of a triangulation of the d-cube, the efficiency of a triangulation of I^d is at most one and the smaller it is, the most efficient is the triangulation. In particular, $\rho_d \leq 1, \forall d$.

In the second chapter of this thesis we present a new construction to triangulate I^d leading to asymptotically good efficiencies. That construction is inspired in the following method due to Haiman [40]:

Let T_k and T_l be triangulations of the regular cubes I^k and I^l, respectively. The product $T_k \times T_l$ of the two triangulations (understood as the pairwise product of their simplices) gives a decomposition of the cube $I^k \times I^l = I^{k+l}$ into $|T_k| \cdot |T_l|$ subpolytopes, each of them isomorphic to the product of simplices $\Delta^k \times \Delta^l$. By Lemma 1.1.14, every triangulation of $\Delta^k \times \Delta^l$ has size $\binom{k+l}{k}$. Hence, refining the subdivision $T_k \times T_l$ in an arbitrary way one gets a triangulation of I^{k+l} of size $|T_k| \cdot |T_l| \cdot \binom{k+l}{k}$, leading to:

Theorem 1.1.20 (Haiman) *For every k and l, ${\rho_{k+l}}^{k+l} \leq {\rho_k}^k {\rho_l}^l$.*

The implication of this is that starting with a triangulation of a cube I^k of a certain efficiency ρ, one can construct a sequence of triangulations of I^{nk} for $n \in \mathbb{N}$ and with exactly that same efficiency. This inmediately suggests the following result:

Corollary 1.1.21 *The sequence $(\rho_n)_{n\in\mathbb{N}}$ converges and*

$$\lim_{n\to\infty} \rho_n \leq \rho_d \qquad \forall d \in \mathbb{N}.$$

Proof: Let us fix $d \in \mathbb{N}$ and $k \in \{1, \ldots, d\}$. Haiman's theorem implies that, for every $i \in \mathbb{N}$,

$$\rho_{k+id} \leq {\rho_k}^{k/(k+id)} {\rho_d}^{id/(k+id)}.$$

Since the right-hand side converges to ρ_d when i grows, the d subsequences of $(\rho_n)_{n\in\mathbb{N}}$ consisting of indices equal modulo d, and hence the whole sequence $(\rho_n)_{n\in\mathbb{N}}$, have upper limit bounded by ρ_d. This holds for every $d \in \mathbb{N}$. An upper limit bounded by every term in the sequence must coincide with the lower limit. □

We denote ρ_∞ and call *asymptotic efficiency* of triangulations of the cube the limit of $(\rho_n)_{n\in\mathbb{N}}$. The corollary says that the asymptotic efficiency of triangulations of the regular cube is at most the smallest effiency ρ_d for every d. In particular, the best upper bound for ρ_∞ prior to our work was $\rho_\infty \leq \rho_7 = 0.840$. In Chapter 2 we prove $\rho_\infty \leq 0.8159$ with a completely novel method.

1.2 Graph Theory

Graph theorists agree to consider Euler's paper [31] (in which the famous problem on the bridges of Königsberg was discussed) to mark the birth of Graph Theory. A number of mathematicians have contributed since then to this field, but it has been during the last century when it has attracted a great interest, increased during the last decades due to the simplicity of representing a graph on a computer and the big number of problems that can be modeled using graphs.

In this section we first present some relations between graphs and polytopes; the first of them is the polytope of subdivisions of a convex polygon and is related to the results of Chapter 3. The others are interesting properties of graphs defined by vertices and edges of polytopes. Finally, we state a result by Tutte about embeddability of an abstract graph as a geometric one, with points of $\mathbb{R}^2$ as vertices and straight segments as edges.

1.2.1 Preliminaries and relations with polytopes

The reader will be familiar with the basic notions we list here. For a general reference on graphs we recommend [15] or [75], and for each statement in this section we specify where a proof can be found.

Definition 1.2.1 A *graph* is a pair of sets $G = (V, E)$ where $V = \{1, \dots, n\}$ is a finite set of *vertices* and $E \subseteq \{\{i, j\} : i, j \in V, i \neq j\}$ is a set of *edges* between the vertices. We will usually denote edges by ij instead of $\{i, j\}$.

By definition, the graphs we are considering are *simple*; they have neither loops (edges joining an element with itself) nor multiple edges. In the next definition we summarize some basic vocabulary. By $\binom{V}{2}$ we denote the set of all unordered pairs of vertices from V:

Definition 1.2.2 Let $G = (V, E)$ be a graph.

- The number of edges in E containing a vertex $v \in V$ is said to be the *degree* of that vertex.
- G is *regular of degree* k if every vertex has degree k.
- G is *complete* if $E = \binom{V}{2}$. The complete graph on n vertices is denoted K_n.
- The graph obtained from G by the *removal* of a vertex set $W \subseteq V$ is defined as the restriction of G to the vertex set $V \backslash W$, that is; $G \backslash W := (V \backslash W, E \cap \binom{V \backslash W}{2})$.
- The set of *neighbors* of a vertex v is defined as $N(v) := \{w \in V : \{v, w\} \in E\}$.
- A graph $G_0 = (V_0, E_0)$ is a *subgraph* of G, denoted $G_0 \subseteq G$, if $V_0 \subseteq V$ and $E_0 \subseteq E$.

We consider three distinguished types of subgraphs:

- A *path* in G between two vertices $v_1, v_k \in V$ (note that $1, k$ are subindices and not necessarily $v_i = i$) is a subgraph $G_{v_1}^{v_k} \subseteq G$ with vertex set $V_{v_1}^{v_k} = \{v_1, \dots, v_k\}$ and edge set $E_{v_1}^{v_k} = \{\{v_1, v_2\}, \{v_2, v_3\}, \dots, \{v_{k-1}, v_k\}\}$.
- A *k-cycle* in G is a subgraph $G_{v_1}^{v_1} \subseteq G$ with $V_{v_1}^{v_1} = \{v_1, \dots, v_k\}$ and $E_{v_1}^{v_1} = \{\{v_1, v_2\}, \{v_2, v_3\}, \dots, \{v_{k-1}, v_k\}, \{v_k, v_1\}\}$.
- A *Y-graph* in G of endpoints $v_1, v_2, v_3 \in V$ is a subgraph composed of a vertex $w \in V \setminus \{v_1, v_2, v_3\}$ and paths $G_{v_1}^w, G_{v_2}^w, G_{v_3}^w$ which are disjoint except for the common point w.

And a special type of graphs which will appear later:

- G is a *bipartite graph* if V is the disjoint union of two sets V_1, V_2 and each edge joins a vertex of V_1 to one of V_2. It is said to be *complete* if its edges are all the possible ones between V_1 and V_2. The complete bipartite graph is denoted $K_{m,n}$, where m and n are the cardinalities of the subsets.

Paths and cycles are often denoted by the sequence of appearance of vertices, that is, $(v_1, v_2, \ldots, v_k)$ or $(v_1, v_2, \ldots, v_k, v_1)$. Another notion which is of interest for us is that of connectivity:

- G is *connected* if for any pair of distinct vertices $v, w \in V$ there is a path G_v^w joining them.
- G is *k-connected* for $k > 1$ if for every $v \in V$ the graph $G \backslash v$ is $(k-1)$-connected.

Thus, an abstract graph $G = (V, E)$ is given as a set of vertices together with a set of unordered pairs of vertices, and no geometry is involved. On the other hand, a *topological embedding* $G \hookrightarrow \mathbb{S}^2$ is a mapping of the vertices V to points in the 2-sphere and of the edges $ij \in E$ to arcs. It is said to be a *plane* or *non-crossing* embedding if no pair of arcs $p_i p_j$ and $p_k p_l$ corresponding to non-adjacent edges $ij, kl \in E, i, j \notin \{k, l\}$ has a point in common.

A *geometric embedding*, or *geometric graph*, $G(S) = (V, E, S)$ is defined as a graph $G = (V, E)$ together with a mapping $i \mapsto p_i \in S$ of the vertices V to a finite set of points $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ in the euclidean plane and of the edges $ij \in E$ to straight line segments $p_i p_j$. Analogously, it is *plane*, or *non-crossing*, if no pair of segments corresponding to non-adjacent edges intersect. Fary's Theorem [33] states that every plane topological embedding can be "stretched" to a plane geometric embedding. On the other hand, every plane geometric embedding trivially induces a plane topological embedding (just projecting to $\mathbb{S}^2$). A graph is said to be *planar* if it has a plane topological or geometric embedding.

The *cells*, or *faces*, of a plane geometric or topological embedding are defined to be the connected components of the complement (in $\mathbb{R}^2$ or $\mathbb{S}^2$) of the union of edges and vertices. A geometric embedding has exactly one unbounded face, usually called *outer* face.

In this thesis we will consider only embeddings of graphs in $\mathbb{R}^2$ or $\mathbb{S}^2$. We will use the notations $G(S)$ and $G \hookrightarrow \mathbb{S}^2$ for geometric and topological embeddings, respectively. But we may abuse notation and use the image set S to refer to the geometric embedding $G(S)$. The next result collects some useful statements; their proofs can be found, for instance, in [77]:

Theorem 1.2.3 *Let G be a graph.*

(i) **(Menger's Theorem)** *G is k-connected if, and only if, between any pair of vertices there are k paths $G_1, \ldots, G_k$ in G which are disjoint except for the endpoints.*

(ii) **(Kuratowski's Theorem)** *G is planar if, and only if, it has no subgraph homeomorphic to the complete graph K_5 or to the complete bipartite graph $K_{3,3}$.*

(iii) **(Whitney's Theorem)** *If G is planar and 3-connected, then all its topological embeddings are equivalent (have the same cells).*

(iv) **(Euler's Theorem)** *If G is connected and planar, for every plane embedding $G(S)$ (or $G \hookrightarrow \mathbb{S}^2$) the numbers c of cells, v of vertices and e of edges are related by*

$$v - e + c = 2$$

Let us consider now the class of geometric *non-crossing* graphs on a given set of vertices $\mathcal{A}$ in the plane $\mathbb{R}^2$. This type of graphs is of interest in Computational Geometry, Geometric Combinatorics, and related areas. In particular, much effort has been directed towards enumeration, counting and optimization on the set of maximal such graphs, that is to say, triangulations. Note that here triangulations are considered with set of vertices equal to $\mathcal{A}$ as advanced in Remark 1.1.12. Since any triangulation of a point set of cardinality n has $n_i + 2n - 3$ edges for n_i the number of points in the interior of $\text{conv}(\mathcal{A})$, a lower bound of $2^{n_i+2n-3} = \Omega(4^n)$ non-crossing graphs on $\mathcal{A}$ is trivial. An upper bound of type $O(c^n)$ for some constant c was first shown in [5]. The best current value for c is $59 \cdot 8 = 472$ [62].

It is clear that the set of non-crossing geometric graphs on a planar point set forms a poset with respect to the inclusion of edges. However, until the development of this thesis their poset structure was only well understood if the points were in convex position. The paper [34] contains several enumerative results about geometric graphs with n vertices in convex position. In particular, it shows that there are $\Theta((6+4\sqrt{2})^n n^{-3/2})$ non-crossing graphs in total and gives explicit formulas for each fixed cardinality.

For points in convex position the non-crossing graphs containing all the convex hull edges are the same as the polygonal subdivisions of the convex n-gon. It is well-known that the poset they form is the opposite of the face poset (except for the empty face) of the $(n-3)$-*associahedron*, a simple polytope of dimension $n-3$ with the following properties:

- Its vertices correspond to all the $\frac{1}{n-1}\binom{2n-4}{n-1}$ different ways of bracketing a string of length $n-1$ or, equivalently, of multiplying an expression $a_1 a_2 \dots a_{n-1}$ when multiplication is not associative.

- Two vertices are adjacent if they correspond to a single application of the associative law.

Thus, vertices of the associahedron correspond to polygonal subdivisions with the maximal number of edges, i.e. triangulations.

Example 1.2.4 *Figure 1.5 depicts the* 2*-associahedron (left), whose vertices correspond to all the* 5 *ways of bracketing the string* 1234*. Hence it is a pentagon, whose face lattice we have already found in Figure 1.2. Neglecting the empty face, this is checked to coincide with the opposite of the poset of polygonal subdivisions of the convex* 5*-gon, which we show in the right picture.*

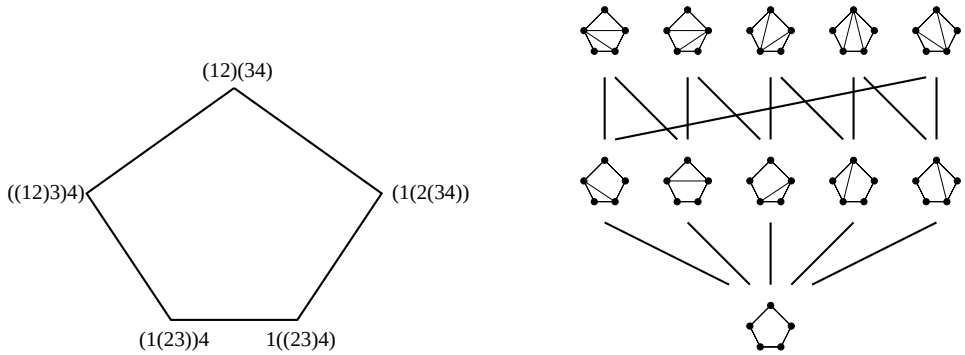

Figure 1.5: Left: The 2-associahedron. Right: Poset of polygonal subdivisions of the convex 5-gon.

The main result in the third chapter of this thesis is a generalization of the associahedron; we give a description of the poset of non-crossing graphs for any set of points, even in non-general position, as a subset of the face poset of a certain polytope we construct.

For a more straightforward relation between graphs and polytopes, observe that the vertices and edges of a polytope P (the 1-skeleton) form an abstract graph, which is called the *graph of* P. The natural reciprocal question about when an abstract graph is the graph of a polytope is solved for 3-polytopes by the following classical result, of which the original proof appears in [66] and a very understandable one can be found at [84]:

Theorem 1.2.5 (Steinitz) *G is the graph of a 3-dimensional polytope if, and only if, it is simple, planar, and 3-connected.*

The "if" part of the theorem is the most difficult to prove. For the "only if part", since polytopal graphs are trivially simple and its planarity can be checked using radial projection to a sphere from an interior point, or a linear projection to a plane from a point beyond a facet, it only remains to apply another well-known result:

Theorem 1.2.6 (Balinski [8]) *The graph of a d-polytope is d-connected.*

1.2.2 Plane embeddings of graphs: Tutte's Theorem

Last, but not least, we introduce the embeddability problem for graphs. The construction we present to finish this section gives an answer to the following question: Given a 3-connected planar graph G, how can we find a geometric embedding in which, in addition, cells are realized as convex polygons?

Graph Drawing is a field with a distinguished history, and plane embeddings have received substantial attention in the literature ([33], [73], [74], [28], [25], [63], [9]).

Also, extensions of graph embeddings from straight-line to pseudo-line segments have been recently considered (see e.g. [55]). It is natural to ask which such embeddings are stretchable, i.e. whether they can be realized with straight-line segments while maintaining some desired combinatorial substructure. Indeed, the primordial planar graph embedding result, the afore mentioned Fary's Theorem [33], is just an instance of answering such a question.

An explanation of the background of the following construction can be easily given by this physical fact: Assume that the edges of G are rubber bands. Take the vertices corresponding to one particular cell F_0 whose boundary has m vertices. Then, pin down these vertices in the plane $\mathbb{R}^2$ in such a way that F_0 is realized as a convex m-gon and that all interior rubber bands are under tension. Voilà; the figure obtained is a plane embedding of G in which all the remaining cells are realized as non-overlapping convex polygons.

In order to model rubber bands, the notion of *weight* $w_{ij} \in \mathbb{R}$ of an edge ij is introduced. Note that since we are considering non-directed edges, the symmetry condition $w_{ij} = w_{ji}$ has to be required.

Definition 1.2.7 Let $G = (V, E)$ be a graph and $w : E \to \mathbb{R}$ be an assignment of weights to the edges. Consider also an assignment of positions $i \mapsto p_i \in \mathbb{R}^2$ for the vertices of G. A vertex $i \in V$ is said to be in *equilibrium* if

$$\sum_{ij \in E} w_{ij}(p_i - p_j) = 0.$$

See Figure 1.6 for an example.

The following important result gives a partial answer to the following question: When can a graph be embedded with straight edges in the plane, in such a way that the cells are convex and do not overlap? (The latter is equivalent to "without crossings").

Theorem 1.2.8 (Tutte) *Let $G = (\{1, \dots n\}, E)$ be a 3-connected planar graph that has a cell $(k+1, \dots, n)$ for some $k < n$. Let $p_{k+1}, \dots, p_n$ be the vertices of a convex $(n-k)$-gon (in this order). Finally, let $w : E' \to \mathbb{R}^+$ be an assignment of positive weights to the internal edges. Then:*

(i) There are unique positions $p_1, \dots, p_k \in \mathbb{R}^2$ for the interior vertices in such a way that all of them are in equilibrium.

(ii) All cells of G are realized as non-overlapping convex polygons.

Although giving a whole proof of this result is out of the purposes of this introductory chapter, we outline in the rest of this section the steps of the proof given in [60], which we slightly modify in Chapter 4 to prove a directed version of Tutte's Theorem. We include here the notions, intermediate results and proofs to which we will refer in that chapter.

Proof of (i): Assume that the positions of the boundary points $p_{k+1}, \ldots, p_n$ are given. Without loss of generality, we may assume $p_n = (0,0)$. Consider the function

$$E(x_1, \ldots, x_k, y_1, \ldots, y_k) = \frac{1}{2} \sum_{ij \in E'} w_{ij}((x_i - x_j)^2 + (y_i - y_j)^2) = \frac{1}{2} \sum_{ij \in E'} w_{ij} \|p_i - p_j\|^2,$$

which is quadratic and non-negative everywhere. Assume that the point $z = (x_1, \ldots, x_k, y_1, \ldots, y_k)$ has at least one entry, say x_l, with large absolute value. This implies that p_l is far away from the boundary point $p_n = (0,0)$. Since the graph G is connected, there is a path joining p_l and p_n; for at least one edge ij in this path the distance $\|p_i - p_j\|$ has to be large.

Thus, for sufficiently large $\alpha > 0$, the condition $|z| > \alpha$ implies $E(z) > E(0)$. Being E quadratic, this implies that it is strictly convex (non-degenerate with a positive definite Hessian) and therefore takes its unique minimum on $\{z : |Z| < \alpha\}$. The assertion follows because the condition $\nabla E = 0$ for a critical point of E:

$$\frac{\partial E}{\partial x_l} = \sum_{ij \in E'} w_{ij}(x_i - x_j) = 0 \quad \text{and} \quad \frac{\partial E}{\partial y_l} = \sum_{ij \in E'} w_{ij}(y_i - y_j) = 0, \quad \forall l \in \{1, \ldots, k\}$$

is exactly the equilibrium condition for the interior vertices. □

Sketch of the proof of (ii): Again, assume that the positions for the boundary points $p_{k+1}, \ldots, p_n$ are given; by (i) the positions for the interior vertices are in turn uniquely determined. Such a point configuration is called an *equilibrium representation* of G.

Example 1.2.9 *In Figure 1.6 an equilibrium representation of a graph G, together with the weights of each interior edge, is depicted. A grid is provided with the aim that the reader can easily check the three interior vertices to be in equilibrium.*

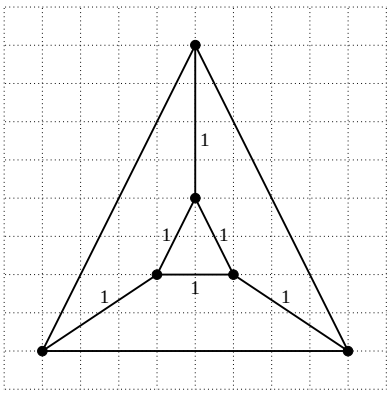

Figure 1.6: A graph in equilibrium.

The *relative interior* of (the convex hull of) a collection of points $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ is defined by

$$\text{relint}(S) := \{\sum_{i=1}^{n} \lambda_i p_i : \sum_{i=1}^{n} \lambda_i = 1 \text{ and } \lambda_i > 0,\ \forall i = 1, \ldots, n\}.$$

Next lemma, whose proof is straightforward using the equilibrium condition, gives a property which turns out to be fundamental in Tutte's Theorem. The statement can be checked again in Figure 1.6:

Lemma 1.2.10 *In an equilibrium representation $S := \{p_1, \ldots, p_n\} \in \mathbb{R}^{2n}$ of the vertices of G , every interior vertex p is in the relative interior of its neighbors; $p \in relint(N(p))$.*

Definition 1.2.11 Let G be a 3-connected, planar graph on n vertices such that those labelled $k+1, \ldots, n$ form, in this order, a cell of G. A point configuration $S := \{p_1, \ldots, p_n\} \in \mathbb{R}^{2n}$ is said to be a *good representation* for G if:

(i) $p_{k+1}, \ldots, p_n$ realize, in this order, a convex $(n-k)$-gon, and

(ii) For $i = 1, \ldots, k$ we have $p_i \in \text{relint}(N(p_i))$.

Lemma 1.2.10 states that equilibrium representations are good representations. Tutte's theorem then follows from this fact and the following assertion:

Theorem 1.2.12 ([60]) *A good representation of a 3-connected planar graph G is a plane geometric embedding of G in which all interior cells are realized as non-overlaping convex polygons.*

The proof is dissected into several smaller claims and has two different parts. In the first one, the following lemma proves that degenerate situations cannot occur, where a vertex p_i is said to be *degenerate* if the set of its neighbors $N(p_i)$ does not affinely span $\mathbb{R}^2$. Its proof can be skipped in a first reading and, as advanced above, appears here for a better understanding of the differences with the proof of the directed Tutte Theorem given in Chapter 4:

Lemma 1.2.13 ([60]) *A good representation S of a 3-connected planar graph G has no degenerate vertices.*

Proof: Before starting the proof, given a line $\ell = \{x : \phi(x) = \alpha\}$ for ϕ a linear functional, a vertex a is defined to be *ℓ-active* if not all its neighbors are also on ℓ.

It is straightforward that the boundary points are non-degenerate, since they are the vertices of an strictly convex polygon. Consider then a degenerate interior vertex $p \in S$, for which there must be a line $\ell = \{x : \phi(x) = \alpha\}$ such that $p \in \ell$ and $N(p) \subset \ell$.

Take a point $q \in S$ which is not on ℓ. Since G is 3-connected, Menger's Theorem 1.2.3.(i) guarantees the existence of three paths A, B, C from p to q having disjoint edges and vertices, except for the endpoints p, q. Since each of them has to leave the line ℓ at some point, they contain at least one ℓ-active vertex each. Denote by a the first ℓ-active point coming after p in the path A. The initial segment of this path is then of the form $A^0 = (p, a_1, \ldots, a_l, a)$ and all its points are on ℓ. Denote by b and c the first ℓ-active points coming after p respectively in B and C. Likewise, let B^0 and C^0 be the corresponding initial segments. Observe that, all together, the edges of A^0, B^0 and C^0 form a Y-graph with endpoints a, b, c, denoted Y^0, which is contained on ℓ.

The next step is to prove that there are also Y-graphs Y^+ and Y^- having the same endpoints a, b, c and whose edges are all respectively above and below the line ℓ. Only the existence of Y^+ is shown, since that of Y^- is proved analogously:

At most two of the boundary vertices can lie on ℓ, therefore among the points a, b, c there are at most two boundary points. Assume that a is interior. For any interior point $q \in S$ the definition of good representation implies $q \in \text{relint}(N(q))$, what implies in turn the following property:

$$\begin{gathered}\text{If there is a } q_- \in N(q) \text{ with } \phi(q) > \phi(q_-),\\ \text{then there is also a } q_+ \in N(q) \text{ with } \phi(q) < \phi(q_+)\end{gathered}$$

which holds as well for boundary points that are not maximal with respect to ϕ. Being a an ℓ-active point, there is either a point $a_+ \in N(a)$ with $\phi(a) < \phi(a_+)$ or a point $a_- \in N(a)$ with $\phi(a) > \phi(a_-)$, in which case the previous property implies the existence of an $a_+ \in N(a)$ with $\phi(a) < \phi(a_+)$, too.

Iterating, a path $A^+ = (a_0^+, \ldots, a_l^+)$ connecting $a = a_0^+$ with a boundary point $a_l^+ = a^+$ maximal with respect to ϕ can be generated. Analogously, paths B^+ and C^+ from b and c to ϕ-maximal boundary points b^+ and c^+ are constructed.

Since there are at most two such ϕ-maximal boundary points, either $a^+ = b^+ = c^+$ or they correspond to two points connected by an edge. Consider now the subgraph G^+ of G containing the paths A^+, B^+, C^+, which is connected. In addition, the three ℓ-active vertices a, b, c have degree one in G^+. Thus, the component G^+ contains a Y-graph Y^+ of endpoints a, b, c having all edges above ℓ. Similarly, there is a Y-graph Y^- with the same endpoints and all edges below ℓ.

Hence, there are three edge-disjoint Y-graphs, namely Y^0, Y^+ and Y^-, having endpoints a, b and c: They form a subgraph homeomorphic to the complete bipartite graph $K_{3,3}$, what contradicts the planarity of G by Theorem 1.2.3.(ii). □

For the second part of the proof of Theorem 1.2.12, a global consistency argument given by a series of claims that we do not list here (we refer the reader to [60]), is used in order to show that all the cells are indeed convex and non-overlapping. □

1.3 Rigidity Theory

It was again Euler who in 1766 conjectured "A closed spatial figure allows no changes, as long as it is not ripped apart" [32], setting the starting point of Rigidity Theory. However, the first major published result [18] is attributed to Cauchy in 1813: the graph of every simplicial 3-polytope is rigid. This was a first step towards the study of Euler's Conjecture, which was finally disproved in 1978 by R. Connelly [21]. Another notorious contribution to the history of rigidity are the works of J.C. Maxwell [52], who studied stresses and forces on graphs and their relations with reciprocal figures.

Although results in Rigidity Theory have been rather sparse from the statement of Euler's conjecture to its resolution, since the 1970s a huge number of researchers have directed their efforts to this field, leading to a number of engineering applications, see for example [42]. Techniques from Rigidity Theory have also been recently applied to answering long-standing problems from robotics or molecular modelling ([22], [69], [12]). Work on understanding some fundamental geometric properties of molecular conformations is currently approached via such techniques [47], with potential applications in protein folding [82], [71].

In this chapter we first give a series of notions and facts, before giving a relation between forces and reciprocal figures which will appear later in our results. Then we give the definition of generic infinitesimal rigidity, which is the fundamental notion for us, and finally we present the graphs with the property of being minimal among those infinitesimally rigid and sketch how they can be constructed. We mainly consider geometric embeddings, although topological embeddings will also appear in Subsection 1.3.2.

1.3.1 The rigidity matrix

A particular geometric embedding $G(S) = (V, E, S)$ of a graph $G = (V, E)$ is usually called a *framework* when referring to its rigidity properties. The *complete framework* for a point set V is that having the complete set $K = \binom{V}{2}$ of edges. As observed before, troughout this thesis we will consider geometric embeddings only into $\mathbb{R}^2$, although most of this section holds for embeddings into a general $\mathbb{R}^m$ as well. In the affine plane (and in the 3-dimensional space), one may think of a framework as a mathematical model for a physical structure in which each vertex i corresponds to a ball joint located at the point p_i and each edge corresponds to a rigid rod connecting the joints at its endpoints.

Obviously, this representation may be used to describe equally rigid structures (like bridges) as moving ones (like machines or organic molecules). Dealing with the distinction between frameworks whose realization is rigid and those which admit movements is the goal of the Rigidity Theory. At a first step, it would be expected that a framework being or not rigid depends both on the graph $G = (V, E)$ and the particular embedding $G(S)$, what makes the difference between the combinatorial and the geometric rigidity. In the rest of the section our goal is to minimize the

importance of the particular embedding and focus mainly on the graph properties. With this aim, we introduce the basic notions needed to understand the results about rigidity presented in the following chapters. Unless otherwise said, for the results in this section we refer the reader to [38].

Definition 1.3.1 Let $G(S)$ be an embedding of a graph G on a finite set of points $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$. The set S is in *general position* and $G(S)$ is a *general embedding* if no 3-subset of S lies on a line.

If we identify an embedding $G(S)$ on the set $S := \{p_1, \ldots, p_n\} \subset \mathbb{R}^2$ with a point $S = (p_1, \ldots, p_n) \in \mathbb{R}^{2n}$, we can measure lenghts of edges by the *rigidity function*

$$\begin{array}{rcl} r : \mathbb{R}^{2n} & \to & \mathbb{R}^{\binom{n}{2}} \\ S & \mapsto & r(S)_{ij} := \|p_i - p_j\|^2 \end{array}$$

which is continuosly differentiable. Then, the *rigidity matrix* $R(S)$ of the embedding $G(S)$ is defined by $r'(S) = 2R(S)$. It is an $\binom{n}{2} \times 2n$ matrix with condensed form

$$R(S) = \begin{pmatrix} p_1 - p_2 & p_2 - p_1 & 0 & 0 & \ldots & 0 & 0 \\ p_1 - p_3 & 0 & p_3 - p_1 & 0 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \ldots & p_{n-1} - p_n & p_n - p_{n-1} \end{pmatrix}$$

We are going to introduce a crucial notion in terms of rows of the rigidity matrix. Recall that these correspond to pairs of vertices and then to edges of the complete graph on V:

Definition 1.3.2 A set of edges $E \subseteq \binom{V}{2}$ is said to be *independent* with respect to an embedding S if the corresponding rows in the rigidity matrix are independent.

Let $\delta(S)$ denote the determinant of some $|E| \times |E|$ minor of the submatrix defined by E. If we permit S to vary over all points in $\mathbb{R}^{2n}$, the solution set to $\delta(S) = 0$ is either all of $\mathbb{R}^{2n}$ or an algebraic hypersurface. The latter is the case if there is an embedding S with respect to which E is independent. Then, the set of all embeddings S for which rows corresponding to E are dependent lies in the intersection $\mathcal{X}_E$ of the algebraic hypersurfaces corresponding to all the $|E| \times |E|$ minors of the submatrix. If E is independent with respect to some embedding, each $\mathcal{X}_E$ is a closed algebraic set of measure zero.

Definition 1.3.3 An embedding S is said to be *generic* if $S \notin \mathcal{X}$, where

$$\mathcal{X} = \cup\{\mathcal{X}_E : \ E \text{ independent for some embedding}\}.$$

Observe that $\mathcal{X}$ is a closed set of measure zero in $\mathbb{R}^{2n}$.

Definition 1.3.4 An edge set $E \subseteq \binom{V}{2}$ is *generically independent* if it is independent with respect to all generic embeddings of V.

This definition, together with that of generic embeddings, makes the following result straightforward:

Lemma 1.3.5 *Let V be a finite set of vertices. If a set of edges $E \subseteq \binom{V}{2}$ is independent with respect to some embedding of V, then it is generically independent.*

We now observe that for a particular embedding S and a set E of edges, considering the minors of the set of rows of $R(S)$ which correspond to E is not a very practical way of checking independence. Instead, one usually looks directly at the dependency relations between rows of the rigidity matrix:

$$\sum_{ij \in E} w_{ij} r_{ij} = 0,$$

for r_{ij} the row corresponding to the edge ij. A set of edges E is independent with respect to the embedding S if, and only if, the only dependency relation satisfied by the corresponding set of rows is the trivial one (that with $w_{ij} = 0$ for all ij).

Definition 1.3.6 A non-zero *stress* on a framework $G(S)$ is an assignment of scalars w_{ij}, not all zero, to the edges E of G in such a way that, for every vertex $i \in V$,

$$\sum_{ij \in E} w_{ij}(p_i - p_j) = 0.$$

For a physical interpretation, the assignment of w_{ij}'s may correspond to the replacement of each edge ij with a spring whose compression or extension is indicated by the value w_{ij}. Note that the last expression is just the previous one considered one column at a time. Then, with respect to the embedding S, a set of edges being independent, that is, their rows in the rigidity matrix having only the trivial dependency, is equivalent to those edges having only the trivial stress.

A graph is said to have a *stress* (sometimes called *self-stress*) if its edges have it on any generic embedding. Analogously, a graph is a *generic rigidity circuit* if for any generic embedding its edges have a non-zero stress but no proper subset of edges has one. In other words, a *circuit* is a minimal dependent set of rows in the rigidity matrix of any generic embedding. Hence, a framework $G(S)$ is a *rigidity circuit* precisely if it has a unique stress and the stress is not zero on any edge.

Example 1.3.7 *The reader may have noticed that, given an embedding $G(S)$ of a graph, having a stress on the edges is equivalent to every vertex (including boundary ones) being in equilibrium with its neighbors. The difference with the situation in Tutte's Theorem is that we now allow negative weights on edges.*

In the next figure we add weights of boundary edges of Figure 1.6. It can be easily checked that the three boundary points are also in equilibrium.

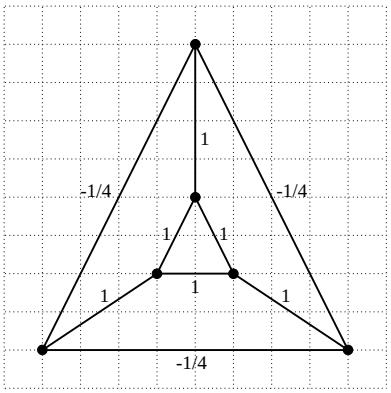

Figure 1.7: A stress in the graph of Figure 1.6.

1.3.2 Stresses and reciprocal graphs

Roughly speaking, in this subsection we show that for those frameworks with stresses such that $w_{ij} \neq 0$ for every edge ij, a dual can always be drawn with an interesting geometric property. In order to do so, we first need a few notions:

Definition 1.3.8 A *spherical framework* is defined to be a pair $(G \hookrightarrow \mathbb{S}^2, G(S))$ composed of a plane topological embedding $G \hookrightarrow \mathbb{S}^2$ and a (possibly not plane) geometric embedding $G(S)$, of a graph G.

A *non-crossing spherical framework* is a plane geometric embedding $G(S)$ together with its trivially corresponding plane topological embedding $G \hookrightarrow \mathbb{S}^2$ (the one given by the inclusion $\mathbb{R}^2 \to \mathbb{S}^2$).

Definition 1.3.9 Given a plane topological embedding $G \hookrightarrow \mathbb{S}^2$ of a 2-connected graph $G = (V, E)$, the set $\underline{E}$ of its *edge patches* is defined to have as elements the ordered 4-tuples $\underline{e} = (i, j; h, k)$ such that the edge from vertex i to vertex j has face h on the right and face k on the left.

Of course, the definition holds as well for plane geometric embeddings; actually we refer the reader to Figure 1.8 (left) in order to check the notion of edge patch. There, edges are denoted by numbers and faces by letters. However, we are going to use this definition in the context of spherical frameworks, where planarity is only guaranteed for the topological embedding.

Definition 1.3.10

- The *dual* $G' \hookrightarrow \mathbb{S}^2$ of a plane topological embedding $G \hookrightarrow \mathbb{S}^2$ was introduced by Poincarè and is the cell decomposition of $\mathbb{S}^2$ having one vertex for each cell of G and in which two vertices are joined by an edge if the corresponding cells are adjacent in G. Observe that the dual of an edge patch $e = (i, j; h, k)$ is $e' = (h, k; j, i)$.

- A spherical framework $(G' \hookrightarrow \mathbb{S}^2, G'(S'))$ is a *reciprocal* of a spherical framework $(G \hookrightarrow \mathbb{S}^2, G(S))$ if $G' \hookrightarrow \mathbb{S}^2$ is the dual of $G \hookrightarrow \mathbb{S}^2$ and for each edge patch $(i, j; h, k) \in \underline{E}$,
$$\langle p_i - p_j, p'_h - p'_k \rangle = 0.$$
That is to say, the edges of the geometric embeddings $G(S)$ and $G'(S')$ are perpendicular.

Figure 1.8 shows the geometric embeddings $G(S)$ (left) and $G'(S')$ (right) of a non-crossing spherical framework and its reciprocal. The points $a, \ldots, d$ in the original represent faces.

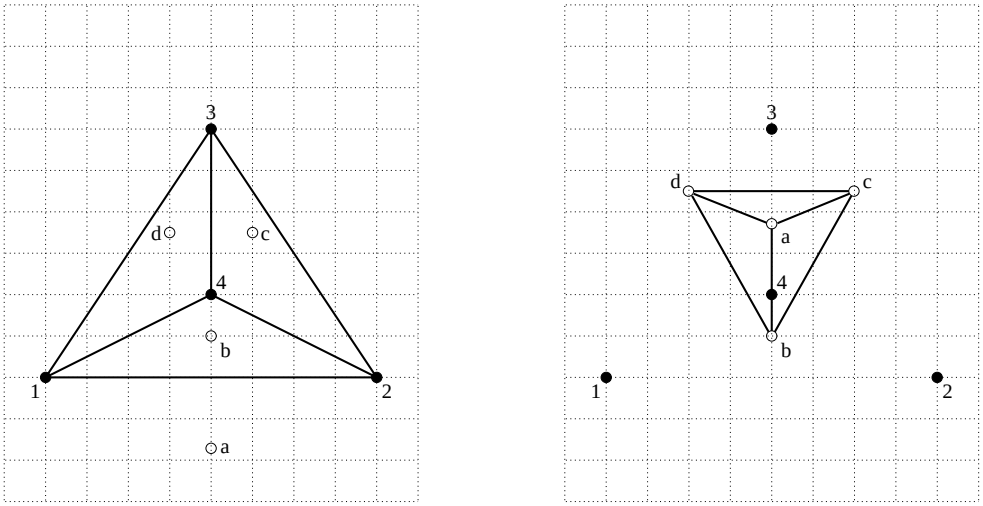

Figure 1.8: Reciprocal geometric embeddings.

Note that the dual of a plane topological embedding always exists. Hence, given a spherical framework, the non-trivial part in the definition of reciprocal is the existence of $G'(S')$. This is why in the next result we abuse of the notation and omit the topological part of spherical frameworks:

Theorem 1.3.11 (Maxwell-Cremona) *For a spherical framework $G(S)$, the following are equivalent:*

(i) *$G(S)$ has a reciprocal $G'(S')$,*

(ii) *$G(S)$ has a stress $\sum_{ij \in E} w_{ij}(p_i - p_j) = 0$ such that $w_{ij} \neq 0, \forall ij \in E$ (non-zero on all edges).*

(iii) *There is a piecewise-linear lifting (may be non-convex and self-intersecting) of $G(S)$ into 3 dimensions, linear on every face.*

Furthermore, if the conditions hold for $G(S)$:

- *Edges with positive/negative stress correspond to valleys/ridges in the lifting of (iii). (Meaning that the slope increases/decreases when crossing the edge).*

- *The conditions also hold for the reciprocal $G'(S')$.*

Proof: We will refer to the equivalence (i)⇔(ii), appeared in [52], as Maxwell's Theorem. We give here its proof, which will be crucial in Chapter 5. For that of the equivalence to the third condition see [24], [43] or [78]. Note that the two sides of a valley (resp. ridge) do not necessarily "go up" (resp. "go down").

(ii)⇒(i) Assume that there is such a stress. Choose an initial point p'_0 for the reciprocal corresponding to a cell F_0 of the original, and define the point corresponding to another cell F_c by $p'_c = p'_0 + (\sum_{e'} w_e d_e)^\perp$ where the sum is over the edge patches $e' = (h, k; j, i)$ of any path from F_0 to F_c in the dual graph, and where $d_e = p_j - p_i$. Observe that this is an inductive process for which the intermediate points between p'_0 and p'_k have been previously obtained. The first task is to prove that, given any two such paths between F_0 and F_c, both lead to the same p'_c, what is true since the sum up one path and back the other follows a sum of face cycles, and each face cycle is zero by the equilibrium condition on its reciprocal vertex. Since no w_{ij} is zero, there is no degenerate edge in the reciprocal.

(i)⇒(ii) Given a reciprocal, the previous construction can be reversed, solving the system $(w_e d_e)^\perp = p'_k - p'_h$ to get back the unique corresponding stress in the original framework. Since the vector sum of edges in every face cycle is zero, the vertices of the original are in equilibrium. □

Observation 1.3.12 Let $G'(S')$ be the reciprocal of $G(S)$ obtained from a certain stress of weights w_{ij}. Let $e = (i, j; h, k)$ denote an edge patch in the original, then:

- $p'_k - p'_h = w_{ij}(p_j - p_i)^\perp$. That is; the straight segment in the reciprocal is perpendicular to the corresponding in the original, scaled by the absolute value of the stress and with direction given by the sign.

- $G'(S')$ has a stress non-zero on all edges defined as $w'_{hk} := 1/w_{ij}$.

Example 1.3.13 *In Figure 1.8 the positions of $a, \ldots, d$ in the reciprocal (right) are obtained according to the previous proof, once defined in the original the stress $w_{14} = w_{24} = w_{34} = 1$ for interior and $w_{12} = w_{23} = w_{13} = -1/3$ for exterior edges. Observe that in the reciprocal the point a corresponding to the exterior face turns out to be above b, what corresponds to the edge $\overline{12}$ having negative stress.*

Remark 1.3.14 Even for non-crossing spherical frameworks, Maxwell's Theorem 1.3.11 only claims the existence of a reciprocal spherical framework, which is not guaranteed to be non-crossing. In Chapter 5 we characterize the conditions for non-crossing spherical frameworks to have a non-crossing reciprocal.

1.3.3 Infinitesimal rigidity and isostatic graphs

In this subsection we present the notion of rigidity we will use in this thesis and an outstanding type of graphs, which are considered to be the fundamental objects in 2-dimensional Rigidity Theory. We focus again on geometric embeddings.

Definition 1.3.15 Let $G(S) = (\{1, \dots, n\}, E, S)$ be a framework and consider an assignment of a 2-dimensional vector v_i to each vertex i. Such an assignment $v \in \mathbb{R}^{2n}$ is said to be an *infinitesimal motion* of the framework if the following homogeneous equation holds for every edge $ij \in E$:

$$\langle v_i - v_j, p_i - p_j \rangle = 0.$$

The set of infinitesimal motions of such a framework is denoted $\mathcal{V}(E)$ and is clearly a linear subspace of $\mathbb{R}^{2n}$.

Each v_i can be interpreted as a velocity vector for the point p_i in the plane. Infinitesimal motions are sometimes called *flexes* in the rigidity literature. For an interpretation, consider the vector assignment v as the initial velocities of a motion for the points p_i, such that they neither compress nor strecht the rods of the framework. This is why a vector v such that

$$\langle v_i - v_j, p_i - p_j \rangle \geq 0,$$

that is, such that it increases or keeps the length of the rods, is said to be an *expansive infinitesimal motion*.

Here comes the precise definition of rigidity we are interested in:

Definition 1.3.16 A framework $G(S) = (\{1, \dots, n\}, E, S)$ is said to be *infinitesimally rigid* if $\mathcal{V}(E)$ equals the subspace $\mathcal{V}(K)$ of infinitesimal motions of the complete framework $G(K)$, which are the translations or rotations of the whole graph.

A graph G is *generically infinitesimally rigid* if it is infinitesimally rigid for every generic embedding S.

Lemma 1.3.17 ([38]) *If $G(S)$ is infinitesimally rigid for some general embedding S of V, then the graph G is generically infinitesimally rigid.*

In the sequel we focus on infinitesimal rigidity in the generic sense and, when referring to a graph as "infinitesimally rigid", we will mean "generically infinitesimally rigid". A graph $G = (V, E)$ is said to be *isostatic* or *minimally (infinitesimally) rigid* if it is (infinitesimally) rigid and its edge set E is independent. Infinitesimal rigidity is then linear algebra, with infinitesimally rigid graphs, stresses and isostatic graphs corresponding to spanning sets, linear dependencies and bases of the row space of the rigidity matrix of the complete graph.

Let us observe that, for each edge ij, the condition $\langle v_i - v_j, p_i - p_j \rangle = 0$ determines a hyperplane in $\mathbb{R}^{2n}$; that of the vectors v which keep the length of the rod $p_i p_j$.

Therefore, a hyperplane arrangement is associated to each framework via its edges. Orienting the hyperplanes according to whether the rods are stretched or shrunk determines the so-called *rigidity matroid*.

To finish the present section, we consider the following type of graphs, which provide a combinatorial characterization of infinitesimal rigidity:

Definition 1.3.18 A graph $G = (V, E)$ is said to be *Laman* if the following conditions hold:

(i) $|E| = 2|V| - 3$, and

(ii) For every subgraph (V', E') with $|V'| \geq 2$ vertices, $|E'| \leq 2|V'| - 3$.

Condition (ii), which is sometimes called the *hereditary Laman property*, can be equivalently restated as:

(ii') Every subset V'' of vertices with $|V''| \leq |V| - 2$ is incident to at least $2|V''|$ edges.

There is a variety of ways to characterize Laman graphs. The one we are most interested in for the purposes of this thesis, since a combinatorial version will be given in Chapter 4, is the *Henneberg construction* [41]. It goes inductively as follows: For a Laman (abstract) graph on n vertices, start with a triangle for $n = 3$. At each step, a new vertex is added in one of the following two ways:

- **Henneberg I** (vertex addition): The new vertex is connected via two edges to two of the old vertices.
- **Henneberg II** (edge splitting): A new vertex is added on some edge, splitting this into two new edges, and then connected to a third vertex. Equivalently, an edge is removed and then a new vertex is connected to its endpoints and to some other vertex.

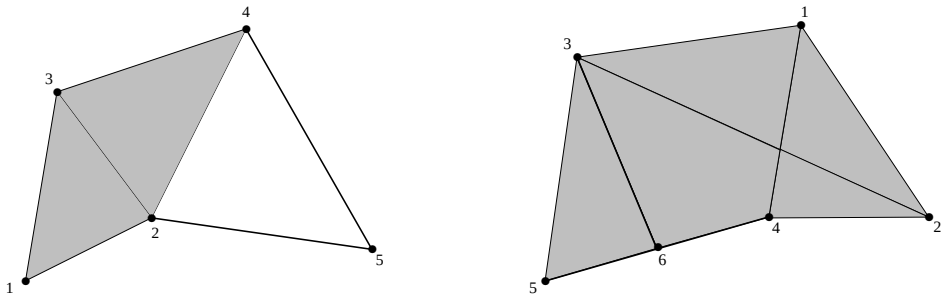

Figure 1.9: Henneberg steps.

Example 1.3.19 *See Figure 1.9 for an example: The points have been added according to their labels and faces obtained previously to the current step are shaded. In the right picture we have chosen a drawing with crossing edges, in order to emphasize that the Henneberg construction holds for* general *and not only planar graphs.*

The next result shows that Laman graphs are the maximal independent sets, and justifies why they are fundamental in 2-dimensional Rigidity Theory. See [50], [38], [80]:

Lemma 1.3.20 ([81]) *Let $G = (V, E)$ be a graph with $|V| \geq 2$. Then, the following are equivalent:*

(i) G is isostatic,

(ii) G is infinitesimally rigid with $|E| = 2|V| - 3$,

(iii) G is independent with $|E| = 2|V| - 3$,

(iv) G is a Laman graph (Laman's Theorem, [50]),

(v) There is a Henneberg construction for G (Henneberg's Theorem, [41]).

In order to conclude this section about rigidity, we include here the following result relating rigidity to polytopes via their graphs. For $d = 3$ it was proved by Dehn [29] in 1916, and the generalization to $d > 3$ is due to Whiteley [79]:

Theorem 1.3.21 *The graph of a simplicial d-polytope P is infinitesimally rigid.*

Let us mention that the $d = 3$ case is an inmediate consequence of the theorem stated by Cauchy [18] in 1813 which we mentioned at the beginning of the section, whose proof turned out to have several minor errors and was completed later (see [67]).

1.4 Pseudo-triangulations of convex bodies and points

Pseudo-triangulations are relatively new objects which have by now been used in many Computational Geometry applications, among them visibility [56, 58, 57, 65], ray shooting [37], kinetic data structures [1, 49] and motion planning for robot arms [69]. In this last paper, Streinu introduced the *minimum* or *pointed pseudo-triangulations* and used them to prove the Carpenter's Rule Theorem (the first proof of which was given shortly before by Connelly et al. [22]).

They have rich combinatorial, rigidity theoretic and polyhedral properties, many of which have only recently started to be investigated, see also [61], [59], [48], [3], [11], [2], [17]. Finding their 3-dimensional counterpart, which is perhaps the main open question about pseudo-triangulations and expansive motions, may lead to efficient motion planning algorithms for certain classes of 3-dimensional linkages, with

potential impact on understanding protein folding processes. We introduce pseudo-triangulations in the context of visibility between convex bodies, where they originally appeared, before turning our attention to pseudo-triangulations of points, those in which we will focus in this thesis.

1.4.1 Pseudo-triangulations of convex bodies

Although in this work we are not dealing with this type of pseudo-triangulations, their importance comes from the fact that precisely in this context Pocchiola and Vegter started the study of pseudo-triangulations in the seminal papers [56] and [57], in which the results of this subsection can be found. A precedent appears in the *geodesic triangulations* of polygons introduced in [19].

Pocchiola and Vegter consider a collection $\mathcal{O} = \{B_1, \ldots, B_n\}$ of pairwise disjoint convex bodies in the plane, which are assumed to be *strictly convex* (meaning that the open line segment joining any two points lies in the interior of the object), to have a *smooth* C^1-differentiable boundary and to be in *general position* in the sense that no three objects share a common tangent line. The following definition collects most of the notions needed in this subsection, and can be checked in Figure 1.10:

Definition 1.4.1 In the situation above,

- A *bitangent* for a pair of convex bodies is a line segment which is tangent to the two objects at its endpoints, in the sense that the line containing that segment is locally a supporting line. There are four such bitangents; two which separate the objects and two which do not.
- The set of *bitangents* of a collection $\mathcal{O} = \{B_1, \ldots, B_n\}$ of pairwise disjoint convex bodies is composed of those line segments which are bitangents of some pair of the objects and do not intersect the others. It is called the *tangent visibility graph.*
- A *pseudo-triangulation* of $\mathcal{O} = \{B_1, \ldots, B_n\}$ is the subdivision of the plane induced by a maximal (with respect to the inclusion) family of pairwise non-crossing bitangents of $\mathcal{O}$.

Lemma 1.4.2 *Every pseudo-triangulation of n convex bodies in general position subdivides* $\mathrm{conv}(\cup_i B_i) \setminus \cup_i B_i$ *into* $2n-2$ *cells, and the number of bitangents is* $3n-3$.

The goal of Pocchiola and Vegter was the study of the tangent visibility graph of a collection $\mathcal{O}$ of n convex bodies. In particular, they used pseudo-triangulations to obtain the first algorithm which constructs the graph in optimal time, $O(n\log(n)+k)$ (for k the graph size), using only linear space $O(n)$. In their algorithm, the pseudo-triangulations counterpart of the well-known notion of *diagonal flip* in triangulations (see, for example [46]) turned out to be crucial:

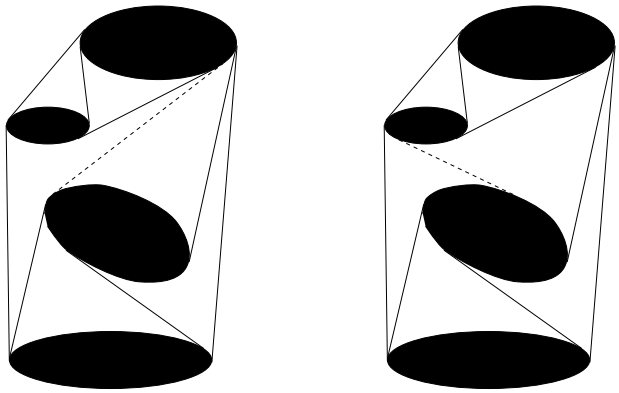

Figure 1.10: Flip between two pseudo-triangulations of a family of convex bodies.

Lemma 1.4.3 *Any two disjoint cells T and T' have exactly one common tangent line. Furthermore, the removal of an interior bitangent gives a cell which can be split in exactly another way by adding a different bitangent. This operation is called* flip*.*

An example of flip can be found between the dashed edges in Figure 1.10.

Corollary 1.4.4 *From the two previous Lemmas, every pseudo-triangulation of $\mathcal{O}$ has exactly $3n-3-n_b$ flips, where n_b is the number of bitangents in the boundary of* $\mathrm{conv}(\cup_i B_i)$*.*

Thus, the *graph of pseudo-triangulations* of $\mathcal{O}$, having pseudo-triangulations as vertices and flips as edges, is regular of degree $3n-3-n_b$. Pocchiola and Vegter show that, in addition, it is connected.

1.4.2 Pseudo-triangulations of points

Consider a set $\mathcal{A}$ of n points in general position in the plane (one of the consequences of the results in Chapter 4 is a natural definition of pseudo-triangulations in presence of collinearities as well). A *triangulation* of $\mathcal{A}$ is a subdivision of the polygon $\mathrm{conv}(\mathcal{A})$ into triangles, in the sense of Remark 1.1.12, assuming that *all* the points of $\mathcal{A}$ are used as vertices.

Definition 1.4.5 Let us call *convex* and *reflex angles* those respectively smaller and greater than π and observe that the general position assumption implies the absence of angles equal to π.

A *pseudo-triangle* is a simple polygon with only three convex vertices (from the interior), called *corners* and joined by three inward convex polygonal chains, called *pseudo-edges* of the pseudo-triangle. See Figure 1.11.(a).

More generally, a *pseudo-quadrilateral*, *pseudo-pentagon*, etc. is a simple polygon with exactly four, five, etc. convex vertices. For them a *diagonal* is defined as a segment contained in the pseudo-polygon and tangent to its boundary at the endpoints. Note that every convex k-gon is a pseudo-k-gon. In particular, every triangle is a pseudo-triangle. Note also that every (perhaps non-convex) polygon is a pseudo-k-gon for some k.

A *pseudo-triangulation of* $\mathcal{A}$ is a geometric non-crossing graph with vertex set $\mathcal{A}$ and which partitions $\text{conv}(\mathcal{A})$ into pseudo-triangles. Every triangulation is a pseudo-triangulation.

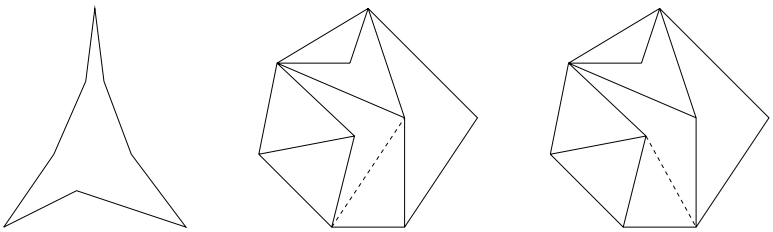

Figure 1.11: (a) A pseudo-triangle. (b) A pointed pseudo-triangulation. (c) The dashed edge in (b) is flipped, giving another pointed pseudo-triangulation.

Since triangulations of $\mathcal{A}$ coincide with the maximal non-crossing graphs on $\mathcal{A}$, they are also the maximal pseudo-triangulations. As is well-known, they all have $2n_v + 3n_i - 3$ edges, where n_v and n_i denote the numbers of points of $\mathcal{A}$ in the boundary and interior of $\text{conv}(\mathcal{A})$, respectively. It turns out that also the pseudo-triangulations with the minimum possible number of edges are very interesting from different points of view:

Definition 1.4.6 A vertex of a geometric graph is called *pointed* if its incident edges span an angle smaller than 180 degrees from that vertex. The graph itself is called *pointed* if all its vertices are pointed.

A *pointed pseudo-triangulation* is that on which every vertex is pointed.

The following statement comes originally from [69]. We sketch here its proof, which can also be found in [61].

Proposition 1.4.7 (Streinu) *Let $\mathcal{A}$ be a planar point set as above. Then:*

(i) Every pseudo-triangulation of $\mathcal{A}$ with n_γ non-pointed vertices and n_ϵ pointed vertices has $2n - 3 + n_\gamma = 3n - 3 - n_\epsilon$ edges and $n - 2 + n_\gamma = 2n - 2 - n_\epsilon$ pseudo-triangles.

(ii) Every pointed and non-crossing graph on $\mathcal{A}$ has at most $2n - 3$ edges, and is contained in some pointed pseudo-triangulation of $\mathcal{A}$.

Proof: The first part comes from Euler's formula, which if we call t and e the numbers of pseudo-triangles and edges, respectively, gives $n - e + t = 1$. We are going to obtain an additional formula by double-counting of the number of convex angles: On the one hand, each of the pseudo-triangles produces exactly three of them. On the other, for each non-pointed vertex there are as many convex angles as incident edges, and for each pointed vertex the number of convex angles is one less than that. Thus, $2e - n_\epsilon = 3t$ and the formula above gives the desired values of e and t.

For the second part, we first observe that we can assume the graph to be connected, since otherwise their components can be analyzed separately. Pointedness means that the number of reflex angles is n. The number of convex angles is at least $3t$, with equality if, and only if, the graph is a pseudo-triangulation. Since the total number of angles sums up to $2e$, we get $2e \geq n + 3t$ which together with Euler's formula implies $e \leq 2n-3$ (and $t \leq n-2$). The last sentence in the statement comes then from two facts: On the one hand, that the addition of *geodesic paths* (those non-crossing paths having shortest length among paths sufficiently close to them) between convex vertices of a polygon keeps the graph pointed and non-crossing, and on the other the fact that there is always some geodesic path through the interior of a polygon if this is not a pseudo-triangle. □

Remarks 1.4.8 Observe that every vertex on the convex hull of $\mathcal{A}$ is trivially pointed. Then, the number n_ϵ of pointed vertices in a pseudo-triangulation runs from n_v, in which case all the pseudo-triangles are indeed triangles and the pseudo-triangulation is actually a triangulation, to n for a pointed pseudo-triangulation.

Part (i) implies that, among pseudo-triangulations of $\mathcal{A}$, pointed ones have the minimum possible number of edges. For this reason they are also called *minimum pseudo-triangulations*.

Part (ii) says that pointed pseudo-triangulations coincide with maximal non-crossing and pointed graphs.

Given a planar point configuration $\mathcal{A}$, it is easy to construct a pointed pseudo-triangulation of it as follows: First, order the points of $\mathcal{A}$ according to their x-coordinate. Then, consider the triangle defined by the three starting points in the ordering. Finally, add iteratively the two tangents joining each new point to the previously constructed pseudo-triangulation. See Figure 1.12. Corollary 1.4.10 below will show that it is not a coincidence that this procedure is a particular case of a Henneberg construction.

Note that a minimum (i.e. pointed) pseudo-triangulation is always "minimal" in the sense that the removal of any edge implies we no longer have a pseudo-triangulation. However, minimal pseudo-triangulations need not be pointed. For an example, consider Figure 1.13; the leftmost pseudo-triangulation is minimum (therefore minimal), the middle one is minimal but not minimum and the rightmost is not minimal and thus neither minimum.

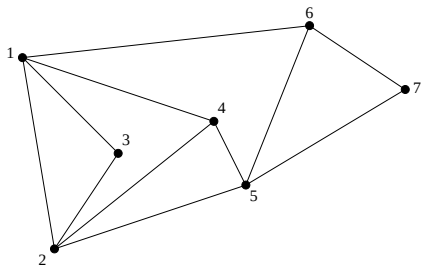

Figure 1.12: Construction of a pointed pseudo-triangulation.

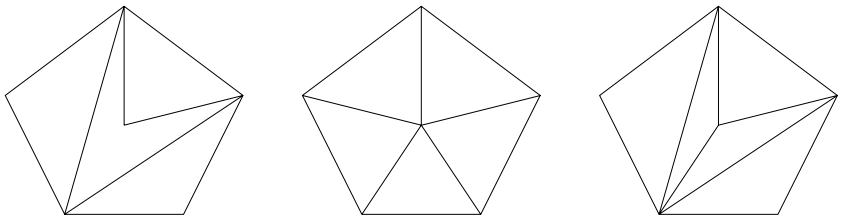

Figure 1.13: Minimum, minimal and not minimal pseudo-triangulations.

The following result, taken from [69], summarizes the above characterizations of pointed pseudo-triangulations and introduces a new one relating them to rigidity:

Theorem 1.4.9 (Characterization of pointed pseudo-triangulations) *Let $G = (V, E)$ a graph embedded on the point set $\mathcal{A} = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^2$. Then, the following properties are equivalent:*

(i) G is a pointed pseudo-triangulation,

(ii) G is a pseudo-triangulation and $|E| = 2n - 3$ (equivalently, the number of pseudo-triangles is $n - 2$),

(iii) G is a minimum pseudo-triangulation, that is, it has the minimum number of edges (equivalently, of pseudo-triangles) among all pseudo-triangulations of $\mathcal{A}$.

Moreover, if any of the above conditions is satisfied, then G fulfills the hereditary Laman property; for every subgraph (V', E') with $|V'| \geq 2$ vertices, $|E'| \leq 2|V'| - 3$.

Proof: (i)⇔(ii)⇔(iii) follow from part (i) of Proposition 1.4.7. The Laman property follows from part (ii) of the same result. □

Corollary 1.4.10 (Streinu) *Pointed pseudo-triangulations are Laman graphs, therefore isostatic.*

In particular, the fact that in addition to be Laman graphs, pointed pseudo-triangulations when stripped of a convex hull edge become expansive mechanisms [69], has proved to be crucial in designing efficient motion planning algorithms for planar robot arms.

1.5 Results and Open problems

According to the problems studied, this thesis can be divided into two parts: The first one deals with the problem of finding efficient triangulations of the hypercube and is carried out in Chapter 2. The rest of the thesis, Chapters 3 to 5, is devoted to the study of planar graphs, pseudo-triangulations and their relations with rigidity. Since the techniques we use are different, both parts can be regarded separately, although the two of them are framed in the context of Geometric Combinatorics.

1.5.1 Efficient triangulations of cubes

Once known a triangulation of efficiency ρ of a certain cube, Haiman's construction gives triangulations of size $\rho^n n!$ of high-dimensional cubes. This turns the problem of deciding how efficient triangulations of cubes can be into the study of triangulations of cubes of small dimensions. However, even this study has turned out to be a hard problem: Triangulations of smallest size ϕ_d, and hence smallest efficiency ρ_d, are only known to date until dimension $d = 7$. And no better values have been found disregarding the smallest condition. In particular, the best asymptotic size of triangulations of cubes until the development of this thesis was $0.840^n n!$, obtained from ρ_7.

In Chapter 2 (joint work with Francisco Santos [54]) we present a construction to triangulate products of two general polytopes. As a particular case, we apply our construction to products of cubes and obtain triangulations of high dimensional cubes whose efficiency is not that of a triangulation of a small dimensional cube, but the "efficiency" of a triangulation of the product of a cube and a simplex, both of small dimensions.

First, we have to define the appropiate notion of efficiency for a triangulation of such a product, which we call *weighted efficiency* (Definition 2.2.1). The main result in the chapter is:

(Theorem 2.2.3) If there is a triangulation of $I^l \times \Delta^{m-1}$ with weighted efficiency ρ_0, then

$$\lim_{n\to\infty} \rho_n \leq \rho_0.$$

As above, this leads us to study the smallest weighted efficiency of triangulations of $I^l \times \Delta^{m-1}$, which we denote by $\rho_{l,m}$. This number can be calculated solving a linear programming problem; in this way we have obtained several values, the smallest one being $\rho_{3,3} = 0.8159$. From this we conclude that:

- For n large enough, the cube I^n can be triangulated with $0.8159^n n!$ simplices.

The natural goal of further work in this part of the thesis should be the computation of better values of $\rho_{l,m}$. Concerning this point, we have to say we are not very enthusiastic about using a "brute force" approach, at least by the moment: We have tried to obtain $\rho_{4,2}$ solving the linear programming problem in a quite powerful computer of the University of California at Davis without any result after weeks of computations. Instead, an interesting approach might be trying to adapt the methods used by Anderson and Hughes [6] to calculate ϕ_6 and ϕ_7, involving *ad hoc* ways of decomposing the system into smaller subsystems. But there are plenty of technicalities and we are not sure about whether they would significatively improve our values, since the efficiency turns out to be a very insensitive parameter.

1.5.2 Non-crossing graphs, rigidity and pseudo-triangulations

The second part of the thesis starts in Chapter 3 (joint work with Francisco Santos [53]). There, we consider a planar point set $\mathcal{A}$ and first study its pseudo-triangulations. Let us denote by n_i and n_v the number of interior points and vertices of $\text{conv}(\mathcal{A})$, respectively. Note that in non-general position, there may appear other points in the boundary of $\text{conv}(\mathcal{A})$.

Before the work presented in this thesis, pseudo-triangulations had only been considered for point sets in general position. There was a well-known notion of *flip* in pointed pseudo-triangulations such that the *graph of pointed pseudo-triangulations*, having these as vertices and flips as edges, is connected and regular of degree $2n_i + n_v - 3$. This graph has been proved by Rote, Santos and Streinu [61] to be the 1-skeleton of a simple polytope $X_f(\mathcal{A})$ of dimension $2n_i + n_v - 3$.

Our first observation was that, still under the general position assumption, one can extend the notion of flip to general pseudo-triangulations, not necessarily pointed (Definition 3.2.6). For the resulting *graph of pseudo-triangulations* we get:

(Proposition 3.2.7) The graph of pseudo-triangulations of $\mathcal{A}$ is connected and regular of degree $3n_i + n_v - 3 = 3n - 2n_v - 3$.

As suggested by this result and the above precedent, we give a positive answer to the question of whether this graph is the 1-skeleton of a simple polytope of dimension $3n_i + n_v - 3$. Our main results are:

(Theorem 3.1.2) Let $\mathcal{A}$ be a finite set of n points in the plane, not all contained in a line. Then, there is a simple polytope $Y_f(\mathcal{A})$ of dimension $2n_i + n - 3$ such that its vertex set is in bijection to the set of all pseudo-triangulations of $\mathcal{A}$. The 1-skeleton of $Y_f(\mathcal{A})$ is the graph of flips between them.

(Theorem 3.1.1) There is a face F of $Y_f(\mathcal{A})$, of dimension $2n_i + n_v - 3$, such that the complement of the star of F in the face-poset of $Y_f(\mathcal{A})$ equals the poset of non-crossing graphs on $\mathcal{A}$ that use all the convex hull edges.

In particular, it is trivial to obtain from this the poset of *all* non-crossing graphs on $\mathcal{A}$, since convex hull edges do not affect to planarity. In addition, the face F equals the polytope $X_f(\mathcal{A})$ of [61].

Theorem 3.1.2 gives an answer to an open question posed in [69]. Furthermore, we show these results and the construction of the polytope $Y_f(\mathcal{A})$ to work even if $\mathcal{A}$ is in non-general position (Section 3.5). This leads us to define non-crossing graphs, pointedness, pseudo-triangles, pseudo-triangulations and flips in presence of collinearities as well (Definitions 3.5.1, 3.5.2 and 3.5.5).

Although it is in Chapters 4 and 5 where rigidity comes to foreground, it also plays a role in the techniques used to construct $Y_f(\mathcal{A})$, and in an interesting by-product of them: Pointed pseudo-triangulations were already proven by Streinu [69] to be minimally rigid graphs. Going further, on the one hand we prove:

(Theorem 3.1.3) Pseudo-triangulations are rigid graphs.

On the other hand, we devote Chapter 4 (joint work with Günter Rote and Francisco Santos) to prove the converse of Streinu's result:

(Theorem 4.1.1) Every planar minimally rigid graph can be embedded in $\mathbb{R}^2$ as a pointed pseudo-triangulation.

Hence, we provide a simple and elegant combinatorial characterization of all the graphs which can be geometrically embedded as pointed pseudo-triangulations, answering an open question posed in [68]:

(Corollary 4.1.2) Given a graph G, the following conditions are equivalent:

(i) G is planar and minimally rigid (isostatic),

(ii) G is a planar Laman graph,

(iii) G can be geometrically embedded as a pointed pseudo-triangulation.

The most famous open question in Rigidity Theory (the *Rigidity Conjecture* [38]) is finding the 3-dimensional counterpart of planar Laman graphs, which we have characterized as pointed pseudo-triangulations. However, finding an appropiate definition of 3-dimensional pseudo-triangulations is not an easy task. It would be interesting to generalize the above construction of $Y_f(\mathcal{A})$ using 3-dimensional rigidity, in order to obtain a polytope whose vertices could be defined to be the pseudo-triangulations of a point set $\mathcal{A}$ in $\mathbb{R}^3$.

It is also natural to wonder about the relation between general rigid graphs (not necessarily isostatic) and general pseudo-triangulations (not necessarily pointed). In particular, whether some or all of the equivalences in the above characterization can

be extended. For this we have a partial answer: We introduce the notion of *combinatorial pseudo-triangulation (c.p.t.)*, defined to be a graph together with a labeling of its angles which induces the combinatorial structure of a pseudo-triangulation (Definition 4.2.1). Furthermore, for c.p.t.'s we define a *generalized Laman property* such that we conjecture:

(Conjecture 4.5.3) Given a graph G, the following conditions are equivalent:

(i) G is planar and rigid,

(ii) G has a c.p.t. labeling with the generalized Laman property,

(iii) G can be geometrically embedded as a pseudo-triangulation.

Our partial answer includes (iii)⇒(i) and (iii)⇒(ii), which are proved in the above mentioned Theorem 3.1.3. On the other hand, we prove (ii)⇒(iii) (Theorem 4.2.4). Nevertheless, we leave open to prove (i)⇒(ii) or (i)⇒(iii), any of which would provide a complete combinatorial characterization of rigid planar graphs.

Finally, in Chapter 5 (joint work with several of the co-authors in [39]) we turn our attention to the conditions for a non-crossing spherical framework $G(S)$ to admit a non-crossing reciprocal. It is well-known that those with a stress non-zero on every edge have a reciprocal $G'(S')$, but nothing was known about when this is non-crossing as well. We give a complete description of the conditions in terms of the shape of the faces of $G(S)$ and the sign patterns of the stress along their boundary edges:

(Theorems 5.2.1 and 5.2.3) Let G be the geometric embedding of a non-crossing spherical framework, with a stress nonzero on all edges. Then the reciprocal of G produced by that stress is non-crossing as well if, and only if, the following holds:

1. The complement of the outer face of G is convex and all its edges have the same sign.

2. All the interior faces of G are either pseudo-triangles or pseudo-quadrilaterals, with one of the following sign patterns for their edges:

   2.*a* A pseudo-triangle in which all the edges in the same pseudo-edge have the same sign and one pseudo-edge has opposite sign to the other two. In other words, there are two sign changes along the boundary of the pseudo-triangle and both occur at corners.

   2.*b* A pseudo-triangle with four sign changes along the boundary, three of which occur at the three corners.

   2.*c* A pseudo-quadrilateral in which all the edges in the same pseudo-edge have the same sign and consecutive pseudo-edges have opposite signs.

Moreover, in the reciprocal graph these four situations correspond respectively to:

1′. A non-pointed vertex with all the incident edges having the same sign.

2.a' A pointed vertex with two sign changes around it, none of which occurs at the big angle.

2.b' A pointed vertex with four sign changes around it, one of which occurs at the big angle.

2.c' A non-pointed vertex with four sign changes around it.

For the particular case in which only pseudo-triangles appear as faces, and if the stress is unique, G is a pseudo-triangulation with exactly one non-pointed vertex. We prove its reciprocal to be again such a pseudo-triangulation:

(Theorem 5.1.3) The reciprocal of a rigidity circuit pseudo-triangulation is again a rigidity circuit pseudo-triangulation.

Chapter 2

Asymptotically efficient triangulations of the d-cube

Consider polytopes P and Q. In practice, we assume P of "low" dimension and Q of "high" dimension. In this chapter, we show how to triangulate the product $P \times Q$ efficiently (i.e., with few simplices) starting with a given triangulation of Q. Our method has a computational part, where we need to compute an efficient triangulation of $P \times \Delta^m$, for a (small) natural number m of our choice. Δ^m denotes the m-simplex. Our procedure can be applied to obtain (asymptotically) efficient triangulations of the cube I^n: We decompose $I^n = I^k \times I^{n-k}$, for a small k. Then we recursively assume we have obtained an efficient triangulation of the second factor and use our method to triangulate the product. The outcome is that using $k = 3$ and $m = 2$, we can triangulate every cube I^n with $O(0.8159^n n!)$ simplices, instead of the $O(0.840^n n!)$ achievable before.

2.1 Introduction

Recall from Definition 1.1.16 that the *size* of a triangulation or dissection T is its number of d-simplices and we denote it $|T|$. It is an open question whether high dimensional cubes admit dissections with less simplices than needed in a triangulation. Actually, the minimum size of dissections of I^7 is unknown, while that of triangulations is known (see below).

The paper [26] describes a general method to obtain the smallest triangulation of a polytope P as the optimal integer solution of a certain linear program. The linear program has as many variables as d-simplices with vertex set contained in the vertices of P exist. That is, $\binom{\#\text{vertices}}{d+1}$ if the vertices of P are in general position and less than that if not. For the d-cube, the direct application of this method is impossible in practice beyond dimension 4 or 5. With a somewhat similar method but simplifying the equations using simmetries of the cube, Anderson and Hughes [6] have calculated the smallest size of a triangulation of the 6-cube and the 7-cube, in a computational *tour-de-force* which involved a problem with 1 456 318 variables

and *ad hoc* ways of decomposing the system into smaller subsystems. The smallest sizes up to dimension 7 are shown in Table 2.1.

We also know from Subsection 1.1.2 that in order to compare sizes of triangulations of cubes in different dimensions the *efficiency* $(|T|/d!)^{\frac{1}{d}}$ of a triangulation T of the d-cube has to be used and that triangulations of maximal efficiency 1 (i.e., unimodular) can be easily constructed in any dimension. On the other extreme, Hadamard's inequality for determinants of matrices with coefficients in $[-1,1]$ implies that the volume of every d-simplex inscribed in the regular d-cube $I^d = [0,1]^d$ is at most $(d+1)^{(d+1)/2}/2^d d!$ Hence, every triangulation has size at least $2^d d!/(d+1)^{(d+1)/2}$ and efficiency at least $2/(d+1)^{\frac{d+1}{2d}} \approx 2/\sqrt{d+1}$.

Following the notation in Subsection 1.1.2, let ϕ_d and ρ_d be the smallest size and efficiency, respectively, of possible triangulations of the cube of dimension d. The number ϕ_d (or some variations for the cases of dissections and/or adding new vertices) is also known as *simplexity* of the d-cube. Recall that $\rho_d = (\phi_d/d!)^{1/d}$.

As we said there, Haiman observed in [40] (see also [51], pages 283–284) that a triangulation of I^{k+l} with $t_k t_l \binom{k+l}{k}$ simplices can be constructed from given triangulations of I^k and I^l with t_k and t_l simplices respectively, and therefore one gets (Corollary 1.1.21):

$$\lim_{i\to\infty} \rho_i \leq \rho_d \qquad \forall d \in \mathbb{N}.$$

What is the limit of this sequence? In particular, is it positive or is it zero? The known values of ρ_d (up to $d = 7$) form a strictly decreasing sequence, as shown in Table 2.1, but it is not even known whether this occurs in general.

Dimension	1	2	3	4	5	6	7	8
Smallest size (ϕ_d)	1	2	5	16	67	308	1493	≤ 11944
Smallest efficiency (ρ_d)	1	1	.941	.904	.890	.868	.840	$\leq .859$
Lower bound Hadamard's ineq.	1	.877	.794	.731	.683	.643	.610	.581
Lower bound Smith [64]	1	1	.941	.889	.833	.789	.751	.718

Table 2.1: Smallest size and efficiency of triangulations of I^d.

Remark 2.1.1 Observe that ρ_d is a very insensitive parameter. Although ϕ_7 equals 1493, which is less than the third part of the maximal size $7! = 5040$, the smallest efficiency $\rho_7 = 0.840$ is far from being one third of the maximal efficiency 1.

Concerning lower bounds, the only significant improvement to Hadamard's inequality has been obtained in [64], where a similar volume argument is used, but with respect to a hyperbolic metric. The last row of Table 2.1 shows the lower bound obtained, translated into efficiency of triangulations. For small dimensions, it is an excellent approximation of the smallest efficiency. Asymptotically, it only increases the lower bound obtained using Hadamard's inequality by a constant factor $\sqrt{3/2}$.

Let us mention that triangulations of the d-cube achieving the smallest size for $d = 2, 3, 4, 5$ can be constructed using *corner-slicing* triangulations, those which include the corners at all the vertices of $I^d = [0,1]^d$ with even (or, equivalently, odd) coordinate sums. A *corner* at a vertex v is the simplex $\mathrm{conv}(\{v, v_1, \ldots, v_d\})$, where $v_1, \ldots, v_d$ are the neighbors of v. The case $d = 2$ is trivial, and that of $d = 3$ can be checked in Figure 1.3 (left). However, the smallest size of corner-slicing triangulations does not coincide with ϕ_d for $d = 6$; Hughes [45] has proved the smallest size of a corner-slicing triangulation of I^6 to be 324.

In this chapter we propose a method to obtain efficient triangulations of a product polytope $P \times Q$ starting from a triangulation of Q and another of $P \times \Delta^{m-1}$, where Δ^{m-1} denotes a simplex of dimension $m-1$ and m is any relatively small number. We apply this with P being a small-dimensional cube and Q a high-dimensional one, iteratively. This allows to obtain asymptotically efficient triangulations of arbitrarily high-dimensional cubes from any (efficient) triangulation of $I^l \times \Delta^{m-1}$. Sections 2.3, 2.4, and 2.5 explain our method, which is first outlined in Section 2.2. If the reader is happy with dissections, we cannot offer better efficiencies for them than for triangulations, but at least he or she can skip Section 2.5, which contains most of the technicalities in this chapter.

The asymptotic efficiency of the triangulations obtained, clearly, depends on how good our triangulation of $I^l \times \Delta^{m-1}$ is. Finding the triangulation of $I^l \times \Delta^{m-1}$ which is optimal for our purposes reduces to an integer programming problem, similar to finding the smallest triangulation of that polytope (actually, it is the same system of equations with a different objective function). Using the linear programming software `CPLEX`, we have solved the system for some values of l and m. The best triangulation we have found is one of $I^3 \times \Delta^2$, with which we obtain

$$\lim_{i\to\infty} \rho_i \leq \sqrt[3]{\frac{44/3}{27}} \approx 0.8159$$

The best bound existing before was $\lim_{i\to\infty} \rho_i \leq \rho_7 = 0.840$. In other words, we prove that (asymptotically) the d-cube can be triangulated with $0.8159^d d!$ simplices, instead of the $0.840^d d!$ achievable before.

It has to be observed that, even if the particular triangulation of $I^3 \times \Delta^2$ that we use was obtained by an intensive computer calculation, once the triangulation is found it is a simple task to check that it is indeed a triangulation and compute the asymptotic efficiency obtained from it. Actually, in Section 2.7 we use the so-called Cayley Trick [44] to do this checking with no need of computers at all. We also use the Cayley Trick to explore the minimum efficiency that can be obtained from the product $I^2 \times \Delta^k$ for any k. We briefly explain the trick in Section 2.6, where we also interpret our whole construction in terms of it.

2.2 Overview of the method and results

We have already described Haiman's method in Subsection 1.1.2. Ours is, in a way, similar: Starting from a triangulation of I^{n-1} and another one of $I^l \times \Delta^{m-1}$, we get one of I^{l+n-1} with the following general method to triangulate $P \times Q$ starting from a triangulation of Q and another one of $P \times \Delta^{m-1}$:

(1) We first show (Sections 2.3 and 2.4) how to obtain triangulations of $P \times \Delta^{n-1}$ from triangulations of $P \times \Delta^{m-1}$, where $n-1 = \dim(Q)$ is supposed to be much bigger than $m-1$. We call our triangulations *multi-staircase triangulations*.

(2) A triangulation of Q induces, as in Haiman's method, a decomposition of $P \times Q$ into polytopes isomorphic to $P \times \Delta^{n-1}$. Each of them can be triangulated using the previous paragraph, although this in principle only gives a dissection of $P \times Q$; if we want a triangulation, we have to apply (1) to all the polytopes $P \times \Delta^{n-1}$ in a compatible way. In Section 2.5 we will show how to do this using an m-coloring of the vertices of Q.

(3) The analysis of the size of the triangulation obtained will also be carried out in Section 2.5.

In step (2), the final size of the triangulation is just the sum of the individual sizes of the triangulations used for the different subpolytopes $P \times \Delta^{n-1}$. In particular, if we are just interested in obtaining dissections, we can take an efficient triangulation of $P \times \Delta^{n-1}$ and repeat it in every subpolytope.

In step (1) the computation of the size is more complicated and to state it in a simple way we introduce the following definitions. When we speak about simplices in a polytope P we will implicitly suppose that its vertices are vertices of P, and we will identify the simplex with its vertex set.

Definition 2.2.1 Let P be a polytope of dimension l and let τ be a simplex of dimension $l+m-1$ in $P \times \Delta^{m-1}$. Let $\{v_1, \ldots, v_m\}$ be the vertices of Δ^{m-1}. Then, τ, understood as a vertex set, decomposes as $\tau = \tau_1 \cup \cdots \cup \tau_m$ with $\tau_i \subset P \times \{v_i\}$.

- We define the sequence $(|\tau_1| - 1, \ldots, |\tau_m| - 1)$ to be the *type* of the simplex τ. The *weight* of a simplex τ of type $(t_1, \ldots, t_m)$ is the number $\frac{1}{\prod_{i=1}^{m} t_i!}$.
- The *weighted size* of a triangulation T of $P \times \Delta^{m-1}$ is $\sum_{\tau \in T} \text{weight}(\tau)$, and the *weighted efficiency* is
$$\sqrt[l]{\frac{\sum_{\tau \in T} \text{weight}(\tau)}{m^l}}$$

With this, our main result can be stated as:

Theorem 2.2.2 *Consider polytopes P of dimension l and Q of dimension $n-1$. Let m be such that $m \leq n$. Given a triangulation T_0 of $P \times \Delta^{m-1}$ of weighted size*

t_0 and a triangulation T_Q of Q, then there are triangulations of $P \times Q$ with size at most

$$|T_Q| t_0 \left(\frac{n}{m} + l\right)^l.$$

Using this method to triangulate $I^{l+n-1} = I^l \times I^{n-1}$ starting from a triangulation of I^{n-1} and another one of $I^l \times \Delta^{m-1}$ we will conclude:

Theorem 2.2.3 *If there is a triangulation T_0 of $I^l \times \Delta^{m-1}$ with weighted efficiency ρ_0, then for every $\epsilon > 0$ and all n bigger than $lm\rho_0/\epsilon$ we have*

$$\rho_{n+l-1}{}^{n+l-1} \leq \rho_{n-1}{}^{n-1}(\rho_0 + \epsilon)^l.$$

As a consequence,

$$\lim_{n\to\infty} \rho_n \leq \rho_0.$$

Observe that the definition of weighted efficiency makes sense even in the case $m = 1$, where all the simplices of $P \times \Delta^0$ have the same type, equal to (l) if $\dim(P) = l$, and the same weight, equal to $1/l!$ In particular, the weighted efficiency of a triangulation of $I^l \times \Delta^0$ is the same as the usual "non-weighted" one. Another common point between efficiency and weighted efficiency is that the weighted efficiency of a triangulation of $I^l \times \Delta^{m-1}$ is always less or equal to 1, and it is 1 if and only if the triangulation is *unimodular*; i.e., if every simplex has volume $1/(l+m-1)!$ (this is proved in Section 2.6).

We now describe the practical results obtained: The last theorem leads us to study the smallest weighted efficiency of triangulations of $I^l \times \Delta^{m-1}$; let us denote it by $\rho_{l,m}$. This number can be calculated minimizing a linear form over the so-called *universal polytope* of all the triangulations of $I^l \times \Delta^{m-1}$.

The definition of this universal polytope for triangulations of an arbitrary polytope P is as follows (see [26] for further details): Let $\Sigma(P)$ be the set of all the simplices of maximal dimension which use as vertices only vertices of P. Given a triangulation T of P, its *incidence vector* $V_T \in \mathbb{R}^{\Sigma(P)}$ has a 1 in the coordinates corresponding to simplices of T and a 0 in the others. The *universal polytope* of P is $\text{conv}\{V_T : T \text{ is a triangulation of } P\}$.

In our case, we want to minimize the weighted efficiency, for what it suffices to calculate the smallest weighted size. Then, we have to minimize over the universal polytope of $I^l \times \Delta^{m-1}$ the linear form having as coefficient of each simplex its weight $1/\prod_{i=1}^m t_i!$ (instead of the one with all coefficients equal to 1, which gives the minimum size triangulation). One of the results in [26] is a description of the vertices of the universal polytope as integer solutions of a certain system of linear inequalities derived from the oriented matroid of P. Thus, our minimization problem is restated as an integer linear programming problem.

In order to apply our method, we have used the program `UNIVERSAL BUILDER` by Jesús A. de Loera and Samuel Peterson [27]. Given as input the vertices of P, the program generates the linear system of equations defining the universal polytope of

P. The output is a file readable by the linear programming software `CPLEX`. We have created a routine that generates our particular objective function. Table 2.2 shows the results obtained in the cases we have been able to solve. Note that $\rho_{l,1} = \rho_l$ and so the column $m = 1$ in Table 2.2 is the same as the second row of Table 2.1.

The fact that $\rho_{1,m} = 1$ for every m reflects that every triangulation of the prism $I \times \Delta^{m-1}$ is unimodular. In the case of $\rho_{2,m}$ we prove (Subsection 2.7.1) that the smallest weighted efficiency is always $\sqrt{\lceil 3m^2/4 \rceil / m^2}$. That is to say, $\sqrt{3/4}$ for even m and $\sqrt{3/4} + \Theta(m^{-2})$ for odd m.

The computation of $\rho_{3,3}$ involved a system with 74 400 variables, whose resolution by `CPLEX` took 37 hours of CPU on a SUN UltraSparc.

$l \setminus m$	1	2	3	≥ 3
1	1	1	1	1
2	1	$\sqrt{\frac{3}{4}} \approx 0.866$	$\sqrt{\frac{7}{9}} \approx 0.8819$	$\geq \sqrt{\frac{3}{4}}$
3	$\sqrt[3]{\frac{5/6}{1}} \approx 0.941$	$\sqrt[3]{\frac{14/3}{8}} \approx 0.8355$	$\sqrt[3]{\frac{44/3}{27}} \approx 0.8159$	

Table 2.2: Values of $\rho_{l,m}$ for $l \leq 2$ or for $l = 3$ and $m \leq 3$.

2.3 Polyhedral subdivision of $P \times \Delta^{k_1+\cdots+k_m-1}$ induced by one of $P \times \Delta^{m-1}$

From Subsection 1.1.1, the *polyhedral subdivisions* of a polytope P are its face-to-face partitions into subpolytopes which only use vertices of P as vertices.

Let P be a polytope of dimension l, let m and $k_1, \ldots, k_m$ be natural numbers and let us call $n := k_1 + \cdots + k_m$. Let $v_1, \ldots, v_m$ be the m vertices of the standard simplex Δ^{m-1} and $v_1^1, \ldots, v_{k_1}^1, \ldots\ldots, v_1^m, \ldots, v_{k_m}^m$ the vertices of Δ^{n-1}. Observe that, implicitly, we have the following surjective map:

$$\begin{array}{ccc} \mathrm{vert}(\Delta^{n-1}) & \to & \mathrm{vert}(\Delta^{m-1}) \\ v_j^i & \mapsto & v_i \end{array}$$

This map uniquely extends to an affine projection $\pi_0 : \Delta^{n-1} \to \Delta^{m-1}$. In turn, this induces a projection

$$\begin{array}{rccc} \pi = 1 \times \pi_0 : P \times \Delta^{n-1} & \to & P \times \Delta^{m-1} \\ (p, a) & \mapsto & (p, \pi_0(a)) \end{array}$$

Given the projection π and a polyhedral subdivision S of the target polytope $P \times \Delta^{m-1}$, it is obvious that the inverse images $\pi^{-1}(B)$ of the subpolytopes of S form a subdivision of $P \times \Delta^{n-1}$ into subpolytopes matching face to face. In a more general projection it would not be true that those subpolytopes use as vertices only vertices of $P \times \Delta^{n-1}$. But it is true in our case:

Lemma 2.3.1 *Let $B \subset P \times \Delta^{m-1}$ be a subpolytope with vert$(B) \subset$ vert$(P \times \Delta^{m-1})$. Let π be the projection considered before. Let be $\tilde{B} := \{(p, v_j^i) : (p, v_i) \in \text{vert}(B)\}$. Then:*

$$\pi^{-1}(B) \subset \text{conv}(\tilde{B}).$$

Proof: Let (p, a) be a point of $\pi^{-1}(B)$, so $(p, \pi_0(a)) \in B$. Let us write a as a convex combination of the vertices of Δ^{n-1}, that is $a = \sum_{i=1}^{m} \sum_{j=1}^{k_i} \lambda_j^i v_j^i$, with $\lambda_j^i \geq 0$, $\forall i, j$ and $\sum_i \sum_j \lambda_j^i = 1$. Without loss of generality, we can suppose that none of the sums $\sum_{j=1}^{k_i} \lambda_j^i$ is zero, otherwise everything "happens" on a face of Δ^{m-1} and we can restrict the statement to that face.

Let $(p_1^1, v_1), \ldots, (p_{l_1}^1, v_1)$, $(p_1^2, v_2), \ldots, (p_{l_2}^2, v_2)$, …, $(p_1^m, v_m), \ldots, (p_{l_m}^m, v_m)$ be the vertices of B, so we have $\tilde{B} = \{(p_h^i, v_j^i) : i = 1, \ldots, m, j = 1, \ldots, k_i, h = 1, \ldots, l_i\}$. We write now $(p, \pi_0(a))$ as convex combination of the vertices of B, that is,

$$(p, \pi_0(a)) = \sum_{i=1}^{m} \sum_{h=1}^{l_i} \mu_h^i \, (p_h^i, v_i),$$

with $\mu_h^i \geq 0$, $\forall i, h$, and $\sum_i \sum_h \mu_h^i = 1$. Observe that $\sum_{j=1}^{k_i} \lambda_j^i = \sum_{h=1}^{l_i} \mu_h^i$ for every i, because $\pi_0(a) = \sum_{i=1}^{m} \sum_{j=1}^{k_i} \lambda_j^i v_i$. Then, it is easy to check that:

$$(p, a) = \sum_{i=1}^{m} \sum_{j=1}^{k_i} \sum_{h=1}^{l_i} \frac{\lambda_j^i \mu_h^i}{\sum_{j=1}^{k_i} \lambda_j^i} (p_h^i, v_j^i) \qquad \text{and} \qquad 1 = \sum_{i=1}^{m} \sum_{j=1}^{k_i} \sum_{h=1}^{l_i} \frac{\lambda_j^i \mu_h^i}{\sum_{j=1}^{k_i} \lambda_j^i}$$

That is, (p, a) is a convex combination of points of $\tilde{B}$, as we wanted to prove. □

Corollary 2.3.2 *Under the previous conditions;*

(i) $\pi^{-1}(B) = \text{conv}(\pi^{-1}(\text{vert}(B)))$.

(ii) $\text{vert}(\pi^{-1}(B)) = \tilde{B}$.

Proof: In both equalities, the inclusion from right to left is trivial. In the second one, observe that if (p, v_j^i) is in $\tilde{B}$, then (p, v_i) is a vertex of B. Therefore, $(p, v_j^i) \in \pi^{-1}(B)$ and it is a vertex of $P \times \Delta^{n-1}$, which implies that it is a vertex of $\pi^{-1}(B)$.

Inclusions from left to right follow from Lemma 2.3.1, because:
(1) $\tilde{B} \subset \pi^{-1}(\text{vert}(B)) \Rightarrow \pi^{-1}(B) \subset \text{conv}(\tilde{B}) \subset \text{conv}(\pi^{-1}(\text{vert}(B)))$.
(2) $\pi^{-1}(B) \subset \text{conv}(\tilde{B}) \Rightarrow \text{vert}(\pi^{-1}(B)) \subset \text{vert}(\text{conv}(\tilde{B})) = \tilde{B}$, where the last equality follows from the fact that all the elements of $\tilde{B}$ are vertices of $P \times \Delta^{n-1}$. □

Corollary 2.3.3 *Every polyhedral subdivision S of $P \times \Delta^{m-1}$ induces a polyhedral subdivision*

$$\tilde{S} := \pi^{-1}(S) = \{\pi^{-1}(B) : B \in S\}$$

of $P \times \Delta^{n-1}$. Furthermore, the vertices of each subpolytope $\pi^{-1}(B)$ in this subdivision are $\tilde{B} = \{(p, v_j^i) : (p, v_i) \in \text{vert}(B)\}$. □

2.4 Triangulation of $P \times \Delta^{k_1+\cdots+k_m-1}$ induced by one of $P \times \Delta^{m-1}$

We will suppose now that the polyhedral subdivision S of $P \times \Delta^{m-1}$ is a triangulation. As in Subsection 1.1.2, a convenient graphic representation of the vertices of the polytope $P \times \Delta^{m-1}$ is as a grid whose rows represent vertices of P and whose m columns represent the vertices $v_1, \ldots, v_m$ of Δ^{m-1}. In order to represent a subset of vertices of $P \times \Delta^{m-1}$ we just mark the corresponding squares. In the same way we can represent the vertices of $P \times \Delta^{k_1+\cdots+k_m-1}$, but now it is convenient to divide the grid horizontally in blocks, each of them corresponding to each k_i and containing the vertices $v_1^i, \ldots, v_{k_i}^i$ of $\Delta^{k_1+\cdots+k_m-1}$.

Figure 2.1 shows how to obtain, with the notation of the previous section, the set $\tilde{B}$ associated to a simplex $B \in S$ in this graphic representation. In the i-th block, rows corresponding to vertices (p, v_i) in B have all its squares marked. Restricting $\tilde{B}$ to that block gives precisely $\mathrm{conv}(\{p_1^i, \ldots, p_{l_i}^i\}) \times \mathrm{conv}(\{v_1^i, \ldots, v_{k_i}^i\})$, with the notation used for the vertices of B in the proof of Lemma 2.3.1.

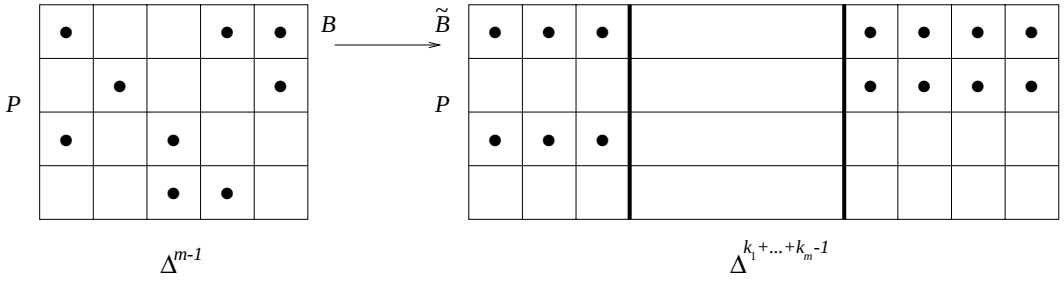

Figure 2.1: How to obtain $\tilde{B}$ from B.

Since B is a simplex, any subset of its vertices forms also a simplex; thus, the restriction of $\tilde{B}$ to each block is a product of two simplices. Let us recall (Subsection 1.1.2) that the *staircase triangulation* of the product of two simplices $\Delta^k \times \Delta^l$ is the one whose $\binom{k+l}{k}$ simplices are all the possible monotone staircases in a grid of size $(k+1) \times (l+1)$. Following this analogy, we define *multi-staircases* as follows:

Definition 2.4.1 Let $\tilde{B} = \pi^{-1}(B)$ be a subpolytope of $P \times \Delta^{k_1+\cdots+k_m-1}$, of the kind obtained in Section 2.3.

- A *multi-staircase* in $\tilde{B}$ is any subset of vertices which restricted to every block forms a monotone staircase.
- The *multi-staircase triangulation* of $\tilde{B}$ is the one which has as simplices the different multi-staircases (see Figure 2.2).

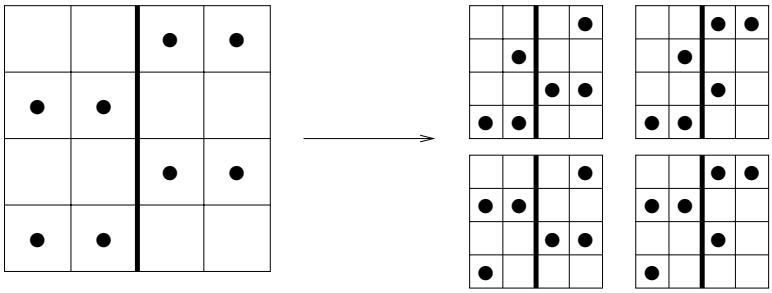

Figure 2.2: The four multi-staircases (right) forming the multi-staircase triangulation of the polytope $\tilde{B}$ in the left.

Lemma 2.4.2 *The multi-staircases indeed form a triangulation of $\tilde{B}$ and taking the multi-staircase triangulations of the different $\tilde{B}$'s obtained from a triangulation of $P \times \Delta^{m-1}$ we get a triangulation of $P \times \Delta^{k_1+\cdots+k_m-1}$, which we call* multi-staircase triangulation.

Proof: It is clear that multi-staircases form full-dimensional simplices in $\tilde{B}$. A way to prove that a collection T of full-dimensional simplices in a polytope $\tilde{B}$ is a triangulation is to show that:

1. They induce a triangulation on one face of $\tilde{B}$.
2. For each full-dimensional simplex in the collection, the removal of any single vertex produces a codimension one simplex either
   - lying on a facet of $\tilde{B}$ and not contained in any other simplex of T, or
   - contained in exactly another full-dimensional simplex of T which is separated from the first one by their common facet.

In our case, the first condition follows by induction on $k_1 + \cdots + k_m$, the base case being $k_1 + \cdots + k_m = m$.

For the second condition, let σ be a multi-staircase, and let (p, v) be a vertex in it. Then, one of the following three things happens (Figure 2.3 gives an example of a multi-staircase):

- If (p, v) is the only point of σ in its column, then no other multi-staircase contains $\sigma \setminus \{(p, v)\}$. Removing (p, v) produces indeed a codimension one simplex contained in a facet $P \times \Delta^{k_1+\cdots+k_m-2}$ of $P \times \Delta^{k_1+\cdots+k_m-1}$.
- If (p, v) is the only point of σ in a row within a block, then no other multi-staircase contains $\sigma \setminus \{(p, v)\}$. Removing (p, v) produces indeed a codimension one simplex in a facet of $\tilde{B}$ of the form $\pi^{-1}(B \setminus \{(p, \pi_0(v))\})$ (remember that $B = \pi(\tilde{B})$ is a simplex).

- If (p, v) is an elbow in the multi-staircase, then removing it leads to a unique different way of completing the multi-staircase. More precisely, let (p', v) and (p, v') be the points of the multi-staircase adjacent to (p, v). Then removing (p, v) and inserting (p', v') produces the other possible multi-staircase. The obvious affine dependency $(p, v) + (p', v') = (p, v') + (p', v)$ implies that the two multi-staircases lie in opposite sides of their common facet.

□

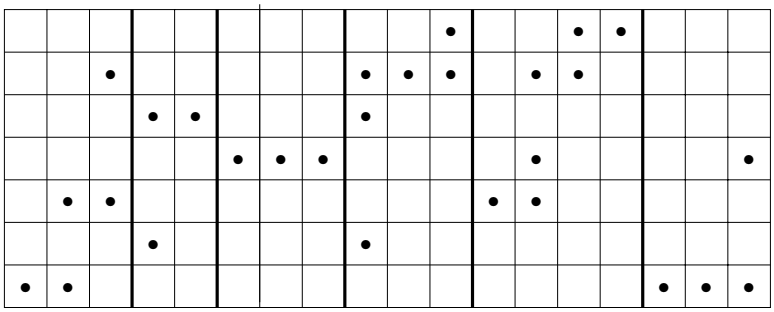

Figure 2.3: An example of a multi-staircase.

Lemma 2.4.3 *Let l_i be the number of vertices of B in the i-th column. Then the multi-staircase triangulation of $\tilde{B}$ has exactly $\prod_{i=1}^{m} \binom{k_i-1+l_i-1}{k_i-1}$ simplices.*

Proof: In each sub-block there are $\binom{k_i-1+l_i-1}{k_i-1}$ possible monotone staircases. □

2.5 A triangulation of $P \times Q$

We will consider polytopes P of dimension l and Q of dimension $n-1$. We assume we are given a triangulation T_Q of Q, which induces a decomposition of $P \times Q$ into cells isomorphic to $P \times \Delta^{n-1}$. Then, any decomposition $n := k_1 + \cdots + k_m$ allows us to apply the procedure of Sections 2.3 and 2.4 to triangulate the cells $P \times \Delta^{n-1}$ starting from a triangulation T_0 of $P \times \Delta^{m-1}$.

There are two important tasks remaining; first, show that the triangulations of the different $P \times \Delta^{n-1}$ can be achieved in a coordinated way to obtain a real triangulation of $P \times Q$ and then, analyze the efficiency of that triangulation. Both of them will be done in this section, using the following trick:

Consider a partition of the vertices of Q into m "colors". Then, in each subpolytope $P \times \Delta^{n-1}$ of $P \times Q$ the vertices of the factor Δ^{n-1} are colored as well, and we use this coloring to construct the projections $\Delta^{n-1} \to \Delta^{m-1}$ we need for each of

them. Then, on each common face (isomorphic to $P \times \Delta^k, k < n-1$) of two of the cells $P \times \Delta^{n-1}$ we get the multi-staircase triangulation induced by the m-coloring of Q restricted to that face. Hence, the triangulations of the different cells $P \times \Delta^{n-1}$ intersect face-to-face, and we obtain a triangulation $T_{P\times Q}$ of $P \times Q$, which we call *multi-staircase triangulation* of $P \times Q$.

In order to analyze the size of $T_{P\times Q}$, we will suppose that the m-coloring of the vertices of Q is chosen at random with a uniform distribution.

For each $\sigma \in T_Q$ let $\sigma_i := \{\text{vertices of } \sigma \text{ colored } i\}$. And for each $\tau \in T_0$, $\tau_i := \{\text{vertices of } \tau \text{ over the } i\text{-th vertex of } \Delta^{m-1}\}$. By Lemma 2.4.3, the triangulation $T_{P\times Q}$ we obtain has

$$\sum_{\sigma\in T_Q}\sum_{\tau\in T_0}\prod_{i=1}^{m}\binom{|\sigma_i|-1+|\tau_i|-1}{|\tau_i|-1}$$

simplices. The expected value of this sum, when the coloring is random, equals the sum of the expected values. So let us fix a pair of simplices $\tau \in T_0$ and $\sigma \in T_Q$ and calculate the expected value of

$$\prod_{i=1}^{m}\binom{|\sigma_i|-1+|\tau_i|-1}{|\tau_i|-1}$$

We call $l_i := |\tau_i|$ and $k_i := |\sigma_i|$. The l_i's are considered constants, while the k_i's are random variables depending on the coloring. They follow a multinomial distribution, in which the probability of the m-tuple $(k_1, \ldots, k_m)$ is $P(k_1, \ldots, k_m) = \frac{n!}{k_1!\cdots k_m! m^n}$.

Since $\binom{x}{n} = x^{\underline{n}}/n!$, where $x^{\underline{n}} := x(x-1)\cdots(x-n+1)$ is the n-th *falling power* of x, we can write:

$$E(\prod_{i=1}^{m}\binom{k_i-1+l_i-1}{l_i-1}) = \frac{E(\prod_{i=1}^{m}(k_i-1+l_i-1)^{\underline{l_i-1}})}{\prod_{i=1}^{m}(l_i-1)!}$$

For the numerator in the RHS of this equation we will use the following result from [83] (Theorem 4.4.4):

Theorem 2.5.1 *Let $X_1, \ldots, X_m$ be scalar random variables. Then, the expected value of $\prod_{i=1}^{m} X_i^{\underline{v_i}}$ is given by the formal calculus*

$$E(\prod_{i=1}^{m} X_i^{\underline{v_i}}) = \left[\ \frac{\partial^{v_1+\cdots+v_m}}{\partial z_1^{v_1}\cdots\partial z_m^{v_m}} E(\prod_{i=1}^{m} z_i^{X_i})\ \right]_{z_1=\cdots=z_m=1}$$

in which extra variables $z_i,\ i = 1, \ldots, m$ appear with formal purposes.

Lemma 2.5.2 *Let $l_1, \ldots, l_m$ be positive integers with $\sum l_i = l+m$ and let $k_1, \ldots, k_m$ be random variables obeying a multinomial distribution with $\sum k_i = n$. Then,*

$$E(\prod_{i=1}^{m}(k_i-1+l_i-1)^{\underline{l_i-1}}) \le \left(l+\frac{n}{m}\right)^{l}$$

Proof: If some l_i equals 1 we can neglect it in the statement; we remove it and will still have $\sum(l_i-1)=l$. Hence, we will assume $l_i \geq 2$ for every i.

We will use the following formula for expectations under a multinomial distribution, again taken from [83] (p. 53, last formula in the proof of Theorem 4.2.1):

$$E(\prod_{i=1}^{m} z_i^{k_i}) = \left(\sum_{i=1}^{m} \frac{1}{m} z_i\right)^n$$

Since $\sum_{i=1}^{m}(l_i-1)=l$ and only the k_i depend on the random process, applying Theorem 2.5.1 to the random variables $X_i := k_i + l_i - 2$ and using the equality above we get:

$$E(\prod_{i=1}^{m}(k_i-1+l_i-1)^{\underline{l_i-1}}) = \left[\ \frac{\partial^l}{\partial {z_1}^{l_1-1}\cdots\partial {z_m}^{l_m-1}} E(\textstyle\prod_{i=1}^{m} z_i^{k_i-1+l_i-1})\ \right]_{\overline{z}=1} =$$

$$= \left[\ \frac{\partial^l}{\partial {z_1}^{l_1-1}\cdots\partial {z_m}^{l_m-1}} \left(\textstyle\prod_{i=1}^{m} z_i^{l_i-2} E(\prod_{i=1}^{m} z_i^{k_i})\right)\ \right]_{\overline{z}=1} =$$

$$= \left[\ \frac{\partial^l}{\partial {z_1}^{l_1-1}\cdots\partial {z_m}^{l_m-1}} \left(\textstyle\prod_{i=1}^{m} z_i^{l_i-2} \left(\sum_{i=1}^{m} \frac{1}{m} z_i\right)^n\right)\ \right]_{\overline{z}=1} \overset{(*)}{\leq} \prod_{i=1}^{m}\left(l_i-2+\frac{n}{m}\right)^{l_i-1} \leq$$

$$\leq \prod_{i=1}^{m}\left(l+\frac{n}{m}\right)^{l_i-1} = \left(l+\frac{n}{m}\right)^{l},$$

where the last inequality comes from $l_i \leq l+1$ (since $\sum(l_i-1)=l$ and $l_i-1\geq 1$), and the one marked with an asterisc needs to be proved. For this we call $F_{l_1,\ldots,l_m}(\overline{z}) := \prod_{i=1}^{m} z_i^{l_i-2}$ and $G_n(\overline{z}) := \left(\sum_{i=1}^{m} \frac{1}{m} z_i\right)^n$. Using that

$$\frac{\partial^k}{\partial x^k} f(x)g(x) = \sum_{j=0}^{k} \binom{k}{j} \frac{\partial^j}{\partial x^j} f(x) \frac{\partial^{k-j}}{\partial x^{k-j}} g(x)$$

we come up to

$$\frac{\partial^{l_i-1}}{\partial {z_i}^{l_i-1}} \left(F_{l_1,\ldots,l_m}(\overline{z})G_n(\overline{z})\right) = \sum_{j=0}^{l_i-1} \binom{l_i-1}{j} (l_i-2)^{\underline{l_i-1-j}} \left(\frac{1}{m}\right)^j n^{\underline{j}}\, F_{l_1,\ldots,j+1,\ldots,l_m}(\overline{z}) G_{n-j}(\overline{z})$$

Therefore, since

$$\frac{\partial^l}{\partial {z_1}^{l_1-1}\cdots\partial {z_m}^{l_m-1}} h(\overline{z}) = \frac{\partial^{l_m-1}}{\partial {z_m}^{l_m-1}} \left(\cdots\left(\frac{\partial^{l_1-1}}{\partial {z_1}^{l_1-1}}\left(h(\overline{z})\right)\right)\cdots\right)$$

and

$$[F_{h_1,\ldots,h_m}(\overline{z})]_{\overline{z}=1} = [G_t(\overline{z})]_{\overline{z}=1} = 1,$$

we get

$$\left[\frac{\partial^l}{\partial {z_1}^{l_1-1}\cdots\partial {z_m}^{l_m-1}} F_{l_1,\ldots,l_m}(\overline{z})G_n(\overline{z})\right]_{\overline{z}=1} =$$

$$= \sum_{j_1=0}^{l_1-1} \cdots \sum_{j_m=0}^{l_m-1} \left(\prod_{i=1}^{m} \binom{l_i-1}{j_i} (l_i-2)^{\underline{l_i-1-j_i}} \left(\frac{1}{m}\right)^{j_i} (n-j_1-\cdots-j_{i-1})^{\underline{j_i}} \right) =$$

$$= \sum_{j_1=0}^{l_1-1} \cdots \sum_{j_m=0}^{l_m-1} \left(\prod_{i=1}^{m} \binom{l_i-1}{j_i} (l_i-2)^{\underline{l_i-1-j_i}} \left(\frac{1}{m}\right)^{j_i} \right) n^{\underline{\sum j_i}} \leq$$

$$\leq \sum_{j_1=0}^{l_1-1} \cdots \sum_{j_m=0}^{l_m-1} \left(\prod_{i=1}^{m} \binom{l_i-1}{j_i} (l_i-2)^{l_i-1-j_i} \left(\frac{1}{m}\right)^{j_i} \right) n^{\sum j_i} =$$

$$= \prod_{i=1}^{m} \left(\sum_{j_i=0}^{l_i-1} \binom{l_i-1}{j_i} (l_i-2)^{l_i-1-j_i} \left(\frac{1}{m}\right)^{j_i} n^{j_i} \right) = \prod_{i=1}^{m} \left(l_i - 2 + \frac{n}{m} \right)^{l_i-1},$$

as wanted. □

This lemma is crucial to prove the two theorems announced in Section 2.2. Indeed, Theorem 2.2.2 is just a version of the following more precise statement:

Lemma 2.5.3 *Consider polytopes P of dimension l and Q of dimension $n-1$. Let m be such that $m \leq n$. Given a triangulation T_0 of $P \times \Delta^{m-1}$ of weighted size t_0 and a triangulation T_Q of Q, the expected size of the multi-staircase triangulation $T_{P\times Q}$ of $P \times Q$ is bounded above by*

$$|T_Q| t_0 \left(\frac{n}{m} + l \right)^l.$$

Proof: Lemma 2.5.2 implies that

$$E\left(\prod_{i=1}^{m} \binom{k_i - 1 + l_i - 1}{l_i - 1} \right) = E\left(\prod_{i=1}^{m} \frac{(k_i - 1 + l_i - 1)^{\underline{l_i-1}}}{(l_i-1)!} \right) \leq \frac{1}{\prod_{i=1}^{m}(l_i-1)!} \left(l + \frac{n}{m} \right)^l.$$

Hence:

$$E\left(\sum_{\sigma \in T_Q} \sum_{\tau \in T_0} \prod_{i=1}^{m} \binom{k_i - 1 + l_i - 1}{l_i - 1} \right) \leq \sum_{\sigma \in T_Q} \sum_{\tau \in T_0} \frac{1}{\prod_{i=1}^{m}(l_i-1)!} \left(\frac{n}{m} + l \right)^l =$$

$$= |T_Q| \left(\sum_{\tau \in T_0} \frac{1}{\prod_{i=1}^{m}(l_i-1)!} \right) \left(\frac{n}{m} + l \right)^l = |T_Q| t_0 \left(\frac{n}{m} + l \right)^l.$$

□

Proof of Theorem 2.2.3: Let $t_0 = \rho_0^l m^l$ be the weighted size of the triangulation in the statement. For any $n \geq \frac{lm\rho_0}{\epsilon}$:

$$\rho_0 + \epsilon \geq \rho_0 \frac{m}{n} \left(\frac{n}{m} + l \right),$$

$$n^l\,(\rho_0+\epsilon)^l \geq \rho_0^l m^l \left(\frac{n}{m}+l\right)^l = t_0\left(\frac{n}{m}+l\right)^l,$$

$$\frac{(n+l-1)!}{(n-1)!}\,(\rho_0+\epsilon)^l \geq t_0\left(\frac{n}{m}+l\right)^l.$$

Theorem 2.2.2, with $Q = I^{n-1}$ and $P = I^l$, tells us that

$$\phi_{n+l-1} \leq \phi_{n-1} t_0 \left(\frac{n}{m}+l\right)^l,$$

or, in other words,

$$\frac{(n+l-1)!{\rho_{n+l-1}}^{n+l-1}}{(n-1)!{\rho_{n-1}}^{n-1}} \leq t_0\left(\frac{n}{m}+l\right)^l.$$

So, we have proved the first part of the theorem:

$$\forall \epsilon > 0, \quad \forall n \geq \frac{lm\rho_0}{\epsilon}, \quad {\rho_{n+l-1}}^{n+l-1} \leq {\rho_{n-1}}^{n-1}(\rho_0+\epsilon)^l.$$

The second part of the statement follows from the first one with arguments similar to those of Corollary 1.1.21. Recursively, we have that:

$$\forall \epsilon > 0, \quad \forall i \in \mathbb{N}, \quad \forall n \geq \frac{lm\rho_0}{\epsilon}, \quad \rho_{n+il} \leq {\rho_n}^{\frac{n}{n+il}}(\rho_0+\epsilon)^{\frac{il}{n+il}}.$$

This implies that for the given l and any fixed n, taking $\epsilon = \frac{lm\rho_0}{n}$ we get

$$\lim_{i\to\infty} \rho_{n+il} \leq \rho_0 + \frac{lm\rho_0}{n}.$$

That is, the sequence of indices congruent to n modulo l has limit bounded by the right hand side. Since we can make n as big as we want and the rest of the right-hand side are constants, the l subsequences of indices modulo l have limit bounded by ρ_0, hence the limit of the whole sequence has this bound. □

2.6 Interpretation of our method via the Cayley Trick

The Cayley Trick allows to study triangulations of a product $P \times \Delta^{n-1}$ as mixed subdivisions of the Minkowski sum $P + \cdots + P$ (n summands). We overview here this method, but the reader should look at [44] for more details.

Let $Q_1, \ldots, Q_m \subset \mathbb{R}^d$ be convex polytopes of vertex sets $\mathcal{A}_i$. Consider their Minkowski sum, defined as

$$\sum_{i=1}^{m} Q_i = \{\, x_1 + \cdots + x_m : x_i \in Q_i \,\}.$$

We understand $\sum_{i=1}^m Q_i$ as a *marked polytope*, whose associated point configuration is $\sum_{i=1}^m \mathcal{A}_i := \{q_1 + \cdots + q_m : q_i \in \mathcal{A}_i\}$. Here a *marked polytope* is a pair $(P, \mathcal{A})$ where P is a polytope and $\mathcal{A}$ is a finite set of points of P including all the vertices. *Subdivisions* of a marked polytope are defined in [36, Chapter 7] (sometimes they are called subdivisions of $\mathcal{A}$). Roughly speaking, they are the polyhedral subdivisions of P which use only elements of $\mathcal{A}$ as vertices (but perhaps not all of them). They form a poset under the refinement relation. The minimal elements are the triangulations of $\mathcal{A}$.

A subset B of $\sum_{i=1}^m \mathcal{A}_i$ is called *mixed* if $B = B_1 + \cdots + B_m$ for some non-empty subsets $B_i \subset \mathcal{A}_i$, $i = 1, \ldots, m$. A *mixed subdivision* of $\sum_{i=1}^m Q_i$ is a subdivision of it whose cells are all mixed. Mixed subdivisions form a subposet of the poset of all subdivisions, whose minimal elements are called *fine mixed*, in which every mixed cell is *fine*, i.e., does not properly contain any other mixed cell.

We call *Cayley embedding* of $\{Q_1, \ldots, Q_m\}$ the marked polytope

$$(\mathcal{C}(Q_1, \ldots, Q_m), \mathcal{C}(\mathcal{A}_1, \ldots, \mathcal{A}_m)) \text{ in } \mathbb{R}^d \times \mathbb{R}^{m-1}$$

defined as follows: Let $e_1, \ldots, e_m$ denote an affine basis in $\mathbb{R}^{m-1}$ and $\mu_i : \mathbb{R}^d \to \mathbb{R}^d \times \mathbb{R}^{m-1}$ be the inclusion given by $\mu_i(x) = (x, e_i)$. Then we define

$$\mathcal{C}(\mathcal{A}_1, \ldots, \mathcal{A}_m) := \cup_{i=1}^m \mu_i(\mathcal{A}_i), \qquad \mathcal{C}(Q_1, \ldots, Q_m) := \operatorname{conv}(\mathcal{C}(\mathcal{A}_1, \ldots, \mathcal{A}_m))$$

Each Q_i is naturally embedded as a face in $\mathcal{C}(Q_1, \ldots, Q_m)$. Moreover, the vertex set of $\mathcal{C}(Q_1, \ldots, Q_m)$ is the disjoint union of the vertices of all the Q_i's. This induces the following bijection between cells in $\mathcal{C}(Q_1, \ldots, Q_m)$ and mixed cells in $Q_1 + \cdots + Q_m$: To each mixed cell $B_1 + \cdots + B_m$ we associate the disjoint union $B_1 \cup \cdots \cup B_m$. To a cell B in $\mathcal{C}(Q_1, \ldots, Q_m)$, we associate the Minkowski sum $B_1 + \cdots + B_m$, where $B_i = B \cap Q_i$.

Theorem 2.6.1 (The Cayley Trick,[44]) *Let $Q_1, \ldots, Q_m \subset \mathbb{R}^d$ be convex polytopes. The bijection just exhibited induces an isomorphism between the poset of all subdivisions of $\mathcal{C}(Q_1, \ldots, Q_m)$ and the poset of mixed subdivisions of $\sum_{i=1}^m Q_i$. In this isomorphism triangulations correspond to fine mixed subdivisions.*

Remark 2.6.2 With the previous definitions, for any polytope P:

$$P \times \Delta^{m-1} = \mathcal{C}(P, \overset{m}{\ldots}, P).$$

In particular, in our context the Cayley Trick provides the following bijections between triangulations and mixed subdivisions:

$$\begin{array}{ccc} \boxed{\text{Triangulation of } P \times \Delta^{m-1}} & & \boxed{\text{Triangulation of } P \times \Delta^{n-1}} \\ \updownarrow \text{ Cayley Trick} & & \updownarrow \text{ Cayley Trick} \\ \boxed{\text{Fine mixed subdivision of } P + \overset{m}{\cdots} + P} & \longrightarrow & \boxed{\text{Fine mixed subdivision of } P + \overset{n}{\cdots} + P} \end{array}$$

Our interest in the Cayley Trick is two-fold. On the one hand, it provides a way to visualize our candidate triangulations of $I^l \times \Delta^{m-1}$ as objects in dimension l, instead of $l+m-1$. We will use this in Section 2.7.

On the other hand, the construction of the previous sections has a simple geometric interpretation in terms of the Cayley Trick. More precisely, the polyhedral subdivision of Corollary 2.3.3 can be obtained as follows: Let S be a polyhedral subdivision of $P \times \Delta^{m-1}$, and S_M the corresponding mixed subdivision of $P+ \overset{m}{\cdots} +P$. Each cell in S decomposes uniquely as

$$\cup_{i=1}^{m} B_i \times \{v_i\},$$

where the B_i are subsets of vertices of P. The corresponding cell in S_M is just $B_1 + \cdots + B_m$. To construct our polyhedral subdivision of $P \times \Delta^{n-1}$ we just need to scale each summand B_i by the integer k_i which tells us how many vertices of Δ^{n-1} correspond to the vertex v_i of Δ^{m-1}.

That is to say, from S_M we construct a mixed subdivision of $P+ \overset{n}{\cdots} +P$ by the formula

$$\widetilde{S_M} = \{B_1+ \overset{k_1}{\cdots} +B_1 + \cdots\cdots + B_m+ \overset{k_m}{\cdots} +B_m : B_1 + \cdots + B_m \in S_M\}.$$

The polyhedral subdivision $\tilde{S}$ of $P \times \Delta^{n-1}$ stated in Corollary 2.3.3 is the one corresponding via the Cayley Trick to the mixed subdivision $\widetilde{S_M}$ of $P+ \overset{n}{\cdots} +P$.

Also the type and weight of a simplex in $P \times \Delta^{m-1}$ have a simple interpretation via the Cayley Trick. With the notation of Definition 2.2.1, let $\tau = \tau_1 \cup \cdots \cup \tau_m$ be a simplex. The corresponding cell τ_M in $P+ \overset{m}{\cdots} +P$ is the Minkowski sum of the simplices $\tau_1, \ldots, \tau_m$, which lie in complementary affine subspaces. Hence, τ_M is combinatorially a product of m simplices, of dimensions $t_1, \ldots, t_m$ where $(t_1, \ldots, t_m)$ is the type of τ. Then the weight of τ represents the volume of τ_M, normalized with respect to the unit parallelepiped in the lattice spanned by the vertices of τ_M. With this we can prove:

Proposition 2.6.3 *Let P be a lattice polytope of dimension l. Let V be its volume, normalized to the unit parallelepiped in the lattice. Then, the weighted size of a triangulation of $P \times \Delta^{m-1}$ is at most $m^l V$, with equality if and only if the triangulation is unimodular (with respect to the lattice).*

In particular, the weighted efficiency of a triangulation of $I^l \times \Delta^{m-1}$ is at most one, with equality for unimodular triangulations.

Proof: If P is a lattice polytope (e.g., a cube), then the lattice spanned by τ_M is a sublattice of the one spanned by the point configuration $P+ \overset{m}{\cdots} +P$, and coincides with it if and only if τ is unimodular. Then, for unimodular triangulations the weighted size is then just the volume of $P+ \overset{m}{\cdots} +P$, normalized to the unit parallelepiped, which equals $m^l V$. For non-unimodular triangulations the weighted size is smaller than that. □

2.7 Computation of $\rho_{l,m}$

In this section we obtain the value of $\rho_{2,m}$ for any m and we will also show triangulations of $I^3 \times \Delta^1$ and $I^3 \times \Delta^2$ providing the values of $\rho_{3,2}$ and $\rho_{3,3}$ stated in Table 2.2.

2.7.1 Smallest weighted efficiency of triangulations of $I^2 \times \Delta^{m-1}$

Here we prove:

Theorem 2.7.1 *The smallest weighted efficiency $\rho_{2,m}$ of triangulations of $I^2 \times \Delta^{m-1}$ is*

$$\rho_{2,m} = \sqrt{\frac{\lceil 3m^2/4 \rceil}{m^2}}.$$

That is, $\sqrt{3/4}$ for even m and $\sqrt{3/4} + \Theta(m^{-2})$ for odd m. A fine mixed subdivision of $I^2 + \overset{m}{\cdots} + I^2$ corresponding to a triangulation of $I^2 \times \Delta^{m-1}$ with that weighted efficiency is given in Figure 2.4.

Let $B_M = B_1 + \overset{m}{\cdots} + B_m$ be a cell in a fine mixed subdivision of $I^2 + \overset{m}{\cdots} + I^2$ (a square of size m). Each B_i must be a simplex coming from the i-th copy of I^2, and in order to have a *fine* mixed subdivision, the different B_i's must lie in complementary affine subspaces. Then, there are the following possibilities:

- B_M is a triangle, obtained as the sum of a triangle in one of the I^l's and a single point in the others. The weight of B_M is $1/2$.
- B_M is a quadrangle, obtained as the sum of two (non-parallel) segments from two of the B_i's and a point in the rest of them. Three types of quadrangles can appear, depending on whether both, one or none of the two segments involved is a diagonal of I^2: a diagonal square (of area 2), a rhomboid, and a square parallel to the axes, of area 1. The weight of B_M is 1.

In particular, the weighted size of the mixed subdivision equals $T/2 + S_1 + S_2 + R$, where T, S_1, S_2 and R denote, respectively, the numbers of triangles, squares of area 1, squares of area 2 and rhombi in the subdivision. Since the total area of $I^2 + \overset{m}{\cdots} + I^2$ is $m^2 = T/2 + S_1 + 2S_2 + R$, we conclude that:

Proposition 2.7.2 *The weighted size of a mixed subdivision of $I^2 + \overset{m}{\cdots} + I^2$ equals $m^2 - S_2$, where S_2 is the number of squares of size 2 in the subdivision.*

With this we can already conclude that the fine mixed subdivisions shown in Figure 2.4 (one for m even and one for m odd) have weighted size equal to $\lceil 3m^2/4 \rceil$, since they have exactly $\lfloor m^2/4 \rfloor$ squares of area 2.

Our task is now to prove that no mixed subdivision can have more than $\lfloor m^2/4 \rfloor$ squares of size 2. For this we use:

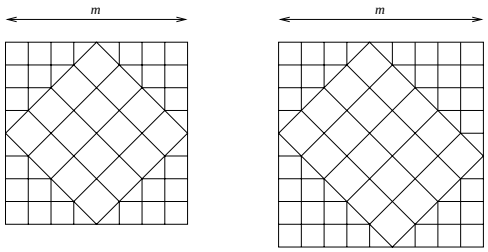

Figure 2.4: Fine mixed subdivisions corresponding to triangulations of $I^2 \times \Delta^{m-1}$ with smallest weighted efficiency for m even (left) or odd (right).

Lemma 2.7.3 *Taking the i-th summands of all the mixed cells $B_1 + \overset{m}{\cdots} + B_m$ in a mixed subdivision of $P_1 + \overset{m}{\cdots} + P_m$ produces a polyhedral subdivision of P_i. If the mixed subdivision was fine, the polyhedral subdivision is a triangulation.*

Proof: By the Cayley Trick, the mixed subdivision of $P_1 + \overset{m}{\cdots} + P_m$ induces a polyhedral subdivision of $\mathcal{C}(P_1, \ldots, P_m)$, and a triangulation if the mixed subdivision is fine. Since the polytope P_i appears as a face in $\mathcal{C}(P_1, \ldots, P_m)$, any subdivision (resp. triangulation) of $\mathcal{C}(P_1, \ldots, P_m)$ induces a subdivision (resp. triangulation) of P_i. That this subdivision is the one obtained taking the i-th summands of all the mixed cells follows from Theorem 2.6.1. □

Proposition 2.7.4 *Let S_M be a fine mixed subdivision of $I^2 + \overset{m}{\cdots} + I^2$. Let a and b be the number of summands I^2's which are triangulated in one and the other possible triangulations of I^2 (i.e., using one or the other diagonal). Then, the number of squares of area 2 in S_M is at most ab and the weighted size of S_M is at least $m^2 - ab$.*

Proof: Each square of area 2 is the Minkowski sum of two opposite diagonals of two copies of I^2 (say, the ith and jth copies), and a point in the other $m-2$ of the copies of I^2. Clearly, the two copies which contribute diagonals have to be triangulated in opposite ways. The only thing which remains to be shown is that the same pair of copies of I^2 cannot contribute two different squares of size 2. For this, observe that "contracting" in every mixed cell of a mixed subdivision all the sumands other than the ith and jth should give a mixed subdivision of $P_i + P_j$. And in a mixed subdivision of two squares there is no room to put two different diagonal squares of area 2. □

In the statement of Proposition 2.7.4, we have that $a + b = m$. In particular, the maximum possible value of ab is $\lfloor m^2/4 \rfloor$. This finishes the proof of Theorem 2.7.1.

2.7.2 Smallest weighted efficiency of triangulations of $I^3 \times \Delta^{m-1}$, for $m = 2, 3$

Here we will try to visualize triangulations of $I^3 \times \Delta^1$ and $I^3 \times \Delta^2$ which minimize the weighted efficiency, which were computed using the integer programming approach sketched in Section 2.2.

Of course, we use the Cayley Trick to decrease the dimension, so we show the corresponding fine mixed subdivision of a Minkowski sum instead of the triangulation itself.

- $\boxed{I^3 \times \Delta^1}$ $\qquad I^3 \times \Delta^1 = \mathcal{C}(I^3, I^3) \longleftrightarrow I^3 + I^3$

  We have to give a subdivision of a 3-cube of size 2. For this we first cut the eight corners of the cube, producing a *cubeoctahedron*, a semi-regular 3-polytope with 6 square and 8 triangular facets. The edges of the cubeoctahedron can be distributed in four "equatorial hexagons" each of which cuts the polytope into 2 halves. Figure 2.5 depicts these two halves for one of the equatorial hexagons. The labels in the vertices are heights, interpreted as follows: Our point configuration is $\{0, 1, 2\}^3$ and the height of point (i, j, k) is just $i + j + k$.

  It turns out that performing three of these four possible halvings, the cubeoctahedron is decomposed into six triangular prisms and two tetrahedra. In Figure 2.5 each half is actually decomposed into three prisms and one tetrahedron. These, together with the eight tetrahedra we have cut from corners, form a fine mixed subdivision of $\{0, 1, 2\}^3$ with 10 tetrahedra and 6 triangular prisms. Its weighted size is then

  $$10 \cdot \frac{1}{6} + 6 \cdot \frac{1}{2} = \frac{14}{3}$$

  and its weighted efficiency $\sqrt[3]{\frac{14/3}{8}}$.

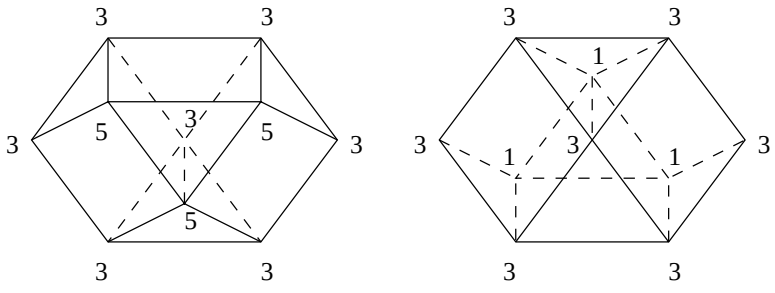

Figure 2.5: Fine mixed subdivision of $I^3 + I^3$ corresponding to a triangulation of $I^3 \times \Delta^1$ with smallest weighted efficiency.

- $\boxed{I^3 \times \Delta^2}$ $\qquad I^3 \times \Delta^2 = \mathcal{C}(I^3, I^3, I^3) \longleftrightarrow I^3 + I^3 + I^3$

  The fine mixed subdivision of $I^3 + I^3 + I^3$ is given in Figure 2.6. It consists of 20 triangular prisms, 16 tetrahedra and 2 parallelepipeds.

  Thus, the weighted size of the triangulation is $20\frac{1}{2} + 16\frac{1}{6} + 2\frac{1}{1} = \frac{44}{3}$, and then the smallest weighted efficiency is

$$\rho_{3,3} = \sqrt[3]{\frac{44/3}{3^3}}.$$

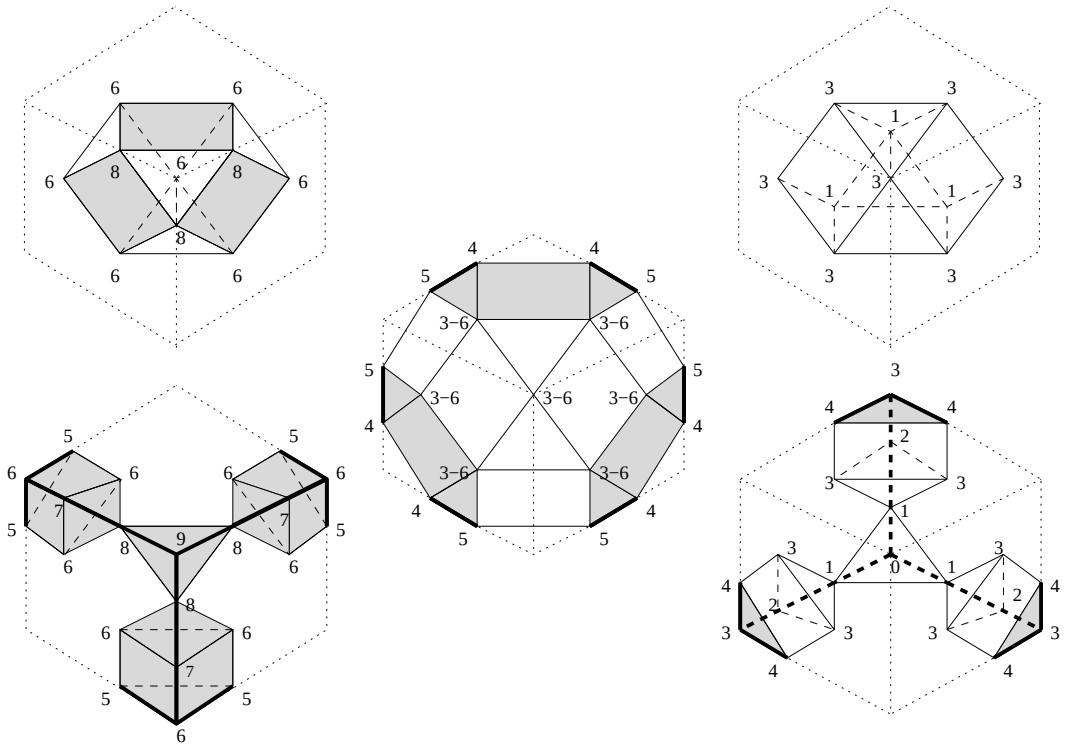

Figure 2.6: Fine mixed subdivision of $I^3 + I^3 + I^3$ corresponding to a triangulation of $I^3 \times \Delta^2$ with smallest weighted efficiency.

Let us explain how to interpret Figure 2.6. Again, each point (i, j, k) in $\{0, 1, 2, 3\}^3$ has been given a height $i + j + k$, which is written next to it. The subdivision is displayed in five parts. The left-top portion in the figure is a half cubeoctahedron exactly as the one in Figure 2.5. The reader has to assume it divided into a tetrahedron and three triangular prisms, as before. The left-bottom portion consists of a corner tetrahedron and three triangular prisms located at corners of the cube, each joined to the half cubeoctahedron by a tetrahedron. So far we have six prisms and five tetrahedra, and the same is got from the right-top and right-bottom portions of the Figure.

In between the two half cubeoctahedra, however, we have now a hexagonal prism decomposed into two triangular prisms and two quadrilateral prisms, all of height $\sqrt{3}$. The hexagonal prism is surrounded by a belt formed by six triangular prisms and six tetrahedra.

The thick edges in Figure 2.6 represent edges of the big 3-cube of side 3, whose three visible facets are drawn by dotted lines. We have also shaded the facets of the subdivision which are contained in those facets of the big 3-cube.

Chapter 3

The polytope of non-crossing graphs on a planar point set

In the whole chapter we focus on *geometric* graphs, that is to say, graphs drawn with straight-line edges in the euclidean plane. We will usually refer to them just as "graphs". For any finite set $\mathcal{A}$ of n points in $\mathbb{R}^2$, we define a $(3n-3)$-dimensional simple polyhedron whose face poset is isomorphic to the poset of "non-crossing marked graphs" with vertex set $\mathcal{A}$, where a marked graph is defined as a geometric graph together with a subset of its vertices. The poset of non-crossing graphs on $\mathcal{A}$ appears as the complement of the star of a face in that polyhedron.

The polyhedron has a unique maximal bounded face, of dimension $2n_i + n - 3$ where n_i is the number of points of $\mathcal{A}$ in the interior of $\text{conv}(\mathcal{A})$. The vertices of this polytope are all the pseudo-triangulations of $\mathcal{A}$, and the edges are flips of two types: the traditional diagonal flips (in pseudo-triangulations) and the removal or insertion of a single edge.

As a by-product of our construction we prove that all pseudo-triangulations are infinitesimally rigid graphs.

3.1 Introduction

Consider a finite point set $\mathcal{A} \subset \mathbb{R}^2$. In this paragraph and in the rest of the chapter n is the total cardinality of $\mathcal{A}$ and n_i denotes the number of points in the interior of $\text{conv}(\mathcal{A})$. For point sets in non-general position we will need to distinguish between vertices of $\text{conv}(\mathcal{A})$ (*extremal* points) and other points in the boundary of $\text{conv}(\mathcal{A})$, which we call *semi-interior* points. We denote n_v and n_s their respective numbers, so that $n = n_i + n_v + n_s$.

We have already noted in Section 1.2 that the poset structure of non-crossing graphs was only well understood if the points were in convex position, when the non-crossing graphs containing all the hull edges are the same as the polygonal subdivisions of the convex n-gon and form the face poset of the $(n-3)$-*associahedron*. In this chapter we generalize the associahedron and construct from $\mathcal{A}$ a polytope

whose face poset contains the poset of non-crossing graphs on $\mathcal{A}$ embedded in a very nice way:

Theorem 3.1.1 *Let $\mathcal{A}$ be a finite set of n points in the plane, not all contained in a line. Let n_i, n_s and n_v be the number of interior, semi-interior and extremal points of $\mathcal{A}$, respectively.*

There is a simple polytope $Y_f(\mathcal{A})$ of dimension $2n_i + n - 3$, and a face F of $Y_f(\mathcal{A})$ (of dimension $2n_i + n_v - 3$) such that the complement of the star of F in the face-poset of $Y_f(\mathcal{A})$ equals the poset of non-crossing graphs on $\mathcal{A}$ that use all the convex hull edges.

This statement deserves some words of explanation:

- Since convex hull edges are irrelevant to crossingness, it is trivial to obtain the poset of *all* non-crossing graphs on $\mathcal{A}$ as the direct product of the poset in the statement and a Boolean poset of rank n_v.
- The equality of posets in Theorem 3.1.1 reverses inclusions. Maximal non-crossing graphs (triangulations of $\mathcal{A}$) correspond to minimal faces (vertices of $Y_f(\mathcal{A})$).
- We remind the reader that the *star* of a face F is the set of faces contained in the union of all the facets (maximal proper faces) containing F. In the complement of the star of F we must include the polytope $Y_f(\mathcal{A})$ itself, which corresponds to the graph with no interior edges.
- If $\mathcal{A}$ is in convex position then $Y_f(\mathcal{A})$ is the associahedron and the face F is the whole polytope, whose star we must interpret as being empty.
- We give a fully explicit facet description of $Y_f(\mathcal{A})$. It lives in $\mathbb{R}^{3n}$ and is defined by the 3 linear equalities (3.1) and the $\binom{n}{2}+n$ linear inequalities (3.4) and (3.5) of Section 3.3, with some of them turned into equalities. (With one exception: for technical reasons, if $\mathcal{A}$ contains three collinear boundary points we need to add extra points to its exterior and obtain $Y_f(\mathcal{A})$ as a face of the polytope $Y_f(\mathcal{A}')$ of the extended point set $\mathcal{A}'$).
- The f_{ij}'s in equations (3.4) and (3.5) and in the notation $Y_f(\mathcal{A})$ denote a vector in $\mathbb{R}^{\binom{n+1}{2}}$. Our construction starts with a linear cone $\overline{Y_0}(\mathcal{A})$ (Definition 3.3.1) whose facets are then translated using the entries of f to produce a polyhedron $\overline{Y_f}(\mathcal{A})$, of which the polytope $Y_f(\mathcal{A})$ is the unique maximal bounded face. Our proof goes by analyzing the necessary and sufficient conditions for f to produce a polytope with the desired properties and then proving the existence of valid choices of f. In particular, Theorem 3.3.7 shows one valid choice. This is essentially the same approach used in [61] for the polytope of pointed non-crossing graphs constructed there. That polytope is actually the face F of the statement of Theorem 3.1.1.

- Our results are valid for point sets in non-general position, an aspect which was left open in [61]. Our definition of non-crossing in non-general position is that if q is between p and r, then the edge pr cannot appear in a non-crossing graph, regardless of whether q is incident to any edge in the graph. This definition has the slight drawback that a graph which is non-crossing in a point set $\mathcal{A}$ may become crossing in $\mathcal{A} \cup \{p\}$, but it has the advantage that it makes maximal crossing-free graphs coincide with the triangulations of $\mathcal{A}$ (with the convention, standard in Computational Geometry, that triangulations of $\mathcal{A}$ are required to use all the points of $\mathcal{A}$ as vertices, recall Remark 1.1.12).

Mapping the instances of a combinatorial object to the vertices of a certain polytope is useful for optimization and enumeration purposes [7]. In the case of triangulations, two such polytopes have been used in the past; the so-called "secondary polytope" [14] and "universal polytope" [26] of the point set. Our construction adds to these two, but it has one advantage: We have an explicit facet description of the polytope. In the secondary polytope, facets correspond to the coarse polygonal subdivisions of $\mathcal{A}$, which have no easy characterization. In the universal polytope, the facet description in [26] gives only a linear programming relaxation of the polytope, which implies that integer programming is needed in order to optimize linear functionals on it.

That the poset in which we are interested is not the whole poset of faces of the polytope $Y_f(\mathcal{A})$ may seem a serious drawback for using it as a tool for enumeration of all the triangulations of a planar point set. But the subposet we are interested in is not just a subposet. Being the complement of the star of a face F has theoretical and practical implications: On the one hand, it implies that the poset is homeomorphic to a ball of dimension $2n_i + n - 4$, since there is a shelling order ending precisely in the facets that contain F. On the other hand, the part of the boundary of $Y_f(\mathcal{A})$ that we are interested in becomes the (strict) lower envelope of a convex polyhedron via any projective transformation that sends a supporting hyperplane of the face F to the infinity.

Observe also that the dimension $2n_i + n - 3$ of the polytope we construct is the minimum possible one because it equals the number of interior edges in every triangulation.

And, actually, the set of *all* the vertices of the polytope $Y_f(\mathcal{A})$ is interesting on its own merits:

Theorem 3.1.2 *The vertex set of the polytope $Y_f(\mathcal{A})$ of Theorem 3.1.1 is in bijection to the set of all pseudo-triangulations of $\mathcal{A}$. The 1-skeleton of $Y_f(\mathcal{A})$ is the graph of flips between them.*

Streinu [69] introduced the *minimum* or *pointed pseudo-triangulations* (p.p.t.'s, for short). In the same paper, she proved certain surprising relations of them to structural rigidity of non-crossing graphs, using them to prove the Carpenter's Rule Theorem (the first proof of which was given shortly before by Connelly et al. [22]). Recall from Remarks 1.4.8 that pointed pseudo-triangulations (p.p.t.'s, for short)

turn out to coincide with the maximal non-crossing and pointed graphs. The paper [61], using these rigid theoretic properties of p.p.t.'s, constructed a polytope $X_f(\mathcal{A})$ whose vertex set is precisely the set of p.p.t.'s of a given point set. Our construction is heavily based on that one, but this chapter is mostly self-contained. Our method not only extends that construction to cover all pseudo-triangulations, it also shows that the method can easily be adapted to point sets in non-general position, with a suitable definition of pseudo-triangulation for such point sets (Definition 3.5.1).

The flips between pseudo-triangulations that appear in our Theorem 3.1.2 are introduced in Section 3.2 (see Definition 3.2.6 and Figure 3.2) for point sets in general position and in Section 3.5 (see Definition 3.5.5 and Figures 3.7 and 3.8) for point sets with collinearities. The definition is new (to the best of our knowledge) but in the case of general position it has independently been considered in [4], where flips between pseudo-triangulations are related to geometric flips between polyhedral terrains.

Our flips restrict to the ones in [69] and [61] when the two pseudo-triangulations involved are pointed, and they are also related to the flips of Pocchiola and Vegter [57] as follows: Recall from Subsection 1.4.1 that Pocchiola and Vegter were interested in pseudo-triangulations of a set $\mathcal{O} := \{o_1, \ldots, o_n\}$ of convex bodies, and they defined a graph of flips between them. That graph is regular of degree $3n-3$. Pocchiola (personal communication) has shown that taking each o_i to be a sufficiently small convex body around each point, our graph is obtained from the one in [57] by contraction of certain edges. In particular, this shows that our graph has diameter $O(n^2)$ since that is the case for the graph in [57].

Our construction has also rigid-theoretic consequences. Recall that when speaking about rigidity we mean infinitesimal rigidity, and that a generically rigid graph (for dimension 2) is a graph which becomes rigid in almost all its straight-line embeddings in the plane (Definition 1.3.16). Generically rigid graphs need at least $2n-3$ edges, because that is the number of degrees of freedom of n points in the plane (after neglecting rigid motions). As said in Section 1.3, generically rigid graphs with exactly $2n-3$ edges are called *isostatic* and they admit the following characterization, due to Laman (see, for example, [38]): they are the graphs with $2n-3$ edges and with the property that any subset of $k \leq n-2$ vertices is incident to at least $2k$ edges. We have already mentioned that, using this characterization, Ileana Streinu [69] proved every pointed pseudo-triangulations to be an isostatic graph (Corollary 1.4.10). Here we prove the following generalization:

Theorem 3.1.3 *Let T be a pseudo-triangulation of a planar point set $\mathcal{A}$ in general position. Let G be its underlying graph. Then:*

(i) G is infinitesimally rigid, hence rigid and generically rigid.

(ii) There are at least $2k+3l$ edges of T incident to any subset of k pointed plus l non-pointed vertices of T (assuming $k+l \leq n-2$).

The property in part (ii) will be called *generalized Laman property* in Chapter 4, where it will play an important role. In the statement, "general position" can actually be weakened to "no three boundary points of $\mathcal{A}$ are collinear". In the presence of boundary collinearities, non-rigid pseudo-triangulations (for our definition) exist. For example, only six of the fourteen pseudo-triangulations of the point set of Figure 3.8 are rigid.

If we recall that a pseudo-triangulation with k non-pointed vertices has exactly $2n-3+k$ edges (see Proposition 1.4.7), Theorem 3.1.3 has the consequence that the space of self-stresses on a pseudo-triangulation has exactly dimension k. This fact follows also from the results of [4].

The structure of the chapter is as follows: We first develop our construction for point sets in general position, in three steps; the combinatorics of the polyhedron we are seeking for is studied in Section 3.2, where we introduce in particular the graph of flips between pseudo-triangulations. Then, the construction of the polytope is given in Section 3.3 but the proof that it has the required properties depends on some assumption which is proved in Section 3.4 using rigidity theoretic ideas. In Section 3.5 everything is adapted to point sets with collinearities.

In closing, we propose two open questions:

Open Question 3.1.4 Can every planar and generically rigid graph be embedded as a pseudo-triangulation? That the answer is positive should be considered a conjecture of the set of authors of [39]. In Chapter 4 we solve the minimal case: planar Laman graphs can be embedded as pointed pseudo-triangulations. The maximal case (combinatorial triangulations can be drawn with convex faces) is also solved, by Tutte's theorem [73]. But the intermediate cases remain open, see Conjecture 4.5.3.

Open Question 3.1.5 Is the poset of non-crossing graphs on $\mathcal{A}$ the poset of a polyhedron? A naive answer would be that the polyhedron can be obtained by just deleting from the facet definition of the polytope $Y_f(\mathcal{A})$ of Theorem 3.1.1 the facets containing F. But we have checked that this does not work even in the simplest non trivial example; a single point in general position in the interior of a quadrilateral. In this example, F itself is a facet but its removal gives a polyhedron with two extra vertices, not present in $Y_f(\mathcal{A})$, and corresponding to graphs with crossings.

3.2 The graph of all pseudo-triangulations of $\mathcal{A}$

All throughout this section, $\mathcal{A}$ denotes a set of n points in general position in the plane, n_i of them in the interior of $\text{conv}(\mathcal{A})$ and n_v in the boundary. Recall the following definitions and proposition from Subsection 1.4.2:

Definition 3.2.1 A *pseudo-triangle* is a simple polygon with only three convex vertices (called *corners*) joined by three inward convex polygonal chains (called *pseudo-edges* of the pseudo-triangle).

A *pseudo-triangulation of* $\mathcal{A}$ is a geometric non-crossing graph with vertex set $\mathcal{A}$ and which partitions $\text{conv}(\mathcal{A})$ into pseudo-triangles.

Definition 3.2.2 A vertex of a geometric graph is called *pointed* if all its incident edges span an angle smaller than 180 degrees from that vertex. The graph itself is called *pointed* if all its vertices are pointed.

A *pointed pseudo-triangulation* is that on which every vertex is pointed.

Proposition 3.2.3 (Streinu) *Let $\mathcal{A}$ be a planar point set as above. Then, every pseudo-triangulation of $\mathcal{A}$ with n_γ non-pointed vertices and n_ϵ pointed vertices has $2n - 3 + n_\gamma = 3n - 3 - n_\epsilon$ edges.*

Let us show that pseudo-triangulations have another crucial property: The existence of a natural notion of flip. Let e be an interior edge in a pseudo-triangulation T of $\mathcal{A}$ and let σ be the union of the two pseudo-triangles incident to e. We regard σ as a graph, one of whose edges is e. One can consider $\sigma \setminus e$ to be a (perhaps degenerate) polygon, with a well-defined boundary cycle; in degenerate cases some edges and vertices may appear twice in the cycle. See an example of what we mean in Figure 3.1, in which the cycle of vertices is $pqrstsu$ and the cycle of edges is $pq, qr, rs, st, ts, su, up$. As in any polygon, each (appearance of a) vertex in the boundary cycle of $\sigma \setminus e$ is either concave or convex. In the figure, there are four convex vertices (corners), namely r, second appearance of s, u and q. Then:

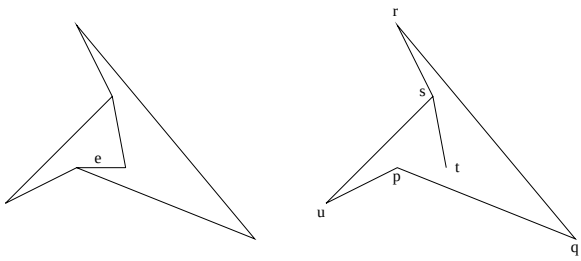

Figure 3.1: The region $\sigma \setminus e$ is a degenerate polygon with four corners

Lemma 3.2.4 *(i) $\sigma \setminus e$ has either 3 or 4 corners.*

(ii) It has 3 corners if and only if exactly one of the two end-points of e is pointed in σ. In this case $T \setminus e$ is still a pseudo-triangulation.

(iii) It has 4 corners if and only if both end-points of e are pointed in σ. In this case $T \setminus e$ is not a pseudo-triangulation and there is a unique way to insert an edge in $T \setminus e$ to obtain another pseudo-triangulation.

Proof: Let v_1 and v_2 be the two end-points of e. For each v_i, one of the following three things occur: (a) v_i is not pointed in σ, in which case it is a corner of the two

pseudo-triangles incident to e and is not a corner of $\sigma \setminus e$; (b) v_i is pointed in σ with the big angle exterior to σ, in which case it is a corner of both pseudo-triangles and of $\sigma \setminus e$ as well, or (c) v_i is pointed with its big angle interior, in which case it is a corner in only one of the two pseudo-triangles and not a corner in $\sigma \setminus e$.

In case (a), v_i contributes two more corners to the two pseudo-triangles than to $\sigma \setminus e$. In the other two cases, it contributes one more corner to the pseudo-triangles than to $\sigma \setminus e$. Since the two pseudo-triangles have six corners in total, $\sigma \setminus e$ has four, three or two corners depending on whether both, one or none of v_1 and v_2 are pointed in σ. The case of two corners is clearly impossible, which finishes the proof of (i). Part (ii) only says that "degenerate pseudo-triangles" cannot appear.

Part (iii) is equivalent to saying that a pseudo-quadrangle (even a degenerate one) can be divided into two pseudo-triangles in exactly two ways. Indeed, these two partitions are obtained drawing the geodesic arcs between two opposite corners. Such a geodesic path consists of a unique interior edge and (perhaps) some boundary edges. □

Cases (ii) and (iii) of the above lemma will define two different types of flips in a pseudo-triangulation. The inverse of the first one is the insertion of an edge, in case this keeps a pseudo-triangulation. The following statement determines exactly when this happens:

Lemma 3.2.5 *Let T be a pseudo-triangle with k non-corners. Then, every interior edge dividing T into two pseudo-triangles makes non-pointed exactly one non-corner. Moreover, there are exactly k such interior edges, each making non-pointed a different non-corner.*

Proof: The first sentence follows from Lemma 3.2.4, which says that exactly one of the two end-points of the edge inserted is pointed (after the insertion). For each non-corner, pointedness at the other end of the edge implies that the edge is the one that arises in the geodesic arc that joins that non-corner to the opposite corner. This proves uniqueness and existence. □

Definition 3.2.6 (Flips in pseudo-triangulations) Let T be a pseudo-triangulation. We call *flips* in T the following three types of operations, all producing pseudo-triangulations. See examples in Figure 3.2:

- (*Deletion flip*). The removal of an edge $e \in T$, if $T \setminus e$ is a pseudo-triangulation.
- (*Insertion flip*). The insertion of an edge $e \notin T$, if $T \cup e$ is a pseudo-triangulation.
- (*Diagonal flip*). The exchange of an edge $e \in T$, if $T \setminus e$ is not a pseudo-triangulation, for the unique edge e' such that $(T \setminus e) \cup e'$ is a pseudo-triangulation.

The *graph of pseudo-triangulations* of $\mathcal{A}$ has as vertices all the pseudo-triangulations of $\mathcal{A}$ and as edges all flips of any of the types.

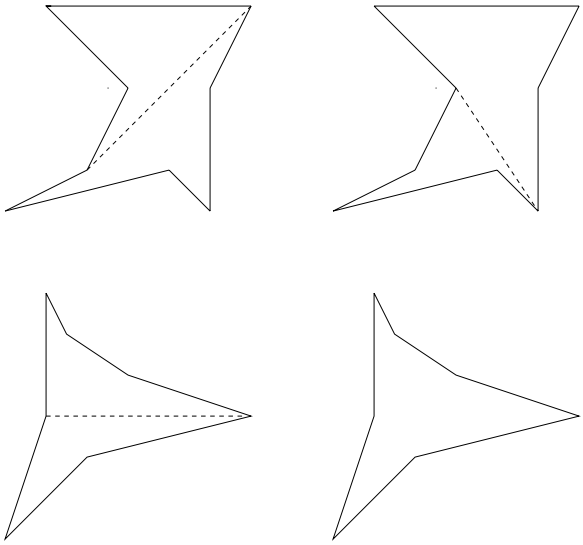

Figure 3.2: Above, a diagonal-flip. Below, an insertion-deletion flip.

Proposition 3.2.7 *The graph of pseudo-triangulations of $\mathcal{A}$ is connected and regular of degree $3n_i + n_v - 3 = 3n - 2n_v - 3$.*

Proof: There is one diagonal or deletion flip for each interior edge, giving a total of $3n - 3 - n_\epsilon - n_v$ by Proposition 3.2.3. There are as many insertion flips as pointed interior vertices by Lemma 3.2.5, giving $n_\epsilon - n_v$.

To establish connectivity, let p be a point on the convex hull of $\mathcal{A}$. The pseudo-triangulations of $\mathcal{A}$ with degree 2 at p coincide with the pseudo-triangulations of $\mathcal{A} \setminus \{p\}$ (together with the two tangents from p to $\mathcal{A} \setminus \{p\}$). By induction, we assume all those pseudo-triangulations to be connected in the graph. On the other hand, in pseudo-triangulations with degree greater than 2 at p all interior edges incident to e can be flipped and produce pseudo-triangulations with smaller degree at e. (Remark: if p is an interior point, then a diagonal-flip on an edge incident to p may create another edge incident to p; but for a boundary point this cannot be the case since p is a corner in the pseudo-quadrilateral $\sigma \setminus e$ of Lemma 3.2.4). Decreasing one by one the number of edges incident to p will eventually lead to a pseudo-triangulation with degree 2 at p. □

Remarks 3.2.8 It is an immediate consequence of Lemma 3.2.4 that every interior edge in a pointed pseudo-triangulation is flippable. As a consequence, this shows

that the graph of diagonal-flips between pointed pseudo-triangulations of $\mathcal{A}$ is regular of degree $2n_i + n_v - 3$, what was a crucial fact in [61].

As another remark, one may be tempted to think that two pseudo-triangulations are connected by a diagonal flip if and only if one is obtained from the other by the removal and insertion of a single edge, but this is not the case: The two pseudo-triangulations of Figure 3.3 are not connected by a diagonal flip, according to our definition, because the intermediate graph $T \setminus e$ is a pseudo-triangulation.

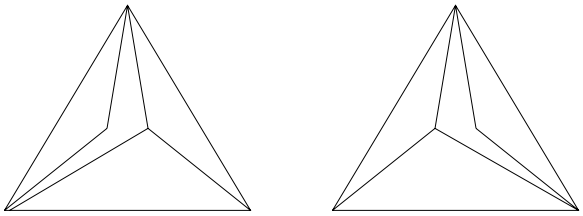

Figure 3.3: These two pseudo-triangulations are not connected by a flip.

3.2.1 Marked non-crossing graphs on $\mathcal{A}$.

As happened with pointed pseudo-triangulations, Proposition 3.2.7 suggests that the graph of pseudo-triangulations of $\mathcal{A}$ may be the skeleton of a simple polytope of dimension $3n_i + n_v - 3$. As a step towards this result we first look at what the face poset of such a polytope should be. The polytope being simple means that we want to regard each pseudo-triangulation T as the upper bound element in a Boolean poset of order $3n - 3 - 2n_v$. Since this number equals, by Proposition 3.2.3, the number of interior edges plus interior pointed vertices in T, we define:

Definition 3.2.9 A *marked graph* on $\mathcal{A}$ is a geometric graph with vertex set $\mathcal{A}$ together with a subset of its vertices, that we call "marked". We call a marked graph *non-crossing* if it is non-crossing as a graph and marks arise only in pointed vertices.

We call a non-crossing marked graph *fully-marked* if it is marked at all pointed vertices. If, in addition, it is a pseudo-triangulation, then we call it a *fully-marked pseudo-triangulation*, abbreviated as *f.m.p.t.*

Marked graphs form a poset by inclusion of both the sets of edges and of marked vertices. We say that a marked graph contains the boundary of $\mathcal{A}$ if it contains all the convex hull edges and convex hull marks. The following results follow easily from the corresponding statements for non-crossing graphs and pseudo-triangulations.

Proposition 3.2.10 *With the previous definitions:*

(i) Every marked pseudo-triangulation of $\mathcal{A}$ with n_γ non-pointed vertices, n_ϵ pointed vertices and n_m marked vertices, has $2n-3+n_\gamma+n_m = 3n-3-n_\epsilon+n_m$ edges

plus marks. In particular, all fully-marked pseudo-triangulations have $3n-3$ *edges plus marks,* $3n-3-2n_v$ *of them interior.*

(ii) Fully-marked pseudo-triangulations of $\mathcal{A}$ *are exactly the maximal non-crossing marked graphs on* $\mathcal{A}$*.*

(iii) **(Flips in marked pseudo-triangulations)** *In a fully-marked pseudo-triangulation of* $\mathcal{A}$*, every interior edge or interior mark can be* flipped*; once removed, there is a unique way to insert another edge or mark to obtain a different fully-marked pseudo-triangulation of* $\mathcal{A}$*. The graph of flips between fully-marked pseudo-triangulations of* $\mathcal{A}$ *equals the graph of pseudo-triangulations of* $\mathcal{A}$ *of Definition 3.2.6.* □

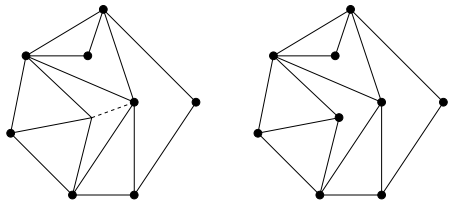

Figure 3.4: Two marked pseudo-triangulations (with marks represented by dots) related by a flip. An edge from the left is switched to a mark on the right.

These properties imply that, if the graph of pseudo-triangulations of $\mathcal{A}$ is to be the 1-skeleton of a simple polytope, then the face poset of that polytope must be (opposite to) the inclusion poset of *non-crossing marked graphs containing the boundary of* $\mathcal{A}$. Indeed, this poset has the right "1-skeleton" and the right lower ideal below every fully-marked pseudo-triangulation (a Boolean lattice of length $3n-3-2n_v$).

3.3 The polyhedron of marked non-crossing graphs on $\mathcal{A}$

In the first part of this section we do not assume $\mathcal{A}$ to be in general position. Only after Definition 3.3.6 we need general position, among other things because we have not yet defined marked non-crossing graphs or pseudo-triangulations for point sets in special position. That will be done in Section 3.5.

The setting for our construction is close to the rigid-theoretic one used in [61]. There, the polytope to be constructed is embedded in the space $\mathbb{R}^{2n-3}$ of all infinitesimal motions of the n points $p_1, \ldots, p_n$. The space has dimension $2n-3$ because the infinitesimal motion of each point produces two coordinates (an infinitesimal velocity $v_i \in \mathbb{R}^2$) but global translations and rotations produce a 3-dimensional subspace

of trivial motions which are neglected. Formally, this can be done by a quotient $\mathbb{R}^{2n}/M_0$, where M_0 is the 3-dimensional subspace of trivial motions, or it can be done by fixing three of the $2n$ coordinates to be zero. For example, if the points p_1 and p_2 do not lie in the same horizontal line, one can take

$$v_1^1 = v_1^2 = v_2^1 = 0.$$

In our approach, we will consider a third coordinate t_i for each point, related to the "marks" discussed in the previous paragraphs, or to pointedness of the vertices.

That is to say, given a set of n points $\mathcal{A} = \{p_1, \dots, p_n\}$ in $\mathbb{R}^2$, we consider the following $(3n-3)$-dimensional space;

$$S := \{(v_1, \dots, v_n, t_1, \dots, t_n) \in \left(\mathbb{R}^2\right)^n \times \mathbb{R}^n : v_1^1 = v_1^2 = v_2^1 = 0\} \subset \mathbb{R}^{3n}. \tag{3.1}$$

In it we consider the following $\binom{n}{2} + n$ linear inequalities

$$H_{ij}^+ := \{(v,t) \in S : \langle p_i - p_j, v_i - v_j\rangle - |p_i - p_j|(t_i + t_j) \geq 0\} \tag{3.2}$$

and

$$H_{0j}^+ := \{(v,t) \in S : t_j \geq 0\}. \tag{3.3}$$

We denote by $H_{i,j}$ and $H_{0,j}$ their boundary hyperplanes.

Definition 3.3.1 We call *expansion cone* of $\mathcal{A}$, and denote it $\overline{Y_0}(\mathcal{A})$, the positive region of the above hyperplane arrangement:

$$\overline{Y_0}(\mathcal{A}) := \bigcap_{i,j \in \{0,1,\dots,n\}} H_{ij}^+$$

When clear from the context we will omit the point set $\mathcal{A}$ and use just $\overline{Y_0}$.

Observe that the equations defining $\overline{Y_0}$ imply that for every i, j:

$$\langle p_i - p_j, v_i - v_j\rangle \geq |p_i - p_j|(t_i + t_j) \geq 0.$$

In particular, the vector $(v_1, \dots, v_n)$ is an expansive infinitesimal motion of the point set, as introduced after Definition 1.3.15.

Lemma 3.3.2 *The polyhedron $\overline{Y_0}(\mathcal{A})$ has full dimension $3n-3$ in $S \subset \mathbb{R}^{3n}$ and it is a pointed polyhedral cone. (Here, "pointed" means "having the origin as a vertex" or, equivalently, "containing no opposite non-zero vectors").*

Proof: The vector (v,t) with $v_i := p_i, t_i := \min_{k,l}\{|p_k - p_l|\}/4$ satisfies all the inequalities (3.2) and (3.3) strictly. In order to obtain a point in S we add to it a suitable infinitesimal trivial motion.

To prove that the cone is pointed, suppose that it contains two opposite nonzero vectors (v,t) and $-(v,t)$. Equivalently, that (v,t) lies in all the hyperplanes $H_{i,j}$ and $H_{0,i}$. That is to say, $t_i = 0$ for every i and

$$\langle v_j - v_i, p_j - p_i \rangle = 0$$

for all i, j. These last equations say that $(v_1, \ldots, v_n)$ is an infinitesimal flex of the complete graph on $\mathcal{A}$. Since the complete graph on every full-dimensional point set is, by definition, infinitesimally rigid, $(v_1, \ldots, v_n)$ is a trivial motion and equations (3.1) imply that the motion is zero. □

An edge p_ip_j or a point p_i are called *tight* for a certain vector $(v,t) \in \overline{Y_0}$ if (v,t) lies in the corresponding hyperplane $H_{i,j}$ or $H_{0,i}$. We call *supporting graph* of any (v,t), and denote it $T(v,t)$, the marked graph of tight edges for (v,t) with marks at tight points for (v,t).

Lemma 3.3.3 *Let $(v,t) \in \overline{Y_0}$. If $T(v,t)$ contains the boundary edges and vertices of a convex polygon, then $v_l = 0$ and $t_l = 0$ for every point p_l in the interior of the polygon. Therefore, $T(v,t)$ contains the complete marked graph on the set of vertices and interior points of the polygon.*

Observe that this statement says nothing about points in the relative interior of a boundary edge, if the polygon has collinear points in its boundary. Indeed, such points may have a non-zero v_l, namely the exterior normal to the boundary edge containing the point.

Proof: The hypotheses are equivalent to $t_i = 0$ and $\langle p_i - p_j, v_i - v_j \rangle = 0$ for all the boundary vertices p_i and boundary edges p_ip_j of the convex polygon. We first claim that the infinitesimal expansive motion $v = (v_1, \ldots, v_n)$ also preserves distances between non-consecutive polygon vertices. Since the sum of interior angles at vertices of an n-gon is independent of coordinates, a non-trivial motion fixing the lengths of boundary edges would decrease the interior angle at some polygon vertex p_i and then its adjacent boundary vertices get closer, what contradicts (3.2). Hence, v is a translation or rotation of the polygon boundary which, by equations (3.1), is zero. On the other hand, if $v_k \neq 0$ for any p_k interior to the polygon, then p_k gets closer to some boundary vertex, what using $t_k \geq 0$ contradicts (3.2) again.

Therefore, $v_l = 0$ for every point p_l enclosed in the polygon (what can be concluded from [61, Lemma 3.2(b)] as well). Then, the equation (3.2) corresponding to p_l and to any point p_i in the boundary of the polygon implies that $t_l \leq 0$. Together with the equation (3.3) corresponding to p_l this implies $t_l = 0$. □

Obviously, $\overline{Y_0}$ is not the polyhedron we are looking for, since its face poset does not have the desired combinatorial structure; it has a unique vertex while $\mathcal{A}$ may have more than only one fully-marked pseudo-triangulation. The right polyhedron for our purposes is going to be a convenient perturbation of $\overline{Y_0}$ obtained by translation of its facets.

Definition 3.3.4 For each $f \in \mathbb{R}^{\binom{n+1}{2}}$ (with entries indexed $f_{i,j}$, for $i, j \in \{0, \ldots, n\}$) we call *polyhedron of expansions constrained by* f, and denote it $\overline{Y_f}(\mathcal{A})$, the polyhedron defined by the $\binom{n}{2}$ equations

$$\langle p_i - p_j, v_i - v_j \rangle - |p_i - p_j|(t_i + t_j) \geq f_{ij} \tag{3.4}$$

for every $p_i, p_j \in \mathcal{A}$ and the n equations

$$t_j \geq f_{0j}, \qquad \forall p_i \in \mathcal{A}. \tag{3.5}$$

From Lemma 3.3.2, we conclude that:

Corollary 3.3.5 *$\overline{Y_f}(\mathcal{A})$ is a $(3n-3)$-dimensional unbounded polyhedron with at least one vertex, for any f.* □

In the rest of this section and in Section 3.4 we assume $\mathcal{A}$ to be in general position. As before, to each feasible point $(v, t) \in \overline{Y_f}$ we associate the marked graph consisting of edges and vertices whose equations (3.4) and (3.5) are tight on (v, t). Similarly, to a face F of $\overline{Y_f}$ we associate the tight marked graph of any of its relative interior points. This gives an (order-reversing) embedding of the face poset of $\overline{Y_f}$ into the poset of all marked graphs of $\mathcal{A}$. Our goal is to show that for certain choices of the constraint parameters f, the face poset of $\overline{Y_f}$ coincides with that of non-crossing marked graphs on $\mathcal{A}$.

Definition 3.3.6 We define a choice of the constants f to be *valid* if the tight marked graph $T(F)$ of every face F of $\overline{Y_f}$ is non-crossing.

The proof that valid choices exist for any point set is postponed to Section 3.4, in order not to interrupt the current flow of ideas. In particular, Corollary 3.4.5 implies that the following explicit choice is valid:

Theorem 3.3.7 *The choice $f_{ij} := \det(O, p_i, p_j)^2$, $f_{0j} := 0$ is valid.*

The main statement in the chapter is then:

Theorem 3.3.8 (The polyhedron of marked non-crossing graphs) *If f is a valid choice of parameters, then $\overline{Y_f}$ is a simple polyhedron of dimension $3n-3$ whose face poset equals (the opposite of) the poset of non-crossing marked graphs on $\mathcal{A}$. In particular:*

(i) Vertices of the polyhedron are in 1-to-1 correspondence with fully-marked pseudo-triangulations of $\mathcal{A}$.

(ii) Bounded edges correspond to flips of interior edges or marks in fully-marked pseudo-triangulations, i.e., to fully-marked pseudo-triangulations with one interior edge or mark removed.

(iii) Extreme rays correspond to fully-marked pseudo-triangulations with one convex hull edge or mark removed.

Proof: By Corollary 3.3.5, every vertex (v,t) of $\overline{Y_f}$ has at least $3n-3$ incident facets. By Proposition 3.2.10, if f is valid then the marked graph of any vertex of $\overline{Y_f}$ has exactly $3n-3$ edges plus marks and is a fully-marked pseudo-triangulation. This also implies that the polyhedron is simple. If we prove that all the fully-marked pseudo-triangulations appear as vertices of $\overline{Y_f}$ we finish the proof, because then the face poset of $\overline{Y_f}$ will have the right minimal elements and the right upper ideals of minimal elements (the Boolean lattices of subgraphs of fully-marked pseudo-triangulations) to coincide with the poset of non-crossing marked graphs on $\mathcal{A}$.

That all fully-marked pseudo-triangulations appear follows from connectedness of the graph of flips: Starting with any given vertex of $\overline{Y_f}$, corresponding to a certain f.m.p.t. T of $\mathcal{A}$, its $3n-3$ incident edges correspond to the removal of a single edge or mark in T. Moreover, if the edge or mark is not in the boundary, Lemma 3.3.3 implies that the edge (of $\overline{Y_f}$) corresponding to it is bounded because it collapses to the origin in $\overline{Y_0}$. Then, this edge connects the original vertex of $\overline{Y_f}$ to another one which can only be the f.m.p.t. given by the flip in the corresponding edge or mark of T. Since this happens for all vertices, and since all f.m.p.t.'s are reachable from any other one by flips, we conclude that they all appear as vertices. □

From Theorems 3.3.7 and 3.3.8 it is easy to conclude the two first statements in the introduction. The following is actually a more precise statement implying both:

Theorem 3.3.9 (The polytope of all pseudo-triangulations) *Let $Y_f(\mathcal{A})$ be the face of $\overline{Y_f}(\mathcal{A})$ defined turning into equalities the equations (3.4) and (3.5) which correspond to convex hull edges or convex hull points of $\mathcal{A}$, and assume f to be a valid choice. Then:*

(i) $Y_f(\mathcal{A})$ is a simple polytope of dimension $2n_i + n - 3$ whose 1-skeleton is the graph of pseudo-triangulations of $\mathcal{A}$. (In particular, it is the unique maximal bounded face of $\overline{Y_f}(\mathcal{A})$).

(ii) Let F be the face of $Y_f(\mathcal{A})$ defined by turning into equalities the remaining equations (3.5). Then, the complement of the star of F in the face-poset of $Y_f(\mathcal{A})$ equals the poset of non-crossing graphs on $\mathcal{A}$ that use all the convex hull edges.

Proof: (i) That $Y_f(\mathcal{A})$ is a bounded face follows from Lemma 3.3.3 (it collapses to the zero face in $\overline{Y_0}(\mathcal{A})$). Since vertices of $\overline{Y_f}(\mathcal{A})$ are f.m.p.t.'s and since all f.m.p.t.'s contain all the boundary edges and vertices, $Y_f(\mathcal{A})$ contains all the vertices of $\overline{Y_f}(\mathcal{A})$. Hence, its vertices are in bijection with all f.m.p.t.'s which, in turn, are in bijection with pseudo-triangulations. Edges of $Y_f(\mathcal{A})$ correspond to f.m.p.t.'s minus one interior edge or mark, which are precisely the flips between f.m.p.t.'s, or between pseudo-triangulations.

(ii) The facets containing F are those corresponding to marks in interior points. Then, the faces in the complement of the star of F are those in which none of the inequalities (3.5) are tight; that is to say, they form the poset of "non-crossing marked graphs containing the boundary edges and marks but no interior marks", which is the same as the poset of non-crossing graphs containing the boundary. □

We now turn our attention to Theorem 3.1.3. Its proof is based in the use of the homogeneous cone $\overline{Y_0}(\mathcal{A})$ or, more precisely, the set $\mathcal{H} := \{H_{ij} : i, j = 1, \ldots, n\} \cup \{H_{0i} : i = 1, \ldots, n\}$ of hyperplanes that define it.

Proof of Theorem 3.1.3: Observe now that the equations defining H_{ij}, when specialized to $t_i = 0$ for every i, become the equations of the infinitesimal rigidity of the complete graph on $\mathcal{A}$. In particular, a graph G is rigid on $\mathcal{A}$ if and only if the intersection

$$(\cap_{ij \in G} H_{ij}) \cap (\cap_{i=1}^{n} H_{0i})$$

equals 0.

This happens for any pseudo-triangulation because Theorem 3.3.8 implies that the hyperplanes corresponding to the $3n - 3$ edges and marks of any fully-marked pseudo-triangulation form a basis of the (dual of) the linear space S. The last part of the statement of (i) comes from Lemma 1.3.17.

To prove part (ii) we only need the fact that the $3n - 3$ linear hyperplanes corresponding to a fully-marked pseudo-triangulation are independent. In particular, any subset of them is independent too. We consider the subset corresponding to the induced (marked) subgraph on the $n - k - l$ vertices other than the k pointed and l non-pointed ones we are interested in. They form an independent set involving only $3(n - k - l)$ coordinates, hence their number is at most $3(n - k - l) - 3$ (we need to subtract 3 for the rigid motions of the $n - k - l$ points, and here is where we need $k + l \leq n - 2$). Since the fully-marked pseudo-triangulation has $3n - 3$ edges plus marks, at least $3k + 3l$ of them are incident to our subset of points. And exactly k marks are incident to our points, hence at least $2k + 3l$ edges are. □

Actually, we can derive some consequences for general planar rigid graphs. Observe that every planar and generically rigid graph G must have between $2n - 3$ and $3n - 3$ edges (the extreme cases being an isostatic graph and a triangulation of the 2-sphere). Hence, we can say that the graph G has $2n - 3 + y$ edges, where and $0 \leq y \leq n - 3$. If the graph can be embedded as a pseudo-triangulation then the embedding will have exactly y non-pointed vertices. In particular, the following statement is an indication that every planar and rigid graph can be embedded as a pseudo-triangulation (see Conjecture 4.5.3):

Proposition 3.3.10 *Let G be a planar and generically rigid graph with n vertices and $2n - 3 + y$ edges. Then, there is a subset Y of cardinality y of the vertices of G such that every set of l vertices in Y plus k vertices not in Y is incident to at least $2k + 3l$ edges, whenever $k + l \leq n - 2$.*

Proof: Consider G embedded planarly in a sufficiently generic straight-line manner. Since the embedding is plane, it can be completed to a pseudo-triangulation T. In particular, the set of edges of G represents an independent subset of $2n-3+y$ hyperplanes of $\mathcal{H}$. But since the graph is rigid, adding marks to all the vertices produces a spanning set of $3n-3+y$ hyperplanes. In between these two sets there must be a basis, consisting of the $2n-3+y$ edges of G plus $n-y$ marks. We call Y the vertices not marked in this basis, and the same argument as in the proof of Theorem 3.1.3 gives the statement. □

It has to be said however, that a planar graph G with a subset Y of its vertices satisfying Proposition 3.3.10 need not be generically rigid. Figure 3.5 shows an example (take as Y any three of the four six-valent vertices).

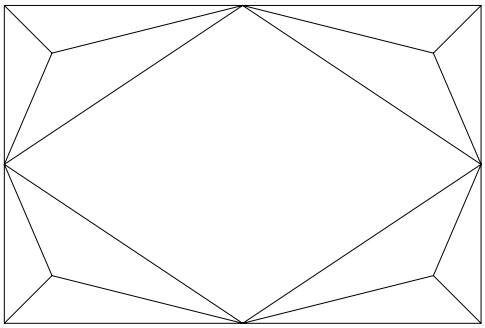

Figure 3.5: A planar graph satisfying the conclusion of Proposition 3.3.10 need not be rigid.

3.4 Valid choices of f

It remains to be proved that valid choices of parameters do exist. In particular, that the choice in Theorem 3.3.7 is valid. Our methods, again inspired on [61], give actually more; a full description of the set of valid choices via a set of $\binom{n}{4}$ linear inequalities, one for each 4-point subset of the n points.

Definition 3.4.1 Let G be a graph embedded on $\mathcal{A}$, with set of edges E and set of marked vertices V. In our context, a *stress* is an assignment of scalars w_{ij} to edges and α_j to marked vertices of G, such that for every $(v,t) \in \mathbb{R}^{3n}$:

$$\sum_{ij \in E} w_{ij}(\langle p_i - p_j, v_i - v_j \rangle - |p_i - p_j|(t_i + t_j)) + \sum_{i \in V} \alpha_i t_i = 0 \tag{3.6}$$

Lemma 3.4.2 *Let $\sum_{i=1}^{n} \lambda_i p_i = 0$, $\sum \lambda_i = 0$, be an affine dependence on a point set $\mathcal{A} = \{p_1, \ldots, p_n\}$. Then,*

$$w_{ij} := \lambda_i \lambda_j \text{ for every } i,j$$

and

$$\alpha_i := \sum_{j:ij\in E} \lambda_i\lambda_j|p_i - p_j| \text{ for every } i$$

defines a stress of the complete graph G on $\mathcal{A}$.

Proof: The condition (3.6) on variables v gives

$$\sum_{ij\in E} w_{ij}\langle p_i - p_j, v_i - v_j\rangle = 0, \text{ for every } v \in \mathbb{R}^{2n} \tag{3.7}$$

what can be equivalently stated as saying that the w_{ij} form a stress, in the sense of Definition 1.3.6, on the underlying (unmarked) graph of G. This is fulfilled by the w_{ij}'s of the statement:

$$\sum_{j\neq i} \lambda_i\lambda_j(p_i - p_j) = \sum_{j=1}^{n} \lambda_i\lambda_j(p_i - p_j) = \lambda_i p_i \sum_{j=1}^{n} \lambda_j - \lambda_i \sum_{j=1}^{n} \lambda_j p_j = 0$$

where last equality comes from λ_i's being an affine dependence. Then, cancellation of the coefficient of t_i in equation (3.6) is equivalent to $\alpha_i = \sum_{j:ij\in E} w_{ij}|p_i - p_j|$. □

Let us consider the case of four points in general position in $\mathbb{R}^2$, which have a unique (up to constants) affine dependence. The coefficients of this dependence are:

$$\lambda_i = (-1)^i \det([p_1, \ldots, p_4]\backslash\{p_i\})$$

Hence, if we divide the w_{ij}'s and α_j's of the previous lemma by the constant

$$-\det(p_1, p_2, p_3)\det(p_1, p_2, p_4)\det(p_1, p_3, p_4)\det(p_2, p_3, p_4)$$

we obtain the following expressions:

$$w_{ij} = \frac{1}{\det(p_i, p_j, p_k)\det(p_i, p_j, p_l)},\ \alpha_i = \sum_{j:ij\in E} w_{ij}|p_i - p_j| \tag{3.8}$$

where, in that of $w_{i,j}$, k and l denote the two indices other than i and j. The reason why we perform the previous rescaling is that the expressions obtained in this way have a key property which will turn out to be fundamental later on, see Figure 3.6:

Lemma 3.4.3 *For any four points in general position, the previous expressions give positive w_{ij} and α_j on boundary edges and points and negative w_{ij} and α_j on interior edges and points.*

Proof: In order to check the part concerning w_{ij}'s we use that $\det(q_1, q_2, q_3)$ is two times the signed area of the triangle spanned by q_1, q_2, q_3: For a boundary edge the two remaining points lie on the same side of the edge, so they have the same sign. For an interior edge, they lie on opposite sides and therefore they have different signs.

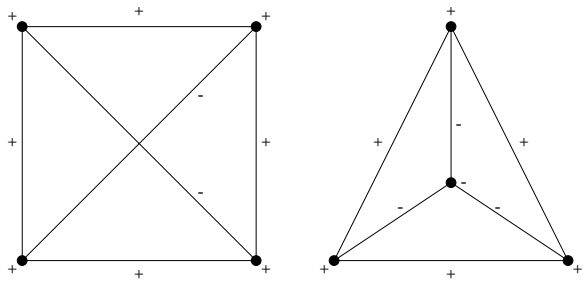

Figure 3.6: The negative parts of these two marked graphs are the excluded minors in non-crossing marked graphs of a point set in general position

For the α_i's, if i is an interior point then all the $w_{i,j}$'s in the formula for α_i are negative and, hence, α_i is also negative. If i is a boundary point then two of the $w_{i,j}$ are positive and the third one is negative. But, since

$$\sum_{j\in\{1,2,3,4\}\setminus i} w_{i,j}(p_i - p_j) = 0,$$

the triangle inequality implies that the two positive summands $w_{ij}|p_i - p_j|$ in the expression of α_i add up to a greater absolute value than the negative one. Hence α_i is positive. □

The previous statement is crucial to us, because no matter whether the four points are in convex position or one of them is inside the convex hull of the other three, the fully-marked pseudo-triangulations on the four points can be characterized as the marked graphs with nine edges plus marks and in which the missing edge or mark is interior (two f.m.p.t's for points in convex position, four of them for a triangle plus an interior point).

We conclude that:

Theorem 3.4.4 *An $f \in \mathbb{R}^{\binom{n+1}{2}}$ is valid if and only if for every four points $\{p_1, p_2, p_3, p_4\}$ of $\mathcal{A}$ the following inequality holds,*

$$\sum_{1\le i<j\le 4} w_{ij} f_{ij} + \sum_{j=1}^{4} \alpha_j f_{0j} > 0 \tag{3.9}$$

where the w_{ij}'s and α_j's are those of (3.8).

Proof: Suppose first that $\mathcal{A}$ has only four points. The polyhedron $\overline{Y_f}(\mathcal{A})$ is nine-dimensional, what implies that for every vertex (v, t) of the polyhedron, the set $T(v, t)$ contains at least nine edges plus marks on those four points. Therefore, $T(v, t)$ is the complete marked graph with an edge or mark removed.

Let us denote by G_k and G_{kl} the complete marked graph with a non-marked vertex k or a missing edge kl, respectively. Recall that by Lemma 3.4.3 the choice of stress on four points has the property that G_k and G_{kl} are fully-marked pseudo-triangulations if and only if α_k and w_{kl} (corresponding respectively to the removed mark or edge) are negative. Let us see that this is equivalent to f being valid:

By the definition of stress (Definition 3.4.1),

$$\sum_{1\le i<j\le 4} w_{ij}(\langle p_i - p_j, v_i - v_j\rangle - |p_i - p_j|(t_i + t_j)) + \sum_{j=1}^{4} \alpha_j t_j$$

equals zero. In the case of G_k, in which every edge and vertex except k are tight, that expression equals

$$\sum_{1\le i<j\le 4} w_{ij} f_{ij} + \sum_{j=1}^{4} \alpha_j f_{0j} + \alpha_k(t_k - f_{0k}).$$

In the case of G_{kl}, where every vertex and edge except kl are tight, it equals

$$\sum_{1\le i<j\le 4} w_{ij} f_{ij} + \sum_{j=1}^{4} \alpha_j f_{0j} + w_{kl}(\langle p_k - p_l, v_k - v_l\rangle - |p_k - p_l|(t_k + t_l) - f_{kl}).$$

Since $\langle p_k - p_l, v_k - v_l\rangle - |p_k - p_l|(t_k + t_l) - f_{kl} \ge 0$ and $t_k - f_{0k} \ge 0$, by (3.4) and (3.5), we conclude that in the first and second cases above, α_k and w_{kl} respectively are negative if, and only if, f is valid.

Now we turn to the case of a general $\mathcal{A}$ and our task is to prove that a choice of parameters f is valid if and only if it is valid when restricted to any four points. Observe that if $\mathcal{A}' \subset \mathcal{A}$ then $Y_f(\mathcal{A}')$ equals the intersection of $Y_f(\mathcal{A})$ with the subspace where $v_i = 0$ and $t_i = 0$ for all $p_i \in \mathcal{A} \setminus \mathcal{A}'$. In particular, the marked graphs on $\mathcal{A}'$ corresponding to faces of $Y_f(\mathcal{A}')$ are subgraphs of marked graphs of faces of $Y_f(\mathcal{A})$. Moreover, non-crossingness of a marked graph on $\mathcal{A}$ is equivalent to non-crossingness of every induced marked graph on four vertices; indeed, a crossing of two edges appears in the marked graph induced by the four end-points of the two edges, and a non-pointed marked vertex appears in the marked graph induced on the four end-points involved in any three edges forming a non-pointed "letter Y" at the non-pointed vertex.

Hence: if f is valid for every four points, then none of the 4-point minors forbidden by non-crossingness appear in faces of $Y_f(\mathcal{A})$ and f is valid for $\mathcal{A}$. Conversely, if f is not valid on some four point subset $\mathcal{A}'$, then the marked graph on $\mathcal{A}'$ corresponding to any vertex of $Y_f(\mathcal{A}')$ would be the complete graph minus one boundary edge or vertex, that is to say, it would not be non-crossing. Hence f would not be valid on $\mathcal{A}$ either. □

Corollary 3.4.5 *For any $a, b \in \mathbb{R}^2$, the choice $f_{ij} := \det(a, p_i, p_j)\det(b, p_i, p_j)$, $f_{0j} := 0$ is valid.*

Proof: Consider the four points p_i as fixed and regard $R := \sum w_{ij} f_{ij} + \sum \alpha_j f_{0j} = \sum w_{ij} f_{ij}$ as a function of a and b:

$$R(a,b) = \sum_{1 \le i < j \le 4} \det(a, p_i, p_j) \det(b, p_i, p_j) w_{ij}.$$

We have to show that $R(a,b)$ is always positive. We actually claim it to be always 1. Observe first that $R(p_i, p_j)$ is trivially 1 for $i \neq j$. Since any three of our points are an affine basis and since $R(a,b)$ is an affine function of b for fixed a, we conclude that $R(p_i, b)$ is one for every $i \in \{1,2,3,4\}$ and for every b. The same argument shows that $R(a,b)$ is constantly 1; for fixed b it is an affine function of a and is equal to 1 on an affine basis. □

3.5 Points in special position

In this section we show that almost everything we said so far applies equally to point sets with collinear points. We will essentially follow the same steps as in Sections 3.2, 3.3 and 3.4. Two subtleties are that our definitions of pointedness or pseudo-triangulations can only be fully justified a posteriori, and that the construction of the polyhedron for point sets with boundary collinearities is slightly indirect: it relies in the choice of some extra exterior points to make collinearities go to the interior.

3.5.1 The graph of all pseudo-triangulations of $\mathcal{A}$

Definition 3.5.1 Let $\mathcal{A}$ be a finite point set in the plane, possibly with collinear points.

- A graph G with vertex set $\mathcal{A}$ is called *non-crossing* if no edge intersects another edge of G or point of $\mathcal{A}$ except at its end-points. In particular, if p_1, p_2 and p_3 are three collinear points, in this order, then the edge p_1p_3 cannot appear in a non-crossing graph, independently of whether there is an edge incident to p_2 or not.
- A *pseudo-triangle* is a simple polygon with only three interior angles smaller than 180 degrees. A *pseudo-triangulation* of $\mathcal{A}$ is a non-crossing graph with vertex set $\mathcal{A}$, which partitions $\text{conv}(\mathcal{A})$ into pseudo-triangles and such that no point in the interior of $\text{conv}(\mathcal{A})$ is incident to more than one angle of 180 degrees.

Figure 3.7 shows the eight pseudo-triangulations of a certain point set. We have drawn them connected by certain flips, to be defined later, and with certain points marked. The graph on the right of the figure is not a pseudo-triangulation because it fails to satisfy the last condition in our definition. Intuitively, the reason why we

do not allow it as a pseudo-triangulation is that we are considering angles of exactly 180 degrees as being reflex, and we do not want a vertex to be incident to two reflex angles.

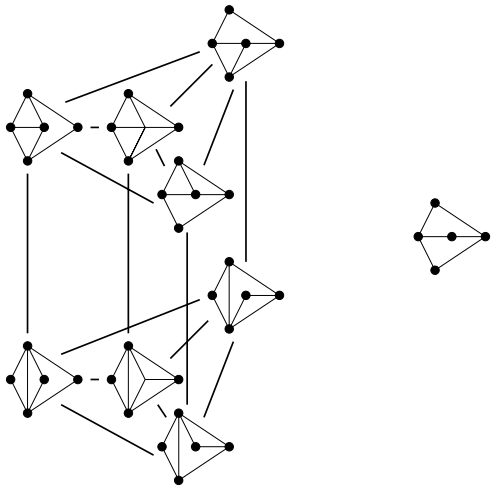

Figure 3.7: The eight pseudo-triangulations of a point set with interior collinearities (left) plus a non-crossing graph with pseudo-triangular faces but which we do not consider a pseudo-triangulation (right)

But if collinearities happen in the boundary of $\text{conv}(\mathcal{A})$, as in Figure 3.8, we treat things differently. The exterior angle of 180 degrees is not counted as reflex, and hence the middle point in a boundary collinearity is allowed to be incident to an interior angle of 180 degrees. The following definition can be restated as "a vertex is pointed if and only if it is incident to a reflex angle", where reflex is meant as in these last remarks.

Definition 3.5.2 A vertex p in a non-crossing graph on $\mathcal{A}$ is considered *pointed* if either (1) it is a vertex of $\text{conv}(\mathcal{A})$, (2) it is semi-interior and not incident to any edge going through the interior of $\text{conv}(\mathcal{A})$ or (3) it is interior and its incident edges span at most 180 degrees.

A *non-crossing marked graph* is a non-crossing graph with marks at some of its pointed vertices. If all pointed vertices are marked we say that the non-crossing graph is *fully-marked*. Marks at interior and semi-interior points will be called *interior marks*.

For example, all the graphs of Figures 3.7 and 3.8 are fully-marked. That is to say, big dots correspond exactly to pointed vertices. Of course, fully-marked pseudo-triangulations are just pseudo-triangulations with marks at all their pointed vertices. Observe that we are calling interior marks and edges exactly those which do not appear in all pseudo-triangulations. From now on, we denote by n_i, n_s and n_v the

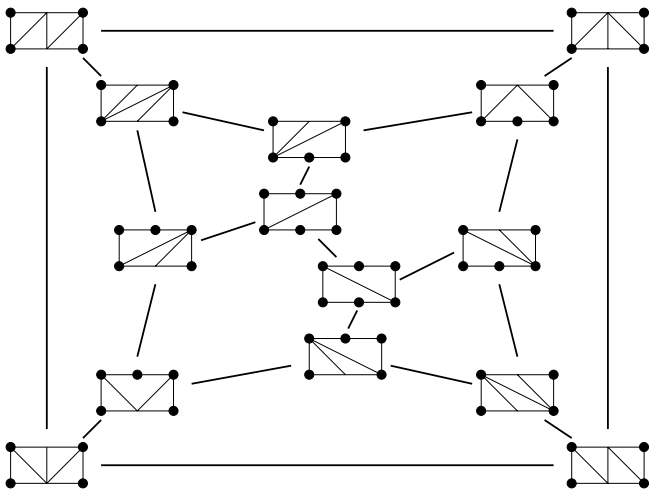

Figure 3.8: The graph of pseudo-triangulations of a point set with boundary collinearities

number of interior, semi-interior and extremal points of $\mathcal{A}$. Finally, $n = n_v + n_s + n_i$ denotes the total number of points in $\mathcal{A}$. The following two statements essentially say that Proposition 3.2.10 is valid for non-generic configurations.

Lemma 3.5.3 *Fully-marked pseudo-triangulations are exactly the maximal marked non-crossing graphs on $\mathcal{A}$. They all have $3n-n_s-3$ edges plus marks and $2n_i+n-3$ interior edges plus interior marks.*

Proof: The first sentence is equivalent to saying that every non-crossing graph G can be completed to a pseudo-triangulation without making any pointed vertex non-pointed. The proof of this is that if G is not a pseudo-triangulation then either it has a face with more than three corners, in which case we insert the diagonal coming from the geodesic between any two non-adjacent corners, or there is an interior vertex with two angles of 180 degrees, in which case we choose to consider one of them as reflex and the other as convex, and insert the diagonal joining the convex angle to the opposite corner of the pseudo-triangle containing it.

To prove the cardinality of pseudo-triangulations, let n_ϵ denote the number of marks. Let us think of boundary collinearities as if they were concave boundary chains in our graph, and triangulate the polygons formed by these chains by adding (combinatorially, or topologically) n_s edges in total. If, in addition, we consider interior angles of 180 degrees or more as reflex and the others as convex, we get a graph with all the combinatorial properties of pseudo-triangulations and, in particular, a graph for which Proposition 3.2.3 can be applied, since its proof is purely combinatorial (a double counting of convex angles, combined with Euler's relation, see Subsection 1.4.2). In particular, the extended graph has $3n-3$ edges plus marks, and the original graph has $3n - 3 - n_s$ of them. Since there are exactly $n_s + n_v$

exterior edges and n_v exterior marks in every pseudo-triangulation, the last sentence follows. □

Lemma 3.5.4 *If an interior edge or mark is removed from a fully-marked pseudo-triangulation then there is a unique way to insert another edge or mark to obtain a different fully-marked pseudo-triangulation.*

Proof: If an edge is removed then there are three possibilities: (1) the removal does not create any new reflex angle, in which case the region obtained by the removal is a pseudo-quadrangle (that is, it has four non-reflex angles), because the two regions incident to it had six corners in total and the number of them decreases by two. We insert the opposite diagonal of it. (2) the removal creates a new reflex angle at a vertex which was not pointed. Then the region obtained is a pseudo-triangle and we just add a mark at the new pointed vertex. (3) the removal creates a new reflex angle at a vertex that was already pointed. This means that after the removal the vertex has two reflex angles, that is to say two angles of exactly 180 degrees each. We insert the edge joining this vertex to the opposite corner of the pseudo-triangle containing the original reflex angle.

If a mark is removed, then the only possibility is: (4) the pointed vertex holding the mark is incident to a unique reflex angle (remember that we consider interior angles of 180 degrees as reflex). We insert the edge joining the vertex to the opposite corner of the corresponding pseudo-triangle. □

Definition 3.5.5 Two fully-marked pseudo-triangulations are said to differ by a *flip* if they differ by just one edge or mark. Cases (1), (2), (3) and (4) in the previous proof are called, respectively, *diagonal flip*, *deletion flip*, *mirror flip* and *insertion flip*.

Of course, our definition of flips specializes to the one for points in general position, except that mirror flips can only appear in the presence of collinearities. An example of a mirror flip can be seen towards the upper right corner of Figure 3.7.

Corollary 3.5.6 *The graph of flips between fully-marked pseudo-triangulations of a planar point set is connected and regular of degree* $2n_i + n - 3$.

The reader will have noticed that the graphs of Figures 3.7 and 3.8 are more than regular of degrees 4 and 3 respectively. They are the graphs of certain simple polytopes of dimensions 4 and 3. (Figure 3.7 is a prism over a simplex).

3.5.2 The case with only interior collinearities

Now we assume that our point set $\mathcal{A}$ has only interior collinearities.

For each $f \in \mathbb{R}^{n+1}$ let $\overline{Y_f}(\mathcal{A})$ be the polyhedron defined in Section 3. Recall that everything we said in that section, up to Corollary 3.3.5, is valid for points in special position. Our main result here is that Theorems 3.3.7 and 3.3.8 hold word by word in the case with no boundary collinearities, except that a precision needs to be made regarding the concept of validity.

Recall that for a given choice of $f \in \mathbb{R}^{\binom{n+1}{2}}$, an edge p_ip_j or a point p_i are called *tight* for a certain $(v,t) \in \mathbb{R}^{3n-3}$ or for a face F of $\overline{Y_f}(\mathcal{A})$ if the corresponding equation (3.4) or (3.5) is satisfied with equality.

Definition 3.5.7 We call *strict supporting graph* of a $(v,t) \in \mathbb{R}^{3n-3}$ (or face F of $\overline{Y_f}(\mathcal{A})$) the marked graph of all its tight edges and points, and denote it $T(v,t)$. We call *weak supporting graph* of a (v,t) or face the marked subgraph consisting of edges and points of $T(v,t)$ which define facets of $\overline{Y_f}(\mathcal{A})$.

A choice of f is called *weakly valid* (resp., *strictly valid*) if the weak (resp., strict) supporting graphs of all the faces of $\overline{Y_f}(\mathcal{A})$ are non-crossing marked graphs.

Observe that from any weakly valid choice f one can obtain strictly valid ones: just decrease by arbitrary positive amounts the coordinates of f corresponding to equations which do not define facets of $\overline{Y_f}(\mathcal{A})$. Hence, we could do what follows only in terms of strict validity and would obtain the same polyhedron. But weak validity is needed, as we will see in Remark 3.5.13, if we want our construction to depend continuously on the coordinates of the point set $\mathcal{A}$.

To obtain the equations that valid choices must satisfy we proceed as in Section 3.4. The crucial point there was that a marked graph is non-crossing if and only if it does not contain the negative parts of the unique stress in certain subgraphs.

Lemma 3.5.8 *Let $\mathcal{A}$ be a point set with no three collinear boundary points. Then, a marked graph on $\mathcal{A}$ is non-crossing if and only if it does not contain any of the following four marked subgraphs: the negative parts of the marked graphs displayed in Figure 3.6 and the negative parts of the marked graphs displayed in Figure 3.9.*

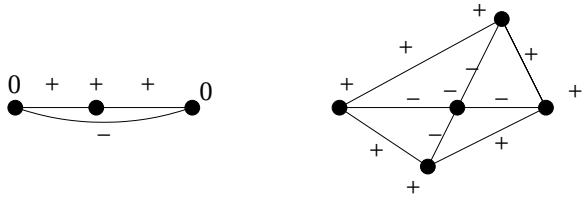

Figure 3.9: The two additional excluded minors for non-crossing marked graphs of a point set with interior collinearities

Proof: Exclusion of the negative parts of the left graphs in both figures are our definition of crossingness for an unmarked graph. An interior vertex is pointed if

and only if none of the negative parts of the right graphs appear. □

Lemma 3.5.9 *The two graphs in Figure 3.9 have a stress with signs as in the figure.*

Proof: For the left part it is easy to show that the following is a stress:

$$w_{12} = \frac{1}{|p_2 - p_1|}, \quad w_{13} = -\frac{1}{|p_3 - p_1|}, \quad w_{23} = \frac{1}{|p_3 - p_2|}, \quad \alpha_1 = \alpha_3 = 0, \quad \alpha_2 = 2.$$

For the right part, observe that, by definition, stresses on a marked graph form a linear space. Let the four exterior points be p_1, p_2, p_3 and p_4, in cyclic order, and let the interior point be p_5. We know three different stresses of the complete graph on these five points: The one we used in Section 3.4 for the four exterior points and the two that we have just introduced for the two collinear triplets. From these three we can eliminate the coordinates of edges p_1p_3 and p_2p_4 and we get a stress with the stated signs. □

Theorem 3.5.10 *Let $\mathcal{A}$ be a point set with no three collinear boundary points. Then, a choice of f is weakly valid if it satisfies equations (3.9) for all quadruples of points in general position plus the following sets of equations:*

- *For any three points p_1, p_2 and p_3 collinear in this order:*

$$\frac{f_{12}}{|p_2 - p_1|} - \frac{f_{13}}{|p_3 - p_1|} + \frac{f_{23}}{|p_3 - p_1|} + 2f_{02} \geq 0 \tag{3.10}$$

- *For any five points as in the right part of Figure 3.9, the following equation where the w_{ij}'s and the α_i's form a stress with signs as indicated in the figure (by convention, w_{ij} equals zero for the two missing edges in the graph):*

$$\sum_{1 \leq i < j \leq 5} w_{ij} f_{ij} + \sum_{j=1}^{5} \alpha_j f_{0j} > 0 \tag{3.11}$$

The choice is strictly valid if and only if, moreover, the equations (3.10) of all collinear triplets are satisfied strictly.

Proof: Equations (3.9) guarantee that no weak or strict tight graph contains the two excluded marked graphs of negative edges and points of Figure 3.6. Equations (3.10) and (3.11) with strict inequality, do the same for the graphs of Figure 3.9. That these equations are equivalent to strict validity is proved exactly as in Section 3.4. The reason why we allow equality in equations (3.10) if we only want a weakly valid choice is that the negative part of the stress consists of a single edge. If the equation is satisfied with equality then the hyperplane corresponding to this edge is a supporting hyperplane of the face of $\overline{Y_f}$ given by the intersection of the three hyperplanes of the positive part of the stress. □

Corollary 3.5.11 *Let $\mathcal{A}$ be a point set with no collinear boundary points. Any choice of f satisfying equations (3.9) for every four points in general position plus the following ones for every collinear triplet is weakly valid:*

$$\frac{f_{12}}{|p_2 - p_1|} - \frac{f_{13}}{|p_3 - p_1|} + \frac{f_{23}}{|p_3 - p_1|} + 2f_{02} = 0 \tag{3.12}$$

In particular, the choices of Corollary 3.4.5 and Theorem 3.3.7 are weakly valid.

Proof: For the first assertion, we need to show that equations (3.11) follow from equations (3.9) and (3.12). But this is straightforward: from our proof of Lemma 3.5.9 it follows that equation (3.11) is just the one obtained substituting in (3.9) the values for w_{13} and w_{24} obtained from the two equations (3.12).

For the last assertion, we already proved in Corollary 3.4.5 that the choices of f introduced there satisfy equations (3.9). It is easy, and left to the reader, to show that they also satisfy (3.12). □

Theorem 3.5.12 (Main theorem, case without boundary collinearities) *Let $\mathcal{A}$ be a point set with no three collinear points in the boundary of conv($\mathcal{A}$), and let f be a weakly valid choice of parameters. Then, $\overline{Y_f}$ is a simple polyhedron of dimension $3n-3$ with all the properties stated in Theorems 3.3.8 and 3.3.9.*

Proof: Recall that if no three boundary points are collinear then every fully-marked pseudo-triangulation (i.e., maximal marked non-crossing graph) has $3n-3$ edges plus marks, exactly as in the general position case (Lemma 3.5.3). In particular, it is still true, for the same reasons as in the general position case, that $\overline{Y_f}(\mathcal{A})$ is simple and all its vertices correspond to f.m.p.t.'s, for any valid choice of f. The rest of the arguments in the proof of Theorem 3.3.8 rely on the graph of flips being connected, a property that we still have. As for Theorem 3.3.9, the face $Y_f(\mathcal{A})$ is bounded because Lemma 3.3.3 still applies. The rest is straightforward. □

Remark 3.5.13 It is interesting to observe that taking the explicit valid choice of f of Theorem 3.3.7 the equations defining $\overline{Y_f}(\mathcal{A})$ depend continuously on the coordinates of the points in $\mathcal{A}$. When three points become collinear, the hyperplane corresponding to the (now) forbidden edge becomes, as we said in the proof of Theorem 3.5.10 a supporting hyperplane of a codimension 3 face of $\overline{Y_f}(\mathcal{A})$. The combinatorics of the polytope changes but maintaining its simplicity. This continuity of the defining hyperplanes would clearly be impossible if we required our choice to be strictly valid for point sets with collinearities.

Example 3.5.14 Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be the two point sets with five points each whose pseudo-triangulations are depicted in Figures 3.10 and 3.11. The first one has three collinear points and the second is obtained by perturbation of the collinearity. These

two examples were computed with the software `CDD+` of Komei Fukuda [35] before we had a clear idea of what the right definition of pseudo-triangulation for points in special position should be. To emphasize the meaning of weak validity, in Figure 3.10 we are showing the weak supporting graphs of the vertices of $\overline{Y_f}(\mathcal{A}_1)$, rather than the strict ones.

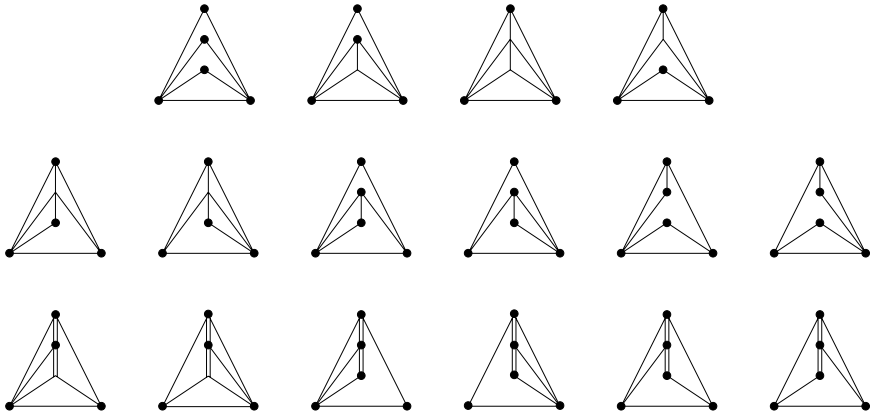

Figure 3.10: The 16 pseudo-triangulations of $\mathcal{A}_1$.

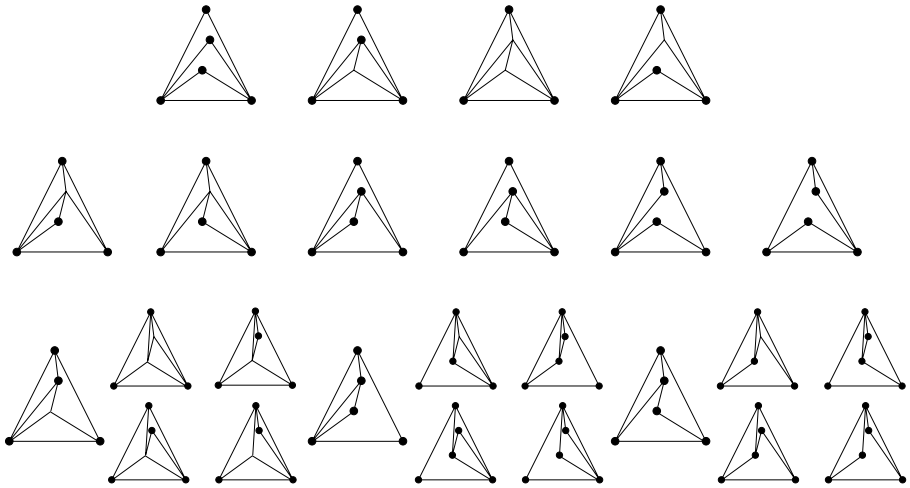

Figure 3.11: The 25 pseudo-triangulations of $\mathcal{A}_2$.

The 10 first pseudo-triangulations are common to both figures (upper two rows). Only the pseudo-triangulations of $\mathcal{A}_1$ using the two collinear edges plus the mark at the central point of the collinearity are affected by the perturbation of the point set. This is no surprise, since these two edges plus this mark are the positive part of the stress involved in the collinearity. At each of these six pseudo-triangulations,

the hyperplane of the big edge is tangent to the vertex of $\overline{Y_f}(\mathcal{A}_1)$ corresponding to the pseudo-triangulation. When the collinearity is perturbed, this hyperplane moves in one of the two possible ways; away from the polyhedron, in which case the combinatorics is not changed, or towards the interior of the polyhedron, in which case the old vertex disappears and some new vertices are cut by this hyperplane. In our case, these two behaviors appear each in three of the six "non-strict" pseudo-triangulations of $\mathcal{A}_1$. When the hyperplane moves towards the interior, four new vertices appear where there was one.

3.5.3 Boundary collinearities

In the presence of boundary collinearities, Lemma 3.5.3 implies a significant difference: with the equations we have used so far, the face of $\overline{Y_0}(\mathcal{A})$ defined by tightness at boundary edges and vertices is not the origin, but an unbounded cone of dimension n_s. Indeed, for each semi-interior point p_i, the vector (v, t) with v_i an exterior normal to the boundary of $\text{conv}(\mathcal{A})$ at p_i and every other coordinate equal to zero defines an extremal ray of that face. As a consequence, the corresponding face in $\overline{Y_f}(\mathcal{A})$ is unbounded.

We believe that it should be possible to obtain a polyhedron with the properties we want by just intersecting the polyhedron of our general definition with n_s hyperplanes. But instead of doing this we use the following simple trick to reduce this case to the previous one. From a point set $\mathcal{A}$ with boundary collinearities we construct another point set $\mathcal{A}'$ adding to $\mathcal{A}$ one point in the exterior of each edge of $\text{conv}(\mathcal{A})$ that contains semi-interior points.

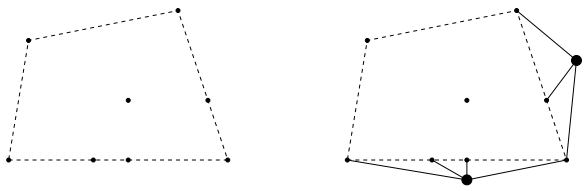

Figure 3.12: The extended point set $\mathcal{A}'$ and the relation between non-crossing marked graphs on $\mathcal{A}$ and $\mathcal{A}'$

Lemma 3.5.15 *A marked graph G on $\mathcal{A}$ is non-crossing if and only if it becomes a non-crossing graph on $\mathcal{A}'$ when we add to it the marks on all points of $\mathcal{A}' \setminus \mathcal{A}$ and the edges connecting each of these points to all the points of $\mathcal{A}$ lying in the corresponding edge of conv$(\mathcal{A})$.*

Proof: Straightforward. □

In particular, we can construct the polyhedron $\overline{Y_f}(\mathcal{A}')$ for this extended point set $\mathcal{A}'$ (taking any f valid on $\mathcal{A}'$), and call $\overline{Y_f}(\mathcal{A})$ the face of $\overline{Y_f}(\mathcal{A}')$ corresponding

to (tightness of) the edges and marks mentioned in the statement of Lemma 3.5.15. Then:

Corollary 3.5.16 (Main theorem, case with boundary collinearities) *$\overline{Y_f}(\mathcal{A})$ is a simple polyhedron of dimension $3n-3-n_s$ with all the properties stated in Theorem 3.3.8.*

Let $Y_f(\mathcal{A})$ be the face of $\overline{Y_f}(\mathcal{A})$ corresponding to the n_v+n_s edges between consecutive boundary points and the n_v marks at vertices of conv$(\mathcal{A})$. Let F be the face of $Y_f(\mathcal{A})$ corresponding to the remaining $n-n_v$ marks. Then, $Y_f(\mathcal{A})$ is a polytope of dimension $2n_i+n-3$ and F is a face of it of dimension $2n_i+n_v-3$. They satisfy all the properties stated in Theorem 3.3.9.

Remark 3.5.17 The reader may wonder about the combinatorics of the polyhedron $\overline{Y_f}(\mathcal{A})$ that one would obtain with the equations of the generic case. Clearly, the tight graphs of its faces will not contain any of the four forbidden subgraphs of Lemma 3.5.8. It can be checked that the maximal marked graphs without those subgraphs all have $3n-3$ edges plus marks and have the following characterization: As graphs they are pseudo-triangulations in which all the semi-interior vertices are incident to interior edges, and they have marks at all the boundary points and at the pointed interior points. In other words, they would be the fully-marked pseudo-triangulations if we treated semi-interior points exactly as interior ones, hence forbidding them to be incident to two angles of 180 degrees and considering them always pointed since they are incident to one angle of 180 degrees.

For example, in the point set of Figure 3.8 there are 6 such graphs, namely the ones shown in Figure 3.13.

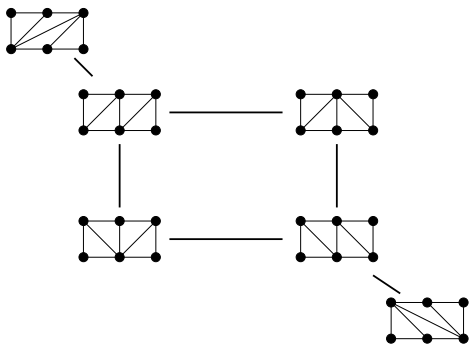

Figure 3.13: The bounded part of $\overline{Y_f}(\mathcal{A})$ for a point set with boundary collinearities

This implies that the polyhedron $\overline{Y_f}(\mathcal{A})$ is still simple. The reason why we prefer the definitions we have given is that the polyhedron no longer has a unique maximal bounded face (it has three in the example of Figure 3.13) and the graph of flips is no longer regular.

Observe finally that the proof of Theorem 3.1.3 given at the end of Section 3.3 is valid for points in special position, without much change: in all cases the hyperplanes corresponding to the edges of a pseudo-triangulation are independent in $\overline{Y_0}(\mathcal{A})$.

Chapter 4

Planar minimally rigid graphs and pseudo-triangulations

Recall from Corollary 1.4.10 that pointed pseudo-triangulations are planar minimally rigid graphs. The main result in this chapter is that the opposite statement is also true, that is, planar minimally rigid graphs admit geometric embeddings as pointed pseudo-triangulations. The paper [39] contains two different proofs of this fact, obtained simultaneously by different subsets of authors. The one we present here should be considered joint work of the author of this thesis with Günter Rote and Francisco Santos.

In order to prove our result, we use a directed version of Tutte's Theorem 1.2.8 to prove that every graph having the same combinatorics as a pseudo-triangulation and fulfilling a generalized Laman property can actually be (geometrically) embedded as a true pseudo-triangulation. The proof of our main result is mainly combinatorial and goes further by proving it to be true even when the graph is under the above mentioned combinatorial constraints. The proofs presented here are constructive and yield efficient embedding algorithms whose worst-case complexity we analyze.

4.1 Introduction

We have already outlined in Section 1.2.2 the importance of Graph Drawing and plane embeddings throughout the history of Graph Theory. Here, we are looking for pointed geometric embeddings of planar minimally rigid graphs; as opposed to traditional planar graph embeddings, in our case the faces are "as non-convex as possible" (pseudo-triangles). Such embeddings with non-convex faces have not been systematically studied and our result links them to rigidity theoretic and combinatorial properties of planar graphs. In addition, it gives a simple and elegant combinatorial characterization of all the graphs which can be embedded as pointed pseudo-triangulations, so that we answer an open question posed in [68].

It has been shown in Section 1.3 that Laman graphs are *the* fundamental objects in 2-dimensional Rigidity Theory since they characterize combinatorially the prop-

erty that a graph, embedded on a generic set of points, is infinitesimally rigid (with respect to the induced edge lengths). We are now interested in *planar* Laman graphs; observe that not all Laman graphs fall into this category. For example, the complete bipartite graph $K_{3,3}$ is Laman but (by Kuratowski's planarity Theorem 1.2.3.(ii)) not planar.

It is a result of Streinu [69] (which we stated as Corollary 1.4.10 in the introductory chapter) that the underlying graphs of all pointed pseudo-triangulations are planar Laman. We prove that the converse is always true, what adds to the already rich body of surprisingly simple and elegant combinatorial properties of pointed pseudo-triangulations:

Theorem 4.1.1 *Every planar Laman graph can be embedded as a pointed pseudo-triangulation.*

Corollary 4.1.2 *Given a graph G, the following conditions are equivalent:*

(i) G is planar and minimally rigid (isostatic),

(ii) G is a planar Laman graph,

(iii) G can be embedded as a pointed pseudo-triangulation.

Proof: (i)⇒(ii) comes from Lemma 1.3.20. The implication (ii)⇒(iii) is the previous result. Finally, (iii)⇒(i) is Corollary 1.4.10. □

The chapter is organized as follows: As advanced above, we prove in fact a result stronger than Theorem 4.1.1 (Theorem 4.2.4), which allows even the fixing *a priori* of the combinatorial information regarding which vertices are convex in each face. This is carried out in Section 4.2, where we introduce combinatorial pseudo-triangulations and the directed version of Tutte's Theorem that we use to embed them. The proof of a key lemma is postponed to Section 4.3. Finally, Section 4.4 is devoted to labeling infinitesimally rigid graphs as combinatorial pseudo-triangulations. In particular, this finishes the proof of Theorem 4.1.1 as a corollary of Theorems 4.2.4 and 4.4.1.

4.2 Embedding combinatorial pseudo-triangulations

The embeddings we consider in the whole chapter are in $\mathbb{R}^2$. We also use topological embeddings, which we regard as projected from the sphere $\mathbb{S}^2$ into the euclidean plane, with edges realized as arcs:

Definition 4.2.1 A *combinatorial pseudo-triangulation*, abbreviated as *c.p.t.*, is a topological plane embedding of a graph together with a labeling of the angles as "big" or "small", such that:

(i) There are exactly three small angles per bounded face,

(ii) There is no small angle in the outer face,

(iii) There is at most one big angle per vertex, and

(iv) Vertices of degree two are incident to exactly one big angle.

Following the analogy with true pseudo-triangulations, the cells of a c.p.t. are called *combinatorial pseudo-triangles*. Thus, *combinatorial corners* and *pseudo-edges* are naturally defined; we may omit the word "combinatorial" when it is clear from the context. Also, the vertices are called *pointed* or *non-pointed* according to whether they have an adjacent big angle or not. A *combinatorial pointed pseudo-triangulation* is one with exactly one big angle at every vertex and is abbreviated as *c.p.p.t.*

Of course, every pseudo-triangulation in general position is a combinatorial pseudo-triangulation, by just labeling "big" the angles greater than 180 degrees and "small" the others. The same arguments that one uses with true pseudo-triangulations in Proposition 1.4.7 show that the number of edges of a combinatorial pseudo-triangulation is exactly $2n_\epsilon + 3n_\gamma - 3$, where n_ϵ and n_γ are the numbers of pointed and non-pointed vertices, respectively (this will be also a consequence of Lemma 4.3.2). Let $n = n_\epsilon + n_\gamma$ be the total number of points:

Definition 4.2.2 We say that a combinatorial pseudo-triangulation G has the *generalized Laman property* or is *generalized Laman* if:

- Every subset of k pointed plus l non-pointed vertices, with $k + l \geq 2$, induces a subgraph with at most $2k + 3l - 3$ edges.

Or, equivalently,

- Every subset of k pointed plus l non-pointed vertices of G, with $k + l \leq n - 2$, is incident to at least $2k + 3l$ edges.

Observation 4.2.3 Observe that for a pointed c.p.t. the generalized Laman property is just the standard Laman property of the underlying graph.

Also, the generalized Laman property implies the condition (iv) in Definition 4.2.1: Since a vertex with all angles labeled small is non-pointed, it has to be incident to at least 3 edges.

The main result in this section is that generalized Laman combinatorial pseudo-triangulations characterize pseudo-triangulations. The converse of the following result is also true since every pseudo-triangulation has the generalized Laman property, by Theorem 3.1.3 of the previous chapter:

Theorem 4.2.4 *Every combinatorial pseudo-triangulation having the generalized Laman property can be embedded as a pseudo-triangulation, with reflex and convex angles respectively for those labeled big and small.*

Note that having the generalized Laman property is not superfluous in the statement, even for pointed c.p.t.'s. Figure 4.1 (left) shows a combinatorial pointed pseudo-triangulation which is not Laman, as the removal of the two most interior vertices shows. Hence it cannot be embedded as a pointed pseudo-triangulation. "Big" labels are represented by a dot and "small" ones by an arc. This notation will hold in the sequel.

In addition, it is worth noticing that for a c.p.p.t., "generalized Laman" is equivalent to "rigid" (by Observation 4.2.3). For non-pointed ones, we conjecture that the generalized Laman property implies rigid, and we prove that every rigid and planar graph can be embedded as a combinatorial pseudo-triangulation.

It is also not true that every rigid c.p.t. has the generalized Laman property, as the right picture in Figure 4.1 shows, removing again the two most interior vertices. However, we also conjecture that a labeling as generalized Laman c.p.t. can always be found for every rigid and planar graph, what implies this could be embedded as a pseudo-triangulation.

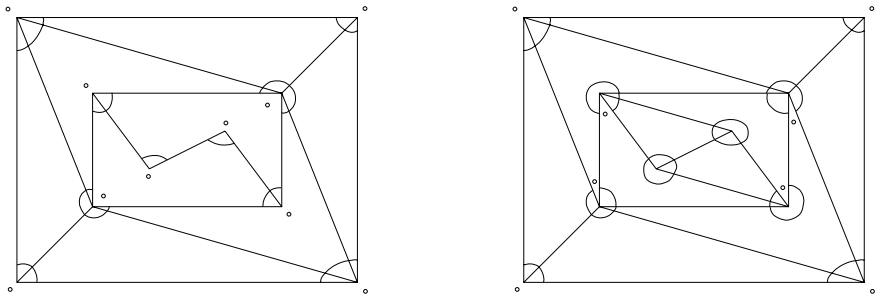

Figure 4.1: Left, a c.p.p.t. which is not a Laman graph. Right, a rigid c.p.t. which is not generalized Laman.

The proof of Theorem 4.2.4 relies on a version of Tutte's Theorem 1.2.8, extended to the directed case. Let us introduce some notions before giving the statement:

Definition 4.2.5

- We define a *partially directed graph* $(V, E, \vec{E})$ to be a graph (V, E) together with an assignment of directions to some of its edges. Thus, edges are allowed to get the two directions, only one, or none. Formally, $\vec{E}$ is a subset of $E \cup (-E)$ where E is the set of edges of G with an arbitrary direction.

- For a given embedding $G(\{p_1, \ldots, p_n\})$ of a partially directed graph $G = (V, E, \vec{E})$ and for a given assignment $w : \vec{E} \to \mathbb{R}$ of weights to the directed edges, a vertex $i \in V$ is said to be in *equilibrium* if

$$\sum_{(i,j)\in\vec{E}} w_{ij}(p_i - p_j) = 0.$$

- A plane embedding of $(V, E, \vec{E})$ is *3-connected to the boundary* if from every interior vertex p there are at least three vertex-disjoint directed paths in $\vec{E}$ ending in three different boundary vertices. Equivalently, if for any interior vertex p and for any pair of forbidden vertices q and r there is a directed path from p to the boundary not passing through q or r.

Theorem 4.2.6 (Directed Tutte Theorem) *Consider a partially directed planar graph $G = (\{1, \ldots, n\}, E, \vec{E})$ planarly embedded with outer face $(k+1, \ldots, n)$ for some $k < n$, and 3-connected to the boundary. Let $p_{k+1}, \ldots, p_n$ be the ordered vertices of a convex $(n-k)$-gon. Let $w : \vec{E}' \to \mathbb{R}$ be an assignment of positive weights to the internal directed edges. Then:*

(i) There are unique positions $p_1, \ldots, p_k \in \mathbb{R}^2$ for the interior vertices such that all of them are in equilibrium in the embedding $G(\{p_1, \ldots, p_n\})$.

(ii) In this embedding, all cells of G are then realized as non-overlapping convex polygons.

Proof: The proof of Tutte's Theorem sketched in Chapter 1 works with only minor modifications. First, in the definition of *good representation* (Definition 1.2.11), each point p_i has to be posed to be in the relative interior of its *out-neighbors*, since this is what the directed equilibrium condition gives. Second, recall that Lemma 1.2.13 proves that in a good representation it is not possible for a vertex p and all its neighbors to lie in a certain line ℓ, using 3-connectedness. The proof can be adapted to use "3-connectedness to the boundary" as follows: consider three vertex disjoint paths from p to the boundary. Call q any of the three end-points, assumed not to lie in the line ℓ. Complete the other two paths to end at q using boundary edges in opposite directions. This produces three vertex-disjoint paths from p to a vertex q not lying on ℓ. The rest needs no change. □

Proof of Theorem 4.2.4. Let T be the combinatorial pseudo-triangulation, topologically embedded in the plane. We first extend the underlying graph of T to become a (topological) triangulation by triangulating those bounded cells of T with more than three vertices. We call G this extended graph and we construct it in such a way that every big angle of T is dissected by at least one new edge of $G \setminus T$ (this can be achieved by triangulating each cell with only two ears, to be placed at small angles). See Figure 4.2, where only small angles are marked (by an arc).

This new graph G is the one to which we will apply Theorem 4.2.6. The set $\vec{E}$ of directed edges is going to have out-degree three at every interior pointed vertex. More precisely, from each interior pointed vertex p we take as out-going directed edges its two *extreme adjacent edges* in T, those defining the big angle, and one of the new edges of $G \setminus T$ incident to p. On the other hand, from each interior non-pointed vertex, we consider all incident edges as out-going.

The key result we need about G is:

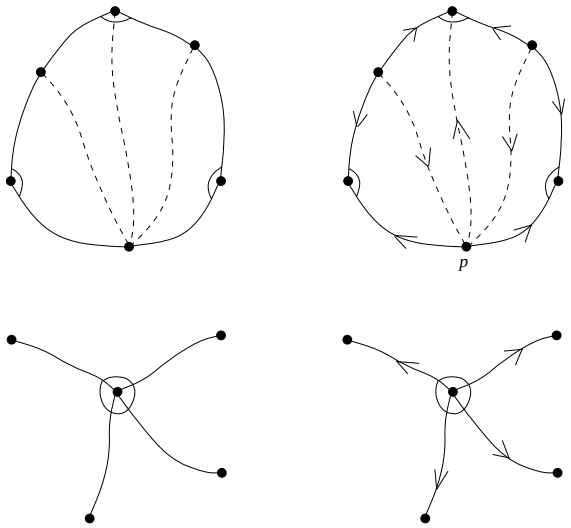

Figure 4.2: Construction of G and the set $\vec{E}$.

Lemma 4.2.7 *The partially directed graph G is 3-connected to the boundary.*

In order not to interrupt the current reasoning, we postpone the proof of this lemma to the next section.

We now apply Theorem 4.2.6 to G, with an arbitrary choice of positive weights for the directed edges and an arbitrary (but convex) choice of positions for the boundary vertices. Besides having a triangulation embedding of G, the equilibrium condition guarantees that every interior vertex is in the relative interior of the convex hull of its out-neighbors. For a pointed vertex; since two of this neighbors are the extreme adjacent vertices, when removing the edges of $G \setminus T$ the angle labeled "big" is indeed greater than 180 and all the others are smaller, so it becomes pointed in the true sense. For a non-pointed vertex, being in the relative interior of all its neighbors makes the angles be smaller than 180 and the vertex be truly non-pointed. In other words, the equilibrium embedding provides the pseudo-triangulation we are looking for. □

Concerning the time analysis, everything in this proof can be done in linear time, except for the computation of the equilibrium position for the interior vertices. In this computation one writes a linear equation for each interior vertex, which says that the position of the vertex is the average of its (out-)neighbors. The position of the boundary vertices is fixed. It has been observed [20, Section 3.4], [39] that the planar "structure" of this system of equations allows it to be solved in $O(n^{3/2})$ time or even in time $O(M(\sqrt{n}))$, where $M(n) = O(n^{2.375})$ is the time to multiply two $n \times n$ matrices.

4.3 3-connectedness of G

This section is devoted to prove Lemma 4.2.7. First, we study c.p.t.'s and show certain property that will be crucial in the proof. We later prove it to be equivalent to the generalized Laman property.

Since a combinatorial pseudo-triangulation is a planar graph embedded in $\mathbb{R}^2$, every subcomplex of it has a well-defined *outer face* (the one containing the outer face of T) and a well-defined "big angle" at every pointed vertex (the one containing the angle labeled "big" at that vertex in T).

Definition 4.3.1 We call *corners* of a connected subcomplex S of a c.p.t. T the boundary vertices of the following two types:

1. Vertices with a big angle in the outer face.
2. Vertices with at least two small angles (of T) in the outer face.

We also define the *boundary* or *contour* of S as being the boundary of its outer face. If S is connected, then its contour is a cycle of edges, except some of them may appear twice in the cycle (e.g., if S is a tree then every edge appears twice).

Lemma 4.3.2 *Let $S \subset T$ be a subcomplex of a c.p.t. and suppose that it is simply connected (that is to say, it has no "holes"; equivalently, its Euler characteristic equals 1).*

Let e, k, l and b denote the numbers of edges, pointed vertices, non-pointed vertices and length of the boundary cycle in S, respectively. Then, the number c_1 of corners of the first type (big angles in the outer boundary) of S is

$$c_1 = e + 3 - 2k - 3l + b.$$

Observe that this statement implies the afore mentioned result that every combinatorial pseudo-triangulation has $2n_\epsilon + 3n_\gamma - 3$ edges, since b equals in this case the number c_1 of outer big angles. It also implies that simply connected subcomplexes of (perhaps not generalized Laman) c.p.t.'s satisfy the inequality $e \leq 2k + 3l - 3$, since clearly b is greater or equal to the number of corners of type 1.

Proof: We classify angles incident to S as big, small or undecided, as follows: angles in bounded faces are big or small, as they were in T. Exterior angles in the boundary cycle are big if they contain a big angle of T, small if they are a small angle of T, and undecided otherwise (i.e., if they contain at least two small angles and no big angle). The number of corners equals the number of undecided plus exterior big angles.

The total numbers of big, small and undecided angles equal, respectively, k, $3f+b-c$ and $c-c_1$, where f is the number of bounded faces and c is the total number of corners of S. Since the total number of angles is $2e$, we get $2e = k + 3f + b - c_1$, or

$$3(e - f) = e + k + b - c_1.$$

Finally, Euler's formula says that $e - f = (k + l) - 1$, which gives $3k + 3l - 3 = e + k + b - c_1$, as desired. □

Corollary 4.3.3 *Every subcomplex S (with at least three vertices) of a combinatorial pseudo-triangulation T with the generalized Laman property has at least three corners.*

Proof: Let us first show that there is no loss of generality in assuming that S is simply connected and that no edge or vertex appears twice in the boundary cycle. If S is not connected, either some connected component has at least three vertices or all the vertices of S are corners. If S is connected but not simply connected, then adding to S the pseudo-triangles, edges and vertices of T that fill the holes of F in does not change the number of corners. If S has an edge e which appears twice in the boundary cycle, this can be removed: If both components of $S \setminus e$ have less than three vertices, then S is a path and all its vertices are corners. If one component has at least three vertices, then it has at least three corners (in $S \setminus e$) and gives two corners to S. The other component clearly gives at least another corner.

Finally, if S has a vertex v which appears twice in the boundary cycle, we can split S into two parts by doubling that vertex. Each part has at least three vertices (because we were already assuming that no edge appears twice in the boundary cycle) and hence each part has at least two corners other than v. In particular S has at least four corners.

Hence, we assume that S is simply connected, has at least three vertices, and has no repeated edges or vertices in the boundary cycle. Equivalently, we are saying that S is a topological ball.

Let k, l, e, and b be the numbers of pointed vertices, non-pointed vertices, edges and boundary vertices of S, respectively. Let V be the set of vertices of S which are either interior to S or boundary vertices, but not corners. V consists of $k+l-c_1-c_2$ vertices, where c_1 and c_2 are the corners of type 1 and 2 of S, respectively. Moreover, it contains all the l non-pointed vertices except for those being corners, which are at most c_2 (corners of type 1 are pointed vertices by definition). Hence, the generalized Laman property implies that the number of edges incident to V is at least

$$2(k + l - c_1 - c_2) + (l - c_2) = 2k + 3l - 2c_1 - 3c_2,$$

two times the total number of points plus the number of non-pointed ones. (Remark: V certainly does not contain the corners of the whole c.p.t., whose number equals b and is at least three. This guarantees that we can apply the generalized Laman condition to V).

On the other hand, the edges incident to V are the $e - b$ interior edges of S plus at most two edges per each boundary non-corner vertex. Hence,

$$2k + 3l - 2c_1 - 3c_2 \leq e - b + 2(b - c_1 - c_2),$$

or, equivalently,

$$c_2 \geq 2k + 3l - e - b.$$

Using Lemma 4.3.2 this gives $c_1 + c_2 \geq 3$. □

Proof of Lemma 4.2.7. We are going to prove the second version of 3-connectedness: We pose any two vertices q and r to be forbidden and our goal is to prove that from every vertex p there is a directed path in G going to the boundary and not passing through q or r.

Let us consider the *directed connected component* of the initial vertex p, defined as the set of all vertices and directed edges of G that can be reached from p without passing through q or r. We define this component as not containing the forbidden points q and r but it may contain edges "arriving to them".

We call N this directed connected component and want to prove that it contains a vertex on the boundary of G. We argue by contradiction. That is, suppose that all the vertices of N are interior to G.

For each interior pointed vertex v, let T_v be the unique pseudo-triangle of T containing the big angle at v. Thus, v is in a concave chain (pseudo-edge) of T_v containing also the two extreme adjacent edges of v. Let S be the union of all pseudo-triangles T_v associated to the pointed vertices v of N and all the pseudo-triangles incident to the non-pointed vertices. This union S covers N. Indeed, every directed edge of G is contained in a pseudo-triangle associated to its source vertex. We apply Theorem 4.3.3 to S, concluding that it has at least three corners. We claim that at least one of them belongs to N. This gives the contradiction, because then there is an edge of $G \setminus T$ jumping out of that corner c (by definition of corner), which means that the pseudo-triangle(s) corresponding to c should have been contained in S and hence c is not a corner of S anymore.

To prove the claim, let $v_1, \ldots, v_k$ be the corners of S which are not in N. We want to prove that $k \leq 2$. For this let $T_1, \ldots, T_k$ be pseudo-triangles in S, each having the corresponding v_i as a corner (there may be more than one valid choice of the T_i's; we just choose one). By definition, some non-corner pointed vertex or some corner non-pointed vertex of each T_i is in the component N. Were there no forbidden points, from a non-corner pointed vertex we could arrive to the three corners of T_i by three disjoint paths: two of them along the concave chain containing the initial point and the third starting with an edge of $G \setminus T$. From a non-pointed corner vertex we could arrive to the other two corners by two disjoint paths: moving out from the vertex to the two neighbors in the incident pseudo-edges and then following along them. In particular, since v_i is not in N, one of the forbidden points must obstruct to one of these paths, which implies that either v_i equals one of the forbidden points q and r or that T_i is the pseudo-triangle of one of the two forbidden points. And, clearly, each of the two forbidden points contributes to only one of the indices i (either as a corner of S or via its associated pseudo-triangle if it is not a corner, but not both). This finishes the proof that $k \leq 2$, and of the Lemma. □

It is interesting to observe that the only property of T used in this proof, other than being a combinatorial pseudo-triangulation, is that every subcomplex has at least three corners. This has the following consequence:

Proposition 4.3.4 *Let T be a combinatorial pseudo-triangulation. The following properties are equivalent:*

(i) T has the generalized Laman property.

(ii) Every subcomplex of T with at least three vertices has at least three corners.

(iii) T can be embedded as a pseudo-triangulation.

Proof: The implication (i)⇒(ii) is Corollary 4.3.3, and the implication (ii)⇒(iii) is our proof of Theorem 4.2.4. Finally, (iii)⇒(i) is Theorem 3.1.3 from the previous chapter. □

The reader can check (the failure of) these three properties in Figure 4.1.

4.4 Laman graphs admit c.p.p.t. labelings.

At this point, there is only one task remaining in order to prove the main Theorem 4.1.1: That it is always possible to give a generalized Laman combinatorial pointed pseudo-triangulation assignment to every planar Laman graph. Hence, the main theorem will be a corollary of the next result and Theorem 4.2.4. This section and the next one are joint work with the co-authors in [39].

Theorem 4.4.1 *For any choice of an outer face, every planar Laman graph has a combinatorial pointed pseudo-triangulation assignment. Every such an assignment has the generalized Laman property.*

Proof: First, let us notice that a planar minimally rigid graph may have several c.p.p.t. assignments. In order to obtain one, we make use of a combinatorial version of the Henneberg construction for planar Laman graphs presented in Chapter 1. This was previously noticed in [39], where a geometric version is used with similar purposes.

The proof is contained in the next two lemmas. It considers any topological plane embedding and follows the Henneberg construction in an inductive process to assign convex and reflex angles in a way naturally induced by the expected tangencies for a true pseudo-triangulation. Note that we first need to assign a consistent outer face to each intermediate graph; this is achieved by the first Lemma using the Henneberg construction in reverse order. The second Lemma performs the inductive assignment, for which the base case of a topological triangle is trivial.

Finally, that every c.p.p.t. assignment has the generalized Laman property is a direct consequence of the graph being Laman, as we have already observed. □

Lemma 4.4.2 *Every topological plane embedding of a Laman graph has a topological Henneberg construction in which:*

(i) All intermediate graphs are planarly embedded, and

(ii) At each step, the topology is changed only on edges and faces involved in the Henneberg step: Either a new vertex is added inside a face of the previous graph (Henneberg I), or inside a face obtained by removing an edge between two faces of the previous graph (Henneberg II).

Proof: By reverse induction. Since a Laman graph has $2n-3$ edges, there exists a vertex v of degree 2 or 3. If it is of degree 2, removing it and its adjacent edges merges two faces into one, and the graph obtained is clearly Laman. If it is of degree 3, let its neighbors be v_1, v_2 and v_3. In the proof of Lemma 1.3.20.(iv)$\Leftrightarrow$(v), Tay and Whiteley [70] prove that the removal of v and the insertion of one of the three edges v_1v_2, v_1v_3 or v_2v_3, say v_1v_2, gives a Laman graph. Topologically, we just delete the edge vv_3 and merge its two adjacent faces into one. The other face incident to v has two boundary edges merged into one. In either case, continue the inductive reduction until $n = 3$. □

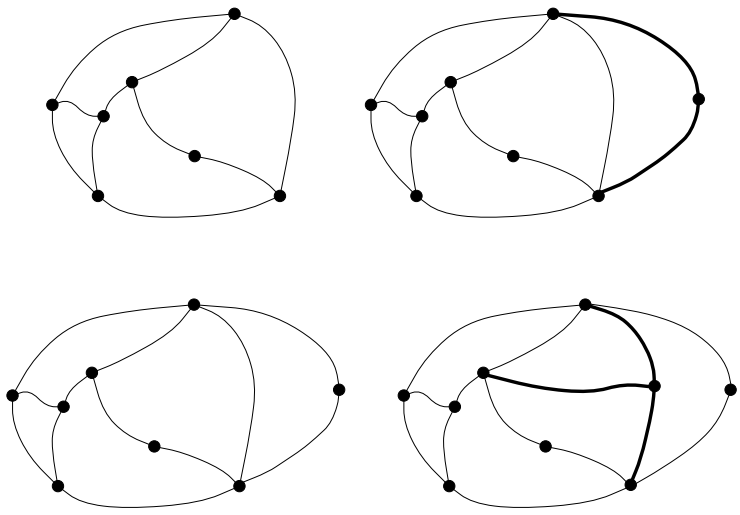

Figure 4.3: Two examples of a topologically planar Henneberg step. Top row: Henneberg I on the exterior face. Bottom row: Henneberg II on an interior face.

Lemma 4.4.3 *Let G_n be a topological plane embedding of a Laman graph obtained from G_{n-1} via a Henneberg step and assume inductively that G_{n-1} is a combinatorial pointed pseudo-triangulation. Then, there also exists a c.p.p.t. assignment for the angles of G_n.*

Proof: Given a topological plane embedding of the Laman graph, we find a c.p.p.t. assignment following the topological Henneberg construction of the previous Lemma and keeping the outer face prescribed there at each intermediate step.

Assume that we have a c.p.p.t. assignment for the angles of G_{n-1} which we now want to extend with a new vertex. A Henneberg I step will add a vertex and split an existing face into two: the faces involved in this step are bounded by simple polygons. A Henneberg II step first removes an edge, thus joining two existing faces into one. Note that, by the previous Lemma, this can always be done maintaining the planarity of the topological embedding and the prescribed outer face.

We have to show that it is always possible to find labels for the newly created angles in such a way that the conditions of Definition 4.2.1 and pointedness of every vertex are maintained:

- **Henneberg I step** (see Figures 4.4 and 4.5): For a Henneberg I step, the new vertex p_n is added inside one of the existing faces F, and joined to points p_i, p_j on the boundary of the face (in a true pseudo-triangulation, they would be tangency points) so that F becomes split into two faces F_1 and F_2. This splits the angles α_i, α_j of F at p_i and p_j into two angles α_{i1}, α_{i2} and α_{j1}, α_{j2}, where subindices are chosen according to the new face in which they lie. In addition, two new angles α_{n1}, α_{n2} appear at p_n. The rest of angles keep their small/big labels, and those for the newly created angles are chosen as follows:

  If F was the exterior face, the chain $p_i p_n p_j$ separates it into the new exterior face F_1 and a new interior one F_2. We label "small" the angles α_{i2}, α_{j2} and α_{n2}, assigning "big" labels to the others.

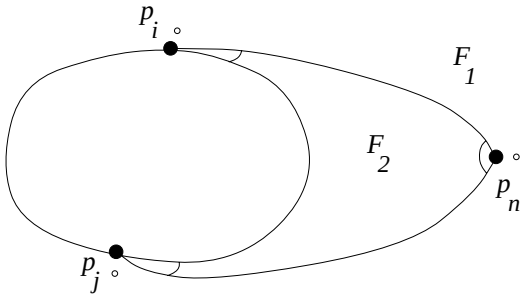

Figure 4.4: Labels for a Henneberg I step when F was the outer face.

Otherwise, F was a combinatorial pseudo-triangle and we consider the following three cases, which can be checked in Figure 4.5, according to whether both, one or none of p_i, p_j were combinatorial corners of F (that is; whether none, one or both of α_i, α_j were big before splitting):

- In the first case, the chain $p_i p_n p_j$ separates F into F_1 containing the third corner and F_2 containing the opposite pseudo-edge. We label small $\alpha_{i1}, \alpha_{i2}, \alpha_{j1}, \alpha_{j2}$ and α_{n2}. The other angle α_{n1} receives "big" label.

- For the second case let us assume that p_i was a corner. According to the position of the two other corners, two subcases arise: First, that $p_ip_np_j$ separates F into F_1 and F_2 containing one corner (different of p_i) each. We label small $\alpha_{i1}, \alpha_{i2}, \alpha_{j1}$ and α_{n2}, big the others. Second, that $p_ip_np_j$ separates F into F_1 containing the other two corners and F_2 containing none (appart from p_i). We label small $\alpha_{i1}, \alpha_{i2}, \alpha_{j2}$ and α_{n2}, the rest big.
- The last case has again two subcases: In the first one, $p_ip_np_j$ separates F into F_1 containing the three corners and F_2 containing none. We assign "small" labels to α_{i2}, α_{j2} and α_{n2}, "big" to the others. In the second subcase, $p_ip_np_j$ separates F into F_1 containing two corners and F_2 containing one and we label small α_{i1}, α_{j2} and α_{n2}, big the rest.

Observe that "a small angle is split into two small angles", "a big angle is split into one small and one big angles" and "at a new point only one angle is labeled big". In this way, every vertex has indeed a big angle. In addition, every face has exactly three small angles since we add three of them, the number of faces increases by one and we always choose to label big "the angle in the face with more small angles".

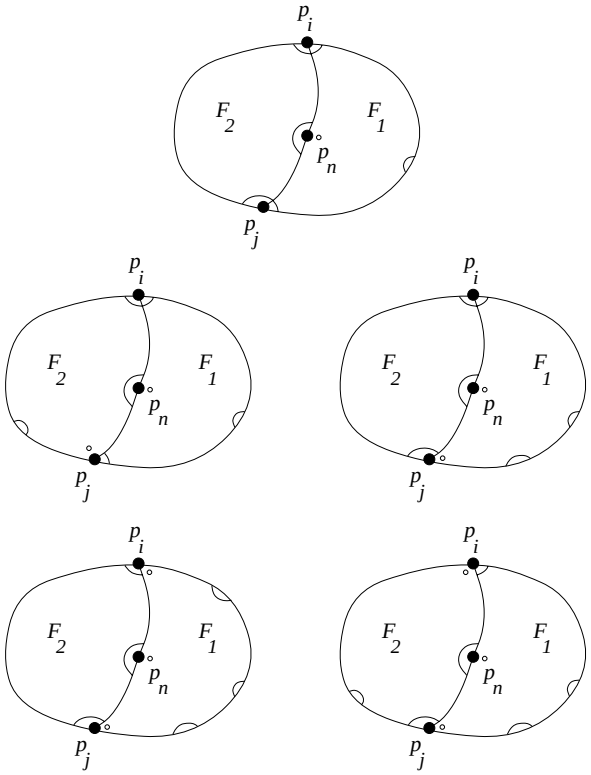

Figure 4.5: Labels for a Henneberg I step when F was a combinatorial pseudo-triangle.

- **Henneberg II step** (see Figures 4.6 and 4.7): Now a new vertex p_n is added on an edge ij which separated two faces F and F'. The new point p_n must be connected to p_i, p_j and a third vertex p_k lying on the boundary of $F \cup F'$, what splits either F or F' into two faces. Let us assume it is F and call them F_1 and F_2 so that $p_i \in F_1$, $p_j \in F_2$ and $p_k, p_n \in F_1 \cap F_2$.

  On the one hand, three new angles appear at p_n; two of them α_{n1}, α_{n2} in F_1 and F_2 respectively, and the other α'_n in F'. In addition, the angle at p_k in F becomes split into α_{k1}, α_{k2} as well. Finally, let us call α_i, α_j and α'_i, α'_j the angles at p_i, p_j respectively in F and F'.

  We have to consider first the case in which F was the outer face: Then, the new edge $p_n p_k$ separates it into the new exterior face F_1 containing p_i and an interior one F_2 containing p_j. Then, we switch the labels of α_j and α'_j (which were respectively big and small), label small $\alpha_{k2}, \alpha_{n2}, \alpha'_n$ and big α_{n1}, α_{k1}.

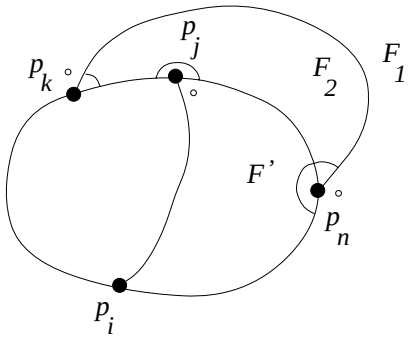

Figure 4.6: Labels for a Henneberg II step when F was the outer face.

  Suppose now that F is a combinatorial pseudo-triangle. We assume without loss of generality that the number of small angles in F_1 is smaller or equal to that in F_2. At a first step we label small the angles α_{n1}, α_{n2} and α_{k1}, and we label big the angle α'_n. We assign to α_{k2} the previous label of α_k in F, and the other angles remain unchanged. See Figure 4.7 (top). Note that F' keeps having exactly three small angles and the total number of small angles in $F_1 \cup F_2$ is then six, since F had three small angles (because it was a combinatorial pseudo-triangle) and we have just added three more. We come up with two possible cases:

  - There are three small angles in both F_1 and F_2: Then, we are done; F_1, F_2 and F' have exactly three small angles and there is exactly one big angle at each point. Observe that this is the case if the exterior face was F', since in this case α_i had to be small (because α'_i in the outer face was big and there is at most one big angle per vertex).
  - There are two small angles in F_1 and four in F_2: This is only possible if α_i was big. Since there was exactly one big angle per point, α'_i was then

small. We switch their labels and do the same with those of α'_n and α_{n2}, so that we have α_i, α'_n small and α'_i, α_{n2} big. This gives exactly three small angles at each of F_1, F_2, F' and exactly one big angle per vertex.

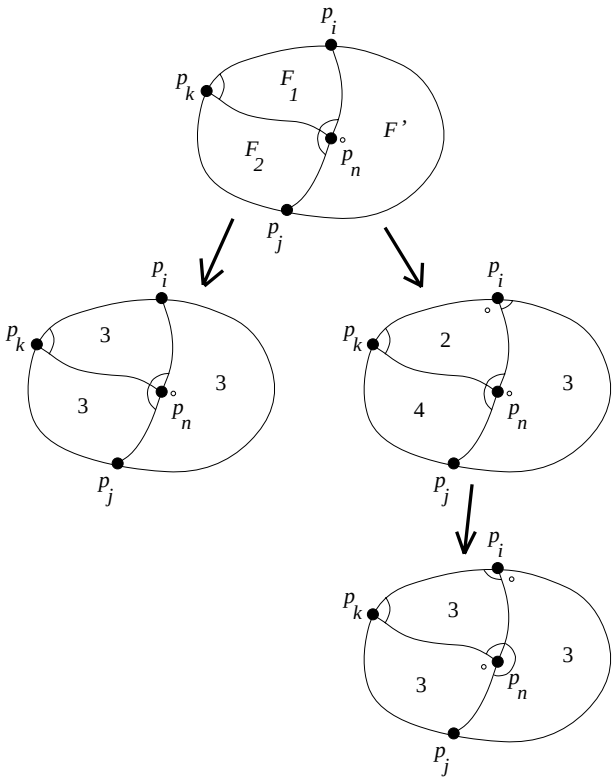

Figure 4.7: Labels for a Henneberg II step when F was a combinatorial pseudo-triangle.

Therefore, the c.p.p.t. assignment for the angles of G_{n-1} can be extended to G_n, as we wanted to prove. □

Concerning complexity, the following lemma, taken from [39] proposes an efficient method of computing a combinatorial pointed pseudo-triangulation:

Lemma 4.4.4 *One can compute a c.p.p.t. assignment for a planar Laman graph in $O(n^{3/2})$ time by solving a bipartite (multi-)matching problem.*

4.5 Infinitesimally rigid graphs admit c.p.t. labelings.

In this last section we prove that, actually, every planar infinitesimally rigid graph admits a labeling as combinatorial pseudo-triangulation. However, we leave open

the problem of whether this c.p.t. labeling can be achieved having the generalized Laman property:

Lemma 4.5.1 *Every planar rigid graph G on $n \geq 3$ vertices has a labeling as combinatorial pseudo-triangulation. The number of non-pointed vertices is $e - (2n-3)$, for e the number of edges.*

Proof: The proof goes by induction on the number of edges: For the base case, start with a minimally rigid spanning subgraph (basis). We already know how to label angles of such a graph to obtain a combinatorial pointed pseudo-triangulation (Theorem 4.4.1). At each inductive step, we have to add one of the remaining edges ij which splits a face F into F_1 and F_2. Let us call α_{i1}, α_{j1} the angles at p_i, p_j in F_1 and α_{i2}, α_{j2} those in F_2.

First, we have to consider the case in which p_ip_j splits the outer face F into the new outer face F_1 and an interior one F_2. Since G has no multiple edges, there has to be a point on the boundary of F lying between p_i and p_j, in F_2 now, and whose angle in F was labeled "big". We switch this label to "small" and label also small the angles α_{i2}, α_{j2}, big α_{i1}, α_{j1}. See Figure 4.8, where a label switched from big to small is represented by an arc traversing a dot.

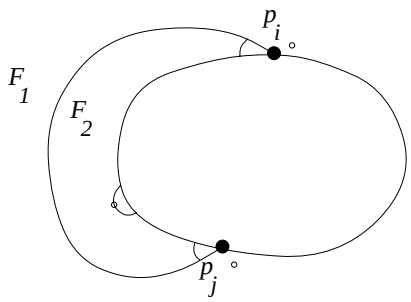

Figure 4.8: c.p.t. labeling when F was the outer face.

Otherwise, p_ip_j separates a combinatorial pseudo-triangle F into two faces F_1, F_2 which we want to be combinatorial pseudo-triangles as well. We again have three cases, depicted in Figure 4.9, according to whether both, one or none of the endpoints p_i, p_j were combinatorial corners of the combinatorial pseudo-triangle F:

- In the first case, the new edge joins two combinatorial corners and separates F into F_1 containing the third one and F_2 containing the opposite pseudo-edge, which since there are not multiple edges, must have a point different from p_i, p_j and with "big" label. We re-label this as small and assign also "small" labels to the four angles $\alpha_{i1}, \alpha_{i2}, \alpha_{j1}, \alpha_{j2}$.

- The new edge p_ip_j in the second case goes from a combinatorial corner, say p_i, to a pseudo-edge and two subcases arise according to the position of the

two other corners: First, that p_ip_j separates F into F_1 and F_2 containing one corner (different of p_i) each. We just have to label small the four new angles. Second, that F is separated into F_1 containing the other two corners and F_2 with no corner (appart from p_i). As above, there is a point different from p_i, p_j and labeled big in the boundary of F_2. We switch its label and label small also the angles α_{i1}, α_{i2} and α_{j2}, and big α_{j1}.

- In the last case the edge joins two pseudo-edges and there are again two subcases: On the one hand, p_ip_j can separate F into F_1 containing the three corners and F_2 containing none of them. Again, a point different from p_i, p_j and labeled big appears in the boundary of F_2. We also switch its label, we label small α_{i2}, α_{j2} and big α_{i1}, α_{j1}. On the other hand, p_ip_j can separate F into F_1 and F_2 containing one and two of the old corners. We just have to label small the two new angles in F_1 and exactly one of those in F_2, for example $\alpha_{i1}, \alpha_{i2}, \alpha_{j1}$ small and α_{j2} big.

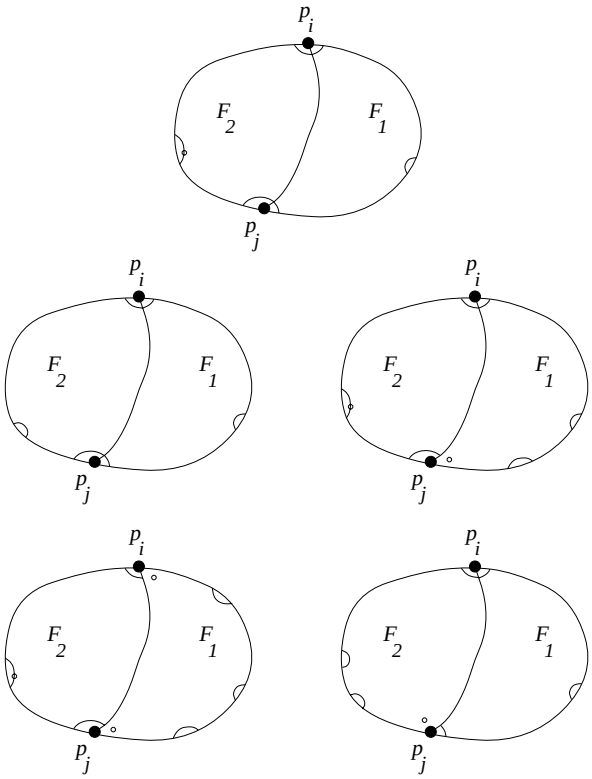

Figure 4.9: c.p.t. labeling when F was a combinatorial pseudo-triangle.

Note that at every case a big angle needs to be sacrificed in order to have 3 small angles in each new face. That is, exactly one pointed vertex loses the incident big angle and becomes non-pointed, as happens when adding a new edge in true

pseudo-triangulations . □

It is not clear from the previous proof how to choose the k pointed and l non-pointed vertices in such a way that the resulting c.p.t. has the generalized Laman property: When a "big" label is switched to "small", a pointed vertex becomes non-pointed but keeps its number of incident edges. However, we believe the following to be true:

Conjecture 4.5.2 *Every planar rigid graph G on $n \geq 3$ vertices has a labeling as* generalized Laman *combinatorial pseudo-triangulation, with $e-(2n-3)$ non-pointed vertices, for e the number of edges.*

This would have as a consequence the following generalization of Corollary 4.1.2, which in particular states that planar and rigid graphs (even if they are not minimally rigid) could be embedded as pseudo-triangulations:

Conjecture 4.5.3 *Given a graph G, the following conditions are equivalent:*

(i) G is planar and rigid,

(ii) G has a generalized Laman c.p.t. labeling,

(iii) G can be embedded as a pseudo-triangulation.

The implication (i)⇒(ii) is the previous conjecture. (ii)⇒(iii) follows from Theorem 4.2.4, and (iii)⇒(i) (as well as (iii)⇒(ii)) from Theorem 3.1.3.

Chapter 5

Rigidity circuits, reciprocal diagrams and pseudo-triangulations

The goal of the present chapter is the study of conditions for the reciprocal of a non-crossing spherical framework $G(S)$ to be non-crossing as well. We give a complete description of the non-crossing spherical frameworks whose reciprocals are non-crossing, in terms of the shape of the faces in the geometric embedding and their signs in the stress. As a consequence, we prove that if the geometric part of a non-crossing spherical framework is a rigidity circuit pseudo-triangulation with exactly one non-pointed vertex, the reciprocal is non-crossing and its geometric part is again a rigidity circuit pseudo-triangulation.

The story of this proof is as follows: the statement was conjectured by Brigitte Servatius during the *Barbados Workshop on Rigidity and Scene Analysis* [1] attended, among others, by six of the authors of [39]. After the workshop, this problem (and that of Chapter 4) were discussed in an informal electronic forum by the nine authors of [39]. The proof presented in this chapter may be difficult to attribute to any specific proper subset of them.

5.1 Introduction

Stresses and forces on graphs and their relations with reciprocal figures were a notorious contribution of J.C. Maxwell [52] to Rigidity Theory, as we mentioned in the introductory chapter. In the last decade, Crapo and Whiteley have resumed the topic and extended the correspondences in Theorem 1.3.11 to the full vector spaces of self stresses, reciprocals and spatial liftings of a plane drawing, see [23] and [24]. Following along the lines of the previous chapter, we present here another relation between a class of graphs and their planarity.

[1] `http://cs.smith.edu/~streinu/Workshops/barbados02.html`

Let us recall the following basic notions which we have already introduced in Section 1.3: A non-zero *stress* (or *self-stress*) on a framework $G(S)$ is an assignment of scalars w_{ij}, not all zero, to the edges E of G in such a way that, for every vertex $i \in V$,

$$\sum_{ij \in E} w_{ij}(p_i - p_j) = 0.$$

Equivalently, it is a dependence relation $\sum_{ij \in E} w_{ij} r_{ij} = 0$ between rows of the rigidity matrix. Recall as well that a *circuit* is a minimal dependent set of rows in the rigidity matrix of any generic embedding. Hence, a framework $G(S)$ is a *rigidity circuit* precisely if it has a unique stress and this is non-zero on all edges.

Definition 5.1.1 A *spherical framework* is a pair $(G \hookrightarrow \mathbb{S}^2, G(S))$ composed of a plane topological embedding $G \hookrightarrow \mathbb{S}^2$ and a (possibly not plane) geometric embedding $G(S)$, of a graph G.

A *non-crossing spherical framework* is a plane geometric embedding $G(S)$ together with its trivially corresponding topological embedding $G \hookrightarrow \mathbb{S}^2$.

We have already seen in Definition 1.3.10 that the notion of dual graph introduced by Poincarè, under the geometric restriction of perpendicularity of edges, leads to:

Definition 5.1.2 A spherical framework $(G' \hookrightarrow \mathbb{S}^2, G'(S'))$ is a *reciprocal* of a spherical framework $(G \hookrightarrow \mathbb{S}^2, G(S))$ if $G' \hookrightarrow \mathbb{S}^2$ is the dual of $G \hookrightarrow \mathbb{S}^2$ and for each edge patch $(i, j; h, k) \in \underline{E}$,

$$\langle p_i - p_j, p'_h - p'_k \rangle = 0.$$

Reciprocal spherical frameworks turn out to exist precisely for those $G(S)$ having a stress nonzero on all edges (Maxwell's Theorem 1.3.11). However, even for non-crossing spherical frameworks the reciprocal is not guaranteed to be a *non-crossing* spherical framework (Remark 1.3.14). For an example, consider the framework in Figure 5.1 (left) together with the stress specified in Figure 1.7. Observe that the position in the reciprocal (right) of the point a corresponding to the exterior face coincides with that of the point e which corresponds to the innermost face. Thus, for example the edges ab and be coincide and therefore the reciprocal has crossings.

In this chapter we focus on non-crossing spherical frameworks, which we regard as completely defined by their geometric part $G(S)$. Since the existence of plane embeddings has attracted so much interest in Graph Drawing (see Section 1.2), our aim in this chapter is to provide a complete characterization of the graphs in this class whose reciprocal $G'(S')$ is again a non-crossing spherical framework, that is to say, a plane geometric embedding. This is carried out in Sections 5.2 and 5.3 (Theorems 5.2.1 and 5.2.3) The characterization is in terms of the faces of $G(S)$ and the signs of the self-stress.

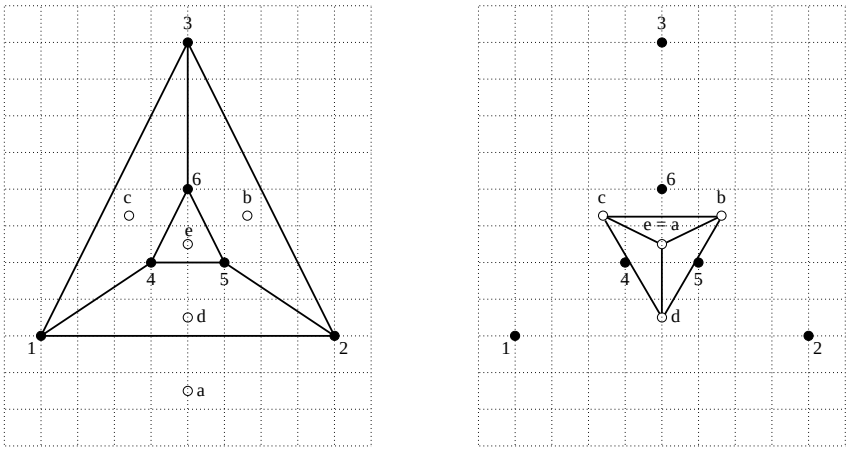

Figure 5.1: Non-crossing spherical framework whose reciprocal has crossings.

Then, in Section 5.4 we turn our attention to those non-crossing spherical frameworks $G(S)$ with a *unique* stress non-zero on all edges, that is, those which are rigidity circuits. We first prove that, in case the reciprocal $G'(S')$ is also non-crossing, it has the same numbers of pseudo-triangles and pseudo-quadrilaterals as the original. Finally, we prove that in a pseudo-triangulation which is a rigidity circuit, the unique stress fulfills the conditions for the reciprocal to be non-crossing, and this turns out to be again a pseudo-triangulation:

Theorem 5.1.3 *The reciprocal of a rigidity circuit pseudo-triangulation with a non-pointed vertex is non-crossing and again a rigidity circuit pseudo-triangulation.*

This result will be a consequence of Theorem 5.4.2 and the above referred characterization. Together with the main result about circuits in Section 5 of [39], which states that any planar generic rigidity circuit can be embedded as a pseudo-triangulation with one non-pointed vertex, Theorem 5.1.3 implies:

Corollary 5.1.4 *Let G be a graph with $2n-2$ edges and minimally dependent in the sense of generic 2-dimensional rigidity. The following are equivalent for G:*

(i) G is planar.

(ii) G can be embedded as a pseudo-triangulation with one non-pointed vertex.

(iii) G can be planarly embedded with a non-crossing reciprocal. (Which, in turn, gets embedded as a pseudo-triangulation with one non-pointed vertex).

Proof: (i)⇒(ii) The previously mentioned result in [39]. (ii)⇒(iii) The above theorem. (iii)⇒(i) Obvious. □

5.2 Non-crossing spherical frameworks with non-crossing reciprocals

From now on, we consider a given configuration S and we denote by G and G' the reciprocal non-crossing spherical frameworks. We are going to work on the geometric embeddings, but having in mind that they are part of spherical frameworks. We assume our geometric graph G to be non-crossing and look for conditions producing a non-crossing reciprocal G'. First, observe that a reciprocal is non-crossing precisely if the face cycles in the original geometric graph G are the vertex cycles in the reciprocal geometric graph G' and vice-versa. For an example see Figure 5.1, where the coincidence of the points corresponding to faces a and e makes the face cycle $beda$ around face c in the original not to be a vertex cycle around the corresponding vertex in the reciprocal.

Then, the reciprocity can in principle be "orientation-preserving" and "orientation-reversing", meaning that from the face cycles of the original to the vertex cycles of the reciprocal the orientation has to be preserved or reversed. But either all orientations are preserved or all are reversed, because the two reciprocal geometric embeddings can be understood as the same decomposition of a 2-sphere as spherical frameworks. Here comes the claimed characterization, in terms of faces and signs of the self-stress, of non-crossing spherical frameworks whose corresponding reciprocal is also non-crossing:

Theorem 5.2.1 *Let G be the geometric embedding of a non-crossing spherical framework, with a stress nonzero on all edges. Then G has a non-crossing reciprocal with the reversed orientation if, and only if, the following holds:*

1. *The complement of the outer face of G is convex and all its edges have the same sign.*

2. *All the interior faces of G are either pseudo-triangles or pseudo-quadrilaterals, with one of the following sign patterns for their edges:*

   2.*a* *A pseudo-triangle in which all the edges in the same pseudo-edge have the same sign and one pseudo-edge has opposite sign to the other two. In other words, there are two sign changes along the boundary of the pseudo-triangle and both occur at corners.*

   2.*b* *A pseudo-triangle with four sign changes along the boundary, three of which occur at the three corners.*

   2.*c* *A pseudo-quadrilateral in which all the edges in the same pseudo-edge have the same sign and consecutive pseudo-edges have opposite signs.*

Moreover, in the reciprocal diagram these four situations correspond respectively to:

1′. *A non-pointed vertex with all the incident edges having the same sign.*

2.a' *A pointed vertex with two sign changes around it, none of which occurs at the big angle.*

2.b' *A pointed vertex with four sign changes around it, one of which occurs at the big angle.*

2.c' *A non-pointed vertex with four sign changes around it.*

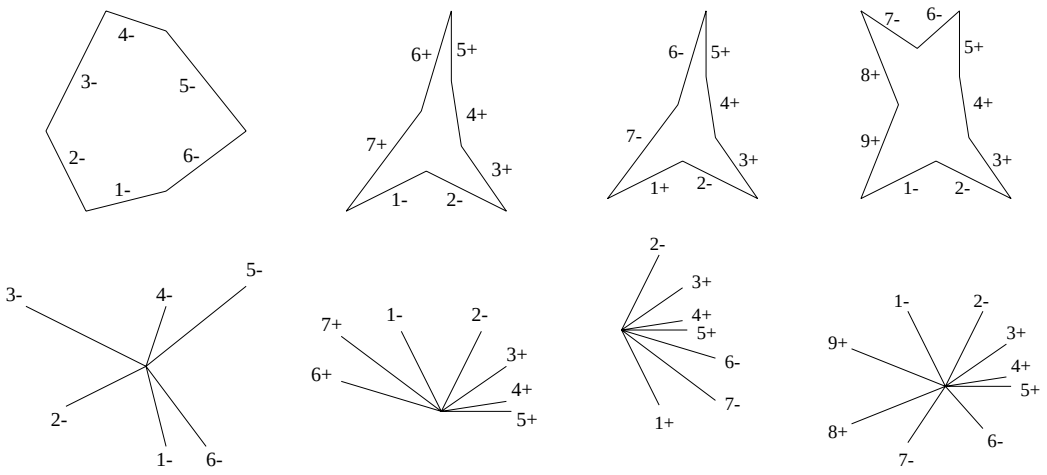

Figure 5.2: Top: Patterns 1, 1.a, 2.b, 2.c. Bottom: Patterns 1', 1.a', 2.b', 2.c'.

In Figure 5.2 we show the conditions for G (top row) to have a non-crossing reciprocal G' with the reversed orientation (bottom row). Recall from the Observation 1.3.12 after the proof of Maxwell's Theorem 1.3.11 that the direction of the reciprocal of an edge is given by its sign. Except for the observation below concerning the outer face, the edges in the original are labeled counter-clockwise, therefore those in the orientation-reversing reciprocal are labeled clockwise:

Observation 5.2.2 The orientation given to the boundary cycle of the outer face is the opposite to the "expected one"; if the outer face is oriented counter-clockwise (meant locally, i.e. around each point), then its boundary (in any reasonable definition, for example the one in homology) will be oriented clockwise.

This is why the conditions for the outer face in Theorem 5.2.1 become conditions for interior faces in the following statement, and vice-versa:

Theorem 5.2.3 *Let G be as in the previous theorem. Then, it is impossible for its reciprocal to be non-crossing and with the same orientation, since for this the two following properties have to hold simultaneously:*

(i) All the interior faces are convex and have no sign changes along their respective boundaries. In particular all the edges in G have the same sign.

(ii) The complement of the outer face is either a pseudo-triangle or a pseudo-quadrilateral, with one of the sign patterns 2.a, 2.b or 2.c. In particular, it has sign changes.

The reader can check the fulfillment and failure of the conditions for the planarity of the reciprocal respectively in Figure 1.8 (stress given in Example 1.3.13) and Figure 5.1 (stress given in Figure 1.7).

5.3 Proof of Theorems 5.2.1 and 5.2.3

Since the proofs are composed of several intermediate results, we collect them in this section. We assume given the G with a self-stress nonzero on all edges, which uniquely induces a drawing of the reciprocal G'. Recall from Observation 1.3.12 that the edges of G' are perpendicular to those of G, scaled by the absolute value of the stress and with direction given by the sign. The condition for the reciprocal to be non-crossing is that the edges around each vertex of the reciprocal graph G' appear in the same order as in the corresponding face of G (if we want to preserve orientation) or in reverse order (if we want to reverse it).

Let us rewrite this condition in terms of a sum of angles: Let F be a face of G, and let $e_1, \ldots, e_k$ be its edges, in the ordering of the "positive" orientation of the euclidean plane. This means counter-clockwise for bounded faces and clockwise for the unbounded one. Let $e'_1, \ldots, e'_k$ be the corresponding reciprocal edges around the reciprocal vertex v_F. Let α_i be the angle between e'_i and e'_{i+1}, measured in the "negative" (clockwise) direction. Then, the reciprocal G' will be considered "non-crossing and with the reversed orientation" at v_F if and only if $\sum_{i=1}^{n} \alpha_i = 2\pi$. Analogously, let $\beta_i = 2\pi - \alpha_i$ be the same angle, but measured in counterclockwise direction. The reciprocal will be considered "non-crossing and with the same orientation as G" at v_F if and only if $\sum_{i=1}^{n} \beta_i = 2\pi$. The whole reciprocal is non-crossing with the same (resp., reversed) orientation if it is non-crossing with the same (resp. reversed) orientation at all the vertices.

Now we have to relate the angles α_i to the corresponding angles in G. Let ϕ_i be the angle between e_i and e_{i+1} measured in the interior of the face F (that is, in clockwise order). Then:

Lemma 5.3.1 *If e_i and e_{i+1} have opposite signs, then $\alpha_i = \phi_i$. If e_i and e_{i+1} have the same sign, then $\alpha_i = \phi_i \pm \pi$, where the plus or minus is there to guarantee α_i to be between 0 and 2π. That is, $\alpha_i = \phi_i + \pi$ at convex angles of F and $\alpha_i = \phi_i - \pi$ at reflex angles of F.*

Proof: Make a picture, or see Figure 5.3. □

According to this lemma, the distribution of signs in the stress of the face F that minimizes $\sum_{i=1}^{n} \alpha_i$ is to give opposite sign to consecutive edges at convex vertices of F (top-left picture in Figure 5.3) and same signs to consecutive edges at reflex vertices of F (bottom-right). This suggests the following definition:

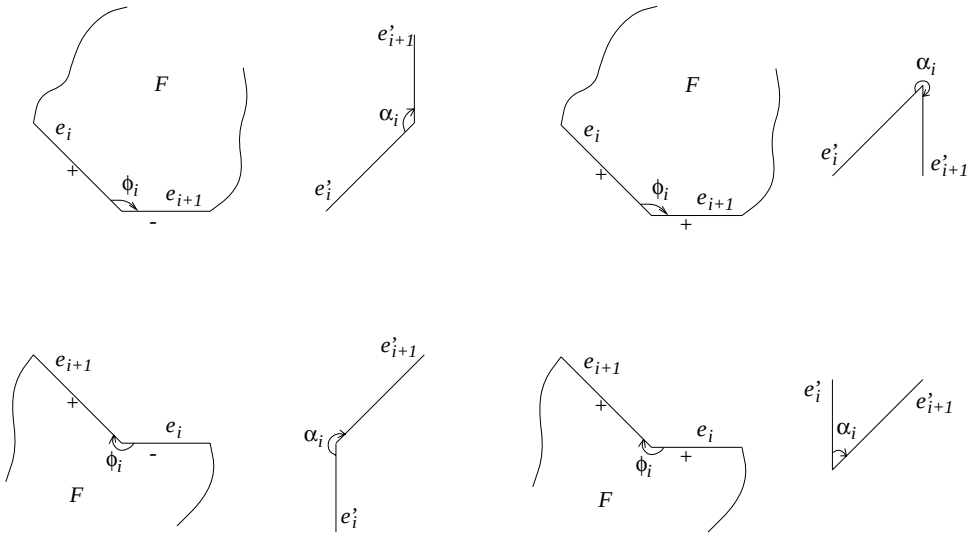

Figure 5.3: Relations between α_i and ϕ_i.

Definition 5.3.2 For a plane framework G with a given self-stress (or at least, a given assignment of signs to its edges) we call angles of *type 1* the big angles ($> \pi$) whose incident edges have the same sign and the small angles ($< \pi$) with sign change (bottom-right and top-left in the figure). We call *type 2* the other two cases (bottom-left and top-right).

For an example, the angles $\widehat{1,2}$ and $\widehat{2,3}$ in the second picture of Figure 5.2 (top) are of type 1. On the other hand, the angles $\widehat{1,2}$ in the third picture and $\widehat{5,6}$ in the second one are of type 2.

Lemma 5.3.3 *Given a bounded face F which is a pseudo-polygon with n edges, k corners and with n_1 and n_2 angles of types 1 or 2 respectively, we have:*

(i) The sum of angles in the reciprocal vertex, measured "reversing orientations" equals $\sum_{i=1}^{n} \alpha_i = (k-2)\pi + n_2\pi$.

(ii) The sum of angles in the reciprocal vertex, measured "preserving orientations" equals $\sum_{i=1}^{n} \beta_i = (n-k+2)\pi + n_1\pi$.

Proof: The total sum of interior angles of F is $\sum_{i=1}^{n} \phi_i = (n-2)\pi$, since that of the k-gon in which it is inscribed is $(k-2)\pi$ and each of the $n-k$ reflex angles adds π. According to Lemma 5.3.1, $\sum_{i=1}^{n} \alpha_i$ equals this amount plus π times the number of convex angles of type 2 and minus π times the reflex angles of type 1. That is to say, $\sum_{i=1}^{n} \alpha_i = (n-2)\pi + n_2\pi - (n-k)\pi$, as stated (note: when adding $n_2\pi$ we add an extra π for each reflex angle of type 2, but the same amount is substracted for that angle in $-(n-k)\pi$, being $n-k$ the number of reflex angles of both types).

This proves part (i). Part (ii) follows from $\alpha_i + \beta_i = 2\pi$ and $n = n_1 + n_2$. □

Corollary 5.3.4 *Suppose that the unbounded face F is a pseudo-polygon with n edges, and with n_1 and n_2 angles of types 1 or 2 respectively. Let k be its number of corners, meaning by this the non-corners of the complement of F. (That is to say, if the boundary of the framework is convex then $k = 0$). Then we have:*

(i) The sum of angles in the reciprocal vertex, measured "reversing orientations" equals $\sum_{i=1}^{n} \alpha_i = (k+2)\pi + n_2\pi$.

(ii) The sum of angles in the reciprocal vertex, measured "preserving orientations" equals $\sum_{i=1}^{n} \beta_i = (n-k-2)\pi + n_1\pi$.

Proof: Same proof as Lemma 5.3.3, except that now "interior angle" means the opposite and, in particular, $\sum_{i=1}^{n} \phi_i = (n+2)\pi$. □

Taking into account that the number of corners is $3 \leq k \leq n$ for an interior face and $0 \leq k \leq n-3$ for the boundary face, making $\sum_{i=1}^{n} \alpha_i = 2\pi$ (respectively, $\sum_{i=1}^{n} \beta_i = 2\pi$) in Lemma 5.3.3 and Corollary 5.3.4 gives Theorem 5.2.1 (respectively, Theorem 5.2.3).

5.4 Rigidity circuits. The case of pseudo-triangulations

In this last section we focus on the class of those non-crossing spherical frameworks which are rigidity circuits. These have a *unique* stress non-zero on all edges and hence a *unique* reciprocal according to the construction in Maxwell's Theorem. They also have the following property: Observe that the removal of an edge from a rigidity circuit produces an isostatic graph, which has $2|V| - 3$ edges by Lemma 1.3.20. Hence, it is clear that in order to be a rigidity circuit, a graph must have $2|V| - 2$ edges. In addition, the stress being non-zero on all edges and the construction of the reciprocal imply that this has the same number of edges (although some of them may overlap if the reciprocal is not crossing-free, as in Figure 5.1).

As a first consequence of Theorems 5.2.1 and 5.2.3 we prove:

Proposition 5.4.1 *Let G be a rigidity circuit non-crossing spherical framework. If the reciprocal G' is also non-crossing, then both have the same number of pseudo-triangles and pseudo-quadrilaterals.*

Proof: On the one hand, the characterization in the theorems says that all the faces of G are pseudo-triangles and pseudo-quadrilaterals. This holds as well for G', since G is non-crossing and its reciprocal. We use Euler's formula $n - e + t + q = 1$ for G, where n, e, t and q are the numbers of vertices, edges, pseudo-triangles and pseudo-quadrilaterals, respectively. A double counting of convex angles as in the proof of

Proposition 1.4.7 gives $2e - n_\epsilon = 3t + 4q$, where n_ϵ is the number of pointed vertices. This, together with the above formula implies $e = 3n - 3 - n_\epsilon - q = 2n - 3 + n_\gamma - q$, for n_γ the number of non-pointed vertices. Since a rigidity circuit has $2n - 2$ edges, we get $q = n_\gamma - 1$. Analogously, we obtain for G' the relation $q' = n'_\gamma - 1$. Finally, the theorems imply $n'_\gamma = q + 1$, so $q' = q$ and the coincidence also in the number of pseudo-triangles is then trivial. □

Let us consider now the case of G being a rigidity circuit pseudo-triangulation. Since this is trivially a non-crossing spherical framework, it is a natural question whether its reciprocal is also non-crossing or not. The next result gives a positive answer, and has Theorem 5.1.3 as a straightforward consequence using Theorems 5.2.1 and 5.2.3:

Theorem 5.4.2 *The unique self-stress in any rigidity circuit pseudo-triangulation with a non-pointed vertex has the sign pattern described by* 1*,* 2.*a and* 2.*b in Theorem 5.2.1. Moreover, the edges in the boundary have opposite sign to those incident to the non-pointed vertex.*

The reciprocal diagram is again a pseudo-triangulation with a non-pointed vertex.

The proof will follow from the next two lemmas:

Lemma 5.4.3 *Along each of the $n - 1$ pseudo-triangles of G there is at least one angle of type 2.*

Proof: If all the angles are of type 1, then along concave chains the same sign holds and a sign change appears at the three corners, which is impossible. □

Lemma 5.4.4 *At each of the $n - 1$ pointed vertices there are at least three angles of type 1.*

Proof: The number of angles of type 1 at a pointed vertex must be odd, because the number of sign changes around it is even. And only one angle of type 1 would mean that either all the edges have the same sign, or there are only two sign changes, one of them at the big angle. Both situations contradict the equilibrium condition for a self-stress. □

Corollary 5.4.5 *In the unique self-stress of any pseudo-triangulation circuit there are exactly $n - 1$ angles of type 2, one in each of the bounded faces, and exactly $3n - 3$ angles of type 1, three of them at each pointed vertex.*

In particular, there are no sign changes at the non-pointed vertex and all the edges of the convex hull have the same sign.

Proof: The total number of angles of type 2 is at least $n-1$, by Lemma 5.4.4 and that of type 1 is at least $3n-3$ by Lemma 5.4.3. But the total number of angles is $4n-4$, because the pseudo-triangulation has 2n-2 edges. In particular, there is no room for angles of type 1 at the non-pointed vertex, or of type 2 in the unbounded face. □

Proof of Theorem 5.4.2: Conditions 1, 2.a and 2.b are exactly the translation of Corollary 5.4.5 (2.a and 2.b describe, respectively, a pseudo-triangle with its type 2 angle at a corner and a non-corner, respectively; this can be checked in Figure 5.2).

To prove that the boundary and the non-pointed vertex get opposite signs we can use Maxwell-Cremona lifting theorem (equivalence (ii)⇔(iii) in Theorem 1.3.11), which implies that the stress induces a lifting with positive edges being ridges and negative edges being valleys. Suppose without loss of generality that the boundary edges are negative. It is easy to observe that a pointed vertex cannot be a maximum height point of the lifting. Hence the maximum height must be attained at the non-pointed vertex, which implies this vertex is incident to some ridges.

Finally, it can be deduced that the reciprocal diagram is a pseudo-triangulation, as follows: by the second part of Theorem 5.2.1 the reciprocal diagram has exactly one non-pointed vertex. In particular, it has exactly $n-1$ big angles and $3n-3$ small angles. Since a planar graph with n vertices and $2n-2$ edges has $n-1$ bounded faces, and since each bounded face needs at least three small angles (corners), we conclude that the $n-1$ bounded faces are pseudo-triangles and the boundary is convex. □

Bibliography

[1] P. Agarwal, J. Basch, L. Guibas, J. Hershberger, and L. Zhang. Deformable free space tilings for kinetic collision detection. In B. R. Donald, K. Lynch, and D. Rus, editors, *Algorithmic and Computational Robotics: New Directions, 4th Int. Workshop on Alg. Found. Robotics (WAFR)*, pages 83–96. A K Peters, 2000.

[2] O. Aichholzer, F. Aurenhammer, H. Krasser, and B. Speckmann. Convexity minimizes pseudo-triangulations. In *Proceedings of the 14th Canadian Conference on Computational Geometry*, pages 158–161, August 2002.

[3] O. Aichholzer, B. Speckmann, and I. Streinu. The path of a pseudo-triangulation. Manuscript, 2002.

[4] O. Aicholzer, F. Aurenhammer, and H. Krasser. Pseudo-triangulations from surfaces and a novel type of flip. Manuscript, 2002.

[5] M. Ajtai, V. Chvátal, M. M. Newborn, and E. Szemerédi. Crossing-free subgraphs. In G. S. A. Rosa and J. Turgeon, editors, *Theory and practice of combinatorics*, volume 60 of *North-Holland Math. Stud.*, pages 9–12. North-Holland, Amsterdam, 1982.

[6] R. Anderson and R. B. Hughes. Simplexity of the cube. *Discrete Math.*, 158:99–150, 1996.

[7] D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Appl. Math.*, 65:21–46, 1996.

[8] M. L. Balinski. On the graph structure of convex polyhedra in n-space. *Pacific J. Math.*, 11:431–434, 1961.

[9] G. D. Battista, P. Eades, R. Tamassia, and I. Tollis. *Graph Drawing.* Prentice Hall, 1999.

[10] A. Below, J. A. de Loera, and J. Richter-Gebert. The complexity of finding small triangulations of convex 3-polytopes. Preprint, 2000. Extended abstract *Finding minimal triangulations of convex polytopes is NP-hard* in Proceedings of SODA 2000.

[11] S. Bespamyatnikh. Enumerating pseudo-triangulations in the plane. In *Proceedings of the* 14*th Canadian Conference on Computational Geometry*, pages 162–166, August 2002.

[12] K. Bezdek and R. Connelly. The Kneser–Poulsen conjecture. To appear in J. reine ang. Math., 2002.

[13] L. J. Billera, R. Cushman, and J. Sanders. The Stanley decomposition of the harmonic oscillator. *Nederl. Akad. Wetensch. Proc. Ser. A*, 91:375–393, 1988.

[14] L. J. Billera, P. Filliman, and B. Sturmfels. Constructions and complexity of secondary polytopes. *Adv. Math.*, 83:155–179, 1990.

[15] B. Bollobás. *Graph Theory*, volume 63 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1979.

[16] G. Brightwell and D. B. West. Partially ordered sets. In K. H. Rosen, editor, *Handbook of Discrete and Combinatorial Mathematics*, chapter 11, pages 717–752. CRC Press, New York, 2000.

[17] H. Brönnimann, L. Kettner, M. Pocchiola, and J. Snoeyink. Counting and enumerating pseudo-triangulations with the greedy flip algorithm. Manuscript, 2001.

[18] A. L. Cauchy. Sur les polygons et les polyèdres. In *Oevres Complètes d'Augustin Cauchy*, 2è Série, pages 26–38, 1905.

[19] B. Chazelle, H. Edelsbrunner, M. Grigni, L. Guibas, J. Erschberger, and M. Sharir. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12:54–68, 1994.

[20] M. Chrobak, M. Goodrich, and R. Tamassia. Convex drawings of graphs in two and three dimensions. In *Proc. 12th Ann. Sympos. Comput. Geom.*, pages 319–328, 1996.

[21] R. Connelly. A counter example to the rigidity conjecture for polyhedra. *Inst. Haut. Étud. Sci. Publ. Math.*, 47:333–335, 1978.

[22] R. Connelly, E. D. Demaine, and G. Rote. Straightening polygonal arcs and convexifying polygonal cycles. *Discrete and Computational Geometry*, Bericht B 02-02:49 pages, February 2002. Preliminary version in Proc. 41st FOCS, 2000, pp. 432-452.

[23] H. Crapo and W. Whiteley. Plane self stresses and projected polyhedra I: the basic pattern. *Structural Topology*, 20:55–78, 1993.

[24] H. Crapo and W. Whiteley. Spaces of stresses, projections and parallel drawings for spherical polyhedra. *Beiträge zur Algebra und Geometrie/Contributions to Algebra and Geometry*, 35(2):259–281, 1994.

[25] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[26] J. A. de Loera, S. Hoşten, F. Santos, and B. Sturmfels. The polytope of all triangulations of a point configuration. *Doc. Math. J. DMV*, 1:103–119, 1996.

[27] J. A. de Loera and S. Peterson. Universalbuilder. Software available at `http://plane.math.ucdavis.edu/~deloera/RECENT_WORK/recent.html`

[28] E. C. de Verdière, M. Pocchiola, and G. Vegter. Tutte's barycenter method applied to isotopies. In *Proceedings of the 13th Canadian Conference on Computational Geometry*, August 2001.

[29] M. Dehn. Über die Starrheit konvexer Polyeder. *Math. Ann.*, 77:466–473, 1916.

[30] S. Eilenberg and N. E. Steenrod. *Foundations of algebraic topology*. Princeton University Press, 1952.

[31] L. Euler. Solutio problematis ad geometriam situs pertinentis. *Commetarii Academiae Scientiarum Imperialis Petropolitanae*, 8:128–140, 1736.

[32] L. Euler. Opera postuma I. *Petropoli*, pages 494–496, 1862.

[33] I. Fary. On straight lines representation of planar graphs. *Acta Sci. Math. Szeged*, 11:229–233, 1948.

[34] P. Flajolet and M. Noy. Analytic combinatorics of non-crossing configurations. *Discrete Math.*, 204:203–229, 1999.

[35] K. Fukuda. Cdd+. Software available at `http://www.cs.mcgill.ca/~fukuda/soft/cdd_home/cdd.html`

[36] I. Gel'fand, M. Kapranov, and A. Zelevinsky. *Multidimensional Determinants, Discriminants and Resultants*. Birkhäuser, Boston, 1994.

[37] M. Goodrich and R. Tamassia. Dynamic ray shooting and shortest paths in planar subdivisions via balanced geodesic triangulations. *J. Algorithms*, 23:51–73, 1997.

[38] J. Graver, B. Servatius, and H. Servatius. *Combinatorial Rigidity*, volume 2 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 1993.

[39] R. Haas, D. Orden, G. Rote, F. Santos, B. Servatius, H. Servatius, D. Souvaine, I. Streinu, and W. Whiteley. Planar minimally rigid graphs and pseudo-triangulations. In *Proc. 19th ACM Symposium on Computational Geometry (SoCG)*, 2003 (To appear).

[40] M. Haiman. A simple and relatively efficient triangulation of the n-cube. *Discrete Comput. Geom.*, 6:287–289, 1991.

[41] L. Henneberg. Die graphische Statik der starren Systeme, Leipzig, 1911. Johnson Reprint 1968.

[42] D. Hilbert and S. Cohn-Vossen. Kinematics. In *Geometry and the Imagination*, chapter 5. Cheslea, New York, 1952.

[43] J. E. Hopcroft and P. J. Kahn. A paradigm for robust geometric algorithms. *Algorithmica*, 7(4):339–380, 1992.

[44] B. Huber, J. Rambau, and F. Santos. The Cayley trick, lifting subdivisions and the Bohne-Dress theorem on zonotopal tilings. *J. Eur. Math. Soc. (JEMS)*, 2(2):179–198, 2000.

[45] R. B. Hughes. Minimum-cardinality triangulations of the d-cube for $d = 5$ and $d = 6$. *Discrete Math.*, 118:75–118, 1993.

[46] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22(1):333–346, 1999.

[47] D. Jacobs, A. Rader, L. Kuhn, and M. Thorpe. Protein flexibility predictions using graph theory. *Proteins*, 44:150–165, 2001.

[48] L. Kettner, D. Kirkpatrick, A. Mantler, J. Snoeyink, B. Speckmann, and F. Takeuchi. Tight degree bounds for pseudo-triangulations of points. *Computational Geometry: Theory and Applications*, 2002 (to appear). Extended abstract in Proc. 13th Canadian Conference on Computational Geometry, pp. 117–120, 2001.

[49] D. Kirkpatrick, J. Snoeyink, and B. Speckmann. Kinetic collision detection for simple polygons. *International Journal of Computational Geometry and Applications*, 12:3–27, 2002.

[50] G. Laman. On graphs and rigidity of plane skeletal structures. *J. Engrg. Math.*, 4:331–340, 1970.

[51] C. W. Lee. Subdivisions and triangulations of polytopes. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 14, pages 271–290. CRC Press, New York, 1997.

[52] J. C. Maxwell. On reciprocal figures and diagrams of forces. *Philos. Mag.*, 4(27):250–261, 1864.

[53] D. Orden and F. Santos. The polytope of non-crossing graphs on a planar point set. Preprint, February 2003. Available at `http://arxiv.org/abs/math.CO/0302126`

[54] D. Orden and F. Santos. Asymptotically efficient triangulations of the d-cube. Accepted for publication in Discrete Comput. Geom., January 2002, revised March 2003.

[55] J. Pach and G. Tóth. Monotone drawings of planar graphs. In P. Bose and P. Morin, editors, *Algorithms and Computation*, volume LNCS 2518, pages 647–653. Springer Verlag, 2002.

[56] M. Pocchiola and G. Vegter. Pseudo-triangulations: Theory and applications. In *Proc. 12th Ann. ACM Sympos. Comput. Geom.*, pages 291–300. ACM Press, May 1996.

[57] M. Pocchiola and G. Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete and Computational Geometry*, 16:419–453, 1996.

[58] M. Pocchiola and G. Vegter. The visibility complex. *Int. J. Comput. Geom. Appl.*, 6:279–308, 1996.

[59] D. Randall, G. Rote, F. Santos, and J. Snoeyink. Counting triangulations and pseudo-triangulations of wheels. In *Proceedings of the 13th Canadian Conference on Computational Geometry*, pages 149–152, August 2001. Available at `http://compgeo.math.uwaterloo.ca/~cccg01/proceedings`

[60] J. Richter-Gebert. *Realization spaces of polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1996.

[61] G. Rote, F. Santos, and I. Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. In J. P. B. Aronov, S. Basu and M. Sharir, editors, *Discrete and Computational Geometry – The Goodman-Polack Festschrift*, volume 25 of *Algorithms and Combinatorics*. Springer Verlag, Berlin, June 2003 (to appear).

[62] F. Santos and R. Seidel. A better upper bound on the number of triangulations of a planar point set. *J. Combin. Theory Ser. A*, 102, April 2003 (to appear). Available at `http://arXiv.org/abs/math.CO/0204045`

[63] W. Schnyder. Embedding planar graphs on the grid. In *Proc. ACM-SIAM Symp. Discrete Algorithms (SODA)*, pages 138–148, 1990.

[64] W. D. Smith. A lower bound for the simplexity of the n-cube via hyperbolic volumes. In K. Fukuda and G. M. Ziegler, editors, *Combinatorics of convex polytopes – European J. Combin.*, volume 21, chapter 1, pages 131–137. Springer Verlag, Berlin, 2000.

[65] B. Speckmann and C. D. Tóth. Allocating vertex π-guards in simple polygons via pseudo-triangulations. In *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms*, 2003 (To appear).

[66] E. Steinitz. Polyeder und Raumeinteilungen. *Encyclopädie der mathematischen Wissenchaften, Band 3 (Geometrie) Teil 3AB12*, pages 1–139, 1922.

[67] E. Steinitz and H. Rademacher. *Vorlesungen über die Theorie der Polyeder*. Springer-Verlag, 1934.

[68] I. Streinu. Combinatorial roadmaps in configuration spaces of simple planar polygons. To appear in S. Basu and L. González-Vega (eds.), Proc. DIMACS Wkshp. *Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science*.

[69] I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *Proc. 41st Ann. Symp. on Found. of Computer Science (FOCS 2000), Redondo Beach, California*, pages 443–453, 2000.

[70] T. S. Tay and W. Whiteley. Generating isostatic graphs. *Structural Topology*, 11:21–68, 1985.

[71] M. F. Thorpe and P. Duxbury, editors. *Rigidity Theory and Applications*. Kluwer Academic, 1999.

[72] M. J. Todd. *The Computation of Fixed Points and Applications*, volume 124 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, 1976.

[73] W. T. Tutte. Convex representations of graphs. *Proc. London Math. Soc.*, 10(38):304–320, 1960.

[74] W. T. Tutte. How to draw a graph. *Proc. London Math. Soc.*, 13(52):743–768, 1963.

[75] W. T. Tutte. *Graph Theory*, volume 21 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, 1984.

[76] B. von Hohenbalken. How to simplicially partition a polytope. Research Paper 79-17, Department of Economics, University of Alberta, Edmonton, 1979.

[77] D. B. West. *Introduction to Graph Theory*. Prentice Hall, Upper Saddle River, 1996.

[78] W. Whiteley. Motions and stresses of projected polyhedra (with a french translation). *Structural Topology*, 7:13–38, 1982.

[79] W. Whiteley. Infinitesimally rigid polyhedra I: statics of frameworks. *Trans. Amer. Math. Soc.*, 285:431–465, 1984.

[80] W. Whiteley. Some matroids from discrete applied geometry. In *Contemporary Mathematics*, pages 171–311. Amer. Math. Soc., 1996.

[81] W. Whiteley. Rigidity and scene analysis. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 49, pages 893–916. CRC Press, New York, 1997.

[82] W. Whiteley. Rigidity of molecular structures: Generic and geometric analysis. In M. F. Thorpe and P. Duxbury, editors, *Rigidity Theory and Applications*, pages 21–46. Kluwer Academic, 1999.

[83] P. Whittle. *Probability via Expectation*. Springer Texts in Statistics. Springer-Verlag, Berlin, 3 edition, 1992.

[84] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.