Dynamics and combinatorics of 2 and 3 dimensional triangulations

Younan, Maher Afif

UNIVERSITÉ DE GENÈVE
Section de Physique
Département de Physique théorique

FACULTÉ DES SCIENCES
Professeur Jean-Pierre ECKMANN
Professeur Peter WITTWER

Dynamics And Combinatorics Of 2 and 3 Dimensional Triangulations

THÈSE

présentée à la Faculté des sciences de l'Université de Genève
pour obtenir le grade de
Docteur ès sciences, mention physique

par

Maher YOUNAN
de
Maghdouché (Liban)

Thèse N° 4738

GENÈVE

Atelier d'impression de l'Université de Genève
2014

UNIVERSITÉ
DE GENÈVE

FACULTÉ DES SCIENCES

Doctorat ès sciences
Mention physique

Thèse de ***Monsieur Maher YOUNAN***

intitulée :

"Dynamics and Combinatorics of 2 and 3 Dimensional Triangulations"

La Faculté des sciences, sur le préavis des Messieurs J.-P. ECKMANN, professeur et directeur de thèse (Département de physique théorique), P. WITTWER, professeur et codirecteur de thèse (Département de physique théorique), Y. VELENIK, professeur (Section de mathématiques), C. RADIN, professeur (University of Texas, Department of Mathematics, Austin, USA) et P. COLLET, professeur (Ecole polytechnique, Centre de physique théorique, Palaiseau, France), autorise l'impression de la présente thèse, sans exprimer d'opinion sur les propositions qui y sont énoncées.

Genève, le 2 décembre 2014

Thèse - 4738 -

Le Doyen

N.B. - La thèse doit porter la déclaration précédente et remplir les conditions énumérées dans les "Informations relatives aux thèses de doctorat à l'Université de Genève".

Remerciements

Ces cinq dernières années de ma vie ont été très enrichissantes, et m'ont permis de grandir en tant que scientifique mais aussi en tant qu'être humain. Je tiens donc, avant tout autre propos, à consacrer ce chapitre à toutes celles et tous ceux qui m'ont aidé et accompagné durant ce parcours.

Tout d'abord, un très grand merci au Professeur Jean-Pierre Eckmann, mon directeur de thèse qui a partagé avec moi bien plus que ses connaissances scientifiques. Je n'oublierai pas de sitôt les pauses café où nous parlions de science mais aussi de politique, d'actualité et de religion. Le Professeur Jean-Pierre Eckmann a toujours été disponible, et je le remercie pour sa direction, ses conseils et toute l'aide qu'il m'a apportée le long de ces années.

Je tiens à remercier le Professeur Peter Wittwer pour avoir accepté d'être mon codirecteur de thèse ainsi que pour toutes les discussions enrichissantes que nous avons eues et tous les conseils et l'aide qu'il m'a apportés.

Je remercie le Professeur Charles Radin de l'Université de Texas, le Professeur Pierre Collet de l'École Polytechnique de Paris et le Professeur Yvan Velenik de l'Université de Genève pour avoir accepté de faire partie de mon jury de thèse et pour avoir partagé avec moi leurs connaissances lors de diverses discussions que nous avons eues durant ces dernières années.

Un merci particulier au Professeur Charles-Edouard Pfister de l'École Polytechnique Fédérale de Lausanne. C'est pendant ses cours que j'ai appris à aimer la physique mathématique.

Ce travail n'aurait pas vu le jour sans le soutien financier du Fonds National Suisse de la Recherche Scientifique (FNS) et du Conseil Européen de la Recherche (ERC).

Sur un plan plus personnel, je tiens à remercier le Département de Physique Théorique de l'Université de Genève, et en particulier les secrétaires Francine Gennai-Nicole et Cécile Jaggi-Chevalley, ainsi que l'ingénieur informatique Andreas Malaspinas.

Je souhaite remercier mes collègues et amis Cyrille Zbinden, Christoph Boeckle, Julien Guillod, Noé Cuneo, Andrea Agazzi et Jimmy Dubuisson pour toutes les discussions et les petites anecdotes qu'on a eues. C'est en grande partie grâce à eux que mon expérience à Genève fut aussi riche.

Je remercie toute ma famille qui m'a toujours soutenu, ma mère Lina, mon père Afif, ma sœur Nathalie et mon frère Nader. Mes pensées vont aussi à tous mes amis. Je tiens à citer Anca Stérie, Marine Vidaud, Ludovic Pirl et la famille Kambly. Sans leur soutien, mes études n'auraient probablement pas été aussi fructueuses.

Finalement, je tiens à remercier toutes les personnes qui liront ce document.

Résumé

Cette thèse traite des triangulations topologiques, *i.e.*, des décompositions en simplexes des sphères S^2 et S^3 de dimension 2 et 3. En particulier, nous nous intéressons à deux aspects de ces triangulations: leur nombre, qui est le sujet du deuxième et troisième chapitre, ainsi que la dynamique de verre qu'on peut avoir sur ces triangulations, traitée dans le premier et le dernier chapitre.

En mécanique statistique hors équilibre, un verre est défini comme étant un solide amorphe qui ne possède pas la périodicité des cristaux. Certains verres ont même une structure similaire à celle des liquides telle qu'on peut leur associer un coefficient de viscosité. Dans cette thèse, nous nous intéressons à un type de modèles de verre très particulier que nous appellons les *verres topologiques*. Ils sont définis de la façon suivante:

L'espace de phase est l'ensemble de toutes les triangulations de la sphère S^d de dimension $d = 2, 3$. Le cas $d = 2$ est résolu dans le premier chapitre tandis que le cas $d = 3$ est le sujet du dernier chapitre. Les mouvements élémentaires de la dynamique sont donnés par les mouvements de Pachner qui conservent le nombre de noeuds, à savoir le mouvement T1 en 2 dimensions et les mouvement 2-3 et 3-2 en 3 dimensions. Etant donné une triangulation A, son énergie est locale, donc une somme des contributions de chaque noeud v:

$$E = \sum_{v \in \mathcal{V}(A)} f(\mathcal{I}(v)) \ ,$$

où $v \in \mathcal{V}(A)$ est un noeud de A, $\mathcal{I}(v)$ est le voisinage de v, *i.e.*, la sous-triangulation de A contenant uniquement les voisins de v et $f(\cdot)$ est une fonction positive qui donne la contribution en énergie de v. Finalement, la dynamique est donnée par un algorithme de Metropolis dont le seul paramètre libre est la tempérture T.

Ces systèmes se comportent comme des verres; en particulier, leur dynamique ralentit énormément quand la tempérture T s'approche de zéro. Par exemple, nous observons que l'énergie s'approche de sa valeur stationnaire à $T \ll 1$ de façon polynomiale et non pas exponentielle. De plus, la fonction de corrélation du temps décroît avec un temps de relaxation τ_r, qu'on peut interpréter comme un coefficient de viscosité, qui croît exponentiellement en fonction de l'inverse de la température.

Dans le deuxième et troisième chapitre, nous étudions en détail une propriété de l'espace de phase 3d, à savoir sa taille: quelle est le nombre de triangulations de S^3 avec t tétrahèdres? La question équivalente en 2d a été résolue par W. Tutte en 1962; sa téchnique est basée sur une variation de la méthode permettant de compter les arbres binaires. Une généralisation simple de cette méthode au cas 3d n'aboutit pas, à cause de la richesse de la topology de S^3. En particulier, il est possible de trianguler la sphère S^3 de telle façon à ce qu'elle contienne une arête nouée.

Etant donné une triangulation arbitraire A de S^3, notre méthode consiste à la décomposer en un arbre de triangulations élémentaires que nous appelons *noyaux*. Un noyau est une triangulation 3d d'un polytope (donc avec bord), sans noeud interne, tel que chaque face interne possède au plus une arête externe (sur le bord). Un tétrahèdre (le simplexe 3d) est un noyau, mais contrairement au cas 2d, il existe des noyaux plus complexes. Nous montrons que si le nombre de noyaux avec t tétrahèdres admet une borne exponentielle en t, alors le nombre de triangulations de S^3 admet une borne exponentielle similaire.

Dans le troisième chapitre, nous réduisons ces noyaux à des triangulations encore plus petites, en identifiant certains noeuds adjacents sur le bord. Nous obtenons ce que nous appelons des *atomes*. Nous montrons que si leur nombre admet une borne exponentielle, alors le nombre de noyaux et par suite, celui de toutes les triangulations de S^3 admet une borne similaire.

L'intérêt de cette décomposition en atomes est le suivant: le nombre de triangulations contenenant une arête nouée d'une façon particulière, comme par exemple en un noeud de trèfle, est très grand. Par contre, un atome contient l'essentiel de l'information topologique. Nous postulons que le nombre d'atomes contenant une arête avec un noeud particulier est très petit, peut-être même égal à un. Démontrer un tel résultat permtettra de résoudre le problème consistant à compter les triangulations de S^3.

Contents

Chapter 1

Introduction

Glasses are among the oldest materials known to man: according to the ancient-Roman historian Pliny (AD 23-79), Phoenician merchants transporting stone discovered glass in the region of Syria around 5000 BC, when the soda they were carrying melted with the sand due to the heat of the fire and formed glass. Interestingly enough, the concept behind glass-making is still unchanged millennia later: the most common way of obtaining a glass is by supercooling a viscous liquid. Nowadays, we define a glass as an amorphous solid that lacks the periodicity of crystals (due to the supercooling). The scientific community has shown an increasing interest in glassy materials these past two decades; the main issues are to understand the enormous slowing down of the dynamics of the liquid as it is cooled, or what is commonly called the liquid-glass transition (see for instance the review article [1] by Stillinger et al.), and the mechanical properties of glasses (see for instance the review article [2] by Rodney et al.). The most common approach, whether by simulating the system or actually experimenting on it, is to apply some force or stress on the material and study its strain and how it reacts.

We want to look at glassy systems from a purely topological point of view. We show that imposing some topological constraints is enough to slow the dynamics at low temperatures so that the system behaves as a glass. In our model, there are no privileged directions so the notion of force does no exist. There is no geometry or coordinates either, so the notions of displacement and velocity cannot be defined. What we mean by topological constraints is the following : we know that the ground state of the system is a crystal. Given a particle, we simply say that its local neighborhood or the set of particles around it should be topologically identical to the local neighborhood of the particle in the crystalline ground state. We use the term *flower* to denote the local neighborhood of a particle (the term star is more common in the literature) and we define a positive measure called the *local energy* that, given a flower, measures how different it is from that of the ground state.

Let us first consider a geometric example in two dimensions, where the particles have coordinates: the hexagonal tiling is a regular tiling of the euclidean space and as such is a good candidate for the ground state of a system with identical particles. The flower of a particle is simply the set of its nearest neighbors. In two dimensions, the flowers are closed polygons and, considering the hexagonal tiling as a ground state, the topological constraint we impose is simply saying that the flower of each particle should be a hexagon or in other

terms, each particle must have six neighbors.

The flower of a particle is defined using the *Voronoi decomposition*: given a configuration, we define around each particle the box of points in space that are closer to this particle than to any other particle (the metric used to define closeness is not important; the most natural one is of course is the Euclidean distance). This box is called the Voronoi cell. We say that two particles are nearest neighbors if and only if their Voronoi cells share a common boundary.

Considering still the example of two dimensions, we see that the Voronoi cell of a particle is a polygon and that there are some points in space where three Voronoi cells meet. Points where four or more Voronoi cells meet are most likely non-existent. We connect two particles if and only if they are nearest neighbors and we get a triangulation called the *Delaunay triangulation*. This result can be easily generalized to three (or even higher) dimensions.

Given a configuration of particles in space, once we obtain its Delaunay triangulation, we can forget the particles' coordinates. We get what we call a *topological triangulation*. It is simply a list of triangles (or tetrahedra in three dimensions) satisfying certain properties. From our point of view, a topological triangulation is the ideal way of representing a configuration of particles since it has all the topological information we need pertaining to the flowers. The phase space we consider is then the set of all simplicial piecewise-linear decompositions of the d-dimensional sphere. We will only consider the cases $d = 2, 3$.

The elementary moves of the dynamics we consider are given by the Pachner transformations in $d = 2, 3$ that conserve the number of particles. More generally, Pachner defined a set of transformations on the triangulations of the d-dimensional sphere S^d (for any $d \geq 2$) called the *Pachner moves* and he showed that any two triangulations of S^d can be transformed into one another by using only these moves. In two dimensions, there is only one Pachner transformation that does not change the number of particles; it is called the *T1 flip*. In three dimensions, there are two such transformations called the *2-3 flip* and the *3-2 flip* respectively. We will see that each of these three moves can be interpreted in a very natural way in terms of moving particles closer or further away from one another.

The first chapter of this thesis deals with a two-dimensional (2d) topological glass. The inspiration for the model comes from [3]. They simulate a mixture of two types of particles, big and small, with the same concentration and a fixed radii ratio. Each particle is coupled to a heat bath and the particles interact with a Lennard-Jones potential, which is highly repulsive at close range (but not infinite so the particles are not hard spheres) and very slightly attractive at long range. They enclose the system in a square box with periodic boundary conditions and they integrate the equations of motion. At high temperature, the system behaves as a liquid but when they start going to low temperatures, the viscosity of the liquid rises super-exponentially and the system behaves as a glass. They observed that at low temperatures, the lowest energy per particle is achieved if small particles have five neighbors and big particles have seven neighbors. Fig. 1.1 shows the Voronoi decomposition of a typical glassy state they find.

As a consequence, the phase space we consider is the set of colored triangulations of S^2: each node can have one of two colors, red or blue. The elementary move is the T1 flip: given two triangles sharing an edge, we remove this edge and replace it with one connecting the

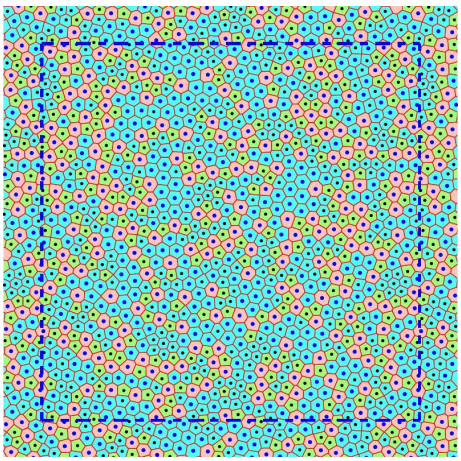

Figure 1.1: This figure is taken from [3]. It shows the Voronoi decomposition of a typical glassy state at low temperature. Each cell is colored according to the type and degree of its particle: green for small particles in pentagons, pink for big particles in heptagons and cyan for small and big particles in hexagons.

two other nodes. The local energy $E(m)$ per node m is $E(m) = |\deg(m) - c_m|^2$, where $\deg(m)$ is the number of edges incident on the node m, called the *degree of node m*, and $c_m = 5, 7$ if the node m is red or blue respectively. The dynamics is given by a simple Metropolis algorithm with one free parameter T we call the *temperature*. This model was first proposed by Eckmann in [4]. A similar study was initiated earlier by Aste and Sherrington [5].

We observe that, at low temperature T, this system behaves as a glass. In particular, we observe the two following indicators:

- We define a measure on our phase space: simply put, the distance between two triangulations $d(T_1, T_2)$ is the fraction of nodes whose flower is not the same in both triangulations T_1, T_2. Using this distance, we define an autocorrelation function $c(\vartheta)$ as the distance between a triangulation T_1 in the stationary state at temperature T and the triangulation T_2 obtained from T after ϑ elementary moves. We observe that $c(\vartheta)$ relaxes exponentially with ϑ with a coefficient τ that grows exponentially with the inverse temperature. This coefficient τ can be viewed as a viscosity coefficient.
- We also observe that the total energy relaxes polynomially in time and not exponentially, as $t^{-\alpha}$. This phenomenon is called the aging process.

We find explicit formulas for τ and α. First, we define a defect as a node with the incorrect degree; for instance, a -1 defect can be a red node with four neighbors or a blue node with six neighbors. In the stationary state at low temperatures T, the only defects present in the system are ± 1.

On a small scale, we see that an isolated defect is stuck; any elementary flip in its neighborhood increases the energy. This is typical of glassy systems. We show that the easiest way, *i.e.*, the way with the smallest increase in the energy of moving defects is through a collective movement of many T1 flips. More precisely, we show that defects move by sending some sort of signal. Furthermore, this signal is one-dimensional with a predefined

trajectory along the nodes of the triangulation. This resembles to some extent the flipping of quadrupolar events observed in many Lennard-Jones glasses (see for instance the works of Lemaitre[6], Tanguy [7] and many others) with a fundamental difference, namely that, in these glasses, the signal travels through elasticity whereas in our model, it is stochastic, *i.e.*, it is simply a Brownian motion.

On a larger scale, we see that the stationary states at low temperatures of our model are actually saddle points in the graph of all configurations: a random T1 flip in a stationary state will most likely increase the energy and the probability of finding a T1 flip that does not decays exponentially with the inverse temperature. In the stationary state at low temperature, the system is stuck in very narrow and steep canyons. This behavior is also typical of glassy systems (see for instance the works of Eckmann [4], Mezard [8] and Stratt [9]) . Moving a defect is equivalent to going from one such canyon to another and it takes a very long time for the system to navigate its way. Furthermore, the number of these canyons becomes exponentially small the further we decrease the temperature.

The relaxation of the autocorrelation function $c(\vartheta)$ is caused by the moving defects. Using the above observations, we derive an explicit formula for this relaxation. The energy relaxation on the other hand is caused by defects meeting one another: if a $+1$ defect and a -1 defect collide, they annihilate and the energy decreases by two. Studying the problem on a larger time scale than that of the emission of the above signals, we show that the movement of isolated defects can be viewed as a 2d random walk with a diffusion constant that decreases with time; this is a direct result of the stochastic nature of the signals' movement. Our model can then be viewed as a dilute gas of two types of particles A, B undergoing the chemical annihilation reaction $A + B \to \emptyset$. This problem was studied in detail by many authors. Using their results, we derive an explicit formula for the relaxation coefficient of the energy.

The next step is to generalize these results to the three-dimensional (3d) case and construct a model of a 3d topological glass. This proved to be a difficult challenge. One of the main reasons is that, contrary to S^2, the topological properties of the triangulations of the 3d sphere S^3 are far from being understood. In the second and third chapters of this thesis, we study one of these important properties, and a long standing problem: What is the number of triangulations of S^3?

In two dimensions, the problem was solved by Tutte in [10]. Using a generalization of the method for counting binary trees, he showed that the number of triangulations of S^2 with n nodes grows exponentially as $(4^4 \cdot 3^{-3})^n$. The main reason his method works is that any 2d triangulation can be constructed in the following manner: starting from a single triangle, we add a new one and we glue them along an edge. We continue adding triangles, one at a time, and gluing each to the already existing triangulation along one or two of its edges.

The problem in three dimensions is to prove the existence of an exponential bound on the number of triangulations of S^3 with t tetrahedra. Such a result could have many applications ranging from models of glassy dynamics such as ours, to other models of statistical mechanics (see for instance the works of Ambjørn, Durhuus and Jónsson [20] on modeling 3d quantum gravity). There have been many attempts over the past three decades to prove this result (see for instance the works of Danaraj and Klee [21], Lickorish [22], Hachimori and Ziegler [23]) but to date, this remains an open problem. Although we do not solve it, we make some

interesting contributions and we strongly believe that our approach is a step in the good direction towards finding a solution.

People tried to generalize Tutte's approach by adding tetrahedra and gluing them along faces. This is called shelling. Using the same ideas as in the 2d case, one can easily show that the number of such triangulations with t tetrahedra grows exponentially in t. The problem is that not all 3d triangulations are shellable. So people defined larger and larger classes of triangulations (shellable $\subset$ constructible $\subset$ collapsible $\subset$ locally constructible) by loosening the restrictions on gluing faces and showed that each of these classes admits exponential bounds in t. The main reason why an exponential bound holds for each of these classes is that they are constructed locally; what this means is that we are not allowed to glue faces that are not adjacent. The last of these attempts was done by Durhuus and Jonsson [24]: they introduced the new move of identifying two external adjacent faces. Starting from what is called a *tree of tetrahedra* (similar to binary trees), they showed that the number of triangulations of S^3 that are obtained by identifying external adjacent faces is exponentially bounded in t and they asked if all triangulations of S^3 can be obtained in this fashion.

This question was answered negatively by Benedetti and Ziegler [25]. They called the triangulations that can be obtained in the fashion described above *locally constructible* (LC) and they showed that a triangulation with a knotted triangle is not LC if the knot is complicated enough. Note that the existence of 3d triangulations with knotted edges was known since at least 1924 (see for instance the construction of Furch [26] known as Furch's ball, also shown in Fig. 1.2).

A triangulation of S^3 is called a *3-sphere*. Removing a node and all tetrahedra having it as a corner, we obtain what is called a *3-ball*. We work with balls. Clearly, if an exponential bound holds for spheres, it also holds for balls. The converse is also true and can be shown by adding a new node and a new tetrahedron for every face of the boundary of the ball; this operation is called *adding a cone over the boundary* of the ball.

Our approach is top down in the sense that we consider balls and we decompose them into their elements we call nuclei using some moves. The first one is the inverse of Durhuus and Jonsson's move of identifying external adjacent faces. Then we introduce a new move of splitting an external node m along a path γ in its flower $\mathcal{I}(m)$. Note that in three dimensions, the flower of a node is by definition a 2d triangulation and the path $\gamma \subset \mathcal{I}(m)$ is simple and splits $\mathcal{I}(m)$ into two separate 2d triangulations. Using these two moves, we show that we can decompose any ball into a tree of nuclei: a nucleus is a triangulation with no internal[1] nodes where every internal face has at most one external edge. Note that a tetrahedron is the most basic nucleus and that any LC ball is decomposed into a tree of tetrahedra. Non-LC balls on the other hand lead to more complex nuclei.

The main result of the second chapter of this thesis is to reduce the question of the existence of an exponential bound on all balls to nuclei. The difficulty of the proof lies in the move of splitting a node along a path γ: although it is a local move, it adds a lot of tetrahedra; more precisely, one new tetrahedron is added for every edge of γ. We show that the number of added tetrahedra can be controlled by carefully choosing the paths along which we split. This allows us to show that if the number of nuclei with t tetrahedra is bounded by C^t, then the number of all balls with t tetrahedra admits a similar exponential bound. Note

[1] A node, edge or face is called external if it is on the boundary; otherwise, it is called internal

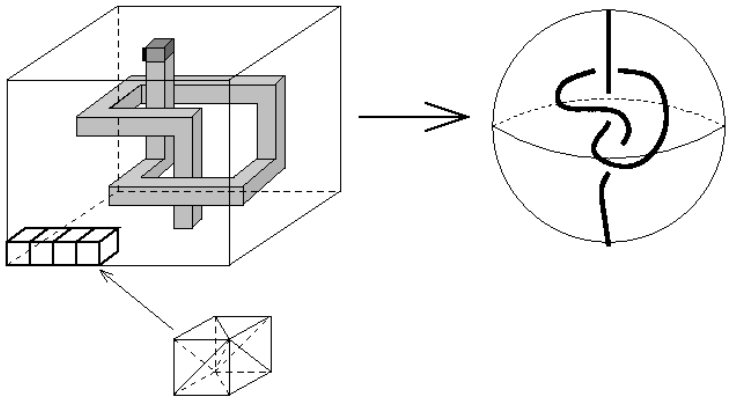

Figure 1.2: This figure is taken from [27]. A "knotted spanning arc" is an arc such as described in the right panel. The left panel indicates how to embed such a knotted arc by one edge. First we prepare a big pile of small cubes and dig a hole from the bottom face, making a knot as in the figure. We stop digging one step before the tunnel goes through the upper face, so that the constructions remains a 3-ball. Finally, we triangulate each cube into six tetrahedra without introducing new vertices.

that our approach always transforms a ball into a ball and as such can be implemented on a computer. This allows for an extensive numeric investigation of these nuclei resulting in several interesting observations.

We summarize the approach to the problem of counting 3d balls we introduced in the second chapter: the balls we do not know how to count yet, *i.e.* the non-LC balls, all have some sort of non-trivial topological properties that cannot be found in two dimensions, mainly knotted edges. Let T be such a non-LC ball; we decompose it into a tree of nuclei T' having the exact same topological properties: if T has a knotted edge with a given knot, say for example the trefoil knot, then T' has exactly one non-trivial nucleus having a trefoil-knotted edge. We show that if the nuclei are exponentially bounded, then so are all the non-LC balls. In the third chapter, we continue this approach of reducing non-LC balls having non-trivial topology into smaller, even simpler balls with the exact same topology. We are motivated by the following observation: given a nucleus with non-trivial topology, say a trefoil-knotted edge, we can split any of its nodes along any path and we obtain a bigger[2] nucleus with the same topology.

We introduce a new move of collapsing an external edge: this is the inverse move of splitting a node along a path. An external edge $e = (a, b)$ is collapsible if the intersection of the flowers of its ends is minimal: $\mathcal{I}(a) \cap \mathcal{I}(b) = \mathcal{I}(e)$, where $\mathcal{I}(e)$ is the flower of the edge e defined as the set of nodes m' and edges e' such that (m', e) is a face and (e', e) is a tetrahedron respectively. Let T' be a non-trivial nucleus. Collapsing any edge of T' yields a smaller nucleus T with the exact same topology as T'. Furthermore, we show that the operations of collapsing edges commute. We define a new class of nuclei called *atoms*:

[2]Bigger refers to the fact that it has more nodes and tetrahedra.

an atom is a nucleus such that no external edge is collapsible. The commutativity of the collapses imply that exactly one atom A can be obtained from a given nucleus T' by collapsing external edges. Moreover, A has the same non-trivial topology as T'.

The main result of the third chapter of this thesis is to show that if the atoms are exponentially bounded, then so are the nuclei and therefore all the 3d balls. This is done by counting the number of nuclei with t tetrahedra that can be obtained from a given atom A with t' tetrahedra by splitting nodes along paths.

The first main idea of the proof is to study the commutativity of splitting moves. We show that two splitting moves commute in all but one case. This allows us to group splittings of nodes into subsequences of commuting moves. The second main trick of the proof is to draw all the paths before performing any splitting. We count the number of ways these initial paths can be drawn. Then, we modify these paths after each subsequence of commuting splitting moves and we count the number of ways these modifications can be made. Combining these results, we show that, given an atom A with t' tetrahedra, the number of nuclei with $t > t'$ tetrahedra that can be obtained from A by splitting nodes along paths is bounded by C^t, where C is some constant.

The number of nuclei with a given topological property, say a trefoil-knotted edge, can be very large. On the other hand, we believe that an atom is the smallest possible ball with a given knot. If such a result were true, it would imply that the number of atoms with a given topological property is very small, maybe even equal to one. Such a result would allow us to bound the number of atoms with t tetrahedra, and consequently the number of all triangulations of S^3 with t tetrahedra, with the number of knots with a given complexity (the crossing number for instance can be considered as a good measure of the complexity of a knot). It was shown by Sundberg [36] that the number of knots with t crossings is exponentially bounded. This would solve the problem of counting triangulations of S^3.

In the final chapter, we construct a model of a 3d topological glass. The phase space is the set of all simplicial piecewise-linear decompositions of S^3 and the elementary moves of the dynamics is given by the 2-3 move and its inverse the 3-2 move. The 2-3 move, also called flipping a face, transforms two tetrahedra into three as follows: we consider two tetrahedra sharing a common face; we remove this face and we connect the two nodes opposite to it by an edge; we obtain three tetrahedra sharing this newly added common edge. The 3-2 move, also called flipping an edge, is simply the inverse move; it transforms three tetrahedra into two.

Contrary to the 2d case, it was shown by Dougherty et al. [38] that the phase space is not irreducible if we only consider the Pachner moves that conserve the number of nodes: they constructed a (LC) triangulation of S^3 with 16 nodes such that no 2-3 and 3-2 moves are possible. Note that a face is flippable, *i.e.*, a 2-3 move is possible, if and only if the target edge is not an edge of the triangulation and an edge is flippable if and only if its degree is three, *i.e.*, it is shared by three faces, and the target face is not a face of the triangulation. Furthermore, they showed that their construct cannot be geometrically realized. This raises another important topological difficulty that was not present in two dimensions: some of the topological triangulations of S^3 cannot be drawn and as a consequence do not represent real configurations of particles, *i.e.*, they cannot be geometrically realized as Delaunay triangulations.

Using some ideas by Santos [39], we show that any two topological triangulations of n nodes that can be geometrically realized as Delaunay are connected by 2-3 and 3-2 flips. Moreover, we show that the set of all these physical triangulations has the *small world property*, *i.e.*, that it contains (at least) exponentially many configurations such that any two configurations can be connected by a polynomially growing number of flips.

The next step is to define the energy. Contrary to the 2d case where the flower of a node is characterized by one integer, the flower of a node in three dimensions is a 2d triangulation of S^2. The first model we consider is identical to the 2d case: the contribution to the energy of a node is a simple function of its degree, or in other terms, we characterize the flower of a node by the number of its vertices. Clearly this constitutes a huge simplification. Nonetheless, the system exhibits glassy behavior; in particular, we observe that the energy relaxes polynomially to its stationary value at some temperature T. Note that in two dimensions, the average degree of a node is fixed by the topology and is equal to 6 whereas in three dimension, it is a free parameter $\bar{d}$ of the model (we only consider mono-colored triangulations). Looking once more at the elementary defects, we see that they are almost identical to the 2d case: a defect is a single node with the incorrect number of neighbors. At low temperature, the only defects present are those whose degree differ from $\bar{d}$ by one. The defects are isolated and their movement can be seen as a 3d random walk with a fixed diffusion constant. When two defects collide, they annihilate. As a consequence, this model has little more to offer than its 2d counterpart.

We want to construct a model that makes full use of the additional degrees of freedom of the 3d flowers. We start by fixing the ground state of the system as a crystal, *i.e.*, as a regular filling of space. Note that this filling must be a triangulation; the cubic lattice for instance is not accepted. The filling cell we choose is the tetrakis hexahedron: it has 14 vertices, 6 of which have a degree of 4 and the remaining 8 have a degree of 6. We show that it is possible to fill space with particles such that the Voronoi cell of each particle is a tetrakis hexahedron. Considering any triangulation, the local energy of each particle is then a measure of how different its Voronoi cell is from a tetrakis hexahedron. Such a definition of the energy is very difficult to implement on a computer, so we simplify it by characterizing a flower $\mathcal{I}(n)$ of a node n, which is a 2d triangulation of S^2, by a vector $\mathbf{e}(n) = (e_3, e_4, \dots)$, where e_k is the number of vertices in $\mathcal{I}$ with degree k. The tetrakis hexahedron's vector is $\mathbf{e}^\star = (0, 6, 0, 8, 0)$ and the energy contribution of a node n is simply the euclidean distance $\|\mathbf{e}(n) - \mathbf{e}^\star\|$.

The dynamics is given by a simple Metropolis algorithm with the temperature being the only free parameter. We observe that the system behaves as a glass; more precisely, we observe that the energy relaxation as a function of time has several polynomial regimes. Unfortunately, the dynamics is very slow at low temperature. Using some additional tricks (improving the algorithm, considering less restrictive energy forms etc...), we can speed things up considerably (by several tens of times) but to little effect and the system never reaches its stationary state in a reasonable amount of time.

One interesting notion we would like to understand is that of a defect: in the 2d model, a defect is a single isolated node with an incorrect flower. In the 3d model, it is easy enough to see that having a single node with an incorrect flower surrounded by nodes with tetrakis hexahedra as flowers is impossible. So one would expect that defects in 3d are strongly connected complexes of several nodes with an incorrect flower each. The simulations on the

other hand seem to show something completely different: even with more than 90% of the nodes having the correct flower, the ”bad” nodes do not cluster; instead they seem to form lines of defects. This shows that the behavior of this type of 3d models of topological glasses is much more complex than we originally thought.

Terminology

Throughout this thesis, we reserve the term **theorem** for mathematical statements which are true under the assumptions and hypotheses formulated therein.

Chapter 2

2D Topological Glass

2.1 Introduction, the model

This chapter deals with a species of a class of models on topological studies of triangulations. Such models have been studied in several contexts 2-d gravitation, froth, [11, and references therein]. The variant we use here was introduced in [4], but it turned out that a very similar study was initiated earlier by Aste and Sherrington [5].

We reconsider here the model which was inspired by [3] and introduced in [4]. For completeness we repeat the definition of the model: We fix a (large) number N of nodes, half of which are red, and the other half blue. These nodes are the nodes of a topological triangulation T of the sphere S^2. The set of all possible such labeled triangulations will be denoted $\mathcal{T}_N$. We define a dynamics on $\mathcal{T}_N$ by the following Metropolis algorithm whose elementary steps are flips (T1 moves): A link is chosen uniformly at random (among the $3N-6$ links). In Fig. 2.2, if the link AB was chosen then the flip consists in replacing it by the link CD. This move is not admissible if the link CD already exists before the move. Otherwise it is admissible. Note that the number of nodes, N, does not change in this model. However, we will be interested in the behavior for $N \to \infty$.

The Metropolis algorithm is based on the energy function E on $\mathcal{T}_N$ which, for any triangulation $T \in \mathcal{T}_N$, is defined as

$$E(T) = \sum_{i \in \text{blue}} (d_i - 7)^2 + \sum_{i \in \text{red}} (d_i - 5)^2 \ ,$$

where d_i is the degree (number of links) of the node i. Thus, this energy favors 7 links for the blue nodes and 5 for the red ones. *Mutatis mutandis*, the detailed definition of the energy is not important for the discussion of the model, and we will stick to this particular form of the energy. Given an admissible flip, compute the energy of the triangulation before and after the flip; this defines

$$\mathrm{d}E = E_{\text{after}} - E_{\text{before}} \ .$$

An admissible flip is performed if either $\mathrm{d}E \leq 0$ or, when $\mathrm{d}E > 0$, with probability $\exp(-\beta \mathrm{d}E)$, where β is the inverse temperature of the system.

Several properties of this model were discussed in [4], but here we study in more detail the dynamical properties of the model. In particular, we introduce a "charge" defined as follows:

Definition 2.1.1. *The* charge *of a red node is defined by $d_i - 5$ and the charge of a blue node is defined by $d_i - 7$. We will say the charge is a* defect *+ if it is +1 and - if it is −1. In general, the color of the charge will not matter and will not be mentioned.*

In principle, all charges between -4 and $O(N)$ can occur, since $d_i \geq 3$, but, obviously, at low temperatures mostly the charges +, 0, and - will come into play.

2.2 Equilibrium and the approximation of the dynamics

The dynamics of the model is given by the Metropolis algorithm. In it, a link is chosen uniformly at random among all possible links. The change of energy induced by the flipping of this link is called $\mathrm{d}E$. If $\mathrm{d}E \leq 0$ the flip is performed, if $\mathrm{d}E > 0$ the flip is performed with probability $p(\mathrm{d}E) = \exp(-\beta \mathrm{d}E)$. This process satisfies detailed balance, and most of the chapter is dealing with the equilibrium properties of this process at low temperatures. Because of the detailed balance, the equilibrium measure μ has the property that the probability to see a given state whose energy is E is proportional to $\exp(-\beta E)$. We use this elementary observation to argue that at low temperature there are only few defects, by which we mean that there are few red nodes whose degree is not 5 and also few blue ones whose degree is not 7. Given that there are few of these "defects", we further assume that the "positions" of these defects are random in the sense that there are no strong conditional expectations: For example, having a defect +1 does not say that there is a defect -1 close-by. The upshot of this way of reasoning, which we corroborate by numerical studies, is that one can approximate the dynamics by just looking at defects.

Indeed, the full dynamics must be described by the evolution of correlation functions. It would have to take into account correlation functions between the charges (and the colors) of, say, the 4 nodes on a pair of triangles sharing an edge. Then, flipping that edge, the correlations of many neighboring triangles would be changed simultaneously, and this would necessitate considering a full hierarchy of correlations (like BBGKY). What we will see is that in this model, these higher order correlation functions do not influence our basic understanding of what is going on.

In contrast, the Euler relations play a small but not totally negligible role for the sizes of the systems we consider.

2.3 Description of the stationary state

It will be useful to define throughout this chapter the natural parameter

$$\varepsilon \equiv e^{-\beta} \ .$$

We are interested in a regime where the density c of charges (which equals E/N) is low but also, where the number $c \cdot N$ of charges is large, so that good statistics and a certain independence of the Euler relations is attained. More precisely, we fix $\varrho \ll 1$ and $D_0 \gg 1$, and require $\varepsilon \leq \varrho$ and $N\varepsilon > D_0$. We furthermore consider the limit of large N.

The main result of this section is summarized in the following proposition:

Proposition 2.3.1. *Consider an equilibrium state at temperature $T \ll 1$ satisfying the above conditions on N and ε.*

1. *At first order in ε, the only charges present in the system are simple defects ± 1. Their density is $2\varepsilon + O(\varepsilon^2)$.*
2. *The distribution of the colors (red or blue) is independent in the limit $\varepsilon \to 0$.*
3. *The distribution of the charges is independent in the limit $\varepsilon \to 0$.*

Remark 2.3.2. *The meaning of $\varepsilon \to 0$ above is that the quantities become more and more decorrelated as $\varepsilon \to 0$ while still maintaining the inequalities $\varepsilon \leq \varrho$ and $N\varepsilon > D_0$.*

2.3.1 Energy of the stationary state

In this paragraph, we will calculate the energy of the stationary state in the limit specified above, as a function of the temperature.

Estimate 2.3.1. *Consider the region $\varepsilon N > D_0$ and $\varepsilon < \varrho$. For sufficiently large D_0 and sufficiently small ϱ the density of charges c is*

$$c \equiv E/N = 2\varepsilon + O(\varepsilon^2) \ .$$

Proof. Assuming equilibrium, by detailed balance, the probability to see a defect of charge ± 1 is $O(e^{-1\cdot\beta}) = O(\varepsilon)$, while the probability to see higher charges is $O(e^{-2^2\beta}) = O(\varepsilon^4)$, by the assumption of equilibrium and the form of the Hamiltonian, since, if $(d_i - 5)^2 > 1$ then it is at least 4.

So it remains to estimate the coefficient in front of the factor ε. There are 4 cases to consider: The number of red nodes with degree 4 or 6, resp. the number of blue nodes with degree 6 or 8. All these cases cost energy 1 per instance, and thus these 4 numbers are equal by the virial theorem.

We also need to estimate the cases with 0 charge, *i.e.*, blue nodes with 7 neighbors and red nodes with 5 neighbors, which appear again equally often, by the virial theorem. Since there are $N/2$ nodes of each color, and each of the colors has 2 states of defect 1 (namely ± 1), we conclude that the expected total number of defects is

$$2 \cdot 2 \cdot \varepsilon \cdot (N/2) = 2\varepsilon N + O(\varepsilon^2) \ . \tag{2.1}$$

□

2.3.2 Distribution of the colors

We next calculate the probabilities that a randomly chosen link connects 2 red (blue) nodes. We denote these probabilities by p_{rr} for red-red, p_{rb} for red-blue and so on. If there are no defects, *i.e.*, at order ε^0, all red nodes have 5 neighbors and all blue nodes have 7. This leads to the following relations:

$$2p_{\mathrm{rr}} + p_{\mathrm{rb}} = 5/6 \ ,$$

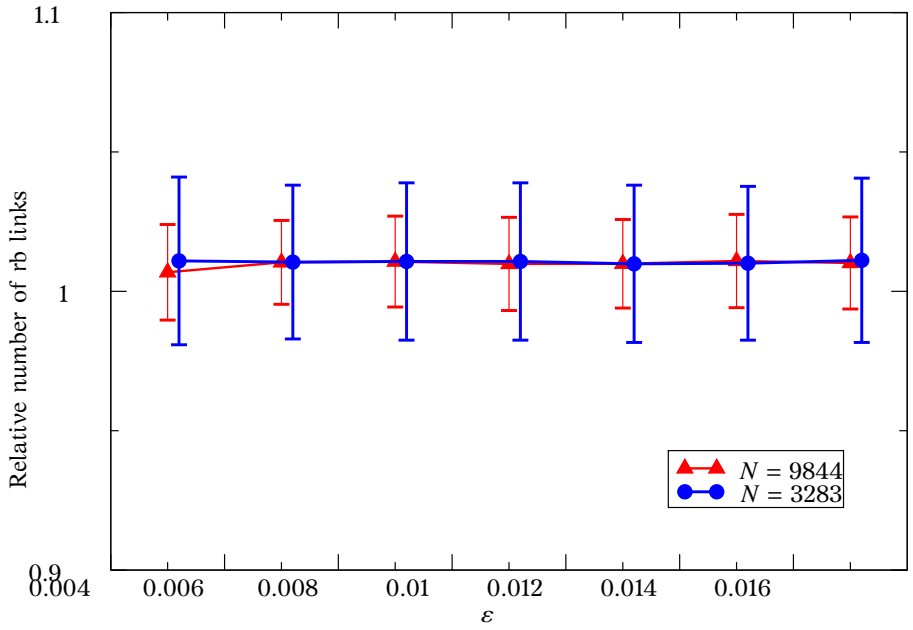

Figure 2.1: Numerical check of relation $p_{\rm rb} = 70/144$ by plotting $p_{\rm rb}/(70/144)$. The error bars are 3 σ and the data for $N = 3283$ are slightly shifted (in the x-direction) for better visibility.

$$2p_{\rm bb} + p_{\rm rb} = 7/6 \ .$$

Assuming that the positions of the colors are uncorrelated, we find that the relative probabilities to find a red-red, resp. blue-blue pair are

$$p_{\rm rr}/p_{\rm bb} = 25/49 \ .$$

This leads to $p_{\rm rr} = 25/144$, $p_{\rm bb} = 49/144$, and $p_{\rm rb} = 70/144$. In Fig. 2.1 we show that numerical simulations confirm this simple approximation to a very high degree of fidelity.

2.3.3 Energy cost of flips

We adopt an approach similar to Sect. 2.3.2. We use the hypothesis that the charges are randomly distributed over the nodes to calculate the probability of finding a link with a given neighborhood of charges and compare it to simulation results. In this case however, given a link ℓ, the neighborhood we consider is the ordered set of all 4 nodes involved in its flipping. For example in Fig. 2.2, this set would be $(c(A), c(B), c(C), c(D))$ where $c(A)$ is the charge of the node A. This choice will be very useful for to study the dynamics later on since it determines the energy cost of flipping a given link:

$$\begin{aligned} {\rm d}E(\ell) &= \sum_{n\in\{A,B\}} (c(n)-1)^2 - (c(n))^2 + \sum_{n\in\{C,D\}} (c(n)+1)^2 - (c(n))^2 \\ &= 4 + 2\left(c(C) + c(D) - c(A) - c(B)\right) \ . \end{aligned} \tag{2.2}$$

It is easy to enumerate all the various cases and the energy cost associated with each of them. We restrict the discussion to those situations where the charges take values in

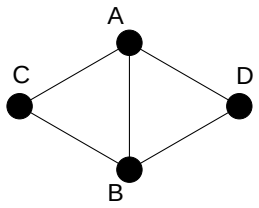

Figure 2.2: Labeling of the corners of 2 adjacent triangles

$\{+1, 0, -1\}$. In principle, there are 3^4 configurations, which are reduced to 36, by symmetry. They are summarized in Table 2.1 (symmetrical cases omitted).

Note that if the defects of the original configuration are bounded by ±1, then dE varies between -4 and 12.

defects	initial state				dE
0	0	0	0	0	4
1	+	0	0	0	2
1	0	0	0	-	2
1	0	-	0	0	6
1	0	0	+	0	6
2	+	+	0	0	0
2	+	0	0	-	0
2	0	0	-	-	0
2	+	-	0	0	4
2	+	0	+	0	4
2	0	-	0	-	4
2	0	0	+	-	4
2	-	-	0	0	8
2	0	-	+	0	8
2	0	0	+	+	8

defects	initial state				dE
3	-	-	+	0	10
3	0	-	+	+	10
3	+	-	0	-	2
3	+	+	+	0	2
3	+	+	0	-	-2
3	+	0	-	-	-2
3	+	0	+	-	2
3	0	-	-	-	2
3	-	-	0	-	6
3	+	-	+	0	6
3	+	0	+	+	6
3	0	-	+	-	6

defects	initial state				dE
4	+	-	-	-	0
4	+	+	+	-	0
4	-	-	+	+	12
4	-	-	-	-	4
4	+	-	+	-	4
4	+	+	-	-	-4
4	+	+	+	+	4
4	-	-	+	-	8
4	+	-	+	+	8

Table 2.1: The energy differences obtained by flipping the link between the first 2 values to a link between the second 2 values, as a function of the number of defects.

2.3.4 The number of local defect configurations

We assume throughout that the number of red (blue) nodes is $n_{\rm r}$ ($n_{\rm b}$) and that $\Delta \equiv n_{\rm r} - n_{\rm b} \in \{0, 1\}$. We denote by $p_\pm$ the probabilities to find charges ±1, respectively. Assuming that there are no other charges (except 0), we can write

$$N \cdot (p_- + p_+) = E \ ,$$
$$N \cdot (p_- - p_+) = 12 - \Delta \ ,$$

where the second equation follows from the Euler formula. In equilibrium, $E = 2N\varepsilon$, by Eq. (2.1), and therefore we get

$$p_\pm = \varepsilon \mp 6/N \pm \Delta/(2N) + \mathcal{O}\left(\varepsilon^2\right) \ . \tag{2.3}$$

We will assume that $N\varepsilon \gg 6$ so that the second term in Eq. (2.3) can be neglected. In a similar way, one can show that

$$p_{\pm 2} = \varepsilon^4 + O(\varepsilon^5) \ ,$$

and combining these we find that the probability of nodes with charge 0 is

$$p_0 = 1 - 2\varepsilon + O(\varepsilon^2) \ .$$

We next consider in more detail what happens in those pairs of triangles where a flip leads to $\mathrm{d}E = 0$. Looking again at Eq. (2.2) we see that the case $\mathrm{d}E = 0$ appears in 3 cases:
Case q_{+-} : One of A or B has charge + and C or D has charge - (and the others, charge 0).
Case q_{++}: A and B charge +, C and D charge 0.
Case q_{--} : A and B charge 0, C and D charge - .

Continuing with the independence assumption, we now look at the probability to find a configuration of type q_{++}, q_{+-} , and q_{--} . Note that there are $6N - 12$ half-links emanating from the nodes, and we are to pair them up randomly. Note that if a site is red, it has 4, 5, 6 outgoing links, depending on whether its charge is −, 0, +, respectively. Similarly, the numbers for a blue node are 6, 7, 8. Therefore, given that there are on average $\varepsilon N/2$ defects of type red-4, red-6, blue-6, blue-8, there will be $4\varepsilon N/2$ links from the red-4, $6\varepsilon N/2$ from red-6 and blue-6, and $8\varepsilon N/2$ from blue-8. The blue-7 and red-5 occur with probability almost 1 and have therefore respectively $7N/2$ and $5N/2$ dangling edges (with a correction factor $1 - O(\varepsilon)$) which we omit throughout. The probabilities to see such dangling edges are the quantities above, divided by $6N - 12$, the total number of dangling edges. We get, omitting higher order terms:

$$\begin{aligned} q_{++} &= (7p_+/6)^2 \cdot p_0^2 = 49\varepsilon^2/36 \ , \\ q_{--} &= (5p_-/6)^2 \cdot p_0^2 = 25\varepsilon^2/36 \ , \\ q_{+-} &= 4\,(5p_-/6)\,(7p_+/6) \cdot p_0^2 = 140\varepsilon^2/36 \ . \end{aligned} \tag{2.4}$$

We also get, by looking at Table 2.1:

$$\begin{aligned} p_{\mathrm{d}E=0} &= q_{++} + q_{--} + q_{+-} = 214\varepsilon^2/36 \ , \\ p_{\mathrm{d}E=2} &= 2\,(7p_+/6 + 5p_-/6) \cdot p_0^3 = 4\varepsilon \ , \\ p_{\mathrm{d}E=4} &= p_0^4 = 1 - O(\varepsilon) \ , \end{aligned} \tag{2.5}$$

The discussion of the other values of $\mathrm{d}E$ shows the limitations due to our closing assumptions: by the virial theorem, in total independence, we would simply have

$$p_{\mathrm{d}E=0} = p_{\mathrm{d}E=8} \text{ and } p_{\mathrm{d}E=2} = p_{\mathrm{d}E=6} \ . \tag{2.6}$$

But we could also have computed the probabilities as above, with the result:

$$\begin{aligned} p_{\mathrm{d}E=-2} &= 2\,(7p_+/6)^2 \cdot (5p_-/6) \cdot p_0 + 2\,(5p_-/6)^2 \cdot (7p_+/6) \cdot p_0 \\ &\approx 3.89\varepsilon^3 \end{aligned} \tag{2.7}$$

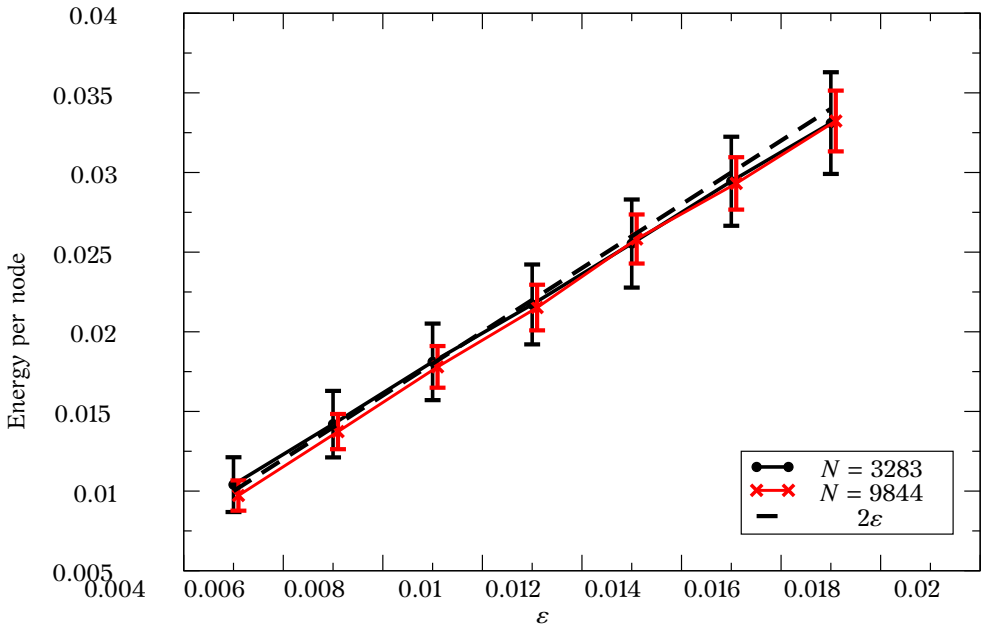

Figure 2.3: Numerical test of the mean energy per node (Estimate 2.3.1) for 950 realizations. The data for $N = 9844$ are slightly shifted for better visibility. Note the excellent fit with the theoretical curve, although the fluctuations are huge, getting better with larger system size (1 standard deviation shown).

instead of $4\varepsilon^3 = p_{\mathrm{d}E=2} \cdot \varepsilon^2$ given by the stationarity assumption, which proves that the distribution of defects is not completely uncorrelated. We will say that the correlation is bounded by $0.1\varepsilon^3$, and can thus be neglected in the limit $\varepsilon \to 0$.

In Figs. 2.3 and 2.4 we show with 2 examples that the numerical simulations confirm these simple approximations to a very high degree of fidelity. Note that in [12], the *uniform* measure on $\mathcal{T}_N$ was considered, and even this leads to correlations of degrees of neighboring nodes.

2.4 Dynamics of the system (at equilibrium)

We can use the results of the previous section to estimate the dynamics of the system under the Metropolis algorithm.

If a flip leads to an energy change $\mathrm{d}E$ then it is accepted in the Metropolis algorithm with probability

$$p_{\mathrm{acceptance}} = \varepsilon^{\max(0,\,\mathrm{d}E)} \quad . \tag{2.8}$$

On the other hand, the probabilities to pick a link with fixed $\mathrm{d}E$ are given by Eq. (2.5) and Eq. (2.6). Multiplying these numbers with the probabilities of Eq. (2.8) leads to an estimate of the probability that the flip in question actually happens. The results are summarized in Table 2.2 (calculated this time with the method of Eq. (2.7)).

Discussion: Inspection of Table 2.2 shows that the events with the highest transition rate are those which cost no energy, followed by those which have an energy cost of ±2. Also

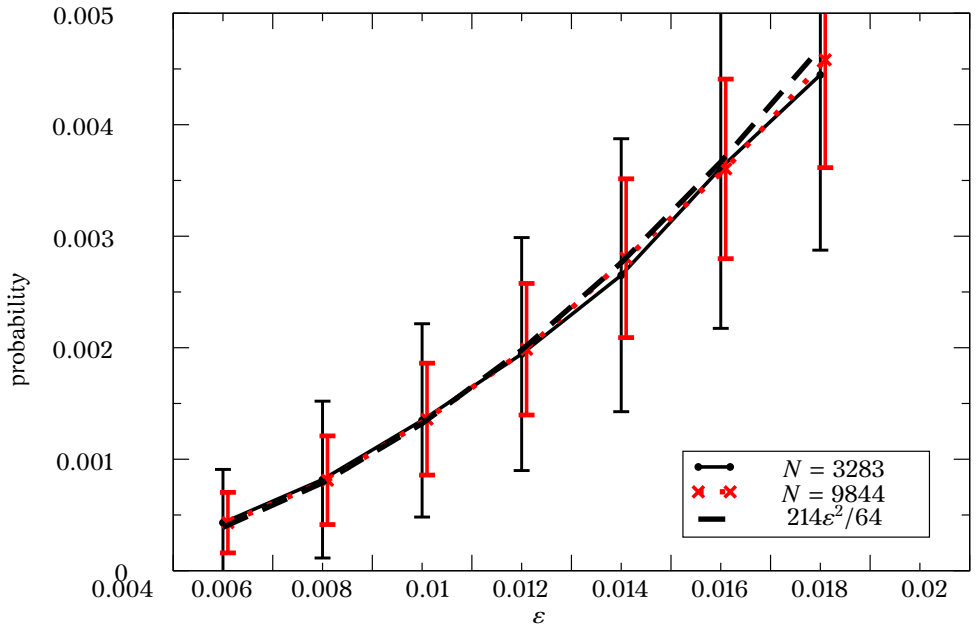

Figure 2.4: Numerical test of $p_{\mathrm{d}E=0}$ (Eq. (2.5)) from 950 realizations. The data for $N = 9844$ are slightly shifted for better visibility. Note the excellent fit with the theoretical curve, although the fluctuations are huge, getting better with larger system size (1 standard deviation shown).

note that the probability to find a link which will lead to a given dE is equal to the quantity in the table times $\varepsilon^{-\max(0,\mathrm{d}E)}$ since then we neglect the Metropolis factor. This leads to a table with the same prefactors, but with a power $\varepsilon^{|\mathrm{d}E-4|/2}$. In particular, *in the steady state, the local landscape is given by the 3rd column of Table 2.2: It is symmetric around* d$E = 4$.

Henceforth, we will only consider the 3 most frequent types of flips (the others are an order ε less probable):

1. Flips which change from 1 defect to 3 of them and which raise the energy by 2. These flips will be called *creation events*.

2. Flips which start from 3 defects and end with 1 defect and which decrease the energy by 2. These flips will be called *annihilation events*. Creation and annihilation events are obviously dual to each other and equiprobable in the stationary state.

3. Flips which do not change the energy, and in which a pair ++, +−, or − is involved. These flips are by far the most probable. We will discuss below in more details the 3 configurations which lead to d$E = 0$.

2.4.1 The most probable flips

As stated above, if $\varepsilon = 1\%$, then over 99% of the flips (which are accepted by the Metropolis algorithm) do not change the energy. It is clear that, in order to understand the dynamics

dE	transition rate	local landscape
-4	$1225/1296 \cdot \varepsilon^4$	$1225/1296 \cdot \varepsilon^4$
-2	$35/9 \cdot \varepsilon^3$	$35/9 \cdot \varepsilon^3$
0	$107/18 \cdot \varepsilon^2$	$107/18 \cdot \varepsilon^2$
2	$4 \cdot \varepsilon^3$	$4 \cdot \varepsilon^1$
4	$1 \cdot \varepsilon^4$	$1 \cdot \varepsilon^0$
6	$4 \cdot \varepsilon^7$	$4 \cdot \varepsilon^1$
8	$107/18 \cdot \varepsilon^{10}$	$107/18 \cdot \varepsilon^2$
10	$35/9 \cdot \varepsilon^{13}$	$35/9 \cdot \varepsilon^3$
12	$1225/1296 \cdot \varepsilon^{16}$	$1225/1296 \cdot \varepsilon^4$

Table 2.2: The probabilities of transitions from the initial state. Data only shown to order ε^4. The second column shows the probabilities to pick a link leading to a given dE. Higher order corrections are omitted.

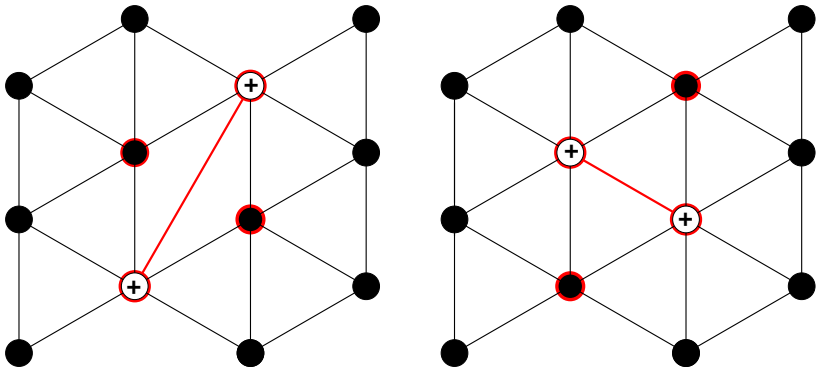

Figure 2.5: A flip from ++00 (on the left) and the resulting triangulation on the right. The affected nodes are supposed to be red, in this example. Note that the result is *again* of the type ++00. Furthermore, *again* with dE = 0 one can flip back. This is reminiscent of "blinkers" in the game of life [13, Chap25].

of the system, one should start by studying these flips.

Looking at Table 2.1 we see that there are 3 candidates for dE = 0 and they all involve only 2 defects. We will now show that the cases of ++00 and 00- are quite different from that of +00- (and its 3 other variants +0-0, ...). In the first case, ++00, which is similar to the case 00-, the local neighborhood looks like in Fig. 2.5. In this case, what happens is a flipping back and forth between the 2 states, with probability $p = (3N - 6)^{-1}$ (the probability to choose the colored link).

The case +00- is illustrated in Fig. 2.6. Here a new, and important phenomenon appears: The pattern, +00- which we will call a *pair*, is recreated, but *at a new position* a distance 1 away from the old one. We will also say that the pair +- *walks* one step.

The more important observation is that the pairs of defects must walk along a *predefined, 1-dimensional path* as shown in Fig. 2.7. This means that *the dE = 0 motion of +- pairs is a* **one***-dimensional random walk in the current triangulation T. This random walk (flipping back and forth on the predefined path) will continue until some other type of event happens.*

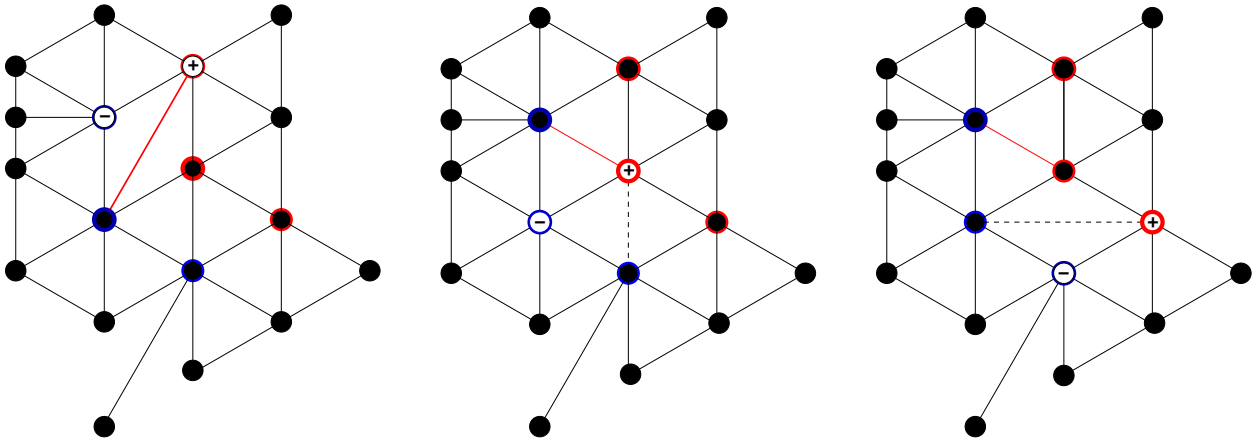

Figure 2.6: Change of pattern in the case +00−. In the sample only the relevant colors are as shown. Note that the effect of the flip is that the 2 defects move (in the picture) *down*. The reverse flip costs nothing $dE = 0$. The second flip (dashed line) moves the defects one step further. Note that this motion must take place one a *predefined, 1-dimensional path*.

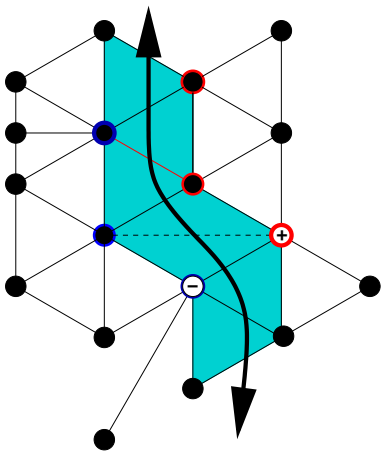

Figure 2.7: The same configuration as in Fig. 2.6 with the *1-dimensional* path of the pair of defects superimposed.

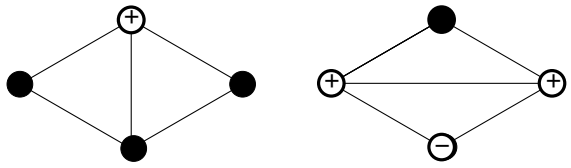

Figure 2.8: A creation event: a +- pair is created from a + defect, which is pushed one step.

2.4.2 Lifetime of pairs

As we have seen, a pair of opposite charges +- can move through the system without energy cost. Its motion is a 1d random walk along a fixed 1-dimensional path. Edges are still chosen randomly and will be flipped if possible and if the Metropolis condition is met in case $dE > 0$. Here we ask about the relative probabilities that a pair disappears, and we will show that *predominantly a pair will die when it collides with a defect.*

We need to compare 3 possibilities of which the first will be seen to be the most probable:

1. The random walk reaches another defect.

2. The pair is destroyed because a creation event involving 1 or 2 of its 2 defects occurs.

3. Two independent random walks meet.

Our earlier discussion says that the concentration of pairs $\left(70\varepsilon^2/36\right)$ is much smaller than the concentration of defects (2ε), implying that the probability of 2 pairs meeting is insignificant when compared to the probability of a pair meeting a defect.

We next estimate the probability of destroying a pair as in case (2). On average, there are 7 links in the neighborhood of a given pair which increase E, and flipping such a link has an energy cost of 2. The probability of this to happen is $7\varepsilon^2/(3N)$. Since the pair moves every $O(N)$ attempted flips, we conclude that, on average, a pair will do $O\left(\varepsilon^{-2}\right)$ steps before it is destroyed as in case (2).

The number of steps needed for case (1) to happen depends obviously on the density of defects. We let ξ denote the average distance between defects (counted in number of links). Since the number of defects is $2\varepsilon N$ and the system is 2-dimensional, we conclude that

$$\xi = O(\varepsilon^{-1/2}) \ .$$

As long as the pair is not destroyed by the mechanism leading to case (2) it can thus do $O(\varepsilon^{-2})$ steps by which time it can visit $O(\varepsilon^{-1})$ defects.

This terminates the comparison of the 3 cases, and shows that a pair has the time to visit a very large number of defects before it is destroyed by the 2 other mechanisms.

2.5 The Geometry of pair-defect collisions

In this section we consider the collisions between a pair and a defect. The discussion is really in two parts: On one hand, we must consider the probability that a collision between

a pair and a defect is initiated. This depends on the density of the defects, and hence on ε. But, once a collision is initiated, we can ask what the effect of the collision is going to be. The next proposition shows that this effect is purely geometrical and independent of ε.

Proposition 2.5.1. *There are 9 topologically different possibilities $Q_i, i = 1, \ldots, 9$ for a collision to be initiated. For each of them, there are 2 purely geometrical constants $P_{\mathrm{m},i} > 0$ and $P_{\mathrm{d},i} > 0$ (depending on i) which tell us the probability that a collision leads to a move ($P_{\mathrm{m},i}$) of a defect (by 1 or 2 links) resp. the deletion of the pair ($P_{\mathrm{d},i}$).*

The remainder of the section deals with the proof of Proposition 2.5.1.

2.5.1 Definition

We will study in detail how collisions move defects. First of all, we will define what we mean by a collision. Assuming that the density of defects is very small, the only collisions we will consider are those involving 3 defects, namely, the pair +- which will collide with a defect + or -.

Definition 2.5.2. *Consider some configuration T. Three defects $D_i, i = 1, 2, 3$ of T are said to be in a* collision *if there are $k \geq 2$ flips (k links of T) that do not increase the energy such that*

1. *The only defects involved in these k flips are $D_i, i = 1, 2, 3$.*
2. *All 3 defects are involved in these k flips.*
3. *At least one of these k flips will move a pair (the others might be any of the 4 cases which do not increase the energy).*

2.5.2 Collision types

In this section, we will describe all possible configurations of a collision and we will show that the probability of any such configuration depends solely on the topology (and not on the temperature).

The third condition of the definition of a collision states that we can always identify a pair; as a result, the set of all possible configurations of a collision can be obtained by taking a pair and placing either a - defect or a + defect in any position where it can interact with one of the pair's 2 defects. As seen in the previous section, a + defect can interact with any defect if and only if both defects are at distance one. Two - defects can interact if and only if they are on opposite corners of 2 adjacent triangles. The last ingredient is that + defects can have a degree of 6 or 8 whereas - defects have a degree of 4 or 6. This yields a systematic method of constructing all possible configurations of a collision: consider a pair and let $\mathcal{U}_1$ be the set of all empty sites (charge 0) which are at distance 1 of any of the pair's 2 defects and $\mathcal{U}_2$ be the set of all empty sites which are opposite to the - defect of the pair. The set of all possible configurations of a collision is obtained by placing a + defect in any of $\mathcal{U}_1$'s sites or a - defect in any of $\mathcal{U}_2$'s sites, as shown in Fig. 2.9 in the case where the + defect is red and the - defect is blue. All in all, we get 9 different configurations of a pair and a defect (symmetrical case omitted).

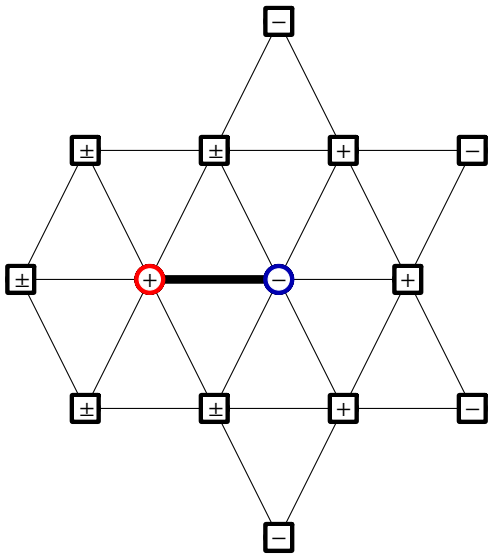

Figure 2.9: A figure showing all possible relative positions of a pair and a defect in collision. The pair is shown as a solid black line.

Assuming that the defects are randomly distributed, it is clear from Fig. 2.9 that the probability of a collision to be of some type $Q_i \in \{1, \ldots, 9\}$ is a temperature independent constant that can be calculated. To prove Proposition 2.5.1 one must study in detail each of the 9 cases. We will study in particular:

- What are the possible outcomes of each collision type and what is the (conditional on having initiated the collision Q_i) probability of each outcome?
- What is the probability (conditional on having initiated the collision Q_i) that a pair pushes a defect?

We can summarize the answers as follows:

- There will always be a defect left over at the end of the collision.
- Finding a pair and a defect at the end of the collision is possible in all 9 cases.
- An annihilation of the pair is possible in 2 of the 9 cases.
- It is possible that the defect is pushed in 8 cases. A defect can be pushed by more than 1 step.
- It is possible that the defect remains in its initial position in all 9 cases.

The relative probabilities of any of the above outcomes only depend on the local geometry. While all the cases have been worked out in detail, we illustrate the discussion for just 2 of them, and this will complete the proof of Proposition 2.5.1.

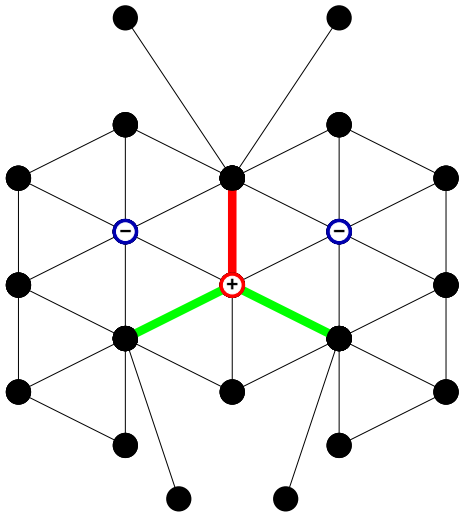

Figure 2.10: A collision where a an annihilation is possible. The green links show the way the pair enters (or exits) the collision. Flipping the red link will cause an annihilation event.

Example 1: A possible annihilation

There are 2 cases where an annihilation might occur. We consider here the case of Fig. 2.10. A +- pair collides with a - defect. For simplicity, assume that the +- pair came from the left. Once the pair and the defect are in collision, there are 3 links whose flipping leads to $dE \leq 0$. Two of these links (the green ones) allow the pair to walk away from (or enter) the collision. Flipping the red link on the other hand causes an annihilation: the pair is destroyed and the defect is pushed by one step. We clearly see that there are 3 possible outcomes:

- The pair exits the collision through the same way it entered (in our case, on the left). The defect remains in its initial position.
- The pair exits the collision through the other green link. The defect moves 2 steps.
- An annihilation event occurs. The pair disappears and the defect moves 1 step.

The (conditional) probability of each outcome is 1/3 and the (conditional) probability that the defect will have moved at the end of the collision is 2/3.

Example 2: A bifurcation

Here, we look at the collision case of Fig. 2.11. No annihilation is possible here and the outcome of the collision is always one pair and one defect. The only relevant question is what is the probability that the defect will have moved at the end of the collision. But the combinatorics is more involved.

The pair enters and may exit the collision through a green link. Flipping a red link on the other hand will not end the collision. Notice that the fifth diagram contains 4 red links

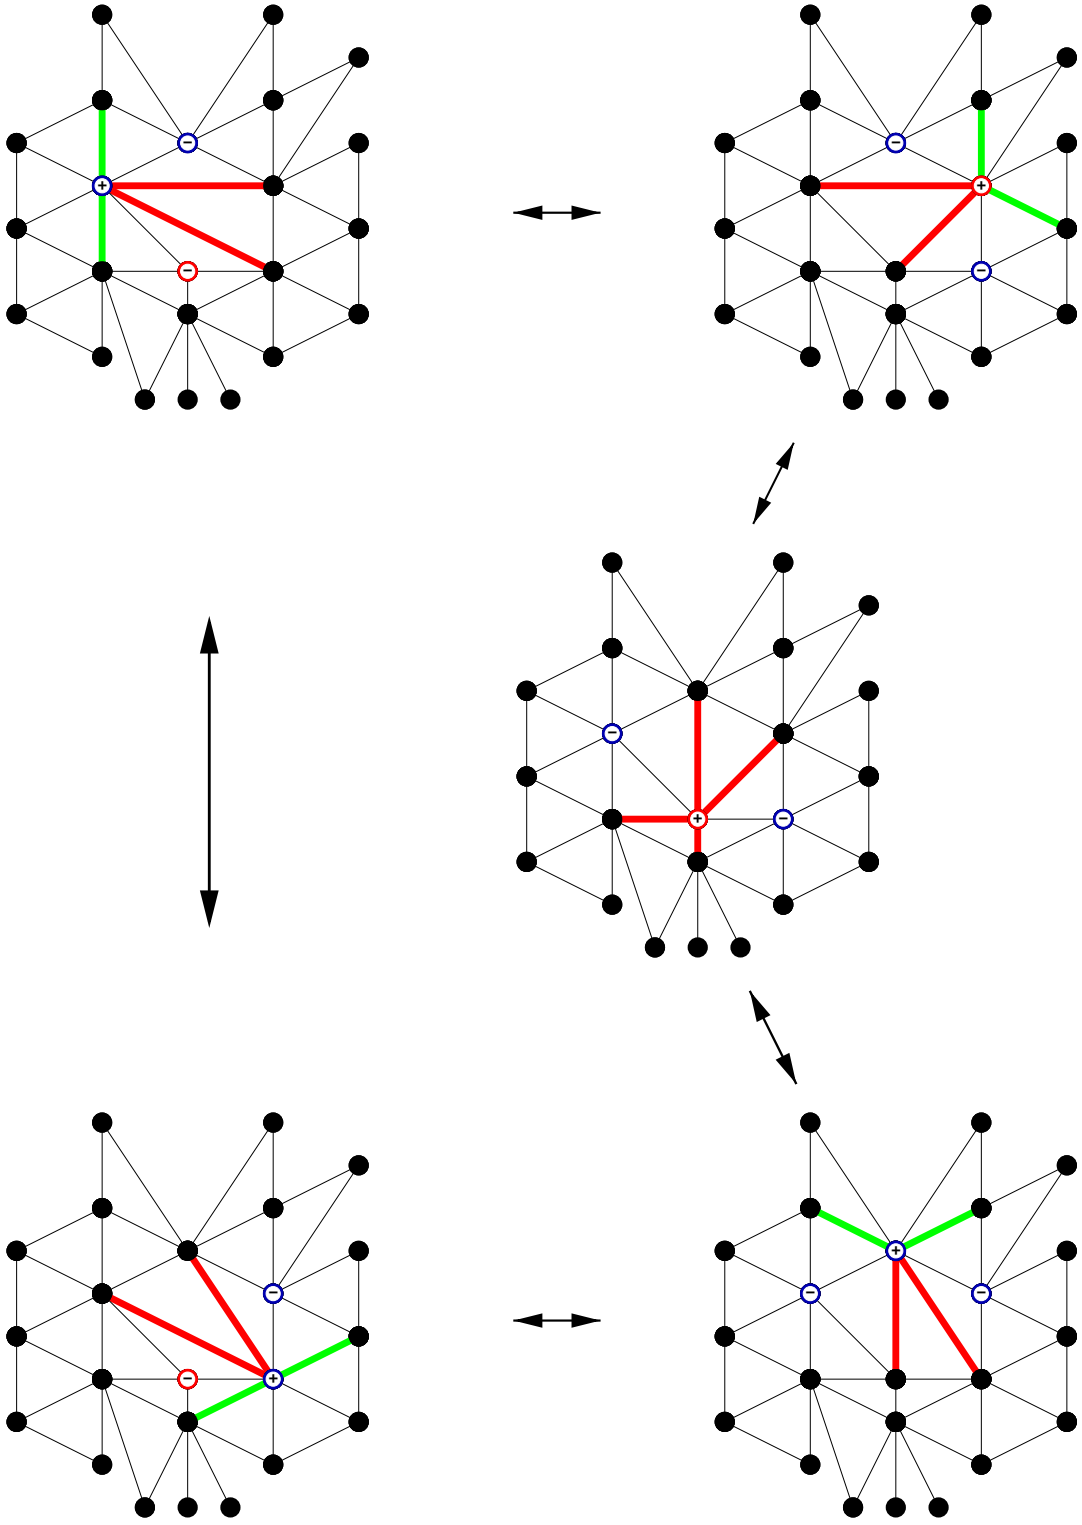

Figure 2.11: The central figure (with only red links) is symmetric along the axis −+−. If we flip the long vertical line, we arrive at the figure top-right. If we flip in it the red link which does not lead back to the center, we arrive at top-left. Flipping the red link which does not lead back to top-right, we arrive at bottom-left, then at bottom-right, and then back to the center. Since the same happens for the two lower links of the center, we see that the local state space is a figure "8" with 9 nodes of which 8 have two exits each. The state space can be symbolized as in Fig. 2.12.

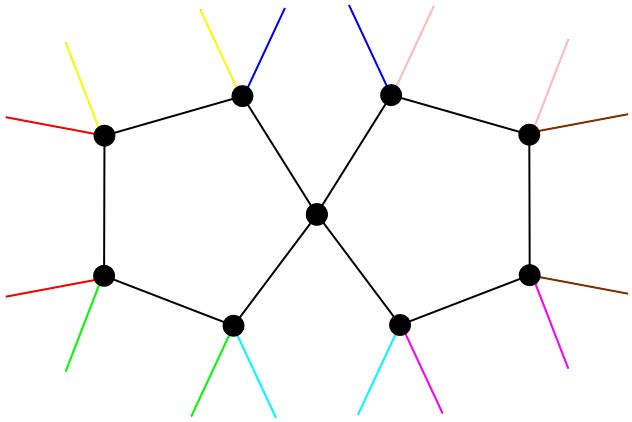

Figure 2.12: Each vertex represents one possible configuration during the collision of Fig. 2.11. Two vertices are linked if one can go from one configuration to the other by flipping a (red) link of Fig. 2.11. The pair enters and exits the collision through one of the 16 dangling links. If these 2 dangling links are the same or if they are of the same color, then the pair does not push the defect, otherwise, it is pushed by 1 step.

and no green ones. Moving a red link will visit the 6 figures sequentially. But moving the two lower red links in the lowest left figure will lead to another circle of five configurations, which is not shown in the figure. This collision case can be represented by a "state diagram" as in Fig. 2.12, where each node represents a state and each link represents the effect of flipping one of the colored links in Fig. 2.11. The pair enters the collision through a dangling link ℓ_1. It can wander around the vertices of the state diagram before exiting through a dangling link ℓ_2.

If $\ell_1 = \ell_2$, then it is as if the collision never occurred. In particular, the defect does not move. Furthermore, if ℓ_1 and ℓ_2 are of the same color, then the defect will remain in its initial position at the end of the collision. Using this remark and the diagrams of Fig. 2.11, one can explicitly compute the (conditional) probability that a pair pushes a defect if the collision is of the above type. This probability will be temperature independent.

The other 7 cases are treated similarly, and this completes the proof of Proposition 2.5.1.

Note that the proof means that collisions lead, on average to a *positive* probability of moving a defect. *This mechanism is the basic reason for the diffusive wandering of the defects in the triangulations. It is mediated by the collision of pairs with the defects. Clearly, if there are no pairs, the defects can not move by this mechanism, but only through much less probable events.*

2.6 Relevant and irrelevant pairs

In Sect. 2.4.2, we have seen that a pair lives long enough to explore its 1D path, before being destroyed by other mechanisms. We now analyze in detail what can happen during this exploration phase.

When a pair is created, it is one step away from its birthplace. It will then perform a

random walk on its predefined 1-dimensional path. Each time it comes back to its birthplace, it can die with probability $p_{\text{death}} = 1/3$ as shown in Fig. 2.10. *If this happens, the triangulation will not have changed.* We will call this an *ineffective pair*. The probability $P_{\text{I}} = P_{\text{I}}(\xi)$ can be estimated as follows:

Assume that a defect X is at a distance ξ from the birthplace of the pair. Then, by extending slightly the gambler's ruin principle [14], *the probability* $P_{\text{R}} = P_{\text{R}}(\xi)$ *that the pair actually can reach* X *is* $(1 + (\xi - 1) \cdot p_{\text{death}})^{-1} = O(1/\xi)$. This implies that the probability for any event implying X when starting from the birthplace depends on ξ, and in the case of many defects, on their average distance (which we call again ξ). Thus,

$$P_{\text{I}} = 1 - O(1/\xi) \ , \quad P_{\text{R}} = O(1/\xi) \ . \tag{2.9}$$

2.7 Time correlations at equilibrium

Here, we estimate the rate of change of triangulations (as a function of time). Since our triangulations are purely topological, we need to define what we mean by the distance between 2 triangulations T_1 and T_2 in $\mathcal{T}_N$ (the space of triangulations of the sphere with N labeled nodes). There are many possible choices, see *e.g.*, [15] many of which lead to equivalent metrics. The one defined below is convenient for our purpose.

Let $\{T_1, T_2\} \subset \mathcal{T}_N$. Consider a node n of T_1. The flower $f(n, T_1)$ of n is defined as the ordered cyclic set of all neighbors of n in T_1. Two flowers are then said to be equal if one can be obtained from the other by a cyclic rotation. We can now define the following metric on $\mathcal{T}_N$:

$$d(T_1, T_2) = \sum_{n=1}^{N} d_n(T_1, T_2) \text{ and}$$
$$d_n(T_1, T_2) = \begin{cases} 0 & \text{if } f(n, T_1) = f(n, T_2) \ , \\ 1 & \text{otherwise} \ . \end{cases}$$

Using this metric, we define the following *time correlation function*:

$$C(\vartheta) = 1 - \frac{d(T(t), T(t + \vartheta))}{N} \ ,$$

where $T(t)$ is the system state at time t. Our result for the decay of this function at equilibrium, *i.e.*, when $t \to \infty$, is as follows:

Proposition 2.7.1. *The correlation function* C *decays like*

$$C(\vartheta) = e^{-\vartheta/\tau_r} \ , \tag{2.10}$$

with a relaxation time τ_r *of the form*

$$\tau_r = O(e^{3\beta}) = O(\varepsilon^{-3}) \ . \tag{2.11}$$

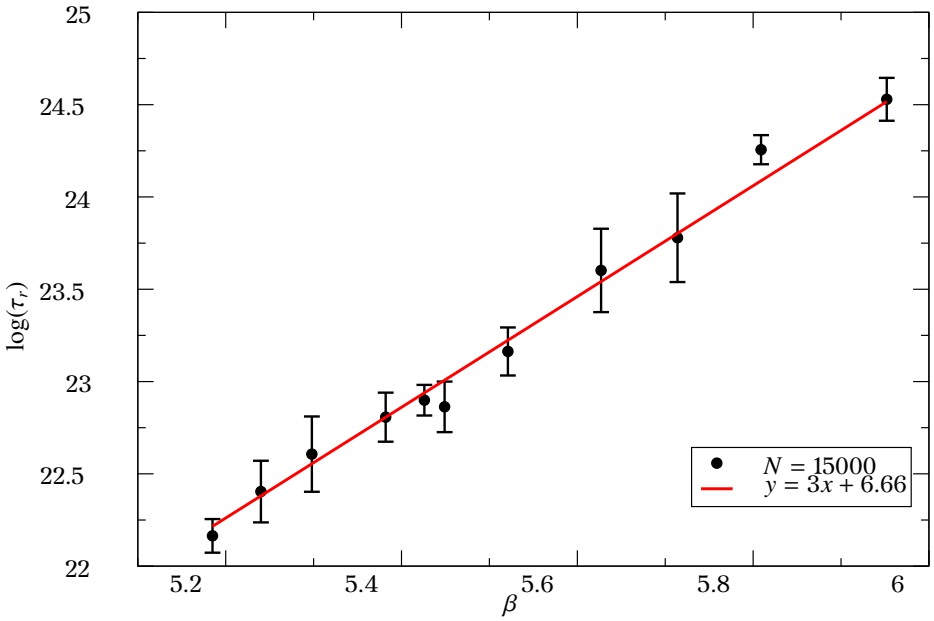

Figure 2.13: Decay rate of correlations at equilibrium. Numerical verification of Eqs. (2.10) and (2.11). The data are averages over 10 runs with $N = 15'000$. The error bars represent 1 standard deviation. The variable β is equal to $-\log(\varepsilon)$. The fits are for C between 0.5 and 0.001.

Proof. The correlation function $C(\vartheta)$ is nothing but the fraction of nodes whose flower is unchanged after ϑ time units. At equilibrium, the number of pairs p was established in Eq. (2.4) to be $p = 70/36 \cdot \varepsilon^2$. On the other hand, the density of defects in equilibrium is $O(\varepsilon)$ and hence, their average distance ξ equals $\xi = O(\varepsilon^{-1/2})$. By the estimates of Sect. 2.6 this means that the effective number of pairs which change the configuration in a permanent way is $O(p \cdot \varepsilon^{-1/2})$. We further saw in Sect. 2.5 that the number of collisions a relevant pair will undergo is a temperature independent constant $\nu = O(1)$. If ξ is the average distance between 2 defects, then, on average, this pair will change, on its way, the flowers of $2\nu\xi$ nodes. At time ϑ, each of these flowers is still unchanged with probability $C(\vartheta)$.

Since the pair makes a 1D random walk, all this happens within an average time interval $\delta\vartheta = \frac{1}{2}\nu^2\xi^2$. This in turn leads to

$$C(\vartheta + \delta\vartheta) = C(\vartheta) - 2pP_{\mathrm{R}}\nu\xi C(\vartheta) \ .$$

In the limit $\vartheta \gg \delta\vartheta$, we find

$$\dot{C}(\vartheta) = -4\frac{pP_{\mathrm{R}}}{\nu\xi}C(\vartheta) \ ,$$

and this leads to Eq. (2.10) with

$$\tau_r = \frac{\nu\xi(\varepsilon)}{4p(\varepsilon)P_{\mathrm{R}}(\xi)} \ . \tag{2.12}$$

Finally, using

$$P_{\mathrm{R}}(\xi) = O(1/\xi) = O\left(\varepsilon^{1/2}\right) \ , \tag{2.13}$$

Eq. (2.11) follows from Equations (2.12) and (2.13). □

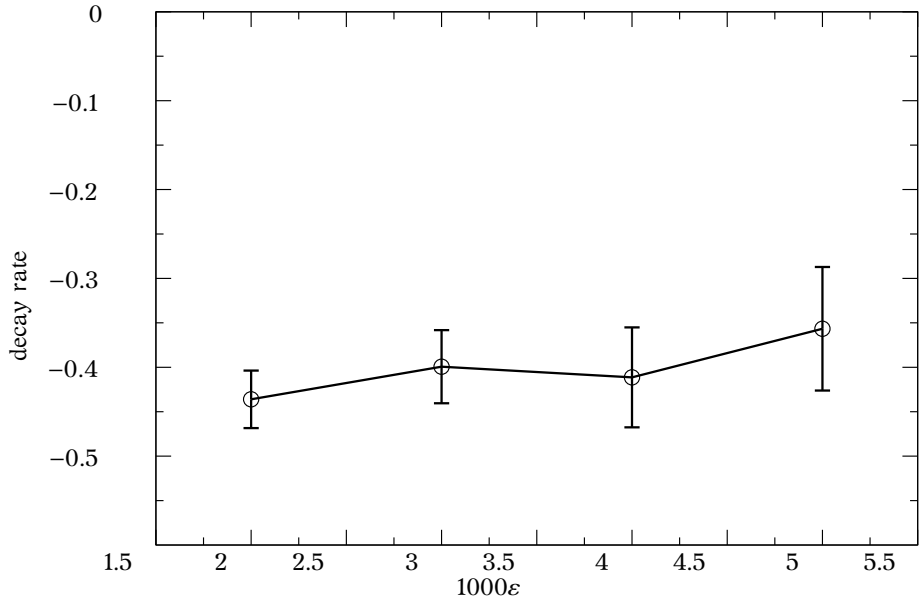

Figure 2.14: The decay rates of several simulations with $N = 15000$ and $\varepsilon = 0.002$ to 0.005.

2.8 The aging process

By the *aging process*, we mean the approach of the energy to its equilibrium value. Since the energy is by and large just the density $d(t)$ of defects we can formulate the result as

Estimate 2.8.1. *Under the assumptions $\varepsilon N > D_0$ and $\varepsilon < \varrho$ one has for the density d of defects:*

$$d(t) = O(\left(\varepsilon^2 t\right)^{-2/5}) \,. \tag{2.14}$$

Note that this result differs from that proposed in [16], where the decay rate was given as $(\varepsilon^2 t)^{-1/2}$. This difference is caused by our observation that the diffusion constant of the defects actually depends on their density, because, if they are rarer, the pairs, which are the only ones able to move them around, need longer to find them.

Power decay rates are extremely hard to distinguish, but we have performed some tests which are illustrated in Fig. 2.14. They give a slight advantage to a decay of -0.4 as compared to -0.5.

Proof. We study the aging process by assuming that, in approach to equilibrium, the system is in a quasistationary state, with charge density $c = E/N$. Here, and in the sequel, time will be in units of $\tau = (3N-6)/2$. Let $d(t)$ and $p(t)$ be the density of defects and pairs respectively. Then, up to terms of order $O(\varepsilon^3)$ one has $c = d + 2p$.

As we will see in this section, the quasistationarity assumption simply means that the relaxation of the energy is a consequence of the *annihilation of colliding defects*. The number of pairs is, up to fluctuations, essentially unchanged during the process we consider.

2.8.1 Three timescales

We saw that a fraction $1 - O(\varepsilon)$ of all occurring flips in the system do not change the energy, and are either motions of pairs or blinkers. Of those, the only relevant ones are the wandering pairs, which induce diffusion of the defects as we have seen in Sect. 2.7. The discussion of the equilibrium probabilities apply also to states close to equilibrium, which is the regime we want to consider now.

The pair dynamics happens on the time scale $\tau_{\text{pair}} = \tau$ and it conserves both the number of pairs $p(t)$ and the number of defects $d(t)$.

The next slower time scale concerns creation and annihilation of pairs. Even though this changes $p(t)$, it conserves $d(t)$. Whenever one of these events happens, defects are pushed around by the pairs with some geometrically defined probability, and this leads to a diffusion, whose constant $D(t)$ measures this second time scale $\tau_{\text{diffusion}} = D^{-1}(t)$.

The third time scale τ_{meeting} is related to collision rate $\gamma(t)$ of defects; $\tau_{\text{meeting}} = \gamma^{-1}(t)$. They undergo a 2D random walk. Sooner or later, 2 defects of opposite charges will meet and will form a new pair which will run on timescale τ until it annihilates. In the regime we consider, *only* this sequence of events (collision and running pair) of the dynamics destroys 2 defects and, as a consequence, is responsible for the relaxation of the energy. Given the 3 time scales, the derivation of the decay rate is now rather straightforward.

2.8.2 The quasistationarity assumption and the density of pairs

By the previous discussion,

$$\tau_{\text{meeting}}(t) \gg \tau_{\text{diffusion}}(t) \gg \tau_{\text{pair}}(t) = \tau = 1 \ .$$

The orders of magnitude of these quantities near equilibrium are

$$\tau_{\text{meeting}}(t) = O(\varepsilon^{-2} d^{-7/2}) \ , \quad \tau_{\text{diffusion}}(t) = O(\varepsilon^{-2} d^{-1/2}), \quad \tau_{\text{pair}}(t) = 1 \ .$$

Consider a system for which, at time 0, $d(0) \gg 1$ and $p(0) \gg 1$. It is clear that the relaxation of pairs is much faster than that of defects. We will assume that pairs are always at equilibrium density, *i.e.*, that creation and destruction rates of pairs are equal and $p(t)$ is independent of t.

Remark 2.8.1. *The above discussion implies that $p(t)$ is constant over time intervals of order $\tau_{\text{meeting}}(t)$. In fact, both creation and annihilation events necessitate the presence of defects so that the creation and destruction rates of pairs will be linear in $d(t)$ at low density. This implies that p depends on t only through the value of $d(t)$. By abuse of notation, we will write $p(d)$ instead of $p(d(t))$.*

The creation rate of pairs is $12 d \varepsilon^2$ and the destruction rate is simply $p(d)/\tau_{\text{lifetime}}$. Therefore, by balancing the rates, we find:

$$p(d) = 12 d \, \tau_{\text{lifetime}} \, \varepsilon^2 \ . \tag{2.15}$$

Since a pair needs to diffuse from one defect to the other in order to annihilate, we estimate that $\tau_{\text{lifetime}} = O\left(\xi^2\right) = O(d^{-1})$. This implies that the density of pairs is $p(d) = O(\varepsilon^2)$.

2.8.3 The diffusion constant of single defects

Repeating the arguments of Sect. 2.7 the average number of collisions ν and the average number of moved defects η are temperature independent constants. The diffusion constant of a defect is simply the probability that a given defect moves by one space unit during one time unit and it is given by

$$D \cdot d = \frac{2p(d)\, P_{\mathrm{R}}(\xi)\eta}{\nu^2 \xi^2} \ ,$$
$$D(d) = O\big(p(d) \cdot P_{\mathrm{R}}(O(d^{1/2}))\big) \ .$$

Using Equations (2.15) and (2.13), this leads to

$$D(d) = O(\varepsilon^2 \cdot d^{1/2}) \ .$$

2.8.4 Collision rate of single defects and relaxation coefficient

The annihilation of 2 diffusive particles $A + B \to \emptyset$ has been studied in depth in [17, 18, 19]. Here, we use the mean field argument of [17], to deduce the collision rates. However, there will also be particle creation. On the other hand, *e.g.*, in [19] creation is indeed considered, but the study is for a fixed substrate, namely the lattice $\mathbb{Z}^2$, while our study is on a more floppy domain.

Given a 2D gas of 2 particles A and B of equal densities $d/2$ such that the diffusion constants $D_A = D_B = D$, it can be deduced from [17] that the collision rate γ is

$$\gamma(d) = O(Dd^3) \ .$$

Extending this identity to a varying diffusion constant, we end up with

$$\dot{d} = -2\gamma(d) = -O(\varepsilon^2 \cdot d^{7/2}) \ ,$$

where we assumed that we are far enough from equilibrium to neglect the creation rate of defects. □

Note that this result differs from that proposed in [16], where the decay rate was given as $(\varepsilon^2 t)^{-1/2}$. This difference is caused by our observation that the diffusion constant of the defects actually depends on their density, because, if they are rarer, the pairs, which are the only ones able to move them around need longer to find them.

Chapter 3

Nuclei and Bounding the Number of 3-Balls

3.1 Introduction

In this chapter, we study the question of the number of triangulations of the 3-ball by tetrahedra. The case of the 2-ball was exactly solved by Tutte in [10]. He showed in particular that the number of rooted triangulations of the 2-sphere with N vertices is $O(1)\,N^{-5/2}(256/27)^N$. It is natural to ask if analogous bounds are true in higher dimension. Such results could have applications in models of Statistical Mechanics (foams [34], quantum gravity [20], or glassy dynamics [3, 5, 4] and the first chapter of this thesis) where the exponential rate of growth can be interpreted as an entropy. The problem of the existence of an exponential bound in 3-dimensions was formulated by Ambjørn, Durhuus and Jónsson in [28] and emphasized by Gromov in [33]: they asked whether the number of triangulations of the 3-sphere is bounded by C^N for some constant C when there are N tetrahedra (facets) in the triangulation. To date, this question remains open. However Pfeifle and Ziegler proved in [35] a super exponential lower bound for the number of triangulations of the 3-ball as a function of the number of vertices. This does not answer negatively Gromov's question (which is in terms of the number of tetrahedra) but makes the problem of proving an exponential bound in terms of the number of tetrahedra even more challenging.

There are several studies in the direction of answering the question, which we summarize now. In [24], Durhuus and Jónsson gave the construction of a class of triangulations for which they could show a bound of the form C^N. These triangulations are obtained by building a tree of tetrahedra, which is obtained by starting from a root tetrahedron and attaching tetrahedra to its faces, and then attaching further tetrahedra to the new open faces. Each tetrahedron is attached to the tree with just one face. It is a common feature of tree-like constructions that they lead to bounds of the form C^N: The prime example in our context is of course the celebrated work of Tutte [10] mentioned above. Coming back to Durhuus and Jónsson, once the tree is constructed, they now collapse adjacent faces of the tree in such a way that at the end of the procedure a triangulation of the 3-sphere is obtained. Their main result says that the number of ways in which to do this is again exponentially bounded. In this way, they construct a set of triangulations of the 3-sphere with tetrahedra which is exponentially bounded. They ask whether these are all possible triangulations.

In a later development, Benedetti and Ziegler [25], show that the Durhuus and Jónsson construction, which they call *locally constructible* (LC), does *not* capture all triangulations of the 3-sphere. Namely, they show that a 3-sphere with a 3-complicated knotted edge (made by tetrahedra) is not LC. They also carefully discuss relations between LC and other classes of constructibility.

In the present chapter, we define a larger class of triangulations, with a construction similar to that of Durhuus and Jónsson, but which uses more general basic elements than the simple tetrahedron, which we call *nuclei*. We prefer to work with 3-balls, and bounds on 3-spheres can be obtained from a bound on triangulations of a tetrahedron. This is usually done by removing a tetrahedron from the 3-sphere (see for example [25, Section 3]).

Nuclei are defined as triangulations of the 3-ball with the following special properties:

1. They have no internal nodes.

2. Internal faces have at most *one* external edge.

Obviously, the tetrahedron is a nucleus. The Furch-Bing ball [26], [29] and [27] and the Bing 2-room house [29] and [27], which are not nuclei, can be reduced by our procedure to one non-trivial nucleus, each. The smallest non-trivial nucleus we know of, given in Table 3.1, has 12 nodes, and 37 tetrahedra, of which 17 have no external face. Nodes are numbered from 1 to 12, and Table 3.1 gives a list of the 37 tetrahedra.

1	**3**	**4**	10	1	3	5	10	**1**	**3**	5	**11**	**1**	**4**	**6**	10	1	5	7	8
1	**5**	**7**	**10**	**1**	5	**8**	**11**	**1**	**6**	7	**8**	1	6	7	10	**2**	3	**5**	**9**
2	3	5	11	2	3	8	9	2	**3**	**8**	**11**	**2**	**5**	**6**	11	**2**	**6**	11	**12**
2	**7**	**10**	11	**2**	**7**	11	**12**	**2**	8	**9**	**10**	2	8	10	11	3	4	9	10
3	**4**	9	**12**	3	**5**	**9**	**10**	**3**	**8**	9	**12**	**4**	**5**	**6**	11	**4**	**5**	**7**	8
4	5	8	11	4	6	10	11	4	7	8	9	**4**	**7**	9	**12**	4	8	9	10
4	8	10	11	6	7	8	9	6	7	9	11	6	7	10	11	**6**	**8**	9	**12**
6	9	11	12	7	9	11	12												

Table 3.1: A nucleus with 12 nodes, and 37 tetrahedra, of which 17 have no external face. If a tetrahedron has an external face, its 3 nodes are shown in boldface.

Our approach is two-fold: Top-down, and bottom-up. In the top-down approach, we define a set of elementary moves which reduce an arbitrary triangulation of the 3-ball into a tree of nuclei, which are glued together by pairs of faces, each such face with 3 external edges. The tree can then be cut into a disjoint union of nuclei by cutting along these faces. The construction always transforms 3-balls to unions of 3-balls, and is thus implementable on a computer.

In the bottom-up approach, we start with any tree whose nodes are arbitrary nuclei, and we construct 3-balls from it by gluing adequate faces together. Not all possible gluings lead to 3-balls, but including also some inadequate gluings still leads to good bounds. Again, the procedure can be programmed on a computer.

Our main result is Theorem 3.5.17. It says that *if the number $\varrho(t, f_s)$ of face-rooted nuclei with t tetrahedra and f_s external faces has a bound of the form $\varrho(t, f_s) \leq C^t$ then the number of rooted triangulations of the 3-ball with t tetrahedra, f external faces and n internal nodes is bounded by C_*^{t+f+n}.*

In particular, since obviously, $f \leq 4t$ and $n \leq 4t$, we would get a bound C^t_{**}.

In summary, our work bounds the number of triangulations in terms of the number of nuclei. Thus, we remain with a new, but hopefully simpler, open question about the problem of exponential growth, namely does the number of face-rooted nuclei with t tetrahedra have an exponential bound in t? While we do not have any mathematical statements about this problem, the methodology of the proof of Theorem 3.5.17 allows for quite extensive numerical experimentation. The most important insight from this experimentation is as follows: It seems that if T is a nucleus with a k-complicated knot (or even braid), then the addition of (at most) k cones and decomposition with our algorithm leads to a tree of *tetrahedra*. Note that the trefoil knot is 1-complicated. Furthermore, Goodrick [32] showed that the connected sum of k trefoil knots is at least k-complicated.

We have analyzed a certain number of classical examples, with the findings summarized in Table 3.2.

Example	knot complication	# of cones added	ref.
Bing 2 room	no knot	1 cone	[29]
1 trefoil	1-complicated	1 cone	[26]
2 trefoils	2-complicated	1 cones	
3 trefoils	3-complicated	2 cones	[25, Figure 3]
4 trefoils	4-complicated	3 cones	
5 trefoils	5-complicated	3 cones	
figure 8	1-complicated	1 cones	
cinquefoil knot	1-complicated	1 cones	

Table 3.2: Experimental upper bound on the number of cones needed to decompose a triangulation into tetrahedra (For the definition of m-complicated, see [25]).

3.1.1 The method

The bounds on the number of triangulations are obtained by studying a set of elementary moves, detailed in Sect. 3.4.1. These moves either decompose the triangulation in two disjoint pieces (by cutting along an interior face with 3 edges on the boundary), or by taking away a tetrahedron with an external face and one internal node. Clearly, this leaves again two 3-balls on which we continue the decomposition. The other operations are "open" a ball along a carefully chosen edge (which we call "split-a-node-along-a-path") or opening one face with 2 external edges. These operations *increase* the number of tetrahedra in the triangulation, but they prepare the moves in which the 3-ball can be cut, and the internal nodes can be eliminated. One of the main novelties of this construction is the observation that this can be done with *few* additional tetrahedra: This follows from a careful analysis of cuts of the 2-dimensional hemisphere attached to any external node. Since this is an important bound, we devote Sect. 3.3 to its proof. In Sect. 3.2, we introduce the (standard) terminology for the pieces of any triangulation. In Sect. 3.4 we combine the 4 moves described above to show how a general triangulation can be decomposed into a set of nuclei. In Sect. 3.5, we perform the bottom-up procedure and show how one bounds the number of triangulations of

the 3-ball in terms of trees whose nodes are (rooted) nuclei, extending in this way the earlier work of [24] and [25].

3.1.2 Comparison with 2d

It is useful to compare our method to what can be done in 2d. In 2d we have a set of triangles. Any triangulation can be obtained in the following way: First, construct a tree of triangles, adding each triangle with only one face to the existing tree. This object has no internal nodes. Now, glue together adjacent faces of the tree, recursively. In this way one can obtain all triangulations of any polygon.

The inverse operation, while intuitively clear, is a little harder to describe, and we just sketch the procedure. Given any internal node x at distance 1 from the polygon, say, connected to node n of the polygon, we can split the edge (n, x) by doubling the node n into a pair n', n'', so that the edges (n', x) and (x, n'') are now external edges and x is *promoted* to an external node. All internal nodes can recursively be brought to the surface in this way. We then have a tree, and the tree can be decomposed into triangles by cutting all internal edges with 2 external nodes. At the end, the basic objects are triangles.

Clearly, therefore, the basic objects in 2d are

2a) internal edges with 2 external nodes

2b) internal nodes (at distance 1) from the polygonal boundary

In 3d, there are many more possibilities, and our procedure will eliminate all those which can be eliminated. The ones which we can deal with are

3a) internal faces with 3 external edges: this corresponds to case 2a) above and will be cut by cut-a-3-face

3b) internal faces with 2 external edges, and therefore one internal edge with 2 external nodes. This resembles 2b) and is dealt with by open-a-2-face.

3c) an internal node x which is the tip of a tetrahedron t whose opposite face is external. One can just eliminate t and x becomes external. This is the second case which corresponds to 2b). We call this C0 later.

3d) an internal node x which is the corner of a face f whose opposite edge is external (but not C0). Again, a sub-case of 2b). This is dealt with split-a-node-along-a-path, and will be called C1.

3e) an internal node x which is the end of an edge e whose opposite end is external (but not C1). Again, a sub-case of 2b). This will be called C2 and reduced to C1 with split-a-node-along-a-path.

The elementary objects are those left over after all these decompositions are performed. In 2d, those objects are just triangles, which makes the counting possible. In 3d these are nuclei. Non-trivial nuclei exist, and they must carry the information about the complications of 3 dimensional topology, since all the other problems have been eliminated. In particular, internal faces of nuclei have 0 or 1 external edges.

3.2 General definitions and notations

3.2.1 Internal and external objects, flowers

We introduce some notation which we apply to triangulations and *tetrahedrizations* (which we also call triangulations when no confusion is possible):

We start with triangulations of S^2. These will have f_s faces, n_s nodes and e_s edges, where the subscript s stands for "surface".

We also consider tetrahedrizations of a ball, which are collections of tetrahedra, whose faces, edges and nodes satisfy the usual topological conditions of piecewise linear triangulations. The boundary of such a tetrahedrization is a triangulation of S^2. We say that t is the number of tetrahedra, f_{tot} the number of faces, e_{tot} the number of edges, and n_{tot} the number of nodes. Faces, edges, and nodes which are not among those of the triangulation of S^2 are called *internal*; the others are called *external*. It will be useful to observe that tetrahedra can have up to 4 external faces, internal faces can have up to 3 external edges, internal edges up to 2 external nodes. We will use the subscript i for internal objects.

Obviously,

$$f_{tot} = f_s + f_i \ , \qquad e_{tot} = e_s + e_i \ , \qquad n_{tot} = n_s + n_i \ .$$

From the Euler relations and trivial geometry, we have the relations

$$\begin{aligned} t - f_{tot} + e_{tot} - n_{tot} &= -1 \ , \\ f_s - e_s + n_s &= 2 \ , \\ 3f_s &= 2e_s \ , \\ 4t &= 2(f_{tot} - f_s) + f_s \ . \end{aligned} \tag{3.1}$$

This leaves us with 3 free variables, which we choose as

$$t, f_s, \text{and } n_i \ .$$

Note that f_s is always even.

Definition 3.2.1. *We use the term* f-vector *for the three variables* $\langle t, f_s, n_i \rangle$ *where* $f_s \geq 4$.

3.2.2 Notation and flowers

- If n_1 and n_2 are 2 distinct nodes, then we denote by (n_1, n_2) the edge connecting the two (if it exists).
- Similarly, if $n_i : i = 1, 2, 3$ are 3 distinct nodes, then (n_1, n_2, n_3) is the face (triangle) with these 3 corners (if it exists).
- If e is an edge and n is a node not in e then (n, e) denotes the face (triangle) with the edge e and the node n (if it exists).

This notation is easily generalized to the case where we consider simplices of dimension 3:

- If n is a node and f is a face not containing n, then (n, f) is the tetrahedron with f as a face and n as the opposite corner (if it exists).

- Similarly, if e is an edge and $n_1, n_2 \notin e$ are 2 distinct nodes then (n_1, n_2, e) is the unique tetrahedron containing all of them (if it exists).
- Finally, if e_1 and e_2 are two edges without common nodes, then (e_1, e_2) is the tetrahedron containing both edges (if it exists).

Paths of nodes connected by edges will be denoted as $\gamma = [n_1, n_2, \ldots, n_k]$ and the union of 2 disjoint paths γ_1, γ_2 (connected by one or both endpoints) will be denoted by $\gamma_1 \cup \gamma_2$.

We next define what we mean by *flowers*. Here, we adapt the conventions to the tetrahedrization of a triangulated sphere S^2. Nodes, edges, and faces are called *external* if they lie entirely in S^2. All others are called *internal*. Consider an external node n_*[1]. We define its 2 flowers:

- The *external flower* $\mathcal{E}(n_*)$ of n_* is the set of all edges e not containing n_* for which (n_*, e) is an external face. Clearly, $\mathcal{E}(n_*)$ is a polygon.
- The *internal hemisphere* $\mathcal{I}(n_*)$ of n_* is the set of all faces f not containing n_* for which (n_*, f) is a tetrahedron. It is easy to see that $\mathcal{I}(n_*)$ is a 2d triangulation whose boundary is the polygon $\mathcal{E}(n_*)$.

We will say that the external flower of an *internal* node x_* is empty. We define the internal (hemi-)sphere $\mathcal{I}(x_*)$ (or simply flower) of x_* as the set of all faces f not containing x_* for which (x_*, f) is a tetrahedron. This is a triangulation of S^2.

We also define the *external flower* $\mathcal{E}(e)$ *of an external edge* e as the 2 nodes n_1, n_2 for which (n_i, e) are 2 external faces. Similarly, the *internal hemisphere* $\mathcal{I}(e)$ *of the external edge* e is defined as the set of all edges e' such that (e, e') is a tetrahedron. By hypothesis, $\mathcal{I}(e)$ is a 1-d triangulation whose boundary is $\mathcal{E}(e)$. Note that there might be internal nodes at distance 1 from e which are not in $\mathcal{I}(e)$.

3.3 Some geometrical considerations: Two-colored paths in a triangulation

Throughout this section all triangulations are 2d triangulations. We describe here properties of paths in a 2d triangulation of a polygon. These properties will play a crucial role when we will bound the effects of moving internal nodes of a 3d triangulation to the surface. However, they are totally independent of the remainder of the chapter.

Lemma 3.3.1. *Let $\mathcal{K}$ be a 2d triangulation with p boundary edges and n internal nodes. Then the number of interior edges in $\mathcal{K}$ is $3n + p - 3$.*

Proof. The proof follows from the Euler relations and is left to the reader. □

Lemma 3.3.2. *Consider a polygon P and let $\mathcal{K}$ be any triangulation whose boundary is P, with $n > 0$ internal nodes. For each node $x \in \mathcal{K} \setminus P$, there are at least 3 simple disjoint paths in the interior of $\mathcal{K}$ connecting it to 3 different points of P.*

[1] We use n_*, m_* and the like for external nodes, and $x_*, y_*, \ldots$ for internal ones.

Proof. Any triangulation of S^2 is 3-connected. Complete $\mathcal{K}$ into a triangulation of S^2 by adding a cone over its boundary. Let m be the apex of the cone. Then there are at least 3 disjoint simple paths connecting x to m, [31]. Any such path must intersect P, and we take the first intersection point. □

We assume now that the nodes of P are labeled.

Definition 3.3.3. *A triangulation $\mathcal{K}$ with $P = \partial\mathcal{K}$ is called an* admissible triangulation *if the following conditions are met:*

K1: The boundary $P = \partial\mathcal{K}$ has at least 2 different labels.

K2: The nodes in $\partial\mathcal{K}$ with a given label form one connected arc of $\partial\mathcal{K}$.

K3: The ends of any internal edge connecting 2 nodes of $\partial\mathcal{K}$ have different labels.

The Fig. 3.1 illustrates the definition.

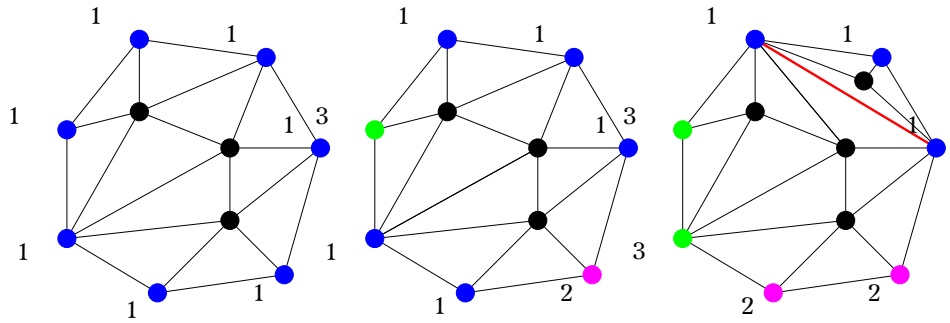

Figure 3.1: An illustration of the conditions K1)–K3). Left: since there is only one label, K1) is violated. Center: The region with label 1 is not connected; K2) is violated. Right: There is an internal edge (red) connecting two nodes with the same label; K3) is violated.

Definition 3.3.4. *An admissible triangulation of a polygon is called a* non-trivial triangulation *if it has at least one internal edge. (The only admissible trivial triangulation is an admissible triangle.)*

We begin with an auxiliary lemma.

Lemma 3.3.5. *Let $\mathcal{K}$ be an admissible triangulation whose boundary is the polygon P. Given two non-adjacent boundary nodes a and b with different labels at least one of the two alternatives below holds:*

1) There is a simple path γ joining a and b without any other node in P,

2) There is an edge (x, y) joining the two pieces of $P \setminus \{a, b\}$.

Postponing the proof of Lemma 3.3.5 we have

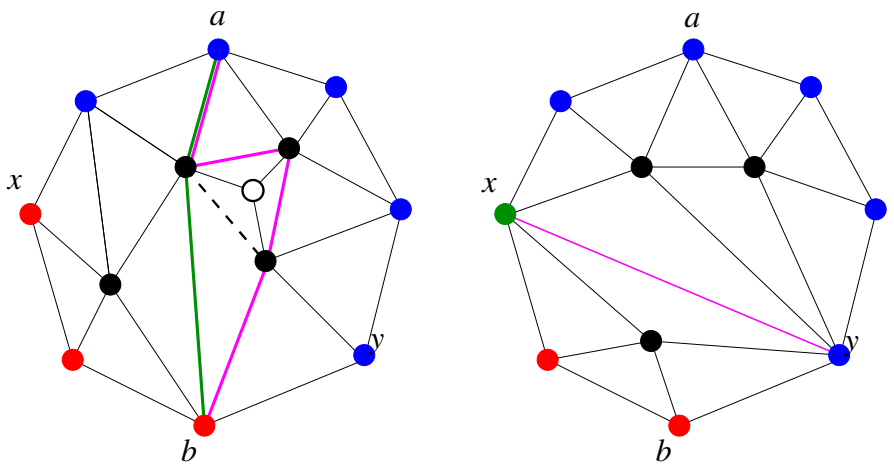

Figure 3.2: The 2 alternatives of finding a path connecting two different labels. Left: There is an interior path between a and b. Right: There is no such path, but then, one can always find an edge connecting two different labels (by K3), (not necessarily the same as a and b). The left panel also illustrates the necessity of choosing a shortest path. For example, choosing the magenta path, the dashed edge will violate K3) in the next step of the procedure.

Proposition 3.3.6. *Assume $\mathcal{K}$ is a non-trivial admissible triangulation with $\partial\mathcal{K} = P$. Then, there exists a simple path γ along internal edges of $\mathcal{K}$ which connects two points in $P = \partial\mathcal{K}$ with* different labels. *It cuts $\mathcal{K}$ in two pieces $\mathcal{K}_{\mathrm{L}}$ and $\mathcal{K}_{\mathrm{R}}$. The path γ can be chosen in such a way that labeling the new boundary piece (namely the interior nodes of γ) in $\mathcal{K}_{\mathrm{L}}$ and $\mathcal{K}_{\mathrm{R}}$ with a label different from the ones used so far, both $\mathcal{K}_{\mathrm{L}}$ and $\mathcal{K}_{\mathrm{R}}$ are again admissible.*

Proof. By admissibility, we know that not all nodes on P have the same label. We distinguish two cases:

- P has more than 3 nodes. In this case, take 2 non-adjacent nodes a and b with different labels and apply Lemma 3.3.5. If 2) holds, then we take γ as the edge connecting x and y. By K3) they have different labels. Otherwise, there is an internal path connecting a and b. We take a shortest path, γ.

- P is a triangle. In this case, $\mathcal{K}$ can be seen as a triangulation of the sphere S^2. In particular, it is 3-connected. We take 2 (adjacent) nodes a and b of P with different labels. Since $\mathcal{K}$ is non-trivial and 3-connected, there are (at least) 3 disjoint simple paths connecting a and b. At most 2 of these paths are on the boundary P, which leaves at least 1 path in the interior of the triangulation $\mathcal{K}$. We take a shortest internal path γ connecting a and b.

In all cases, we find 2 boundary nodes with different labels and a shortest internal path γ connecting them. Cutting along the path γ, we obtain 2 pieces $\mathcal{K}_{\mathrm{L}}$ and $\mathcal{K}_{\mathrm{R}}$. If γ is just one edge then inspection shows that K1)–K3) hold. If not, K1) and K2) are obviously true by construction; we have to show that K3) is also true. Giving a new label, say L, to the interior nodes of γ, we have to show that there are no edges connecting any two non-consecutive nodes with label L. But if there were, the path would not be minimal. □

Proof of Lemma 3.3.5. The reader may want to look at Fig. 3.2. Assume 1) does not hold. This means that one cannot draw 3 disjoint paths between a and b, as the middle one would satisfy 1). We can take the two disjoint paths to go along the two boundary segments between a and b. By Menger's theorem [31], and since a and b are not adjacent, there must then be 2 nodes x and y (other than a or b) such that all paths from a to b must pass through at least one of them. Since the boundary paths are candidates, we see that x and y are in P, one per arc connecting a and b. Consider now the path from a to b along P which goes through x. Modify it so that instead of going through x it goes through the flower of x. We get a new path from a to b which does not go through x. This means that the new path goes through y implying that y is in the flower of x. Thus, x and y are connected by an edge.

This completes the proof. □

3.4 Part I: Reducing any triangulation into a set of nuclei

3.4.1 The elementary moves

In this section we define the elementary moves which transform any triangulation into a (set of) nuclei. The first two moves, which we call *open-a-2-face* and *cut-a-3-face*, are used to transform any triangulation with no internal nodes into a set of nuclei, and the third and fourth move, which we call *remove-1-tetra* and *split-a-node-along-a-path*, are used to remove all internal nodes of a triangulation.

Henceforth, T will denote a triangulated 3-ball with f-vector $\langle t, f_{\rm s}, n_{\rm i}\rangle$.

Cut-a-3-face

Let (n_1, n_2, n_3) be an internal face with its 3 edges on the surface ∂T of T. Then, it cuts the 3-ball T into 2 distinct parts. We simply separate these 2 parts and we get 2 "smaller" 3-balls, as shown in Fig. 3.3.

If $\langle t, f_{\rm s}, n_{\rm i}\rangle$, $\langle t_1, f_{1,{\rm s}}, n_{1,{\rm i}}\rangle$ and $\langle t_2, f_{2,{\rm s}}, n_{2,{\rm i}}\rangle$ are the f-vectors of the initial ball and the 2 new ones, then we have

$$t = t_1 + t_2 \ , \qquad f_{\rm s} = f_{1,{\rm s}} + f_{2,{\rm s}} - 2 \ , \qquad n_{\rm i} = n_{1,{\rm i}} + n_{2,{\rm i}} \ .$$

Open-a-2-face

Consider 3 external nodes n_*, n_1, n_2 of T which form a triangle (n_*, n_1, n_2). We assume that (n_*, n_1, n_2) is an internal face, with (n_1, n_2) an internal edge, and the two other edges external. Let $\mathcal{I}$ and $\mathcal{E}$ be the internal and external flower of the external node n_*. As we have already stated, $\mathcal{I}$ is a triangulation with $\partial\mathcal{I} = \mathcal{E}$. By hypothesis, the edge (n_1, n_2) divides $\mathcal{I}$ into 2 distinct sets of faces. The operation open-a-2-face consists in removing n_* and all tetrahedra attached to it, replacing it by $n_{*,1}$ and $n_{*,2}$ and attaching each of these 2 new nodes to all faces of one of the two parts of $\mathcal{I}$, see Fig. 3.4. *This operation transforms a triangulation of the 3-ball into a triangulation of the 3-ball.* If $\langle t, f_{\rm s}, n_{\rm i}\rangle$ and $\langle t', f'_{\rm s}, n'_{\rm i}\rangle$ are the f-vectors of the initial ball and the resulting ball, then we have

$$t' = t \ , \qquad f'_{\rm s} = f_{\rm s} + 2 \ , \qquad n'_{\rm i} = n_{\rm i} \ .$$

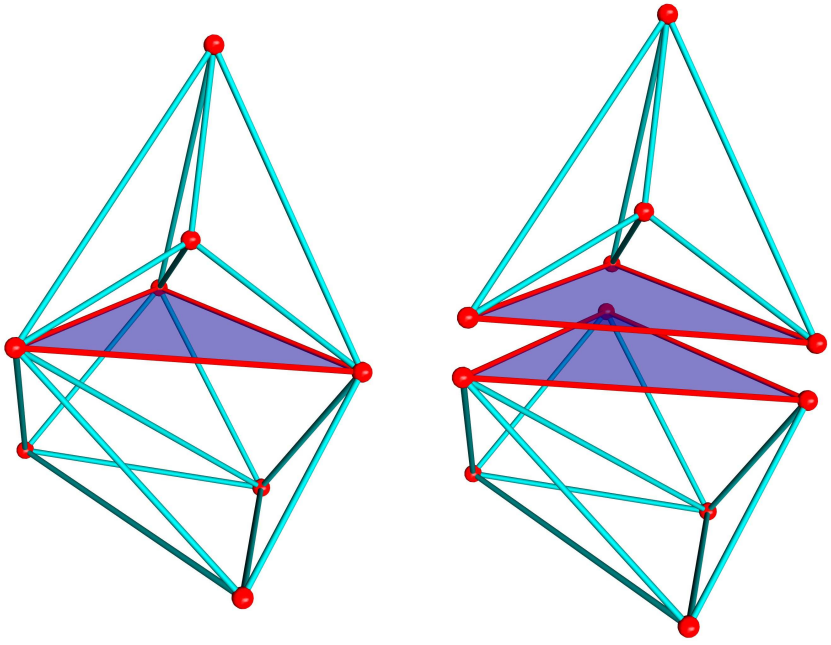

Figure 3.3: Sketch of cut-a-3-face.

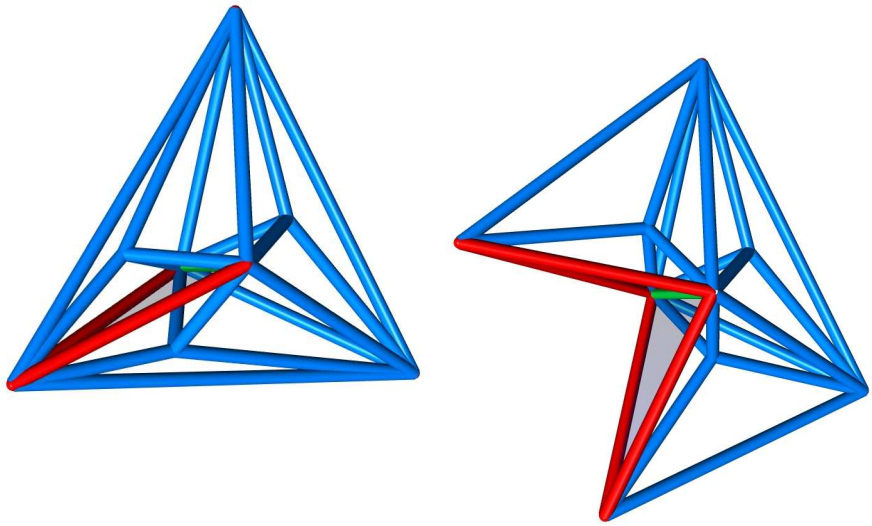

Figure 3.4: Sketch of open-a-2-face.

We will say that the f-vector changes by $\langle 0, +2, 0\rangle$.

Remove-1-tetra

Definition 3.4.1. *A* removable tetrahedron *is any tetrahedron t with one internal node and one external face.*

The operation remove-1-tetra is as follows: let $t_* = (x_*, n_1, n_2, n_3)$ be a removable tetrahedron with internal node x_* and external face (n_1, n_2, n_3). We simply remove t_* and its external face; the internal node x_*, the 3 internal edges and the 3 internal faces of t_* all become external. The f-vector $\langle t, f_s, n_i\rangle$ changes to $\langle t-1, f_s+2, n_i-1\rangle$; the change of f-vector is $\langle -1, 2, -1\rangle$.

Split-a-node-along-a-path, hemispheres and pieces

Consider an external node n_* of T and its internal hemisphere $\mathcal{I} = \mathcal{I}(n_*)$, see Fig. 3.5 for an illustration. By definition of a triangulation, $\mathcal{I}$ is a 2d triangulation of a polygon.

Definition 3.4.2. *A* splitting path γ *is any simple path in $\mathcal{I}$ which connects two different vertices on $\partial\mathcal{I}$ and contains no edge of $\partial\mathcal{I}$.*

Let γ be a splitting path. Clearly it divides $\mathcal{I}$ into 2 pieces $\mathcal{K}_L$ and $\mathcal{K}_R$ with $\mathcal{I} = \mathcal{K}_L \cup \mathcal{K}_R$ and $\mathcal{K}_L \cap \mathcal{K}_R = \gamma$.

The move split-a-node-along-a-path γ is defined as follows:

1. Remove the node n_* and all tetrahedra having n_* as a corner
2. Add 2 new nodes $n_{*,L}$ and $n_{*,R}$
3. For each face $f_* \in \mathcal{K}_L$ add the tetrahedron $(n_{*,L}, f_*)$
4. For each face $f_* \in \mathcal{K}_R$ add the tetrahedron $(n_{*,R}, f_*)$
5. For each edge $e \in \gamma$ add the tetrahedron $(n_{*,L}, n_{*,R}, e)$

Note that by construction, one of the nodes on $\partial\mathcal{K}_L$ is $n_{*,R}$, and the edges in $\mathcal{K}_L$ originating from $n_{*,R}$ reach (the image of) γ. Analogous statements hold for $\mathcal{K}_R$.

Definition 3.4.3. *In the construction above, we refer to $\mathcal{K}(n_{*,L}) = \mathcal{K}_L$ as the left* piece*. It is simply the subtriangulation obtained from the hemisphere $\mathcal{I}(n_{*,L})$ after removing the cone connecting $n_{*,R}$ to every node of γ. Similarly, we define the right piece $\mathcal{K}(n_{*,R}) = \mathcal{K}_R$.*

Remark 3.4.4. *Hemispheres $\mathcal{I}$ and pieces $\mathcal{K}$ will play an important role in our construction. Some statements will be given for hemispheres, others for pieces and so it is important to be able to distinguish between the two definitions.*

Remark 3.4.5. *A splitting path γ is always associated with a hemisphere $\mathcal{I}$ and not with a piece $\mathcal{K}$. We will see that, under some conditions, a simple path $\tilde{\gamma}$ connecting two nodes of the boundary of a piece $\mathcal{K}$ can be extended into a splitting path γ.*

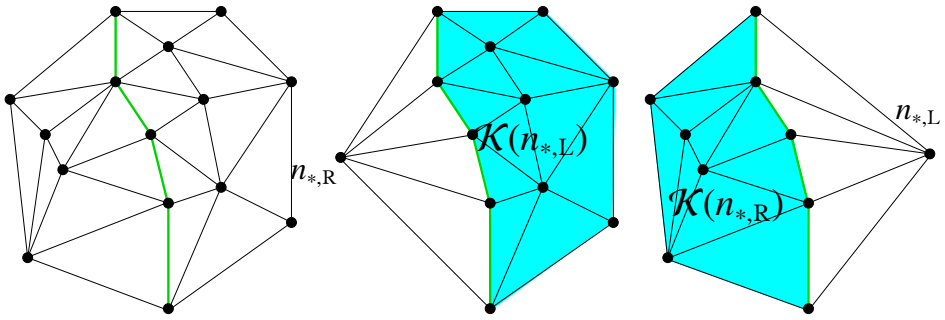

Figure 3.5: The left panel shows the internal hemisphere $\mathcal{I}(n_*)$ of n_*. The node n_* is not shown, but is connected to every node of $\mathcal{I}(n_*)$. We split n_* into $n_{*,\mathrm{L}}$ and $n_{*,\mathrm{R}}$ along the green path γ. The other 2 panels show the internal hemispheres $\mathcal{I}(n_{*,\mathrm{L}})$ and $\mathcal{I}(n_{*,\mathrm{R}})$ of the 2 new nodes ($n_{*,\mathrm{L}}$ in the center panel, $n_{*,\mathrm{R}}$ in the right one (again, they are not shown but connected to every node we draw)). Notice that each internal node of the green path γ is at distance 1 from $n_{*,\mathrm{R}}$ resp. $n_{*,\mathrm{L}}$ Also, the edges originating in $n_{*,s}$, $s \in \{\mathrm{L},\mathrm{R}\}$ have been added during the split.

Definition 3.4.6. *The new nodes $n_{*,\mathrm{R}}$ and $n_{*,\mathrm{L}}$ are called the* children *of n_*.*

Lemma 3.4.7. *The move split-a-node-along-a-path transforms a 3-ball into a 3-ball. The f-vector $\langle t, f_\mathrm{s}, n_\mathrm{i}\rangle$ is mapped to $\langle t+|\gamma|, f_\mathrm{s}+2, n_\mathrm{i}\rangle$, where $|\gamma|$ is the number of edges in γ.*

The f-vector changes by $\langle|\gamma|,+2,0\rangle$. In particular, the number of tetrahedra *increases*. But we will show that this increase can be controlled.

Proof. The count of the f-vector is as follows: Removing and adding the tetrahedra in steps 1,3,4 above does not change their number. The number of external faces increases by two, namely the two external faces sharing the new edge $(n_{*,\mathrm{L}}, n_{*,\mathrm{R}})$. And each internal face (n_*, e) which connected n_* to an edge e in γ gives rise to a new tetrahedron $(n_{*,\mathrm{R}}, n_{*,\mathrm{L}}, e)$. There are $|\gamma|$ such faces and so the f-vector is seen to change by $\langle|\gamma|,+2,0\rangle$, as asserted.

□

3.4.2 Summary

In the sequel, we want to bound the effect of removing internal nodes, since our building blocks are the nuclei, which do not have any internal nodes. Eliminating the internal nodes will cost the addition of tetrahedra, and the issue here is how many are needed to obtain a ball without internal nodes. Internal nodes disappear when we perform the remove-1-tetra operation, and only then.

Before starting the bounds proper, we explain here the point of our construction, based on the evolution of the f-vectors $\langle t, f_\mathrm{s}, n_\mathrm{i}\rangle$. Open-a-2-face costs a change $\langle 0,2,0\rangle$, and split-a-node-along-a-path costs $\langle|\gamma|,2,0\rangle$, where $|\gamma|$ is the length of the path along which we cut. In principle, each path γ might have a length proportional to the number of nodes, which in turn would imply that the sum of the lengths of all paths exceeds $O(n_{\mathrm{tot}}^2)$. So one needs a strategy which improves this naive bound.

While we cut, new external edges appear, and also, new external edges appear when we remove a tetrahedron which costs $\langle -1, 2, -1\rangle$. But it is only this operation which reduces the number of internal nodes. So, there are two opposing tendencies. One is the preparation of promoting an internal node into an external one, and it *adds* many tetrahedra, and the other is remove-1-tetra, which reduces the number of internal nodes by 1.

The real issue is thus to bound the *number of added tetrahedra per removed internal node*. We will perform this bound in terms of the number $e_{\rm i}$ of internal edges. Our main result is Corollary 3.4.14 which says that the number of internal edges grows by no more than $C_\Delta(t + n_{\rm i})$. The Euler relations Eq. (3.1) allow to express t as a function of $e_{\rm s}$, $f_{\rm s}$, and $n_{\rm i}$,

$$t = e_{\rm i} - n_{\rm i} + f_{\rm s}/2 - 1 \ . \tag{3.2}$$

Therefore, and since $n_{\rm i} < 4t$, $f_{\rm s} = 2n_{\rm s} - 4 < 4e_{\rm i}$ and $e_{\rm i} < 6t$, Corollary 3.4.14 implies that the elimination of all $n_{\rm i}$ internal nodes leads to an f-vector of the form

$$\langle t, f_{\rm s}, n_{\rm i}\rangle \to \langle t', f'_{\rm s}, 0\rangle \ ,$$

with

$$e'_{\rm i} < e_{\rm i} + 5C_\Delta \cdot t < (6 + 5C_\Delta) \cdot t = C/4 \cdot t \ ,$$

and therefore

$$f'_{\rm s} < C \cdot t \ , \quad t' < C \cdot t \ ,$$

with a finite constant C which is *independent* of the triangulation.

3.4.3 Removing internal nodes

This is the most difficult, and novel, part of our construction.

Definitions and strategy

Given any triangulation, the *depth* D_x of a node x is the minimal number of connected edges needed to reach the boundary, starting from x. We also say that the triangulation has maximal depth $d_{\max} = \max_{n\in T} D_n$. Our strategy consists in reducing the depth of all internal nodes by 1. The depth will be reduced by working on all internal nodes of depth 1 and moving them to the surface. If the maximal depth of a triangulation was $d_{\max}$ it will end up being of depth $d_{\max} - 1$. We repeat this procedure until no internal nodes remain.

So, consider now an internal node x_* of depth 1. It comes in 3 flavors which we call *C0*, *C1* and *C2*:

C0: x_* is the internal node of a removable tetrahedron.

C1: x_* is not of type C0 but is in a face (x_*, n_*, m_*) where (n_*, m_*) is an external edge.

C2: x_* is neither of type C0 nor C1.

Obviously, in case C0, we would just remove the tetrahedron, promoting x_* to the surface. The other cases are more complicated and need a careful estimate in terms of edges and faces which appear when the nodes are brought to the surface. No node will ever disappear in these constructions, but some will be doubled (split) and edges will be added.

So we begin with a triangulation T_0 of maximal depth $d_{\max} > 0$. We define $\mathcal{L}_\ell = \{n : D_n = \ell\}$, the set of nodes at depth ℓ in the triangulation T_0. These *original depths* should be viewed as labels assigned to each node. If a node is split, its children inherit the label. If a node comes closer to the surface in the procedure below, its *actual depth* decreases, but its label (the original depth) does *not* change.

We then iterate the following 4 steps until no internal nodes remain, for $\ell = 0, \dots d_{\max} - 1$. Each iteration transforms the ball T_ℓ into a new ball $T_{\ell+1}$, such that the nodes of $\mathcal{L}_{\ell+1}$ are external.

Assume iteration $\ell - 1$ is completed: Then we say that in T_ℓ the nodes of $\mathcal{L}_\ell$ have become external. Given such a node $n_* \in \mathcal{L}_\ell$, we consider its hemisphere $\mathcal{I}(n_*)$ (in T_ℓ).

An internal node $x_* \in \mathcal{I}(n_*)$ of type C2 can be *promoted* to an internal node of type C1 by drawing a path $\gamma \subset \mathcal{I}(n_*)$ that goes through it and splitting n_* into $n_{*,\mathrm{L}}$ and $n_{*,\mathrm{R}}$ along γ. Then, $n_{*,\mathrm{R}}$ is among the edges $\mathcal{E}(n_{*,\mathrm{L}})$ and therefore $(n_{*,\mathrm{L}}, n_{*,\mathrm{R}}, x_*)$ is a face with the external edge $(n_{*,\mathrm{L}}, n_{*,\mathrm{R}})$, so that x_* is now of type C1.

In a similar manner, a node $x_* \in \mathcal{I}(n_*)$ of type C1 can be promoted into an internal node of type C0 by drawing a path $\gamma \subset \mathcal{I}(n_*)$ which contains the edge (x_*, y). Here, y is the external node of the face (x_*, n_*, y) which defined x_* as a node of type C1. Splitting n_* along γ, the tetrahedron $(n_{*,\mathrm{L}}, n_{*,\mathrm{R}}, y, x_*)$ becomes removable.

Finally, any internal node of type C0 can be made external by simply removing one tetrahedron.

Thus, to move all nodes of depth 1 to the surface we proceed in 4 steps (3 sweeps).

- **Step 1 (Sweep C2→C1)** : We promote all the x_* of type C2 in the following order: For each $n_* \in \mathcal{L}_\ell$, we promote all internal nodes of $\mathcal{I}(n_*)$ of type C2 into internal nodes of type C1. We will show that this can be done in such a way that every internal edge of the triangulation $\mathcal{I}(n_*)$ belongs to at most 1 of the splitting paths (as defined in Sect. 3.4.1).

  When this first step is complete, all internal nodes initially at depth 1 will be of type C1 or C0. There appears a new set $\mathcal{M}_\ell$ of external nodes containing the nodes of $\mathcal{L}_\ell$ which were not split and new external nodes obtained by the splitting.

- **Step 2 (Sweep C1→C0)** : We promote all the x_* of type C1 in the following order: For each $n_* \in \mathcal{M}_\ell$, we promote all promotable internal nodes of $\mathcal{I}(n_*)$ of type C1 into internal nodes of type C0.[2] We will show that this can be done in such a way that every internal edge of the triangulation $\mathcal{I}(n_*)$ belongs to at most 1 of the splitting paths (as defined in Sect. 3.4.1).

- **Step 3 (Sweep C0→external)** : Finally, we make each node of type C0 external by removing one tetrahedron.

[2] A node x_* of type C1 can be promoted to C0 only if it is connected to a node of $\mathcal{E}(n_*)$. This might not be true for all n_* for which $x_* \in \mathcal{I}(n_*)$ but for every x_* there are at least two n_* for which it is promotable.

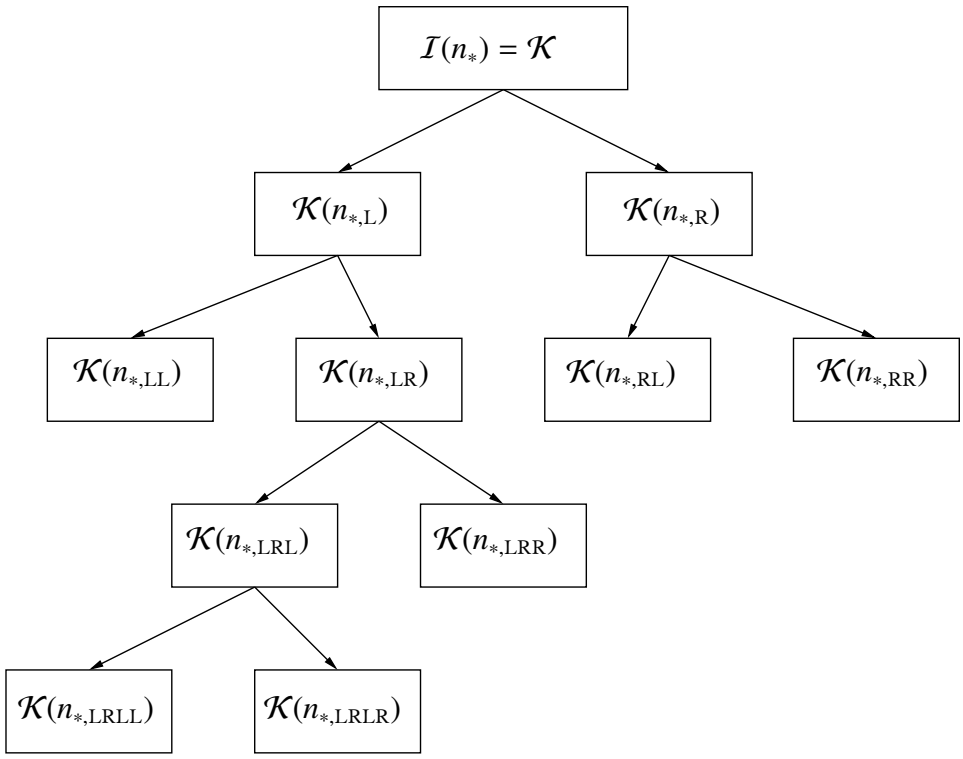

Figure 3.6: An example of a binary tree of pieces associated with the hemisphere of an external node n_* containing 6 triangles.

- **Step 4** : At this point every internal node has been moved up one level of depth.

  We have now a new triangulation $T_{\ell+1}$ of the ball, of maximal depth $d_{\max} - \ell - 1$.

Remark 3.4.8. *The reader should be aware that the 4 steps are repeated until no internal nodes remain. However, these steps are not independent, and the proof of our bound does depend on the precise definition of $\mathcal{L}_\ell$.*

Since $d_{\max} < \infty$, the procedure will end after a finite number of iterations of C2→C1→C0. We number these steps from $\ell = 0$ to $\ell = d_{\max} - 1$.

Reducing C2-nodes to C1-nodes

Given an external node $n_* \in \mathcal{L}_\ell$, we now describe in detail the algorithm which promotes the internal nodes of type C2 in $\mathcal{I} = \mathcal{I}(n_*)$ to type C1. This is achieved by a succession of carefully chosen moves of type split-a-node-along-a-path.

Each of these cuts produces a "left" and a "right" piece, which are then cut again into left and right pieces, until only triangles remain. The pieces will be called $\mathcal{K}_{\mathbf{S}} = \mathcal{K}(n_{*,\mathbf{S}})$, where $\mathbf{S}$ is a sequence of letters L and R which designate the successive choices of left and right. They are all admissible.

Thus, we construct a binary tree of pieces (see Fig. 3.6). In detail:

1. Label the nodes of $\partial\mathcal{I}$ from -1 to $-|\partial\mathcal{I}|$.

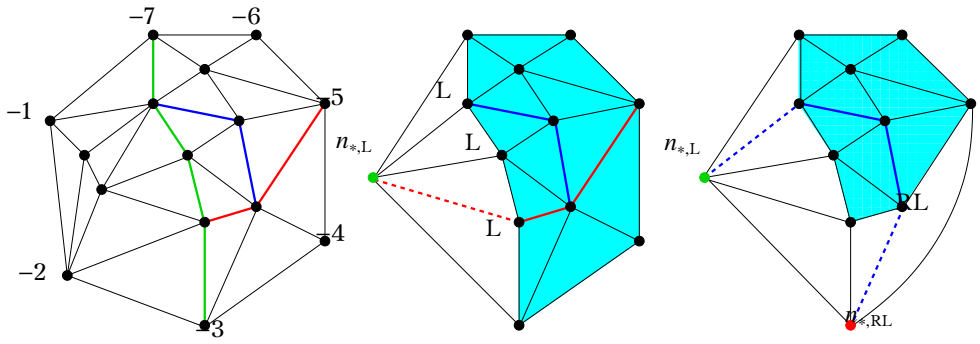

Figure 3.7: The left panel shows the internal flower $\mathcal{I}(n_*)$ of n_*. We split it in succession along the green, red and blue paths. We first split n_* into $n_{*,\mathrm{R}}$ and $n_{*,\mathrm{L}}$ along the green path. The middle panel shows $\mathcal{I}(n_{*,\mathrm{R}})$ and the green node is $n_{*,\mathrm{L}}$. The shaded region is $\mathcal{K}_{\mathrm{R}}$ and the new labels are L. One end of the red path has a label which is a negative integer. while the other has the label L and must therefore be connected to $n_{*,\mathrm{L}}$. We obtain the cutting path γ_{R}, and after the cut, we obtain two pieces $\mathcal{K}_{\mathrm{RR}}$ and $\mathcal{K}_{\mathrm{RL}}$. In the third panel we show the hemisphere of $n_{*,\mathrm{RR}}$. The blue path has labels L and RL at its extremities, which must therefore be connected to $n_{*,\mathrm{L}}$ and $n_{*,\mathrm{RL}}$. This defines the cutting path γ_{RR}. Note that it always suffices to add at most 2 dashed segments.

2. The hemisphere $\mathcal{I}$ is an admissible triangulation in the sense of Definition 3.3.3. Proposition 3.3.6 implies the existence of a shortest path γ which connects two nodes of $\partial\mathcal{I}$ (with different labels). We choose this path γ.

3. After splitting along this path, $\mathcal{I}$ is divided in two pieces, as shown in Fig. 3.5. The two pieces are called $\mathcal{K}_{\mathrm{L}}$ and $\mathcal{K}_{\mathrm{R}}$. The splitting has replaced n_* by $n_{*,\mathrm{L}}$ and $n_{*,\mathrm{R}}$ and $\mathcal{I}(n_{*,\mathrm{L}})$ is actually just $\mathcal{K}_{\mathrm{L}}$ with the cone between $n_{*,\mathrm{R}}$ and γ added. This also means that $n_{*,\mathrm{R}}$ is in the external flower $\mathcal{E}(n_{*,\mathrm{L}})$ of $n_{*,\mathrm{L}}$. Analogous terminology is used for the other half. At this point, $\mathbf{S}$ is equal to L or R, and we continue with $\mathbf{S} = \mathrm{L}$ (and do later $\mathbf{S} = \mathrm{R}$).

4. If $\mathcal{K}_{\mathbf{S}}$ is a triangle, we are done (for this branch of the tree).

5. Label all nodes on $\partial\mathcal{K}_{\mathbf{S}}$ which had no label with the label $\hat{\mathbf{S}}$, where $\hat{\mathbf{S}}$ is obtained from $\mathbf{S}$ by exchanging the last letter, cf. Fig. 3.7. In this way, the newly labeled nodes are connected to $n_{*,\hat{\mathbf{S}}}$ in $\mathcal{I}(n_{*,\mathbf{S}})$.

6. Considering $\mathcal{K}_{\mathbf{S}}$, Proposition 3.3.6 implies the existence of a new shortest path $\tilde{\gamma}_{\mathbf{S}}$ which connects two nodes of $\partial\mathcal{K}_{\mathbf{S}}$ with different labels.

7. We extend the path $\tilde{\gamma}_{\mathbf{S}}$ as follows:

   - If the end of $\tilde{\gamma}_{\mathbf{S}}$ has a negative label, we do nothing.
   - If the label of the end of $\tilde{\gamma}_{\mathbf{S}}$ is some sequence $\mathbf{S}'$ and if the node $n_{*,\mathbf{S}'}$ still exists, *i.e.*, has not been split yet, we connect the end to $n_{*,\mathbf{S}'}$ by one edge.

- If the label of the end of $\tilde{\gamma}_{\mathbf{S}}$ is some sequence $\mathbf{S}'$ and if the node $n_{*,\mathbf{S}'}$ was previously split, then at least 1 of its children has replaced it in the flower of $n_{*,\mathbf{S}}$ (see Fig. 3.8). In this case, we connect the end to any one of the corresponding children.

Doing this for both ends we obtain a path $\gamma_{\mathbf{S}}$.

8. Perform a split-a-node-along-a-path on $\gamma_{\mathbf{S}}$ and continue with step 4 for the pieces $\mathcal{K}_{\mathbf{SL}}$ and $\mathcal{K}_{\mathbf{SR}}$.

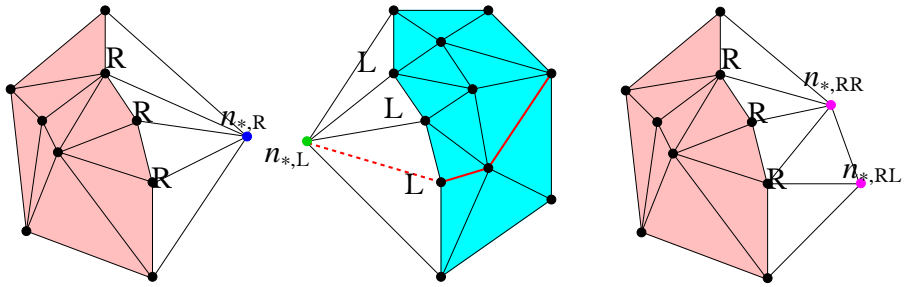

Figure 3.8: The first 2 panels show the hemispheres of $n_{*,\mathrm{L}}$ and $n_{*,\mathrm{R}}$ obtained by splitting the node n_* of Fig. 3.7 along the green path. Note that $n_{*,\mathrm{R}}$ is in $\mathcal{E}(n_{*,\mathrm{L}})$ and $n_{*,\mathrm{L}}$ is in $\mathcal{E}(n_{*,\mathrm{R}})$. The third panel shows what happens to $\mathcal{I}(n_{*,\mathrm{L}})$ if we split $n_{*,\mathrm{R}}$ along the red path of the second panel into $n_{*,\mathrm{RR}}$ and $n_{*,\mathrm{RL}}$. Note that at least 1 child of $n_{*,\mathrm{R}}$ (in this case both of them) is still in $\mathcal{E}(n_{*,\mathrm{L}})$.

Remark 3.4.9. *The boundary of a hemisphere $\mathcal{I}_{\mathbf{S}} = \mathcal{I}(n_{*,\mathbf{S}})$ is composed of 2 types of nodes:*

1. *Nodes with a negative label which are part of the original boundary $\mathcal{E}$.*
2. *The children $n_{*,\mathbf{S}'}$ of the original node n_*, where $\mathbf{S}'$ is a sequence of R's and L's.*

The boundary of a piece $\mathcal{K}_{\mathbf{S}}$, which is a sub-triangulation of $\mathcal{I}_{\mathbf{S}}$, is also composed of 2 types of nodes:

1. *Nodes with a negative label which are part of the original boundary $\mathcal{E}$, and therefore they are part of the boundary of the hemisphere $\mathcal{I}_{\mathbf{S}}$ as well.*
2. *The other nodes whose label is some sequence S'. These nodes satisfy the following two conditions:*
   (a) *All nodes of $\partial\mathcal{K}_{\mathbf{S}}$ with the same label form a connected arc of $\partial\mathcal{K}_{\mathbf{S}}$.*
   (b) *If a node $y \in \partial\mathcal{K}_{\mathbf{S}}$ has the label $\mathbf{S}'$, then y is an internal node of the hemisphere $\mathcal{I}_{\mathbf{S}}$ seen as a 2d triangulation. Furthermore, the node $n_{*,\mathbf{S}'}$ (or at least 1 of its children if $n_{*,\mathbf{S}'}$ was previously split, (see Fig. 3.8)) is in $\partial\mathcal{I}_{\mathbf{S}}$ and the edge connecting y and $n_{*,\mathbf{S}'}$ (or its corresponding child) is an internal edge of the triangulation $\mathcal{I}_{\mathbf{S}}$.*

Theorem 3.4.10. *The algorithm promotes all of the internal nodes of $\mathcal{I}$ of type C2 (and depth 1) into nodes of type C1. Furthermore, every edge in $\mathcal{I} \setminus \partial\mathcal{I}$ is in at most one path. No new nodes of type C2 at depth 1 are created.*

Proof. We first check that the different steps of the algorithm can be performed. The steps 1–3 follow from the definition of split-a-node-along-a-path. Steps 4 and 5 need no verification. Step 6 relies on Proposition 3.3.6, which implies the existence of a (shortest) path $\tilde{\gamma}_{\mathbf{S}}$, cutting the admissible piece $\mathcal{K}_{\mathbf{S}}$ into two admissible pieces $\mathcal{K}_{\mathbf{SL}}$ and $\mathcal{K}_{\mathbf{SR}}$.

In step 7, we need to make sure that the path $\gamma_{\mathbf{S}}$ connects two *different* nodes of $\mathcal{E}(n_{*,\mathbf{S}})$ which is also $\partial\mathcal{I}(n_{*,\mathbf{S}})$, to be distinguished from $\partial\mathcal{K}(n_{*,\mathbf{S}})$. The whole construction of labels has been done with this aim in mind. Note that if a node u has a negative label, we do nothing because any node u with a negative label is part of the original boundary $\mathcal{E}(n_*)$, implying that if $u \in \mathcal{I}(n_{*,\mathbf{S}})$, then $u \in \mathcal{E}(n_{*,\mathbf{S}})$ for any child $n_{*,\mathbf{S}}$ of n_*. On the other hand, if the label is the sequence $\mathbf{S}'$, then by construction (step 5), u is connected to $n_{*,\mathbf{S}'}$ (or to one of its children) with one edge. Since the labels are different by construction, the path γ is a splitting path, and therefore a cut along it is possible. In step 8, we need to verify that the cut can indeed be done, and that the algorithm can be applied to the children of the $\mathcal{K}$ which was just cut. But this is the content of Proposition 3.3.6, which shows that the cut can be done in such a way that the children are admissible in the sense of Definition 3.3.3.

Since new paths are always constructed in the interior of $\mathcal{K}$, and the $\mathcal{K}$'s are cut along them, it is obvious that no edge (of the original hemisphere $\mathcal{I}(n_*)$) is covered by more than one path.

Finally, to finish the proof, we note that every node of $\mathcal{I} \setminus \partial\mathcal{I}$ belongs to at least 1 path $\tilde{\gamma}_{\mathbf{S}}$, since the only pieces remaining at the end of the algorithm are triangles. In particular, this implies that every internal node of type C2 in $\mathcal{I}$ is promoted to a node of type C1. □

Reducing C1-nodes to C0-nodes

Let T be a triangulation of a ball. Consider an external node n_* of T and let $\mathcal{I} = \mathcal{I}(n_*)$ be its internal hemisphere. Furthermore, assume that all nodes of $\mathcal{I}$ are either external (with regard to T) or internal of type C0 or C1 but not C2. We now describe the algorithm which promotes all the internal nodes of type C1 of $\mathcal{I}$ to internal nodes of type C0. The approach is somewhat different from that of the previous section. Indeed promoting an internal node x of type C2 to an internal node of type C1 is done by splitting some external node n_* along a path going through x. However, let $x \in \mathcal{I}(n_*)$ be an internal node of type C1 and let (x, y, n_*) be an internal face which defines x as C1; by hypothesis, $y \in \partial\mathcal{I}$. Promoting x to an internal node of type C0 is done by splitting n_* along a path which contains the edge (y, x). So we have to make sure such a path exists.

For every internal node x of type C1 in $\mathcal{I}(n_*)$ we choose one of the $y \in \partial\mathcal{I}$ for which (x, y, n_*) is an internal face and call it $y(x)$. We define

$$\mathcal{Y} = \{(x, y(x)) \mid x \text{ is C1 } \} .$$

We will eliminate elements in the list $\mathcal{Y}$ by iterating an algorithm similar to the one in the previous section, until none are left. A binary tree of left and right pieces will be formed in the process (see Fig. 3.6).

At the first step of this algorithm, this tree only contains one element, namely the hemisphere $\mathcal{I}$. We will form a tree of $\mathcal{K}$'s as before, starting at $\mathcal{K} = \mathcal{I}$.

The algorithm starts with steps 1 and 2 below, and then repeats the other steps until it stops.

1. Pick an edge $(x, y) = (x, y(x)) \in \mathcal{Y}$.

2. By hypothesis, $y \in \partial\mathcal{I}(n_*)$. By Lemma 3.3.2 there is a second, disjoint, simple path connecting x to a node $z \in \partial\mathcal{I}(n_*)$, $z \neq y$. This defines a splitting path γ connecting 2 distinct nodes y and z of $\partial\mathcal{I}(n_*)$. Similarly to the previous section, we split n_* along γ into $n_{*,\mathrm{R}}$ and $n_{*,\mathrm{L}}$. We add the 2 new pieces $\mathcal{K}(n_{*,\mathrm{R}})$ and $\mathcal{K}(n_{*,\mathrm{L}})$ as two leaves of $\mathcal{K}$ in the tree. We remove the edge (x, y) from the list $\mathcal{Y}$. Note that the path γ might promote a second internal node x' of type C1 into a node of type C0, if the edge (x', z) is in the list $\mathcal{Y}$ and in the path γ. In that case, both edges (x, y) and (x', z) are removed from $\mathcal{Y}$.

3. If the list $\mathcal{Y}$ is empty, we are done.

4. Pick an edge $(x, y) \in \mathcal{Y}$.

5. Find the piece $\mathcal{K}(n_{*,s_1,\ldots,s_k})$, where $s_i \in \{\mathrm{L}, \mathrm{R}\}$, among the leaves of the binary tree which contains the edge (x, y). We use the abbreviations $\mathbf{S} = s_1, \ldots, s_k$ and $n_{*,\mathbf{S}}$. The edge (x, y) belongs to exactly one piece.[3]

6. Observe that the node y is in $\partial\mathcal{I}(n_{*,\mathbf{S}}) \cap \partial\mathcal{I}(n_*)$.[4] The edge (x, y) gives us the first simple path connecting x to $\partial\mathcal{I}(n_{*,\mathbf{S}})$ since by construction, it is an internal edge of $\mathcal{K}(n_{*,s_1,\ldots,s_k})$. We still need to find the other part of the splitting path $\gamma_{\mathbf{S}}$:
   • If x is in the interior of $\mathcal{K}(n_{*,\mathbf{S}})$, by Lemma 3.3.2 there is a second independent path connecting x to a node $z \in \partial\mathcal{K}(n_{*,\mathbf{S}})$, $z \neq y$.

   If z is also in $\partial\mathcal{I}(n_{*,\mathbf{S}})$ we have found a $\gamma_{\mathbf{S}}$ along which we can cut. Note that in this case, the path $\gamma_{\mathbf{S}}$ might promote a second node x' of type C1; this happens if $z \in \partial\mathcal{I}(n_*)$ and (x', z) is an edge of $\gamma_{\mathbf{S}}$.

   If $z \notin \partial\mathcal{I}(n_{*,\mathbf{S}})$, the path $\gamma_{\mathbf{S}}$ is obtained by adding the edge which connects z to the tip of the cone.[5]
   • If x is not in the interior of $\mathcal{K}(n_{*,\mathbf{S}})$, $\gamma_{\mathbf{S}}$ is found by connecting x to a tip of one of the cones attached to $\mathcal{K}(n_{*,\mathbf{S}})$[6] (see Footnote 5).

7. We split $n_{*,\mathbf{S}}$ along the path $\gamma_{\mathbf{S}}$ and add the 2 new pieces $\mathcal{K}(n_{*,\mathbf{S}\mathrm{R}})$ and $\mathcal{K}(n_{*,\mathbf{S}\mathrm{L}})$ to the tree as leaves of $\mathcal{K}(n_{*,\mathbf{S}})$. Note that $\mathcal{K}(n_{*,\mathbf{S}})$ is no longer a leaf of the tree and will never

[3]Note that the only edges which are common to more than one piece are the edges of the paths along which we already cut. Since (x, y) is still in the list $\mathcal{Y}$, it cannot be such an edge.

[4]By hypothesis, $y \in \partial\mathcal{I}(n_*)$ and therefore also $y \in \partial\mathcal{I}(n_{*,\mathbf{S}}) \cap \partial\mathcal{I}(n_*)$.

[5]The distance between $\partial\mathcal{I}(n_{*,\mathbf{S}})$ and any node in $\partial\mathcal{K}(n_{*,\mathbf{S}})$ is at most 1, see Fig. 3.7. The node z belongs to a path $\gamma_{\mathbf{S}'}$ along which we already cut. This implies that z is connected to $n_{*,\mathbf{S}'\mathrm{L}}$ or $n_{*,\mathbf{S}'\mathrm{R}}$, called the *tip of the cone* associated with z.

[6]Note that the node y is *not* on a tip of a cone but is on the original boundary $\partial\mathcal{I}(n_*)$, guaranteeing that $\gamma_{\mathbf{S}}$ is not a closed loop.

be encountered in the remaining steps of the algorithm. Finally, we remove the edge (x, y) (and eventually (x', z) if x' is also promoted by γ_S) from the list $\mathcal{Y}$.

8. We continue with step 3.

The algorithm stops when all internal nodes of type C1 of $\mathcal{I}(n_*)$ have been promoted to C0. Since each branch of the tree is used at most once and since we never cut along the boundary of any $\mathcal{K}_S$ we have shown:

Theorem 3.4.11. *The algorithm promotes all of the internal nodes of $\mathcal{I}$ of type C1 (and depth 1) into nodes of type C0. Furthermore, every edge in $\mathcal{I} \setminus \partial\mathcal{I}$ is in at most one path, and no new nodes of type C1 or C2 are created.*

Change of the f-vector by removing all internal nodes

We now bound the change of the f-vector which results from transforming a triangulation of maximal depth $d_{\max}$ to one with maximal depth $d_{\max} - 1$, until no internal nodes remain.

Unfortunately, the f-vector alone is not good enough for efficient bounds, since new edges will appear in the construction, and we need to keep track not only on the total number of internal edges as the procedure continues, but also how many there are on each (current) depth. By definition an edge is either on one depth or connects two adjacent depths, and a face also connects at most 2 depths.

Given the original triangulation T, the bookkeeping will be done by associating with each node x_* its *original depth* $d(x_*)$.

As we are going to split nodes, we also define $d(x_{*,\mathrm{R}}) = d(x_{*,\mathrm{L}}) = d(x_*)$, and similarly for all further splittings.

Definition 3.4.12. *For each $d : 0 \le d \le d_{\max}$ we set:*

- *a_d as the number internal edges (x, y) with $d(x) = d$ and $d(y) = D(y) = d + 1$.*
- *b_d as the number of internal edges (x, y) with $d(x) = d$ and $d(y) = d$.*
- *f'_d is the number of internal faces (x, y, z) with $d(x) = d$ and $d(y) = d + 1$ (this implies $d(z) \in \{d, d + 1\}$).*

We also say that $a_{-1} = a_{d_{\max}} = 0$ and $f'_{-1} = f'_{d_{\max}} = 0$. Note that all these constants will not change during the iterations since they are counters of the initial triangulation.

As every node is connected to nodes of the same depth or to depths differing by at most 1, the following obvious relations hold:

$$\sum_d (a_d + b_d) = e \ , \quad \sum_d f'_d \le f_{\mathrm{i}} \ , \tag{3.3}$$

where e is the number of internal edges, and f_{i} is the number of internal faces.

Moving nodes to the surface causes the creation of new edges and faces. Our study is based on a careful bound of this growth in terms of the counters of the initial triangulation introduced in Definition 3.4.12. Let Δ_ℓ denote the increase of the number of internal edges obtained by performing the steps C2→C1→C0→external at iteration ℓ.

Proposition 3.4.13. *There is a constant C' such that*

$$\begin{aligned} \Delta_\ell &\leq C' \ (a_\ell + a_{\ell-1} + b_\ell + f'_\ell) \ , \text{ for } \ell > 0 \ , \\ \Delta_0 &\leq C' \ (a_0 + b_0 + n_{\mathrm{s}}) \ . \end{aligned} \tag{3.4}$$

Corollary 3.4.14. *Eliminating all internal nodes of a triangulation T with f-vector $\langle t, f_{\mathrm{s}}, n_{\mathrm{i}} \rangle$ leads to a total increase Δ of internal edges which is bounded by*

$$\Delta \leq C \ (t + n_{\mathrm{i}}) \ .$$

Proof of the corollary. From Eq. (3.1) we deduce $f_{\mathrm{i}} = 2t - f_s/2$ and $e = t + n_{\mathrm{i}} - f_{\mathrm{s}}/2 + 1$. Also, $n_{\mathrm{s}} = f_{\mathrm{s}}/2 + 2$. Using Eq. (3.3) and the proposition, we get

$$\Delta = \sum_{\ell \geq 0} \Delta_\ell \leq C' \ (2e + f_{\mathrm{i}} + n_{\mathrm{s}}) \ ,$$

from which the assertion follows (the coefficient of f_{s} is negative and the additive constants can be bounded since $1 \leq t$). □

The proof of Proposition 3.4.13 will take up most of this subsection. We proceed as follows:

- Bound Δ_ℓ in terms of the number of edges in the hemispheres of nodes which are split (Lemma 3.4.15).
- Two terms appear: the first for the internal edges in the hemispheres, the other for the external edges (Lemma 3.4.17).
- Bound the first term with the counters a_ℓ, b_ℓ, f'_ℓ (Lemma 3.4.17 and Lemma 3.4.18).
- Bound the second term with the first term and the number of external nodes n_{s} (Lemma 3.4.19).

We decompose $\Delta_\ell = \Delta_{\ell,0} + \Delta_{\ell,1} + \Delta_{\ell,2}$, where the first term comes from the sweep C2→C1 at iteration ℓ, and the second from C1→C0. The third term coming from the sweep C0→external only decreases internal nodes and edges, so $\Delta_{\ell,2} \leq 0$ and we do not take that into account.

We define $\mathcal{L}_\ell = \{n : d(n) = \ell\}$, and, after having performed C2→C1 we define $\mathcal{M}_\ell$ as the set of those nodes of $\mathcal{L}_\ell$ which have not been split, as well as the children of those which have been split.

Lemma 3.4.15. *We have the bounds*

$$\begin{aligned} \Delta_{\ell,0} \ &\leq \ 2 \sum_{n_* \in \mathcal{L}_\ell} \#\left(\text{internal edges in } \mathcal{I}(n_*)\right) \ , \\ \Delta_{\ell,1} \ &\leq \ 2 \sum_{n_* \in \mathcal{M}_\ell} \#\left(\text{internal edges in } \mathcal{I}(n_*)\right) \ . \end{aligned}$$

Proof. Cutting along a path γ adds $|\gamma|-1$ internal edges to the triangulation (see Lemma 3.4.7).

The extension of the path $\tilde{\gamma}_S$ to a splitting path γ_S adds at most 2 to its length (see step 7 for the case C2→ C1, and step 6 for the case C1→C0). Therefore, $|\gamma_S| - 1 \le 2|\tilde{\gamma}_S|$, since $|\tilde{\gamma}_S| \ge 1$.

For a given node n_*, all the paths $\tilde{\gamma}_S$ are drawn in its hemisphere $\mathcal{I}(n_*)$ and each edge of this hemisphere is used in at most 1 path $\tilde{\gamma}_S$ (see Theorem 3.4.10 and 3.4.11) so that for a given n_*, we have

$$\sum_S |\tilde{\gamma}_S| \le \#(\text{internal edges in } \mathcal{I}(n_*)) \ .$$

Summing over all splittable nodes yields the claim. (This is the crucial bound, which has become possible through our careful cutting procedures.) □

We next bound the number of internal edges in the $\mathcal{I}(n_*)$.

Definition 3.4.16. *We define the numbers $a_{\ell,i}$ depending on whether we are before the C2→C1 sweep ($i = 0$) or after it but before the C1→C0 sweep ($i = 1$), and after that sweep ($i = 2$), all at iteration ℓ. The numbers $a_{\ell,i}$ are defined as the number of edges (x, y) with $d(x) = \ell$ and $d(y) = \ell + 1$, at iteration ℓ and before the sweep determined by i. Analogously, we define $b_{\ell,i}$ and $f'_{\ell,i}$.*

We also let $\hat{a}_{\ell,0}$ be the number of edges at the beginning of iteration ℓ with $d(x) = \ell - 1$ and $d(y) = \ell$. (This is not the same as $a_{\ell-1,0}$, which is defined for iteration $\ell - 1$.) Also, $\hat{a}_{\ell,1}$ is defined after sweep C2→C1.

Note that these numbers change with i, and depend on ℓ, since new edges are being added.

Lemma 3.4.17.

$$\Delta_{\ell,0} \ \le \ 6a_{\ell,0} + 12b_{\ell,0} + 6\hat{a}_{\ell,0} + 2\sum_{n_* \in \mathcal{L}_\ell} (|\mathcal{E}(n_*)| - 3) \ , \tag{3.5a}$$

$$\Delta_{\ell,1} \ \le \ 6a_{\ell,1} + 12b_{\ell,1} + 6\hat{a}_{\ell,1} + 2\sum_{n_* \in \mathcal{M}_\ell} (|\mathcal{E}(n_*)| - 3) \ . \tag{3.5b}$$

Proof. Starting from the relations of Lemma 3.4.15, we use Lemma 3.3.1 for the hemisphere of n_* which is a polygon with $p = |\mathcal{E}(n_*)|$ sides. Summing over $\mathcal{L}_\ell$ or $\mathcal{M}_\ell$, and since the only nodes x present in the hemisphere $\mathcal{I}(n_*)$ of a node n_* with $d(n_*) = \ell$ are such that $d(x) \in \{\ell - 1, \ell, \ell + 1\}$, the assertion follows. The factor $12 = 2 \cdot 6$ takes into account that the edge (n_*, m_*) can appear both in $\mathcal{I}(n_*)$ and in $\mathcal{I}(m_*)$ (if n_* and m_* are in $\mathcal{L}_\ell$ or $\mathcal{M}_\ell$. □

So there remains to bound the terms on the r.h.s. of Eq. (3.5) in terms of the counters of the initial triangulation. This is done in the next lemmas.

The effect of the sweeps C2→C1 →C0 at iteration ℓ is summarized by

Lemma 3.4.18. *One has, for each $\ell \ge 0$:*

$$a_{\ell,0} = a_\ell \ , \quad b_{\ell,0} = b_\ell \ , \quad f'_{\ell,0} = f'_\ell \ , \tag{3.6a}$$

$$a_{\ell,1} \le a_{\ell,0} + 2f'_{\ell,0} \ , \quad a_{\ell,2} \le a_{\ell,1} + 2f'_{\ell,1} \ , \tag{3.6b}$$

$$f'_{\ell,1} \le 7 f'_{\ell,0} \ , \tag{3.6c}$$
$$\hat{a}_{\ell,0} \le a_{\ell-1} + 16 f'_{\ell-1} \ . \tag{3.6d}$$

Proof of Lemma 3.4.18. Proof of Eq. (3.6a): Let $\ell' < \ell$. During the iteration $\ell' < \ell$, we split nodes n_* of $\mathcal{L}_{\ell'}$ and $\mathcal{M}_{\ell'}$ into $n_{*,\mathbf{S}}$ with $d(n_{*,\mathbf{S}}) = \ell'$. This implies that every internal edge added during the iteration ℓ' has an end $n_{*,\mathbf{S}}$ at $d(n_{*,\mathbf{S}}) = \ell' < \ell$. But $a_{\ell,0}$ and $b_{\ell,0}$ count internal edges with both ends at depth $d(x) \ge \ell > \ell'$. The same argument also holds for $f'_{\ell,0}$. This means that at the beginning of the iteration ℓ, the values of the counters $a_{\ell,0}$, $b_{\ell,0}$ and $f'_{\ell,0}$ are still identical to the constants of the original triangulation.

Proof of Eq. (3.6b): We need to bound the added number of internal edges $(n_{*,\mathbf{S}}, y)$ in the sweep C2→C1 with $n_* \in \mathcal{L}_\ell$, $y \in \mathcal{I}(n_*)$, and $d(y) = \ell + 1$. By construction, this number is bounded by the number of paths $\gamma_{\mathbf{S}}$ which go through such a node y in the 2d triangulation $\mathcal{I}(n_*)$. Furthermore, by Theorem 3.4.10, each edge of $\mathcal{I}(n_*)$ is used in at most one path $\gamma_{\mathbf{S}}$. We deduce that, for two such nodes n_* and y, the number of added internal edges of type $(n_{*,\mathbf{S}}, y)$ is bounded by the degree of the edge (n_*, y) (the number of faces containing the edge). Summing these degrees for all such edges (n_*, y) is bounded by $2f'_{\ell,0}$. The same argument proves the second relation.

Proof of Eq. (3.6c): To prove this relation, we need to bound the added number of internal faces $(n_{*,\mathbf{S}}, y, z)$ in the sweep C2→C1 at iteration ℓ when $n_{*,\mathbf{S}}$ is obtained from splitting some $n_* \in \mathcal{L}_\ell$ and y is such that $d(y) = \ell + 1$. But each added internal face $(n_{*,\mathbf{S}}, y, z)$ requires the addition of the internal edge $(n_{*,\mathbf{S}}, y)$. Furthermore, by definition of the move split-a-node-along-a-path, each new internal edge is added along with three internal faces. We deduce from Eq. (3.6b) that $f'_{\ell,1} - f'_{\ell,0} \le 3(a_{\ell,1} - a_{\ell,0}) \le 6f'_{\ell,0}$.

Proof of Eq. (3.6d): The reason this proof is tricky is that during the previous iteration $\ell - 1$, new internal edges satisfying $d(x) = \ell - 1$ and $d(y) = \ell$ have been added: the proof of this relation at iteration ℓ therefore involves the other relations of Lemma 3.4.18 at iteration $\ell - 1$.

When $\ell = 0$, Eq. (3.6d) obviously holds since both sides are 0. So we now assume $\ell > 0$ and that the conclusion of Lemma 3.4.18 holds for $\ell - 1$. We know that the sweep C0→external does not add any internal edge. From this we deduce that $\hat{a}_{\ell,0} \le a_{\ell-1,2}$. The inequality Eq. (3.6d) then follows from the relations of Lemma 3.4.18 at iteration $\ell - 1$ we have already proved:

$$\begin{aligned} \hat{a}_{\ell,0} \le a_{\ell-1,2} &\le a_{\ell-1,1} + 2f'_{\ell-1,1} \\ &\le a_{\ell-1,0} + 2f'_{\ell-1,0} + 14f'_{\ell-1,0} \\ &\le a_{\ell-1} + 16 f'_{\ell-1} \ , \end{aligned}$$

which is the bound we seek.

□

We finally need to discuss the terms $|\mathcal{E}(n_*)| - 3$ of Lemma 3.4.17.

Lemma 3.4.19. *One has for each $\ell \ge 0$:*

$$\sum_{n_* \in \mathcal{L}_\ell} (|\mathcal{E}(n_*)| - 3) \le \delta_{\ell,0} \cdot 3n_{\mathrm{s}} \ ,$$

$$\sum_{n_* \in \mathcal{M}_\ell} (|\mathcal{E}(n_*)| - 3) \leq \delta_{\ell,0} \cdot 3n_{\mathrm{s}} + \Delta_{\ell,0} \ .$$

Proof. The external degree of n_* is always 3 for those nodes which have been promoted to the surface by removing a tetrahedron (those which were at the surface at level $\ell = 0$ can of course have higher degree). Since the only way to lower the depth of a node is by removing a tetrahedron, we get

$$\begin{aligned} \sum_{n_* \in \mathcal{L}_\ell} (|\mathcal{E}(n_*)| - 3) &= \delta_{\ell,0} \sum_{n_* \in \mathcal{L}_0} (|\mathcal{E}(n_*)| - 3) \\ &= \delta_{\ell,0} \cdot (2e_{\mathrm{s}} - 3n_{\mathrm{s}}) \\ &\leq \delta_{\ell,0} \cdot 3n_{\mathrm{s}} \ , \end{aligned}$$

since the set $\mathcal{L}_0$ is the set of all external nodes.

The sum over $\mathcal{M}_\ell$ is more delicate, since the external degree of a node $n_* \in \mathcal{M}_\ell$ can be larger than 3. However, if we split a node $n_{*,\mathbf{S}}$ into $n_{*,\mathrm{SR}}$ and $n_{*,\mathrm{SL}}$, then the external degrees satisfy

$$(|\mathcal{E}(n_{*,\mathrm{SL}})| - 3) + (|\mathcal{E}(n_{*,\mathrm{SR}})| - 3) = (|\mathcal{E}(n_{*,\mathbf{S}})| - 3) + 1 \ . \tag{3.7}$$

We know that nodes of $\mathcal{M}_\ell$ are the children of nodes in $\mathcal{L}_\ell$. Each node n_* of $\mathcal{L}_\ell$ is split along a binary tree into a set $\{n_{*,\mathbf{S}}\}_{\mathbf{S}}$ (some nodes of $\mathcal{M}_\ell$ are nodes of $\mathcal{L}_\ell$ which were not split; in this case the binary tree has no vertices). What Eq. (3.7) means is that each vertex of this binary tree adds 1 to the sum $\sum_{n_* \in \mathcal{M}_\ell} (|\mathcal{E}(n_*)| - 3)$. Therefore,

$$\sum_{n_* \in \mathcal{M}_\ell} (|\mathcal{E}(n_*)| - 3) \leq \sum_{n_* \in \mathcal{L}_\ell} (|\mathcal{E}(n_*)| - 3) + \#\,(\text{splits in C2}\to\text{C1}) \ .$$

Since each split adds at least one internal edge, we have a bound

$$\sum_{n_* \in \mathcal{M}_\ell} (|\mathcal{E}(n_*)| - 3) \leq \sum_{n_* \in \mathcal{L}_\ell} (|\mathcal{E}(n_*)| - 3) + \Delta_{\ell,0} \ .$$

□

Proof of Proposition 3.4.13. We start from Lemma 3.4.17. The internal edges at the current point of the construction are all the edges with a node with $d(\cdot) = \ell$. These come in three types: Those connecting $(\ell, \ell+1)$ (the $a_{\ell,0}$) those connecting (ℓ, ℓ) (the $b_{\ell,0}$), and those between $(\ell, \ell-1)$ (the $\hat{a}_\ell$). Note that the last of these quantities is *not* equal to $a_{\ell-1}$ because the number of edges $(\ell, \ell-1)$ might have changed in iteration $\ell - 1$. Lemma 3.4.18 provides, however, the necessary bounds.

Using Eq. (3.6a), Eq. (3.6d), and Lemma 3.4.19, we immediately get the bound on $\Delta_{\ell,0}$ we seek:

$$\begin{aligned} \Delta_{\ell,0} &\leq 6a_\ell + 12b_\ell + 6(a_{\ell-1} + 16f'_{\ell-1}) + 6n_{\mathrm{s}} \cdot \delta_{\ell,0} \\ &\leq 96(a_\ell + b_\ell + a_{\ell-1} + f'_{\ell-1} + n_{\mathrm{s}} \cdot \delta_{\ell,0}) \ . \end{aligned} \tag{3.8}$$

All internal edges added during the sweep C2→C1 at iteration ℓ have by construction an end x with $d(x) = \ell$. Therefore, the total number of internal edges added during the sweep C2→C1 at iteration ℓ, $\Delta_{\ell,0}$, is also given by

$$(a_{\ell,1} + b_{\ell,1} + \hat{a}_{\ell,1}) - (a_{\ell,0} + b_{\ell,0} + \hat{a}_{\ell,0}) = \Delta_{\ell,0} \ .$$

Using this relation and Lemma 3.4.19, starting from Lemma 3.4.17, we have:

$$\begin{aligned}\Delta_{\ell,1} &\leq 12(a_{\ell,1} + b_{\ell,1} + \hat{a}_{\ell,1}) + \delta_{\ell,0} \cdot 6n_{\mathrm{s}} + 2\Delta_{\ell,0} \\ &\leq 12(a_{\ell,0} + b_{\ell,0} + \hat{a}_{\ell,0} + \Delta_{\ell,0}) + \delta_{\ell,0} \cdot 6n_{\mathrm{s}} + 2\Delta_{\ell,0} \ .\end{aligned}$$

Using Lemma 3.4.18, we get:

$$\Delta_\ell = \Delta_{\ell,0} + \Delta_{\ell,1} \leq 12(a_\ell + b_\ell + a_{\ell-1} + 16f'_{\ell-1}) + 15\Delta_{\ell,0} + 6n_{\mathrm{s}} \cdot \delta_{\ell,0} \ .$$

Replacing $\Delta_{\ell,0}$ by Eq. (3.8) yields the result we seek. □

3.4.4 Reducing a triangulation with no internal nodes into a set of nuclei

Let T be any triangulation. In the previous section, we described an algorithm which transforms T into a new triangulation T' with no internal nodes. We now systematically apply the moves cut-a-3-face and open-a-2-face on every internal face of T' with less than 2 internal edges. We end up with a collection of triangulations $\{N_i\}$ satisfying the following properties:

- All nodes of any such N_i are external.
- All internal faces of any such N_i have at least 2 internal edges.

Any triangulation satisfying these two conditions is called a *nucleus*.

3.5 Part II: Bounding the number of triangulations

We showed that any triangulation can be reduced into a collection of nuclei using four moves. For the moment, we proceed without using the move cut-a-3-face. This implies that any triangulation can be transformed into a "tree of nuclei" (the formal definition of a tree of nuclei will be given later on) using the three remaining moves. Equivalently, this shows that any triangulation can be constructed from a tree of nuclei, using the inverse of these three moves. Bounding the number of trees of nuclei, and then bounding the number of ways one can perform the inverse moves on such a tree yields a bound on the total number of triangulations.

3.5.1 Rooted triangulations

We define what we mean by a rooted triangulation T and we show that one can label all external nodes of T. In the sequel, we use a particular labeling described below.

Definition 3.5.1. *A* rooted triangulation (T, F) *of the 3-ball is a triangulation T with one labeled external face F. This labeled face is called the* root. *The three nodes of the root are always labeled 0, 1, and 2.*

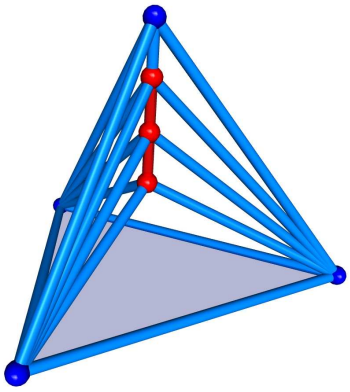

Figure 3.9: The Christmas tree with $m = 3$ internal nodes. This triangulation can be rooted in more than one way.

Remark 3.5.2. *We will only consider rooted triangulations. This means for instance that talking about the Christmas tree $T_m, m > 1$ makes no sense, since there is more than one such rooted triangulation. The exceptions are of course symmetric triangulations T such as the tetrahedron.*

Definition 3.5.3. *Consider the boundary of a rooted triangulation (T, F) and let $\mathcal{N}_s$ be the set of all external nodes. We define a particular labeling $h(\cdot) : \mathcal{N}_s \mapsto \mathbb{N} \cup \{0\}$ of all external nodes.*

The labeling is defined as follows: the root is labeled as $(0, 1, 2)$. Any labeled edge can be seen as an element $(a, b) \in \mathbb{Z}_+^2$ with $a < b$.[7] We consider the lexical order on $\mathbb{Z}_+^2$. We start with the node 0. Its external flower is a 1d triangulation of the circle S^1 and it contains the edge $(1, 2)$ by definition. This edge determines the direction in which we label all unlabeled nodes of the external flower of node 0.

Next, we consider the external flower of node 1 and we look for the smallest labeled edge in the sense of the above ordering. In this case, this edge is $(0, 2)$. This edge fixes the direction in which we label all the yet unlabeled nodes of the external flower of node number 1. Notice that every labeled node is part of a face along with 2 already labeled nodes. This implies that the external flower of any labeled node contains a smallest labeled edge and as such can be directed.

We continue with all the nodes in their natural order until all external nodes of T are labeled.

[7] We use the notation $\mathbb{Z}_+ = \{0, 1, 2, \dots\}$.

3.5.2 Trees of nuclei

Since we work with rooted triangulations, from now on, we will only use rooted nuclei, namely:

Definition 3.5.4. *A* nucleus *is a* ***rooted*** *triangulation with no internal nodes such that every internal face has at most one external edge.*

Rooted trees of nuclei and planar rooted trees

Let $\mathcal{N}$ be the set of all nuclei and $\mathcal{N}_{t,f}$ be the subset of all nuclei with t tetrahedra and f external faces.

Definition 3.5.5. *A rooted triangulation T is called a* rooted tree of nuclei *if all nodes of T are external and all internal faces of T have 0, 2, or 3 internal edges. (In other words, no internal face has 2 external edges.)*

In other words, a rooted tree of nuclei is simply a rooted triangulation which is obtained by gluing sequentially nuclei along pairs of their external faces. This is done in such a way that each nucleus is glued to an external face (a, b, c) of its parent through its root; 0 is identified with a, 1 with b and 2 with c. Once the tree is built, the external nodes are renumbered in the sense of Definition 3.5.3.

Since all external faces of a rooted triangulation are ordered, this defines a bijection between rooted trees of nuclei (T, F) and rooted planar trees with colored vertices in the following manner:

- Each nucleus of the triangulation (T, F) is represented by a colored vertex.
- The root-vertex of the planar tree represents the nucleus with the root F, *i.e.*, with the face $(0, 1, 2)$.
- Each internal face of the triangulation with three external edges is shared by two nuclei and hence it is represented in the tree by an edge connecting the corresponding two colored vertices.
- Since the internal faces with three external edges are ordered, this induces an order of the edges of the planar tree, say from left to right.

Hypothesis on the number of rooted nuclei

We next show how the question of exponential growth can be reformulated. We show that if there are not "too many" different types of nuclei, then there is indeed an exponential bound on the number of triangulations, when expressed in terms of the number of tetrahedra.

Hypothesis 3.5.6. *There is a finite constant $K_1 > 1$ such that the number $\varrho(t, f_s)$ of face-rooted nuclei with f-vector $\langle t, f_s, 0\rangle$ is bounded by K_1^t.*

In order to alleviate the notation, from now on, we will denote f_s by f.

Lemma 3.5.7. *For any nucleus $N \in \mathcal{N}_{t,f}$ one has $f \leq t + 3$.*

Proof. If N is a tetrahedron, the assertion is obvious. If N is non-trivial each tetrahedron of N can have at most 1 external face, since otherwise it would have an internal face with more than one external edge. □

The number of rooted trees of nuclei

We use the classical method for counting planar ordered trees, generalized to the case of a multitude of different nodes, which are the face-rooted nuclei.

Definition 3.5.8. *Let $A_{v,t,f}$ be the number of rooted trees of nuclei with $v > 0$ nuclei, t tetrahedra and f external faces. We define $A_{0,t,f} = \delta_{t,0}\,\delta_{f,0}$.*

Our main bound is:

Proposition 3.5.9. *Under the Hypothesis 3.5.6 there is a K_2, with $2 < K_2 < \infty$ such that for all t, f, one has*

$$\sum_v A_{v,t,f} \leq K_2^t \ .$$

Proof. Consider a tree of nuclei, and let N be the nucleus containing the root F and assume that $N \in \mathcal{N}_{t_0,f_0}$. Removing N from the tree leads to $f_0 - 1$ rooted trees of nuclei, some of which may be empty. We let v_i, t_i, and f_i denote the counters for the branch i. Note that if a branch i has 0 nuclei, *i.e.*, if $v_i = 0$, then, obviously, $t_i = f_i = 0$. Thus, we get the relations:

$$\sum_{i=1}^{f_0-1} v_i = v - 1 \ , \qquad \sum_{i=1}^{f_0-1} t_i = t - t_0 \ , \qquad \sum_{i=1}^{f_0-1} \Big(\delta_{f_i>0}(f_i - 1) + \delta_{f_i=0}\Big) = f - 1 \ . \tag{3.9}$$

In the sequel, we denote by $\sum'_{v,t,f,t_0,f_0}$ the sum over the set

$$\{v_i, t_i, f_i \ \big| \ i = 1, \ldots, f_0 - 1 \ , v_i \geq 0, t_i \geq 0, f_i \geq 0 \text{ and satisfying Eq. (3.9)}\} \ .$$

This observation allows us to write a recursive relation

$$A_{v,t,f} = \delta_{v,0}\,\delta_{t,0}\,\delta_{f,0} + \sum_{t_0>0, f_0 \geq 4} \varrho(t_0, f_0) \sum\nolimits'_{v,t,f,t_0,f_0} \prod_{i=1}^{f_0-1} A_{v_i,t_i,f_i} \ . \tag{3.10}$$

Fix $M \in \mathbb{Z}_+$, and assume that v, t, f satisfy $3v + 3t + f \leq M$. By Eq. (3.9), we deduce

$$3v_i + 3t_i + f_i \leq 3v - 3 + 3t - 3t_0 + f \leq M - 1 \ .$$

We define

$$A_M(s) = \sum_{3v+3t+f \leq M} A_{v,t,f}\, s^{3v+3t+f} \ .$$

Clearly, $A_0(s) = 1$ for all s, $A_M(0) = 1$ for all $M \geq 0$, and for a fixed s, $A_M(s)$ is an increasing sequence in M.

Multiplying Eq. (3.10) by $s^{3v+3t+f}$ and summing, we get, using Eq. (3.9):

$$\begin{aligned} A_M(s) = 1 + \sum_{3v+3t+f\le M} \sum_{t_0=1}^{t} \sum_{f_0=4}^{f} \varrho(t_0,f_0)\ s^{3+3t_0+1-\sum_{i=1}^{f_0-1}(\delta_{f_i>0}-\delta_{f_i=0})} \\ \times \sum\nolimits'_{v,t,f,t_0,f_0} \prod_{i=1}^{\ell} A_{v_i,t_i,f_i} s^{3v_i+3t_i+f_i}\ . \end{aligned} \tag{3.11}$$

Using Lemma 3.5.7, we have

$$\begin{aligned} 3+3t_0+1-\sum_{i=1}^{f_0-1}(\delta_{f_i>0}-\delta_{f_i=0}) &\ge 3+3t_0+1-(f_0-1)\cdot 1+0 \\ &\ge 5+3t_0-f_0 = 2(t_0+3-f_0)+t_0+f_0-1 \\ &\ge t_0+f_0-1\ . \end{aligned}$$

Restricting to $0 \le s \le 1$, this implies

$$s^{3+3t_0+1-\sum_{i=1}^{f_0-1}(\delta_{f_i>0}-\delta_{f_i=0})} \le s^{t_0+f_0-1}\ . \tag{3.12}$$

Using now the Hypothesis 3.5.6, *i.e.*, $\varrho(t,f) \le K_1^t$, we get from Eq. (3.11) and Eq. (3.12):

$$A_M(s) - A_M(0) \le \sum_{t_0=0}^{M} (sK_1)^{t_0} \sum_{f_0-1=0}^{M} \prod_{i=1}^{f_0-1} sA_{M-1}(s) \le \frac{1-(sK_1)^{M+1}}{1-sK_1} \frac{1-(sA_{M-1}(s))^{M+1}}{1-(sA_{M-1}(s))}\ .$$

Restricting s further to $s \le 1/(2K_1)$ this leads to

$$A_M(s) - A_M(0) \le 2\frac{1-(sA_{M-1}(s))^{M+1}}{1-(sA_{M-1}(s))}\ .$$

Fix $s^* = \min(0.1, 1/(2K_1))$ and consider the map $F : x \mapsto 1 + 2/(1 - s^* \cdot x)$. One easily checks that F maps the interval $[1,5]$ to itself. Furthermore, we have $s^* \cdot x \le 1$ for $x \in [1,5]$. Starting with $x = A_0(s^*) = 1$ we conclude that for all M one has $A_M(s^*) \le 5$. This implies that the monotone sequence $A_M(s^*)$ converges as $M \to \infty$ and thus

$$A_{v,t,f} \le 5 \cdot (s^*)^{-3v-3t-f}\ .$$

Summing over v and using $v \le t$ and $f \le 4t$ we complete the proof. □

3.5.3 Bound on triangulations

Having discussed the number of trees, we now study the number of ways these trees can be made into triangulations by identifying faces and nodes. This process is patterned after the work of [24] and [25].

Our bounds are based on using the inverses of the moves open-a-2-face, remove-1-tetra, and split-a-node-along-a-path. Since we are only interested in the bound, we will allow for inverse moves which do not necessarily lead to 3-balls.

Remark 3.5.10. *While we over-count the number of triangulations, by allowing for moves which may not lead to 3-balls, we can in fact formulate precise conditions which guarantee that after each move, a 3-ball is obtained. These conditions are spelled out in Lemmas 3.5.11 and 3.5.14. This actually allows for efficient programming of the inverse operations.*

Bounding the number of rooted triangulations with no internal nodes

Let $\mathcal{R}_{t,f}$ be the set of all rooted trees of nuclei with t tetrahedra and f external faces and let $\mathcal{T}_{t,f,0}$ be the set of all rooted triangulations with t tetrahedra, f external faces and no internal nodes. In this section, we will define the inverse move of open-a-2-face and we will use it to count the number of rooted triangulations with no internal nodes.

The inverse operation of open-a-face, which we will call *identification of 2 adjacent external faces* or simply *identification* when there is no ambiguity, is to identify two adjacent external faces, satisfying some conditions. Indeed, identifying any two adjacent external faces might lead to a complex which is not a triangulation. For instance, assume that (n_1, n_2, m_1) and (n_1, n_2, m_2) are two adjacent external faces such that there exists a node x adjacent to both m_1 and m_2. After identifying the two faces, we obtain a complex with a double edge $(x, m_1) = (x, m_2)$.

Lemma 3.5.11. *Consider a triangulation T. Let (a, b) be an external edge and let x, y be its two opposite external nodes (so that (x, a, b) and (y, a, b) are external faces). Assume that the following conditions are satisfied:*

- *The nodes x and y are not connected by an edge.*
- *The only nodes m such that (m, x) and (m, y) are edges are the two nodes a and b.*

Then, one can identify the two external nodes x and y as well as the two external faces sharing (a, b). This operation transforms a 3-ball to a 3-ball, and will be called identification (of two adjacent external faces).

Proof. The proof is left to the reader. □

Proposition 3.5.12. *Under Hypothesis 3.5.6, there is a constant K_3 such that for all t and f one has*

$$|\mathcal{T}_{t,f,0}| \leq K_3^t \ .$$

Proof. Let $T \in \mathcal{T}_{t,f,0}$ be any rooted triangulation with no internal nodes. Using repetitively the move open-a-2-face on T transforms it into a rooted triangulation T' with no internal nodes such that each internal face has 0, 1 or 3 external edges. In other words, T' is a rooted tree of nuclei. Equivalently, given a rooted tree of nuclei T' with t' tetrahedra and f' external faces, one can count the number of ways one can identify two adjacent external faces, *without any conditions guaranteeing ballness*. Summing this number over all rooted trees of nuclei gives us an upper bound on the number of rooted triangulations with no internal nodes.

We count the number of $T \in \mathcal{T}_{t,f,0}$ obtained by identification from a rooted tree of nuclei T' with t' tetrahedra and f' external faces. This means that we identify $D = (f' - f)/2$ pairs of adjacent external faces.

We first observe that choosing a pair of adjacent external faces is equivalent to choosing an external edge. We then note that some faces which are not adjacent in T' might become adjacent after some identifications are done. This means that we have a sequence $e_1, e_2, \dots, e_\ell$ with $e_i \geq 1$ and $\sum_i e_i = D$ which is defined as follows:

- e_1 is the number of external edges (or equivalently of pairs of adjacent external faces) of T' which are identified.
- e_2 is the number of pairs of faces which were not adjacent in T' but became so after the first series of e_1 identifications. However, each identification of two adjacent external faces creates exactly two new pairs of adjacent external faces, implying that $e_2 \leq 2e_1$.
- e_i is defined by analogy from the e_{i-1} identifications, implying that $e_i \leq 2e_{i-1}$.

This leads to the following bound (recall that the number of external faces f' is bounded by $4t$):

$$|\mathcal{T}_{t,f,0}| \leq \sum_{f'>f} |\mathcal{R}_{t,f'}| \sum_{\ell=1}^{D\equiv(f'-f)/2} \sum_{\sum_{i=1}^{\ell} e_i = D, e_i \geq 1} \binom{3f'/2}{e_1}\binom{2e_1}{e_2}\cdots\binom{2e_{\ell-1}}{e_\ell}.$$

Since $\binom{a}{b} \leq 2^a$, we find, using Proposition 3.5.9 to bound $|\mathcal{R}_{t,f'}|$,

$$\begin{aligned}
|\mathcal{T}_{t,f,0}| &\leq \sum_{f'>f} |\mathcal{R}_{t,f'}| 2^{(5f'/2-f)} \sum_{\ell=1}^{D\equiv(f'-f)/2} \sum_{\sum_{i=1}^{\ell} e_i = D, e_i \geq 1} 1 \\
&\leq \sum_{f'>f} |\mathcal{R}_{t,f'}| 2^{(5f'/2-f)} \sum_{\ell=1}^{D\equiv(f'-f)/2} \binom{D-1}{\ell-1} \\
&\leq \sum_{f'>f} |\mathcal{R}_{t,f'}| 2^{(3f'-3f/2)} \\
&\leq \sum_{f'=f+2}^{4t} K_2^t K_2^{3f'} \\
&\leq K_2^{13t} = K_3^t ,
\end{aligned}$$

where $K_3 = K_2^{13}$.

The proof is complete. □

Bounding the number of rooted triangulations (internal nodes included)

In this section, we define the inverse moves of remove-1-tetra and split-a-node-along-a-path and we use them to count the number of rooted triangulations.

Definition 3.5.13. *We define the inverse move of remove-1-tetra, which we call* adding a tetrahedron: *Consider a triangulation T. Let x be an external node with external degree equal to 3 and let a_1, a_2 and a_3 be its external neighbors, i.e., (x, a_i) is an external edge. Adding a tetrahedron then consists in adding the face (a_1, a_2, a_3) and the tetrahedron (x, a_1, a_2, a_3).*

We define the inverse move of split-a-node-along-a-path.

Lemma 3.5.14. *Consider a triangulation T. Let $e = (a, b)$ be an external edge. Assume that the following conditions are satisfied:*

- *For each node m such that (m, a) and (m, b) are edges, (m, e) is a face.*
- *For each edge e' such that (e', a) and (e', b) are faces, (e', e) is a tetrahedron.*
- *There are no faces f such that (f, a) and (f, b) are both tetrahedra.*

Then, one can collapse the two nodes a and b, and the result is again a 3-ball. This move is called collapse of an external edge *or simply* collapse.

Proof. The proof is left to the reader. □

Notation 3.5.15. *The above three conditions of a collapse are written in a concise way as*

$$\mathcal{I}(a) \cap \mathcal{I}(b) = \mathcal{I}(e) \ .$$

In Sect. 3.4.3, we described an algorithm which transforms any triangulation with f-vector $\langle t, f, n\rangle$ into a triangulation with f-vector $\langle t', f', 0\rangle$. We have the following lemma:

Lemma 3.5.16. *There is a constant $K_4 > 0$ such that the f-vectors $\langle t, f_s, n_i\rangle$ and $\langle t', f'_s, 0\rangle$ satisfy the following linear relation:*

$$t' \leq K_4 t \ , \qquad f'_s \leq K_4 t \ , \tag{3.13}$$

Proof. Let e, e' be the number of internal edges of both triangulations. By Corollary 3.4.14, we have $e' - e \leq C_\Delta(t + n_i)$. Using Eq. (3.2) and the obvious relations $n_i \leq 4t$, $f_s < 2n_s < 4e_i$ and $e_i < 6t$, (and analogous ones with primes) the result follows. □

This proves that any triangulation in $\mathcal{T}_{t,f,n}$ can be obtained from a triangulation with no internal nodes in $\mathcal{T}_{t',f',0}$ with a series of carefully chosen collapses and additions of tetrahedra, with t, f, n, t', f' satisfying Eq. (3.13).

We can now use a similar approach to that of the previous section. It is clear that choosing a triplet of external faces for the move add-1-tetrahedron is equivalent to choosing an external node x, and that choosing a couple of external nodes for collapse is equivalent to choosing an external edge.

3.5.4 Combining the bounds

Before we state our main result, we recall the

Hypothesis 3.5.6 1. *There is a finite constant $K_1 > 1$ such that the number $\varrho(t, f)$ of face-rooted nuclei with f-vector $\langle t, f\rangle$ is bounded by K_1^t.*

Theorem 3.5.17. *Under Hypothesis 3.5.6 one has the bound: There is a finite constant C such that the number of rooted triangulations with f-vector $\langle t, f, n\rangle$ is bounded by*

$$|\mathcal{T}_{t,f,n}| \leq C^t \ . \tag{3.14}$$

Proof. Consider a rooted triangulation $T \in \mathcal{T}_{t,f,n}$ with t tetrahedra, f external faces and n internal nodes. We showed that T can be obtained from a rooted triangulation $T' \in \mathcal{T}_{t',f',0}$ by a series of carefully chosen collapses and additions of tetrahedra.

Note that the algorithm of Sect. 3.4.3 which transforms T into T' can always be stopped when the last internal node of T is removed. This implies that, in the inverse construction we are doing now, we must start by adding tetrahedra to T', and not by collapsing external edges. So the first step is to choose n_1 external nodes (of external degree 3) out of the $f'/2+2$ external nodes of T', and to insert a tetrahedron on each of them with one tip at the node. We call this "covering the node."

This reduces the number of external edges from $3f'/2$ to $3(f'/2 - n_1)$. Then, we choose m_1 external edges and we collapse them. The possibility to simultaneously collapse $m_1 > 1$ edges is justified as follows:

Any labeled triangulation can be viewed as a list of tetrahedra L_t satisfying certain conditions (a face is shared by no more than 2 tetrahedra etc...). From this point of view, collapsing an external edge e is simply the operation where we remove from L_t all the tetrahedra of $\mathcal{E}(e)$ (Lemma 3.5.14 guarantees that these conditions remain true, *i.e.*, that the resulting list of tetrahedra is indeed a triangulation). Let e_1 and e_2 be two collapsible edges. The construction implies that the order in which we collapse them is irrelevant and so, the idea that we simultaneously collapse m_1 edges makes sense. One should pay attention to the case where we collapse two edges $e_1 = (a, b_1)$ and $e_2 = (a, b_2)$ when $(b_1, b_2) = e_3$ is also an edge. In this case, all tetrahedra sharing one of the three edges are removed simultaneously. Clearly, this yields the same result regardless of the order in which we collapse e_1 and e_2.

The next step is to choose n_2 external nodes among the new possibilities which appear after performing the first series of coverings and collapses, and cover them. For each external edge e, we can associate four nodes: the two endpoints of e and the two nodes x_1, x_2 such that (x_i, e) is an external face. Assume that x is one of the n_2 chosen external nodes. The fact that x appeared after the first series implies that x is either one of the four nodes associated with one of the m_1 collapsed edges (note that these four nodes become three after the collapse), or that there is a node y among the first n_1 nodes such that (x, y) was an external edge (before covering y with a tetrahedron). But each such y has exactly 3 external neighbors. This implies that $n_2 \leq 3m_1 + 3n_1$ and the number of ways to choose these nodes is bounded by

$$\binom{3(m_1 + n_1)}{n_2}.$$

Continuing in this way, we choose m_2 external edges and we collapse them. Let e be such an edge. Again, e was not among the first m_1 edges. This implies that there must be a node x of the series of n_2 covered external nodes such that (e, x) formed an external face before covering x with a tetrahedron. But for each such x there are exactly three external edges satisfying this condition. We deduce that $m_2 \leq 3n_2$.

We continue adding tetrahedra and collapsing edges. This leads to two sequences

$n_i, m_i, i = 1, \dots, \ell$, with $\ell \le n$, satisfying:

$$\begin{aligned} 1 \le n_i \ , \qquad 0 \le m_i \le 3n_i \ , \qquad \sum_{i=1}^{\ell} n_i = n \ , \\ 1 \le n_i \le 3n_{i-1} + 3m_{i-1} \ , \quad i > 1 \ , \\ \sum_{i=1}^{\ell} (2n_i + 2m_i) + f = f' \ . \end{aligned} \tag{3.15}$$

The last identity follows because each collapse of an external edge and each covering of an external node (by a tetrahedron) reduces the number of external faces by 2. Note that some, or all, of the m_i's might be equal to zero. Using Eq. (3.15) we get a bound

$$\begin{aligned} |\mathcal{T}_{t,f,n}| \le \sum_{t',f'} |\mathcal{T}_{t',f',0}| \sum_{\ell=1}^{n} \sum_{\sum_{i=1}^{\ell} n_i = n, n_i \ge 1} \ \sum_{\sum_{i=1}^{\ell} m_i = (f'-f)/2 - n, m_i \ge 0} \\ \times \binom{f'/2+2}{n_1} \binom{3(n_1+m_1)}{n_2} \cdots \binom{3(n_{\ell-1}+m_{\ell-1})}{n_\ell} \\ \times \binom{3f'/2}{m_1} \binom{3n_1}{m_2} \cdots \binom{3n_{\ell-1}}{m_\ell} \ , \end{aligned}$$

where the sum over t', f' is restricted by Eq. (3.13). Bounding each binomial by a power of 2 and using Proposition 3.5.12, Eq. (3.15), Eq. (3.13), and $n_s = f_s/2 + 2$, we get, as in the proof of Proposition 3.5.12,

$$|\mathcal{T}_{t,f,n}| \le \sum_{t',f' \le K_4 t} K_3^{t'} \le C^t \ .$$

This shows Eq. (3.14) and completes the proof. □

Chapter 4

From Nuclei to Atoms

4.1 Introduction

Atoms, defined below, are a subset of nuclei. We define $\mathcal{A}(t)$ as the set of atoms with t tetrahedra.

Hypothesis 4.1.1. *Throughout this chapter, we assume that the number $A(t) = |\mathcal{A}(t)|$ of atoms with t tetrahedra is bounded by $A(t) \leq C_1^t$, where $C_1 > 1$ is some constant.*

Let $N(t)$ be the number of nuclei with t tetrahedra. The purpose of this chapter is to prove the following result:

Theorem 4.1.2. *Under Hypothesis 4.1.1, there is a constant $C > 1$ such that $N(t) \leq C^t$.*

In the previous chapter, we introduced the move of splitting a node n into an external edge $e = (n_R, n_L)$ along a path γ (we will simply call it a *splitting*). We also introduced the inverse move of collapsing the external edge e (the identification of the nodes a and b of e) into a single node n (in this case, the path γ is simply the flower of e).

Atoms are defined as follows:

Definition 4.1.3. *Given a nucleus T', we start collapsing its external edges until no more collapsible edges remain. The resulting triangulation is called* atom.

Any triangulation can be seen as a list of tetrahedra verifying certain conditions. From this point of view, we showed that collapsing an edge e of the triangulation is equivalent to removing from the corresponding list all tetrahedra having e as an edge. From this we showed that the order in which we collapse the edges is irrelevant, hence the following lemma:

Lemma 4.1.4. *Starting with a given nucleus, we always end up with the same atom regardless of the order in which we collapse the edges.*

To put this result in a more general context, recall that Benedetti and Ziegler [25] showed that the main reason why an exponential bound has not been established for all triangulations of S^3 is the existence of triangulations with knotted edges. Several methods have been proposed to construct such triangulations (Bing [29], Lutz [30], etc...). From our point of

view, each such triangulation leads to what we call a nucleus. Consider for instance a knot, say the trefoil knot. One can construct a triangulated ball with a trefoil-knotted edge in several ways. Each such construction leads to a distinct nucleus. Furthermore, given any nucleus, splitting a node along a path leads to a nucleus having the same topological properties. The number of nuclei with a trefoil-knotted edge is large (and the number of balls with such an edge is even larger). On the other hand, and because of Lemma 4.1.4, it seems that the number of atoms with a trefoil-knotted edge is small, maybe even equal to one. If one can prove such a result for all knots, then one can bound the number of atoms with t tetrahedra with the number of knots with certain complexity (the crossing number can be considered as a measure of the complexity of a knot, or maybe the number of generators of the fundamental group of the complement). It was shown by Sundberg [36] that the number of knots with t crossings is exponentially bounded. Using such results and Theorem 4.1.2, one can hope to obtain an exponential bound for all triangulations of S^3.

4.1.1 Notation

Nuclei are obtained from atoms by splitting external nodes along paths.

Definition 4.1.5. *Let A be an atom and let $\mathcal{N}_A(n,t)$ be the set of all nuclei with n nodes and t tetrahedra obtained from A by splittings.*

$\mathcal{N}_A(n,t)$ is simply the set of nuclei with n nodes and t tetrahedra such that collapsing all possible edges gives A. By the above, we deduce that if $A \neq A'$ are 2 distinct atoms, then $\mathcal{N}_A(n,t) \cap \mathcal{N}_{A'}(n,t) = \emptyset$. We can then write:

$$N(t) = \sum_{t' \leq t} \sum_{A \in \mathcal{A}(t')} \sum_{n \leq 4t} |\mathcal{N}_A(n,t)| \ .$$

Define $N_A(t)$ as the number of nuclei with t tetrahedra obtained from the atom A by splittings, *i.e.*, $N_A(t) = \sum_{n \leq 4t} |\mathcal{N}_A(n,t)|$. From this, we deduce that proving Theorem 4.1.2 is equivalent to proving the following theorem:

Theorem 4.1.6. *Let A be an atom. Then there is a constant $C > 1$ such that $N_A(t) \leq C^t$.*

Notation 4.1.7. *Let T be a nucleus (or an atom). The splitting of a node $n \in T$ along a path γ is denoted by $S(n,\gamma)$ and the result is the nucleus $S(n,\gamma)T$.*

Consider an atom A with t' tetrahedra and n' nodes. We can then write:

$$\mathcal{N}_A(n,t) = \sum_{\mathbf{n},\gamma} S(n_k,\gamma_k) S(n_{k-1},\gamma_{k-1}) \cdots S(n_1,\gamma_1) A \ ,$$

where $\sum_{\mathbf{n},\gamma}$ is the sum over all possible sequences of $k \equiv n - n'$ nodes $\mathbf{n} = (n_1, \ldots, n_k)$ and paths $\gamma_1, \ldots, \gamma_k$ such that $\sum_{i=1}^{k} |\gamma_i| = t - t'$, where $|\gamma_i|$ is the number of edges in the path $|\gamma_i|$.

In other terms, we need to count the number of such ordered sequences of splittings $S(n_k,\gamma_k) S(n_{k-1},\gamma_{k-1}) \ldots S(n_1,\gamma_1)$ under the condition that if 2 sequences applied to A give the same result, then they are counted once. In particular, we need to look at the commutativity of the splittings; for example, if all splittings commute, then any 2 sequences that differ only by a permutation should be counted as one.

4.1.2 Structure of the proof

In Sect. 4.3, we show that 2 arbitrary splittings commute in all but one case; the non-commuting case is illustrated in Sect. 4.3.1. The main idea of the proof is to draw all the paths before performing any splitting. The number of ways this can be done is counted in Lemma 4.2.3 and Remark ⋆ of Sect. 4.4.3. Then we start splitting. The number of ways of choosing which nodes to split, taking into account the non-commutativity of the splittings, is counted in Sect. 4.4.2 and Sect. 4.4.3. Of course, once we start splitting nodes, some of the paths we drew need to be modified. The number of ways this can be done is counted in Sect. 4.4.3, Sect. 4.4.3 and Sect. 4.4.3. Finally in Sect. 4.4.4, putting everything together, using some general well known results (Sect. 4.2), we prove Theorem 4.1.6 and therefore Theorem 4.1.2.

4.2 Some general results

In this section we recall some classical results we will be using.

4.2.1 Combinatorics

Lemma 4.2.1. *Let N be an integer. The number of ordered sequences $(a_1, \dots, a_M)$ of M integers $a_i \in \{0, 1, \dots\}$ such that $\sum_{i=1}^{M} a_i = N$ is given by $\binom{N+M-1}{N} < 2^{N+M}$.*

We now consider a 2d triangulation $\mathcal{I}$ with d nodes and e internal[1] edges. We want to divide it into k parts along the e internal edges; we say that a division into parts is *admissible* if it satisfies the following restrictions:

- R1 Each part is connected in the sense that any two nodes of a part are connected by a path of edges in the same part.
- R2 No part has a hole in it, *i.e.*, the boundary of any part is homeomorphic to a circle.
- R3 The division is planar. In particular, two parts cannot intersect one another.
- R4 The intersection of two parts is by definition their shared *boundary*. These boundaries must be connected paths, not necessarily simple.

In Fig. 4.1, we illustrate these restrictions. Note that internal edges of $\mathcal{I}$ can belong to more than two parts while still verifying the planarity condition, as seen in the second panel of Fig. 4.1. We define the multiplicity of an internal edge as the number of times it appears in the boundaries. For example, the green edge of the second panel of Fig. 4.1 has multiplicity 2.

We would like to count the number of such divisions.

Remark 4.2.2. *The purpose of splitting, and as a consequence the strategy used in this chapter differs completely from those of the previous chapter. There, the purpose of splitting is to promote every internal node into an external one without adding too many tetrahedra;*

[1] internal and external nodes, edges and faces in the 3d case as defined in the same manner as in the previous chapter

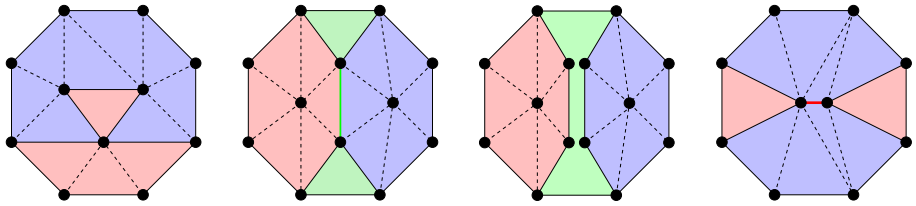

Figure 4.1: This figure illustrates the restrictions on the division of Sect. 4.2.1 of a 2d triangulation. The first panel shows a valid division into two parts such that the path forming the boundary between them is not simple. In the second panel, we divide the triangulation into three parts colored red, green and blue. Note that the bold green vertical edge in the middle of the panel belongs to the green part, making it connected. This division should be interpreted as the third panel: the blue and red parts do not share a boundary; they are separated by the green part. The bold green edge belongs to both red-green and green-blue boundaries. We say that such an edge has multiplicity 2. The last panel shows a invalid division: since the bold red edge belongs to the red part, the red and blue parts intersect one another.

thus, the strategy is to carefully choose the paths along which we split in such a way that every internal edge of every hemisphere belongs to at most 1 path (in each of the sweeps C2→C1→C0). In this chapter, given an integer t and an atom A with $t' < t$ tetrahedra, the purpose is to count all possible ways of splitting nodes of A such that the end triangulation has t tetrahedra.

Lemma 4.2.3. *Consider a 2d triangulation $\mathcal{I}$ with n nodes and e internal edges and let ℓ be an integer. We divide $\mathcal{I}$ into k parts verifying the above conditions R1-R4, such that the total length of the boundaries of the parts is ℓ (each edge is counted according to its multiplicity). The number of these divisions is bounded by $C^{\ell+n}$ where C is some constant.*

Proof. We start by choosing the boundaries among the e internal edges of $\mathcal{I}$, then we count the number of ways the parts can be chosen. By 2d Euler, one can show that $e = 3n - 2p - 3 < 3n$, where p is the number of external edges of $\mathcal{I}$. To each such edge e', we assign a multiplicity $m(e') \geq 0$ which counts the number of boundaries this edge is part of. By definition of ℓ, we have $\sum_{e'} m(e') = \ell$. By Lemma 4.2.1, the number of ways of choosing these multiplicities and therefore the number of ways of choosing the boundaries is bounded by $2^{\ell+3n}$. Note that the boundaries chosen in this manner have no restrictions so we are clearly over-counting.

Consider such a collection of boundaries. It divides $\mathcal{I}$ into k' parts satisfying R1-R4 (if not, then it is a bad collection and we just ignore it and pick another one, *i.e.*, such a bad collection of boundaries is among the ones we overcount). For example, looking at the second panel of Fig. 4.1, one would get four parts up to this point, and three parts for the first panel.

We draw $\mathcal{I}$ and we remove all edges e' of multiplicity 0. Let $e' = (a, b)$ be an internal edge with multiplicity $m(e') > 0$. Assume that $m(e') = 2$. We draw it as two separate edges joining the nodes a and b. In other terms, e' in itself forms one additional part (see Fig. 4.2).

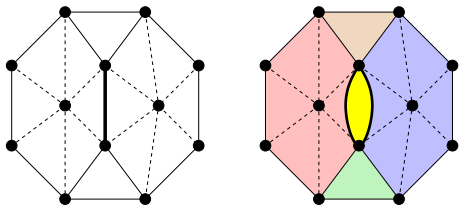

Figure 4.2: The first panel reproduces the division of the second panel of Fig. 4.1. The dashed edges are not used in the division and hence their multiplicity is null. The straight edges have multiplicity 1 and the bold edge has multiplicity 2. The second panel shows the 5 parts induced by the choice of the edges and their multiplicities. Note that the bold edge of multiplicity 2 forms in itself the yellow part.

More generally, each edge e' with multiplicity $m(e') \geq 2$ forms $m(e') - 1$ additional part(s) that we add to the already existing k' parts. We get a total of k'' parts verifying R1-R4. For example, k'' is 3 for the first panel of Fig. 4.1 and 5 for the second one, as seen in Fig. 4.2. If $k'' < k$, then we ignore the current collection of boundaries and we pick another one. Finally, we need to count the number of ways of identifying $k'' - k$ parts. This is done as follows:

- We consider all ℓ edges of the boundaries (each edge according to its multiplicity). We consider every internal node where four or more of these ℓ edges meet. Let n' be such a node and let $u = u(n') \geq 4$ be the number of these edges having n' as corner. We have $\sum_{n'} u(n') \leq 2\ell$.
- By planarity, these $u(n')$ edges can be ordered, say clockwise. Two consecutive edges form the corner of a part. To each such corner, we associate a distinct vertex. We get u vertices. Two vertices are connected if and only if the corresponding corners belong to the same part. In other terms, to each such node n', we draw some sort of dual graph depicting the parts around it. We identify two parts as one if the corresponding vertices are connected. The planarity condition R3 of the division implies that these "local dual graphs" must be planar. For example, considering once more Fig. 4.1, we see that in the first panel, there is only one node with $u(n') \geq 4$ and there are two such nodes in the second panel (the ends of the bold green edge); the diagrams associated with these nodes are respectively:

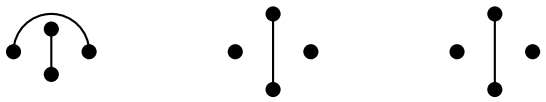

The number of planar graphs with u vertices is bounded by C^u [37]. Multiplying all these bounds for every internal node n', we get a bound of the form

$$C^{\sum_{n'} u(n')} \leq C^{2\ell} \ .$$

Since the identifications are done locally around each node, some of them might lead, once the entire triangulation is considered, to a division that violates the conditions R1-R4, which in turn leads to over-counting. However, it is clear from the construction that any valid division can be obtained in the manner described above.

□

4.2.2 Rearranging sequences of operators

Let $\mathcal{S} = \mathcal{S}_1 = A_n \cdots A_1$ be an ordered sequence of operators. We define the following subsequences $\{\mathcal{K}_i\}_{i \leq I}$ of operators:

The first subsequence is defined by

$$\mathcal{K}_1 = \{A_i \in \mathcal{S}_1 : \forall j < i \text{ such that } A_j \in \mathcal{S}_1, [A_i, A_j] = 0\} .$$

It is simply the set of all operators of $\mathcal{S}_1$ that commute with every operator to their right in $\mathcal{S}_1$. By definition, we have $A_1 \in \mathcal{K}_1$. Furthermore, all operators of $\mathcal{K}_1$ commute with one another.

The other subsequences $\ell \geq 2$ are defined recursively:

- We consider the ordered subsequence $\mathcal{S}_\ell = \mathcal{S}_{\ell-1} \setminus \mathcal{K}_{\ell-1}, \ell \geq 2$. For instance, to get $\mathcal{S}_2$, we simply remove from the original sequence $\mathcal{S}_1$ all the operators of $\mathcal{K}_1$ while maintaining the original order of the operators.
- $\mathcal{K}_\ell$ is then given by all the operators of $\mathcal{S}_\ell$ that commute with every operator to their right in of $\mathcal{S}_\ell$:

$$\mathcal{K}_\ell = \{A_i \in \mathcal{S}_\ell : \forall j < i \text{ such that } A_j \in \mathcal{S}_\ell, [A_i, A_j] = 0\} .$$

The maximal index I is of course defined as the largest integer i such that $\mathcal{K}_i \neq \emptyset$. Since each subsequence $\mathcal{K}_\ell, \ell \leq I$ contains at least one operator, namely the first operator of $\mathcal{S}_\ell$, we have $I \leq n$. By construction, every operator of the original sequence $\mathcal{S}_1$ belongs to exactly one subsequence $\mathcal{K}_i, i = 1, \ldots, I$.

We now consider the new sequence $\mathcal{S}'$ of operators given by

$$\mathcal{S}' = \mathcal{K}_I \cdots \mathcal{K}_1 .$$

Lemma 4.2.4. *We have $\mathcal{S}' = \mathcal{S}$. Furthermore, $\mathcal{S}'$ has the following properties:*

P1 $\mathcal{S}'$ is obtained from $\mathcal{S}$ by changing the order of pairs of commuting operators.

P2 The operators within each subsequence $\mathcal{K}_i, i = 1, \ldots, I$ commute with one another.

P3 For every $\ell > 1$ and for every operator $X \in \mathcal{K}_\ell$, there is an operator $Y \in \mathcal{K}_{\ell-1}$ such that $[X, Y] \neq 0$.

Proof. We prove each point of the lemma.

- P2 is immediate from the construction of the blocks $\mathcal{K}_i$.
- P1 is proved recursively. We start by moving all the operators of $\mathcal{K}_1$ to the right of $\mathcal{S} = \mathcal{S}_1$. Note that moving operators to the right is done by commuting pairs. By construction, every operator $X \in \mathcal{K}_1$ commutes with all the operators to its right. We get a new sequence $\mathcal{S}'_2 = \mathcal{S}_2 \cdot \mathcal{K}_1$. Recall that $\mathcal{S}_2 = \mathcal{S}_1 \setminus \mathcal{K}_1$ and the order of the operators in $\mathcal{S}_2$ is the same as that in $\mathcal{S}_1$. By construction we have $\mathcal{S}'_2 = \mathcal{S}$ since we only changed the order of commuting pairs.
  Then we move all operators of $\mathcal{K}_2$ to the right of $\mathcal{S}_2$. We obtain $\mathcal{S}'_3 = \mathcal{S}_3 \cdot \mathcal{K}_2 \cdot \mathcal{K}_1$ and by construction, since all operators of $\mathcal{K}_2$ commute with all the operators to their right in $\mathcal{S}_2$, we have $\mathcal{S}'_3 = \mathcal{S}$. We do this for all $\ell \leq I$ and we get $\mathcal{S}'_I = \mathcal{K}_I \cdots \mathcal{K}_1 = \mathcal{S}$.

- Let $\ell > 1$ and consider $X \in \mathcal{K}_\ell \subset \mathcal{S}_\ell$ such that $X \notin \mathcal{K}_{\ell-1}$. The sequence $\{\mathcal{S}_\ell\}_{\ell \leq I}$ is decreasing in the sense that every operator of $\mathcal{S}_\ell$ is an operator of $\mathcal{S}_{\ell-1}$. Since $X \notin \mathcal{K}_{\ell-1}$, by definition of $\mathcal{K}_{\ell-1}$, there is an operator Y to the right of X in $\mathcal{S}_{\ell-1}$ such that $[X, Y] \neq 0$. Since $X \in \mathcal{K}_\ell$, it commutes with every operator to its right in $\mathcal{S}_\ell$, implying that $Y \notin \mathcal{S}_\ell$. We conclude that $Y \in \mathcal{S}_{\ell-1} \setminus \mathcal{S}_\ell = \mathcal{K}_{\ell-1}$, which proves P3.

□

4.3 Commutativity of the splittings

We saw that external edges can be collapsed in any order we want, *i.e.*, let $C(e_1)$ and $C(e_2)$ be the operations of collapsing edges e_1 and e_2 respectively and let T be a triangulation, then

$$C(e_2)C(e_1)T = C(e_1)C(e_2)T \text{ for all collapsible edges } e_1, e_2 \text{ of } T \ .$$

Remark 4.3.1. *There is one special case where e_2 becomes collapsible only after e_1 is collapsed, i.e., that $C(e_2)T$ is not a ball (it contains a double edge). Even in this case, $C(e_1)C(e_2)T$ is once again a ball and furthermore it is the same ball as $C(e_2)C(e_1)T$.*

Splitting a node along a path is the inverse operation of collapsing external edges. From the above relation, one expects that the splittings should also commute, and indeed such a result is true in all but one case (which is closely related to Remark 4.3.1) but less obvious for the simple reason that the operation of collapsing an edge requires only one parameter, namely the edge in question, but the operation of splitting a node requires an additional parameter, namely the path along which we split.

In what follows, we examine all cases of 2 successive splittings and show that they commute in all but one case. The notation is the following: we first split the node n in some triangulation, say T_1, along the path γ_n; we obtain another triangulation, say T_2. Then we split m in T_2 along the path γ_m. The splitting is denoted by $S(n, \gamma_n)$ and similarly for m. Note that m can be a child of n (in which case we will say that $n = m$) and that γ_m is drawn in T_2.

Definition 4.3.2. *We say that 2 splittings* commute *if there are 2 paths γ'_n and γ'_m such that*

$$S(m, \gamma_m)S(n, \gamma_n) = S(n, \gamma'_n)S(m, \gamma'_m) \ ,$$

i.e., splitting n along γ_n then m along γ_m yields the same result as splitting m along γ'_m then n along γ'_n.

Before we start examining the cases, we recall the following fact

Remark 4.3.3. *An external edge $e = (a, b)$ is collapsible if $\mathcal{I}(a) \cap \mathcal{I}(b) = \mathcal{I}(e)$.*

The only case where the splittings do not commute is the following:

Lemma 4.3.4. *Let $n \neq m$ be 2 nodes such that (n, m) is an internal edge. We split n into n_{R} and n_{L} along a path γ_n such that $m \in \gamma_n$. Then we split m along a path γ_m such that:*

- *Both nodes n_{R} and n_{L} are in γ_m.*

- *The nodes n_{R} and n_{L} are not consecutive in γ_m, i.e., $(n_{\mathrm{R}}, n_{\mathrm{L}})$ is not an edge of γ_m.*

These 2 splittings do not commute and there are no other cases of non-commutativity.

Proof. Let $e = (a, b)$ be an external edge. We recall that the condition $\mathcal{I}(a) \cap \mathcal{I}(b) = \mathcal{I}(e)$ of Remark 4.3.3 actually means the following 2 things:

- For every node c such that (c, a) and (c, b) are edges, (c, e) is a face, **and**
- For every edge e' such that (e', a) and (e', b) are faces, (e', e) is a tetrahedron.

We also recall that splitting a node n into n_{R} and n_{L} along a path γ_n results in the following:

- Let c be a node. Then

$$c \text{ is a node of } \gamma_n \Leftrightarrow (n_{\mathrm{R}}, n_{\mathrm{L}}, c) \text{ is a face .}$$

- Let e' be an edge. Then

$$e' \text{ is an edge of } \gamma_n \Leftrightarrow (n_{\mathrm{R}}, n_{\mathrm{L}}, e') \text{ is a tetrahedron .}$$

The general approach to proving commutativity is as follows:

- We consider a triangulation T.
- We split n into n_{R} and n_{L} along γ_n. By construction, the external edge $(n_{\mathrm{R}}, n_{\mathrm{L}})$ satisfies Remark 4.3.3.
- We split m into m_{R} and m_{L} along γ_m. By construction, the external edge $(m_{\mathrm{R}}, m_{\mathrm{L}})$ satisfies Remark 4.3.3.
- We check that the edge $(n_{\mathrm{R}}, n_{\mathrm{L}})$ still satisfies Remark 4.3.3. Let γ'_n be its flower. We collapse it into a node we call n.
- We check that the edge $(m_{\mathrm{R}}, m_{\mathrm{L}})$ still satisfies Remark 4.3.3. Let γ'_m be its flower. We collapse it into a node we call m.
- We check that the final triangulation we obtain is the same as the original one T.

The non-commutativity arises when one of the 3 checks fails. We show that this can happen in only one case.

We assume that the first check fails, *i.e.*, that after splitting m along γ_m, the external edge $(n_{\mathrm{R}}, n_{\mathrm{L}})$ becomes non-collapsible. This means that there is a node c such that (c, n_{L}) is an edge, (c, n_{R}) is an edge and $(c, n_{\mathrm{R}}, n_{\mathrm{L}})$ is not a face, or that there is an edge e' such that (e', n_{R}) is a face, (e', n_{L}) is a face and $(e', n_{\mathrm{R}}, n_{\mathrm{L}})$ is not a tetrahedron. However, before the splitting of m, $(n_{\mathrm{R}}, n_{\mathrm{L}})$ was collapsible. This means that the node c or the edge e' we are looking for are respectively the new nodes $c = m_i, i = \mathrm{R}, \mathrm{L}$ or the new edge $e' = (m_{\mathrm{R}}, m_{\mathrm{L}})$. We will see that these 2 conditions are always satisfied simultaneously, *i.e.*, such a node c exists if and only if such an edge e' exists. Furthermore, this happens in only one case which is none other than the case we are looking for.

We first assume that the first condition is true, *i.e.*, that such a node exists, say $c = m_{\mathrm{R}}$. Recall that the path γ_m splits the flower $\mathcal{I}(m)$ into 2 parts associated with m_{R} and m_{L} respectively. Since both $(m_{\mathrm{R}}, n_{\mathrm{L}})$ and $(m_{\mathrm{R}}, n_{\mathrm{R}})$ are edges, we deduce that $m \in \gamma_n$ implying that

(n_R, n_L) is an edge of $\mathcal{I}(m)$. We also conclude that both nodes n_R and n_L belong to the part of $\mathcal{I}(m)$ associated with m_R. But since (m_R, n_R, n_L) is not a face, we deduce that the edge (n_R, n_L) belongs to the other part of $\mathcal{I}(m)$ which is associated with with m_L. This proves that the nodes n_R, n_L belong to the intersection of the two parts of $\mathcal{I}(m)$, *i.e.*, $n_R, n_L \in \gamma_m$, but are not consecutive in γ_m (otherwise (n_R, n_L) would also be in γ_m).

We now assume that $e' = (m_R, m_L)$ exists. Since (m_R, m_L, n_R) and (m_R, m_L, n_L) are faces, we deduce that $m \in \gamma_n$ and that both nodes n_R and n_L are in γ_m. Since (m_R, m_L, n_R, n_L) is not a tetrahedron, we deduce that (n_R, n_L) is not an edge of γ_m and therefore n_R and n_L are non-consecutive in γ_m.

From the above, we deduce that the first of the 3 commutativity checks fails if and only if $m \in \gamma_n$ and n_R and n_L are non-consecutive nodes of γ_m. Next, we assume that the first check passes. We show that the next 2 checks always pass as well.

We begin with the second check and we assume it fails. This means that when we collapse (n_R, n_L) into n, the edge (m_R, m_L) which was collapsible becomes non-collapsible. Since there are no new edges and the only new node in the triangulation is n, we deduce that (n, m_R) and (n, m_L) are both edges but (n, m_R, m_L) is not a face. But we know that the first check passed; in particular, this implies that (n_R, n_L, m_R, m_L) was a tetrahedron and thus, after the collapse of (n_R, n_L) into n, (n, m_R, m_L) is a face. We deduce that if the first check passes, then so does the second.

Finally we come to the third check. Since a triangulation is a list of tetrahedra (satisfying certain conditions) and from this point of view, collapsing an edge e is simply the operation of removing all tetrahedra having e as an edge. It is then obvious that if we start from a triangulation T and we split n and m into (n_R, n_L) and (m_R, m_L) respectively, then we collapse these 2 edges, we get the original triangulation T regardless of the order of the collapses.

In summary, we showed that if the first check passes, then so do the other 2 and the splittings do commute. Non-commutativity arises only if the first check fails. This only happens in the case described in Lemma 4.3.4. □

We now illustrate the non-commuting case with an explicit example:

4.3.1 A non-commuting example

We split n along a path γ_n into n_R and n_L. Then we split m along γ_m. The paths are such that $m \in \gamma_n$ and the nodes n_R, n_L are non-consecutive node of γ_m. The case is illustrated in Fig. 4.3.

There are 4 sets of panels. The first set contains 2 hemispheres: the upper one is the hemisphere of the red node n and the lower one that of the blue node m. The path γ_n is the red path. We split the red node along the red path into a red node and a hollow-red node. Since the red path goes through the blue node, both new red nodes are in the hemisphere of the blue node, as shown in the second set of 3 panels (red top-left, hollow-red top-right and blue bottom). Then we split the blue node along the blue path. Notice that both new red nodes are non-consecutive in the blue path. We obtain the third set of 4 panels. It is clear that the edge e = (red , hollow-red) does not satisfy Remark 4.3.3, since the hollow-blue node is in the intersection of the hemispheres of both nodes but not in the flower of e. If we try to collapse e, we get the fourth set of 3 panels, belonging to a configuration which is

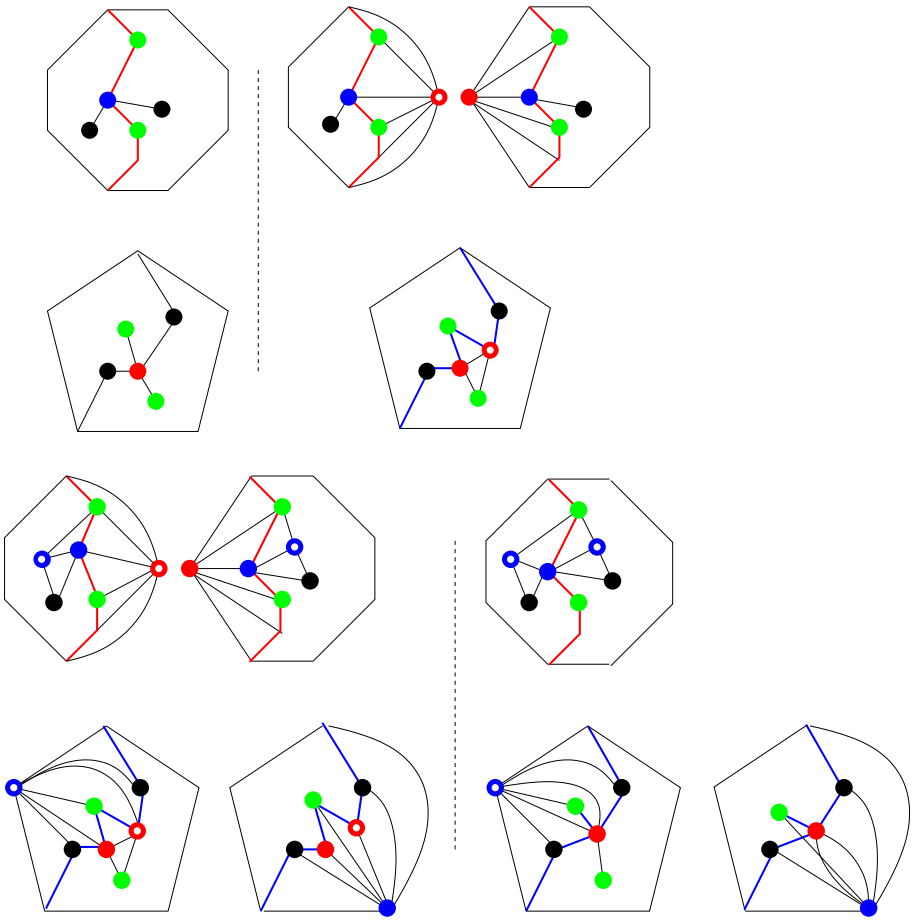

Figure 4.3: Illustration of Sect. 4.3.1. See the discussion there for an explanation of the different panels.

a not a valid triangulation, since, in particular, it contains the double-edge (hollow-blue,red) (looking at the hemisphere of the red node shown in the upper panel of the fourth set, we see that the same hollow-blue node appears twice, meaning that there are 2 distinct edges linking it to the red node).

In other terms, we showed that the triangulation of the third set of panels cannot be obtained from that of the first set of panels by splitting first the blue node and then the red node.

The 3-dimensional flowers are illustrated in Fig. 4.4.

We now illustrate 2 interesting commuting cases with an example each; the purpose of these examples is to show the relation between the new paths $\gamma'_i, i = \mathrm{L}, \mathrm{R}$ and the old ones γ_i. The trivial cases are not illustrated. It is left to the reader to check that in the trivial cases, we have $\gamma'_i = \gamma_i, i = \mathrm{R}, \mathrm{L}$.

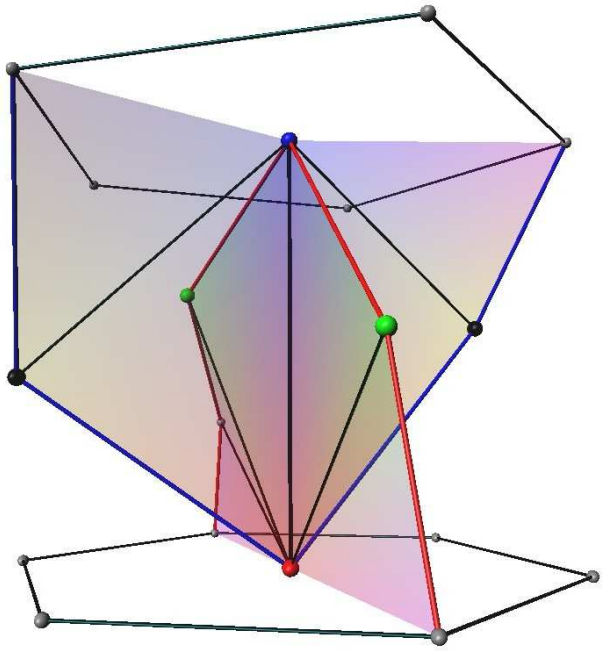

Figure 4.4: 3-dimensional illustration of the first set of 2 panels of Fig. 4.3.

4.3.2 The case $n \neq m$, $m \in \gamma_n$ and n_R, n_L are consecutive in $\gamma(m)$

We split n along γ_n into n_R and n_L. Then we split m along γ_m. We assume that $m \in \gamma_n$ and that the nodes n_R and n_L are consecutive nodes of γ_m. See Fig. 4.5 for an illustration of the following discussion.

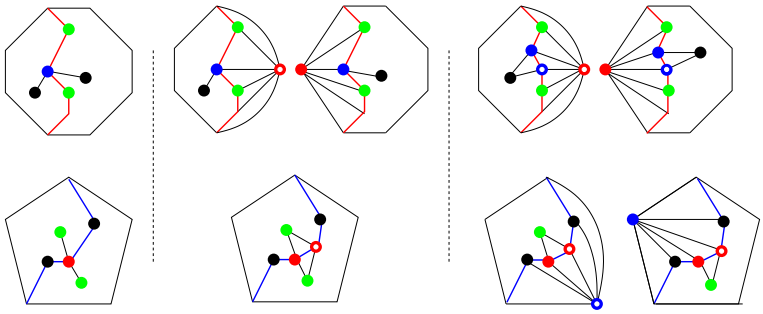

Figure 4.5: Illustration of Sect. 4.3.2. See the discussion there for an explanation of the different panels.

We proceed in the same manner as before. We split the red node n along the red path γ_n then the blue node m along the blue path γ_m.

Remark 4.3.5. *The blue path γ_m must be drawn after the splitting of n. However it can be drawn in the original hemisphere, i.e., before this first splitting, and that, by definition since the blue path must go through both new red nodes consecutively, there is only one way to extend it once we split the red node.*

We first split the red node (upper-left panel) along the red path into a red (left upper-middle panel) and a hollow-red (right upper-middle panel) node. The hemisphere of the blue node changes as shown in the lower-middle panel. Notice that the blue path goes through both red nodes consecutively. Remark 4.3.5 simply means that the blue path can be drawn in the original flower (lower left panel) and that there is a unique way of extending it into the blue path of the lower-middle panel. Next, we split the blue node along the blue path to get the flowers of the right panels.

In Fig. 4.6, we first collapse the edge (red,hollow-red). Note that its flower, drawn in magenta, satisfies the conditions of Remark 4.3.3. Then we collapse the edge (blue,hollow-blue); its flower is drawn in cyan and it also satisfies Remark 4.3.3. The final triangulation is identical to the original one of Fig. 4.5. Note that the blue and cyan paths are identical; the same is true for the red and magenta paths. In other terms, in this case, taking into account Remark 4.3.5, we have $\gamma'_i = \gamma_i, i = \mathrm{R}, \mathrm{L}$.

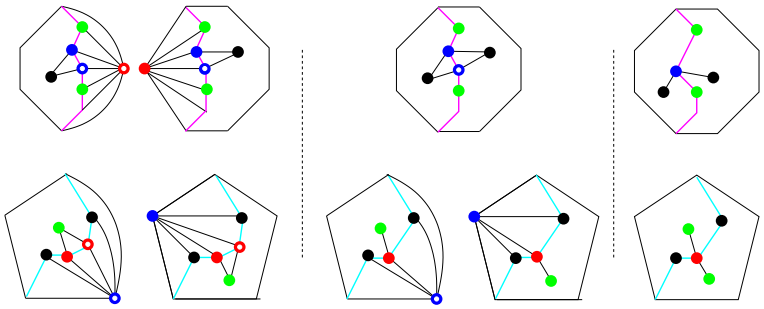

Figure 4.6: Illustration of Sect. 4.3.2. See the discussion there for an explanation of the different panels.

4.3.3 The case $n = m$ and the paths intersect

The case is illustrated in Fig. 4.7 and Fig. 4.8. The first panel shows the hemisphere of the red node. We split it along the red path into a red and a blue node, then we split the blue node along the blue path into a blue and a hollow-blue node.

In Fig. 4.8, we collapse the edge (red,blue) into a single node we call again red; notice that the edge satisfies Remark 4.3.3 and its flower is drawn in magenta. Then we collapse the edge (red,hollow-blue), which again satisfies Remark 4.3.3; its flower is drawn in cyan. We get the same triangulation as the starting one, *i.e.*, splitting first along the red path then along the blue path is equivalent to splitting first along the cyan path then the magenta path.

Clearly, in this case, changing the order of splittings changes the paths: $\gamma'_i \neq \gamma_i$. It is left to the reader to check that this is the only case where such a thing happens (one still needs to check that the trivial cases where the paths are disjoint have this property).

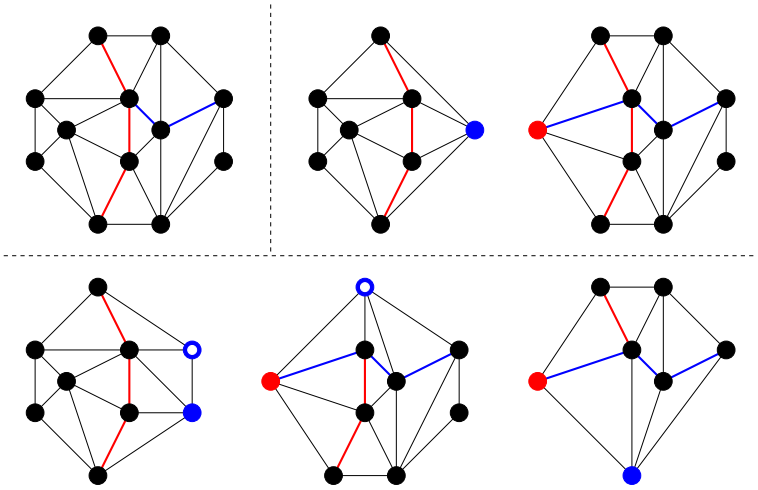

Figure 4.7: Illustration of Sect. 4.3.3. See the discussion there for an explanation of the different panels.

4.4 The strategy

4.4.1 Introduction

Starting from an atom A, we split nodes repeatedly to get a nucleus T. Throughout this process, we keep track of the parent of each node. This naturally leads to the following definition:

Definition 4.4.1. *We say that two nodes x and y of a triangulation T are* related *if they both descend from the same node.*

This notion is actually an equivalence relation; we use it to divide the set of all nodes of the nucleus T into disjoint subsets or equivalence classes. All related nodes of each equivalence class are children of one node in the atom A.

Definition 4.4.2. *Let x be a node of the nuclei T. We denote by $\tilde{x}$ the parent of x in the atom A. We define the* original hemisphere *(or* original flower*) of x as the hemisphere of $\tilde{x}$ in the atom A. It is denoted by $\mathcal{I}(\tilde{x})$.*

From this point on, we will use the tilde to denote nodes and paths in the atom A.

Let x be a node. Assume that we split it into x_1 and x_2. Then we split x_1 into x_3 and x_4 and so on. We get a collection $\{x_i\}_{i=1\ldots r}$ of nodes. By definition, all these nodes are related since they are all children of x. We have the following corollary of Lemma 4.3.4:

Corollary 4.4.3. *The splittings of related nodes commute.*

Proof. We split x into x_1 and x_2. By definition of the splitting move, (x_1, x_2) is an external edge of the triangulation. Recursively, one can show that if x_i, x_j are two related nodes, then (x_i, x_j) is an external edge of the triangulation. But Lemma 4.3.4 shows that (x_i, x_j) must be internal if the splittings are to be non-commuting. □

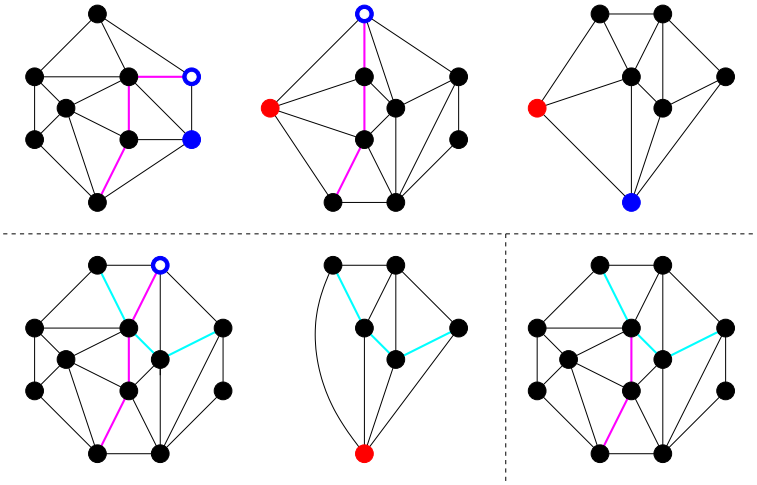

Figure 4.8: Illustration of Sect. 4.3.3. See the discussion there for an explanation of the different panels.

We consider an atom A with t' tetrahedra and n' nodes. Our purpose is to count the number of different sequences of splittings $S(n_k,\gamma_k)\cdots S(n_1,\gamma_1)$ such that $\sum_{i=1}^{k}|\gamma_i| = t - t'$ and $k = n - n'$ with n, t two given constants (representing the number of nodes and tetrahedra of the target nuclei obtained from the atom A). Two sequences are equivalent, *i.e.*, they are the same and therefore should be counted once, if one can be obtained from the other by changing the order of splittings that commute, leading to the same nuclei.

We consider such a sequence $S(n_k,\gamma_k)\cdots S(n_1,\gamma_1)$. Using Lemma 4.2.4, we rearrange the order of the splittings. We obtain an equivalent sequence $S_k\cdots S_1$. This sequence is divided into m subsequences of lengths $k_1, k_2, \ldots, k_m$ where $\sum_i k_i = k$. We consider one of these subsequences with a given index r. Its contains k_r splittings resulting in a total of ℓ_r added tetrahedra. By construction, we have

$$\sum_{r=1}^{m} k_r = n - n' \text{ and } \sum_{r=1}^{m} \ell_r = t - t' \ . \tag{4.1}$$

Some of these k_r splittings might be of related nodes. Without loss of generality (since all splittings within any such subsequence commute), we assume that within each subsequence $r = 1,\ldots,m$, splittings of related nodes are successive. Let X_r be the set of unrelated nodes of the subsequence r, *i.e.*, X_r is the set of equivalence classes in the subsequence r, the equivalence relation being of course the one of Definition 4.4.1. Let $q_r = |X_r|$ be its cardinality. We represent each class of X_r by the parent node in the atom A. Let $\tilde{x}$ represent such a parent. We define $k_r(\tilde{x})$ as the number of splittings of the children of $\tilde{x}$ in the subsequence r, and $\ell_r(\tilde{x})$ as the corresponding number of added tetrahedra. We immediately get

$$\sum_{\tilde{x}\in X_r} k_r(\tilde{x}) = k_r \text{ and } \sum_{\tilde{x}\in X_r} \ell_r(\tilde{x}) = \ell_r \ . \tag{4.2}$$

Remark 4.4.4. *The children of a node $\tilde{x} \in A$ can appear in more than one subsequence of splittings. For each node $\tilde{x}$ of the atom A, we define $k(\tilde{x}) = \sum_r k_r(\tilde{x})$ as the total number of*

splittings of all children of $\tilde{x}$ in every subsequence r, and $\ell(\tilde{x}) = \sum_r \ell_r(\tilde{x})$ as the corresponding number of added tetrahedra.

To count the number of distinct sequences $S_k \cdots S_1$, we first need to count the number of ways the sets X_r can be chosen (Sect. 4.4.2). Then, given a node x in the subsequence r, we need to count the number of ways x and its related nodes in the subsequence r can be split (Sect. 4.4.3).

Remark 4.4.5. *In what follows, several bounds of the form C^ℓ will appear, where C is some constant. I will not use a different notation for every new constant. Instead, all of them will be called C.*

4.4.2 Choice of the sets X_r

Recall that X_r is defined as the set of unrelated nodes split in the subsequence r and q_r is its cardinality. The first subsequence of q_1 unrelated nodes can be chosen randomly among all the nodes of A, thus we get the first binomial $\binom{n'}{q_1} \leq 2^{n'}$. We split these nodes (the number of ways these splittings can be done is counted in the following section). Recall that for every node $\tilde{x} \in X_r, r \geq 1$, $\ell_r(\tilde{x})$ is defined as the total number of added tetrahedra in all splittings of $\tilde{x}$ and its children in the subsequence r.

We come now to X_2. By P3 of Lemma 4.2.4, every one of the second subsequence of splittings does not commute with at least one of the previous subsequence. And we saw that the only case where two splittings do not commute is when each node belongs to the path along which we split the other. This implies that the second subsequence of q_2 unrelated nodes are chosen among the nodes of the paths of the first subsequence. But by construction of the splitting move, the length of a path along which we split is bounded by the number of added tetrahedra, which in our case is given by the sequence $\{\ell_1(\tilde{x})\}_{\tilde{x}\in X_1}$. Thus, the number of choices of the second set X_2 is bounded by

$$\binom{\sum_{\tilde{x}\in X_1} \ell_1(\tilde{x})}{q_2} \leq 2^{\sum_{\tilde{x}\in X_1} \ell_1(\tilde{x})} .$$

Doing the same for every set X_r, using Eq. (4.1) and Eq. (4.2), we see that the number of ways the sets X_r can be chosen is bounded by

$$2^{n'} \cdot 2^{\sum_{r,\tilde{x}} \ell_r(\tilde{x})} \leq 2^{n'+t-t'} .$$

Remark 4.4.6. *Choosing X_r does not fix the choice of the nodes we split in the subsequence r. Saying that $\tilde{x} \in X_r$ means that in the subsequence r, we need to split some of the children of $\tilde{x}$ $k_r(\tilde{x})$ times.*

4.4.3 Choice of the splittings

Fix an index $r \leq m$, consider a node $\tilde{x} \in A$ and let $\{x_i\}$ denote its children in the subsequence indexed r. We want to count the number of ways these children $\{x_i\}$ can be chosen as well as the number of ways of splitting them a total of $k_r(\tilde{x})$ times in such a way that $\ell_r(\tilde{x})$ tetrahedra are added.

We first need to choose which of the children we split; then, for each chosen x_i, we need to choose a splitting path in its flower. The following problems appear:

- P1 Assume that some node x is a child of $\tilde{x}$ in the subsequence r. We split x into $x_{\mathrm{R}}, x_{\mathrm{L}}$. Both these nodes are also children of $\tilde{x}$ and their splitting commutes with that of x, as seen in Sect. 4.3.3. This means that both nodes x_{R} and x_{L} can belong to the same subsequence r, in which case we need to choose among them. Each time we split, the total number of these children increases by one. This means that, if we naively count the number of ways the children $\{x_i\}$ can be chosen, we get a bound of the form $k_r(\tilde{x})!$.
- P2 Assume that we split some node x into $x_{\mathrm{R}}, x_{\mathrm{L}}$. By definition of the splitting move, the degree of one of the two children, say x_{R}, might be larger than the degree of x (by at most one). This implies that the number of ways of splitting x_{R} is larger than that of splitting x (since its flower is larger). Each time we split, the size of the flower might increase. A naive estimate would not work here either.
- P3 Similarly to P2, when a node x is split along a path γ, the degree of every node of γ increases by one. A naive estimate would not work here either.

We start with P2. We discuss P3 in Sect. 4.4.3 and P1 in Sect. 4.4.3. Finally, we prove Theorem 4.1.6 in Sect. 4.4.4.

Solving P2

Consider a node $\tilde{x} \in X_r$ and let $u = k_r(\tilde{x})$ be its multiplicity and $\ell_r(\tilde{x})$ the corresponding number of added tetrahedra. We denote by $S(x_u, \gamma_u) \cdots S(x_1, \gamma_1)$ the successive splittings of its children $\{x_i, 1 = 2 \dots u\}$. We saw in Sect. 4.3.3 that these splittings commute and that the paths are not the same depending on the order in which the splittings are done. We fix an order for the splittings, *i.e.*, we first split x_1 along γ_1, then x_2 along γ_2 and so on and so forth.

The reason P2 arises is that the path γ_i is drawn in the hemisphere of x_i which might be bigger than the original hemisphere of $\tilde{x}$ if it was split in a way that increases its degree. The idea is then to draw all the paths γ_i in the original hemisphere $\mathcal{I}(\tilde{x})$. But by doing so, some paths might need an extension to become splitting paths; what we mean by this is the following: looking at Fig. 4.7, one can draw the blue path in the original hemisphere (the upper-left panel) but one needs to extend it as shown in the lower-middle panel before splitting.

Our strategy is then to draw all paths in the original hemisphere. By Lemma 4.2.3 and what follows (Remark ⋆), this can be done in at most $C^{d(\tilde{x})+\ell_r(\tilde{x})}$ ways, where $d(\tilde{x})$ is the degree of $\tilde{x}$ in the atom A. We clearly have

$$\sum_{\tilde{x} \in A} d(\tilde{x}) = 2\#(\text{edges of } A) \leq 12t' \ . \tag{4.3}$$

Remark ⋆. Lemma 4.2.3 counts the number of ways $\mathcal{I}(\tilde{x})$ can be divided into parts but what we want is to bound the number of ways the paths can be chosen. The purpose of this remark is to show that both bounds are the same. More precisely, we show that if 2 (ordered) collections of paths come from the same division into parts, then splitting along the

first collection gives the same result as splitting along the second and hence, both collections should be counted once.

We consider a valid division of $\mathcal{I}(\tilde{x})$ into $k' = k_r(\tilde{x}) + 1$ parts such that the total length of the boundaries of the parts is $l \leq \ell_r(\tilde{x})$. By C2, the boundary of each part is a closed polygon. Each edge of this boundary can be shared by other parts or it can be an external edge in $\mathcal{I}(\tilde{x})$. We consider a part such that the set of edges on its boundary that are external in $\mathcal{I}(\tilde{x})$ form a connected path (the green part of the second panel of Fig. 4.1 is not such a part, the red one is). Since each part is connected (R1) and since the division is planar (R3), there is always at least one such part. We define γ_1 as the subset of its boundary that is shared by other parts, *i.e.*, that is not on the exterior of $\mathcal{I}(\tilde{x})$. By R4 and the choice of the part in question, γ_1 is connected. We remove this first part from $\mathcal{I}(\tilde{x})$ and we repeat the same discussion to get γ_2. All in all, we obtain an ordered collection of $k_r(\tilde{x})$ connected paths with a total length of l.

However, given a valid division, its is clear that different choices of the parts we remove lead to different paths and that for a given division into parts, the number of choices of the paths can be large. A naive count easily leads to superexponential bounds.

Let $\{\gamma_i\}$ and $\{\gamma'_i\}$ be two such choices of paths obtained from the same division into parts. We consider the explicit example given in Fig. 4.7 and Fig. 4.8: we can see that the first panel of Fig. 4.7 and the last panel of Fig. 4.8 are of the same division of a triangulation but the choice of the paths is different. However, we saw in Sect. 4.3.3 that both choices lead to the same result; more precisely, we saw that each choice of paths corresponds to a choice of the order of the splittings and since the splittings of related nodes commute, both choices lead to the same result.

Applying this result recursively for pairs of splittings of related nodes, one sees that splitting along $\{\gamma_i\}$ or $\{\gamma'_i\}$ leads to the same result, since both choices come from the same division. Another way of seeing this is as follows: recall the move of splitting a node y into $y_\mathrm{R}, y_\mathrm{L}$ along a path γ is defined by splitting $\mathcal{I}(y)$ into two parts along γ and assigning each part with one of two new nodes. If the division of $\mathcal{I}(\tilde{x})$ into parts is given, then the end result will always be the same regardless of the order and choice of the paths.

All in all, the number of distinct choices of paths is indeed given by Lemma 4.2.3.

We draw all paths in the original hemisphere of $\tilde{x}$. These paths are simple but not splitting. We need to count the number of ways we can extend them into splitting paths. In the case of Fig. 4.7 where we only have two paths, there is only one way to extend the blue path, namely to connect it the new red node. In the following lemma, we count the number of ways the extensions can be done.

Lemma 4.4.7. *Consider a subsequence of splittings denoted by r. Let $\tilde{x} \in X_r$ and let $\{\tilde{\gamma}_i\}$ denote all the paths, drawn in its original hemisphere, along which we want to split $\tilde{x}$ in the subsequence r. The number of ways these paths can be extended on their ends into splitting paths is bounded by $C^{k_r(\tilde{x})+\ell_r(\tilde{x})}$, where $k_r(\tilde{x})$ is the number of paths and $\ell_r(\tilde{x})$ is the total length of the extended paths.*

Proof. Assume that all paths $\tilde{\gamma}_i$ are drawn in the original hemisphere $\mathcal{I}(\tilde{x})$ of x (the tilde is used to denote paths in the original hemisphere which might need an extension to become splitting paths). If any of the two ends of a path $\tilde{\gamma}_j$ is not in $\partial\mathcal{I}(\tilde{x})$, then that path needs to

be extended before splitting along it. We denote the extended path by γ_j. Of course, γ_j can no longer be drawn in the original hemisphere of x. The purpose of this proof is to count the number of ways all these extensions can be done.

Consider a node $z \in \mathrm{int}\mathcal{I}(\tilde{x}) \subset A$; assume that z is the end of at least one of the paths $\{\tilde{\gamma}_i\}$; let $J = J(z)$ denote the set of indices j such that $z \in \tilde{\gamma}_j$ ($\tilde{\gamma}_j$ contains z) and K the subset of J such that z is an end of $\tilde{\gamma}_k$ for every $k \in K$ ($\tilde{\gamma}_k$ ends at z). All the paths $\tilde{\gamma}_k, k \in K$ need an extension at their end z. We look at the problem from the point of view of the hemisphere of z (in the atom A). Before any splitting, since $z \in \mathrm{int}\mathcal{I}(\tilde{x})$, then $\tilde{x} \in \mathrm{int}\mathcal{I}(z)$. Let $j \in J$ and let $x_{j\mathrm{R}}, x_{j\mathrm{L}}$ denote the two new nodes obtained from splitting x_j along γ_j. By definition of the move split-a-node-along-a-path, both new nodes as well the edge connecting them are in $\mathrm{int}\mathcal{I}(z)$ (since $(z, x_{j\mathrm{R}})$ and $(z, x_{j\mathrm{L}})$ are both internal edges and $(z, x_{j\mathrm{R}}, x_{j\mathrm{L}})$ is an internal face). Furthermore, after all splittings are done, $\mathcal{I}(z)$ remains a 2d triangulation. This means that the node $\tilde{x}$, which before any splitting was in the interior of the hemisphere of z, is replaced in this hemisphere $\mathcal{I}(z)$, after all splittings are done, by a triangulation of the nodes $\{x_{j\mathrm{R}}, x_{j\mathrm{L}}, j \in J\}$. Deciding on a set of extensions for all $\tilde{\gamma}_k, k \in K$ is then equivalent to picking a triangulation of the nodes $\{x_{j\mathrm{R}}, x_{j\mathrm{L}}, j \in J\}$: $\tilde{\gamma}_k$ is extended to $x_{j\mathrm{R}}$ on its z end if and only if $(x_{j\mathrm{R}}, x_{k\mathrm{R}}, x_{k\mathrm{L}})$ is a triangle in the flower of z. The easiest non-trivial example is shown in Fig. 4.9 and Fig. 4.10. The possible number of these triangulations is bounded by $C^{|J|}$. Multiplying all these bounds for every such node z leads to a bound of the form $C^{\sum_z |J(z)|}$. We can then write

$$\sum_z |J(z)| = \sum_z \sum_{\tilde{\gamma}_i} \chi_{\tilde{\gamma}_i}(z) \ ,$$

where $\chi_{\tilde{\gamma}_i}(z) = 1$ if and only if the path $\tilde{\gamma}_i$ contains the node z. Changing the order of the two sums, and since the total number of nodes in all (extended) paths (also counting the multiplicities of the edges) is $\ell_r(\tilde{x}) + k_r(\tilde{x})$, we get a bound of the form $C^{k_r(\tilde{x})+\ell_r(\tilde{x})}$ for the total number of possible extensions of all paths. □

Remark 4.4.8. *Note that in Lemma 4.4.7, given a subsequence of splittings r and a node $\tilde{x} \in X_r$, we count the number of ways of extending on their ends all paths $\tilde{\gamma}_i$ of $\mathcal{I}(\tilde{x})$ along which we split in the subsequence r. But we saw that a node $\tilde{x}$ can belong to more than 1 subsequence of splittings. Let $\tilde{\gamma}_j$ be a path in in $\mathcal{I}(\tilde{x})$ along which we split $\tilde{x}$ (actually, we split one of its children) in another subsequence $r' > r$. Such paths might also need an extension on their ends. The number of ways this can be done for all such paths and for all later subsequences $r' > r$ is counted in Sect. 4.4.3.*

Solving P3

The above gives a solution to the problem P2. We now look at P3. We consider two nodes x, y. We first split x into $x_{\mathrm{R}}, x_{\mathrm{L}}$ along γ_x and we assume $y \in \gamma_x$. This increases the degree of y by one. Looking at the flower $\mathcal{I}(y)$, we see that the splitting replaces x by the edge $(x_{\mathrm{R}}, x_{\mathrm{L}})$. Next, we want to split y along γ_y, which is drawn in $\mathcal{I}(y)$ after the splitting of x. We will denote by $\mathcal{I}(\tilde{y})$ the original hemisphere of y before x is split and by $\tilde{\gamma}_y \subset \mathcal{I}(\tilde{y})$ the path obtained by identifying x_{R} with x_{L} in $\mathcal{I}(y)$, *i.e.*, $\tilde{\gamma}_y$ is the path drawn in the original hemisphere of y before splitting x. There are 4 cases to consider:

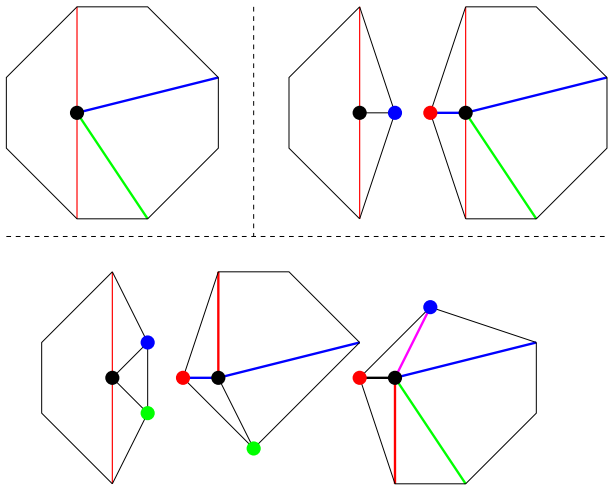

Figure 4.9: Illustration of the proof of Lemma 4.4.7. The upper-left panel shows the hemisphere of the red node. We split along the red path into a new node we also call red (upper-middle) and a blue node (upper-right). Next we want to split the blue node along the blue path. Notice that the blue path of the upper-left panel needs an extension on the black node end and that there is only one way to extend it by connecting the black node to the red one. We split the blue node along the extended blue path into a blue node (lower middle) and a green one (lower right). Next, we want to split the green node along the green path into a green and a brown node. However, there are now two possible ways of extending the green path: the first towards the blue node, shown in magenta; the second towards the red node, shown in black.

1. The path γ_y avoids both new nodes x_R, x_L. In this case, it is clear that $\gamma_y = \tilde{\gamma}_y$ can be drawn in the original hemisphere of y. Splitting x along γ_x does not change $\tilde{\gamma}_y$.
2. The path γ_y contains only one of the the two new nodes, say x_R. In this case, γ_y is obtained from $\tilde{\gamma}_y$ by simply renaming x into x_R. And so we have $\gamma_y = \tilde{\gamma}_y$.
3. The path γ_y contains the edge (x_R, x_L). This is the case of Sect. 4.3.2. The two paths γ_y and $\tilde{\gamma}_y$ are not the same. But $\tilde{\gamma}_y$ is still simple and γ_y is obtained from it by removing x and inserting the edge (x_R, x_L) in its stead.
4. The path γ_y contains both nodes x_R, x_L but not the edge (x_R, x_L). This is the only case where splittings do not commute. $\tilde{\gamma}_y$ can be drawn in the original hemisphere of y before splitting x, but it is no longer simple. It can contain a loop or the same edge twice (see the blue path of the lower right panel of Fig. 4.3). In this case, splitting x unfolds the double edge or the loop.

We see that in all four cases, the path $\tilde{\gamma}_y$ can be drawn in the original hemisphere $\mathcal{I}(\tilde{y})$ before splitting x; it can be non-simple (the fourth case) and it might need an extension (the third case). Splitting x adds a new edge to the flower $\mathcal{I}(\tilde{y})$, namely (x_R, x_L). Using this new edge, we can extend and unfold $\tilde{\gamma}_y$ into γ_y. We need to count the number of ways this can be done. The simplest example is illustrated in Fig. 4.11.

Fig. 4.11 shows that in the case where we only have two splittings (one for x, one for y), the number of ways of choosing the paths γ_x in $\mathcal{I}(x)$ and γ_y in $\mathcal{I}(y)$ is bounded by four

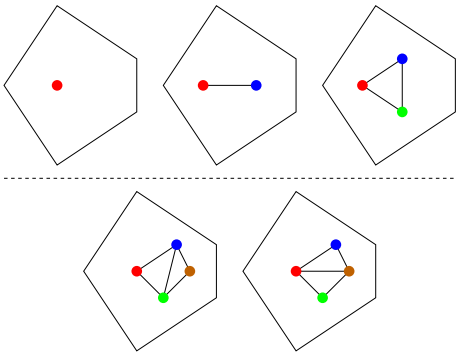

Figure 4.10: Illustration of the proof of Lemma 4.4.7. In this figure, we show the effect of the splits of Fig. 4.9 from the point of view of the hemisphere of the black node. The upper-left panel is before any splitting. It simply shows the red node in the hemisphere of the black node. Splitting the red node into a red node and a blue node along the red path gives the upper-middle panel (the paths colored red blue and green are those shown in Fig. 4.9). Splitting the blue node into a blue and a green node along the blue path gives the upper-right panel: there is only one way of extending the blue path, by connecting the black and red nodes. By definition of the splitting move, this means that (red,black,blue,green) is a tetrahedron, hence the triangle (red,blue,green) in the upper-right panel. Finally, we split the green node along the green path. But before doing so, one needs to extend the green path and there are two possible ways: connecting to the blue node (the magenta edge) gives the lower-left panel, since (black,blue,green,brown) is a tetrahedron; connecting to the red node on the other hand (the black edge) gives the lower-right panel. Each possibility corresponds to a triangulation of the 4 nodes black,blue,green,brown.

times the number of ways of choosing γ_x in $\mathcal{I}(x)$ and $\tilde{\gamma}_y$ in $\mathcal{I}(\tilde{y})$. More importantly, we can see how the extension of $\tilde{\gamma}_y$ into γ_y is done:

- The path γ_x contains y and it is simple (since we split along it). Let a, b be the two nodes adjacent to y in γ_x (the green and brown nodes of Fig. 4.11). By definition, both nodes are present in $\mathcal{I}(\tilde{y})$.
- Splitting x along γ_x replaces the edges $(x,a),(x,b)$ with a triangulation of four nodes $x_{\mathrm{R}}, x_{\mathrm{L}}, a, b$. The choice of the triangulation is fixed by the choice of γ_x.
- If any (or both) of the two edges $(x,a),(x,b)$ is part of $\tilde{\gamma}_y$, then it must be replaced in γ_y by a path in this new triangulation. This is true for **only** these two edges. For instance, looking at the second and third panels, we see that we have no choice but to replace the edge (black,red) of $\tilde{\gamma}_y$ with the edge (black,hollow-red) of γ_y, since it is not one of the two aforementioned edges.

We generalize this result to the case where we have u simple paths $\{\gamma_{x,i}\}$ in $\mathcal{I}(x)$ containing y and v paths (not necessarily simple) $\{\tilde{\gamma}_{y,j}\}$ in $\mathcal{I}(\tilde{y})$ containing x. Let $\{a_i\}$ denote all nodes adjacent to y in all the paths $\{\gamma_{x,i}\}$. Since these paths are simple, there are at most $2u$ such nodes a_i. By definition of the splitting move, the flower $\mathcal{I}(y)$ is obtained from $\mathcal{I}(\tilde{y})$ by replacing all edges $\{(x,a_i)\}$ with a triangulation of the nodes $\{x_i\},\{a_i\}$. The number of these nodes is at most $3u$. The triangulation in itself is determined by the choice of the u

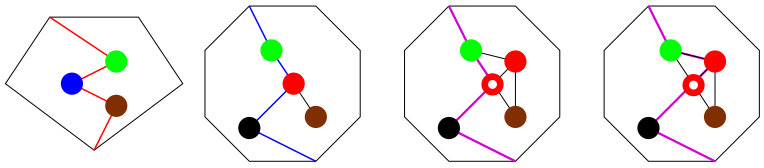

Figure 4.11: We illustrate a simple case of Sect. 4.4.3. The first panel shows the flower of the red node x; the red path is γ_x and the blue node is y. The second panel shows the original flower of y before x is split along the red path. The blue path is $\tilde{\gamma}_y$. Note that the green (brown) disc in both flowers represents the same node. Also note that the blue path $\tilde{\gamma}_y$ goes through the green node (it can also go though both green and brown nodes). We split x along the red path γ_x into a red and a hollow-red node. The last two panels show how γ_y, represented by the magenta path, is obtained: it can be the same as $\tilde{\gamma}_y$ (Case 2 of Sect. 4.4.3) or it can contain an additional edge. If the blue path $\tilde{\gamma}_y$ contains both the green and brown nodes, then γ_y can be obtained from it in four different ways.

paths $\{\gamma_{x,i}\}$ and the extensions on their ends (see Lemma 4.4.7). We now look at the v paths $\{\tilde{\gamma}_{y,j}\}$. They divide $\mathcal{I}(\tilde{y})$ into $v+1$ parts. It is possible that the paths $\{\tilde{\gamma}_{y,j}\}$ no longer form a division of $\mathcal{I}(y)$ into parts. We extend and unfold them by dividing the new triangulation into parts. By Lemma 4.2.3, and since the new triangulation has at most $3u$ nodes, this can be done in at most $C^{\ell'+u}$ ways, where ℓ' is the total number of added edges to the paths, *i.e.*, $\ell' = \sum_j |\gamma_{y,j}| - |\tilde{\gamma}_{y,j}|$.

We summarize this section. We denote by P3 the following problem: we want to count the number of ways of choosing γ_x in $\mathcal{I}(x)$ and γ_y in $\mathcal{I}(y)$. This number clearly depends on the size of each hemisphere. But if $y \in \gamma_x$, then splitting x increases the size of $\mathcal{I}(y)$. If we consider $n - n'$ successive splittings, naively counting the choice of the paths leads to superexponential bounds.

We solve P3 by saying that all paths are drawn in the original hemispheres in the atom A but the paths are no longer simple; they may contain loops and double edges (however, there is at least one simple path). Each splitting of a node leads to some extensions and unfolding of these paths. We need to count the number of ways these paths are chosen and the number of ways they are extended and unfolded.

Let x be a node and let $\{\gamma_{x,i}\}$ be the set of paths (not necessarily simple) dividing $\mathcal{I}(x)$ into parts. We denote by $\ell(x) = \sum_i |\gamma_{x,i}|$ the total length of these paths (each edge is counted according to its multiplicity). By Lemma 4.2.3 and Remark $\star$, these paths can be chosen in $C^{\ell(x)+d(x)}$ ways, where $d(x)$ is the degree of x. Note that Lemma 4.2.3 is still valid for non-simple paths and the choice of the paths still has the correct bound.

We consider all the simple paths of $\{\gamma_{x,i}\}$ and we split x along them (the other non-simple paths are left for later subsequences of splittings. By the time we want to split along them, they will have become simple). Let $\ell_r(x)$ denote the total length of these simple paths (the index r refers to the current subsequence of splittings). Let y be an internal node of $\mathcal{I}(x)$

and let $u_r(x,y)$ be the number of simple paths $\gamma_{x,i}$ containing y. Clearly, we have

$$\sum_{y\in\mathcal{I}(x)} u_r(x,y) \leq 2\ell_r(x) \ .$$

The flower $\mathcal{I}(\tilde{y})$ of y (before splitting x) is also divided into parts by paths. Let $\{\tilde{\gamma}_{y,j}\}$ denote these paths. They can be chosen in $C^{\ell(y)+d(\tilde{y})}$ ways, where $d(\tilde{y})$ is the degree of y before splitting x. We saw that once the splittings of x along all the simple paths of $\{\gamma_{x,i}\}$ are done, then the paths $\{\tilde{\gamma}_{y,j}\}$ need to be extended by $\ell'_r(x,y)$. The number of ways these extensions can be done is bounded by $C^{\ell'_r(x,y)+u_r(x,y)}$. This holds true for every internal node y of $\mathcal{I}(x)$. Multiplying this bound for every $y \in \mathcal{I}(x)$, defining $\ell'_r(x) = \sum_y \ell'_r(x,y)$, we get the following result:

Lemma 4.4.9. *We assume that the (original) hemisphere of every node of the atom A is divided into parts by paths. Consider a node $\tilde{x}$ of X_r. We split it along all simple paths in its (original) hemisphere and let $\ell_r(\tilde{x})$ be their total length. These splittings of $\tilde{x}$ extend all paths of every hemisphere of A containing $\tilde{x}$ by a total of $\ell'_r(\tilde{x})$ edges. The number of ways these extensions can be chosen is bounded by $C^{\ell'_r(\tilde{x})+\ell_r(\tilde{x})}$.*

Furthermore, we have

$$\sum_{r,\tilde{x}} \ell'_r(\tilde{x}) \leq t - t' \ , \tag{4.4}$$

since the total length of all extended path cannot exceed $t - t'$.

P2 again

This section deals with the second part of the problem P2 raised in Remark 4.4.8. As in the previous sections, we start by illustrating the problem with an explicit example, shown in Fig. 4.12.

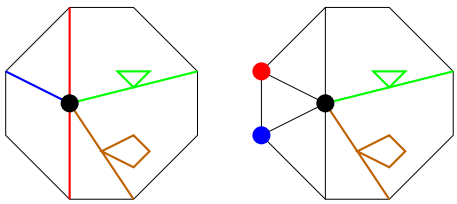

Figure 4.12: Illustration of Sect. 4.4.3. See the discussion there for an explanation of the different panels.

We consider a node $\tilde{x}$ and we assume there are 4 paths in its hemisphere: the blue and red paths are simple whereas the green and brown are not. This is shown in the first panel of Fig. 4.12. We split $\tilde{x}$ into three nodes, say red blue and yellow, along the blue and red paths (after extending the blue path as seen in Lemma 4.4.7). Assume that the part of $\mathcal{I}(\tilde{x})$ containing the brown and green paths is associated with the yellow node. The second panel of Fig. 4.12 shows the hemisphere of this yellow node. The red and blue nodes are the children of $\tilde{x}$ obtained after the splittings along the red and blue paths. It is clear that at

least one of the 2 paths of the second panel (or even both of them) needs an extension on its end.

More generally, consider a node $\tilde{x} \in X_r$, assume that all subsequences of splittings $r'' < r$ are done and let $\{\gamma_{x,i}\}$ denote all remaining[2] paths of $\mathcal{I}(\tilde{x})$. These paths $\{\gamma_{x,i}\}$ are of 2 types: the simple ones are the paths along which we split $\tilde{x}$ and its children in the subsequence r. The non-simple ones are those reserved for later subsequences $r'' > r$. Let $\ell_r(\tilde{x})$ denote the total length of the simple ones. Lemma 4.4.7 shows that the number of ways of extending these simple paths on their ends is bounded by $C^{\ell_r(\tilde{x})}$ ways (since $k_r(\tilde{x}) \leq \ell_r(\tilde{x})$). We split along these extended simple paths. After these splittings are done, some of the non-simple paths will need an extension on their ends, as seen in Fig. 4.12. We need to bound the number of ways these extensions can be done.

This discussion is very similar to that of the previous section. Using the same arguments of Sect. 4.4.3, one can show that the number of ways these extensions can be done is bounded by $C^{\ell''_r(\tilde{x})+\ell_r(\tilde{x})}$, where $\ell''_r(\tilde{x})$ is the total length of the extensions. Thus, we have the following result:

Lemma 4.4.10. *We assume that the (original) hemisphere of every node of the atom A is divided into parts by paths. Consider a node $\tilde{x}$ of X_r. We split it along all simple paths in its (original) hemisphere and let $\ell_r(\tilde{x})$ be their total length. These splittings of $\tilde{x}$ extend all remaining non-simple paths of $\mathcal{I}(\tilde{x})$ by a total of $\ell''_r(\tilde{x})$ edges. The number of ways these extensions can be chosen is bounded by $C^{\ell''_r(\tilde{x})+\ell_r(\tilde{x})}$.*

Clearly, the total lengths of all extensions is bounded by $t - t'$ and we have

$$\sum_{r,\tilde{x}} \ell''_r(\tilde{x}) \leq t - t' \ . \tag{4.5}$$

This bound can also be obtained in a simpler way: let z be an internal node of $\mathcal{I}(\tilde{x})$ such that z is the end of at least one of the non-simple paths (associated with the sequences $r'' > r$); for example, z is the black node of Fig. 4.12. Let $u_r(x, z)$ denote the number of simple paths going through z in $\mathcal{I}(\tilde{x})$; by definition, we have $\sum_z u_r(x, z) \leq 2\ell_r(x)$. Splitting along these $u_r(x, z)$ paths creates $2u_r(x, z)$ new children of $\tilde{x}$, denoted by $\{x_i\}$. It is easy to see that the extensions should be done exclusively along the edges (z, x_i) and that it is sufficient to assign a multiplicity m for each of these edges: the edge (z, x_i) is present $m(z, x_i)$ times in the extensions. Doing this for every z, since all these multiplicities sum up to the total length of the extensions $\ell''_r(x)$, and since their numbers $2u_r(x, z)$ sum up to $4\ell_r(x)$, using Lemma 4.2.1, we get the bound $C^{\ell''_r(\tilde{x})+\ell_r(\tilde{x})}$.

Solving P1

We consider a node $\tilde{x} \in X_r$. We split its hemisphere $\mathcal{I}(\tilde{x})$ into $k_r(\tilde{x}) + 1$ parts by drawing $k_r(\tilde{x})$ paths, not necessarily simple, of total length smaller than $\ell_r(\tilde{x})$. We still need to address the first problem P1, namely that knowing the paths γ_i along which we split does not always tell us which of $\tilde{x}$'s children we actually split. Recall that an edge might belong to more than one path and in particular, two paths might be identical.

[2] If $\tilde{x}$ belongs to some $X_{r''}, r'' < r$, then it was split along some paths of $\mathcal{I}(\tilde{x})$ in the subsequence $r'' < r$.

The first scenario here is when all paths are distinct, *i.e.*, any two paths differ by at least one edge. The discussion can be reduced to the case where we have one pair of paths, since the splittings are successive; we consider a node x and two paths $\gamma_1 \neq \gamma_2$ in its hemisphere $\mathcal{I}(x)$. We split x into x_R and x_L along γ_1 and we ask the question: knowing that the next splitting is along γ_2, which of the two new nodes should be split? By construction, the paths cannot cross one another and γ_1 splits $\mathcal{I}(x)$ in two parts. The only common edges of these two parts are those of γ_1. Since $\gamma_2 \neq \gamma_1$, only one of these two parts of $\mathcal{I}(x)$ contains γ_2 and the answer to the above question is known: we split the child which is associated with the part of $\mathcal{I}(x)$ containing γ_2. In other terms, if all paths are distinct, then it is sufficient to list only the paths along which we split.

Consider now a node x and two identical paths $\gamma_1 = \gamma_2$ in its hemisphere. We first split x into x_R and x_L along γ_1. But now γ_2 can be drawn in both new hemispheres. Fig. 4.13 shows that, in this case, splitting x_R along γ_2 yields the same result as splitting x_L along γ_2 and as a consequence, the choice of which child is split is irrelevant.

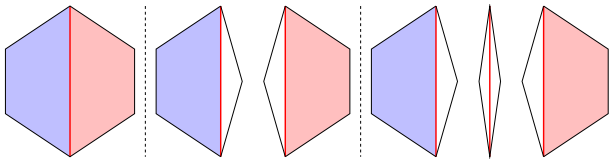

Figure 4.13: The first panel shows the hemisphere of x. We split x along the red path into x_R and x_L. The hemispheres are shown in the second set of panels. Splitting x_R or x_L along the red paths leads to the same result shown in the third set of panels.

4.4.4 Putting everything together

We summarize the results of the previous sections to prove Theorem 4.1.6:

1. We consider every node $\tilde{x}$ in the atom A. For each such node, we pick two integers $k(\tilde{x})$ and $\ell(\tilde{x})$. These integers must verify

$$\sum_{\tilde{x}} k(\tilde{x}) = n - n' \text{ and } \sum_{\tilde{x}} \ell(\tilde{x}) = t - t' \ .$$

   By Lemma 4.2.1, each of the two sequences can be chosen in C^t ways, since $n \leq 4t$.
2. For every node $\tilde{x}$ of A, We divide its flower into $k(\tilde{x})+1$ parts such that the total length of the boundaries is bounded by $\ell(\tilde{x})$. The number of ways this can be done is given by Lemma 4.2.3 and is bounded by $C^{d(\tilde{x})+\ell(\tilde{x})}$, where $d(\tilde{x})$ is the degree of $\tilde{x}$, verifying Eq. (4.3).
3. Each division into parts gives us several equivalent ordered collection of paths (not necessarily simple). We pick one such ordered collection (this is equivalent to fixing an order for the splittings of $\tilde{x}$'s children, see Remark ⋆ of Sect. 4.4.3).
4. We choose the sets X_r. In Sect. 4.4.2, we saw that this can be done in C^t ways, since $n' < n' < 4t$.

5. Let $q = \sum_r |X_r| \leq n - n'$. We choose 2 sequence of q integers $\ell'_r(\tilde{x}), \ell''_r(\tilde{x})$ verifying Eq. (4.4) and Eq. (4.5) respectively. By Lemma 4.2.1, they can be chosen in C^t ways.
6. We start with the first subsequence of splittings $r = 1$. We consider the nodes $\tilde{x}$ as prescribed by the choice of X_1.
   - For every $\tilde{x}$, we only consider the paths in its flower $\mathcal{I}(\tilde{x})$ that are simple (the other non-simple paths are left for later subsequences of splittings $r > 1$). Let $k_1(\tilde{x})$ and $\ell_1(\tilde{x})$ denote the number of these simple paths and their total length respectively.
   - The number of ways of extending these paths into splitting paths is given by Lemma 4.4.7 and is bounded by $C^{k_1(\tilde{x})+\ell_1(\tilde{x})}$.
   - In Sect. 4.4.3, we saw that once the choice of the paths is done, then the choice of which of $\tilde{x}$'s children are split is fixed. We split along these extended paths.
   - These splittings of $\tilde{x}$ and its children extend and unfold the paths of every hemisphere of A by a total of $\ell'_1(\tilde{x}) + \ell''_1(\tilde{x})$. The number of ways this can be done is given by Lemma 4.4.9 and Lemma 4.4.10 and is bounded by $C^{\ell_1(\tilde{x})+\ell'_1(\tilde{x})+\ell''_1(\tilde{x}}$, where $\ell'_1(\tilde{x}), \ell''_1(\tilde{x})$ are the integers chosen in the fifth point.

   Doing this for every $\tilde{x} \in X_1$ ends the first subsequence $r = 1$ of splittings.
7. We start the next subsequence $r \geq 2$ of splittings . We consider the nodes $\tilde{x}$ of X_r. Some of the non-simple paths we ignored in the previous subsequences $r' < r$ were unfolded and extended into simple paths by the previous splittings. We consider all simple paths (including the new ones). Note that their order is still fixed by the above choice of the third point. At this point, there is no freedom on the choice of these paths. They were unfolded and extended in the previous splittings. Let $k_r(\tilde{x})$ and $\ell_r(\tilde{x})$ denote the number of these simple paths and their total length respectively. We extend these paths at their ends (Lemma 4.4.7) and we split along them. The number of ways this can be done is again bounded by $C^{k_r(\tilde{x})+\ell_r(\tilde{x})}$. The splittings of each $\tilde{x} \in X_r$ and its children extend and unfold all remaining paths of every hemisphere by a total of $\ell'_r(\tilde{x}) + \ell''_r(\tilde{x})$. The number of ways this can be done is given by Lemma 4.4.9 and Lemma 4.4.10 and is bounded by $C^{\ell_r(\tilde{x})+\ell'_r(\tilde{x})+\ell''_r(\tilde{x})}$, where $\ell'_r(\tilde{x}), \ell''_r(\tilde{x})$ are the 2 sequences of integers of the fifth point.
8. We repeat the last step until no more paths are left.

Putting everything together, using Eq. (4.1), Eq. (4.2), Eq. (4.3), Eq. (4.4), Eq. (4.5), and since $n' < n < 4t$ and $t' < t$, Theorem 4.1.6 and hence Theorem 4.1.2 follow.

Chapter 5

3D Topological Glass

5.1 Introduction

In this last chapter, we return to the problem of constructing a model of a three-dimensional (3d) topological glass. Similarly to the 2d model (see [5], [4] and the first chapter of this thesis), the phase space we consider is the set of all simplicial piecewise-linear decompositions of the sphere S^3, simply known as *triangulations* or 3-spheres. The elementary moves are given by the two Pachner moves in three dimensions that conserve the number of particles; they are called the 2-3 flip and 3-2 flip. The energy is local, and the contribution $f(v)$ of each node v is a function of its local neighborhood $\mathcal{I}(v)$ we call flower: $f(v) = f(\mathcal{I}(v))$. Finally, the dynamics is given by a simple Metropolis algorithm with one free parameter we call the temperature.

Contrary to 2-spheres, the topological properties of triangulations of S^3 are far from being well understood. In the previous two chapters, we discussed in detail one of these properties, namely the number of 3-spheres with t tetrahedra. We defined a subset of small, elementary triangulations we called atoms and we showed that if the number of these atoms is exponentially bounded, then so is the number of all 3-spheres. Furthermore, we expressed our belief that the number of these atoms can be bounded by the number of knots and links with a given complexity, which is known to be exponential.

Another interesting question concerns the reducibility of the phase space under the 2-3 and 3-2 flips: Pachner's result [40] applied to S^3 states that any two 3-spheres can be transformed into one another by a series of four types of flips called Pachner moves. However, we only consider two of these moves, namely the 2-3 and 3-2 flips which conserve the number of particles. Dougherty, Faber and Murphy [38] showed that the set of all 3-spheres is reducible under these two flips, by constructing a triangulation such that no 2-3 or 3-2 flip is possible. Furthermore, they showed that their construct cannot be realized geometrically in S^3, which points up another interesting property of 3-spheres: the vast majority of (topological) triangulations of S^3 are not geometric, *i.e.*, cannot be drawn in S^3 (or $\mathbb{R}^3$) (see [42, 35] and [43, 44]).

Using some ideas by Santos [39], we show any two Delaunay triangulations with n nodes, *i.e.*, any two topological triangulations of S^3 that can be realized geometrically as the dual of a Voronoi decomposition of point sets of n particles (with coordinates), are connected by at most $O(n^5)$ 2-3 and 3-2 flips.

We then proceed to constructing two different models of 3d topological glasses. The first example is motivated by [45] and is the simplest possible generalization of the 2d model introduced in [5, 4] and discussed in the first chapter of this thesis: we say that the energy contribution of a node v is a function of its degree, which is the number of nodes of the flower $\mathcal{I}(v)$. Since the 2-3 and 3-2 flips change the degree of only two nodes (contrary to the T1 flip in two dimensions which changes the degree of four nodes), this model turns out to be even simpler, in a way, than its 2d counterpart.

The second model we introduce makes full use of the additional degrees of freedom of 3d flowers, which are, by definition, triangulations of S^2. We start by fixing the ground state of the system as a crystal, *i.e.*, a regular filling of space, which is also a triangulation. The crystal we choose is the one such that the flower of every node is a tetrakis hexahedron: it has 14 vertices, 6 of which have a degree of 4 and the remaining 8 have a degree of 6. We show that it is possible to fill space with particles such that the Voronoi cell of each particle is a tetrakis hexahedron. Note that such a crystalline structure is known to appear in fluorite systems [46] and rarely in diamonds [47]. Considering any triangulation, the energy contribution of a node is then a measure of how different its flower is from a tetrakis hexahedron. Such a definition of the energy is very difficult to implement on a computer, so we simplify it by characterizing a flower $\mathcal{I}(v)$ of a node v by a sequence of integers $\{e_k, k \geq 3\}$, where e_k is the number of vertices in $\mathcal{I}$ with degree k. The ground state sequence is $e_4^\star = 6$, $e_6^\star = 8$, $e_k^\star = 0$ $\forall k \notin \{6, 8\}$ and the energy contribution of a node n is simply the euclidean distance $\|\mathbf{e}(n) - \mathbf{e}^\star\|$.

The dynamics is given by a simple Metropolis algorithm with the temperature being the only free parameter. We observe that the system's dynamics goes through a tremendous slowing down as it approaches its stationary state at low temperature. Using some additional ideas (improving the algorithm, considering less restrictive energy forms and defining a new equivalent set of elementary moves), we can speed things up considerably (by several tens of times) but to little effect and the system never reaches its stationary state in a reasonable amount of time.

As a consequence, our results are only numerical. We observe that the energy relaxation undergoes several polynomial regimes as it approaches its stationary state value. One interesting notion we would like to understand is that of an elementary defect: in the 2d model, we showed that the appropriate definition of an elementary defect is as a single isolated node with an incorrect flower. In the 3d model, it is easy enough to see that having a single node with an incorrect flower surrounded by nodes with tetrakis hexahedra as flowers is topologically impossible. So one would expect that elementary defects in 3d are clusters of several nodes with an incorrect flower each. But the simulations seem to indicate something completely different: even with less than 10% of the nodes having an incorrect flower, they do not cluster; instead they seem to form lines of defects, which shows that the behavior of this type of 3d models of topological glasses is much more complex than we originally thought.

5.1.1 Notations and definitions

We recall some notations and definitions we introduced in the previous chapters. We start with the notion of f-vector:

Any triangulation of S^3 can be characterized by four integers called the *f-vector*: it has n nodes, E edges, F faces and t tetrahedra. They must satisfy the following two relations:

- Euler's theorem for decompositions of S^3 reads: $n - E + F - t = 0$.
- By definition of a triangulation, each face is shared by exactly two tetrahedra and each tetrahedron has exactly four faces, thus $4t = 2F$.

Combining these two relations, we get:

$$F = 2t \qquad \text{and} \qquad E = n + t \ , \tag{5.1}$$

leaving only two free parameters, n the number of nodes and t the number of tetrahedra.

Next, we recall some notations:

- Let n_1, n_2 be two adjacent nodes. We denote by (n_1, n_2) the edge connecting them. Similarly, (n_1, n_2, n_3) denotes the face having the nodes $n_i, i = 1, 2, 3$ as corners.
- Let n be a node and e an edge not having n as an end. Then (n, e) denotes the face containing both e and n. If such a face exists, then we say that the node n is *opposite* to the edge e and reciprocally, the edge e is *opposite* to the node n.
- More generally, let s, s' be two simplices of dimension $d + d' \leq 3$ (nodes are simplices of dimension 0, edges and faces have dimension 1 and 2 respectively). Assume that $s \cap s' = \emptyset$, *i.e.*, that both simplices do not share a corner. Then (s, s') denotes the simplex of dimension $d + d'$ having both s and s'. If such a simplex exists, then s and s' are called *opposite one to the other*.

We now recall the previously introduced notions of flowers and degrees:

Definition 5.1.1. *We consider a triangulation T of S^3. Let s be a simplex of T of dimension $d = 0, 1, 2$. We define its* flower *$\mathcal{I}(s)$ as the set of all simplices s' of T of dimension $d' = 2 - d$ such that (s, s') is a simplex of T. By definition of a triangulation, $\mathcal{I}(s)$ is a triangulation of the d'-dimensional sphere $S^{d'}$.*

In detail, we have:

- Let f be a face. Its flower is the set of all nodes n' opposite to it. It is a triangulation of S^0, *i.e.*, it consists of exactly two nodes. This simply means that a face is shared by exactly two tetrahedra.
- Let e be an edge. Its flower $\mathcal{I}(e)$ is the set of all edges e' opposite to it. It is a triangulation of S^1, *i.e.*, it is a closed polygon. The degree of e, denoted by $\deg(e)$, is simply the number of edges of this polygon or equivalently, the number of faces having e as an edge.
- Let m be a node. Its flower $\mathcal{I}(m)$ is the set of all faces f' opposite to it. It is a triangulation of S^2 with n_f faces, n_e edges and n_v nodes. The degree $\deg(m)$ of the node m is defined as n_v or equivalently, the number of edges having m as an end. Note that using Euler's theorem on S^2 and a simple geometric relation, we have $n_f = 2\deg(m) - 4$ and $n_e = 3\deg(m) - 6$.

Note that in the literature, flowers are more commonly referred to as *stars*.

Remark 5.1.2. *In this chapter, and contrary to the previous two, we only consider 3-spheres. In other terms, the triangulations we consider do not have a boundary and all nodes, edges and faces are internal.*

Remark 5.1.3. *It is easy enough to see that the minimal degree for an edge is 3 and for a node is 4.*

5.2 The phase space, the elementary moves and the dynamics

The phase space is the set of all simplicial piecewise-linear decompositions of S^3, simply called triangulations. In the previous two chapters, we discussed in detail the question of its size. In this section we introduce the elementary moves of our models as well as the general form of the energy and we discuss some additional properties of triangulations of S^3.

5.2.1 Topological, geometric and Delaunay triangulations

It is known that, contrary to the 2d case, not all topological triangulations can be geometrically realized. Consider a (topological) triangulation of S^3. We say that it is *geometric* if it can be realized as the boundary of a convex four-dimensional polytope. In two dimensions, Steinitz's theorem [41] implies that any triangulation of S^2 can be realized as the boundary of a convex 3d polyhedron. In 3d, it is known that the number of triangulations of S^3 with n nodes grows at least as $2^{n^{1.25}}$ [42, 35] and the number of 4-polytopes is at most $2^{20n\log n}$ [43, 44]. This means that the majority of triangulations of S^3 (for large n) cannot be drawn, *i.e.*, cannot be geometrically realized in S^3 or $\mathbb{R}^3$. As a consequence, recalling that only Delaunay triangulations can be used to represent positions of particles, almost all triangulations of S^3 do not have a physical interpretation.

5.2.2 The elementary moves and reducibility of the phase space

Definition of the elementary moves

The elementary moves we use in our model are called the 2-3 and 3-2 flips. They are shown in Fig. 5.1.

The *2-3 flip* transforms two tetrahedra sharing a common face f into three tetrahedra sharing a common edge e as follows: assume that $(a,b,c) = f$ are the three corners of f and d,e are its two opposite nodes. If (d,e) is already an edge of the triangulation, then we say that the face f is *not topologically flippable*. If not, the the 2-3 flip is done by removing f and both tetrahedra sharing it, adding the edge (d,e) and the three tetrahedra (d,e,a,b), (d,e,b,c) and (d,e,c,a).

The *3-2 flip* is simply the inverse move of the 2-3 flip. We say that an edge e is *not topologically flippable* if its degree is not minimal or if the target face f, whose three edges are those of the flower $\mathcal{I}(e)$, already exists.

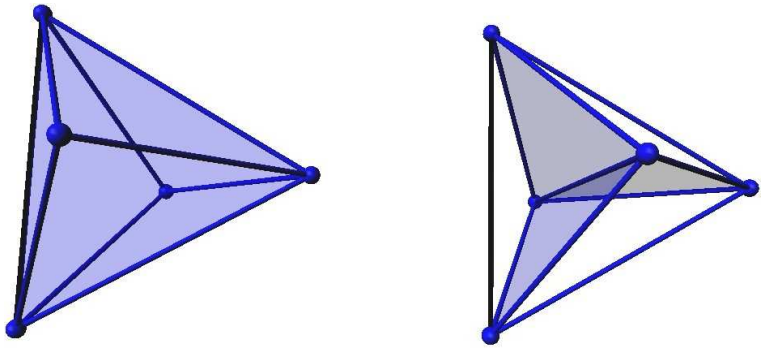

Figure 5.1: The left panel shows two tetrahedra sharing a common shaded face. The right panel shows three tetrahedra sharing a common edge such that each pair of tetrahedra shares a common shaded face. The 2-3 flip transforms the tetrahedra of the left panel into those of the right panel. The 3-2 flip is the inverse move; it transforms the tetrahedra of the right panel into those of the left one.

Reducibility of the phase space

Pachner showed [40] that any two triangulations of S^3 can be connected, *i.e.*, can be transformed into one another, by a series of four types of moves called Pachner moves in three dimensions. Two of these moves, the 2-3 and 3-2 flips shown above, do not change the number of nodes. The other two, called the 1-4 and its inverse the 4-1 flip do. The 1-4 flip divides a tetrahedron t into four by adding one node and a new tetrahedron for every face of t.

In our model, we only allow Pachner moves that do not change the number of nodes; so the question we ask is: can any two triangulations of S^3 with n nodes be connected by 2-3 and 3-2 flips? The answer to the equivalent question in two dimensions is yes: it is known that any two triangulations of S^2 with n nodes can be connected using only T1 flips (see for instance [4]).

The question in three dimensions was answered by Dougherty, Faber and Murphy in [38]. They constructed a triangulation of S^3 such that any two nodes are connected by an edge ,*i.e.*, the 1-skeleton of the triangulation is the complete graph, and every edge has its degree larger than three. Since all pairs of nodes are adjacent, no 2-3 flip is possible. Since no edge has minimal degree, no 3-2 flip is possible and this construct forms an isolated component of our phase space.

The phase space of our model is then reducible and to date, the number of connected components it has is unknown. We show that all Delaunay triangulations are in one component. Furthermore, any two Delaunay triangulations with n nodes can be connected by at most $O(n^5)$ 2-3 and 3-2 flips.

5.2.3 All Delaunay triangulations are connected

We show that all Delaunay triangulations with n nodes of the 3-sphere S^3 are connected by at most $O(n^5)$ 2-3 and 3-2 Pachner moves. This proof is based on ideas by Santos [39].

Note that the result is wrong in the euclidean flat space $\mathbb{R}^3$ because of the boundary of the triangulations (one needs two additional moves to deal with tetrahedra on the boundary. We will not develop this point further although we will point out the part of the proof that does not work in $\mathbb{R}^3$). Aside from this, the proof is identical in both spaces, with the usual definitions of straight line, ball, circle and plane in S^3. As such, and for simplicity, the proof is done in $\mathbb{R}^3$.

Remark 5.2.1. *We show that any two Delaunay triangulations of S^3 can be connected by 2-3 and 3-2 flips, i.e., that all Delaunay triangulations of S^3 are in one connected component of the phase space. We do not say that this component only contains Delaunay triangulations. Actually, using the algorithms of the previous two chapters of this thesis, we can show that this component contains non-LC spheres.*

We start with some definitions:

Notation 5.2.2. *The boundary operator is denoted by ∂.*

Definition 5.2.3. *Let $X_n \subset \mathbb{R}^3$ be a set of $n > 3$ points in $\mathbb{R}^3$. We say that the points of X_n are* coplanar *if they belong to one plane in $\mathbb{R}^3$.*

Definition 5.2.4. *Let $X_n \subset \mathbb{R}^3$ be a set of $n > 4$ points in $\mathbb{R}^3$. We say that the points of X_n are* cospherical *if there exists a point $o \in \mathbb{R}^3$ such that all points of X_n are equidistant to o.*

Definition 5.2.5. *Let $X_n \subset \mathbb{R}^3$ be a set of $n > 3$ points in $\mathbb{R}^3$. We say that the points of X_n are* cocircular *if they are coplanar and cospherical.*

Definition 5.2.6. *Let $t \subset \mathbb{R}^3$ be a tetrahedron. The unique sphere $S(a,b,c,d)$ passing through all four corners a,b,c,d of t is called the* circumsphere *of t. The* circumball *of t is the closed set containing t and whose boundary is the circumsphere of t.*

Definition 5.2.7. *Let $X \subset \mathbb{R}^3$ be a finite set of points in $\mathbb{R}^3$. We say that the points of X are* in general position *if no four points are coplanar and no five points are cospherical.*

Definition 5.2.8. *Let $X \subset \mathbb{R}^3$ be a finite set of points in $\mathbb{R}^3$ and let T be a triangulation whose nodes are the points of X. We say that T has the* empty sphere property *if the circumball of each tetrahedron contains no points of X other than the four corners of the tetrahedron.*

Let $X \subset \mathbb{R}^3$ be a finite set of points in general position in $\mathbb{R}^3$. The *Delaunay* triangulation of X can be defined in two equivalent ways (see for instance [48] or any standard book of your choice on Delaunay triangulations):

Proposition 5.2.9. *Let $X \subset \mathbb{R}^3$ be a finite set of points in general position in $\mathbb{R}^3$. Then, there is a unique triangulation T of X satisfying the empty sphere property. This triangulation is called the* Delaunay triangulation *of X and it is the dual of the Voronoi decomposition of X.*

Remark 5.2.10. *The Delaunay triangulation of a point set X whose nodes are not in general position might not be well defined.*

We now define the 2-3 and 3-2 geometrical Pachner moves. Note that, contrary to the previous topological definitions, the geometry now plays an important role, hence the need for new slightly more restrictive definitions:

Definition 5.2.11. *Let T be a triangulation and let $t_1 = (a,b,c,d)$, $t_2 = (a,b,c,e)$ be two tetrahedra sharing a common face (a,b,c) such that the six other faces of t_1, t_2 form a convex polytope. The* geometric 2-3 flip *is the operation where t_1, t_2 are replaced by three tetrahedra (a,b,d,e), (b,c,d,e), (c,a,d,e).*

Definition 5.2.12. *Let T be a triangulation and let $t_1 = (a,b,d,e)$, $t_2 = (b,c,d,e)$, $t_3 = (c,a,d,e)$ be three tetrahedra sharing a common edge (d,e) such that each of the three pairs of tetrahedra shares a common face and the six remaining faces form a convex polytope. The* geometric 3-2 flip *is the operation where t_1, t_2, t_3 are replaced by two tetrahedra (a,b,c,d), (a,b,c,e).*

Remark 5.2.13. *From a topological point of view, the moves are still the same. However, not all topologically possible flips are geometrically possible, hence the need for the convexity of the "outer polytope" in the two previous definitions. Two counter-examples are shown in Fig. 5.2 and Fig. 5.3.*

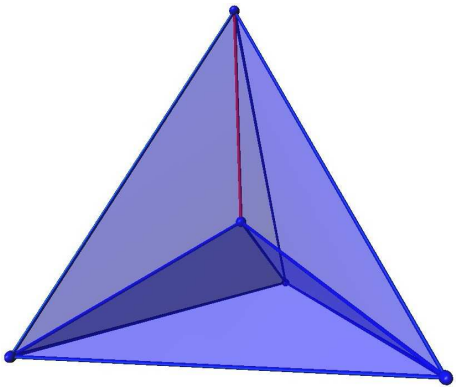

Figure 5.2: This figure shows three tetrahedra sharing a common red edge. The outer polytope formed of the six shaded faces is not convex. Clearly, the red edge is not geometrically flippable.

We now prove the main result of this section:

Theorem 5.2.14. *Let $X = \{x_1, \ldots, x_n\}, X' = \{x'_1, \ldots, x'_n\}$ be two point sets in general position in S^3 such $x_1 \neq x'_1$ and $x_\ell = x'_\ell \ \forall \ell > 1$. Let T, T' be the Delaunay triangulations of X, X'. Then, there is a path $\gamma(s), s \in [0,1]$ connecting x_1 to x'_1 such that T can be transformed into T' by moving the first node m (and all incident edges, triangles and tetrahedra) along γ and doing at most $O(n^4)$ Pachner moves. Furthermore, let $T(s)$ denote the triangulation when the first node is at $\gamma(s)$, then, one of the two following statements holds:*

P1 Either $T(s)$ satisfies the empty sphere property,
P2 Or there is one circumsphere containing m and exactly four other nodes.

Remark 5.2.15. *Saying that P1 or P2 holds implies that for every $s \in [0,1]$, the interior of any circumball of $T(s)$ is empty, i.e., it has no nodes. The only way the empty sphere property is violated is if there is exactly one circumsphere with exactly five nodes instead of four.*

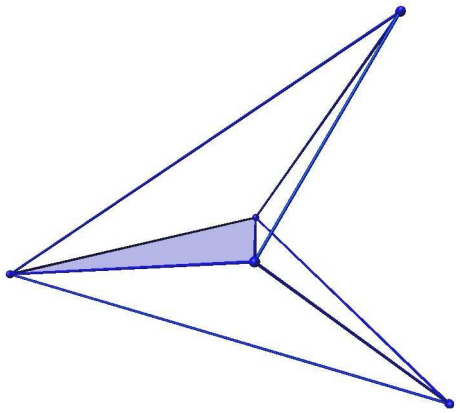

Figure 5.3: This figure shows two tetrahedra sharing a common shaded face. The outer polytope is not convex and the 2-3 flip of the shaded face is not geometrically doable.

We will need the following three lemmata:

Lemma 5.2.16. *Let $B_1 \neq B_2$ be two balls such that $\partial B_1 \cap \partial B_2$ is the circle C. The plane P supporting C cuts the the two balls into two parts each. Let $B \subset B_1, B' \subset B_2$ be two of these four parts lying on the same side of the plane P. Then either $B \subset B'$ or $B' \subset B$.*

Proof. This lemma is a simple geometrical observation and the proof is left to the reader. □

Lemma 5.2.17. *Let T be a triangulation of a given point set Y of S^3 such that P1 or P2 holds. Let a, b, c, d, e be five cospherical nodes of T such that $t = (a, b, c, d)$ is a tetrahedron of T. If the two nodes d, e are on either side of the plane supporting the face (a, b, c), then (a, b, c, e) is a tetrahedron of T.*

Proof. Since the triangulation is in S^3, every face is internal, *i.e.*, every face is shared by two tetrahedra. Let e' be the second node opposite to the face (a, b, c), the first being of course the node d. Assume that $e' \neq e$; let B_1 be the ball whose boundary contains the five nodes a, b, c, d, e (and whose interior contains the tetrahedron t) and let B_2 be the circumball of the tetrahedron (a, b, c, e'). Since no circumsphere contains more than five nodes, the node $e' \notin \partial B_1$ therefore $B_1 \neq B_2$. Since there is at most one sphere with five nodes, the node $e \notin \partial B_2$. The nodes e' and d are on either side of the plane supporting the face (a, b, c) implying that the nodes e' and e are on the same side of this plane. Applying Lemma 5.2.16 to B_1 and B_2, we deduce that either $e' \in \text{int}B_1$ or $e \in \text{int}B_2$, contradicting P1 and P2. We deduce that $e' = e$. □

Remark 5.2.18. *The above lemma is the only part of the proof that does not work in $\mathbb{R}^3$. Note that the result becomes true in $\mathbb{R}^3$ if we assume that the face (a, b, c) is internal in T.*

Lemma 5.2.19. *Let C be a finite collection of circles and let x, x' be any two points in $\mathbb{R}^3 \setminus C$. Then, there is a point $y \in \mathbb{R}^3 \setminus C$ such that the straight segments $[x, y]$ and $[x', y]$ are both in $\mathbb{R}^3 \setminus C$.*

Proof. Consider a plane $P \in \mathbb{R}^3$ such that $x, x' \notin P$. For each point $z \in \mathbb{R}^3$, we define its projection $\pi(z)$ as the intersection of the straight lines (x, z) and (x', z) with P. For any given point z, $\pi(z)$ is a finite set of points. This implies that the projection $\pi(C)$ of C cannot cover P. Let y be any point of $P \setminus \pi(C)$, then the straight segments $[x, y]$ and $[x', y]$ are both in $\mathbb{R}^3 \setminus C$. □

Proof of Theorem 5.2.14. The proof is done in $\mathbb{R}^3$ under the assumption that every face is internal. Let m denote the node of T at position x_1 and let A, B be the sets of all subsets of $X \setminus \{x_1\}$ with three and four nodes respectively. We define the following sets:

- The set
$$C_1 = \bigcup_{t \in A} C(t) \ ,$$
where $C(t)$ is the circumcircle associated with the triplet t.
- The set
$$C_2 = \bigcup_{t,t' \in B} S(t) \cap S(t') \ ,$$
where $S(t) \cap S(t')$ is the intersection of the circumspheres associated with the quadruplets t, t'.
- The set
$$C = C_1 \cup C_2 \ .$$

The set C is a finite collection of circles. Lemma 5.2.19 gives us a continuous path $\gamma(s), s \in [0, 1]$ connecting $x_1 = \gamma(0)$ and $x_1' = \gamma(1)$ such that γ is the concatenation of two straight segments and $\gamma \cap C = \emptyset$.

We start moving m along γ. All other nodes are fixed, and the only circumspheres moving are those corresponding to the tetrahedra having m as a corner. At some point, the empty sphere property gets violated and one or more of the moving spheres will have m and more than three other nodes.

Let S be such a sphere. We first show that S cannot have more than five nodes. This is obvious because all nodes but m are fixed, and at $s = 0$, the general position property is satisfied so that no more than four nodes $x_\ell, \ell > 1$ are cospherical. With m moving, this implies that no more than five nodes are cospherical.

Next, we show that for any $s \in [0, 1]$, there is at most one such sphere S. Assume that for some value $s \in [0, 1]$, there are two spheres $S \neq S'$ with five nodes each. Since at $s = 0$, the general position property is satisfied and since all nodes but m are fixed, we deduce that $m \in S \cap S'$, contradicting the fact that $\gamma \cap C = \emptyset$.

Let s_* be the smallest value of the parameter s such that:

- For $s \leq s_*$, either P1 or P2 holds.
- For any (small enough) $\varepsilon > 0$, placing m at $\gamma(s_* + \varepsilon)$ violates both P1 and P2. This means that moving m beyond $\gamma(s_*)$ with even an infinitesimal value ε will lead one of the nodes at position $x_\ell, \ell > 1$ into the interior of one of the moving balls.

This definition makes sense because a ball is a closed set. It implies that for every $s \leq s_*$, the triangulation $T(s)$ is well defined and it satisfies either P1 or P2.

For a point to cross into the interior of a ball, it must be at its boundary. This implies that at $s = s_*$, we have exactly one sphere S with m and four other nodes. We will show that the five nodes of S form two or three tetrahedra and that a Pachner move is always (geometrically) possible. By construction, we know that m and three of the four other nodes of S form a tetrahedron t (because the only moving balls are those associated with a tetrahedron having m as a corner). Let (a, b, c) be a face of $t = (a, b, c, d)$ such that the other two nodes d, e of S are on either side of the plane supporting it. It is easy enough to see that such a face always exists. More precisely, one can see that, since all five nodes are cospherical, there are only two options: either there is one such face of t, or there are two such faces of t. Applying Lemma 5.2.17, we deduce that (a, b, c, e) is a tetrahedron. But S has only five nodes, hence the two following possible outcomes:

- The five nodes of S form two tetrahedra sharing a common face (which is none other than (a, b, c)).
- The five nodes of S form three tetrahedra sharing a common edge such that each of the three pairs of tetrahedra shares a common face (satisfying the same property of the face (a, b, c) we described above).

The next step is to show that in both cases, the "outer polytope" formed by these five nodes is convex, *i.e.*, we need to show that the cases of Fig. 5.2 and Fig. 5.3 cannot happen.

- Assume that the five nodes of S form three tetrahedra (a, b, d, e),(b, c, d, e),(c, a, d, e) sharing a common edge (d, e). In this case, the outer polytope p is formed by the faces (a, b, d), (a, b, e), (a, c, d), (a, c, e), (b, c, d), (b, c, e). Since all five nodes are cospherical, we deduce that the outer polytope is convex.
- Assume that the five nodes of S form two tetrahedra (a, b, c, d),(a, b, c, e) sharing a common face (a, b, c). The outer polytope p is formed by the faces (a, b, d), (a, b, e), (a, c, d), (a, c, e), (b, c, d), (b, c, e). Assume that p is not convex. This implies that two of p's six faces, say (a, b, d) and (a, b, e) form an angle $\alpha \geq \pi$. If $\alpha = \pi$, then the four nodes a, b, d, e are cocircular, which contradicts $\gamma \cap C = \emptyset$. We deduce that $\alpha > \pi$ and that c, d are on either side of the plane supporting the face (a, b, e). Assuming that (a, b, e) is internal and applying Lemma 5.2.17, we deduce that (a, b, d, e) is a tetrahedron, contradicting the fact that a, b, c, d, e form only two tetrahedra.

All in all, at $s = s_*$, the outer polytope p is convex. For any small enough ε, and by continuity of γ, p remains convex if m is placed at $\gamma(s_* + \varepsilon)$ [1]. This implies that, when m is placed at position $\gamma(s_* + \varepsilon)$, a Pachner move is possible in both situations. One easily checks that this flip restores P1. Therefore, for ε small enough, the triangulation $T(s_* + \varepsilon)$ is obtained from $T(s_*)$ by moving m to $\gamma(s_* + \varepsilon)$ and either doing a 2-3 or a 3-2 Pachner move so that P1 is restored.

Next, we define s_{**} as the smallest value of $s > s_*$ such that

- For $s_* < s \leq s_{**}$, either P1 or P2 holds.
- For any (small enough) $\varepsilon > 0$, placing m at $\gamma(s_{**} + \varepsilon)$ violates both P1 and P2.

[1] This result is certainly true for the only polytope with six faces.

The same arguments hold. We repeat the same discussion until m reaches x'_1 at $s = 1$. By construction, $T(1)$ satisfies either P1 or P2. But the nodes of $T(1)$ are in general position, implying that $T(1) = T'$ is the Delaunay triangulation of X'.

We still need to show that the number of flips is at most $O(n^4)$. Recall that a flip only occurs if there is a sphere with five nodes, and that there is only one node moving. This means that a flip only occurs when the path γ intersects the circumsphere of a quadruplet $t \in B$. Since B contains $\binom{n-1}{4} = O(n^4)$ elements, and since the path γ is the concatenation of two straight lines, and since a straight line intersects a sphere at most twice, we deduce that the number of flips is bounded by $2 \cdot 2 \cdot \binom{n-1}{4} = O(n^4)$. □

Corollary 5.2.20. *Let $X = \{x_1, \dots, x_n\}, X' = \{x'_1, \dots, x'_n\}$ be two point sets in general position in S^3. Let T, T' be the Delaunay triangulations of X, X'. Then, there are n paths $\gamma_i(s), s \in [0,1]$ connecting x_i to x'_i for every $i = 1, \dots, n$ such that T can be transformed into T' by moving each node along the corresponding path and doing at most $O(n^4)$ Pachner moves (per path). Furthermore, let $T(i+s), i \leq n, s \in [0,1]$ denote the triangulation when the first $i-1$ nodes have already been moved and the i-th node is at $\gamma_i(s)$, then, one of the two following statements holds:*

P1 Either $T(i+s)$ satisfies the empty sphere property,
P2 Or there is at least one circumsphere containing the i-th node and four others.

Proof. Apply Theorem 5.2.14 n times. Each application gives a path γ_i and moves one node. □

Remark 5.2.21. *From the proof of Theorem 5.2.14, we deduce that the number of flips needed is bounded by $O(n^5)$.*

5.2.4 Local energy and the dynamics

Similarly to the 2d model, the energy is local: $E(A) = \sum_{v \in \mathcal{V}(A)} f(\mathcal{I}(v))$, where $\mathcal{V}(A)$ is the set of all nodes of the triangulation A, $\mathcal{I}(v)$ is the flower of node v and $f(\cdot)$ is a positive (integer) function on the set of all triangulations of S^2.

Since the number of triangulations of S^2 is quite large, such an energy form is difficult to implement. We will use the following simplification: let $\mathcal{I}(v)$ be the triangulation of S^2 representing the flower of a node v; we define the sequence of integers $e_k(v), k \geq 3$ as the number of nodes of $\mathcal{I}(v)$ with degree k and we assume that $f(\mathcal{I})$ is a function of these integers:

$$E(A) = \sum_{v \in \mathcal{V}(A)} f(\{e_k(v), k \geq 3\}) \ .$$

Note that this is simplification since many triangulations of S^2 can share the same set of integers $\{e_k\}$.

Finally, the dynamics are given by a simple Metropolis algorithm with one free parameter we call the temperature.

5.3 The 3d model

In the first part of this section, we consider a simple example of a 3d topological glass, where we say that the energy contribution of a node v is a function of its degree $\deg(v)$. Then, we construct a model of a 3d glass that makes use of the additional degrees of freedom of the flowers $\mathcal{I}(v)$, which are triangulations of S^2.

5.3.1 A first simple example of a 3d topological glass

Let v be a node. Its energy contribution $f(v)$ is a function of the flower $\mathcal{I}(v)$, which is a triangulation of S^2. The simplest possible form for $f(v)$ is by characterizing $\mathcal{I}(v)$ with a single integer, namely the number of its nodes, also known as the degree $\deg(v)$ of v; the energy of a triangulation A is then

$$E(A) = \sum_{v \in \mathcal{V}(A)} |\deg(v) - \bar{d}|^2 \ ,$$

where $\bar{d}$ is a constant of the model. Contrary to the 2d case, the number of edges and thus the average degree of a node is not fixed by the topology and the total number of nodes; as a consequence, this model has two free parameters: the (exponential of the inverse) temperature ε and the average degree at equilibrium $\bar{d}$.

Remark 5.3.1. *In [45], the authors consider a 3d binary mixture of large and small particles with equal concentration and a fixed radii ratio, each coupled to a heat bath, interacting via a Lennard-Jones potential and they perform a molecular dynamics simulation, as in [3] which served as the motivation for the 2d model. Once again, they define a quasi-species as a particle of a given size (big or small) having a certain number of neighbors, and they find that, on average at low temperature, the lowest energy per quasi-species is achieved when small particles have 10 neighbors and big particles have 14 neighbors. That corresponds, in our model, to taking* $\bar{d} = 12$ *or alternatively, to defining* $f(v) = \deg(v) - \bar{d}_v$*, where* $\bar{d}_v = 10, 14$ *for red and blue (or small and big) nodes respectively.*

As in the 2d model, we once again define a defect as a node with an incorrect degree. In the stationary state, only defects of charge ± 1 can be found and the system can be modeled as a 3d dilute gas of two types of particles.

Looking at the elementary moves around isolated defects, we see that, contrary to the 2d case, and since the 2-3 and 3-2 flips only change the degree of two nodes, an isolated defect can move without increasing the energy: let v be an isolated -1 defect, that is a node with $\bar{d} - 1$ neighbors surrounded by nodes with degree $\bar{d}$, let f be a flippable face opposite to v and let v' be the other node (of degree $\bar{d}$) opposite to f. Performing a 2-3 flip on f, we see that the degree of v becomes $\bar{d}$ and that of v' becomes $\bar{d} + 1$. In other terms, the -1 charge "jumps" from a node v to another node v' such that v, v' are opposites to some face f, without increasing the energy, but with a change of the defect's sign. If the node v' happens to have a -1 charge before the flip, then both charges disappear and the energy decreases by 2.

Similarly, using the 3-2 flip, we can see that a +1 defect jumps from a node v to an adjacent node v' while changing sign, without increasing the energy. If a flippable edge e connects two +1 charges, flipping it annihilates both charges and the energy decreases by 2.

As a consequence, and from the dynamics' point of view, such a system is simpler than the 2d model, since defects can freely jump between neighboring nodes with no energy cost (albeit with a change of sign). When two defects (of the same sign) collide, they annihilate and the energy decreases by 2.

5.3.2 Construction of the 3d model

We start by fixing the ground state of the system as a space filling crystal. We require that the crystal satisfies the following condition: its Delaunay triangulation is well defined. For instance, simply placing nodes on a cubic lattice is not acceptable since the flower of every node is not a triangulation. For simplicity, we also require that the Delaunay flowers of all nodes are identical, *i.e.*, we have one type of particles. The following is a well known fact in crystallography:

Theorem 5.3.2. *It is possible to fill the space $\mathbb{R}^3$ with particles such that the flower of every node of the Delaunay triangulation is a tetrakis hexahedron.*

A tetrakis hexahedron is the dual of the truncated octahedron, an Archemidean solid. It can be seen as a cube with square pyramids on each of its six faces. It has 14 vertices, 6 of which has a degree of 4 and the remaining 8 have a degree of 6. It is shown in Fig. 5.4.

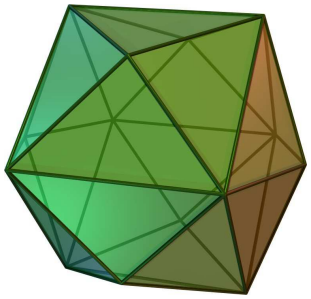

Figure 5.4: This figure shows the tetrakis hexahedron. It is taken from Wikipedia [49].

Illustration of Theorem 5.3.2. Consider the hexagonal close packing (hcp) lattice; its formed using two layers or planes, say in the directions x and y of a Cartesian set of coordinates, each tessellated with an hexagonal lattice, such that each vertex of a layer is placed above (or below, in the direction z) the center of gravity of three vertices of the other layer, in such a way that the four vertices form a regular tetrahedron. The two layers are repeatedly translated in the direction z to fill the entire space. The Delaunay triangulation of the hcp lattice is not well defined (since the vertices are not in general position): the first panel of Fig. 5.5 and Fig. 5.6 shows the flower of a vertex of the hcp lattice; it is a polyhedron with 12 vertices called a triangular orthobicupoly [50]. Notice that it contains four pairs of triangles sharing an edge and three pairs of squares sharing an edge. The idea is to move

one layer, say the layer containing the blue vertices, in such a way that the 6 squares get divided into triangles. The reader can check that this can be accomplished by moving the blue layer, by any small enough distance, in the direction orthogonal to one of the three pairs of squares, as shown in the second panel of Fig. 5.5 and Fig. 5.6: the two squares of the pair in question get divided into four triangles each by the addition of two red nodes, the other four squares of the hcp flower get divided into two triangles each and the flower becomes a tetrakis hexahedron. By symmetry of the hcp lattice, the same thing happens to the flower of every vertex. □

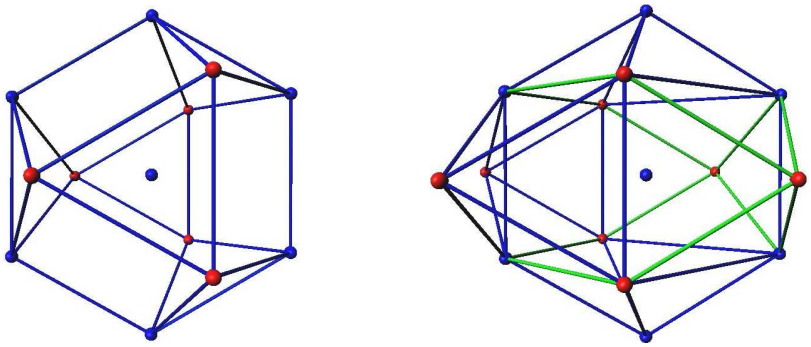

Figure 5.5: The first panel shows the flower of the blue central vertex in an hcp lattice. Vertices are colored according to their layer. Notice that the flower contains three pairs of adjacent squares and as a consequence, the Delaunay triangulation of the hcp lattice is not well defined. The second panel shows how the flower transforms if we move the blue layer to the right. The green edges are the new edges created by the translation. The two squares to the right get divided into four triangles each by the addition of two red nodes, the other four squares get divided into two triangles each and the flower becomes a tetrakis hexahedron.

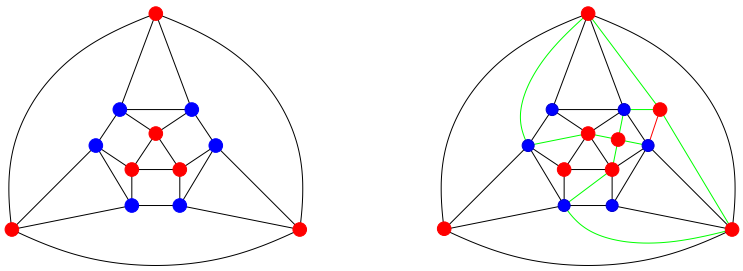

Figure 5.6: The first panel shows the flower a blue node in the hcp lattice (the colors are the same as in Fig. 5.5). It is not a triangulation. The second panel shows the flower of the same node after shifting the blue layer.

The integers $\{e_k, k \geq 3\}$ associated with a tetrakis hexahedron are $e_4^* = 6$, $e_6^* = 8$ and $e_k^* = 0 \ \forall k \notin \{4, 6\}$. We define the energy contribution of a node v as the euclidean distance

of its flower to the tetrakis hexahedron:

$$f(\mathcal{I}(v)) = \sum_{k\geq 3} \|e_k(v) - e_k^*\|^2 \ .$$

Remark 5.3.3. *There are several triangulations of S^2 with the same sequence of integers $\{e_k\}$ as the tetrakis hexahedron. A numerical investigation using all three 2d Pachner moves leads us to believe that there are only four such triangulations. In what follows, we always check that the stationary state of our system is indeed formed by tetrakis hexahedra.*

Such a model behaves as glass; its dynamics slow down tremendously when we lower the temperature. In the following section, we show a few ideas we used to speed things up.

New energy form

The first idea is to redefine the energy in a less restrictive way to give the system more room to perform flips. We consider several energy forms $f(\{e_k\})$ under the condition that $f(\{e_k\}) = 0$ if and only if $e_k = e_k^*\ \forall k \geq 3$. The one observed with the fastest dynamics is the following:

$$f(\{e_k(v)\}) = \|\deg(v) - 14\|^2 + h(\bar{k}(v)) + \varphi(e_4(v) - e_6(v)) \ ,$$

where $\deg(v) = \sum_k e_k(v)$ is the degree of node v, $\varphi(e_4(v) - e_6(v))$ is the jump function and is equal to 1 if $e_4(v) \geq e_6(v)$ and is null otherwise, $\bar{k}(v)$ is the maximal value of $k \notin \{4, 6\}$ such that $e_k(v) > 0$ and $h(\bar{k}(v)) = 1$ if $\bar{k}(v) = 3$ and $\|\bar{k}(v) - 6\|^2$ otherwise.

The reader can easily check, using Euler's theorem on triangulations of S^2, that this energy form satisfies the above condition $f(v) = 0 \Leftrightarrow e_k(v) = e_k^*\ \forall k \geq 3$.

New flips

This idea is motivated by the following observation: in the stationary state of the system at low temperatures, one expects that almost all flowers are tetrakis hexahedra. As a consequence, the degree of any edge is almost surely 4 or 6 and the degree of any node is almost surely 14 implying that the total number of edges is $E \approx 7N$. This means that, at low temperatures, the number of topologically flippable edges is exponentially small and the system will most likely flip faces. But flipping a face increases the total number of edges, leading the system away from the stationary value of $E \approx 7N$. As a consequence, the system's dynamics slow tremendously. We can speed things up by defining new elementary moves as a combination of 2-3 and 3-2 flips.

We start with edges of degree 4, as seen in Fig. 5.7: let (a, b) be an edge of degree 4 and let $(n_i, n_{i+1}), i = 1 \ldots 4$ be its flower (define $n_5 = n_1$). If n_1, n_3 are not adjacent, we can perform a 2-3 flip on the face (a, b, n_2) (or equivalently on (a, b, n_4)) and add the edge (n_1, n_3). The degree of (a, b) becomes 3 and it becomes topologically flippable. We can then perform a 3-2 flip and remove (a, b). The degree of (n_1, n_3) becomes 4.

Alternatively, if n_2, n_4 are not adjacent, we first perform a 2-3 flip of the face $(a, b, n1)$ (or equivalently on $(a, b, n3)$) followed by a 3-2 flip of the edge (a, b).

If n_1, n_3 are not adjacent, we define the $(a, b) \rightarrow (n_1, n_3)$ 4-4 flip as the operation of removing the edge (a, b) of degree 4 and flower $(n_i, n_{i+1}), i = 1 \ldots 4$ and adding the edge

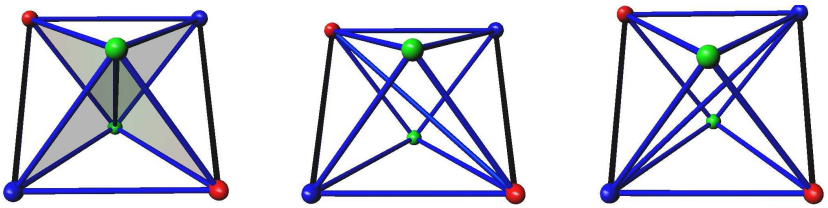

Figure 5.7: This figure shows all three related 4-4 flips $(a,b) \leftrightarrow (n_1,n_3) \leftrightarrow (n_2,n_4)$. The nodes a,b are green, the nodes n_1,n_3 are red and the node n_2,n_4 are blue.

(n_1,n_3) of degree 4 and flower (a,n_2), (n_2,b), (b,n_4) and (n_4,a). If n_2,n_4 are not adjacent, we define the $(a,b) \to (n_2,n_4)$ 4-4 flip as the operation of removing the edge (a,b) of degree 4 and flower $(n_i,n_{i+1}), i = 1\ldots4$ and adding the edge (n_2,n_4) of degree 4 and flower (a,n_1), (n_1,b), (b,n_3) and (n_3,a). Note that an edge of degree 4 can be flipped in up to two different ways; contrary to the 3-2 move, we now need to specify the edge we flip as well as the target edge in order to define the flip. Also note that the 4-4 flip does not change the number of edges. Furthermore, all three 4-4 flips $(a,b) \leftrightarrow (n_1,n_3) \leftrightarrow (n_2,n_4)$ are related in way very similar to the 2d T1 flip $(a,b) \leftrightarrow (c,d)$. Finally, since this new move is a combination of 2-3 and 3-2 flips, adding it to the allowed elementary moves does not change the connected component of the phase the system evolves in.

We can define flips of edges of higher degree in the same manner: an edge (a,b) of degree 5 and flower $(n_i,n_{i+1}), i = 1\ldots5$ can be flipped in up to 5 ways; for instance, if n_1,n_3 and n_1,n_4 are not adjacent, we define the $(a,b) \to (n_1,n_3),(n_1,n_4)$ 5-6 flip by first performing a 2-3 flip on the face (a,b,n_2) (which adds the edge (n_1,n_3) and lowers the degree of (a,b) from 5 to 4) followed by 2-3 flip on the face (a,b,n_5) (which adds the edge (n_1,n_4) and lowers the degree of (a,b) to 3) followed by a 3-2 flip on the edge (a,b).

Edges of degree 6 can be flipped in up to twelve ways. In our simulations, we restrict the elementary moves to the 2-3, 3-2, 4-4 and 5-6 flips only.

New algorithm

The variation on the Metropolis algorithm we discuss in this section is a classic one. It becomes much faster than the usual Metropolis algorithm when the probability of performing a flip is very small, which is precisely our case (due to the high energy cost of the flips and the low temperature).

Consider a triangulation A and let $\mathcal{E}$ be the set of all possible flips of A: faces and edges of degree 3 are counted once each; edges of degree 4 and 5 are counted twice and five times respectively: $|\mathcal{E}| = F + E_3 + 2E_4 + 5E_5$, where F is the total number of faces and E_i is the total number of edges of degree i.

For every flip $e \in \mathcal{E}$, we define $\Delta(e)$ as the energy difference $\Delta(e) = E(A') - E(A)$, where A' is the triangulation obtained from A by performing the flip e. Note $\mathcal{E}$ as defined above contains topologically unflippable faces and edges; by definition, the energy cost associated with a topologically impossible flip is infinity.

Remark 5.3.4. *The new algorithm we define here is an example of an absorbing Markov chain*

with the transient states being the flips e with $\Delta_e = \infty$ and the absorbing states or the exits being the topologically possible flips e with Δ_e finite.

We define the following partition of $\mathcal{E}$:

- The set C_0 is the set of all flips $e \in \mathcal{E}$ such that $\Delta(e) \leq 0$. Its cardinality is n_0.
- The set $C_\Delta, \Delta \in \mathbb{N} \setminus \{0\}$ is the set of all edges $e \in \mathcal{E}$ such that $\Delta(e) = \Delta$. Its cardinality is n_Δ.
- The set C_∞ is the set of all flips $e \in \mathcal{E}$ that are not topologically possible.

The old algorithm is the following:

- Define $\alpha_\Delta = \varepsilon^\Delta, \Delta \in \mathbb{N} \cup \{0\}$, where $\varepsilon = e^{-\beta}$, β being the inverse temperature is a fixed parameter of the algorithm. We also define $\alpha_\infty = 0$.
- Pick a flip $e \in \mathcal{E}$ with uniform probability $|\mathcal{E}|^{-1}$. If $e \in C_\Delta$, flip it with probability α_Δ, otherwise do nothing.
- Repeat the previous step.

The problem with the classic Metropolis algorithm is that almost all flips have a high energy cost (for the choice of local energy we make); more precisely, when the system is close to its stationary state at low temperature, the smallest Δ such that $n_\Delta \neq 0$ is always 2 in the 2d model and at least five times bigger in the 3d model so that the simulation spends a lot of time ($\varepsilon^{-\Delta} \gg 1$ on average) trying unsuccessfully to perform a flip.

The idea is then to compute the probability distribution $p(e, \Delta)$ of performing a flip e of energy cost Δ as well as the probability distribution $q(t)$ of the waiting time before performing the flip. We then pick a flip e with probability $p(e, \Delta)$, we flip it, we pick a time t with probability $q(t)$ and we add it to the current time.

Clearly, this new algorithm leads to the same result as the old Metropolis one without the waiting time at each flip. Note that the set of flips $\mathcal{E}$ as well as the probabilities $p(e, \Delta)$ and $q(t)$ change with the triangulation A and as a consequence need to be updated after every flip. Fortunately, the elementary moves we use are local so that the flips e' that need updating after performing a flip e are those "around" it. Since at low temperatures the degree of every node is close to 14, the number of flips e' that need updating after every flip e is relatively small which makes this new algorithm much faster than the old one.

We now compute the probabilities $q(t)$ and $p(e, \Delta)$:

Lemma 5.3.5. *With the above definitions, the probability $p(e, \Delta)$ that the next performed flip is e with energy cost Δ and the probability $q(t)$ of waiting t Metropolis time units before actually flipping e are given by:*

$$p(e, \Delta) = |\mathcal{E}|^{-1} \cdot \alpha_\Delta \cdot \delta_{\Delta, \Delta(e)} \text{ and}$$
$$q(t) = (1 - \alpha(\mathcal{E}))^{t-1} \cdot p(e, \Delta) \ ,$$

where $1 - \alpha(\mathcal{E})$ is probability of not performing a flip during one Metropolis time step:

$$\alpha(\mathcal{E}) = \sum_{\Delta, e} p(e, \Delta) = |\mathcal{E}|^{-1} \sum_{\Delta \geq 0} \alpha_\Delta n_\Delta(\mathcal{E}) \ .$$

The proof is elementary and is left for the reader. Note that $\{\alpha_\Delta\}_\Delta$ are constants and the only dependence on the system state is in $|\mathcal{E}|$ the total number of flips and the sequence $n_\Delta(\mathcal{E}) = |C_\Delta|$, the total number of flips with energy cost Δ.

Finally, we deduce that replacing the classic Metropolis algorithm with the following one will not change the probability distributions of the evolution of the system:

1. Compute the constants α_Δ for every $\Delta \geq 0$.
2. Let A_0 be the initial triangulation and $t_0 = 0$ be the initial time.
3. Compute (simultaneously) the partition $\{C_\Delta\}$ of $\mathcal{E}(A)$ as well as $\alpha(\mathcal{E})$.
4. Pick a set C_Δ with probability $\alpha_\Delta \cdot |C_\Delta| \cdot \mathcal{E}^{-1}$. Note that the set C_∞ as well as all empty sets $C_{\Delta'}$ will never get picked.
5. Randomly pick a flip $e \in C_\Delta$. This insures that every flip e is chosen with probability given by Lemma 5.3.5.
6. Pick an integer t with probability $q(t)$ given by Lemma 5.3.5. Add $t+1$ to the current time (1 time step for the flip, t steps waiting).
7. Transform A into A' by flipping e; locally update the partition $\{C_\Delta\}$ and $\alpha(\mathcal{E})$.
8. Go to step 3.

5.3.3 Results

The system as defined above shows a tremendous slowing down of its dynamics as it approaches its stationary state at low values of ε. We can say with certainty that it behaves as a glass. Unfortunately, the simulations at low values of ε do not converge in a reasonable amount of time.

The initial state is prepared as follows: we construct an hcp lattice with $n = 8000$ nodes, we shift one of its two layers as in Sect. 5.3.2 and we impose periodic boundary conditions; we obtain a Delaunay triangulation of the 3d torus such that the flower of every node is a tetrakis hexahedron. Then we heat it up (we use the ground state as a starting point of a simulation with a very high value of ε). This guarantees that the initial state of our simulations is in the same connected component as the Delaunay ground state. Note that we use the torus and not the sphere S^3 to guarantee that the ground state has energy 0.

Fig. 5.8 shows a typical behavior of the energy: the number of nodes is $n = 8000$ and the temperature is given by $\varepsilon = 0.036$. It is clear that the system has not reached its stationary state by the end of the simulation. We use a $\log - \log$ scale where a polynomial curve is represented by a straight line; we can see that the energy relaxation as a function of time exhibits three polynomial regimes during the above simulation.

An interesting question is about the notion of elementary defects: the system arranges the nodes such that their flowers are tetrakis hexahedra; at low temperature, large defect-free regions start to appear, where every node has the correct flower. When two such regions meet, mistakes, *i.e.*, nodes such that their flower is not a tetrakis hexahedron, will appear. In both the 2d model and the first 3d model of Sect. 5.3.1, we saw that these mistakes can be easily corrected and we defined the elementary defect as a single isolated node with an incorrect flower.

In the 3d model of Sect. 5.3.2, one can easily see that having a single node with an incorrect flower surrounded by a defect-free region is topologically impossible. One would

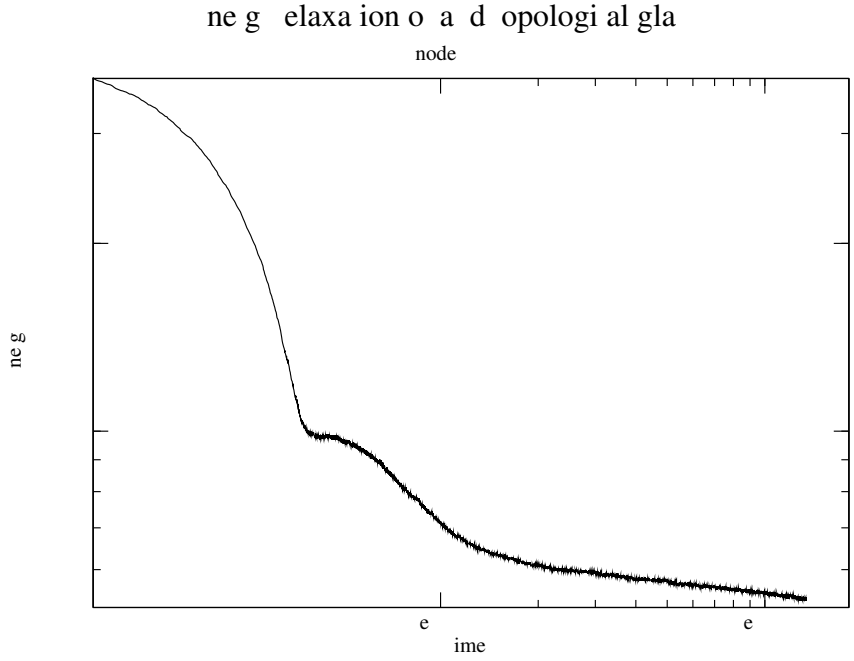

Figure 5.8: This graph shows a typical behavior of the energy of our 3d model as a function of time. The number of nodes is 8000 and the temperature is $T = 0.3 \Leftrightarrow \varepsilon = 3.6\%$. We use a log – log scale. Notice that the energy relaxation exhibits several polynomial regimes.

then expect that an elementary defect is an isolated cluster of nodes with an incorrect flower each.

However, the simulations seem to show something completely different: Fig. 5.9 shows the graph of nodes with an incorrect flower of a triangulation A with N=8000 near the stationary state at $\varepsilon = 0.1$.

The triangulation A has the following properties:

- Out of a total of 8000 nodes, 7205 nodes (90.1%) have a tetrakis hexahedron as their flower.
- Out of a total of 56477 edges, 47169 edges (83.5%) connect two nodes with a correct flower and 2776 edges (5%) connect two nodes with an incorrect flower.

We see that there are no isolated clusters of defects; on the contrary, nodes with an incorrect flower seem to form lines of defects, which would indicate that the behavior such models of 3d topological glasses is much more complicated than we originally thought.

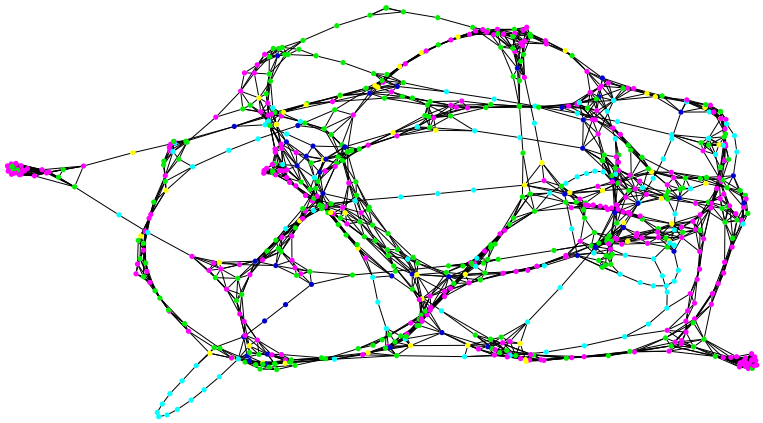

Figure 5.9: We show the "graph of defects" of a triangulation A with 8000 nodes near the stationary state at $\varepsilon = 0.1$. Each vertex of the graph represents a node of A with an incorrect flower; there are 795 such nodes (9.9%). A vertex is colored according to the energy contribution of the corresponding node: magenta for $f(v) = 1$, green, yellow, cyan and blue for $f(v) = 2, 3, 4$ and ≥ 5 respectively. Two vertices are connected if the corresponding nodes are connected by an edge in A; the above graph contains 2776 edges (5% of the total edges of A). The graph is drawn using a spring-embedder. Note that there are no isolated clusters of defects. On the contrary, the vertices seem to form lines of defects.

Index

Bibliography

[1] F. H. Stillinger and P. G. Debenedetti. *Annu. Rev. Condens. Matter Phys.* 4 (2013), 263–285.

[2] D. Rodney, A. Tanguy and D. Vandembroucq. *Model. Simul. Mater. Sci. Eng.* 19 (2011), 083001.

[3] E. Aharonov, E. Bouchbinder, H. G. E. Hentschel, V. Ilyin, N. Makedonska, I. Procaccia, and N. Schupper. Direct identification of the glass transition: Growing length scale and the onset of plasticity. *Europhysics Letters* 77 (2007), 56002.

[4] J.P. Eckmann. *J. Stat. Phys.* 129 (2007), 289–309.

[5] T. Aste and D. Sherrington. *J. Phys. A* 32 (1999), 7049–7056.

[6] A. Lemaitre, C. Caroli. *Phys. Rev. Lett.* 103 (2009), 065501.

[7] A. Tanguy, F. Leonforte, J.-L. Barrat. Plastic response of a 2D Lennard-Jones amorphous solid: Detailed analysis of the local rearrangements at very slow strain rate. *Eur. Phys. J. E* 20 (2006), 355–364.

[8] M. Mezard. First steps in glass theories, in *More Is Different*, M. P. Ong and R. N. Bhatt, Princeton Univerity Press, 2011.

[9] C. Wang and R. M. Stratt. *J. Chem. Phys.* 12 (2007), 224503.

[10] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.* 14 (1962), 21–38.

[11] G. Schliecker. *Adv. Physics* 51 (2002), 1319–1378.

[12] C. Godrèche, I. Kostov and I. Yekutieli. *Phys. Rev. Lett.* 69 (1992), 2674–2677.

[13] E.R. Berlekamp, J.H. Conway and R.K. Guy. *Winning ways for your mathematical plays. Vol. 2*, Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, 1982 Games in particular.

[14] W. Feller *An introduction to probability theory and its applications. Vol. I*, John Wiley and Sons, Inc., New York, 1957 2nd ed.

[15] L. Davison and D. Sherrington. *J. Phys. A* 33 (2000), 8615–8625.

[16] D. Sherrington, L. Davison, A. Buhot and J.P. Garraham. *J. Physics: Condens, Matter* 14 (2002), 1673–1682.

[17] D. Toussaint and F. Wilczek. *J. Chem. Phys.* 78 (1983), 2642–2647.

[18] B.P. Lee and J. Cardy. *J. Stat. Phys.* 80 (1995), 971–1007.

[19] M. Sasada. *ALEA Lat. Am. J. Probab. Math. Stat.* 7 (2010), 277–292.

[20] J. Ambjørn, B. Durhuus, and T. Jónsson. *Quantum geometry*. Cambridge Monographs on Mathematical Physics, Cambridge University Press, 1997.

[21] G. Danaraj and V. Klee. *Annals of Discrete Math* 2 (1978), 33–52.

[22] W. B. R. Lickorish. *Europ. J. Combinatorics* 12 (1991), 527–530.

[23] M. Hachimori and G. Ziegler. *Math. Z.* 235 (2000), 159–171.

[24] B. Durhuus and T. Jónsson. Remarks on the entropy of 3-manifolds. *Nuclear Phys. B* 445 (1995), 182–192.

[25] B. Benedetti and G. M. Ziegler. On locally constructible spheres and balls. *Acta Math.* 206 (2011), 205–243.

[26] R. Furch. Zur Grundlegung der kombinatorischen Topologie. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 3 (1924), 69–88.

[27] M. Hachimori. Simplicial complex library. `http://infoshako.sk.tsukuba.ac.jp/~HACHI/math/library/index_eng.html`.

[28] J. Ambjørn, B. Durhuus, and T. Jónsson. Three-dimensional simplicial quantum gravity and generalized matrix models *Modern Phys. Lett. A* 6 (1991), 1133–1146.

[29] R. H. Bing. Some aspects of the topology of 3-manifolds related to the Poincaré conjecture. In: *Lectures on modern mathematics, Vol. II* (New York: Wiley, 1964), 93–128.

[30] F. H. Lutz. Small examples of nonconstructible simplicial balls and spheres. *SIAM J. Discrete Math.* 18 (2004), 103–109.

[31] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*, Heidelberg: Springer, 2010, fourth edition.

[32] R. E. Goodrick. Non-simplicially collapsible triangulations of I^n. *Proc. Cambridge Philos. Soc.* 64 (1968), 31–36.

[33] M. Gromov. Spaces and questions. *Geom. Funct. Anal.* (2000), 118–161GAFA 2000 (Tel Aviv, 1999).

[34] M. O. Magnasco. Two-dimensional bubble rafts. *Philosophical magazine B* 65 (1992), 895–920.

[35] J. Pfeifle and G. M. Ziegler. Many triangulated 3-spheres. *Math. Ann.* 330 (2004), 829–837.

[36] C. Sundberg. *Pacific J. Math* 182 (1998), 329–358.

[37] O. Giménez and M. Noy. The number of planar graphs and properties of random planar graphs. *DMTCS proc. AD* (2005), 147–156.

[38] R. Dougherty, V. Faber and M. Murphy. Unflippable Tetrahedral Complexes. *Discrete Comput. Geom.* 32 (2004), 309–315.

[39] F. Santos. Geometric bistellar flips: the setting, the context and a construction. *In International Congress of Mathematicians. Vol. III* (2006).

[40] U. Pachner. P.L. homeomorphic manifolds are equivalent by elementary shellings. *European J. Combin.* 12 (1991), 129–145.

[41] G. M. Ziegler. *Lectures on polytopes*, chapter 4. Springer, volume 152, 1995.

[42] N. Alon. The number of polytopes, configurations and real matroids. *Matematika* 33 (1986), 62–71.

[43] J. E. Goodman and R. Pollack. There are asymptotically far fewer polytopes than we thought. *Bull. Amer. Math. Soc.* 14 (1986), 127–129.

[44] G. Kalai. Many triangulated spheres. *Discrete Comput. Geom.* 3 (1988), 1–14.

[45] E. Lerner, I. Procaccia and J. Zylberg. Statistical mechanics and dynamics of a 3-dimensional glass-forming system. *Phys. Rev. Lett.* 102 (2009), 125701.

[46] H. G. Dill and B. Weber. Variation of color, structure and morphology of fluorite and the origin of the hydrothermal F-Ba deposits at Nabburg-Wolsendorf, SE Germany *Neues Jahrbuch für Mineralogie - Abhandlungen* 187 (2010), 113–132.

[47] R. Tappert and M. C. Tappert. *Diamonds in nature : a guide to rough diamonds*. Springer, 2011, p. 25-26.

[48] M. de Berg, O. Cheong ,M. van Kreveld and M. Overmars. *Computational geometry: algorithms and applications*. Springer-Verlag, 3rd ed., 2008.

[49] Wikipedia. Tetrakis hexahedron. `http://en.wikipedia.org/wiki/Tetrakis_hexahedron`.

[50] D. Wells. *The Penguin Dictionary of Curious and Interesting Geometry*. London: Penguin, 1991, p. 53-54.