Configuration Spaces of Graphs

Daniel Lütgehetmann

Master's Thesis

Berlin, October 2014

Supervisor: Prof. Dr. Holger Reich

Selbstständigkeitserklärung

Hiermit versichere ich, dass ich die vorliegende Masterarbeit selbstständig und nur unter Zuhilfenahme der angegebenen Quellen erstellt habe.

____

Daniel Lütgehetmann

Erstkorrektor: Prof. Dr. Holger Reich
Zweitkorrektor: Prof. Dr. Elmar Vogt

Contents

Introduction

Let Γ be a graph and n a natural number. We want to understand the ordered configuration space $\mathrm{Conf}_n(\Gamma)$ consisting of all n-tuples $\mathbf{x} = (x_1, \ldots, x_n)$ of elements in Γ such that $x_i \neq x_j$ for $i \neq j$, endowed with the subspace topology induced by the inclusion $\mathrm{Conf}_n(\Gamma) \subset \Gamma^n$. The element x_i is called the i-*th particle of the configuration* $\mathbf{x}$. Configuration spaces of topological spaces have been intensively studied and come up very often in different fields of mathematics and physics. Letting the symmetric group Σ_n act by permuting the n particles gives a free Σ_n-action. Taking rational cohomology of $\mathrm{Conf}_n(\Gamma)$ yields for each k a Σ_n-representation $H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ which we want to understand for growing number of particles n.

In the case where X is a topological manifold of dimension at least 2 (and additionally connected, orientable and of finite type), Thomas Church found a nice asymptotic description of the representations $H^k(\mathrm{Conf}_n(X); \mathbb{Q})$ as n tends to infinity (compare [Chu12]); he showed that this sequence of representations is *representation stable*, a concept introduced by him and Benson Farb ([CF13]). Our aim was to find a similar description in the case of graphs.

The quotient of $\mathrm{Conf}_n(\Gamma)$ by the Σ_n-action defines the *unordered configuration space* $\mathrm{UConf}_n(\Gamma)$. Świątkowski constructed in [Św01] for any *finite* graph Γ a finite, non-positively curved cube complex $\mathrm{UK}_n\Gamma$ that is a deformation retract of $\mathrm{UConf}_n(\Gamma)$. We will generalize this construction to the ordered configuration space of locally finite graphs. More concretely, we will construct a non-positively curved, locally finite, finite-dimensional cube complex $K_n\Gamma$ and embed it into $\mathrm{Conf}_n(\Gamma)$. This complex has a Σ_n-action which is cellular and the embedding is equivariant with respect to this action. The main result in the second chapter is the following:

Theorem 2.3. *For each locally finite graph Γ the finite-dimensional cube complex $K_n\Gamma$ is an equivariant deformation retract of $\mathrm{Conf}_n(\Gamma)$. If the graph is finite, then $K_n\Gamma$ consists of finitely many cells.*

Abrams constructed in [Abr00] a complex with similar properties, but the complex $K_n\Gamma$ is much smaller and therefore better suited for concrete calculations.

Returning to our problem of understanding the cohomology we therefore may investigate the cohomology of $K_n\Gamma$ instead of the whole configuration space, which is much easier. Furthermore, we may assume that Γ is connected since the configuration space of a disconnected space is easily understood in terms of configuration spaces of its connected components.

Fact. *The complex* $K_n\Gamma$ *has only cells of dimension* $\leqslant \min\{b, n\}$ *where* b *is the number of* branched vertices *of* Γ, *i.e. vertices which have valency at least 3 (or* $b = 1$ *for* $\Gamma \cong S^1$*). If* Γ *is finite, then* $K_n\Gamma$ *has only finitely many cells in each dimension.*

Hence, if Γ is locally finite we have that $K_n\Gamma$ is a deformation retraction of $\mathrm{Conf}_n(\Gamma)$ and therefore that $H^k(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is trivial for $k > \min\{b, n\}$. If Γ is additionally finite then $\dim H^k(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is finite for each k. Note that this bound is sharp if n is big enough, more explicitly we have for $n \geqslant 2b$ that $H^b(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is non-trivial, compare Proposition 3.6. The zeroth cohomology $H^0(\mathrm{Conf}_n([0,1]);\mathbb{Q})$ is the regular representation and hence its dimension is $n!$, so we cannot expect for general graphs Γ that the dimension of $H^k(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ grows polynomially for all k.

Since $K_n\Gamma$ is non-positively curved and complete we get:

Fact. *If* Γ *is locally finite and has at least one vertex of valency* $\geqslant 3$, $K_n\Gamma$ *is an Eilenberg-MacLane space of type* $K(\pi_1(\mathrm{Conf}_n(\Gamma)), 1)$. *Since* $K_n\Gamma$ *is finite-dimensional, the fundamental group* $\pi_1(\mathrm{Conf}_n(\Gamma))$ *is torsion free.*

Calculations suggest that $H^k(\mathrm{Conf}_n(\Gamma);\mathbb{Z})$ has no torsion for any k, but we do not know how to prove this.

There are only two connected graphs Γ for which $\mathrm{Conf}_n(\Gamma)$ is not path-connected: the unit interval and S^1. Indeed, if we have a vertex of valency at least three we can bring the particles in any given ordering by "reparking" them. The two cases where the configuration space is not path connected are easily calculated explicitly, so it suffices for our general approach to consider graphs whose configuration space is path-connected.

If we are dealing with a finite graph Γ then our complex $K_n\Gamma$ allows us to explicitly compute the *generalized Euler characteristic with values in* $K_0(\mathbb{Q}\Sigma_n)$:

$$\chi^{\Sigma_n}(\mathrm{Conf}_n(\Gamma)) = \sum_{i=0}^{\infty}(-1)^i\left[H^i(\mathrm{Conf}_n(\Gamma);\mathbb{Q})\right] \in K_0(\mathbb{Q}\Sigma_n),$$

where $[M]$ is the isomorphism class of the Σ_n-representation M in $K_0(\mathbb{Q}\Sigma_n)$. This sum is well-defined as all occurring modules are finite-dimensional and only finitely many of them are non-trivial.

In the third chapter we will show (where $\chi(?)$ denotes the ordinary Euler characteristic):

Theorem 3.3. *Let* Γ *be a finite graph, then* $\chi^{\Sigma_n}(\mathrm{Conf}_n(\Gamma))$ *is a direct sum of* $\chi(\mathrm{UConf}_n(\Gamma))$ *copies of the regular representation.*

The analogous theorem is true in a more general setting, the argument uses only that the space is homotopy equivalent to a finite free Σ_n-CW complex, compare Proposition 3.4.

Consider the graph Y^l_k defined as the graph with one essential vertex, k leaves and l petals such that $2l + k \geqslant 3$:

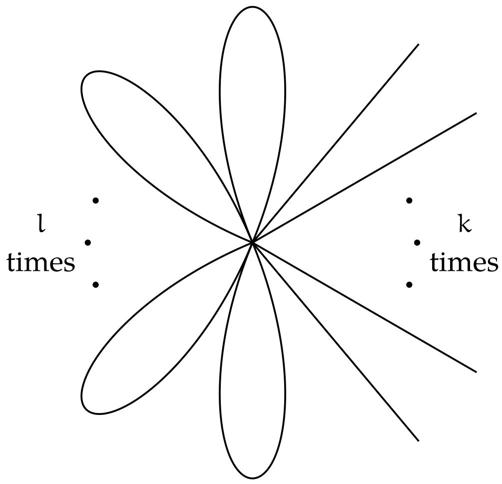

Then the generalized Euler characteristic yields a complete characterization of $H^1(\mathrm{Conf}_n(Y^l_k); \mathbb{Q})$: the zeroth cohomology is trivial by path-connectedness and since there is only one vertex of valency $\geqslant 3$ the i-th cohomology is trivial for all $i \geqslant 2$. From this one can show (the Euler characteristic $\chi(\mathrm{UConf}_n(Y^l_k))$ is non-positive):

$$H^1(\mathrm{Conf}_n(Y^l_k); \mathbb{Q}) \cong \mathbb{Q} \oplus (\mathbb{Q}\Sigma_n)^{\oplus(-\chi(\mathrm{UConf}_n(Y^l_k)))}.$$

Note that this describes the first cohomology *as Σ_n-representation*, the $\mathbb{Q}$-summand is the trivial representation.

In summary, if we have a finite connected graph with at most one branched vertex we understand $H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ as Σ_n-representation for all k. For graphs with more than one branched vertex the cohomology is not anymore determined by its Euler characteristic. Our result about the generalized Euler characteristic yields the following:

Fact. *If* $\chi(\mathrm{UConf}_n(\Gamma))$ *is not eventually zero, then there exists some* $k > 0$ *such that* $\dim H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ *grows factorially.*

To see this, calculate the dimension of $\chi^{\Sigma_n}(\mathrm{Conf}_n(\Gamma))$, then this grows factorially. Since we have a finite, fixed number of summands in this generalized Euler characteristic, the dimension of at least one of the summands also has to grow factorially.

We also have an upper bound for finite graphs: for each $i \geqslant 0$ the number of i-cells is at most $n!$ times a polynomial of degree $e(\Gamma) - 1$, where $e(\Gamma)$ is the number of unoriented edges of Γ, so this gives an upper bound for the i-th cohomology.

To check whether $\chi(\mathrm{UConf}_n(\Gamma))$ is eventually zero is very easy if one has a concrete graph since it is for large n a polynomial of degree at most $e(\Gamma) - 1$, the explicit formula is given in Chapter 3.

Although our cube complex is a finite CW complex–which is a very good situation for cohomology calculations–the number of cells grows factorially, so the involved combinatorics are very elaborate. Therefore, we were not able to deduce further restrictions on the cohomology.

As mentioned earlier, Theorem 3.3 holds for every topological space X for which $\mathrm{Conf}_n(X)$ is equivariantly homotopy equivalent to a finite CW complex with a free cellular Σ_n-action. For example, this is the case for $\mathrm{Conf}_n(M)$ with M a smooth manifold of dimension $\geqslant 2$, so the Euler characteristic $\chi(\mathrm{Conf}_n(M))$ grows factorially. Since there is only a finite number of non-trivial modules $H^k(\mathrm{Conf}_n(M); \mathbb{Q})$, this seems contradictory to Church's result in [Chu12] that the dimension of $H^k(\mathrm{Conf}_n(M); \mathbb{Q})$ grows polynomially in n for each k. However, the difference to the case of graphs is that the number of summands grows with n and is not fixed from above by some constant. Therefore, the dimension of the k-th cohomology can grow polynomially in each fixed degree k whilst at the same time the Euler characteristic grows factorially.

The high complexity of the Σ_n-representations $H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ for fixed k therefore results from the fact that all regular representations in the generalized Euler characteristic have to come from a constant finite number of representations.

In the last chapter, we take a different approach to the subject by imitating a technique which was extremely successful in the study of configuration spaces of topological manifolds, which involves spectral sequences and calculating sheaf cohomology instead of singular cohomology. At first, this looks very promising since in the case of manifolds it really helped to know the configuration spaces *locally*, and because every graph is locally star-shaped we have this knowledge also for graphs. Nevertheless, the resulting sheaves are far more complicated and–at least at the moment–we do not see a way to calculate them. The main problem here is that contrary to manifolds, graphs do not look locally everywhere the same: we have two different kinds of points, namely branched vertices and all other points. The configuration spaces of small neighborhoods of these points are then homeomorphic to the configuration space of the unit interval or a star-shaped graph with at least three edges. These spaces are however quite different, which results in complicated sheaves.

We can calculate the stalks of all occurring sheaves, but we are not able to distribute these stalks over easy sheaves in order to calculate the sheaf cohomology. In the manifold case the stalks were distributed over constant sheaves over *closed* subsets, which allowed to compute singular cohomology instead of sheaf cohomology. In our case this is not possible because there is always an *open* set on which the stalks are much more complicated than on the rest of the space, so we would need to define some of these constant sheaves over open subsets. Then, however, the stalks at points in the boundary of these open subsets are very hard to compute, which is why we were not able to find such a decomposition.

Additionally, we did some calculations with the cube complex $K_n\Gamma$ which produced a short list of some graphs and the exact dimension of their integral cohomology. This list can be found in the Appendix. Due to its shortness the list is not useful to give intuition or reveal patterns but rather to have a quick test whether a conjecture is valid at least for simple graphs and few particles or not. Note that none of the calculated $\mathbb{Z}$-modules had any torsion.

Acknowledgements

I would like to thank Holger Reich for suggesting such an interesting topic and supporting me in every possible way over the past years. I am very grateful to him and Elmar Vogt for answering all my questions and sharing their mathematical understanding.

Chapter 1

Preliminaries

In this chapter we will introduce the necessary definitions and notations. However, the reader is assumed to be familiar with the usual notions from algebra and topology, such as rings, modules, homotopy, homology and the like. The first aim is to introduce the main actors of this thesis: configuration spaces. For further information about these spaces we refer the reader to the book [FH01].

1.1 Configuration Spaces

Let X be a topological space. We want to model the situation where a finite number of particles moves around in X without touching each other. If we consider exactly n distinguishable particles in X, where n is a natural number, then the corresponding space is called the *n-th (ordered) configuration space of X*. The underlying set of this space consists of all injections $\underline{\mathbf{n}} := \{1, 2, \dots, n\} \hookrightarrow X$. To define a topology, consider the obvious injection from this space to the cartesian product X^n, mapping an injection $\mathbf{x}\colon \underline{\mathbf{n}} \hookrightarrow X$ to the tuple $(\mathbf{x}(1), \mathbf{x}(2), \dots, \mathbf{x}(n))$. Through this map, we endow the configuration space with the subspace topology. If we denote by

$$\Delta_{\text{fat}} = \big\{(x_1, x_2, \dots, x_n) \in X^n \,|\, \text{there exist } i \neq j \text{ with } x_i = x_j\big\}$$

the *fat diagonal* of X^n, then the image of our embedding is $X^n - \Delta_{\text{fat}}$.

Notation. For a topological space X the n-th (ordered) configuration space is denoted by $\mathrm{Conf}_n(X)$.

The configuration space is not only a topological space but also a right Σ_n-space in the sense of the following definition.

Definition. For a topological space Y and a topological group G we say that $Y = (Y, \rho)$ is a *(left-/right-) G-space* if ρ is a continuous (left/right) group action on Y

$$\rho\colon G \times Y \to Y. \qquad \triangle$$

Here, the group is Σ_n, the *symmetric group on* n *elements*. A permutation $\sigma \in \Sigma_n$ acts on $\mathrm{Conf}_n(X)$ by precomposition, i.e. $\mathbf{x}\colon \underline{\mathbf{n}} \to X$ is sent by σ to

$$\sigma^*(\mathbf{x}) = \mathbf{x} \circ \sigma\colon \underline{\mathbf{n}} \to \underline{\mathbf{n}} \to X.$$

Since multiplication in Σ_n is given by composition (so for $\sigma, \nu \in \Sigma_n$ we have that $\sigma \cdot \nu = \sigma \circ \nu$) this defines a continuous right action on X. The continuity is clear since Σ_n is finite and therefore endowed with the discrete topology. If for any $\mathbf{x} \in \mathrm{Conf}_n(X)$ we have $\mathbf{x}\sigma = \mathbf{x}$, then σ is the identity by injectivity of $\mathbf{x}$, so the action is free.

Remark. We denote the quotient of $\mathrm{Conf}_n(X)$ by the free action of Σ_n by

$$\mathrm{UConf}_n(X) := \mathrm{Conf}_n(X)/\Sigma_n$$

and call it the n*-th unordered configuration space*. Whenever the space X is Hausdorff the action is automatically properly discontinuous by the finiteness of Σ_n. In this case, the map $\mathrm{Conf}_n(X) \to \mathrm{UConf}_n(X)$ is a covering map with deck transformation given by the action of Σ_n on $\mathrm{Conf}_n(X)$ (compare [Hat02, Proposition 1.40, p.72], you additionally need to argue why a deck transformation of this covering is uniquely determined by the image of an arbitrary point even when $\mathrm{Conf}_n(\Gamma)$ is disconnected, but that is obvious). △

Sometimes it is more convenient to consider a "coordinate free" description of the ordered configuration space, namely to consider for a finite set S the space $\mathrm{Conf}_S(X)$ consisting of all injections $S \to X$. For every bijection between two finite sets S and T we get a bijective map between $\mathrm{Conf}_S(X)$ and $\mathrm{Conf}_T(X)$ by precomposition. $\mathrm{Conf}_{\underline{\mathbf{n}}}(X)$ is by definition the same as $\mathrm{Conf}_n(X)$ and hence carries a topology, on every other $\mathrm{Conf}_S(X)$ we put the unique topology that makes all these bijective maps into homeomorphisms. This is possible since every two bijections $S \to \underline{\mathbf{\#S}}$ differ by a permutation and permutations act on $\mathrm{Conf}_{\#S}(X)$ by homeomorphisms. Thus, the spaces $\mathrm{Conf}_S(X)$ for different S of a fixed cardinality are homeomorphic, so it suffices to understand $\mathrm{Conf}_n(X)$. Nevertheless, this description adds flexibility to our notation. By precomposition, we get an action of $\mathrm{Aut}(S)$, the group of bijections $S \to S$, on $\mathrm{Conf}_S(X)$, which therefore coincides with the Σ_n-action in the case $S = \underline{\mathbf{n}}$.

Definition. A *graph* Γ is a 1-dimensional CW-complex. We denote by $E = E(\Gamma)$ the set of open 1-cells and call it the *open edges* of Γ. The set of *vertices* of Γ consists of its 0-cells, denoted by $V = V(\Gamma)$. A graph is *finite* if E and V are finite sets. △

In this thesis we try to understand the (co-)homology of configuration spaces of graphs. Since we want to talk about Σ_n-representations we will calculate the (co-)homology with rational coefficients: over $\mathbb{Q}$ every finite-dimensional Σ_n-representation decomposes into irreducibles. We will restrict ourselves to the cohomology as we will see that over the rationals this is isomorphic to the homology.

1.2 Reduction to Connected Spaces

We will only handle connected graphs since it is easy to understand the configuration space of a topological space in terms of the configuration spaces of its connected components. Let

$$X = \coprod_{\lambda \in \Lambda} X_\lambda$$

with X_λ connected for all $\lambda \in \Lambda$ and $\Lambda^{\underline{n}}$ be the set of maps $\underline{\mathbf{n}} \to \Lambda$, then we claim to get a homeomorphism

$$\begin{aligned} \psi\colon \coprod_{\phi \in \Lambda^{\underline{n}}} \prod_{\lambda \in \operatorname{im}\phi} \mathrm{Conf}_{\phi^{-1}(\lambda)}(X_\lambda) &\xrightarrow{\cong} \mathrm{Conf}_n(X) \\ (\mathbf{x}_\lambda)_{\lambda \in \operatorname{im}\phi} &\mapsto \big(i \mapsto \mathbf{x}_{\phi(i)}(i)\big). \end{aligned} \tag{1.1}$$

Fixing one ϕ it is clear that this is a continuous map and by the universal property of the disjoint union this yields the continuity of the whole map. Now we describe the inverse map, so let $\mathbf{x} \in \mathrm{Conf}_n(X)$. Define $\phi_{\mathbf{x}}\colon \underline{\mathbf{n}} \to \Lambda$ to map i to λ for $\mathbf{x}(i) \in X_\lambda$. The image of $\mathbf{x}$ under ψ^{-1} is then $(\mathbf{x}_\lambda)_{\lambda \in \operatorname{im}\phi_{\mathbf{x}}}$ where $\mathbf{x}_\lambda\colon \phi_{\mathbf{x}}^{-1}(\lambda) \hookrightarrow \underline{\mathbf{n}} \xrightarrow{\mathbf{x}} X_\lambda$, which is well-defined as $\mathbf{x}(i)$ lies in X_λ for all $i \in \phi_{\mathbf{x}}^{-1}(\lambda)$ by definition of $\phi_{\mathbf{x}}$. It is not hard to see that the second map is continuous and the inverse map of ψ, so the latter one is indeed a homeomorphism. By the Künneth formula (compare [Bre93, Theorem 1.6, p. 320]) this induces a decomposition of $H^k(\mathrm{Conf}_n(X);\mathbb{Q})$ into a direct sum of tensor products of $H^i(\mathrm{Conf}_S(X_\lambda);\mathbb{Q})$ for different i, λ, S.

We have a canonical Σ_n-action on the left hand side of (1.1): for $(\mathbf{x}_\lambda)_{\lambda \in \operatorname{im}\phi}$ define $(\mathbf{x}_\lambda)\sigma$ to be the element $(\mathbf{x}_\lambda\sigma) \in \prod_{\lambda \in \operatorname{im}\phi\circ\sigma} \mathrm{Conf}_{(\phi\circ\sigma)^{-1}(\lambda)}(X_\lambda)$ where $\mathbf{x}_\lambda\sigma$ is obtained by precomposition with the restriction of σ to $(\phi\circ\sigma)^{-1}(\lambda)$. With the described action, the map ψ is Σ_n-equivariant. This implies that the homeomorphism descends to the quotient, yielding a decomposition of $\mathrm{UConf}_n(X)$ into products of unordered configuration spaces of the X_λ. Therefore, the cohomology of $\mathrm{UConf}_n(X)$ decomposes analogously to the cohomology of the ordered configuration space.

For a general graph we thus may handle each connected component individually, which is why we restrict our attention to connected graphs in the sequel.

1.3 The Rational Cohomology as Σ_n-Representation

We do not want to simply calculate the dimension of the cohomology of our configuration spaces but rather understand it with the additional structure given by the Σ_n-action. The next two sections are based upon [FH91] and aim to recall a few aspects of representation theory that we will need later.

Let G be a group, a G-representation over a field k is a k-vector space V with a group homomorphism $\rho\colon G \to \mathrm{Aut}_k(V)$, where $\mathrm{Aut}_k(V)$ is the automorphism group of V. By abuse of notation we will write gv instead of $\rho(g)(v)$ for $g \in G$ and $v \in V$ and omit the map ρ by talking about the representation V.

Note that a G-representation V induces a left kG-module, where kG is the *group ring*, i.e. the ring of all finite linear combinations $\lambda_1 g_1 + \cdots + \lambda_i g_i$ with $\lambda_j \in k$ and $g_j \in G$ for all j, where the multiplication is the k-bilinear extension of the group multiplication. As underlying set for this kG-module take the same vector space V, an element of the group ring then acts in the obvious way:

$$(\lambda_1 g_1 + \cdots + \lambda_i g_i) \cdot v := \lambda_1 \cdot g_1 v + \cdots + \lambda_i \cdot g_i v.$$

An easy verification yields that this makes V into a kG-module. Conversely, every kG-module V leads to a G-representation by just restricting the module multiplication to the group elements (considered as trivial linear combinations $1_k \cdot g$). In summary, we may switch between these two points of view.

By functoriality of cohomology the right action of the symmetric group described in the previous section induces a left action on the $\mathbb{Q}$-vector space $H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$. Hence, we get a Σ_n-representation for each $k \geqslant 0$, which is–as we will see–finite dimensional. The goal is therefore to understand the k-th cohomology *as a Σ_n-representation*.

Note that since the homeomorphism ψ from the last paragraph is Σ_n-equivariant the induced isomorphism on cohomology is also Σ_n-equivariant, hence the cohomologies of the two spaces are isomorphic as Σ_n-representations.

1.4 Some Representation Theory of the Symmetric Group

Since our goal is the rational representation theory of Σ_n, all our groups will assumed to be finite and our field k will always have characteristic 0.

A vector subspace U of a representation V is called *subrepresentation* if it is invariant under the action of G, that means $gu \in U$ for all $u \in U$. A representation is called *irreducible* if every proper subrepresentation is trivial. For two representations V, W the direct sum with the obvious action and the tensor product with the diagonal action $g(u \otimes v) = gu \otimes gv$ are again representations. The *trivial representation* is the 1-dimensional vector space k^1 with $gu = u$ for all $g \in G$ and $u \in k^1$, the *regular representation* as a vector space is generated by e_g for $g \in G$, the action of $h \in G$ on such a generator is given by $he_g = e_{hg}$, which determines the representation.

It is well-known that every finite-dimensional representation over $\mathbb{C}$ splits into a direct sum of irreducible representations in a unique fashion, namely the occurring irreducibles and their multiplicities are equal for any two sum decompositions.

To return to the symmetric group, we note that the irreducible representations of Σ_n are in one-to-one correspondence with partitions of n, i.e. decreasing sequences $\lambda_1 \geqslant \lambda_2 \geqslant \ldots \lambda_k \geqslant 1$ with $\lambda_1 + \lambda_2 + \cdots + \lambda_k = n$. Such partitions are often visualized as *Young diagrams*. A diagram consists of a table of boxes, where for each summand of the partition there is one row with λ_i boxes: the partition $(4, 2, 1)$ for example corresponds to the diagram . The irreducible representation corresponding to the partition λ will be written as V_λ.

Example. (i) For all $n \in \mathbb{N}$ we have the 1-dimensional Σ_n-representation $V = \mathbb{Q}$ on which every element σ acts as the identity, it is called the *trivial representation*. The partition corresponding to this representation is (n).
(ii) Another 1-dimensional representation of Σ_n is the *sign representation*, which is the vector space $\mathbb{Q}$ where the action of an element σ is given by multiplication with the sign of the permutation. The partition for this representation is $(1, 1, 1, \dots, 1)$, or in the language of Young diagrams one single row with n boxes.

Special about Σ_n-representations is that all irreducible representations can be defined even over the rationals, so the decomposition into irreducibles also works over $\mathbb{Q}$. Hence, we calculate our cohomology with rational coefficients so that we lose as few information as possible while having the nice representation theory of Σ_n.

Chapter 2

A Deformation Retraction

Now let Γ be a *connected* graph, then we construct a deformation retract $K_n\Gamma$ of $\mathrm{Conf}_n(\Gamma)$ generalizing the ideas of Świątkowski in [Św01] where a deformation retraction for the *unordered* configuration space was constructed. The disconnected case can be handled by applying the deformation retraction in each connected component individually, see Section 1.2.

We call a vertex of valency one *free*, one of valency two *inessential* and all other vertices *branched*. Since the configuration space depends only on the topology of the space we may assume that all vertices are essential, the only case where we cannot assume this is a graph homeomorphic to the 1-sphere: in this case we require it to have exactly one vertex, which will then be called branched despite the fact that it has valency two. Let B be the set of branched vertices and E_{or} the set of oriented edges in Γ: every 1-cell $|s|$ of Γ has two distinct orientations s and $-s$ and we have obvious maps

$$\iota, \tau \colon E_{or} \to V(\Gamma)$$

such that $\iota(s) = \tau(-s)$. We say that $\iota(s)$ is the initial, $\tau(s)$ the terminal vertex of the oriented edge s. For an oriented edge $s \in E_{or}$ we call $|s| \in E(\Gamma)$ the *underlying unoriented edge* of s. We fix an arbitrary orientation E_{std} of the edges of Γ and call it the *standard orientation*, i.e. we choose $E_{std} \subset E_{or}$ in such a way that for all $e \in E$ there exists exactly one $s \in E_{std}$ with $|s| = e$. For an unoriented edge e we denote by $+e$ the oriented edge $s \in E_{std}$ such that $|s| = e$; for an oriented edge s we write $-s$ for the same underlying edge with the opposite orientation. This choice of standard orientation is only necessary for easier description, two different choices yield the same deformation retraction.

To define the combinatorial structure of our cube complex we need the following definition.

Definition. A *poset* is a set P equipped with a partial order "$\preceq$". We write $F \prec G$ if

$F \preceq G$ and $F \neq G$. It is called *graded* if it admits a rank function rank: $P \to \mathbb{N} \cup \{0\}$ that satisfies

(i) $F \prec G$ implies $\operatorname{rank}(F) < \operatorname{rank}(G)$ for all $F, G \in P$

(ii) if $F \prec G$ and there exists no $H \in P$ such that $F \prec H \prec G$ then $\operatorname{rank}(F) = \operatorname{rank}(G) - 1$.

This rank function then yields the obvious decomposition $P = (P^{(0)}, P^{(1)}, \ldots)$ where $P^{(k)}$ is the set of all elements of rank k. △

We mostly will suppress the rank function and instead use the notation $P^{(k)}$ for all k-faces. For every poset $P = (P^{(0)}, P^{(1)}, \ldots)$ we always implicitly add $P^{(-1)} = \{\varnothing\}$ which fulfills $\varnothing \preceq F$ for all $F \in P$. Define for any face $F \in P$ the subposet

$$P_{\preceq F} := \{G \in P \mid G \preceq F\}.$$

If $P_{\preceq F} \cap P_{\preceq F'}$ is finite, denote by $F \cap F'$ the maximal element G (with respect to the partial order) with $G \preceq F$ and $G \preceq F'$.

A morphism of posets $\phi\colon P \to Q$ is one that preserves the partial order, i.e. $F \preceq G$ implies $\phi(F) \preceq \phi(G)$. It is a morphism of graded posets if it preserves the grading, namely $\operatorname{rank} \phi(F) = \operatorname{rank} F$ for all $F \in P$. The categories of posets and graded posets are denoted by **Poset** and **GrPoset**, respectively.

2.1 The Poset $P_n\Gamma$

In this section we will construct a graded poset $P_n\Gamma = (P_n^{(0)}\Gamma, P_n^{(1)}\Gamma, \ldots)$ which will serve as a combinatorial model for the cube complex $K_n\Gamma$, an element in $P_n^{(k)}\Gamma$ will represent a k-cube.

The idea is that we do not need to know the exact positions of the n particles in the graph, it suffices to know how they are distributed among the branched vertices and (open) edges. More explicitly we can assign to a point $\mathbf{x} = (i \mapsto x_i)$ in the configuration space a map $f_{\mathbf{x}}\colon E \cup B \to \operatorname{Tup}_n(\underline{\mathbf{n}})$. Here, for a set W and a number $n \in \mathbb{N}_0$ we denote by $\operatorname{Tup}_n(W)$ the set of ordered k-tuples of elements in W for all $k \leqslant n$:

$$\operatorname{Tup}_n(W) = \bigcup_{k=0}^{n} W^k.$$

We write a k-tuple T as usual as $(w_1, \ldots, w_k)$, its length $|T|$ is defined to be k. In the definition of $\operatorname{Tup}_n(W)$ the convention is that $W^0 = \{\varnothing\}$, so the *empty tuple* $\varnothing$ written as $()$ is contained in $\operatorname{Tup}_n(W)$ for arbitrary n and W.

For every $e \in E$ the tuple $f_{\mathbf{x}}(e)$ describes the sequence of particles occupying e in the order given by the standard orientation. The image $f_{\mathbf{x}}(b)$ of a branched vertex b

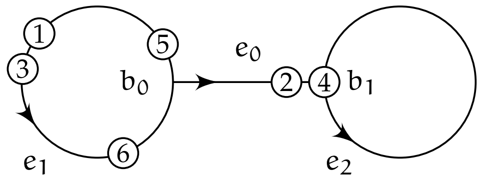

c	b_0	b_1	e_0	e_1	e_2
$f_{\mathbf{x}}(c)$	()	(4)	(2)	(5, 1, 3, 6)	()

Figure 2.1: The map $f_{\mathbf{x}}$ for the shown $\mathbf{x} \in \mathrm{Conf}_6(\Gamma)$.

has length at most 1: it is given by (j) if $b = x_j$ or the empty tuple if no such j exists. For an example of such a correspondence see Figure 2.1.

The 0-faces of the poset will be exactly those $f\colon E \cup B \to \mathrm{Tup}_n(\underline{\mathbf{n}})$ such that there exists an $\mathbf{x} \in \mathrm{Conf}_n(\Gamma)$ with $f = f_{\mathbf{x}}$. We picture such a 0-cell by putting the particles equidistantly on the edges and vertices according to the map f. A 1-face corresponds to a combinatorial movement between two 0-faces, i.e. the movement of one particle from an edge to a branched vertex or vice versa and fixing all other particles. This movement can be described by the number of the moving particle and the edge the particle is moved off with the orientation corresponding to the direction of the movement. This data is given by a map $f^{\mathrm{mov}}\colon E_{\mathrm{or}} \to \mathrm{Tup}_1(\underline{\mathbf{n}})$, compare Figure 2.2. A k-face now corresponds to moving k distinct particles in this fashion, no two towards the same branched vertex, i.e. we need k oriented edges with pairwise distinct terminal vertices to describe it.

After this motivation we will make the definition of the poset precise and afterwards discuss how this leads to the deformation retraction we announced earlier.

Formally, the set $P_n^{(k)}\Gamma$ consists of all pairs $F = (f, f^{\mathrm{mov}})$ of two maps $f\colon E \cup B \to \mathrm{Tup}_n(\underline{\mathbf{n}})$ and $f^{\mathrm{mov}}\colon E_{\mathrm{or}} \to \mathrm{Tup}_1(\underline{\mathbf{n}})$ satisfying certain properties. To describe these properties more easily we set

$$S_F = \{s \in E_{\mathrm{or}} \mid f^{\mathrm{mov}}(s) \neq \varnothing\}$$

and write $f(a)_i$ for the i-th component of $f(a)$. Such a tuple $F = (f, f^{\mathrm{mov}})$ is an element of $P_n^{(k)}\Gamma$ if:

(i) the cardinality of S_F is exactly k

(ii) $\tau(s) \in B$ for all $s \in S_F$ and $\tau(s) \neq \tau(s')$ for all $s \neq s' \in S_F$

(iii) $|f(b)| \leqslant 1$ for all $b \in B$ and $|f(\tau(s))| = 0$ for all $s \in S_F$

(iv) every $z \in \underline{\mathbf{n}}$ appears exactly once, either as value of f^{mov} or as some $f(a)_i$

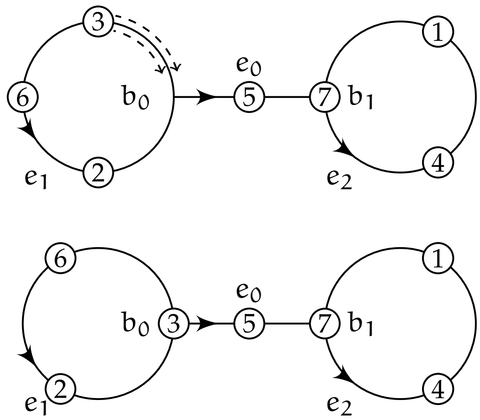

c	b_0	b_1	e_0	e_1	e_2
$f(c)$	()	(7)	(5)	(6, 2)	(4, 1)

Figure 2.2: The two 0-faces corresponding to the configurations drawn in the two graphs are connected via the 1-face $(f, f^{mov}\colon E_{or} \to \mathrm{Tup}_1(\underline{\mathbf{n}}))$, where f is defined in the table at the bottom; f^{mov} sends $-e_1$ to (3) and the remaining oriented edges to the empty tuple $()$. Note that the moving particle is not counted by f.

In our interpretation the map f distributes $n-k$ particles onto edges and branched vertices in the correct order with respect to the standard orientation of the edges; every oriented edge $s \in S_F$ stands for an additional particle moving from the unoriented edge $|s|$ towards $\tau(s)$, yielding n particles in total by (iv). Condition (ii) assures that the particles move to distinct branched (as opposed to free) vertices, the third condition guarantees that they are not moving to a vertex which is already occupied by some other particle and that there is at most one particle at each branched vertex. Note that $P_n^{(k)}\Gamma$ is empty for all $k > \min\{|B|, n\}$ by the first and second condition: we cannot have more than $|B|$ oriented edges with distinct branched endpoints and the set S_F has cardinality at most n by definition.

To introduce the partial ordering, we define the operation $*$ on tuples by

$$(a_1, a_2, \ldots, a_k) * (b_1, b_2, \ldots, b_l) := (a_1, \ldots, a_k, b_1, \ldots, b_l).$$

We say for $F = (f, f^{mov}), G = (g, g^{mov})$ that $F \prec G$ if $S_G = S_F \cup \{s\}$ for $s \notin S_F$, exactly one of the following conditions holds:

(i) $f(|s|) = g(|s|) * g^{mov}(s)$ if s is $|s|$ with the standard orientation or $f(|s|) = g^{mov}(s) * g(|s|)$ else

(ii) $f(\tau(s)) = g^{mov}(s)$

and additionally f and g as well as f^{mov} and g^{mov} agree on all other edges and vertices. In the first case we get from G to F by putting the particle moving along s onto the edge $|s|$. We have to take care of the orientation and put the number of the particle at the beginning or at the end of our *ordered* list of particles occupying $|s|$. The second case represents the moving particle being on the branched vertex, here we do not have the problem of ordering since every branched vertex is occupied by at most one particle. Recall that the moving particles are not described by f and g but rather by f^{mov} and g^{mov}. In case (i) we write $G_{-s} = (g_{-s}, g^{mov}_{-s})$ for F, in the second case F is denoted by $G_{+s} = (g_{+s}, g^{mov}_{+s})$. We extend this relation to a partial ordering by fixing transitivity, which completes the definition of our poset $(P_n\Gamma, \preceq)$.

We claim that $P_n\Gamma$ is the face poset of a cube complex $K_n\Gamma$. In the next two sections we give an explicit construction of a cube complex out of a cube poset–which is a graded poset with a special property–and then show that $P_n\Gamma$ is indeed a cube poset.

2.2 Cube Complexes and Cube Posets

Let's first examine for $k \geqslant 0$ the k-dimensional cube denoted by $\mathcal{C}^k = [-1, 1]^k$ (endowed with the standard topology) and its face poset $\mathcal{PC}^k$. For $k = 0$ this means $\mathcal{C}^0 = \{pt\}$. We will denote by **Cube** the category with objects $\mathcal{C}^k$ for $k \geqslant 0$ and morphisms all *cube maps* $\phi\colon \mathcal{C}^l \to \mathcal{C}^k$, i.e. continuous maps that send each face of $\mathcal{C}^l$ linearly onto some face of $\mathcal{C}^k$.

Now fix a $k \geqslant 1$. Taking the subsets $F_i^{\pm} = \{(x_1, \dots, x_{i-1}, \pm 1, x_{i+1}, \dots, x_k) \in \mathcal{C}^k\}$ yields all $(k-1)$-faces. Hence, we have $2k$ faces of dimension $k-1$, one for each choice of i and plus or minus. More generally, fixing $k-l$ coordinates to ± 1 gives for each combination an l-face linearly isomorphic to an l-cube. Thus, the faces are in bijection to elements in $\{-, 0, +\}^k$, where $\pm$ at position i stands for fixing the i-th coordinate to -1 or $+1$, respectively; a zero stands for no restriction in this coordinate. The dimension of the face represented by such a sequence is the number of unfixed coordinates, i.e. the number of zeros. The partial ordering in $\mathcal{PC}^k$ translates to $p \prec p'$ for $p, p' \in \{-, 0, +\}^k$ if and only if p can be obtained from p' by replacing some (at least one) of the zeros by $-$ or $+$. Observe how every l-face F of $\mathcal{C}^k$ is canonically isomorphic to $\mathcal{C}^l$ by the projection onto the non-fixed coordinates, therefore the subposet $\mathcal{PC}^k_{\preceq F}$ is isomorphic to $\mathcal{PC}^l$.

We now want to describe this category in a more combinatorial way. Every cube $\mathcal{C}^k$ is uniquely characterized by its poset, so we try to get a unique characterization of the morphisms, too. A morphism $\psi\colon \mathcal{C}^k \to \mathcal{C}^l$ in **Cube** maps by construction each face of $\mathcal{C}^k$ to a face in $\mathcal{C}^l$. Hence, this induces a map $\mathcal{P}(\psi) = \psi_*\colon \mathcal{PC}^k \to \mathcal{PC}^l$ in the obvious way. Such a linear map ψ is uniquely determined by the images of its 0-faces, so no two different maps yield the same map on posets. We call a

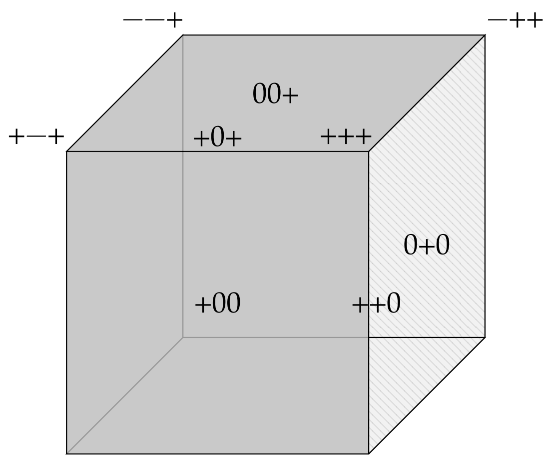

Figure 2.3: The faces of the 3-cube, observe how $000 \succ +00 \succ +0+ \succ +-+$ is a chain of faces obtained by replacing zeros by + or −, where the face 000 is the whole cube.

morphism $\phi\colon \mathcal{PC}^k \to \mathcal{PC}^l$ a *cube poset map* if there exists a morphism ψ in **Cube** such that $\phi = \mathcal{P}(\psi)$. This yields the category **PCube** with objects $\mathcal{PC}^k$ for $k \geqslant 0$ and morphisms the described cube poset maps. It is easily checked that $\mathcal{P}$ is functorial, i.e. $\mathcal{P}(\psi \circ \psi') = \mathcal{P}(\psi) \circ \mathcal{P}(\psi')$ as well as $\mathcal{P}(\mathrm{id}) = \mathrm{id}$, which proves that the morphisms in **PCube** are closed under composition and that $\mathcal{P}\colon$ **Cube** $\to$ **PCube** is a covariant functor. By construction, these assignments are bijective and for completely formal reasons we get an inverse functor $\mathcal{K}$ with $\mathcal{P} \circ \mathcal{K} = \mathrm{id}$ and $\mathcal{K} \circ \mathcal{P} = \mathrm{id}$. We therefore proved the following proposition:

Proposition 2.1. *The functors*

$$\mathcal{P}\colon \mathbf{Cube} \leftrightarrows \mathbf{PCube}\colon \mathcal{K}$$

are inverse to each other. In particular, the categories **Cube** *and* **PCube** *are isomorphic.*

Consider some isomorphism $\psi\colon \mathcal{C}^k \to \mathcal{C}^k$ in **Cube**, then this gives by functoriality an isomorphism $\mathcal{P}(\psi)$. Note that ψ in fact only permutes the coordinates and multiplies some of them by -1. The map $\pi^n_k\colon \mathcal{C}^n \to \mathcal{C}^k$ that projects onto the first k coordinates induces a map $\mathcal{P}(\pi^n_k)$ which we call a *collapse* (clearly, we need to have $n \geqslant k$). A third type of morphism in **PCube** is induced by the injection $\iota^n_k\colon \mathcal{C}^k \to \mathcal{C}^n, (x_1, \dots, x_k) \mapsto (x_1, \dots, x_k, 1, 1, \dots, 1)$ for some $k \leqslant n$, which will be called a *face injection*. Now obviously all concatenations of these three kinds of morphisms yield again morphisms, but it is indeed the case that *every* morphism in **PCube** can be written like this. To see this, simply check what it means for a map $\psi\colon \mathcal{C}^n \to \mathcal{C}^k$ to map each face of $\mathcal{C}^n$ linearly onto some face of $\mathcal{C}^k$.

The next step is to generalize this correspondence to more complicated topological spaces built by gluing cubes together:

Definition (Cube Complex, [BH99, Definition I.7.32])**.** A cube complex K is the quotient of a disjoint union of cubes $X = \bigsqcup_{\lambda \in \Lambda} \mathcal{C}^{k_\lambda}$ by an equivalence relation $\sim$ such that the quotient map $p \colon X \to X/\!\sim\, = K$ maps each cube injectively into K and we only identify faces of the same dimensions by an isometric homeomorphism. △

Remark. In the original definition by Bridson and Häfliger two cubes cannot be identified along more than one face, so in particular between two vertices there cannot be two distinct 1-cubes connecting them. This, however, happens in the complex we want to describe, so we need this slight generalization. △

As all cubes map injectively into the quotient, we may talk about *a cube in* K.

Definition. A continuous map $\psi \colon K_1 \to K_2$ between two cube complexes K_1, K_2 is called *cube map* if it maps every cube in K_1 linearly onto a cube in K_2. We denote the category with cube complexes as objects and cube maps as morphisms by **CubeComplex** △

By construction there exist canonical forgetful functors from **CubeComplex** to **Top**, the category of topological spaces and continuous maps, and **CW**, the category of CW complexes and cellular maps. We now give a combinatorial description of this category.

Definition. A graded poset P is called a *cube poset* if for every $F \in P^{(k)}$ we have that $P_{\preceq F}$ is isomorphic as a graded poset to $\mathcal{P}\mathcal{C}^k$. A poset morphism $\phi \colon P_1 \to P_2$ between two cube posets is called a *cube poset map* if the dotted map in

$$\begin{array}{ccc} (P_1)_{\preceq F} & \xrightarrow{\ \phi\ } & (P_2)_{\preceq \phi(F)} \\ \downarrow & & \downarrow \\ \mathcal{P}\mathcal{C}^{\operatorname{rank} F} & \dashrightarrow & \mathcal{P}\mathcal{C}^{\operatorname{rank} \phi(F)} \end{array}$$

is a morphism in **PCube** for every $F \in P_1$ and some (hence any) choice of poset isomorphisms for the vertical arrows. Note that although we have graded posets, we do not require a cube poset to be a graded poset morphism. The category of cube posets and cube poset maps is denoted by **CubePoset**. △

Let K be a cube complex with the notation from the definition, then we get a poset $\mathcal{P}(K)$ by taking the posets $\mathcal{P}(\mathcal{C}^{k_\lambda})$ and identifying those elements that correspond to faces that get glued together. Note that the identifications are compatible with the partial ordering on the posets, so this uniquely defines a poset. Since cubes are mapped injectively into the complex and faces are identified only isometrically the poset is a cube poset. The k-cubes in K are now in bijective correspondence to the elements of rank k in $\mathcal{P}(K)$.

A cube map $\psi\colon K_1 \to K_2$ between cube complexes sends every cube in K_1 to a cube in K_2, so define $\mathcal{P}(\psi)(F)$ to be the element in $\mathcal{P}(K_2)$ which represents the image of the cube represented by F. This is a cube poset map since every restriction to some $F \in \mathcal{P}(K_1)$ is induced by a cube map, namely the appropriate restriction of ψ. We immediately see that this is functorial and therefore this defines a functor $\mathcal{P}$: **CubeComplex** $\to$ **CubePoset**.

Observe that considering a single cube $\mathcal{C}^k$ as cube complex all the new definitions coincide with our previous ones.

2.3 The Geometric Realization of a Cube Poset

This section is devoted to the construction of the generalization of the inverse functor $\mathcal{K}$ giving a topological realization of cube posets and cube poset maps. Let P be a cube poset, we fix the isomorphisms $\rho_F\colon P_{\preceq F} \to \mathcal{P}\mathcal{C}^{\operatorname{rank} F}$ arbitrarily, we will later see that our topological space does not depend upon these choices.

Let X in the definition of cube complexes be

$$X = \bigsqcup_{F \in P} \mathcal{C}^F,$$

where $\mathcal{C}^F$ is a copy of $\mathcal{C}^{\operatorname{rank} F}$. We now need to define the equivalence relation $\sim$ on X; we do this by describing isomorphisms between faces of cubes along which we identify. For $F \in P$ every $G \preceq F$ corresponds uniquely to a face $\mathcal{C}^F_G$ of the cube $\mathcal{C}^F$ through the map ρ_F. Every such $\mathcal{C}^F_G$ for G of rank k is canonically linearly isomorphic to $\mathcal{C}^k$ by the projection map π^F_G onto those coordinates where $\rho_F(G)$ has zeros. For every pair $F \neq F' \in P$ and every non-trivial $G \in P_{\preceq F} \cap P_{\preceq F'}$ we identify $\mathcal{C}^F_G$ with $\mathcal{C}^{F'}_G$ in the following way:

Restricting the map ρ_F to $P_{\preceq G}$ we get a map $P_{\preceq G} \to \mathcal{P}\mathcal{C}^{\operatorname{rank} F}$ that yields a poset isomorphism $\overline{\rho}_F\colon \mathcal{P}_{\preceq G} \to \mathcal{P}\mathcal{C}^k$ by postcomposing with $\mathcal{P}(\pi^F_G)\colon \mathcal{P}\mathcal{C}^{\operatorname{rank} F} \to \mathcal{P}\mathcal{C}^k$, where $k = \operatorname{rank} G$. This yields a cube poset isomorphism

$$\zeta^{F,F'}_G = \left(\overline{\rho}_{F'} \circ \overline{\rho}_F^{\,-1}\right)\colon \mathcal{P}\mathcal{C}^k \to \mathcal{P}\mathcal{C}^k$$

because all involved maps are cube poset maps. The map ζ can be thought of as translating the "A-coordinates" of a face into its "B-coordinates". In Figure 2.4 we give an example of such an isomorphism.

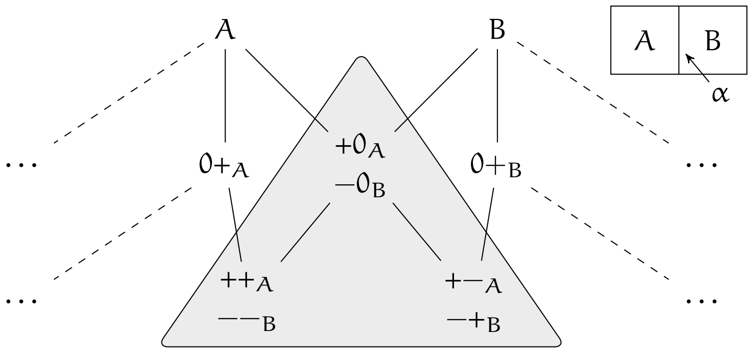

the subposet $P_{\preceq A\cap B}$

$$\zeta(0) = 0 \qquad \zeta(+) = - \qquad \zeta(-) = +$$

Figure 2.4: We see a part of the poset corresponding to two 2-cubes A, B which are identified along a common edge α. For both cubes, we chose an isomorphism to $\mathcal{PC}^2$, the label $XX_{A/B}$ stands for the fact that this element maps to XX under $\rho_{A/B}$, respectively. At the bottom we wrote down the map $\zeta = \zeta_\alpha^{A,B}$ for this situation.

Therefore we may identify along the induced isometric homeomorphism:

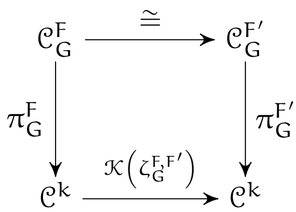

For every $H \prec G$ this construction yields the isomorphism we constructed for G restricted to the face $\mathcal{C}_H^F$ of $\mathcal{C}_G^F$ in the source and $\mathcal{C}_H^{F'}$ in the target. This is seen by observing that the following diagram commutes, where pr is the projection onto those coordinates where $\rho_F(H)$ has zeros, i.e. it forgets those coordinates that we have to fix to get from G to H:

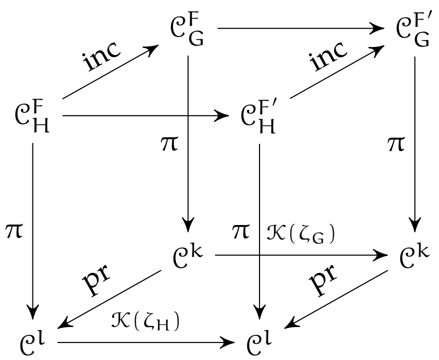

The square on the bottom commutes since we have that restricted to $\mathcal{PC}^k_{\preceq H}$ the maps $\zeta_G^{F,F'}$ and $\zeta_H^{F,F'}$ coincide. To check that, simply write down the corresponding diagrams and use that $\pi_H^F = \mathrm{pr} \circ \pi_G^F$ and therefore $\mathcal{P}(\pi_H^F) = \mathcal{P}(\mathrm{pr}) \circ \mathcal{P}(\pi_G^F)$. Hence, faces get identified if and only if they are representatives of the same element in the poset.

It is readily checked that these identifications fulfill the requirements to describe a cube complex, and this complex is denoted by $\mathcal{K}(P) := X/\sim$. We will mean by $\mathcal{C}^F$ the copy of $\mathcal{C}^{\mathrm{rank}\,F}$ as well as its image in the quotient, depending on the context.

If we have chosen a different poset isomorphism ρ_F' for some $F \in P$ we get a graded poset isomorphism $\rho_F \circ (\rho_F')^{-1}$ that defines a linear automorphism of $\mathcal{C}^F$. Precomposing every map from $\mathcal{C}^F$ and $P_{\preceq F}$ with these two maps our construction yields the same topological space as before. Hence, our cube $\mathcal{C}^F$ only needs to be replaced by a homeomorphic one to get from the definition with ρ_F to the definition with ρ_F', which shows that the resulting topological space did not change at all. Therefore, all choices of ρ_F lead to the same topological space.

Let $\phi\colon P \to Q$ be a cube poset map, we need to define a continuous map $\mathcal{K}(\phi)\colon \mathcal{K}(P) \to \mathcal{K}(Q)$. Since ϕ is a cube poset map the restriction of it to the subposet $P_{\preceq F}$ is induced by a unique linear map $\mathcal{C}^F \to \mathcal{C}^{\phi(F)}$. Through the quotient map into $\mathcal{K}(Q)$ we get a map $\mathcal{C}^F \to \mathcal{K}(Q)$. Note that this map sends every cell $\mathcal{C}^H$ corresponding to some $H \preceq F$ linearly to the cell $\mathcal{C}^{\phi(H)}$ in $\mathcal{K}(Q)$. If now two faces of $\mathcal{C}^F$ and $\mathcal{C}^{F'}$ get identified we have to check that our defined maps are compatible with the identifications. The restriction of these maps to the appropriate faces is in both cases linear and is therefore uniquely determined by the images of the 0-cells. As a result of our previous remark that every cell $\mathcal{C}^H$ gets mapped to $\mathcal{C}^{\phi(H)}$, the images of these 0-cells are independent of F, so our map descends to the quotient and yields a cube map $\mathcal{K}(\phi)$.

The functoriality of $\mathcal{K}\colon$ **PCube** $\to$ **Cube** shows that we really get a covariant functor $\mathcal{K}\colon$ **CubePoset** $\to$ **CubeComplex** since this was the only ingredient that we used. By construction, the cubes are mapped onto each other in such a way that $\mathcal{P}$ and $\mathcal{K}$ are mutually inverse functors, which shows the following proposition:

Proposition 2.2. *The functors*

$$\mathcal{P}\colon \textbf{CubeComplex} \leftrightarrows \textbf{CubePoset}\colon \mathcal{K}$$

are inverse to each other. In particular, the category of cube complexes is isomorphic to the category of cube posets.

2.4 The Cube Complex $K_n\Gamma$

All that is left to check for the concrete poset $P_n\Gamma$ is that it is a cube poset. Let $F = (f, f^{\mathrm{mov}}) \in P_n^{(k)}\Gamma$ be a k-face, meaning that S_F has k elements. We give a bijection to the face poset of the k-cube. Order the edges in S_F arbitrarily, then F corresponds to the zero vector $(0, 0, \dots, 0)$ of length k. For every $s \in S_F$ we get two descendants $F_{\pm s} = (f_{\pm s}, f^{\text{mov}}_{\pm s})$, one by putting the particle moving along s onto the terminal vertex of s, which we denote by $+$, and one by putting it on the edge $|s|$, which we denote by $-$. Descending l steps in the poset corresponds to taking l distinct s_i in S_F and deciding whether we put each corresponding particle on the edge or on the terminal vertex. For each such s_i we replace the 0 in $(0, 0, \dots, 0)$ at the i-th place by a $+$ or $-$. This is visualized in Figure 2.5. This correspondence is clearly bijective and therefore yields a poset isomorphism ρ_F as desired.

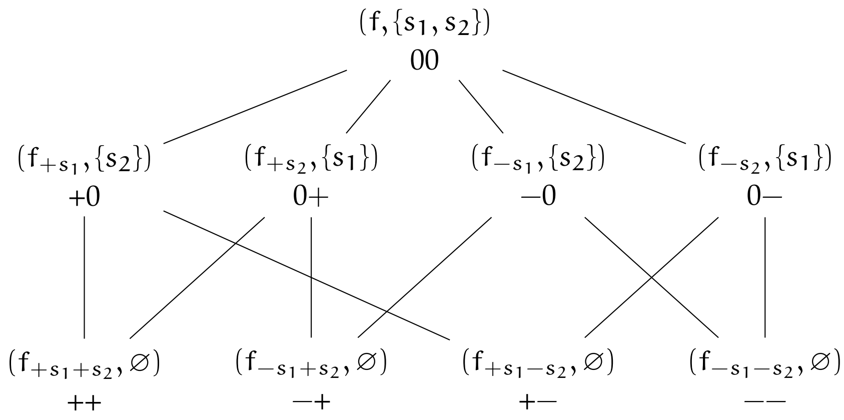

Figure 2.5: The subposet of a 2-face F in $P_n\Gamma$ and the poset of the 2-cube $\mathcal{C}^2$ combined. We wrote each $G \preceq F$ simply as (g, S_G) because this uniquely determines G as a face of F.

Definition. For a graph Γ and a natural number n we define $K_n\Gamma$ to be $\mathcal{K}(P_n\Gamma)$. Its k-cubes are in one-to-one correspondence with the elements $P_n^{(k)}\Gamma$ and the relation "$\preceq$" in $P_n\Gamma$ translates to "is contained in". △

Note that the construction of our isomorphisms ρ_F yields linear isomorphisms

$$\mathcal{C}^F \xrightarrow{\cong} [-1, 1]^{\{s_1, s_2, \dots, s_k\}} \xrightarrow{\cong} [0, 1]^{\{s_1, s_2, \dots, s_k\}}$$

for every k-face F in $P_n\Gamma$ with $S_F = \{s_1, s_2, \dots, s_k\}$; the last map is in each coordinate $t \mapsto \frac{t+1}{2}$.

The rest of this chapter contains the proof of the following result:

Theorem 2.3. *For each locally finite graph* Γ *the finite-dimensional cube complex* $K_n\Gamma$ *is an equivariant deformation retract of* $\mathrm{Conf}_n(\Gamma)$. *If the graph is finite, then* $K_n\Gamma$ *consists of finitely many cells.*

A graph is called *locally finite* if the valency of every vertex is finite.

2.5 The Embedding into the Configuration Space

We first construct an embedding $\iota\colon K_n\Gamma \to \mathrm{Conf}_n(\Gamma)$ of our cube complex into the ordered configuration space.

From the combinatorial description as cube complex we know how the particles are distributed over the edges and vertices. Hence, by putting them equidistantly onto the corresponding edges, we get for every 0-cell a well-defined image under ι. To handle a 1-cell $F = (f, f^{mov})$ with $S_F = \{s\}$ we linearly move the particle $f(s)$ on $|s|$ towards $\tau(s)$ while keeping all particles on $|s|$ equidistant. For all higher-dimensional cells we do this for multiple edges at the same time.

To be more precise consider the interval $[0, 1]$, a natural number $k \in \mathbb{N}_0$ and two numbers $t_-, t_+ \in [0, 1]$. We want to distribute k particles on $[0, 1]$ such that any two adjacent particles have the same distance c between each other. Furthermore, we want to achieve that the distance between 0 and the first particle as well as the distance between 1 and the last particle is $t_- \cdot c$ or $t_+ \cdot c$, respectively, see Figure 2.6. More concretely, we define $D_i(k, t_-, t_+) := (t_- + i - 1) \cdot c \in [0, 1]$, where $c = (t_- + k - 1 + t_+)^{-1}$. Note that this is only well defined if $k > 1$ or one of t_- and t_+ is nonzero. Thus, the map

$$\begin{aligned} D_?(k, -, -)\colon [0,1]^2 &\to \mathrm{Conf}_k([0,1]) \\ (t_-, t_+) &\mapsto D_?(k, t_-, t_+) \end{aligned}$$

is well-defined (in the case $k = 1$ remove $(0, 0)$ from the domain) and one easily verifies that it is also continuous and injective.

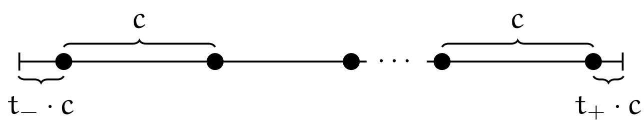

Figure 2.6: The configuration $D_?(k, t_-, t_+)$.

For every oriented edge s of Γ we now get a unique linear orientation preserving isomorphism $[0, 1] \to \bar{s} \subset \Gamma$ (orientation preserving means mapping 1 to $\tau(s)$), then denote by $D_i^s(k, t_-, t_+) \in \Gamma$ the image of $D_i(k, t_-, t_+)$ under this map. This yields k distinct elements except when $t_- = t_+ = 0$ and additionally either $k = 1$ or s is a loop and $k \geqslant 2$. Hence, this yields a continuous map $D_?^s(k, -, -)\colon [0, 1]^2 \to \mathrm{Conf}_k(\Gamma)$. Of course, we remove $(0, 0)$ again if it is necessary.

Now let $F = (f, f^{mov}) \in P_n\Gamma$ be a k-face, then we want to construct an embedding

$$\iota_F\colon \mathcal{C}^F \to \mathrm{Conf}_n(\Gamma).$$

Given $x \in \mathcal{C}^F$ and $i \in \underline{\mathbf{n}}$ we need to describe the position of the i-th particle of the configuration $\iota_F(x)$ in Γ. By the identification of $\mathcal{C}^F$ with $[0,1]^{S_F}$ from the last section we see that x is represented by a map $\nu_x \colon S_F \to [0,1]$. Now for $e \in E(\Gamma)$ and $\alpha \in \{+,-\}$ set

$$t_x^F(\alpha e) = \begin{cases} 1 - \nu_x(\alpha e) & \text{if } \alpha e \in S_F \\ 1 & \text{otherwise.} \end{cases}$$

We will omit the F from the notation whenever this does not lead to confusion.

To give the position of the i-th particle we modify our map f in the following way: for every branched vertex $b \in B$ we define $\hat{f}(b) = f(b)$; for every edge $e \in E(\Gamma)$ we define

$$\hat{f}(e) = f^{\mathrm{mov}}(-e) * f(e) * f^{\mathrm{mov}}(+e).$$

This modification represents putting all moving particles onto the edges.

Now we either have $\hat{f}(b) = (i)$ for some $b \in B$ or $\hat{f}(e)_j = i$, so we can define

$$\iota_F(x) \colon i \mapsto \begin{cases} b & \text{if } \hat{f}(b) = (i) \\ D_j^{+e}\big(|\hat{f}(e)|, t_x(-e), t_x(+e)\big) & \text{if } \hat{f}(e)_j = i. \end{cases}$$

It is easy to see that the map is well-defined, i.e. that for $i \neq j$ we get distinct elements of Γ: two points coincide only if either $t_x(s) = 0$ and $|f(\tau(s))| = 1$, so the particle moving along s maps to an occupied branched vertex, or $t_x(s_1) = t_x(s_2) = 0$ for $\tau(s_1) = \tau(s_2)$, so the particles moving along these edges are mapped to the same branched vertex. This is the case since the $D_j(k, t_-, t_+)$ are always k distinct points, so intersections have to happen at the branched vertices. Now if $t_x(s) = 0$ we have $s \in S_F$ by definition of t_x, so s, s_1 and s_2 lie in S_F and the definition of F contradicts these two cases. Therefore, the map is well-defined.

Since $t_x(s)$ is continuous in x for each s and $\hat{f}$ does not change inside of F, the map ι_F is continuous by the continuity of $D_j^s(k,-,-)$ for fixed k. It is clear that it is an injective map since for any $x \neq x' \in \mathcal{C}^F$ the maps ν_x and $\nu_{x'}$ differ on some edge s_j and hence $t_x(s_j) \neq t_{x'}(s_j)$.

Proposition 2.4. *Let* Γ *be an arbitrary graph. The maps* ι_F *for* $F \in P_n\Gamma$ *described above descend to the quotient yielding a continuous map*

$$\iota \colon K_n\Gamma \to \mathrm{Conf}_n(\Gamma).$$

Proof. Let $G = (g, g^{\mathrm{mov}}) \prec F$ be a face of F, then we again consider the identification of $\mathcal{C}^G$ and $[0,1]^{S_G}$. The identification of $\mathcal{C}^G$ and a face of $\mathcal{C}^F$ corresponds to an embedding $\omega \colon [0,1]^{S_G} \hookrightarrow [0,1]^{S_F}$: for each s which is in S_F but not in S_G the corresponding particle $f^{\mathrm{mov}}(s)$ in the face G is fixed either on the branched vertex or on the edge $|s|$. In the coordinate s we then define our embedding to be constantly 1 in the first and constantly 0 in the second case. In the coordinates $s \in S_F \cap S_G$ we map by the identity.

Now let $x \in [0,1]^{S_G}$, then we have to show that $\iota_G(x) = \iota_F(\omega(x))$. Let $i \in \underline{\mathbf{n}}$, we handle the two cases in the definition of $\iota_G(x)$ individually.

First Case: $\hat{g}(b) = i$. By the definition of the partial order "$\prec$" we either have $f^{mov}(s) = i$ for some $s \in S_F$ with $\tau(s) = b$ and therefore $t_{\omega(x)}(s) = 0$ by definition of ω or we have $\hat{f}(b) = i$. In both situations this implies $\iota_F(\omega(x)) = b = \iota_G(x)$.

Second Case: $\hat{g}(e)_j = i$. We always have $\{\pm e\} \cap S_G \subset \{\pm e\} \cap S_F$, and when this is even an equality then the claim is true because it then appears that $\hat{g}(e) = \hat{f}(e)$ and $t^G_x(\pm e) = t^F_{\omega(x)}(\pm e)$. There are three other possible cases: the set $\{\pm e\} \cap (S_F - S_G)$ contains only $-e$, it contains only $+e$ or it is equal to $\{\pm e\}$.

If it contains only $-e$ then $t^G_x(-e) = 1$ and $t^G_x(+e) = t^F_{\omega(x)}(+e)$. In this situation there are only two possibilities for $t^F_{\omega(x)}(-e)$: it is either 1 or 0. If it is equal to 1, then $\hat{f}(e) = \hat{g}(e)$ and $t^F_{\omega(x)}(-e) = t^G_x(-e)$, so we are done. Otherwise, we have $\hat{f}(e)_{j+1} = i$ and $|\hat{f}(e)| = |\hat{g}(e)| + 1$, so the claim follows from the fact

$$D_{j+1}(k, 0, t) = D_j(k-1, 1, t) \quad \text{for all } k > 1, t \in [0,1] \text{ and } 1 \leqslant j < n,$$

which is seen by plugging in the definitions.

The argument is analogous when $\{\pm e\} \cap (S_F - S_G) = \{+e\}$ of $\{\pm e\}$, here one additionally has to use the identity

$$D_j(k, t_-, 0) = D_j(k-1, t_-, 1). \qquad \square$$

Our next goal is to define the retraction from the configuration space onto our cube complex in the case of a locally finite graph Γ; in this discussion we also show that our map ι is injective. With this knowledge we have the following result which allows us to think of $K_n\Gamma$ as a subspace of $\mathrm{Conf}_n(\Gamma)$:

Corollary. *If Γ is locally finite, then the map $\iota\colon K_n\Gamma \to \mathrm{Conf}_n(\Gamma)$ is a closed embedding.*

Proof. Since ι is injective we only need to see that it is a local homeomorphism in order to prove that it is an embedding. Since Γ is locally finite our cube complex $K_n\Gamma$ is also locally finite, which means that for every point $x \in K_n\Gamma$ there exists an open neighborhood of x which is contained in a finite subcomplex Z_x. The (still injective) map $\iota|_{Z_x}$ is now a homeomorphism onto its image because Z_x is compact and $\mathrm{Conf}_n(\Gamma)$ is Hausdorff, so ι is a local homeomorphism.

Now let $\mathbf{x} = (x_1, \dots, x_n) \in \mathrm{Conf}_n(\Gamma) - \iota(K_n\Gamma)$, then we can choose an open neighborhood $U = (U_1, \dots, U_n) \subset \mathrm{Conf}_n(\Gamma)$ of $\mathbf{x}$ such that every U_i meets only finitely many edges of Γ. This is possible due to the local finiteness of Γ. That also implies that U meets only finitely many cells of $\iota(K_n\Gamma)$, so by making U smaller we can arrange that U lies in the complement of $\iota(K_n\Gamma)$. This shows that $\iota(K_n\Gamma)$ is closed in $\mathrm{Conf}_n(\Gamma)$. $\square$

2.6 The Deformation Retraction

We finally construct the deformation retraction, i.e. a map

$$r\colon \operatorname{Conf}_n(\Gamma) \to \operatorname{Conf}_n(\Gamma)$$

such that $r \circ \iota = \iota$ and $r \simeq \mathrm{id}_{\operatorname{Conf}_n(\Gamma)}$ relative $K_n\Gamma$. For this, from now on we need to assume that Γ is locally finite. For $\mathbf{x} \in \operatorname{Conf}_n(\Gamma)$ the idea is to move the particles only inside the edges they occupy in such a manner that they become equidistant. We have to be careful at both ends of the edges because restricted to our cube complex the map has to be the identity.

For all points $\mathbf{x} \in \operatorname{Conf}_n(\Gamma)$ we define $\hat{f}_{\mathbf{x}}$ and $t_{\mathbf{x}}(s)$ for all $s \in E_{or}$ in such a way that if $F = (f, f^{mov}) \in K_n\Gamma$ is the smallest cube containing some $x \in K_n\Gamma$ we recover the data used for defining $\iota(x)$, i.e. we have $\hat{f}_{\iota(x)} = \hat{f}$ and $t_{\iota(x)}(s) = t^F_x(s)$ for all s. The data $\hat{f}_{\mathbf{x}}$ and $t_{\mathbf{x}}(s)$ determines a cube $F(\mathbf{x})$ in $K_n\Gamma$ and an element $r(\mathbf{x}) \in F(\mathbf{x})$ (see below for more details), in formulas we will define

$$r(\mathbf{x})\colon i \mapsto \begin{cases} b & \text{if } \hat{f}_{\mathbf{x}}(b) = (i) \\ D_j^{+e}\big(|\hat{f}_{\mathbf{x}}(e)|, t_{\mathbf{x}}(-e), t_{\mathbf{x}}(+e)\big) & \text{if } \hat{f}_{\mathbf{x}}(e)_j = i. \end{cases} \tag{2.1}$$

The similarity to the definition of ι looks very promising, so let's get to the details.

Endow Γ with a metric such that every edge has length 1 and fix some ordered configuration $\mathbf{x}$, then Γ is subdivided into segments (of positive length) by its vertices and the particles of $\mathbf{x}$. For every oriented edge s denote by $d_{\mathbf{x}}(s)$ the length of the last (with respect to the orientation) segment contained in $|s|$.

First, we define $\hat{f}_{\mathbf{x}}$. For an open edge e denote by e^0 the closure of e minus those vertices which are branched in Γ, then for each such edge we get an ordered sequence of particles on e^0 with respect to the standard orientation. This defines the map $\hat{f}_{\mathbf{x}}$ on E, the definition on B is clear. If $\mathcal{C}^{(f, f^{mov})}$ is the smallest cube containing $x \in K_n\Gamma$ then the map $\hat{f}_{\iota(x)}$ is the same as $\hat{f}$ defined earlier.

This was the easier part, now we motivate and define $t_{\mathbf{x}}(s)$ for all oriented edges s. For this, we first look again at our cube complex. Let x be an element of $K_n\Gamma$ and $\mathcal{C}^{(f, f^{mov})}$ the smallest cube containing x. Consider the particles $\iota(x)$ in Γ and fix some oriented edge s, then we want to reconstruct $t_x(s)$ from $\iota(x)$. We have to consider three different cases:

<u>s is occupied by $\geqslant 2$ particles.</u> If c denotes the (constant!) distance between the particles on $|s|$ we have by construction $t_x(-s) \cdot c = d_{\iota(x)}(-s)$ and $t_x(s) \cdot c = d_{\iota(x)}(s)$, compare Figure 2.6. Furthermore the lengths of all segments have to sum up to 1, so

we get $d_{\iota(x)}(-s) + (|\hat{f}(|s|)| - 1) \cdot c + d_{\iota(x)}(s) = 1$ and therefore

$$c = \frac{1 - d_{\iota(x)}(-s) - d_{\iota(x)}(s)}{\left|\hat{f}(|s|)\right| - 1}.$$

Thus, we are able to calculate $t_x(s) = d_{\iota(x)}(s) \cdot c^{-1}$ by only knowing the number of particles on the edge s and the lengths $d_{\iota(x)}(\pm s)$.

<u>s is occupied by 1 particle.</u> In the case where s is occupied by exactly one particle we only know the ratio

$$\frac{t_x(s)}{t_x(-s)} = \frac{d_{\iota(x)}(s)}{d_{\iota(x)}(-s)},$$

but since in this case one of $t_x(\pm s)$ is by definition equal to 1 we get

$$t_x(s) = \min\left\{1, \frac{d_{\iota(x)}(s)}{d_{\iota(x)}(-s)}\right\}.$$

<u>s is occupied by 0 particles.</u> In this case we always have $t_x(s) = 1$.

The formulas above recover $t_x(s)$ for all $x \in K_n\Gamma$ and we want to use these formulas to define $t_{\mathbf{x}}(s)$ for *arbitrary* $\mathbf{x} \in \mathrm{Conf}_n(\Gamma)$. However, these numbers will not be automatically between 0 and 1, so we need to adapt the formulas slightly. Furthermore, we have to make sure that a particle does not move towards an occupied vertex and that no two particles move towards the same branched vertex, which is not guaranteed from the formula. If at least two particles approach a vertex b, i.e. if our formulas give $t_{\mathbf{x}}(s) < 1$ for multiple edges $s \in E_{or}(b) := \{s \in E_{or} \mid \tau(s) = b\}$ with the same terminal vertex b, we choose the edge s_{min} with the smallest value $t_{\mathbf{x}}(s_{min})$ and simply redefine $t_{\mathbf{x}}(s')$ to be 1 for the remaining edges $s' \in E_{or}(b) - \{s_{min}\}$ (for details see below), which fixes the second problem.

As an intermediate step, we define for $\mathbf{x} \in \mathrm{Conf}_n(\Gamma)$ and $s \in E_{or}$ with k the number of particles on $|s|^0$:

$$\delta_{\mathbf{x}}(s) = \begin{cases} \min\left\{1, d_{\mathbf{x}}(s) \cdot \frac{|\hat{f}_{\mathbf{x}}(|s|)|-1}{1-d_{\mathbf{x}}(-s)-d_{\mathbf{x}}(s)}\right\} & \text{for } k \geqslant 2 \\ \min\left\{1, \frac{d_{\mathbf{x}}(s)}{d_{\mathbf{x}}(-s)}\right\} & \text{for } k = 1 \\ 1 & \text{for } k = 0. \end{cases}$$

As mentioned earlier, this is almost the correct definition, we just have to make sure that the movement of the particles is allowed in the cube complex, so we set

$$t_{\mathbf{x}}(s) = \begin{cases} 1 & \text{if } \tau(s) \text{ free or } \hat{f}_{\mathbf{x}}(\tau(s)) \neq \varnothing \\ \min\left\{1, \frac{\delta_{\mathbf{x}}(s)}{\min\{\delta_{\mathbf{x}}(s') \mid s \neq s' \text{ and } \tau(s)=\tau(s')\}}\right\} & \text{else.} \end{cases}$$

Note that the minimum is not taken over the empty set since $\tau(s)$ is not free, so there exists another edge s' with the same terminal vertex. Furthermore, this minimum is attained since there are only finitely many particles. Comparing all oriented edges ending in some branched vertex one sees that the quotient is only smaller than 1 for the edge with smallest $\delta_{\mathbf{x}}$, so we indeed have ensured that every branched vertex is approached by at most one particle. By defining $t_{\mathbf{x}}(s)$ to be 1 if $\hat{f}_{\mathbf{x}}(\tau(s))$ is not empty we ensured that no particle moves towards an occupied vertex.

The main step towards proving Theorem 2.3 is the following proposition:

Proposition 2.5. *If* Γ *is locally finite, then the map* $r\colon \operatorname{Conf}_n(\Gamma) \to \operatorname{Conf}_n(\Gamma)$ *defined in* (2.1) *is a deformation retraction onto* $K_n\Gamma$.

Proof. The map r is easily seen to be continuous: if no particle occupies any branched vertex, the maps $\delta_{\mathbf{x}}(s)$ and therefore $t_{\mathbf{x}}(s)$ are continuous in $\mathbf{x}$ since this is true for $d_{\mathbf{x}}(s)$, so we may use the same argument as in the discussion of the map ι. If some particles occupy branched vertices, then $d_{\mathbf{x}}(s)$ is not continuous anymore for some s, but for those edges and the particles occupying them it is easy to verify continuity by looking at the construction.

Let $\mathbf{x}$ be an ordered configuration, then we want to read off an element of $K_n\Gamma$ which maps under ι to $r(\mathbf{x})$: to determine the face we have to give $F = (f, f^{\mathrm{mov}})$. For every $s \in E_{\mathrm{or}}$ such that $t_{\mathbf{x}}(s) < 1$ define $f^{\mathrm{mov}}(s)$ to be the last particle on s, for all other s set $f^{\mathrm{mov}}(s) = ()$. The map f on the set of branched vertices is defined in the obvious way: if b is occupied by vertex j, then we set $f(b) = (j)$, otherwise $f(b) = ()$. We set $f(e)$ to be $\hat{f}_{\mathbf{x}}(e)$ minus those elements contained in $f^{\mathrm{mov}}(\pm e)$.

Now by the similarity of the definition of r and ι it is clear that the element $x \in \mathcal{C}^F \cong [0,1]^{S_F}$ defined by $(1 - t_{\mathbf{x}}(s))_{s \in S_F}$ maps under ι to $r(\mathbf{x})$. Hence, the image of r is contained in $\iota(K_n\Gamma)$.

If (f, f^{mov}) is the smallest face containing $x \in K_n\Gamma$, we recall that $\hat{f}_{\iota(x)}(s) = \hat{f}(s)$ and $t_{\iota(x)}(s) = t_x(s)$ by construction (observe that in this case $t_{\iota(x)}(s) = \delta_{\iota(x)}(s)$), and since the particles occupying branched vertices didn't move we have $r(\iota(x)) = \iota(x)$.

An analogous argument shows the injectivity of ι: every element $x \in K_n\Gamma$ has a unique inclusion-wise smallest face it is contained in. This gives a unique description given by this face $F = (f, f^{\mathrm{mov}})$ and an element of $[0,1]^{S_F}$ where no coordinate is equal to 0 or 1. Indeed, if we would have a 0 or 1 in the s-coordinate then we would get the strictly smaller face F_{-s} or F_{+s}, respectively, containing x. The uniqueness is clear and since we can reconstruct all the data from the image point $\iota(x)$ as described above, injectivity follows.

To see the homotopy between the identity and r note that the retraction does not change the combinatorial distribution of the particles on the branched vertices. Hence, we may just use the linear homotopy on each edge e adjusting the distances between the particles properly:

Fix a point $\mathbf{x} \in \mathrm{Conf}_n(\Gamma)$ and an edge e, then this gives us a bijection $\eta\colon \hat{f}_{\mathbf{x}}(e) \to \underline{\mathbf{k}}$ for $k = \left|\hat{f}_{\mathbf{x}}(e)\right|$ indicating in which order the particles appear on the edge with respect to the standard orientation. Let $x_1, \ldots, x_k \in e^0$ be the positions of the particles of $\mathbf{x}$ occupying e^0 in the ordering according to the standard orientation of e, then if the j-th particle of $\mathbf{x}$ sits on e^0 its exact position is given by $x_{\eta(j)}$. Now let $y_1, \ldots, y_k \in e^0$ denote the positions of the particles of $r(\mathbf{x})$, then the exact position of a particle j of $r(\mathbf{x})$ sitting on e^0 is given by $y_{\eta(j)}$ for *the same* η. For a particle j occupying e^0 the homotopy is given by

$$t \mapsto t \cdot x_{\eta(j)} + (1-t) \cdot y_{\eta(j)},$$

for a particle on a branched vertex we take the constant homotopy. This fits together to a well defined homotopy between the identity on $\mathrm{Conf}_n(\Gamma)$ and r fixing $\iota(K_n\Gamma)$. □

Remark. In the paper [Św01] the formula for what corresponds in their notation to $\delta_{\mathbf{x}}(s)$ is slightly incorrect, as the resulting map is not really a retraction. There it was assumed that the particles and the branched vertices divide each edge into pieces of the same length, i.e. that even the first and the last piece have the same length as those in the middle. Consequently, there was argued that $c = (n+1)^{-1}$ and therefore the formula got slightly simpler, it would correspond in our notation to $\delta_{\mathbf{x}}(s) = d_{\mathbf{x}}(s) \cdot \left(\left|\hat{f}_{\mathbf{x}}(|s|)\right| + 1\right)$. However, the additional assumption is obviously not true for $\iota(x)$ as soon as x is not a 0-cell of $K_n\Gamma$. △

The cube complex $K_n\Gamma$ additionally admits a continuous cellular Σ_n-right action in the following way: let $\sigma \in \Sigma_n$, then for every k-cell (f, f^{mov}) in $P_n\Gamma$ we get the k-cell $(f, f^{\mathrm{mov}})\sigma := (\sigma^{-1} \cdot f, \sigma^{-1} \cdot f^{\mathrm{mov}})$, where Σ_n acts on vectors entrywise. This yields a poset automorphism denoted by σ_* and therefore induces a self-homeomorphism $\mathcal{K}(\sigma_*)$ of $K_n\Gamma$. We have $\mathcal{K}((\sigma\nu)_*) = \mathcal{K}(\nu_* \circ \sigma_*) = \mathcal{K}(\nu_*) \circ \mathcal{K}(\sigma_*)$ as well as $\mathcal{K}(\mathrm{id}) = \mathrm{id}$, so the map

$$\begin{aligned} \Sigma_n &\to \mathrm{Homeo}(K_n\Gamma) \\ \sigma &\mapsto \mathcal{K}(\sigma_*), \end{aligned}$$

where $\mathrm{Homeo}(X)$ is the group of self-homeomorphisms for any topological space X, defines a continuous right action. It is cellular because it is induced by a graded poset automorphism. The embedding $\iota\colon K_n\Gamma \to \mathrm{Conf}_n(\Gamma)$ as well as the retraction r are Σ_n-equivariant by construction, which concludes Theorem 2.3.

The action of Σ_n on $K_n\Gamma$ and on $\mathrm{Conf}_n(\Gamma)$ is clearly free and therefore properly discontinuous (because Σ_n is finite and $\mathrm{Conf}_n(\Gamma)$ is Hausdorff), so we get covering maps

$$K_n\Gamma \to UK_n\Gamma := K_n\Gamma/\Sigma_n$$

and

$$\mathrm{Conf}_n(\Gamma) \to U\mathrm{Conf}_n(\Gamma)$$

with deck transformation given by the action of Σ_n. Since ι, r and the homotopy constructed above are Σ_n-equivariant we get an embedding $UK_n\Gamma \to UConf_n(\Gamma)$ and a deformation retraction from $UConf_n(\Gamma)$ to $UK_n\Gamma$. These maps and the space $UK_n\Gamma$ are exactly those constructed in [Św01].

2.7 The Path Metric on $K_n\Gamma$

In this section we want to define a metric on the cube complex $K_n\Gamma$ in the case where Γ is a locally finite graph. The main reference for the rest of this chapter is [BH99].

Let $K = X/\sim$ be a path-connected cube complex, where X is a disjoint union of unit cubes. Every k-cube $\mathcal{C}$ in X comes with a standard metric induced from the Euclidean metric on $\mathbb{R}^k$. By injectivity of the projection map from X to K this induces a metric the cube $\mathcal{C}$ as a subset of K, denoted by $d_{\mathcal{C}}$.

Definition ([BH99, Definition I.7.4, p. 99])**.** An m-*string in* K *from* x *to* y is a sequence $\sigma = (x_0, x_1, \dots, x_m)$ of points in K such that $x = x_0, y = x_m$ and for each i there exists a cube $\mathcal{C}_i$ in K containing x_i and x_{i+1}. The *length* of such an m-string is defined as

$$l(\Sigma) := \sum_{i=0}^{m-1} d_{\mathcal{C}_i}(x_i, x_{i+1}). \qquad \triangle$$

Now we define the pseudometric d on K by

$$d(x, y) := \inf\{l(\Sigma) \mid \Sigma \text{ is a string from } x \text{ to } y\}.$$

A key property of this pseudometric is the following:

Theorem 2.6 ([BH99, Theorem I.7.50, p. 118])**.** *If the number of isometry classes of cubes in* K *is finite, then* (K, d) *is a complete geodesic metric space.*

Corollary. *If* K *is finite-dimensional, then* (K, d) *is a complete geodesic metric space. In particular,* $K_n\Gamma$ *and* $UK_n\Gamma$ *equipped with the described path metric are complete and geodesic for any* n *and arbitrary graphs* Γ*.*

For this to make sense we need another definition.

Definition. A *geodesic space* is a metric space where any two points x and y can be connected by a *geodesic segment*, i.e. by a path α such that $d(\alpha(s), \alpha(t)) = |s - t|$ for all $s, t \in [0, 1]$. $\triangle$

Remark. Recall that Bridson and Häfliger have a slightly more restrictive definition of cube complexes, so we cannot directly use their theorems. However, in their proofs they do not use the condition that two different cubes get identified along at most one face, so they work even in this more general setting. If we only would be interested in the statement that d is a metric we could instead use Corollary I.5.28 on page 69, which gives an explicit criterion for arbitrary quotients of metric spaces. $\triangle$

Now we have two topologies on K: the quotient topology and the one induced by the path metric. The first question one has to ask is answered by the next result:

Proposition 2.7. *If K is finite-dimensional and locally finite, then the topology induced by the path metric coincides with the quotient topology on K.*

Proof. We have to show that any basis element of $\mathcal{T}_{path}$ is open in $\mathcal{T}_{quot}$ and vice versa. The open balls with respect to the path metric form a basis of the topology $\mathcal{T}_{path}$, so let $x \in K$ and $\varepsilon > 0$ be arbitrary, then we have to show that for $y \in B_\varepsilon(x)$ we find an open set with respect to the quotient topology which is contained in that ball. For this, choose a string between x and y of length $\varepsilon_0 < \varepsilon$. Now by local finiteness we can find $U \in \mathcal{T}_{quot}$ containing y such that for every $z \in U$ there exist preimages $\tilde{y}, \tilde{z} \in X$ of y and z, respectively, contained in one single cube such that their distance inside this cube is strictly smaller than $\varepsilon - \varepsilon_0$. Clearly, for every such $z \in U$ we find a string between x and z of length strictly less than ε, so $\mathcal{T}_{path} \subset \mathcal{T}_{quot}$.

Conversely, let U be open with respect to the quotient topology and $x \in U$. There exists an open set $V \subset U$ containing x and meeting only finitely many cells of K by the local finiteness. Now for every preimage $\tilde{x}$ of x we define $\varepsilon(\tilde{x}) < 1/3$ such that the $\varepsilon(\tilde{x})$-ball around $\tilde{x}$ is still contained in the preimage of V. The minimum ε of these $\varepsilon(\tilde{x})$ is strictly positive and the ε-ball around x with respect to the path metric is contained in V. Hence, the two topologies agree. □

Now this is in particular true for $K_n\Gamma$ and $UK_n\Gamma$ for a locally finite graph Γ and an arbitrary n. Since the covering map $K_n\Gamma \to UK_n\Gamma$ sends cubes to cubes via isometric homeomorphisms, we get:

Lemma 2.8. *If Γ is locally finite, then the covering map $K_n\Gamma \to UK_n\Gamma$ is a local isometry with respect to the path metrics on the two cube complexes.*

2.8 The Curvature of $K_n\Gamma$

Our aim in this section is to show that the cube complex $K_n\Gamma$ is non-positively curved. But first, we define what that means.

A *geodesic triangle* Δ in a metric space X is given by three points $x, y, z \in X$ and three geodesic segments connecting them pairwise. Given such a geodesic triangle Δ, we define the *comparison triangle* $\overline{\Delta}$ to be a triangle in $\mathbb{R}^2$ with the same side lengths.

Definition. A geodesic triangle Δ satisfies the CAT(0)-inequality if the distance of any two points on the three geodesic segments of Δ is smaller or equal to the distance of the corresponding points in the comparison triangle. If each geodesic triangle in a geodesic metric space X satisfies the CAT(0)-inequality, then it is called a CAT(0)-space. A space is called *non-positively curved* if it is locally a CAT(0)-space. △

Świątkowski showed in [Św01] that $UK_n\Gamma$ is a non-positively curved cube complex by using Gromov's Lemma. His proof works also in the case where Γ is locally finite instead of finite. We could repeat this argument for the complex $K_n\Gamma$, but since we know that the covering map $K_n\Gamma \to UK_n\Gamma$ is a local isometry we immediately have by the locality of the definition:

Corollary. *The spaces* $K_n\Gamma$ *and* $UK_n\Gamma$ *endowed with their respective path metrics are non-positively curved cube complexes.*

The Cartan-Hadamard Theorem ([BH99, Theorem II.4.1, p. 193]) now tells us that the universal covering of a complete connected non-positively curved metric space is a $\mathrm{CAT}(0)$-space and in particular contractible ([BH99, Corollary II.1.5, p. 161]). This yields:

Theorem 2.9. *If* Γ *is a locally finite graph with at least vertex of valency* $\geqslant 3$*, then* $K_n\Gamma$ *and* $UK_n\Gamma$ *are Eilenberg-MacLane spaces of type* $K(\pi, 1)$ *for arbitrary* n*.*

A well-known fact is that if a $K(\pi, 1)$ is finite-dimensional, then π is torsion-free. To show this, observe that the homology of a group of finite order is non-trivial in infinitely many dimensions (compare [Bro82, Equation 3.1, p. 35]) and that group homology is isomorphic to the singular homology of the classifying space (compare [Bro82, Proposition 4.1, p. 36]).

Thus, the fundamental group of $K_n\Gamma$ is torsion-free. Calculations indicate that in (co-)homology there is no torsion either, but we do not know how to show this.

Chapter 3

The Homology of $\mathrm{Conf}_n(\Gamma)$

We now want to use the deformation retraction from the last chapter to understand the Σ_n-representations $H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ for arbitrary $k \geqslant 0$ and growing n. Our embedding $\iota\colon K_n\Gamma \to \mathrm{Conf}_n(\Gamma)$ induces isomorphisms

$$\iota_*\colon H_k(K_n\Gamma; R) \xrightarrow{\cong} H_k(\mathrm{Conf}_n(\Gamma); R)$$
$$\iota^*\colon H^k(\mathrm{Conf}_n(\Gamma); R) \xrightarrow{\cong} H^k(K_n\Gamma; R)$$

for all $k \geqslant 0$ and an arbitrary ring R. The map ι is Σ_n-equivariant, so the induced isomorphisms are really isomorphisms of $R\Sigma_n$-modules. With the universal coefficient theorem for cohomology (compare [tom08, Theorem 17.4.4, p. 418]) we get the following isomorphism

$$H^k(K_n\Gamma; \mathbb{Q}) \xrightarrow{\cong} \mathrm{Hom}_{\mathbb{Q}}(H_k(K_n\Gamma; \mathbb{Q}), \mathbb{Q}).$$

Since this isomorphism is natural, it is additionally Σ_n-equivariant. This also yields an equivariant isomorphism between $H^k(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ and the dual of homology $\mathrm{Hom}_{\mathbb{Q}}(H_k(\mathrm{Conf}_n(\Gamma); \mathbb{Q}), \mathbb{Q})$, where the action of Σ_n on the latter one is given by precomposition of the induced (right-)action on homology. If Γ is finite then we can find a non-canonical isomorphism $H_k(K_n\Gamma; \mathbb{Q}) \xrightarrow{\cong} \mathrm{Hom}_{\mathbb{Q}}(H_k(K_n\Gamma; \mathbb{Q}), \mathbb{Q})$ since both vector spaces are finite-dimensional, but this is only an isomorphism of vector spaces. The question whether this is an isomorphism of Σ_n-representations cannot be asked because the former vector space is a *right* $\mathbb{Q}\Sigma_n$-module. We can make it into a left $\mathbb{Q}\Sigma_n$-module by involution (induced by $g \mapsto g^{-1}$), but the map is still not a morphism of left $\mathbb{Q}\Sigma_n$-modules. Nevertheless, we restrict our attention to cohomology.

We only treat the ordered configuration space because the cohomology of the unordered configuration space can then be reconstructed:

Proposition 3.1. *Let* G *be a finite group and* X *be a CW complex with a free, cellular* G*-action, i.e. cells get mapped homeomorphically onto cells. Then the projection map* $\pi\colon X \to X/G$ *induces an isomorphism*

$$\pi^*\colon H^k(X/G;F) \xrightarrow{\cong} H^k(X;F)^G,$$

for any field F *whose characteristic does not divide* $|G|$*.* $H^k(X;F)^G$ *is the submodule of* $H^k(X;F)$ *consisting of all* G*-invariant elements.*

Proof. By Maschke's theorem, every FG module is projective, so in particular the trivial representation F is projective. For every FG-module V we have a canonical isomorphism of functors from **FG**−**Mod** to **F**−**Mod** between $\mathrm{Hom}_{FG}(F, ?)$ and $?^G$, which is given by

$$\begin{aligned} \mathrm{Hom}_{FG}(F,V) &\xrightarrow{\cong} V^G \\ \phi &\mapsto \phi(1). \end{aligned}$$

As F is a projective FG-module, this shows that the functor $?^G$ sending an FG-module V to $V^G \subset V$ and a morphism $f\colon V \to W$ to the restriction $f^G\colon V^G \to W^G$ is exact.

If $\mathrm{map}(A, B)$ denotes the set of maps between two sets A and B and S is a (left or right) G-set, then the map $\mathrm{pr}\colon S \to S/G$ induces an isomorphism

$$\begin{aligned} \mathrm{pr}^*\colon\ \mathrm{map}(S/G,F) &\xrightarrow{\cong} \mathrm{map}(S,F)^G \\ f &\mapsto f \circ \mathrm{pr}, \end{aligned}$$

where the G-action on $\mathrm{map}(S,F)$ is given by precomposition with the left/right-translation map. The inverse map sends $h\colon S \to F$ to $(S/G \ni Gx \mapsto h(x) \in F)$.

If now $X(n)$ denotes the G-set of n-cells of X and

$$0 \to \mathrm{Hom}_{\mathbb{Z}}(C_0(X),F) \xrightarrow{\delta^0_X} \mathrm{Hom}_{\mathbb{Z}}(C_1(X),F) \xrightarrow{\delta^1_X} \cdots$$

is the cellular cochain complex, then we get the following isomorphism of chain complexes:

$$\begin{array}{ccccccc}
0 & \longrightarrow & \mathrm{Hom}_{\mathbb{Z}}(C_0(X),F)^G & \xrightarrow{(\delta^0_X)^G} & \mathrm{Hom}_{\mathbb{Z}}(C_1(X),F)^G & \xrightarrow{(\delta^1_X)^G} & \cdots \\
 & & \downarrow \cong & & \downarrow \cong & & \\
 & & \mathrm{map}(X(0),F)^G & & \mathrm{map}(X(1),F)^G & & \cdots \\
 & & \mathrm{pr}^* \uparrow \cong & & \mathrm{pr}^* \uparrow \cong & & \\
 & & \mathrm{map}(X(0)/G,F) & & \mathrm{map}(X(1)/G,F) & & \cdots \\
 & & \uparrow \cong & & \uparrow \cong & & \\
0 & \longrightarrow & \mathrm{Hom}_{\mathbb{Z}}(C_0(X/G),F) & \xrightarrow{\delta^0_{X/G}} & \mathrm{Hom}_{\mathbb{Z}}(C_1(X/G),F) & \xrightarrow{\delta^1_{X/G}} & \cdots
\end{array}$$

Note that this indeed commutes. The k-th homology of the last row is $H^k(X/G;F)$, so it remains to show that the homology of the first row is isomorphic to $H^k(X;F)^G$.

Since $?^G$ is exact, it preserves kernels and cokernels. As $\operatorname{im} f = \ker(B \twoheadrightarrow \operatorname{coker}(f))$ for $f\colon A \to B$, it also preserves images, so we have in particular the following identities:

$$\ker\left(f^G\right) = (\ker f)^G \qquad \operatorname{im}\left(f^G\right) = (\operatorname{im} f)^G$$

They yield after applying $?^G$ to the short exact sequence

$$0 \to \operatorname{im} \delta_X^{k-1} \to \ker \delta_X^k \to H^k(X;F) \to 0$$

the short exact sequence

$$0 \to \operatorname{im}\left(\left(\delta_X^{k-1}\right)^G\right) \to \ker\left(\left(\delta_X^{k}\right)^G\right) \to H^k(X;F)^G \to 0,$$

which proves the claim. □

Corollary. *We have an isomorphism*

$$\pi^*\colon H^k(\mathrm{UConf}_n(\Gamma);\mathbb{Q}) \xrightarrow{\cong} H^k(\mathrm{Conf}_n(\Gamma);\mathbb{Q})^{\Sigma_n}.$$

The deformation retraction from the last chapter furthermore shows that the module $H^k(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is trivial for any locally finite graph Γ as soon as $k > \min\{b, n\}$, where b is the number of branched vertices of Γ, so we have only finitely many modules that we have to understand for a growing number of particles. To approach this problem we think of the canonical maps

$$f_n\colon \mathrm{Conf}_n(\Gamma) \to \mathrm{Conf}_{n-1}(\Gamma)$$

which simply forget the last particle. They induce maps

$$H^k(\mathrm{Conf}_1(\Gamma)) \to H^k(\mathrm{Conf}_2(\Gamma)) \to H^k(\mathrm{Conf}_3(\Gamma)) \to \cdots$$

The most naive guess would be that this sequence satisfies *homological stability*, i.e. that the maps f_n^* are all isomorphisms for n big enough. But even in the simple case that Γ is the unit interval we will see that the dimension of the zeroth cohomology will be $n!$, so there is no chance that the maps f_n^* are isomorphisms.

Now we have the additional information of the Σ_n-action on the vector spaces, but comparing two representations V_n and V_{n+1} requires that they are representations of the same group, which they are not. The key idea is to consider certain representations of Σ_n and Σ_m as being "the same representation" even for $n \neq m$; this will be covered in the next section.

3.1 Representation Stability

In [CEF12] the concept of representation stability and finitely generated FI-modules is treated. We want to shortly review these notions and discuss how one could use this theory to analyze our situation.

To make the irreducible representations of Σ_n somehow independent of n we define for a partition λ of k the *padded* partition $\lambda[n] = (n-k, \lambda_1, \lambda_2, \ldots)$ of n. For this to be a well-defined partition we need $n - k \geqslant \lambda_1$. We now define $V(\lambda)_n$ to be the irreducible representation $V_{\lambda[n]}$, then every irreducible Σ_n-representation is of that form for a unique λ. In the language of Young diagrams this corresponds simply to forgetting the first row, for example we have

$$V_{(3,2)} = V((2))_5.$$

By this notation we have an irreducible representation $V(\lambda)$ for *all* big enough n at the same time, so we are now able to say that an irreducible representation of Σ_n is equal to an irreducible representation of Σ_m in the sense that they come from the same partition λ in the form of $V(\lambda)_n$ and $V(\lambda)_m$.

A sequence of Σ_n-representations (V_n, ψ_n) with linear maps $\psi_n \colon V_n \to V_{n+1}$ is called *consistent* if for all $\sigma \in \Sigma_n$ the following diagram commutes:

$$\begin{array}{ccc} V_n & \xrightarrow{\psi_n} & V_{n+1} \\ \sigma \downarrow & & \downarrow \sigma \\ V_n & \xrightarrow{\psi_n} & V_{n+1} \end{array}$$

The action of σ on V_{n+1} is through the inclusion $\Sigma_n \to \Sigma_{n+1}$ by fixing the last element $n+1$.

Definition. A consistent sequence of finite-dimensional Σ_n-representations (V_n, ψ_n) is called *(uniformly) representation stable* if for all n bigger than some fixed N we have that

(i) ψ_n is injective

(ii) the Σ_{n+1}-span of $\operatorname{im}\psi_n$ is equal to V_{n+1}

(iii) the decomposition $V_n = \bigoplus_\lambda c_{\lambda,n} V(\lambda)_n$ with $0 \leqslant c_{\lambda,n} \leqslant \infty$ is independent of $n \geqslant N$. △

Note that since $V(\lambda)_n$ is only defined for $n \geqslant |\lambda| + \lambda_1$ there are only finitely many partitions λ for which $c_{\lambda,n}$ is nonzero. This means that–for big n–we have a finite decomposition $V_n = \bigoplus_\lambda c_\lambda V(\lambda)_n$ where c_λ does not depend on n. We will only consider finite-dimensional $\mathbb{Q}$-vector spaces V_n, and in this case the concept of representation stability can be expressed in the language of FI-modules, which we now introduce.

Denote by **FI** the category with finite sets as objects and injections as morphisms.

Definition. An *FI-module* V *over a commutative ring* R is a covariant functor

$$V\colon \mathbf{FI} \to \mathbf{R{-}Mod},$$

where $\mathbf{R{-}Mod}$ is the category of R-modules. △

This functor encodes a sequence of Σ_n-representations because for every $n \in \mathbb{N}_0$ we get an R-module $V_n := V(\underline{\mathbf{n}})$ and for every permutation $\sigma\colon \underline{\mathbf{n}} \to \underline{\mathbf{n}}$ we get a linear automorphism $V(\sigma)$ of V_n. We have maps $V_n \to V_{n+1}$ induced by the canonical inclusions $\underline{\mathbf{n}} \hookrightarrow \underline{\mathbf{n+1}}$, so we get a consistent sequence (V_n) of Σ_n-representations.

Definition. An FI-module V is said to be *finitely generated* if there exists a finite set $S \subset \oplus_n V_n$ such that no proper sub-FI-module of V contains all elements of S. △

The following theorem connects the theory of FI-modules to representation stability.

Theorem 3.2 ([CEF12, p. 8])**.** *An FI-module V over a field of characteristic 0 is finitely generated if and only if the sequence (V_n) of Σ_n-representations is uniformly representation stable and each V_n is finite-dimensional.*

One major result is that for a finitely generated FI-module the dimension of V_n is eventually polynomial, i.e. there exists a polynomial p and some $N \in \mathbb{N}$ such that $\dim V_n = p(n)$ for all $n \geqslant N$.

Example 3.1. Let X be a topological space, then we want to investigate $\mathrm{Conf}_S(X)$ for finite sets S. The assignment of $\mathrm{Conf}_S(X)$ to a finite set S defines a contravariant functor $\mathrm{Conf}_?(X)\colon \mathbf{FI} \to \mathbf{Top}$ by sending $i\colon S \to T$ to $i^*\colon \mathrm{Conf}_T(X) \to \mathrm{Conf}_S(X)$ defined as $i^*((x_t)_{t\in T}) = \left(\left(x_{i(s)}\right)_{s\in S}\right)$. Composing with the (also contravariant) cohomology functor in some fixed degree we get an FI-module $H^k(\mathrm{Conf}_?(X); \mathbb{Q})$ over the rationals.

For manifolds, the situation is very nice: if X is a connected, oriented topological manifold (of finite type) of dimension at least two, then Theorem 6.2.1 in [CEF12] shows that the FI-module $H^k(\mathrm{Conf}_?(X); \mathbb{Q})$ is finitely generated.

Encouraged by this example one could hope that something similar happens with the FI-module $H^k(\mathrm{Conf}_?(\Gamma); \mathbb{Q})$ for a graph Γ, but this is even in the simplest cases far from true.

3.2 The Generalized Euler Characteristic of $\mathrm{Conf}_n(\Gamma)$

The representation ring $R_k(G)$ of a group G and a field k is defined to be $K_0(kG)$, i.e. the group completion of the monoid where the elements are isomorphism classes of finitely generated projective kG-modules and addition is the direct sum. The multiplication is the tensor product of kG-modules. We try to compute the Euler characteristic of $\mathrm{Conf}_n(\Gamma)$ with values in $R_{\mathbb{Q}}(\Sigma_n)$, which is defined to be

$$\chi^{\Sigma_n}(X) = \chi^{\Sigma_n}(X;\mathbb{Q}) = \sum_{i=0}^{\infty}(-1)^i\left[H^i(X;\mathbb{Q})\right] \in R_{\mathbb{Q}}(\Sigma_n)$$

for all Σ_n-spaces X for which the sum $\sum_{i=0}^{\infty}\dim_{\mathbb{Q}} H^i(X;\mathbb{Q})$ is finite and $[M]$ is the isomorphism class of the $\mathbb{Q}\Sigma_n$-module M in $K_0(\mathbb{Q}\Sigma_n)$. We need the condition on X for the well-definedness of this element since it implies that each individual $H^i(X;\mathbb{Q})$ is finite dimensional and that only a finite number of these modules are non-zero. Hence, in this whole chapter we restrict ourselves to finite graphs. Furthermore, all $\mathbb{Q}\Sigma_n$-modules are projective since by Maschke's theorem $\mathbb{Q}\Sigma_n$ is a semisimple ring (compare [CR62, Theorem 10.8, page 41]). As cohomology is homotopy invariant, this also holds for the generalized Euler characteristic χ^{Σ_n}, so it suffices to calculate $\chi^{\Sigma_n}(K_n\Gamma)$.

This generalized Euler characteristic can be calculated, the result is the following (where $\chi(Y)$ is the ordinary Euler characteristic of Y):

Theorem 3.3. *Let Γ be a finite graph, then $\chi^{\Sigma_n}(\mathrm{Conf}_n(\Gamma))$ is a direct sum of $\chi(\mathrm{UConf}_n(\Gamma))$ copies of the regular representation.*

Since this is true in a more general setup, we will prove the following proposition that implies this result.

Proposition 3.4. *Let G be a finite group, Y a finite free G-CW complex and k a field whose characteristic does not divide the order of G (e.g. $k = \mathbb{Q}, \mathbb{R}, \mathbb{C}, \ldots$), then*

$$\chi^G(Y;k) = \chi(Y/G)\cdot[kG] \in K_0(kG).$$

A G-*CW complex* is a CW complex together with a cellular G-action, i.e. an action that maps cells homeomorphically onto cells.

Proof. The assumptions on k are only necessary to ensure that Maschke's theorem holds. We first show that $\chi^G(Y;k)$ can be calculated by using the cochain complex of $\mathbb{Q}G$-modules

$$0 \xrightarrow{\delta^{-1}} C^0(Y;k) \xrightarrow{\delta^0} C^1(Y;k) \xrightarrow{\delta^1} \cdots \xrightarrow{\delta^{k-2}} C^{k-1}(Y;k) \xrightarrow{\delta^{k-1}} C^k(Y;k) \xrightarrow{\delta^k} 0.$$

By the first isomorphism theorem, we have $\operatorname{im}\delta^i \cong C^i(Y;k)/\ker\delta^i$, and by definition $H^i(Y;k) \cong \ker\delta^i/\operatorname{im}\delta^{i-1}$ as kG-modules. Since every kG-module is semisimple, complements of submodules exist, so we get for each $i \in \mathbb{Z}$

$$\left[C^i(Y;k)\right] = \left[C^i(Y;k)/\ker\delta^i\right] + \left[\ker\delta^i\right] = \left[\operatorname{im}\delta^i\right] + \left[H^i(Y;k)\right] + \left[\operatorname{im}\delta^{i-1}\right].$$

This yields the following calculation in $K_0(kG)$:

$$\begin{aligned}\sum_{i=0}^{\infty}(-1)^i\left[C^i(Y;k)\right] &= \sum_{i=0}^{\infty}(-1)^i\left(\left[\operatorname{im}\delta^i\right] + \left[H^i(Y;k)\right] + \left[\operatorname{im}\delta^{i-1}\right]\right)\\ &= \chi^G(Y;k) + \underbrace{\left[\operatorname{im}\delta^{-1}\right]}_{=0} + \sum_{i=0}^{k-1}\left(\left[\operatorname{im}\delta^i\right] - \left[\operatorname{im}\delta^i\right]\right) + (-1)^k\underbrace{\left[\operatorname{im}\delta^k\right]}_{=0}\\ &= \chi^G(Y;k).\end{aligned}$$

Clearly, the G-orbits of i-cells of Y correspond bijectively to the i-cells of Y/G, and each such orbit is a G-set isomorphic to G itself (considered as a G-set) by the freeness of the action. Hence, if we denote by n_i the number of i-cells of Y/G we get that $C^i(Y;k)$ is a free kG module on n_i generators, i.e.

$$\left[C^i(Y;k)\right] = n_i \cdot [kG].$$

Since $\chi(Y/G) = \sum_{i=0}^{\infty}(-1)^i n_i$ this shows the claim. □

Proof of Theorem 3.3. This follows from the last proposition, the only thing we have to show is that the action of Σ_n on $K_n\Gamma$ is free.

If $\sigma \neq 1$ is not the identity, then there exists some $j \in \underline{\mathbf{n}}$ such that $\sigma(j) = k \neq j$, and therefore for any $x \in K_n\Gamma$ we have $\iota(\sigma x)_j = \iota(x)_k \neq \iota(x)_j$, and since the map $\iota\colon K_n\Gamma \hookrightarrow \operatorname{Conf}_n(\Gamma)$ is an injection we have that $\sigma x \neq x$. Hence, no cell gets mapped to itself, so the action is free. □

3.3 Graphs With At Most One Branched Vertex

The propositions in this section characterize the cohomology of configuration spaces of graphs with at most one branched vertex *as* $\mathbb{Q}\Sigma_n$*-modules*.

Example 3.2. Let I be the unit interval $[0,1]$, then there is no branched vertex, so the only nontrivial cohomology group is $H^0(\operatorname{Conf}_n(I);\mathbb{Q})$ and coincides with the generalized Euler characteristic. The ordinary Euler characteristic of $\operatorname{Conf}_n(I)$ is the number of 0-cells in K_nI, which can be calculated by multiplying $n!$ with the number of 0-cells in UK_nI. Since there is exactly one such 0-cell, we get

Proposition.

$$H^i(\mathrm{Conf}_n([0,1]);\mathbb{Q}) \cong \begin{cases} \mathbb{Q}\Sigma_n & \text{if } i = 0 \\ 0 & \text{else} \end{cases}$$

Note that already this simple example produces an FI-module that is not finitely generated since the dimension is equal to $n!$, which does not grow polynomially.

Example 3.3. Consider the graph S^1 which consists of one vertex and one edge forming a loop. Then in the construction of our cube complex we agreed to call this vertex branched, despite of the fact that it has valency two. Therefore, our cube complex has only 0-cells and 1-cells. In UK_nS^1 we have two 0-cells, one where the branched vertex is occupied and one where it is not. We also have two 1-cells, one for each orientation of the edge. Hence, $\chi(K_nS^1) = n! \cdot (2-2) = 0$. If we choose some standard orientation on the edge then for every 0-cell in K_nS^1 we get an ordering on the set $\underline{\mathbf{n}}$. Moving along 1-cells changes this ordering only by a shift, so for every n-cycle we get exactly one connected component. The action of Σ_n permutes these cycles, so we have the following:

Proposition. *The* $\mathbb{Q}\Sigma_n$*-modules* $Y = H^0(\mathrm{Conf}_n(S^1);\mathbb{Q})$ *and* $H^1(\mathrm{Conf}_n(S^1);\mathbb{Q})$ *are isomorphic, the character*

$$\chi_Y(\sigma) = \#\{\tau \ n\text{-cycle} \mid \sigma\tau = \tau\}.$$

determines these modules uniquely. Their dimension is $(n-1)!$ *(as* $\mathbb{Q}$*-vector spaces).*

For an introduction to the concept of characters see [FH91].

Example 3.4. Consider the graph Y_k^l which is constructed by attaching k leaves and l loops to a single vertex and assume $k + 2l \geqslant 3$, compare Figure 3.1. All graphs with a single branched vertex of valency $\geqslant 3$ are homeomorphic to such a graph, hence after this example we have covered all graphs with at most one branched vertex.

Proposition 3.5.

$$\left[H^i\left(\mathrm{Conf}_n\left(Y_k^l\right);\mathbb{Q}\right)\right] \cong \begin{cases} [\mathbb{Q}] & \text{for } i = 0 \\ [\mathbb{Q}] - \chi\left(\mathrm{UConf}_n\left(Y_k^l\right)\right) \cdot [\mathbb{Q}\Sigma_n] & \text{for } i = 1 \\ 0 & \text{for } i \geqslant 2 \end{cases}$$

where $[\mathbb{Q}]$ *denotes the isomorphism class of the trivial representation. More explicitly, we have*

$$\left[H^1(\mathrm{Conf}_n(Y_k^l);\mathbb{Q})\right] = [\mathbb{Q}] + \frac{(n+k+l-2)!}{n!(k+l-1)!}(n(k+2l) - 2n - (k+l) + 1)[\mathbb{Q}\Sigma_n]$$

in $R_{\mathbb{Q}}(\Sigma_n)$.

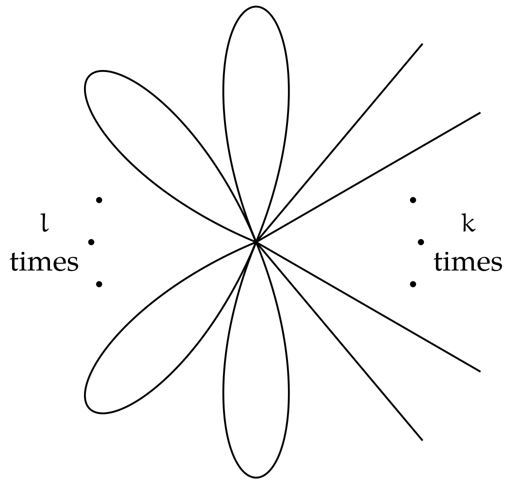

Figure 3.1: The graph Y_k^l with k leaves and l petals.

Proof. The zeroth cohomology group is the trivial 1-dimensional representation since the graph is connected and contains a vertex of valency at least three, the i-th cohomology for $i \geqslant 2$ vanishes since we only have one branched vertex and therefore the complex $K_n\, Y_k^l$ is only 1-dimensional. With Theorem 3.3 the first part of the proposition follows.

The factor in front of the regular representation is the same as $\chi(UK_nY_k^l)$, which we will now compute. There are two different kinds of 0-cells in $UK_nY_k^l$: one where the branched vertex is occupied and one where it is not. In the first case, we can distribute the remaining $n-1$ particles onto the $k+l$ edges arbitrarily, in the second case we have to distribute n particles. A simple combinatorial argument shows that the number of possibilities to distribute n particles onto K edges without ordering is $\binom{n+K-1}{K-1}$. To see this, put $n+K-1$ elements in a row and mark $K-1$ of them, then the marked elements divide the n unmarked elements into K parts. Thus, we have

$$\dim_{\mathbb{Q}} C^0(UK_nY_k^l;\mathbb{Q}) = \binom{(n-1)+(k+l)-1}{(k+l)-1} + \binom{n+(k+l)-1}{(k+l)-1}.$$

Every 1-cell has exactly one oriented edge s where $\tau(s)$ is the branched vertex, so there are exactly $k+2l$ choices. Since then there cannot be a particle on the branched vertex, we have to distribute the other $n-1$ particles onto the $k+l$ edges, leading to

$$\dim_{\mathbb{Q}} C^1(UK_nY_k^l;\mathbb{Q}) = (k+2l)\cdot\binom{(n-1)+(k+l)-1}{(k+l)-1}.$$

Altogether we get

$$\chi(\mathrm{UConf}_n(Y_k^l)) = \binom{n+(k+l)-1}{(k+l)-1} + (1-(k+2l))\binom{n+(k+l)-2}{(k+l)-1},$$

and rewriting the binomial coefficient with the identity

$$\binom{A}{B} = \frac{A!}{(A-B)!B!}$$

we get the given formula. □

Therefore, we have

$$\dim H^1(\mathrm{Conf}_n(Y_k^l);\mathbb{Q}) = 1 + \frac{(n+k+l-2)!}{(k+l-1)!}(n(k+2l) - 2n - (k+l) + 1),$$

which is for $l = 0$ the same result as given in [Ghr99]. Note that in particular for $n = 2$ and $k \geqslant 3$ we have at least one non-trivial cycle.

3.4 The General Case of Finite Graphs

Now let Γ be an arbitrary *finite* connected graph with $|B| > 1$ branched vertices. Then denote by k the number of unoriented edges and by k_b the number of oriented edges ending in $b \in B$. We only have to determine the number of i-cells in $UK_n\Gamma$ for all i in order to identify the generalized Euler characteristic. Since we are interested in the behavior for $n \to \infty$ we may assume $n \geqslant |B|$, the formulas for smaller n are obtained analogously. Fix $i \leqslant |B|$, then we need to choose i distinct branched vertices and for each of them an oriented edge ending in this vertex. For this, we have $\kappa_S^i := \sum_{\{b_j\}} k_{b_1} k_{b_2} \cdots k_{b_i}$ possibilities, the sum is taken over all subsets of B with cardinality i. For $i = 0$ this is evaluates to $\kappa_S^0 = 1$. As a result, this is a polynomial in the k_b and independent of n. The other $n - i$ particles have to be distributed onto the k edges and the remaining $|B| - i$ branched vertices. This can be done in

$$\kappa_n^i := \sum_{j=0}^{|B|} \binom{|B|-i}{j} \cdot \binom{(n-i-j)+k-1}{k-1}$$

different ways, which is a polynomial in n of degree $k - 1$. Hence, we have

Corollary. *For a finite connected graph with at least two branched vertices the map*

$$n \mapsto \chi(UK_n\Gamma) = \chi(U\mathrm{Conf}_n(\Gamma))$$

is given for $n \geqslant |B|$ *by a polynomial* p_Γ *of degree at most* $k - 1$*, where* k *is the number of unoriented edges of* Γ*. This yields*

$$\chi^{\Sigma_n}(\mathrm{Conf}_n(\Gamma)) = p_\Gamma(n) \cdot [\mathbb{Q}\Sigma_n],$$

so the number of regular representations in the generalized Euler characteristic is polynomial in n*. Furthermore, for each* j *the dimension of* $H^j(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ *is bounded from above by the product of* $n!$ *and a polynomial of degree at most* $k - 1$*.*

Additionally, we know that our configuration space is connected, so the zeroth cohomology $H^0(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is the trivial representation. In general, however, the module $H^i(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is not trivial for $i > 1$, so the Euler characteristic is not sufficient anymore to determine all cohomology groups.

For concrete calculations, it is useful to use the following formula:

$$\begin{aligned}
\chi(UK_n\Gamma) &= \sum_{i=0}^{|B|}(-1)^i \kappa_S^i \kappa_n^i \\
&= \underbrace{\frac{(n+k-1-|B|)!}{(k-1)!n!}}_{=:p_{|B|,k}(n)} \sum_{i=0}^{|B|}(-1)^i \kappa_S^i \kappa_n^i \frac{n!(k-1)!}{(n+k-1-|B|)!} \\
&= p_{|B|,k}(n) \sum_{i,j=0}^{|B|} \underbrace{(-1)^i \kappa_S^i \binom{|B|-i}{j}}_{\text{independent of } n} \frac{(n+k-1-(i+j))!}{(n+k-1-|B|)!} \frac{n!}{(n-(i+j))!}
\end{aligned}$$

Remark. The analogue of Theorem 3.3 is by Proposition 3.4 true whenever $\mathrm{Conf}_n(X)$ is homotopy equivalent to a finite Σ_n-CW complex, which is the case for example for compact smooth manifolds of dimension $\geqslant 2$. To construct this, take a finite CW structure on X (compare [Whi40]) yielding a finite CW structure on X^n, apply barycentric subdivision twice to the latter one and remove all cells meeting the fat diagonal. Since the cells are products of cells of X the cellular free Σ_n-action is clear.

The cohomology of these spaces is known to be representation stable, so the dimension of $H^j(\mathrm{Conf}_n(X);\mathbb{Q})$ grows polynomially for every fixed j. This seems contradictory to our result that the ordinary Euler characteristic grows factorially, but the cohomological dimension of $\mathrm{Conf}_n(X)$ grows in this case with n (by cohomological dimension we mean the smallest $m \geqslant 0$ such that the j-th cohomology is trivial for all $j > m$). Hence, the factorial growth in the alternating sum is achieved by adding more and more summands, which individually grow only polynomially. In the case where Γ is a finite graph, the cohomological dimension is for n big enough exactly $|B|$, the number of branched vertices. Therefore, all complexity has to come from finitely many summands and whenever $\chi(\mathrm{Conf}_n(X))$ is not eventually zero, at least one of the terms $\dim H^j(\mathrm{Conf}_n(X);\mathbb{Q})$ has to grow factorially. △

3.5 Applications

From this knowledge we can draw some conclusions. The first one is that for every graph satisfying the assumptions in the corollary with the additional property that $\chi(UK_n\Gamma) = 0$ holds only for finitely many n there exists an $i \geqslant 1$ such that the i-th cohomology is not representation stable. Indeed, assume that the cohomology is representation stable in all degrees, then the dimension of the cohomology grows polynomially, but since the cohomology is only for $i < |B|$ non-trivial the Euler characteristic would have to grow polynomially, too, which it does not.

Proposition 3.6. $H^b(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ *is non-trivial for every finite graph* Γ *as soon as* n *is at least* $2b$*, where* b *denotes the number of branched vertices.*

Proof. In the last section it was shown that the configuration space $\mathrm{Conf}_2(Y_k^0)$ contains for any $k \geqslant 3$ a non-trivial 1-cycle. From this it follows that a small neighborhood around a branched vertex looks like Y_k^0 for some k, so we have a non-trivial cycle which involves only two particles. For each branched vertex, choose such a non-trivial cycle making sure that all of them are disjoint. Then the product of these cycles yields a non-trivial element in the b-th homology since there are no $(b+1)$-cells, and by the isomorphism discussed earlier the b-th cohomology is non-trivial, too. □

In his dissertation Abrams asked the following question:

Question ([Abr00, Question 5.2]). *Are there graphs* Γ *such that* $C_n(\Gamma)$ *is homeomorphic to the surface of genus 2?*

The space $C_n(\Gamma)$ is a combinatorial version of the configuration space of a graph Γ. This graph can even be a colored graph, which gives a lot more spaces $C_n(\Gamma)$, but we do not want to discuss this construction here. Interesting for us is the case of an uncolored graph which is subdivided sufficiently such that $C_n(\Gamma)$ is a deformation retract of $\mathrm{Conf}_n(\Gamma)$. For more information on when this is the case we refer to Abrams dissertation.

Now assume that we have a finite graph Γ such that $C_n(\Gamma)$ is homotopy equivalent to the surface of genus 2 called S_2 and also a deformation retract of $\mathrm{Conf}_n(\Gamma)$. Then the configuration space is homotopy equivalent to S_2 and therefore we have

$$H^0(\mathrm{Conf}_n(\Gamma); \mathbb{Q}) \cong \mathbb{Q}, \quad H^1(\mathrm{Conf}_n(\Gamma); \mathbb{Q}) \cong \mathbb{Q}^4, \quad H^2(\mathrm{Conf}_n(\Gamma); \mathbb{Q}) \cong \mathbb{Q}$$

and all other cohomology groups are trivial. The Euler characteristic of $\mathrm{Conf}_n(\Gamma)$ is therefore -2. This already implies that n is at most 2, and since we have $\mathrm{Conf}_1(\Gamma) \cong \Gamma$ the only possibility is $n = 2$, so the generalized Euler characteristic of $\mathrm{Conf}_n(\Gamma)$ is the additive inverse of the regular representation $\mathbb{Q}\Sigma_2$. Now consider $K_2\Gamma$, which is homotopically the same as $\mathrm{Conf}_2(\Gamma)$, then there are no k-cells for $k \geqslant 3$. Since the space is connected and we need 2-dimensional cells our graph has at least two branched vertices and is connected itself. By the discussion in the previous chapters we can talk about homology instead of cohomology, which we will do as the description of cycles is simpler than the one of cocycles. It is easy to see that a 2-dimensional cycle in this space is a torus: on each 2-cell there is a unique direction in which the first particle moves and the second one is fixed. Moving along this direction now gives a loop in Γ described by the movement of the first particle. In the direction perpendicular to this the first particle is fixed and the second one moves, so this gives us a second loop. It is clear that these two loops are disjoint and hence we have a torus corresponding to the product of these two loops.

But this means that our graph Γ contains the dumbbell graph 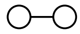 as a subgraph. However, this implies that we have in fact two different non-trivial 2-cycles, the one described above and the one where the particles are swapped. These two 2-cycles are definitely not homologous since there are no 3-cells. Hence, we have $\dim H^2(\mathrm{Conf}_2(\Gamma); \mathbb{Q}) \geqslant 2$, which contradicts the homotopy equivalence to S_2.

Since we did not cover Abrams construction we cannot answer his question completely, but we can say that *if* there exists such a graph then $C_n(\Gamma)$ is not even homotopy equivalent to $\mathrm{Conf}_n(\Gamma)$. Especially, the graph cannot be simple, since then for 2 particles $C_n(\Gamma)$ is a deformation retract of the configuration space, as is shown in Abrams dissertation.

3.6 Next Steps

From this point on, one could tackle the problem combinatorially since we know a finite cube complex with Σ_n-action that is an equivariant deformation retraction. The problem is that the formulas get very involved and the number of cells grows factorially in n, so calculations with a computer reach their limit very soon. In the appendix we present some results from such calculations, but even for very small graphs we do not get any further than seven particles.

For every n the regular representation contains *all* irreducible Σ_n-representations, so we cannot get such a nice description as we are used to in the world of representation stability. Nevertheless, it would be nice to have some kind of asymptotic behavior, in particular one could ask the following questions:

Questions. *(i) Is there any degree* $i > 0$ *such that* $\dim H^i(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ *grows polynomially?*

(ii) Does the number of regular representations in $H^i(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ *grow polynomially?*

(iii) If one disregards all regular representations from the i*-th cohomologies, is the remaining sequence representation stable? Or does one have an alternative asymptotic description for this sequence?*

(iv) Are there graphs such that $\chi(\mathrm{Conf}_n(\Gamma)) = 0$ *for many* n*? If yes, are there some of them where the growth rate of all* $H^i(\mathrm{Conf}_n(\Gamma); \mathbb{Q})$ *is strictly smaller than "polynomial times factorial"?*

Disregarding the regular representations means that two modules are considered equivalent if by adding a number of regular representation to the first module it becomes $\mathbb{Q}\Sigma_n$-isomorphic to the second one. From our calculations (we included some of them in the appendix) one would probably guess that the answer to the first two questions is negative in general, but at the moment we do not see how to prove this. The polynomial asked for in Question (ii) could have at most dimension $k-1$, where k is the number of edges in Γ, since the dimension of the homology $H^i(\mathrm{Conf}_n(\Gamma);\mathbb{Q})$ is bounded from above by a polynomial of degree at most $k-1$ multiplied with $n!$ by Theorem 3.3. If you now look at the calculations you do not find such a polynomial that works for all n in the list, so we can at least say that *if* the answer to this question is positive, then the polynomial growth happens only for a subsequence or for $n \gg 0$.

We do not think that the fourth question has a positive answer either, but probably we will not be able to deduce that directly from our results; although the formula for the polynomial p_Γ is relatively easy to compute for concrete graphs, to test whether it can be constantly zero for any graph Γ other than S^1 is difficult. In order to guess what the answer to the third question is one would probably need to do calculations for bigger n, but with the tools we used the cohomology is not computable anymore for any bigger number of particles.

Chapter 4

A Sheaf-Theoretic Approach

In [Chu12], Thomas Church showed that the cohomology of configuration spaces of manifolds is representation stable. The used technique–which is partly due to Burt Totaro's paper [Tot96]–will be applied to the case where the manifold is replaced by a finite graph. From the last chapter we know that our sequence of representations is not representation stable, but we still can try to gain some insight using these methods. Since this technique showed the desired result for arbitrary manifolds one could hope that it also works in our case, but we will see that the situation is much more complicated and sadly there is a point where we get stuck. Hopefully, someone will eventually come up with a solution, which is why we present our approach nevertheless.

4.1 Sheaf Theory

We shortly introduce the notions of sheaves and sheaf cohomology since we are going to use them later. All statements in this section that we use without proof can be found in [Bre97]. Denote for a topological space X by $\mathbf{Open}(\mathbf{X})$ the category whose objects are open subsets of X and where the only morphisms are the inclusions of open subsets.

Definition. A *presheaf (with values in abelian groups)* is a contravariant functor $\mathscr{F}\colon \mathbf{Open}(\mathbf{X}) \to \mathbf{Ab}$ such that $\mathscr{F}(\varnothing) = \{0\}$, where $\mathbf{Ab}$ is the category of abelian groups. For inc: $V \to U$ we denote $\mathscr{F}(\text{inc})$ by $\text{res}_{U,V}$ or $\text{res}^{\mathscr{F}}_{U,V}$. The elements $s \in \mathscr{F}(U)$ are called *sections of* $\mathscr{F}$ *over* U, we denote $\text{res}_{U,V}(s)$ as $s|_V$ and call it the section s restricted to V. A morphism between two presheaves $\alpha\colon \mathscr{F} \to \mathscr{G}$ is a natural transformation of functors, so for every $U \subset X$ open we have a morphism of abelian groups $\alpha(U)\colon \mathscr{F}(U) \to \mathscr{G}(U)$ such that for every open $V \subset U$ the following diagram commutes:

$$\begin{array}{ccc} \mathscr{F}(U) & \xrightarrow{\alpha(U)} & \mathscr{G}(U) \\ {\scriptstyle \mathrm{res}^{\mathscr{F}}_{U,V}}\big\downarrow & & \big\downarrow{\scriptstyle \mathrm{res}^{\mathscr{G}}_{U,V}} \\ \mathscr{F}(V) & \xrightarrow[\alpha(V)]{} & \mathscr{G}(V) \end{array}$$

△

Definition. A presheaf $\mathscr{F}$ over X is a *sheaf (of abelian groups)* if it satisfies the following two conditions for any open covering $(U_i)_{i \in I}$ of any open subset $U \subset X$:

(i) (locality) if $s, s' \in \mathscr{F}(U)$ are sections such that $s|_{U_i} = s'|_{U_i}$ for all $i \in I$, then $s = s'$

(ii) (gluing) if for all $i \in I$ there exists a section $s_i \in \mathscr{F}(U_i)$ such that $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ for all $i, j \in I$, then there exists a section $s \in \mathscr{F}(U)$ such that $s|_{U_i} = s_i$ for all i.

A *morphism of sheaves* is just a morphism of the underlying presheaves. △

Note that the element s in the gluing condition is unique by the locality condition. We can replace **Ab** by other categories like **R**−**Mod** for an arbitrary ring R or **Set**.

Example 4.1. Let A be an abelian group. Equip it with the discrete topology and set

$$\underline{A}(U) = \{f : U \to A \mid f \text{ locally constant}\}.$$

This defines a sheaf $\underline{A}$ that is called the *constant sheaf*.

For $x \in X$ denote by $I_x = \{U \subset X \mid U \text{ open}, x \in U\}$, then we define the *stalk* of a sheaf $\mathscr{F}$ at x to be

$$\operatorname{colim}_{U \in I_x} \mathscr{F}(U).$$

A morphism $\alpha : \mathscr{F} \to \mathscr{G}$ induces a morphism of abelian groups $\alpha_x : \mathscr{F}_x \to \mathscr{G}_x$ for all $x \in X$. Additionally, we get for all $x \in U \subset X$ open a map $\mathscr{F}(U) \to \mathscr{F}_x$ sending a section s to the element s_x. For every presheaf $\mathscr{F}$ we can define a sheaf $\mathscr{F}'$ and a morphism $\mathrm{inc} : \mathscr{F} \to \mathscr{F}'$ such that for every morphism $f : \mathscr{F} \to \mathscr{G}$ to a *sheaf* $\mathscr{G}$ there exists a unique morphism $f' : \mathscr{F}' \to \mathscr{G}$ such that $f = f' \circ \mathrm{inc}$:

$$\begin{array}{ccc} \mathscr{F} & \xrightarrow{f} & \mathscr{G} \\ {\scriptstyle \mathrm{inc}}\big\downarrow & \nearrow{\scriptstyle \exists! f'} & \\ \mathscr{F}' & & \end{array}$$

More explicitly, we set

$$\mathscr{F}'(U) = \left\{ (s_x, x \in U) \;\middle|\; \begin{array}{c} \text{for all } x \in U \text{ there exists a } V \subset U \text{ open with } x \in V \text{ and} \\ t \in \mathscr{F}(V) \text{ s.t. for all } y \in V : t_y = s_y \end{array} \right\}$$

for all open subsets $U \subset X$. This is a sheaf satisfying the mentioned universal property; we furthermore have $\mathrm{inc}_x \colon \mathscr{F}_x \to \mathscr{F}'_x$ is an isomorphism for all $x \in X$. We say that $\mathscr{F}'$ is the *sheafification* of $\mathscr{F}$.

Denote by $\mathbf{Sh}(X)$ the category whose objects are sheaves on the topological space X and whose morphisms are sheaf morphisms.

Example. Let $f \colon X \to Y$ be a continuous map, then we define

(i) the *direct image functor* $f_* \colon \mathbf{Sh}(X) \to \mathbf{Sh}(Y)$ by

$$f_*\mathscr{F}(U) = \mathscr{F}(f^{-1}(U)),$$

(ii) the *inverse image functor* $f^{-1} \colon \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$: it sends a sheaf $\mathscr{G}$ on Y to the sheafification of the presheaf

$$U \mapsto \mathrm{colim}_{V \supset f(U)} \mathscr{G}(V),$$

where the colimit is taken over all open sets $V \subset Y$ containing $f(U)$.

The direct image functor is right adjoint to the inverse image functor, so we have a natural isomorphism

$$\mathrm{Hom}_{\mathbf{Sh}(X)}(f^{-1}\mathscr{G}, \mathscr{F}) \xrightarrow{\cong} \mathrm{Hom}_{\mathbf{Sh}(Y)}(\mathscr{G}, f_*\mathscr{F}).$$

For two sheaves $\mathscr{F}, \mathscr{G}$ we define the direct sum $\mathscr{F} \oplus \mathscr{G}$ open set by open set, i.e. $(\mathscr{F} \oplus \mathscr{G})(U) = \mathscr{F}(U) \oplus \mathscr{G}(U)$ for all open sets U. It is immediately checked that this is again a sheaf. The kernel and cokernel of a sheaf morphism α are taken again open set by open set; the resulting presheaf for the kernel is already a sheaf, so we will denote it by $\ker(\alpha)$, the presheaf for the cokernel needs sheafification and the corresponding sheaf will be denoted by $\mathrm{coker}(\alpha)$.

Fact. *The category* $\mathbf{Sh}(X)$ *is an abelian category and we have that a sequence of sheaves* $0 \to \mathscr{F} \to \mathscr{G} \to \mathscr{H} \to 0$ *is exact if and only if the same holds for the sequences of abelian groups* $0 \to \mathscr{F}_x \to \mathscr{G}_x \to \mathscr{H}_x \to 0$ *for all* $x \in X$.

For every fixed open subset $U \subset X$ we have the *evaluation functor* $\mathrm{ev}_U \colon \mathbf{Sh}(X) \to \mathbf{Ab}$ sending $\mathscr{F}$ to $\mathrm{ev}_U(\mathscr{F}) := \mathscr{F}(U)$. It is a left-exact functor, so it takes exact sequences $0 \to \mathscr{F} \to \mathscr{G} \to \mathscr{H}$ to exact sequences $0 \to \mathscr{F}(U) \to \mathscr{G}(U) \to \mathscr{H}(U)$. The functor ev_X is called the *global section functor*.

To define sheaf cohomology, we first note that $\mathbf{Sh}(X)$ has enough injectives, i.e. for every sheaf $\mathscr{F}$ there exists an injective sheaf $\mathscr{I}$ and a monomorphism $\mathscr{F} \hookrightarrow \mathscr{I}$. An object I in an abelian category $\mathbf{C}$ is said to be *injective* if the hom functor

$$\mathrm{Hom}_{\mathbf{C}}(-, I) \colon \mathbf{C} \to \mathbf{Sets}$$

is exact. Hence, we get for every sheaf $\mathscr{F}$ an injective resolution $\mathscr{I}^\bullet$, namely an exact sequence

$$0 \to \mathscr{F} \to \mathscr{I}^0 \to \mathscr{I}^1 \to \mathscr{I}^2 \to \cdots$$

such that $\mathscr{I}^k$ is injective for all k. We apply the global section functor ev_X to the sequence $\mathscr{I}^\bullet = (\mathscr{I}^0 \to \mathscr{I}^1 \to \cdots)$ of injective objects (adding zeros to the left of $\mathscr{I}^0$) and take the homology of the resulting chain complex, then this defines the *sheaf cohomology of* X *with values in* $\mathscr{F}$. This process of taking an injective resolution, applying some left-exact functor F to it and taking the i-th homology yields the so-called i-th right derived functor of F applied to our sheaf $\mathscr{F}$. In other words:

Definition. Let $\mathscr{F}$ be a sheaf on X. The *sheaf cohomology of* X *with values in* $\mathscr{F}$ is defined to be the i-th right derived functor of the global section functor, i.e.

$$H^i_{\mathrm{sh}}(X; \mathscr{F}) := R^i \mathrm{ev}_X(\mathscr{F}). \qquad \triangle$$

This cohomology has the following useful property comparing it to singular cohomology:

Fact. *If* X *is a locally contractible space, then we have*

$$H^i_{\mathrm{sh}}(X; \underline{A}) \cong H^i_{\mathrm{sing}}(X; A)$$

for any abelian group A.

This is the point at which we can return to our problem since we now have seen that we can try to compute sheaf cohomology instead of singular cohomology.

4.2 The Spectral Sequence

To calculate the cohomology of $\mathrm{Conf}_n(X)$ for some topological space X with rational coefficients we therefore have to analyze the functor

$$\mathrm{ev}_{\mathrm{Conf}_n(X)} \colon \mathbf{Sh}(\mathrm{Conf}_n(X)) \to \mathbf{Ab}.$$

We have an inclusion map inc: $\mathrm{Conf}_n(X) \to X^n$ which yields the commuting diagram

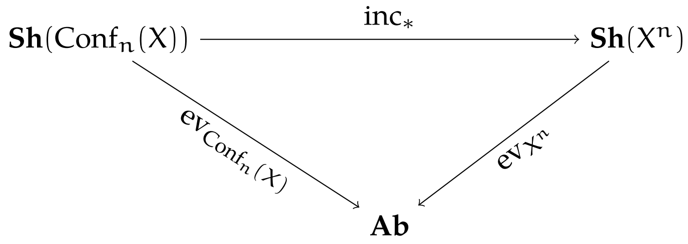

Since the space X^n is a lot easier to understand than our configuration space, we want to describe the right derived functors of $\mathrm{ev}_{\mathrm{Conf}_n(X)}$ in terms of the right derived functors of ev_{X^n} and inc_*. This description is given by the *Grothendieck spectral sequence*, which we now introduce. The definitions and facts about spectral sequences are taken from [McC01].

Definition. A *differential bigraded module* over a ring R is a collection of R-modules $\{E^{p,q}\}$ where p and q are integers, together with an R-linear mapping $d\colon E^{*,*} \to E^{*,*}$ called *differential* of bidegree $(s, 1-s)$ or $(-s, s-1)$ for some $s \in \mathbb{Z}$ with $d \circ d = 0$. △

For the map d to have bidegree (r, s) means that it is a collection of maps $d_{p,q}\colon E^{p,q} \to E^{p+r,q+s}$. We denote by $H^{p,q}(E^{*,*}, d)$ the homology of $E^{*,*}$ with respect to the boundary maps d at the point (p, q), so for the bidegree (r, s) we have

$$H^{p,q}(E^{*,*}, d) := \ker(d\colon E^{p,q} \to E^{p+r,q+s}) / \mathrm{im}(d\colon E^{p-r,q-s} \to E^{p,q}).$$

Definition. A *spectral sequence* is a collection of differential bigraded R-modules $\{E_r^{*,*}, d_r\}$ where $r \in \mathbb{N}$; the differentials are either all of bidegree$(-r, r-1)$ (for a spectral sequence of *homological type*) or all of bidegree $(r, 1-r)$ (for a spectral sequence of *cohomological type*) and for all p, q, r the module $E_{r+1}^{p,q}$ is isomorphic to $H^{p,q}(E_r^{*,*}, d_r)$. △

From now on we restrict ourselves to spectral sequences of cohomological type. We say that a spectral sequence $E_r^{p,q}$ converges to the final page $E_\infty^{p,q}$ if for all p, q there exists an r such that $d_r^{p-r,q+r-1}$ and $d_r^{p,q}$ are zero and $E_s^{p,q} \cong E_\infty^{p,q}$ for all $s \geqslant r$. If the final page is not given by a bigraded module but rather by a graded module E_∞^n then we mean by convergence that there exists a filtration $F^\bullet E_\infty^n$ such that the spectral sequence converges to

$$E_\infty^{p,q} := F^p E_\infty^{p+q} / F^{p+1} E_\infty^{p+q}.$$

A filtration $F^\bullet M$ of a module is a descending sequence of submodules $F^p M \supset F^{p+1} M$ of M indexed by $p \in \mathbb{Z}$ such that $\bigcap_p F^p M = \{0\}$ and $\bigcup_p F^p M = M$. Now we can examine the Grothendieck spectral sequence:

Theorem 4.1 ([McC01, Thm 12.10, p. 514]). *Let* $\mathbf{C_i}$ *be abelian categories for* $i = 1, 2, 3$. *Suppose the functors* $F\colon \mathbf{C_1} \to \mathbf{C_2}$ *and* $G\colon \mathbf{C_2} \to \mathbf{C_3}$ *are covariant,* G *is left exact and* F *takes injective objects in* $\mathbf{C_1}$ *to* G*-acyclic objects in* $\mathbf{C_2}$*. Then there is a spectral sequence with*

$$E_2^{p,q} \cong (R^p G)(R^q F(A)),$$

converging to $R^*(G \circ F)(A)$ *for* A *in* $\mathbf{C_1}$.

Corollary. *For a sheaf* $\mathscr{F}$ *the Grothendieck spectral sequence for* $\mathrm{ev}_{X^n} \circ \mathrm{inc}$ *has* E_2*-page*

$$E_2^{p,q} = R^p(\mathrm{ev}_{X^n}) \circ R^q(\mathrm{inc}_*)(\mathscr{F})$$

and converges to $H^*_{sh}(\mathrm{Conf}_n(X); \mathscr{F})$.

Proof of Corollary. We only need to show that inc_* takes injective sheaves to ev_{X^n}-acyclic objects. An object I in an abelian category $\mathbf{C}$ is called F-acyclic for a functor $F\colon \mathbf{C} \to \mathbf{D}$ to some category $\mathbf{D}$ if all higher right derived functors of F vanish, i.e. if $R^i F(I) = 0$ for all $i > 0$. Clearly, any injective object is F-acyclic for any functor F since we then have an injective resolution of length one; this fact is all we are going to use about acyclic objects.

Let $\mathscr{I}$ be an injective sheaf on $\mathrm{Conf}_n(X)$, then this is equivalent to saying that the hom functor

$$\mathrm{Hom}_{\mathbf{Sh}(\mathrm{Conf}_n(X))}(-, \mathscr{I})\colon \mathbf{Sh}(\mathrm{Conf}_n(X)) \to \mathbf{Set}$$

is exact. We now have by adjointness of inc^{-1} and inc_* a natural isomorphism

$$\mathrm{Hom}_{\mathbf{Sh}(X^n)}(-, \mathrm{inc}_* \mathscr{I})) \xrightarrow{\cong} \mathrm{Hom}_{\mathbf{Sh}(\mathrm{Conf}_n(X))}(\mathrm{inc}^{-1}(-), \mathscr{I}).$$

Since $f^{-1}\colon \mathbf{Sh}(Y) \to \mathbf{Sh}(X)$ is always exact for any continuous map $f\colon X \to Y$ we see that this composition is also exact. Hence, $\mathrm{inc}_* \mathscr{I}$ is injective and therefore ev_{X^n}-acyclic. □

If we take the sheaf $\mathscr{F}$ to be the constant sheaf $\underline{\mathbb{Q}}$, then this spectral sequence converges to the sheaf cohomology $H^*_{sh}(\mathrm{Conf}_n(X); \underline{\mathbb{Q}})$, so the remaining steps are to understand $R^p(\mathrm{ev}_{X^n}) \circ R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ for arbitrary p, q and then try to get some information from the spectral sequence. In particular, we now know that the i-th rational cohomology of $\mathrm{Conf}_n(X)$ is a subquotient of

$$\bigoplus_{p+q=i} R^p(\mathrm{ev}_{X^n}) \circ R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}}).$$

We have by definition of sheaf cohomology

$$R^p(\mathrm{ev}_{X^n})(R^q(\mathrm{inc}_*))(\underline{\mathbb{Q}}) = H^p_{sh}(X^n; R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})).$$

The problem now is that $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ is not necessarily a constant sheaf anymore, so we cannot compute this cohomology topologically. In [Chu12] and [Tot96] the successful approach for manifolds is to decompose this sheaf into several simpler sheaves. This is possible because the cohomology of configuration spaces of $\mathbb{R}^k$ for all $k \geqslant 0$ is known: every point on a manifold has a small neighborhood homeomorphic to $\mathbb{R}^k$, so locally one already knows the cohomology and one easily computes the stalks. Then, the stalks are "distributed" over a bunch of sheaves that are constant on a closed subset and zero on the rest. More explicitly, one takes one of the *generalized diagonals* of X^n for these closed subsets, which consists of all points $(x_1, \dots, x_n)$ such that $x_i = x_j$ for a fixed set of pairs $\{(i, j)\}$ (see below for a precise definition). The cohomology of these sheaves is computable topologically as they are constant sheaves over some closed subspace of X, and thus the complete E_2-page is known (except for the morphisms). This already suffices to prove representation stability with an explicit stable range, which is a very deep result.

This great success motivates us to decompose the spectral sequence in the case where X is not a manifold but a graph. The situation looks promising since locally, every graph is either a line or a star-shaped graph with one branched vertex. We know the cohomology of the configuration spaces of these types of graphs from the previous chapter, so we may easily compute all stalks. However, it turns out to be a problem that most graphs do not possess the property of being locally everywhere of the same shape, there are basically two different kinds of points: branched vertices and all other points. The configurations around these points behave very differently in these two cases, so we have to expect our sheaves to be much more complicated than in the situation of manifolds. Before we investigate these sheaves, we try to simplify the situation.

If we restrict ourselves to locally finite graphs Γ, a possible simplification could be to use the deformation retract $K_n\Gamma$ and its natural inclusion $j\colon K_n\Gamma \hookrightarrow \mathrm{Conf}_n(\Gamma) \hookrightarrow \Gamma^n$. This is a closed inclusion since j is again an embedding and $K_n\Gamma$ is closed in the Hausdorff space Γ^n by a proof analogous to the one in Chapter 2. In this case we have the following result:

Proposition 4.2. *If* $j\colon Z \to X$ *is a closed inclusion and* $\mathscr{F}$ *is a sheaf on* Z *then we have a natural isomorphism* $H^p_{sh}(X; j_*\mathscr{F}) \cong H^p_{sh}(Z; \mathscr{F})$.

Proof. First note that in this case j_* is an exact functor which is easily checked stalk by stalk. Use in that argument that $(j_*\mathscr{F})_{j(p)} = \mathscr{F}_p$ if $p \in Z$ and zero else. Let

$$0 \to \mathscr{F} \to \mathscr{I}^0 \to \mathscr{I}^1 \to \cdots$$

be an injective resolution, then we get by exactness of j_* the injective resolution (recall that j_* preserves injectivity)

$$0 \to j_*\mathscr{F} \to j_*\mathscr{I}^0 \to j_*\mathscr{I}^1 \to \cdots .$$

Now by definition we have $j_*\mathscr{I}^k(X) = \mathscr{I}^k(Z)$ and the result follows. □

The proof also shows that $R^k(j_*) = 0$ for all $k > 0$ by exactness of j_*. Hence, the only possibly non-trivial modules in the spectral sequence are

$$R^p(\mathrm{ev}_{\Gamma^n}) \circ R^0(j_*)(\mathbb{Q}) = H^p_{\mathrm{sh}}(\Gamma^n; j_*(\underline{\mathbb{Q}})) \cong H^p_{\mathrm{sh}}(K_n\Gamma; \underline{\mathbb{Q}}).$$

In conclusion, the E_2-page of the spectral sequence consists of one row of $H^p_{\mathrm{sh}}(K_n\Gamma; \underline{\mathbb{Q}})$ for all p and zeros everywhere else, so we gain no additional information from this page. This leaves us no other choice than using the original configuration space, which we will now do.

4.3 The Stalks

In this section we compute all stalks of $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ defined above. To do this, we first examine the so-called higher direct image functor $R^q(\mathrm{inc}_*)$.

Lemma 4.3. *Let $\psi\colon U \to X$ be a continuous map and $\mathscr{F}$ a sheaf on U, then the functor $R^q(\psi_*)(\mathscr{F})$ is the sheafification of the presheaf $H^q_{\mathrm{sh}}(\psi^{-1}(-); \mathscr{F})$.*

To see this, one just has to work out the definitions: construct an injective resolution $\mathscr{I}^\bullet$, apply ψ_* and take the q-th homology. Taking homology in the category of chain complexes of sheaves (not sheaf cohomology!) means taking it set by set–which is done by the presheaf $H^q_{\mathrm{sh}}(\psi^{-1}(-); \mathscr{F})$–and then sheafifying.

This result is good because stalks don't change in the process of sheafification, so we can calculate the stalks of $H^q_{\mathrm{sh}}(\mathrm{inc}^{-1}(-); \underline{\mathbb{Q}})$ and for every open U we have $H^q_{\mathrm{sh}}(\mathrm{inc}^{-1}(U); \underline{\mathbb{Q}}) \cong H^q_{\mathrm{sing}}(\mathrm{inc}^{-1}(U); \mathbb{Q})$, so we are back in the business of topology.

Let Γ be an arbitrary graph (not necessarily locally finite) and $\mathbf{x} \in \Gamma^n$ fixed, then we want to calculate the stalk of $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ at $\mathbf{x}$. Let $\mathbf{x} = (x_1, \ldots, x_n)$, then choose for every i an open neighborhood U_i in Γ such that

(i) U_i contains a branched vertex b if and only if $b = x_i$,

(ii) if $x_i = x_j$ then $U_i = U_j$, else $U_i \cap U_j = \varnothing$.

Then $U_{\mathbf{x}} = U_1 \times U_2 \times \cdots \times U_n$ is an open neighborhood of $\mathbf{x}$. The family of such neighborhoods is *cofinal* in the sense that every open neighborhood of $\mathbf{x}$ contains a neighborhood of this kind. Therefore,

$$R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})_{\mathbf{x}} \cong H^q_{\mathrm{sing}}\big(\mathrm{inc}^{-1}(U_{\mathbf{x}}); \mathbb{Q}\big) = H^q_{\mathrm{sing}}(\mathrm{Conf}_n(\Gamma) \cap (U_1 \times \cdots \times U_n); \mathbb{Q}).$$

Define the *valency* of a point $y \in \Gamma$ as follows: take a connected open neighborhood V of y such that $V - \{y\}$ contains no vertex of Γ, then the valency is the number of connected components of $V - \{y\}$. This defines a map $\mathrm{val}\colon \Gamma \to \mathbb{N}$ sending a point $y \in \Gamma$ to its valency.

We may consider $\mathbf{x}$ as a map $\mathbf{x}\colon \underline{\mathbf{n}} \to \tilde{\mathbf{x}} = \{x_1, \ldots, x_n\}$ sending i to x_i, then $\#\big(\mathbf{x}^{-1}(\{x\})\big)$ is the number of elements x_i that are equal to x. Note that $\tilde{\mathbf{x}}$ is $\mathbf{x}$ considered as an unordered subset of Γ and has cardinality at most n. Let $Y_k = Y_k^0$ be the star-shaped tree with k edges and one non-free vertex of valency k. Note that Y_1 and Y_2 are homeomorphic.

Choose for each i a homeomorphism $\xi_i\colon U_i \to Y_{\mathrm{val}(x_i)}$ and an embedding $\zeta_i\colon Y_{\mathrm{val}(x_i)} \hookrightarrow \Gamma$ such that:

(i) ζ_i sends the branched vertex to x_i

(ii) $x_i = x_j$ implies $\xi_i = \xi_j$ and $\zeta_i = \zeta_j$

(iii) the images of ζ_i and ζ_j are disjoint for $x_i \neq x_j$.

We consider $Y_{\mathrm{val}(x_i)}$ through ζ_i as a subspace of Γ and hence, the maps ξ_i induce a homeomorphism

$$\mathrm{Conf}_n(\Gamma) \cap U_{\mathbf{x}} \xrightarrow{\cong} \mathrm{Conf}_n(\Gamma) \cap \prod_{i=1}^{n} Y_{\mathrm{val}(x_i)} \xrightarrow{\cong} \prod_{x \in \tilde{\mathbf{x}}} \mathrm{Conf}_{\mathbf{x}^{-1}(x)}\big(Y_{\mathrm{val}(x)}\big).$$

We now have by the Künneth formula

$$\begin{aligned} H^q_{\mathrm{sing}}(\mathrm{Conf}_n(\Gamma) \cap U_{\mathbf{x}}; \mathbb{Q}) &\cong H^q_{\mathrm{sing}}\left(\prod_{x \in \tilde{\mathbf{x}}} \mathrm{Conf}_{\mathbf{x}^{-1}(x)}\big(Y_{\mathrm{val}(x)}\big); \mathbb{Q}\right) \\ &\cong \bigoplus_{\substack{\alpha\colon \tilde{\mathbf{x}} \to \mathbb{N}_0 \\ |\alpha| = q}} \bigotimes_{x \in \tilde{\mathbf{x}}} H^{\alpha(x)}_{\mathrm{sing}}\Big(\mathrm{Conf}_{\mathbf{x}^{-1}(x)}\big(Y_{\mathrm{val}(x)}\big); \mathbb{Q}\Big), \end{aligned} \tag{4.1}$$

where $|\alpha| = \sum_{x \in \tilde{\mathbf{x}}} \alpha(x)$. All these maps are Σ_n-equivariant with respect to the obvious Σ_n-actions (compare Section 1.2).

We have reduced the problem to understanding $H^i_{\mathrm{sing}}(\mathrm{Conf}_n(Y_k); \mathbb{Q})$, which we already did in the last chapter. There we showed that in the case of $i = 0$ this cohomology is the Σ_n-representation $\mathbb{Q}\Sigma_n$ for $k = 1$ or 2 and the trivial representation $\mathbb{Q}$ for all other $k \geqslant 0$. If we have $i = 1$ and $k \geqslant 3$ then we have

$$H^1_{\mathrm{sing}}(\mathrm{Conf}_n(Y_k); \mathbb{Q}) \cong \mathbb{Q} \oplus \frac{(n+k-2)!}{n!(k-1)!}(nk - 2n - (k-1)) \cdot \mathbb{Q}\Sigma_n.$$

In all other cases the cohomology vanishes, so in the formula above we only have to consider those $\alpha\colon \tilde{\mathbf{x}} \to \mathbb{N}_0$ whose image is contained in $\{0, 1\}$.

Altogether, this discussion yields a precise formula for all stalks of our sheaf $R^q(\mathrm{inc}_*)(\mathbb{Q})$. The concrete calculation, however, is in most cases pretty elaborate.

4.4 Decomposition of $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ into Easier Sheaves

We want to understand how the stalks $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})_x$ could be distributed over finitely many sheaves whose cohomology is easy enough to be computed or estimated. In the case of manifolds the sheaf is decomposed into sheaves whose support lies on *generalized diagonals*, see below for a definition. Restricted to their support, these sheaves are constant, so the cohomology is computable topologically.

In our case, however, we do not see any possibility of decomposing the sheaf into constant sheaves. The purpose of the rest of this section is to illustrate the difficulties that appear. Of course, these problems do not arise for every graph, for example calculating the sheaves for the graph consisting of only one point is certainly no challenge. Therefore, we simply assume that our graph Γ is sufficiently complicated.

Take the ordinary diagonal, i.e. the subset

$$\Delta := \{(x, x, \dots, x) \mid x \in \Gamma\} \subset \Gamma^n.$$

We want to examine the sheaf $\mathscr{F}^q := R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ restricted to Δ. Let $q = 0$, then the formulas above show

$$\mathscr{F}^0_{(x,\dots,x)} \cong \begin{cases} \mathbb{Q} & \text{if } x \text{ is branched} \\ \mathbb{Q}\Sigma_n & \text{otherwise.} \end{cases}$$

This is a sheaf whose stalks are always $\mathbb{Q}^{n!}$ except at the finite closed set of branched vertices. Hence, for $n > 1$ the sheaf is not constant. There is no way to decompose it into finitely many constant sheaves over closed subsets of Δ: the support of a sheaf is always closed, so there is no possibility that the stalk at every point arbitrarily close to a branched vertex has dimension $n! > 1$.

For $q = 1$ the situation is better: then $\mathscr{F}^1$ is a sum of skyscraper sheaves, one for each branched vertex.

If we now have a disjoint decomposition $J_1 \cup \dots \cup J_k = \underline{\mathbf{n}}$ into non-empty sets J_i then we can look at $\mathscr{F}^q$ restricted to the *generalized diagonal corresponding to* $\{J_i\}$, namely

$$\Delta_{\{J_i\}} = \{(x_1, \dots, x_n) \mid x_j = x_k \text{ if } j, k \in J_i \text{ for some } i\} \subset \Gamma^n.$$

As soon as $k < n$, the situation is very similar to the case of the ordinary diagonal: we cannot decompose $\mathscr{F}^0$ into finitely many constant sheaves over closed subsets of $\Delta_{\{J_i\}}$. This problem even affects the case $q = 1$: in Equation (4.1) we see that on some generalized diagonals the zeroth cohomology appears because each α meets 1 exactly once. This is the case for all generalized diagonals corresponding to $\{J_i\}$ where at least two of the sets J_i contain two or more elements.

In conclusion, it would be a huge step towards a proof to find a nice decomposition of sheaves $\mathscr{G}$ on Γ which satisfy

$$\mathscr{G}_x \cong \begin{cases} \mathbb{Q} & \text{if } x \text{ is branched} \\ \mathbb{Q}^k & \text{otherwise.} \end{cases}$$

for a fixed $k > 1$ such that we can calculate the sheaf cohomology of the individual summands. Unfortunately, we did not manage to do this. The difficult combinatorics make it unlikely that one is able to calculate this cohomology by explicit techniques like flasque resolutions or Čech cohomology, so at the moment we do not know how one should proceed from here on.

4.5 Sketch of the Proof for Manifolds

This last section sketches the rest of the proof in the case where X is a topological manifold of dimension at least 2 which is connected, orientable and of finite type. For more details see [Chu12] and [Tot96].

The first goal is to decompose $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ into locally constant sheaves. Let $x \in X^n$ be a point, then every open neighborhood of x contains a neighborhood of x which is a product of n balls, each of them homeomorphic to $\mathbb{R}^d$, where d denotes the dimension of the manifold X. By considerations similar to those in the last section we thus may assume that this neighborhood intersected with $\mathrm{Conf}_n(X)$ is homeomorphic to

$$\mathrm{Conf}_I(\mathbb{R}^d) := \mathrm{Conf}_{i_1}(\mathbb{R}^d) \times \mathrm{Conf}_{i_2}(\mathbb{R}^d) \times \cdots \times \mathrm{Conf}_{i_k}(\mathbb{R}^d)$$

for some $1 \leqslant k \leqslant n$ and some partition $I = \{i_1, \ldots, i_k\}$ of $\underline{\mathbf{n}}$ into k subsets.

The configuration space of Euclidean space $\mathbb{R}^d$ is well-known, we will use the description of $H^*_{\mathrm{sing}}(\mathrm{Conf}_m(\mathbb{R}^d); \mathbb{Z})$ as algebra given in Theorem 4.2 of [FH01, p. 105] throughout the rest of this section. In particular, this tells us that $H^q_{\mathrm{sing}}(\mathrm{Conf}_I(\mathbb{R}^d); \mathbb{Q})$ can only be non-trivial for q divisible by $(d-1)$. Thus, for q not of this form every stalk of $R^q(\mathrm{inc}_*)(\underline{\mathbb{Q}})$ is zero, so the whole sheaf is trivial.

The top-dimensional cohomology of $\mathrm{Conf}_I(\mathbb{R}^d)$ is in degree $(n-k)(d-1)$, where k is the size of I. We want to see that every class in $H^{r(d-1)}_{\mathrm{sing}}(\mathrm{Conf}_I(\mathbb{R}^d); \mathbb{Q})$ is represented by some top-dimensional cohomology class, even for $r < n-k$, and by refining the partition this is indeed possible (even with $\mathbb{Z}$ coefficients):

Lemma 4.4 ([Tot96, Lemma 3, p. 5]). *Let $I = \{i_1, \ldots, i_k\}$ be a partition of $\underline{\mathbf{n}}$ and $r \geqslant 0$, then the canonical map*

$$\bigoplus_J H^{r(d-1)}_{\mathrm{sing}}\big(\mathrm{Conf}_J(\mathbb{R}^d); \mathbb{Z}\big) \to H^{r(d-1)}_{\mathrm{sing}}\big(\mathrm{Conf}_I(\mathbb{R}^d); \mathbb{Z}\big),$$

where the sum is taken over all partitions J of $\underline{\mathbf{n}}$ into $n-r$ subsets refining I, is an isomorphism.

Note that in the source of this map we only have top-dimensional cohomology classes, so every element in the target can be written as a top-dimensional cohomology class with respect to some refined partition, as claimed. Now we have that $H^{(n-1)(d-1)}_{\mathrm{sing}}\big(\mathrm{Conf}_n(\mathbb{R}^d);\mathbb{Q}\big) \cong \mathbb{Q}^{(n-1)!}$, so it follows by the Künneth formula

$$H^{r(d-1)}_{\mathrm{sing}}\big(\mathrm{Conf}_J(\mathbb{R}^d);\mathbb{Q}\big) \cong \mathbb{Q}^{c_J},$$

where $c_J := (j_1-1)!\cdots(j_{n-r}-1)!$ for a partition J of $\underline{\mathbf{n}}$ into $n-r$ subsets.

Corollary. $R^{r(d-1)}(\mathrm{inc}_*)\big(\underline{\mathbb{Q}}\big)_x \cong \bigoplus_J \mathbb{Q}^{c_J}$ *for* $0 \leqslant r < n$, *otherwise it is trivial.*

If now $\underline{\mathbb{Q}}^{c_J}_J$ denotes the pushforward of the locally constant sheaf $\underline{\mathbb{Q}^{c_J}}$ on the generalized diagonal

$$X^{n-r}_J := \big\{(x_1,\dots,x_n) \in X^n \mid x_i = x_j \text{ if } i \text{ and } j \text{ belong to the same element of } J\big\}$$

via the inclusion $X^{n-r}_J \subset X^n$, then one can show that

$$R^{r(d-1)}(\mathrm{inc}_*)\big(\underline{\mathbb{Q}}\big) \cong \bigoplus_{|J|=n-r} \underline{\mathbb{Q}}^{c_J}_J$$

for all $0 \leqslant r < n$. To do this, one constructs a map which induces an isomorphism on stalks, we won't discuss this here.

This shows for all $0 \leqslant r < n$:

$$\begin{aligned} H^p_{\mathrm{sh}}\Big(X^n; R^{r(d-1)}(\mathrm{inc}_*)(\underline{\mathbb{Q}})\Big) &\cong H^p_{\mathrm{sh}}\Big(X^n; \oplus_J \underline{\mathbb{Q}}^{c_J}_J\Big) \\ &\cong \bigoplus_J H^p_{\mathrm{sh}}\Big(X^n; \underline{\mathbb{Q}}^{c_J}_J\Big) \\ &\cong \bigoplus_J H^p_{\mathrm{sh}}\big(X^{n-r}_J; \underline{\mathbb{Q}^{c_J}}\big) \\ &\cong \bigoplus_J H^p_{\mathrm{sing}}\big(X^{n-r}_J; \mathbb{Q}^{c_J}\big) \\ &\cong \bigoplus_J H^p_{\mathrm{sing}}\big(X^{n-r}_J; \mathbb{Q}\big) \otimes \mathbb{Q}^{c_J}. \end{aligned}$$

The second isomorphism comes from the fact that the sum is finite, the third one uses Proposition 4.2, the fourth isomorphism uses that the sheaves are locally constant. This discussion shows the following proposition:

Proposition 4.5. *The* E_2*-page of the Grothendieck spectral sequence for* $\mathrm{ev}_{X^n} \circ \mathrm{inc}$ *converging to* $H^*_{\mathrm{sh}}(\mathrm{Conf}_n(X);\underline{\mathbb{Q}})$ *has the following form:*

$$E^{p,q}_2 \cong \begin{cases} \bigoplus_{|J|=n-r} H^p_{\mathrm{sing}}\Big(X^{n-r}_J;\mathbb{Q}\Big) \otimes \mathbb{Q}^{c_J} & \text{if } q = r(d-1) \text{ for } 0 \leqslant r < n \\ 0 & \text{else.} \end{cases}$$

Hence, it looks as follows:

$$
\begin{array}{cccc}
\vdots & \vdots & \vdots & \\
r(d-1) & \bigoplus_J H^0_{\mathrm{sing}}\left(X_J^{n-r};\mathbb{Q}\right)\otimes\mathbb{Q}^{c_J} & \bigoplus_J H^1_{\mathrm{sing}}\left(X_J^{n-r};\mathbb{Q}\right)\otimes\mathbb{Q}^{c_J} & \cdots \\
\vdots & \vdots & \vdots & \\
1\cdot(d-1) & \bigoplus_J H^0_{\mathrm{sing}}\left(X_J^{n-1};\mathbb{Q}\right) & \bigoplus_J H^1_{\mathrm{sing}}\left(X_J^{n-1};\mathbb{Q}\right) & \cdots \\
0 & H^0_{\mathrm{sing}}(X^n;\mathbb{Q}) & H^1_{\mathrm{sing}}(X^n;\mathbb{Q}) & \cdots
\end{array}
$$

Note that everything we did worked equivariantly (with the accordingly defined actions on the different spaces involved). Hence, we get a very explicit description in terms of the cohomology of the generalized diagonals which can be used to prove representation stability.

To do this, one looks at the map $E_2(n) \to E_2(n+1)$ induced by the map $\mathrm{Conf}_{n+1}(X) \to \mathrm{Conf}_n(X)$ forgetting the last point, through which one turns all modules and E_i-pages into FI-modules. From the explicit description one then can see that the FI-modules $E^{(0,d-1)}$ and $E^{(p,0)}$ for $p \geqslant 0$ are finitely generated. As for fixed n the bigraded algebra $E_2(n)$ is generated by $E^{(0,d-1)}(n)$ and $E^{(p,0)}(n)$ for $p \geqslant 0$, the FI-module E_2 is finitely generated, too.

Since the infinity page E_∞ is a subquotient of the E_2-page, it is also finitely generated and the underlying sequence of representations is representation stable. For more details, we refer the reader to [Tot96], [Chu12] and [CEF12].

Bibliography

[Abr00] Aaron David Abrams. *Configuration Spaces and Braid Groups of Graphs*. PhD thesis, UC Berkley, 2000.

[BH99] Martin R. Bridson and André Häfliger. *Metric Spaces of Non-Positive Curvature*. Springer, 1999.

[Bre93] Glen E. Bredon. *Topology and geometry.* New York: Springer-Verlag, 1993.

[Bre97] Glen E. Bredon. *Sheaf theory. 2nd ed.* New York, NY: Springer, 1997.

[Bro82] Kenneth S. Brown. *Cohomology of groups.* Springer-Verlag., 1982.

[CEF12] Thomas Church, Jordan S. Ellenberg, and Benson Farb. FI-modules: a new approach to stability for Sn-representations. arXiv:1204.4533 [math.RT], 2012.

[CF13] Thomas Church and Benson Farb. Representation theory and homological stability. *Adv. Math.*, 245:250–314, 2013.

[Chu12] Thomas Church. Homological stability for configuration spaces of manifolds. *Invent. Math.*, 188(2):465–504, 2012.

[CR62] Charles W. Curtis and Irving Reiner. *Representation theory of finite groups and associative algebras.* Interscience Publishers New York, reprint of the 1962 original edition, 1962.

[FH91] W. Fulton and J. Harris. *Representation Theory: A First Course*. Graduate Texts in Mathematics / Readings in Mathematics. Springer New York, 1991.

[FH01] Edward R. Fadell and Sufian Y. Husseini. *Geometry and topology of configuration spaces.* Berlin: Springer, 2001.

[Ghr99] Robert Ghrist. Configuration spaces and braid groups on graphs in robotics. arXiv:math/9905023, 1999.

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge: Cambridge University Press, 2002.

[McC01] John McCleary. *A user's guide to spectral sequences. 2nd ed.* Cambridge: Cambridge University Press, 2nd ed. edition, 2001.

[PM] Paweł Pilarczyk and Marian Mrozek. Chomp homology calculation. `http://chomp.rutgers.edu/Projects/Computational_Homology/OriginalCHomP/software/`.

[tom08] Tammo tom Dieck. *Algebraic topology.* Zürich: European Mathematical Society (EMS), 2008.

[Tot96] Burt Totaro. Configuration spaces of algebraic varieties. *Topology*, 35(4):1057–1067, 1996.

[Whi40] J.H.C. Whitehead. On C^1-complexes. *Ann. Math. (2)*, 41:809–824, 1940.

[Św01] Jacek Świątkowski. Estimates for homological dimension of configuration spaces of graphs. *Colloquium Mathematicum*, 89(1):69–79, 2001.

Appendix

Here we give the results of some explicit calculations of $H^i(\mathrm{Conf}_n(\Gamma);\mathbb{Z})$ for different Γ. None of the results contains any torsion, which is why the tables only contain the ranks of the free $\mathbb{Z}$-modules. The calculations were done with the help of the wonderful homology calculation software CHomP [PM] by Paweł Pilarczyk and Marian Mrozek.

$\Gamma =$ 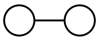							
n	1	2	3	4	5	6	7
rank H^1	2	7	25	111	601	3 783	27 077
rank H^2	0	2	12	86	720	6 662	67 396
$\Gamma =$ 							
rank H^1	0	3	31	247	1 911	15 531	?
rank H^2	0	0	0	6	230	5 450	?
$\Gamma =$ 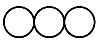							
rank H^1	3	13	55	255	1 331	7 983	?
rank H^2	0	2	30	374	4 450	53 342	?
$\Gamma =$ 							
rank H^1	2	5	13	32	70	?	?
rank H^2	0	0	0	7	189	?	?