Bull. Aust. Math. Soc. (First published online 2025), page 1 of 14*
doi:10.1017/S000497272400131X
*Provisional—final page numbers to be inserted when paper edition is published

FORBIDDEN PARTITION CONFIGURATION SPACES OF GRAPHS

JAMES DOVER

(Received 26 November 2024; accepted 3 December 2024)

Abstract

We propose a generalised version of configuration spaces defined by disallowing combinations of simultaneous collisions among the n points determined by a family of forbidden partitions. In the case where the underlying space is a finite graph, we construct a cubical complex with the same homology as this configuration space.

2020 *Mathematics subject classification*: primary 55R80; secondary 55N99, 57Q70.

Keywords and phrases: configuration spaces, cell complexes, discrete Morse theory, Mayer–Vietoris spectral sequence.

1. Introduction

The ordered configuration space, introduced in [6], models the possible arrangements of n particles in an underlying space X, where two particles are not permitted to occupy the same position simultaneously. The notion can be broadened to allow for certain collisions but not others. The no-k-equal configuration space (see for example [3, 5]) allows up to $k-1$ particles to occupy the same position at one time. In a coloured configuration space (the 'exotic' configuration spaces in [2, 8]), each particle is assigned a colour, and a collision is allowed or disallowed based on the colour composition of the particles involved. In these variations, it is still always the case that a configuration is disallowed due to a single collision, meaning a set of particles occupying one position.

In this note, we generalise these ideas with a version of configuration space in which a specific arrangement of particles can be disallowed based on the presence and compositions of separate simultaneous collisions. For example, this new version allows us to define a no-k-collision configuration space, where an arrangement is disallowed if it includes k or more pairs of colliding particles. More generally, if a collision between certain particles introduces some amount of interference, the configurations may be limited by the total amount of interference in the entire system.

Any collection of simultaneous collisions induces a partition on the set of particles. Our generalised configuration space is defined by specifying an appropriate collection of forbidden partitions. The space consists of all arrangements where the collisions do not conform to any of these forbidden partitions.

After defining this new forbidden partition configuration space, we will consider the case where the underlying space is a finite graph. Here, we adapt the approach used in [5] to construct a cubical complex as a discrete model for the homology, generalising the models of [1] for the nth configuration space and [5] for the no-k-equal space.

For $n \in \mathbb{N}$, let $N = \{1, \ldots, n\}$. A partial order on the partitions of N is given by the rule: $P \leq Q$ if and only if, for all $i, j \in N$, we have $(i \sim_P j \Rightarrow i \sim_Q j)$.

Let $\mathcal{P}$ be a collection of partitions of N such that, for partitions P and Q of N if $P \in \mathcal{P}$ and $P \leq Q$, then $Q \in \mathcal{P}$. We say $\mathcal{P}$ is an upper set of partitions and always assume it does not contain the discrete partition $\{\{1\}, \{2\}, \ldots, \{n\}\}$.

Let X be a topological space. For any configuration $x = (x_1, \ldots, x_n) \in X^n$, there is an induced partition of N, $P = \pi(x)$, given by $i \sim_P j$ if and only if $x_i = x_j$.

We define $\mathrm{Conf}_n(X, \mathcal{P}) := \{x \in X^n \mid \pi(x) \notin \mathcal{P}\}$, the $\mathcal{P}$-configuration space of X.

Example 1.1 (No-k-equal configurations). For a given $k \geq 2$, let $\mathcal{P}$ be the collection of partitions P such that there exists $i \in N$ with $|[i]_P| \geq k$. Then, $\mathrm{Conf}_n(X, \mathcal{P})$ is the no-k-equal space of X.

Example 1.2 (Coloured configurations). Assign to each point indexed by N one of m colours and, for a subset $A \subseteq N$, let $c(A) \in \mathbb{N}^m$ denote the vector counting the number of points with indices in A having each colour. Let $\mathcal{I}$ be an ideal of colour count vectors in permitted collisions. Define $\mathcal{P}$ as the collection of partitions P such that there exists $i \in N$ with $c([i]_P) \notin \mathcal{I}$. Then, $\mathrm{Conf}_n(X, \mathcal{P})$ gives the coloured configuration space from [2].

Example 1.3 (Mass configurations). If x_i represents the location in X of a particle with nonnegative mass m_i, let $M \geq \max\{m_i \mid i \in N\}$. Define $\mathcal{P}$ as the collection of partitions P such that there exists $i \in N$ with $\sum_{j \in [i]_P} m_j > M$. Then, $\mathrm{Conf}_n(X, \mathcal{P})$ gives the space of configurations of n particles, where the total mass of particles located at any point of X is at most M.

Example 1.4. Let $N = \{1, 2, 3, 4\}$ and $\mathcal{P}$ be the set of partitions $\{\{1, 2\}, \{3, 4\}\}$ and $\{\{1, 2, 3, 4\}\}$. Then, $\mathrm{Conf}_4(X, \mathcal{P})$ consists of all (x_1, x_2, x_3, x_4) except those of the form (a, a, b, b). In words, collisions are allowed as long as there is no collision of x_1 with x_2 at the same time as x_3 with x_4.

Example 1.5 (No-k-collision configurations). For a given $k \geq 1$, let $\mathcal{P}$ be the collection of partitions P such that the number of parts in P is $\leq N - k$. Then, $\mathrm{Conf}_n(X, \mathcal{P})$ is the space of all configurations having fewer than k pairs of coincident points $x_i = x_j, i < j$.

Example 1.6 (Configurations with an interference threshold). For each pair $i < j$, let $f_{ij} \geq 0$ be the amount of interference created if x_i coincides with x_j. Let $M > 0$

be the maximum amount of interference allowed in the system. Let $\mathcal{P}$ be the collection of partitions P such that $\sum_{i\sim_P j} f_{ij} > M$. Then, $\mathrm{Conf}_n(X, \mathcal{P})$ gives the space of configurations with an allowed level of interference.

In the remaining sections, we study the forbidden partition configuration space $\mathrm{Conf}_n(\Gamma, \mathcal{P})$ of a graph Γ and introduce a cubical complex $\mathcal{D}_n(\Gamma, \mathcal{P})$ as a discrete model. We will define open covers for both that induce homological equivalence between the two spaces via the associated Mayer–Vietoris spectral sequence (see [4]).

2. $\mathcal{P}$-configuration space of a graph

Let $N = \{1, \dots, n\}$ and $\mathcal{P}$ be an upper set of partitions of N.

Let Γ be a finite, connected graph with no loops and at least one vertex of degree ≥ 2. Each edge is given a length of 1, inducing the path metric d on Γ. We give Γ^n and its subspace $\mathrm{Conf}_n(\Gamma, \mathcal{P})$ the metric $\delta(x, y) = \max\{d(x_i, y_i) \mid i \in N\}$.

An *open star* V_t in Γ consists of a single vertex v_t of degree ≥ 2 (called the *central vertex*) together with all of its incident edges, excluding their second vertices if they also have degree ≥ 2. Let $\mathcal{V} = \{V_1, V_2, \dots, V_s\}$ be the collection of open stars covering Γ. (As degree 1 vertices are considered part of the open star, we also consider them part of the interior of an edge.)

We will now construct an open cover of the $\mathcal{P}$-configuration space $\mathrm{Conf}_n(\Gamma, \mathcal{P})$. Let σ be a choice of open star V_t for each $i \in N$. Let $\pi(\sigma)$ be the partition of N induced by σ: $i \sim j$ if and only if σ assigns i and j to the same open star. Choose a refinement $P \leq \pi(\sigma)$ which is not in $\mathcal{P}$. Choose a partial order ω of N that restricts to a total order of the indices assigned by σ to each open star V_t and groups together members of equivalence classes of P. (If $a \sim b$ and c is between a and b in the total order, then $a \sim c$.) We use $\omega_t(i)$ (or simply $\omega(i)$ if t is understood) to denote the ith index in the total order on the star V_t.

A *condition* is a choice λ of compatible σ, P and ω as above. A configuration $x = (x_1, \dots, x_n) \in \mathrm{Conf}_n(\Gamma, \mathcal{P})$ *conforms* to $\lambda = (\sigma, P, \omega)$ if:

- for each $i \in N$, x_i lies in the open star $\sigma(i) = V$ with central vertex v;
- if x_i and x_j lie on the same edge of Γ and $\sigma(i) \neq \sigma(j)$, then $d(x_i, v) < d(x_j, v)$, where v is the central vertex of $\sigma(i)$;
- for a given open star V, if $i < j$, $\omega(i) \nsim_P \omega(j)$ and $x_{\omega(i)}$ and $x_{\omega(j)}$ lie on the same edge, then $d(x_{\omega(i)}, v) < d(x_{\omega(j)}, v)$.

We note that, in a configuration conforming to λ, two points x_i and x_j are only allowed to be equal if i and j are assigned to the same star and $i \sim_P j$.

Let U_λ be the set of configurations conforming to λ, and $\mathcal{U}$ be the collection $\{U_\lambda\}_\lambda$.

PROPOSITION 2.1. *$\mathcal{U}$ is a finite, open cover of* $\mathrm{Conf}_n(\Gamma, \mathcal{P})$.

PROOF. There are finitely many possible conditions λ, so $\mathcal{U}$ is finite.

Given any $x \in \mathrm{Conf}_n(\Gamma, \mathcal{P})$, we may define a condition λ as follows. Assign i to the open star whose central vertex is closest to x_i; if x_i is at the midpoint of an edge between two vertices of degree ≥ 2, choose either open star for i and for all j for which $x_j = x_i$. Let $P = \pi(x)$. Define the total order on the indices assigned to V_t by increasing distance from v_t; if multiple x_j have the same distance from v_t, choose any order for those indices that groups together indices of coincident points. By construction, $x \in U_\lambda$.

We show that each U_λ is open. For $x \in U_\lambda$, define $\epsilon_0 = \min\{d(x_i, x_j) \mid x_i \neq x_j\}$, and define $\epsilon_t = \min\{d(x_i, v_t) \mid i \nsim_P \omega_t(1)\}$ for each $t \in \{1, \dots s\}$. That is, for each open star V_t, ϵ_t is the minimum distance to its central vertex of points in the configuration which are not in the same part of P as the first index in the order ω_t. Now let $\epsilon = \min\{1, \epsilon_0, \epsilon_1, \dots, \epsilon_s\}$. We show the open ball $B(x, \epsilon/3)$ is contained in U_λ. Let y be a configuration with $\delta(y, x) < \epsilon/3$. Then, $d(y_i, x_i) < \epsilon/3$ for all $i \in N$. We need to show that y conforms to λ.

Since $d(y_i, x_i) < \epsilon/3 < 1$, then y_i and x_i must lie on the same edge or on two edges that share a vertex. If they lie on the same edge, then y_i lies in the same open star $\sigma(i)$ as x_i. If they lie on two edges that share the central vertex of $\sigma(i)$, then y_i lies in $\sigma(i)$. Assume that they lie on two adjacent edges that do not share the central vertex of $\sigma(i)$. Then, they share a vertex with degree ≥ 2 which is the central vertex v_t of some open star $V_t \neq \sigma(i)$. Since i is not assigned to V_t by σ, i cannot be in the same part of P as $\omega_t(1)$, as P is a refinement of $\pi(\sigma)$. Therefore, $d(x_i, v_t) \geq \epsilon_t$. However, $d(y_i, x_i) \geq d(x_i, v_t)$ since the shortest part between y_i and x_i passes through v_t. This contradicts $d(y_i, x_i) < \epsilon/3$. Therefore, y_i lies on $\sigma(i)$.

Assume y_i and y_j lie on the same edge of Γ and $\sigma(i) \neq \sigma(j)$. Then, $x_i \neq x_j$, so $d(x_i, x_j) \geq \epsilon$ and consequently $d(y_i, y_j) > \epsilon/3$. If x_i is not on this same edge, it must be on another edge of V_t, so $d(y_i, v_t) < \epsilon/3$. (If the two edges share both of their vertices and the shortest path between x_i and y_i were to pass through the other vertex, we would have $d(x_i, y_i) > \epsilon$, which is a contradiction.) Meanwhile, $d(y_j, v_t) > 2\epsilon/3$ because $d(x_j, v_t) \geq \epsilon_k \geq \epsilon$. Therefore, y_i is strictly closer to the central vertex of $\sigma(i)$ than is y_j. If x_j is not on the same edge as y_j and y_i, then by a similar argument, y_j must be closer to the other vertex of the edge than is y_i, so therefore, y_i is strictly closer to v_t than is y_j. The final case is that x_i and x_j are both on this edge, and since $\sigma(i) \neq \sigma(j)$ and x conforms to λ, it must be the case that $d(x_i, v_t) < d(x_j, v_t)$. Since they are on the same edge, $d(x_i, x_j) = d(x_j, v_t) - d(x_i, v_t) \geq \epsilon$. We have $d(y_i, v_t) < d(y_i, x_i) + d(x_i, v_t) < \epsilon/3 + d(x_i, v_t)$ and $d(x_j, v_t) < d(y_j, v_t) + \epsilon/3$, so $d(y_j, v_t) > d(x_j, v_t) - \epsilon/3 > d(x_i, v_t) + \epsilon/3 > d(y_i, v_t)$. Thus, y_i is strictly closer to the central vertex of $\sigma(i)$ than is y_j.

Finally, for an open star V_t, let $i < j$ with $\omega(i) \nsim_P \omega(j)$. Suppose $y_{\omega(i)}$ and $y_{\omega(j)}$ lie on the same edge. If $x_{\omega(i)}$ and $x_{\omega(j)}$ lie on this same edge, then by a similar argument as above, $d(x_{\omega(i)}, v_t) < d(x_{\omega(j)}, v_t)$. We know $\omega(j) \nsim_P \omega(1)$. Therefore, $x_{\omega(j)}$ and $y_{\omega(j)}$ must lie on the same edge as the distance from $x_{\omega(j)}$ to any vertex of degree ≥ 2 is at least ϵ, so $d(y_{\omega(j)}, v_t) > 2\epsilon/3$. If $x_{\omega(i)}$ is on a different edge from $y_{\omega(i)}$, then $d(y_{\omega(i)}, v_t) < \epsilon/3$, which means $d(y_{\omega(i)}, v_t) < d(y_{\omega(j)}, v_t)$.

Therefore, y conforms to λ, so U_λ is open. □

For a set Λ of conditions, let $U_\Lambda = \bigcap_{\lambda \in \Lambda} U_\lambda$. Given a nonempty Λ, let Q be the common refinement of the partitions P^λ, $\lambda \in \Lambda$: $i \sim_Q j$ if and only if $i \sim_{P^\lambda} j$ for all

$\lambda \in \Lambda$. For points in U_Λ, $x_i = x_j$ implies $i \sim_Q j$. Also, note that if $x_i = v_t$, where v_t is a vertex with degree ≥ 2, $\sigma^\lambda(i) = V_t$ for all $\lambda \in \Lambda$, and if $j \nsim_Q i$, then there exists $\lambda \in \Lambda$, where either $\sigma^\lambda(j) \neq V_t$ or j comes after i in the total order induced by ω^λ. Therefore, in U_Λ, the only points allowed to move onto a central vertex v_t are those x_i where $i \sim_Q \omega_t^\lambda(1)$ for all $\lambda \in \Lambda$. We call these points which are allowed to move onto a central vertex Type 1 under Λ.

If $\sigma^\lambda(i) = V_t$ for all $\lambda \in \Lambda$ but $i \nsim_Q \omega_t^\lambda(1)$ for some $\lambda \in \Lambda$, then x_i is called Type 2 under Λ. Within U_Λ, such a point is not allowed to move onto the central vertex and cannot leave the open star V_t so is restricted to a single edge within any connected component of U_Λ.

Lastly, if $\sigma^\lambda(i) \neq \sigma^\mu(i)$ for some $\lambda, \mu \in \Lambda$, then x_i is called Type 3 under Λ. Within U_Λ, such a point must remain in the interior of an edge connecting v_t and v_u, the central vertices of $\sigma^\lambda(i)$ and $\sigma^\mu(i)$, respectively, and it must remain strictly between the Type-2 points assigned to V_t and V_u.

THEOREM 2.2. *For nonempty Λ, each connected component of U_Λ is contractible.*

PROOF. Let A be a connected component of U_Λ. Within A, the only points that are able to move between edges are Type 1. We construct a homotopy in two stages. First, we move all Type-1 points onto the central vertex of their designated open star via straight-line homotopy. This is possible because $i \sim_Q \omega^\lambda(1)$ for all $\lambda \in \Lambda$, so x_i is not obstructed by any x_j with $i \nsim_Q j$. Denote by B the resulting subspace of A. Within B, each Type-1 point must remain on its vertex, and no point can enter or leave any edge.

Consider an edge e of an open star V_t and assume that at least one point of the configurations in B lies on e. There are two cases to consider.

In the first case, the other vertex of e has degree 1. Then, the only points on e are Type 2. Let m be the number of equivalence classes of Q that are represented by the Type-2 points lying on e. These equivalence classes are ordered by the distance of their closest representative to v_t, since ω^λ groups members of equivalence classes of P^λ together in its total order. Let y_e be the configuration of points on e where the points of the kth equivalence class of Q lie together at a distance of k/m away from v_t.

In the second case, e connects v_t to another central vertex v_u. There may be points of Type 2 which are assigned to V_t, points of Type 2 which are assigned to V_u and points of Type 3 which lie in between. As before, the equivalence classes of Q represented by points on e can be ordered by the distance of their closest representative to v_t. (Here, we use the facts that $d(x_i, v_t) = 1 - d(x_i, v_u)$ and equivalence classes represented by Type-2 points assigned to V_u can be ordered by their distances from v_u.) Let m_1, m_2 and m_3 denote the numbers of equivalence classes of Q represented by the points of Type 2 assigned to V_t, Type 3 and Type 2 assigned to V_u, respectively, and $m = m_1 + m_2 + m_3$. Now define the configuration y_e of points on e where the points of the kth equivalence class of Q lie together at a distance of $k/(m+1)$ away from v_t. (Note that this will give the same configuration y_e when defined for the open star V_u, since the kth equivalence class from v_t is the $(m-k+1)$th equivalence class from v_u.)

Now let y be the configuration of points that restricts to y_e on each edge e of Γ (and all Type-1 points lie on their designated central vertex). Use a straight-line homotopy to move the configurations of B to y. If x_i and x_j lie on the same edge with $i \nsim_Q j$ with $d(x_i, v_t) < d(x_j, v_t)$, then $d(y_i, v_t) < d(y_j, v_t)$, so the straight-line homotopy will not create any collisions of points which are not in the same class of Q. □

3. The discrete model

Next, we define a finite cubical complex which will serve as our discrete model $\mathcal{D}_n(\Gamma, \mathcal{P})$ for the $\mathcal{P}$-configuration space.

First, subdivide each edge e of Γ between two vertices of degree ≥ 2 into $n+1$ segments of length $1/(n+1)$, and subdivide each edge having a vertex of degree 1 into n segments of length $1/n$. We will call this subdivided graph Γ'. To avoid confusion, we will refer to the cells τ_i of Γ' as 'nodes' and 'segments', and reserve the terms 'vertices' and 'edges' for the cells of Γ. So, 'central vertices' still refer only to original vertices of Γ and the 'open stars' refer to the original stars of Γ, which are now subdivided in Γ'. The subdivision gives a cubical structure to Γ^n with (closed) cells of the form $\tau = (\tau_1, \ldots, \tau_n)$ with $\tau_i \in V(\Gamma') \cup E(\Gamma')$. We now define the subcomplex

$$\mathcal{D}_n(\Gamma, \mathcal{P}) = \left\{ \tau \;\middle|\; \forall P \in \mathcal{P}, \exists j \in N, \bigcap_{i \in [j]_P} \tau_i = \emptyset \right\}.$$

Let $D_\lambda = \mathcal{D}_n(\Gamma, \mathcal{P}) \cap U_\lambda$ and $D_\Lambda = \bigcap_{\lambda \in \Lambda} D_\lambda$.

Let F be the face poset of $\mathcal{D}_n(\Gamma, \mathcal{P})$ and let F_Λ be the subposet of cells with nonempty intersection with D_Λ.

THEOREM 3.1. *The geometric realisation of the order complex ΔF_Λ is a deformation retract of D_Λ.*

PROOF. Since $\mathcal{D}_n(\Gamma, \mathcal{P})$ is a regular cell complex, it is homeomorphic to the geometric realisation of ΔF. The 0-simplices of ΔF correspond to the elements of F (which are cells of $\mathcal{D}_n(\Gamma, \mathcal{P})$). A k-simplex is a chain of $k+1$ cells ordered by inclusion, and the points of the simplex are convex combinations of its vertices. A homeomorphic inclusion of $|\Delta F|$ into $\mathcal{D}_n(\Gamma, \mathcal{P})$ is obtained by choosing a point in the interior of a cell τ for the 0-simplex corresponding to τ. This map is then defined linearly on the simplices.

For each cell $\tau = (\tau_1, \tau_2, \ldots, \tau_n) \in F_\Lambda$, we show that $D_\Lambda \cap \tau$ is convex. Suppose $x, y \in D_\Lambda \cap \tau$. Let $x_i, y_i \in \tau_i$, so $z_i^t = (1-t)x_i + ty_i \in \tau_i$ for any $t \in [0,1]$, so $z^t \in \tau$. If τ_i and τ_j lie on the same edge e with vertex v of Γ and $i \nsim_Q j$, then $d(x_i, v) < d(x_j, v)$ implies $d(y_i, v) < d(y_j, v)$ since both x and y conform to all $\lambda \in \Lambda$. Therefore, $d(z_i, v) = (1-t)d(x_i, v) + td(y_i, v) < (1-t)d(x_j, v) + td(y_j, v) = d(z_j, v)$. (If $i \nsim_Q$, and τ_i and τ_j lie on two edges e_i and e_j that meet at v with τ_i containing v, then if either x_i or y_i on v, then i is Type 1 and assigned to v under Λ, and neither x_j nor y_j is allowed to be on v since $j \nsim_Q i$.) Therefore, for all $t \in [0,1]$, z_j^t must lie in the interior of e_j and z_i^t is either on e_i or at v. Therefore, z^t conforms to all $\lambda \in \Lambda$.

Since $D_\Lambda \cap \tau$ is convex, it must lie in a connected component A of U_Λ, so every Type-2 and Type-3 point is restricted to one edge of Γ. Type-1 points might be able to move between edges within an open star. For each edge e of Γ, let m_e be the number of equivalence classes of Q whose points could lie on e in the connected component A of U_Λ. These equivalence classes are ordered along e. (The Type-1 class assigned to the initial vertex of e first, then Type-2 classes assigned to the initial vertex, then, if the terminal vertex is also a central vertex of Γ, Type-3 classes, and Type-2 and Type-1 classes assigned to the terminal vertex.)

For each τ, we will choose a point in the interior of τ that lies in D_Λ. Each τ_i is either a node or a segment of Γ'. If τ_i is an original vertex from Γ of degree ≥ 2, then i must be of Type 1 in Λ. For points x_i of Types 2 and 3 under Λ, τ_i must be a segment or an interior node on an edge e (possibly an original vertex of degree 1). We will now define a configuration x^τ in the interior of τ in D_Λ.

- If τ_i is a node, put x_i^τ at that node.
- If τ_i is a segment on edge e and i belongs to the kth of the m_e classes, put x_i^τ at $k/(m_e + 1)$ along the length of τ_i from its initial node (the direction consistent with the direction chosen for e).

As defined, x_i^τ is contained in the interior of τ_i, so $x^\tau \in \tau$. Also, if $x_i^\tau = x_j^\tau$, then $i \sim_Q j$, and the order of the equivalence classes along each edge of Γ have been preserved, so x^τ conforms to all $\lambda \in \Lambda$. Therefore, $x^\tau \in D_\Lambda$.

We will first define a continuous map $f : D_\Lambda \to |\Delta F_\Lambda|$ and then create a straight-line homotopy from the identity map to f. For $x \in D_\Lambda$, let τ be the minimal cell of F_Λ containing x. For each x_i, we will define a point $y_i \in \tau_i$ and then define $f(x) = y = (y_1, \ldots, y_n) \in |\Delta F_\Lambda|$.

We note that, if $\tau \in F_\Lambda$ and τ_i is a segment from node a to node b, then the face of τ with τ_i replaced with a is only in F_Λ if i is in the first equivalence class on τ_i (or Type 1 and assigned to a if a is a degree ≥ 2 vertex of Γ). Likewise with b and the last equivalence class on τ_i. Furthermore, if i is in the first class on τ_i and j is in the last class on the preceding segment τ_j, then the face obtained by replacing both τ_i and τ_j with a is only in F_Λ if $i \sim_Q j$. Therefore, a configuration $y \in \tau$ is only in $|\Delta F_\Lambda|$ if each of its points y_i either lies at x_i^τ or strictly between x_i^τ and a node of τ_i that i is free to occupy (in which case, i has 'claimed' that node, preventing any y_j with $j \nsim_Q i$ from also claiming it).

If τ_i is a node, then $x_i = x_i^\tau$ and we define $y_i = x_i$.

Suppose τ_i is a segment with initial node a and terminal node b. If i is a member of an equivalence class of Q which is neither the first nor last with representatives in τ_i, then define $y_i = x_i^\tau$.

We now suppose i is a member of the first of multiple equivalence classes of Q with representatives in τ_i.

In the case that a is an original vertex of Γ with degree ≥ 2, if i is of Type 1 and assigned to a, then define y_i to be the point closer to a out of x_i and x_i^τ, and if i is not of that type, define $y_i = x_i^\tau$.

If a is not an original vertex of Γ, then the segment τ_i has a preceding segment on its edge. If other members of $[i]_Q$ are assigned to that preceding segment or if it contains no points at all in its interior, then define y_i to be the point closer to a out of x_i and x_i^τ.

Finally, suppose the last equivalence class with representatives assigned to that preceding segment is $[j]_Q$, where $j \nsim_Q i$.

Let $d_i = \min\{d(x_i, a)/d(x_i^\tau, a), 1\}$ and let $d_{[i]} = \min\{d_k \mid k \in [i]_Q, \tau_k = \tau_i\}$. Similarly, define d_j and $d_{[j]}$. Set $M = \max\{d_{[i]}, d_{[j]}\}$. If $d_i \geq d_{[j]}$, set $y_i = x_i^\tau$. If $d_i < 1 = d_{[j]}$, set $y_i = x_i$. Otherwise, we have $d_i < d_{[j]} = M < 1$, and we set $y_i = (d_i/M)x_i^\tau + (1 - d_i/M)a$.

If i is a member of the last of multiple equivalence classes of Q in τ_i, y_i can be defined symmetrically. (If τ_i is a segment whose terminal node b is a degree-1 vertex of Γ, then set y_i to be the point closer to b out of x_i and x_i^τ.)

The last possibility is if all points in τ_i belong to the same equivalence class (so it is both the first and the last class on τ_i). Then, if x_i lies before x_i^τ, define y_i according to the rules for the first equivalence class on a segment, and if it lies after, use the rules for the last equivalence class. If $x_i = x_i^\tau$, both rules agree that $y_i = x_i^\tau$.

We set $f(x) = y = (y_1, y_2, \ldots, y_n)$. Then, f is continuous on each intersection $D_\Lambda \cap \tau$, so f is continuous. Furthermore, f fixes each point of $|\Delta F_\Lambda|$. The homotopy $H : D_\Lambda \times [0, 1] \to D_\Lambda$ given by $H(x, t) = (1 - t)x + tf(x)$ is therefore a deformation retraction of D_Λ onto $|\Delta F_\Lambda|$. □

We will employ the notion of Morse matchings from [7] applied to simplicial complexes. For a simplicial complex K, let M be a collection of pairs (ρ, ψ), where $\rho < \psi \in K$ with ρ a codimension-1 face of ψ, such that no simplex of K appears in more than one pair. A path is a sequence $\rho_1 < \psi_1 > \rho_2 < \psi_2 > \cdots > \rho_{t-1} < \psi_{t-1} > \rho_t$, with $t \geq 2$, where for each i, $(\rho_i, \psi_i) \in M$ and ρ_i a codimension-1 face of ψ_{i-1} not equal to ρ_{i-1}. If for all such paths, $\rho_t \neq \rho_1$, then M is acyclic and is called a *Morse matching* on K.

We will prove that each connected component of D_Λ is contractible by showing that each connected component of the simplicial complex ΔF_Λ collapses to a point by defining Morse matchings that collapse it onto a sequence of subcomplexes, using the following lemma (a more general version for cell complexes can be found in [9, Theorem 11.13]).

LEMMA 3.2. *Let K be a finite simplicial complex with a Morse matching whose critical simplices form a subcomplex L. Then, K collapses simplicially to L.*

In preparation for the Morse matchings, we will define some important terminology. First, we choose a direction on every edge of Γ. Assume any degree-1 vertex is the terminal vertex of its edge.

Let K be a connected component of ΔF_Λ. Since (the realisation of) K lies within a connected component of U_Λ, each member of a Type-2 or Type-3 equivalence class of Q is restricted to a single edge, while a member of a Type-1 class may lie on any edge in its open star. As before, the equivalence classes that may intersect e are ordered by their distance from the initial vertex v. We write $C \prec C'$ if C and C' are equivalence classes on e with C closer to the initial vertex.

Let e be a fixed edge with $r > 0$ equivalence classes having members that may lie on e in this connected component of U_Λ. Suppose C is such an equivalence class with q_C members that may lie on e. We will define a *zone* Z_C on e by specifying a collection of consecutive nodes in e and the zone consists of those nodes and all segments in e that are incident to them.

If C is the only class on e and Type 1, define Z_C by the $q_C + 1$ consecutive nodes (and their incident segments) starting at the vertex to which the class has been assigned. Otherwise, if C is the first class on e and Type 1, define Z_C by the first $q_C + 1$ nodes of e starting at the initial vertex. If there is no Type-1 class assigned to the initial vertex, designate the initial vertex and its incident segments as a zone Z for an 'empty' first class. For $k < r$, if C is the kth class on e, define Z_C to be the q_C nodes immediately following the last node of the $(k-1)$st zone. For the rth (last) class C on e, if C is not Type 1 and assigned to the terminal vertex of e, define Z_C the same way as the previous zones; if C is Type 1 and assigned to the terminal vertex of e, define Z_C to be the last $q_C + 1$ nodes of e (including the terminal vertex). There are enough nodes on e to define these zones, but there may be some nodes which are not assigned to a zone (either after the rth class's zone or between the zones for the $(r-1)$st and rth classes). We will consider this set of consecutive nodes (and their incident segments) to be another zone for another 'empty' class. The zones are ordered in the same way as the classes on e (with 'empty classes' inserted in the appropriate place). We write $Z_C \prec Z_{C'}$ if $C \prec C'$. Note that each pair of consecutive zones share only a single segment.

For a segment α on e, denote its initial and terminal nodes by $\alpha(0)$ and $\alpha(1)$, respectively. The cells (nodes and segments) of e are totally ordered by their distance from the initial vertex of e, with the understanding that $\alpha(0) \prec \alpha \prec \alpha(1)$ for a segment α.

Let $\tau \in F_\Lambda$. Suppose τ_i lies on e. We say that an interior node a on e with $a \succ \tau_i$ is *unobstructed for* i if $a \cap \tau_j = \emptyset$ for any $\tau_j \succ \tau_i$. Likewise, an interior node a on e with $a \prec \tau_i$ is unobstructed for i if $a \cap \tau_j = \emptyset$ for any $\tau_j \prec \tau_i$. If a and $\tau_i = b$ are adjacent nodes joined by the segment α and a is unobstructed for i, then we can define a new cell $\tau' \in F_\Lambda$ by setting $\tau'_j = \tau_j$ for all j with $\tau_j \neq b$ and choosing $\tau'_j \in \{b, \alpha\}$ for each j with $\tau_j = b$.

THEOREM 3.3. *For nonempty* Λ, *each connected component of* D_Λ *is contractible.*

PROOF. By the previous theorem, each connected component of D_Λ has a connected component of $|\Delta F_\Lambda|$ as a deformation retract. We will show each connected component of ΔF_Λ is contractible.

Let J be a connected component of ΔF_Λ. For an edge e of Γ, we will now define a Morse matching M^1_e on the simplices of J. A k-simplex ρ in J is a chain $\tau^0 \subset \tau^1 \subset \cdots \subset \tau^k$, where each $\tau^j = (\tau^j_1, \ldots \tau^j_n) \in F_\Lambda$. We will create pairs $\rho \subset \rho'$, where ρ' is a $(k+1)$-simplex.

For $\rho = \tau^0 \subset \cdots \subset \tau^k$, consider the maximal cell τ^k. Let $C = C(\rho)$ be the last equivalence class with $\tau^k_i \in Z_B$ for some $i \in C$ and $B \prec C$. Let $\beta = \beta(\rho)$ be the last node or segment on e with $i \in C$ and $\tau^k_i = \beta$ such that there exists an interior node $a \succ \tau^k_i$ in Z_C, which is unobstructed for i. Let $I = I(\rho)$ be the set of all $i \in C$ with $\tau^k_i = \beta$.

Note that, since no class after C has members in any zone before its own, then if Z_B is the first zone occupied by members of C, no zone after Z_B up to and including Z_C can contain any members of classes other than C. Since at least one member of C is in Z_B, there is at least one node in Z_C that is unobstructed for at least one $i \in C$ with τ_i^k lying before it.

If β is a node $\alpha(0)$, we define ρ' by adding τ^{k+1} to the end of the chain, where $\tau_i^{k+1} = \alpha$ for all $i \in I$ and $\tau_i^{k+1} = \tau_i^k$ for all $i \notin I$. This new cell τ^{k+1} is still in F_Λ because β is the last cell with an unobstructed node in Z_C after it, implying that $\alpha(1)$ must be unobstructed for $i \in I$. Note that $C(\rho') = C(\rho)$, replacing $\alpha(0)$ with α cannot cause the class to leave a zone, and $\beta(\rho') = \alpha$. Similarly, if $\beta = \alpha$ is a segment and there exist $i \in C$ with $\tau_i^k = \alpha(0)$, we may define $\tau_i^{k+1} = \alpha$ for such i, and $\tau_i^{k+1} = \tau_i^k$ for all other i, and add this τ^{k+1} to the end of ρ to produce ρ'.

If $\beta = \alpha$ is a segment and there are no $i \in C$ with $\tau_i^k = \alpha(0)$, find the last τ^j in ρ such that $\tau_i^j = \alpha(0)$ for some $i \in I$. Define τ' by setting $\tau'_i = \alpha$ for all $i \in I$ with $\tau_i^j = \alpha(0)$ and $\tau'_i = \tau_i^j$ for all other i. We obtain ρ' by inserting τ' in the chain immediately after τ^j. (If $\tau^{j+1} = \tau'$ already, then ρ is paired by this rule with the $(k-1)$-simplex obtained by removing τ^{j+1}.) Note that τ' is in F_Λ: $\tau' \subseteq \tau^{j+1} \in F$, so $\tau' \in F$, and τ' has all the same cells as τ^j except for a subset of one class C which changed from a node to an incident segment, so τ' still has nonempty intersection with D_Λ.

If $\beta = \alpha$ is a segment and no $\tau_i^j = \alpha(0)$ for any $i \in I$, instead find the first τ^j in ρ with $\tau_i^j = \alpha$ for some $i \in I$. Define τ' by setting $\tau'_i = \alpha(1)$ for all $i \in I$ and $\tau'_i = \tau_i^j$ for all other i. Then, ρ' is obtained by inserting τ' immediately before τ^j (or removing it if $\tau^{j+1} = \tau'$). Note that $\tau' \in F_\Lambda$: $\alpha(1) \subset \alpha$, so $\tau' \in F$, and since only cells for members of C are being changed from α to $\alpha(1)$ and C must be the last equivalence class on α, τ' has nonempty intersection with D_Λ.

Note that if $\rho \subset \rho'$ is a pair, $C(\rho) = C(\rho')$ and $I(\rho) \subseteq I(\rho')$. If $\beta(\rho) = \alpha(0)$, then $\beta(\rho') = \alpha$. Otherwise, if $\beta(\rho) = \alpha$, then $\beta(\rho') = \alpha$. In every case, $\beta(\rho) \preceq \beta(\rho')$.

Note also, the cell τ' advances some members of $C(\rho)$ one step along e from their position in τ^j (replacing a node $\alpha(0)$ with its succeeding segment α, or replacing a segment α with its terminal node $\alpha(0)$), and all other coordinates remain in the same position.

Also, if $\tau^1 \subset \tau^2$, then any $\tau_i^1 \subseteq \tau_i^2$.

Next, we consider how C may change for ρ obtained by deleting the maximal cell $\tau^{\max}(\psi)$ from a chain ψ. Note that $\tau^{\max}(\rho)$ is a face of $\tau^{\max}(\psi)$ obtained by replacing some segments with nodes. The zone in which a given coordinate i is located can only change if $\tau^{\max}(\psi)_i$ is a segment α, where two zones overlap and $\tau^{\max}(\rho)_i$ is one of its nodes. That coordinate switches from lying in both zones to lying in only one. Therefore, if $\tau^{\max}(\psi)_i$ does not lie in any zone prior to $Z_{[i]}$, then neither does $\tau^{\max}(\rho)_i$. If $\tau^{\max}(\psi)_i$ is the first segment of $Z_{[i]}$ and $\tau^{\max}(\rho)_i$ is its terminal node, then i has exited the lower zone. Thus, $C(\rho) \preceq C(\psi)$.

Now assume $C(\rho) = C(\psi) = C$. There are a few ways that β and I could change. A segment in Z_C appearing in $\tau^{\max}(\psi)$ could be replaced in $\tau^{\max}(\rho)$ by its terminal node, causing its initial node to become unobstructed for its preceding segment and node (appearing in $\tau^{\max}(\rho)$), in which case $\beta(\rho)$ would be that preceding segment or node, $\beta(\rho) \succ \beta(\psi)$ and $I(\rho) \cap I(\psi) = \emptyset$. If $\beta(\psi) = \alpha$ is a segment, all of its appearances in $\tau^{\max}(\psi)$ could be replaced with $\alpha(0)$ in $\tau^{\max}(\rho)$, in which case, $\beta(\rho) = \alpha(0)$ and $I(\psi) \subseteq I(\rho)$. If some of its appearances are replaced with $\alpha(1)$ in $\tau^{\max}(\rho)$, then $\beta(\rho) \preceq \alpha(1)$. Now, $\alpha(1)$ may or may not have unobstructed nodes of Z_C ahead of it. If it does, $\beta(\rho) = \alpha(1)$ and $I(\rho) = I(\psi)$. If it does not, $\beta(\rho) \preceq \alpha(0)$ and there exists at least one coordinate $i \in I(\psi)$ such that $i \notin I(\rho)$.

Note that if ρ is obtained from ψ by deleting its maximal cell, $I(\psi) \nsubseteq I(\rho)$ implies $\tau^{\max}(\rho)$ assigns a node a to a coordinate i for which $\tau^{\max}(\psi)_i = \alpha$ with $\alpha(1) = a$. If ρ and ψ have the same maximal cell, then $I(\psi) = I(\rho)$.

We now show M_e^1 is acyclic. Suppose we have a path

$$\rho_1 < \psi_1 > \rho_2 < \psi_2 > \cdots > \rho_{t-1} < \psi_{t-1} > \rho_t,$$

where $t \geq 3$, and for $1 \leq r \leq t-1$, ρ_r is paired with ψ_r in the matching, and ρ_{r+1} is a codimension-1 face of ψ_r not equal to ρ_r.

The chain ψ_r is obtained from ρ_r by inserting some $\tau'(r)$ based on the maximal cell $\tau^{\max}(\rho_r)$ which determines the class $C(\rho_r)$, cell $\beta(\rho_r)$, and coordinate set $I(\rho_r)$ for the pairing. Here, ρ_{r+1} is obtained by removing some cell other than $\tau'(r)$ from ψ_r. Therefore, $\tau'(r+1) \neq \tau'(r)$.

Suppose the path is a cycle, meaning $\rho_1 = \rho_t$. Recall that $C(\rho_r) = C(\psi_r)$ and $C(\psi_r) \geq C(\rho_{r+1})$ for all $1 \leq r \leq t-1$. Therefore, since the path is a cycle, we must have $C(\psi_r) = C(\rho_{r+1})$ for all $1 \leq r \leq t-1$.

Suppose $I(\psi_r) \nsubseteq I(\rho_{r+1})$ for some $1 \leq r \leq t-1$. Then, ρ_{r+1} must be obtained by removing $\tau^{\max}(\psi_r)$ from the chain, and there is a coordinate j such that $\tau^{\max}(\rho_{r+1})_j$ is the terminal node $\alpha(1)$ of the segment $\alpha = \tau^{\max}(\psi_r)_j$. Therefore, $\tau_j = \alpha(1)$ for all cells τ in ρ_{r+1}. The matching cannot introduce a preceding segment in a coordinate, so there is no way to regain $\tau^{\max}(\psi_r)$, which contradicts the path being a cycle. Therefore, we may assume $I(\psi_r) \subseteq I(\rho_{r+1})$. We also know $I(\rho_r) \subseteq I(\psi_r)$ for all $1 \leq r \leq t-1$. Since we have a cycle, we can conclude that $I(\rho_r) = I(\psi_r) = I$ for all $1 \leq r \leq t-1$.

Now, for all $i \in I$ and $1 \leq r \leq t-1$, we must have $\beta(\rho_r) = \tau^{\max}(\rho_r)_i$ and $\beta(\psi_r) = \tau^{\max}(\psi_r)_i$. We have $\beta(\rho_r) \preceq \beta(\psi_r)$, and $\beta(\psi_r)$ must be a segment α_r. If $\beta(\psi_r) \succ \beta(\rho_{r+1})$, then ρ_{r+1} is obtained by removing $\tau^{\max}(\psi_r)$ and $\beta(\rho_{r+1}) = \tau^{\max}(\rho_{r+1})_i = \alpha_r(0)$. This forces the conclusion that $\tau^{\max}(\psi_r) = \tau'(r)$, so it could not have been removed from ψ_r in a valid path. Therefore, $\beta(\psi_r) \preceq \beta(\rho_{r+1})$. Since we have a cycle, this means β must be the same segment α for all chains in the path.

We know $\tau'(1)$ is obtained by replacing $\alpha(0)$ with α if any τ^j in ρ_1 includes $\alpha(0)$ in any of its coordinates, or by replacing α with $\alpha(1)$ otherwise. In the first case, ρ_2 must be obtained from ψ_1 by removing the predecessor to $\tau'(1)$, a cell with $\alpha(0)$ in at least one coordinate i, where $\tau_i^{\max} = \alpha$, because any other removal would result in ρ_2 being paired with the lower dimensional chain obtained by removing $\tau'(1)$. Such a cell cannot be regained later in the path because every cell in the chain has $\beta = \alpha$, so any new cell will replace all $\alpha(0)$ of an existing cell with α or α with $\alpha(1)$. Likewise, in the second case, ρ_2 must be obtained from ψ_1 by removing the successor to $\tau'(1)$, a cell with α in at least one coordinate i, where $\tau'_i = \alpha$. Such a cell cannot be regained later in the path because $\alpha(0)$ does not appear in any coordinate, so any new cell introduced by the matching will be replacing all α in an existing cell with $\alpha(1)$. Therefore, the path cannot be a cycle.

Thus, M_e^1 is a Morse matching. The only simplices in J which are unpaired in M_e^1 are those chains $\rho = \tau^1 \subset \cdots \subset \tau^k$, where no zone Z_B on e contains any τ_i^k with $[i] \succ B$. These simplices form a subcomplex K: the order complex of cells $\tau \in J^0$ such that no τ_i in e lies in a zone Z_B of e with $[i] \succ B$.

We can define a similar Morse matching M_e^2 on K, where we consider the maximal cell τ^k of a simplex ρ. Let $C = C(\rho)$ be the first equivalence class such that $\tau_i^k \in Z_B$ for some $i \in C$ and $C \prec B$, and $\beta = \beta(\rho)$ be the first node or segment on e with $i \in C$ and $\tau_i^k = \beta$ such that there exists an interior node $a \prec \tau_i^k$ in Z_C which is unobstructed for i, and let $I = I(\rho)$ be the set of all $i \in C$ with $\tau_i^k = \beta$. As before, no zone before Z_B down to and including Z_C can contain any member of a class other than C. From here, the matching M_e^2 may be defined similarly to M_e^1, and it is similarly acyclic. Its critical simplices are those chains in K of cells τ for which no τ_i in e lies in a zone Z_B with $[i] \prec B$. However, since these chains are already in K, that means each τ_i on e lies in $Z_{[i]}$. This forms a subcomplex L to which K now simplicially collapses.

The pairings M_e^1 and M_e^2 are determined by and only affect the τ_i in e. Therefore, if f is another edge of Γ that contains points in U_Λ, the matchings M_f^1 and M_f^2 induce matchings on the resulting subcomplex L to which J collapses. Collapsing via the two matchings for one edge of Γ one at a time, we arrive at a subcomplex H consisting of chains of cells τ for which, if τ_i lies on an edge e, then τ_i lies in zone Z_B if and only if $i \in B$.

We will use $Z(i)$ to refer to the (closed) zone associated to i. If i belongs to a class B which is not Type 1, then $Z(i)$ is the interval from the first node to the last node of Z_B on the edge containing x_i in U_Λ. If i belongs to a Type-1 class B, then $Z(i)$ is the union of the intervals from the central vertex to the last node of Z_B on each edge of the open star.

Now, if $\tau \in H^0$, each τ_i belongs to $Z(i)$. Conversely, we will show that any cell τ such that each τ_i belongs to $Z(i)$ is in H^0.

Since $Z(i) \cap Z(j) = \emptyset$ if $i \nsim_Q j$, if $\bigcap_{i\in C} \tau_i \neq \emptyset$, then C must be a subset of a class of Q. Suppose there exists a partition P such that $\bigcap_{i\in C} \tau_i \neq \emptyset$ for all classes C of P. Then, P is a refinement of Q. Since $Q \notin \mathcal{P}$, we have $P \notin \mathcal{P}$, so τ is a cell of $\mathcal{D}_n(\Gamma, \mathcal{P})$. Since the zones are ordered in accordance with the orders of classes on each edge dictated by Λ, τ has nonempty intersection with U_Λ. Therefore, $\tau \in F_\Lambda$, and thus it is in H^0 since it is unpaired in all of the Morse matchings.

It follows that H is the order complex of the face poset of the cubical complex $\{(\tau_1, \ldots, \tau_n) \mid \tau_i \in Z(i)\}$, which is contractible as the product of contractible zones. □

4. The equivalence

Finally, we will show that $\mathrm{Conf}_n(\Gamma, \mathcal{P})$ and $\mathcal{D}_n(\Gamma, \mathcal{P})$ have isomorphic homology by using Mayer–Vietoris spectral sequences associated to an open cover of a space.

Let X be a topological space with a countable open cover $\mathcal{W}$. Let $\Lambda = \{\lambda_1, \ldots, \lambda_p\}$ be a collection of indices such that $W_\Lambda = \bigcap_{i=1}^{p} W_{\lambda_i} \neq \emptyset$. For such Λ, we denote by $C_q(W_\Lambda)$ the singular q-chains in W_Λ. We can then define the double complex $E^0_{p,q} = \bigoplus_{|\Lambda|=p} C_q(W_\Lambda)$. The vertical differentials are given by the usual boundary operator ∂. The horizontal differential $\delta : E^0_{p,q} \to E^0_{p-1,q}$ is induced by inclusion. For $\Lambda' \subset \Lambda$, if c_Λ is a q-chain in W_Λ, denote by $c_{\Lambda'}$ its inclusion in $C_{\Lambda'}$. Then, for $\Lambda = \{\lambda_1, \ldots, \lambda_p\}$, define

$$\delta(c_\Lambda) = \sum_{i=1}^{p} (-1)^{i-1} c_{\Lambda - \lambda_i}.$$

We now restate without proof the relevant result [5, Proposition 2.1.9] for the Mayer–Vietoris spectral sequence associated to the open cover $\mathcal{W}$ of X.

PROPOSITION 4.1. *The double complex* $E^1_{p,q} = \bigoplus_{|\Lambda|=p} H_q(W_\Lambda)$ *for* $p, q \geq 0$ *with*

$$d_1 = \delta_* : E^1_{p,q} \to E^1_{p-1,q}$$

forms a spectral sequence converging to $H_*(X)$.

We have open covers $\{D_\lambda\}$ and $\mathcal{U} = \{U_\lambda\}$ of $\mathcal{D}_n(\Gamma, \mathcal{P})$ and $\mathrm{Conf}_n(\Gamma, \mathcal{P})$, respectively, indexed by the same conditions λ.

THEOREM 4.2. *The inclusion of* D_Λ *into* U_Λ *induces a homotopy equivalence for all* Λ.

PROOF. We have seen that both U_Λ and D_Λ are disjoint unions of their connected components which are all contractible via deformation retraction. Each component of D_Λ is contained in one component of U_Λ. Within such a component of U_Λ, there is a path between the two retract points. We therefore need to show a one-to-one correspondence between the connected components of U_Λ and D_Λ.

As described previously, the choice of Λ determines which points are designated as Type 1, 2 and 3. Type-1 points are permitted to lie on their designated vertex, while Types 2 and 3 cannot move between edges. Furthermore, within each edge, points must remain ordered by equivalence class. We will therefore argue that the connected components of both spaces are in one-to-one correspondence with the possible assignments of Type-2 and Type-3 points to edges permissible by Λ.

In the proof of Theorem 2.2, we showed that each connected component deformation retracts to a configuration where each point of Type 1 lies on its assigned vertex and points of Type 2 and Type 3 coincide with their equivalence class members on each edge, with the classes ordered in accordance with Λ. Within U_Λ, any two such configurations can be joined by a path, so there is only one connected component for each assignment.

In the proof of Theorem 3.3, we showed that each component deformation retracts to a cubical complex of 'zones'. The zones are determined by the edge assignments, so there is only one connected component for each permissible assignment. □

THEOREM 4.3. *$H_*(\mathrm{Conf}_n(\Gamma, \mathcal{P}))$ is isomorphic to $H_*(\mathcal{D}_n(\Gamma, \mathcal{P}))$.*

PROOF. We have the spectral sequence associated to the open cover $\mathcal{U}$ of $\mathrm{Conf}_n(\Gamma, \mathcal{P})$ which converges to $H_*(\mathrm{Conf}_n(\Gamma, \mathcal{P}))$, and we have the spectral sequence associated to the open cover $\{D_\lambda\}$ of $\mathcal{D}_n(\Gamma, \mathcal{P})$ which converges to $H_*(\mathcal{D}_n(\Gamma, \mathcal{P}))$. Since D_Λ is homotopically equivalent to U_Λ, for every set Λ, by the inclusion map and the d_1 differentials are induced by inclusion, we have an isomorphism for each p, q between the $E^1_{p,q}$ entries in the two spectral sequences commuting with d_1. The two sequences are therefore equivalent, so $H_*(\mathrm{Conf}_n(\Gamma, \mathcal{P}))$ is isomorphic to $H_*(\mathcal{D}_n(\Gamma, \mathcal{P}))$. □

References

[1] A. Abrams, *Configuration Spaces and Braid Groups of Graphs*, PhD Thesis (UC Berkeley, 2000).

[2] Y. Baryshnikov, 'Euler characteristics of exotic configuration spaces', *Sém. Lothar. Combin.* **84B** (2020), Article no. 20, 12 pages.

[3] Y. Baryshnikov, C. Klivans and N. Kosar, 'On the topology of no *k*-equal spaces', *Adv. Appl. Math.* **149** (2023), Article no. 102538.

[4] L. Caputi, D. Celoria and C. Collari, 'From the Mayer–Vietoris spectral sequence to überhomology', *Proc. Roy. Soc. Edinburgh Sect. A*, to appear. Published online (2 October 2023).

[5] S. Chettih, *Dancing in the Stars: Topology of Non-k-Equal Configuration Spaces of Graphs*, PhD Thesis (University of Oregon, 2016).

[6] E. Fadell and L. Neuwirth, 'Configuration spaces', *Math. Scand.* **10** (1962), 111–118.

[7] R. Forman, 'Morse theory for cell complexes', *Adv. Math.* **134** (1998), 90–145.

[8] S. M. Gusein-Zade, 'Generating series of the classes of exotic unordered configuration spaces', *Sib. Math. J.* **64** (2023), 62–66.

[9] D. Kozlov, *Combinatorial Algebraic Topology*, Algorithms and Computation in Mathematics, 21 (Springer, Berlin–Heidelberg, 2008).

JAMES DOVER, Department of Mathematics,
University of North Georgia, 82 College Circle, Dahlonega, GA 30597, USA
e-mail: james.dover@ung.edu