Chapter 2

Spaces: Complexes

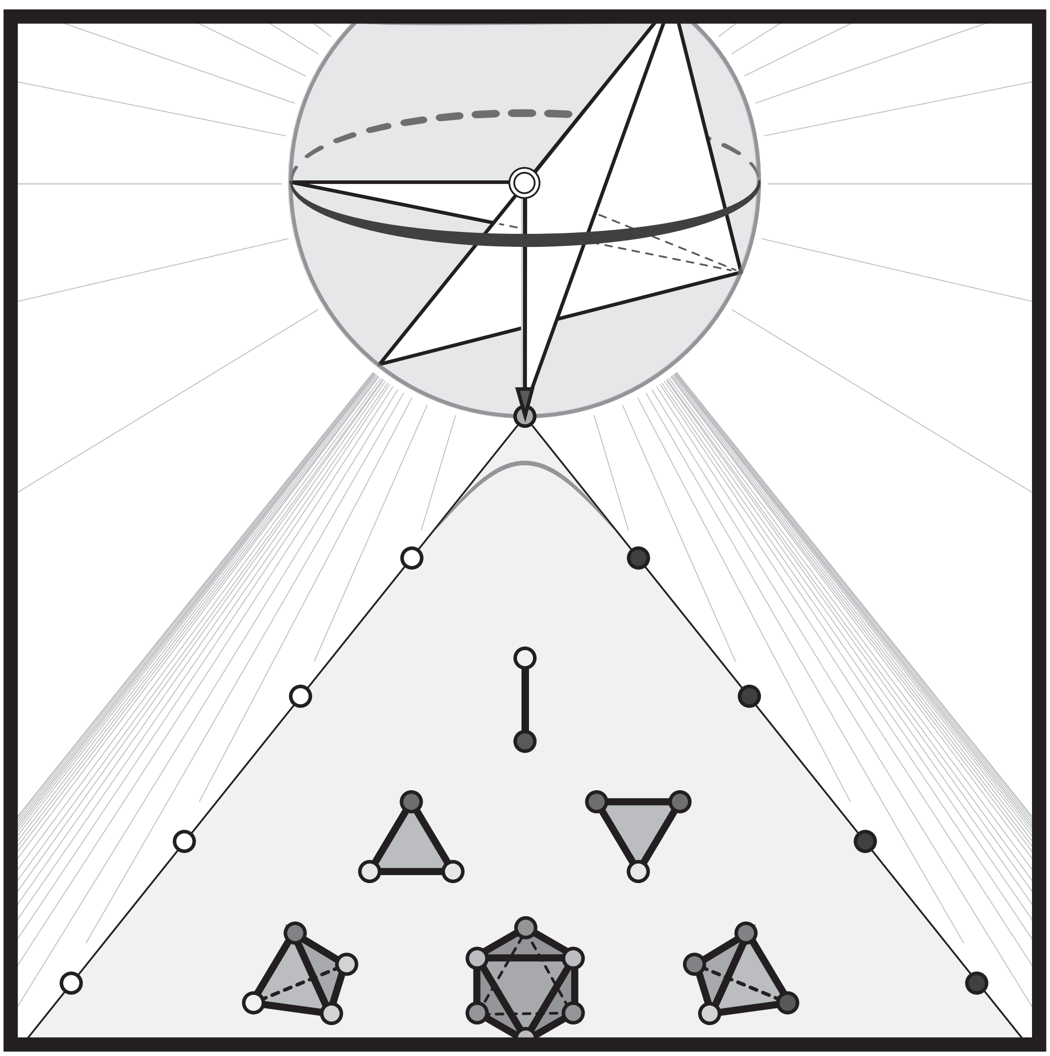

Parts assemble into wholes. A change in focus from smooth spaces patched together with diffeomorphisms to complexes assembled from a bin of simple pieces motivates the eventual transition to algebraic assemblages of spaces and invariants.

2.1 Simplicial and cell complexes

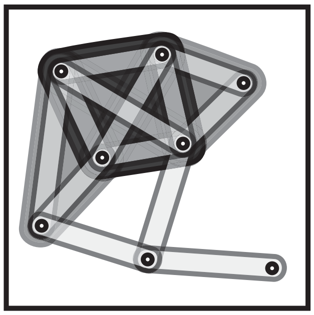

Let S be a discrete set. An **abstract simplicial complex** is a collection X of finite subsets of S, closed under restriction: for each $\sigma \in X$, all subsets of σ are also in X. Each element $\sigma \in X$ is called a **simplex**, or rather a k-simplex, where $|\sigma| = k+1$. Given a k-simplex σ, its **faces** are the simplices corresponding to all subsets of $\sigma \subset S$. For example, if $S = \{s_i\}_i$, the subcollection $\{s_2, s_5, s_7\}$ is a 2-simplex with three 1-simplex faces, three 0-simplex faces, and, of course, the empty-set face.

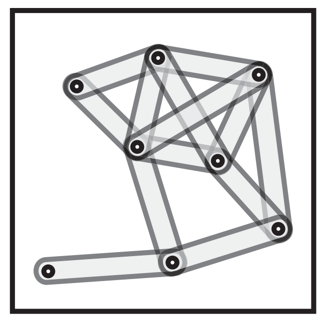

The most familiar simplicial complexes are abstract **graphs**, simplicial complexes without simplices of dimension higher than one. An abstract graph is often presented as a pair (V, E), where V is the set of **vertices** (or 0-simplices) and E is a collection of distinct unordered pairs in V. These **edges** are the 1-simplices of the graph. Abstract graphs are ideal for expressing pairwise relations between objects. Relations of a higher order point to simplicial models. One class of elementary but illuminating examples occurs naturally in the form of independence of objects. Given a finite collection of objects $\mathcal{O} = \{x_\alpha\}$, an **independence system** is a collection of unordered subsets of $\mathcal{O}$ declared *independent*. Any such system must be closed with respect to restriction (any subset of an independent set is also independent) and thus defines a simplicial complex: the **independence complex** $\mathcal{I}_\mathcal{O}$ is the abstract simplicial complex on vertex set $\mathcal{O}$ whose k-simplices are collections of $k+1$ independent objects.

Example 2.1 (Statistical independence) ⊙

Independence occurs in multiple contexts, including linear independence of a collection of vectors or linear independence of solutions to linear differential equations. More subtle examples include statistical independence of random variables: recall that the random variables $\mathcal{X} = \{X_i\}_1^n$ are statistically independent if their probability densities f_{X_i} are jointly multiplicative, *i.e.*, the probability density $f_\mathcal{X}$ of the combined random variable $(X_1, \ldots, X_n)$ satisfies $f_\mathcal{X} = \prod_i f_{X_i}$. The independence complex of a collection of random variables compactly encodes statistical dependencies. ⊙

As with the setting of manifolds, one should rapidly metabolize the formal definition and progress to drawing pictures. One topologizes a simplicial complex X as a quotient space built from topological simplices. The **standard k-simplex** is:

$$\Delta^k := \left\{ x \in [0,\infty)^{k+1} : \sum_{i=0}^{k} x_i = 1 \right\}.$$

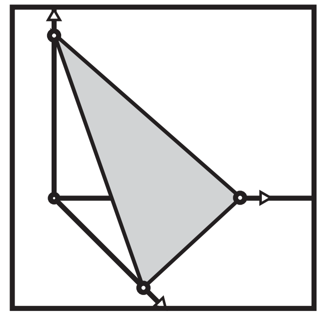

The faces of Δ^k are copies of Δ^j for $j < k$ via restriction of $\mathbb{R}^{k+1}$ to coordinate subspaces; via this restriction, one sees how these faces are *attached* along the boundary of Δ^k. One imagines building an abstract simplicial complex X into a space by producing one copy of Δ^k for each k-simplex of X, then attaching these together via faces. This ill-named **geometric realization** of X is performed inductively as follows. Define the **k-skeleton** of X, $k \in \mathbb{N}$, to be the quotient space:

$$X^{(k)} := \left(X^{(k-1)} \bigcup \coprod_{\sigma : \dim \sigma = k} \Delta^k \right) \Big/ \sim, \qquad (2.1)$$

where $\sim$ is the equivalence relation that identifies faces of Δ^k with the corresponding combinatorial faces of σ in $X^{(j)}$ for $j < k$.

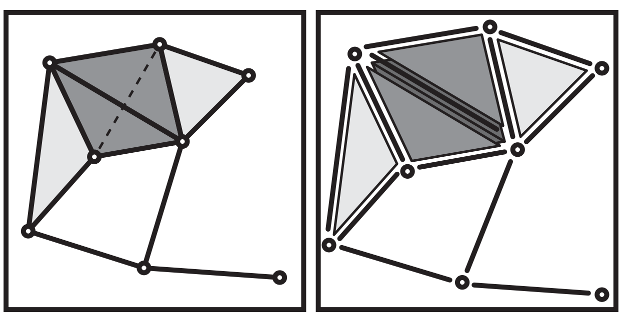

Thus, $X^{(0)} = S$ is a discrete set. The 1-skeleton $X^{(1)}$ is a space homeomorphic to a collection of vertices (S) and edges connecting vertices: a topological graph. By construction, $X^{(k)} \supseteq X^{(k-1)}$ and one therefore identifies X with

$$X = \bigcup_{k=0}^{\infty} X^{(k)}.$$

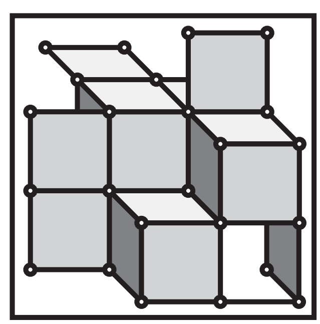

The abuse of notation is intentional; X (the space) will be conflated with X (the abstract simplicial complex). As a space, X is given the **weak topology**: a subset $U \subset X$ is open if and only if $U \cap X^{(k)}$ is open for all k.

A simplicial complex is one of many possible structures for building a space. For example, one may take the basic k-dimensional building block to be a k-cube I^k, where $I = [0,1]$. The faces of the cube are obviously smaller-dimensional cubes; attaching maps between faces identify cubes and are likewise clear. The resulting class of **cubical complexes** are natural in many applications. Digital cameras store data in pixels on a 2-dimensional cubical complex. Numerical analysis and finite element methods work well on 3-dimensional cubes (or *voxels*) arranged in a lattice. Cubes and simplices are examples of compact convex polytopes: cell complexes built from arbitrary compact convex polytopes are likewise easy to work with.

More general still is the class of spaces referred to as **CW complexes**. These are built inductively, as per Equation (2.1) with basis cells closed balls B^k, and for attaching maps *arbitrary* continuous functions from the boundary ∂B^k to $X^{(k-1)}$. Though this tends to be among the best class of complexes for doing topology, the reader may find it simpler to work at first with **regular cell complexes**, whose cells are balls and whose boundary maps are *homeomorphisms* of ∂B^k to their images in $X^{(k-1)}$.

Example 2.2 (Unstable manifolds) ⊙
The flexibility of a CW complex is sometimes needed. Consider the case of a smooth vector field V on a compact 2-manifold $S \subset \mathbb{R}^3$. Assume $V = -\nabla h$ is the gradient field for some height function $h\colon S \to \mathbb{R}$ (so that flowlines of V *flow downhill*), and that there are a finite number of nondegenerate critical points of h (*i.e.*, V is transverse to the zero-section of T_*S), so that each is either a source, sink, or saddle of V. Recall from §1.4 that each fixed point p has a stable ($W^s(p)$) and unstable ($W^u(p)$) manifold, whose intersection is p. By nondegeneracy (and the Stable Manifold Theorem [258]), these are in fact submanifolds of S.

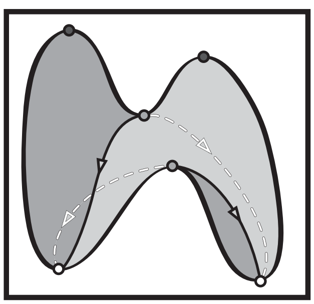

Consider the collection of unstable manifolds $\{W^u(p) : p \in \mathrm{Fix}(V)\}$. These partition the surface S into points (unstable manifolds of sinks), open line segments (unstable manifolds of saddles), and open discs (unstable manifolds of sources). These together define a CW structure on S, where: the 0-skeleton $S^{(0)}$ is the union of (unstable manifolds of) sinks; $S^{(1)}$ is the union of unstable manifolds over sinks and saddles obtained by *gluing* intervals to the sinks via the dynamics; and $S^{(2)} = S$ glues the discs defined by unstable manifolds of sources along the 1-skeleton. This simple example can be greatly expanded and forms the basis of **Morse theory**, to come in Chapter 7. ⊙

The distinction between various types of cell complexes is intricate, and the beginner should beware getting lost at this juncture.This text uses simplicial and cubical complexes frequently; regular cell complexes often; and general CW complexes only occasionally. The phrase *for a cell complex* implies that it holds in the fully general CW setting. The remainder of this chapter is a menagerie of interesting simplicial and cellular complexes that have been found useful in applications to date.

2.2 Vietoris-Rips complexes and point clouds

Consider a discrete subset $\mathcal{Q} \subset \mathbb{R}^n$ of Euclidean space. If one receives $\mathcal{Q}$ as a collection of data points – a **point cloud** – sampled from a submanifold, it may be desirable to reconstruct the underlying submanifold from the point cloud. A graph based on a point cloud $\mathcal{Q}$ can be efficacious in discerning shape (as any child doing a dot-to-dot can attest), but a simplicial complex may improve matters.

Choose a constant $\epsilon > 0$. The **Vietoris-Rips complex** of scale ϵ on $\mathcal{Q}$, $\mathrm{VR}_\epsilon(\mathcal{Q})$, is the simplicial complex whose simplices are all those finite collections of points in

$\mathcal{Q}$ of pairwise-distance $\leq \epsilon$ (which is clearly closed under subsets). The *hope* is that $\mathrm{VR}_\epsilon(\mathcal{Q})$ gives a good approximation to the structure underlying the point cloud. Note the difficulty in specifying what that means, as well as choosing the *correct* ϵ to best approximate. Although this complex is based on points in $\mathbb{R}^n$, there may be simplices of dimension much larger than n (when sufficiently many points are crowded together).

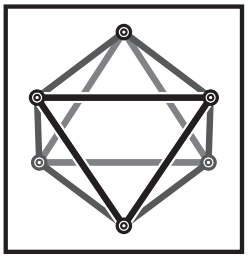

The degree to which a Vietoris-Rips complex VR (of some fixed but unwritten ϵ) accurately captures the topology of the point cloud is not obvious. Consider a point cloud $\mathcal{Q}$ in $\mathbb{R}^n$ and the projection map $\mathcal{S}\colon \mathrm{VR} \to \mathbb{R}^n$ taking the vertices $\mathrm{VR}^{(0)}$ to $\mathcal{Q}$ and taking a k-simplex of VR to the convex hull of the associated vertices in $\mathcal{Q}$. The image of $\mathcal{S}$ in $\mathbb{R}^n$ is called the **shadow**, Sh, of the Vietoris-Rips complex. There is no hope of VR and Sh being homeomorphic, as the domain of the projection $\mathcal{S}$ is likely to have higher dimension than the codomain. Even homotopy equivalence is too much to ask, as it is possible to arrange data points in $\mathcal{Q}$ in $\mathbb{R}^2$ so as to have a simplicial sphere $\mathbb{S}^k$ for arbitrary $k > 1$ which projects in Sh to a convex set: such bubbles are artifacts of the Vietoris-Rips complex and not representative the the point cloud topology. Nevertheless the Vietoris-Rips complex *seems* to capture the topology of large-scale holes correctly: see [64].

2.3 Witness complexes and landmarks

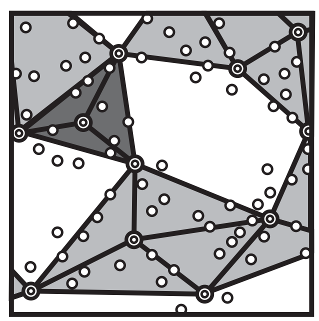

The Vietoris-Rips complex of a large data set may be too unwieldy to store and manipulate. To address this concern, the **witness complex** constructions of Carlsson and de Silva [86] greatly reduce size at the potential expense of topological precision. This has become an intricate subject, with a variety of related definitions built to suit different data sets. The following is a elementary version, suitable for an introduction.

Given a (large) point cloud of nodes $\mathcal{Q}$ in a Euclidean space $\mathbb{R}^n$, choose a (small) set of **landmarks** $\mathcal{L} \subset \mathcal{Q}$. The **weak witness complex**, $W_\mathcal{L}$, is an abstract simplicial complex on the vertex set $\mathcal{L}$ defined as follows. A subcollection of points $S \subset \mathcal{L}$ is a simplex in the witness complex if and only if for each subset $T \subset S$ there is a **weak witness** $x_T \in \mathcal{Q}$ with every point in T closer to x_T than to any point in $\mathcal{L} - T$. Simple examples indicate that for a reasonable choice of landmarks, the witness complex provides a decent topological approximation to the underlying structure of the point cloud. More recent generalizations to *strong* and *parameterized* witness complexes span the gap between computability and rigorous recovery of the correct topological type [55].

2.4 Flag complexes and networks

The Vietoris-Rips complex is an example of a more general class of simplicial complexes that fill a frame. One of the advantages of this construction is its parsimonious representation with regards to input and storage of the data. To encode a typical simplicial complex into a form that computational software can manage, it is necessary to list all the simplices. In contrast, note that the Vietoris-Rips complex is entirely specified by knowing the vertices $\mathcal{Q}$ and the edges (depending on ϵ). The remaining, higher-dimensional, simplices are all accounted for: the graph determines the complex.

Given any graph, the **flag complex** (or **clique complex**) is the maximal simplicial complex having the graph as its 1-skeleton. Otherwise said, whenever there appears to be the outline of a simplex, one fills it in. The Vietoris-Rips complex is the flag complex associated with the distance-ϵ proximity graph of a point cloud vertex set $\mathcal{Q}$.

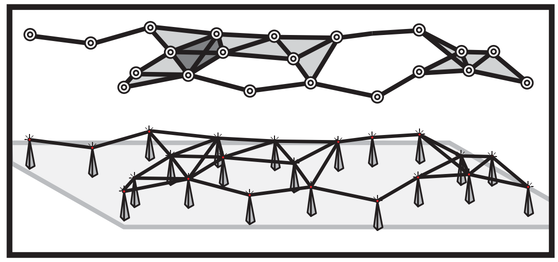

Flag complexes also arise as higher-dimensional models of communications networks. Consider a collection of points $\mathcal{Q} \subset \mathbb{R}^2$ that represent the locations of nodes in a wireless communications network. Ignoring the details of the communications protocols, consider the following simple model of communication: nodes broadcast their unique identities, received by nodes which are within range; these neighbors then establish communication links and in so doing assemble an *ad hoc* network.

Due to irregularities in transmission characteristics, ambient characteristics of the domain, and signal bounce, the communication links formed may not be solely a function of geometric distance between nodes. Nevertheless, forming the flag complex associated with these edges may recover some of the rough structure implicit in the global network. In §5.6, a method for addressing coverage-type problems is considered using these flag complexes.

2.5 Čech complexes and random samplings

The problem of building topologically accurate simplicial models is classical. In the context of the Vietoris-Rips complexes of §2.2, it is frustrating to have higher-dimensional features appearing. One way to circumvent this is to consider the higher-dimensional simplices of the Vietoris-Rips complex more carefully. Given a point cloud $\mathcal{Q}$ in $\mathbb{R}^n$ and a length parameter $\epsilon > 0$, define the **Čech complex** $\check{C}_\epsilon$ to be the simplicial complex built on $\mathcal{Q}$ as follows. A k-simplex of $\check{C}_\epsilon$ is a collection of $k+1$ distinct elements x_i of $\mathcal{Q}$ such that the net intersection of *diameter* ϵ balls at the x_i's is nonempty. The Čech complex is a subcomplex of the Vietoris-Rips complex VR_ϵ; it is sometimes a proper subcomplex. The disadvantages of a Čech complex are its storage requirements (one cannot simply store the 1-skeleton and fill the rest in) and its construction (one must check many intersections to build the full complex). These are compensated for by a topological accuracy: the Čech complex $\check{C}_\epsilon$ is *always* homotopic to the union of diameter-ϵ balls about $\mathcal{Q}$ (see Theorem 2.4).

Niyogi, Smale, and Weinberger [241] have used this property of Čech complexes to prove a result about randomly sampled point clouds. Consider a smooth submanifold $M \subset \mathbb{R}^n$. Let $\mathcal{Q}$ be a collection of points sampled at random from M. It is intuitively obvious that for a sufficiently dense sampling, the union of properly sized balls about the points is homotopic to M: this union of balls can be captured by the Čech complex of $\mathcal{Q}$. The ability of the Čech complex to approximate M is conditional upon the density of sampling and the manner in which M is geometrically embedded in $\mathbb{R}^n$. Let τ denote the **injectivity radius** of M: the largest number such that all rays orthogonal to M of length τ are mutually nonintersecting.

Theorem 2.3 ([241]). *Assume a collection $\mathcal{Q}$ of points on a smooth compact submanifold $M \subset \mathbb{R}^N$, where M has injectivity radius τ. If the density of $\mathcal{Q}$ is such that the minimal distance from any point of M to $\mathcal{Q}$ is less than $\epsilon/2$ for $\epsilon < \tau\sqrt{3/5}$, then the Čech complex $\check{C}_{2\epsilon}(\mathcal{Q})$ deformation retracts to M.*

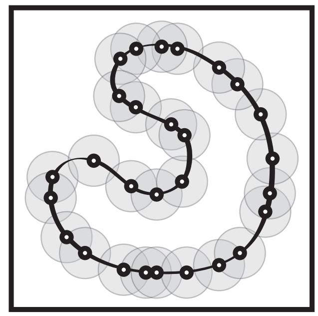

Additional results, such as criteria for these hypotheses for uniformly randomly distributed points on M (with or without various types of noise) are also available and relevant to the statistics of point clouds. This result has been useful in validating data analysis from experiments in, *e.g.*, materials science [82, 83].

2.6 Nerves and neurons

As the Vietoris-Rips complex is a metric-ball version of the more general flag complex construction, so also is the Čech complex the metric-ball version of a more general object. Given a collection $\mathcal{U} = \{U_\alpha\}$ of (say, compact) subsets of a topological space X, one builds the **nerve** of $\mathcal{U}$, $\mathcal{N}(\mathcal{U})$, as follows. The k-simplices of $\mathcal{N}(\mathcal{U})$ correspond to nonempty intersections of $k+1$ distinct elements of $\mathcal{U}$. For example, vertices of the nerve correspond to elements of $\mathcal{U}$; edges correspond to pairs in $\mathcal{U}$ which intersect nontrivially. This definition respects faces: the faces of a k-simplex are obtained by removing corresponding elements of $\mathcal{U}$, leaving the resulting intersection still nonempty. Examples of nerves based on, *e.g.*, convex subsets of $\mathbb{R}^n$, suffice to suggest the following classical generalization of the result for Čech complexes.

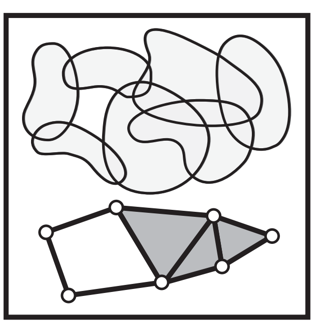

Theorem 2.4 (Nerve Lemma). *If $\mathcal{U}$ is a finite collection of open contractible subsets of X with all non-empty intersections of subcollections of $\mathcal{U}$ contractible, then $\mathcal{N}(\mathcal{U})$ is homotopic to the union $\cup_\alpha U_\alpha$.*

One of the more recent and interesting applications of nerves is in neuroscience. The work of Curto and Itskov [80] considers how neural activity can represent external environments. In particular, the authors consider the impacts of external stimuli on certain **place cells** in the dorsal hippocampus of rats. These cells experience dynamic electrochemical activity which is known to strongly correlate to the rat's location in physical space. Each such cell group is assumed to determine a specific location in space, and the collection of such location patches, or **place fields**, forms a collection $\mathcal{U}$ satisfying the hypotheses of the Nerve Lemma, assuming that place fields exist and are stable, omni-directional, and have firing fields that are convex. Curto and Itskov argue that these assumptions are generally satisfied for place fields of dorsal hippocampal place cells recorded from a freely foraging rat in a familiar open field environment.

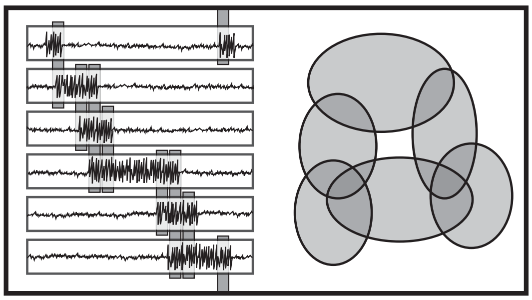

In an experiment, the place cells are monitored, with activity recorded as a time-series. These time series show occasional bursts of activity, such **spike trains** being common data forms in neuroscience. Cell groups are identified by correlating spike train firings that occur at the same time (within a window of time based on some parameter chosen appropriately). The correlation of spiking activity provides the intersection data of the place fields. This, in turn, allows for the computation of the nerve of the place fields based on spike train correlation. These investigations suggest that rats may build a structured representation of their external environment within an abstract space of stimuli. Of note is the lack of metric data – the construction of the physical environment is purely topological and can be effected without reference to coordinates.

2.7 Phylogenetic trees and links

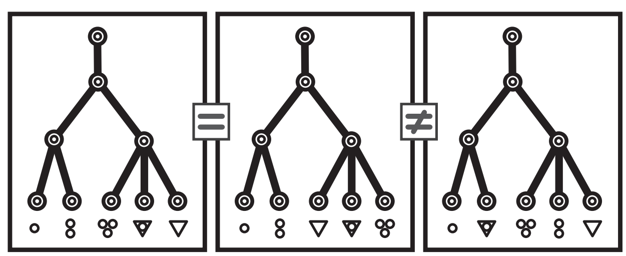

Mathematical biology is ripe for topological tools, given the noisy nature of Nature. In addition to the neuroscience example above is an analogue of configuration spaces associated to genetic data in the form of **phylogenetic trees** – data structures for organizing and comparing taxonomic and genetic data. The relevant objects of study are rooted, end-labeled, metric tress. A **tree** is a graph without cycles. The **leaves** of a tree are vertices of degree 1 (only one edge is attached), and a **root** is a dedicated leaf. A phylogenetic tree is a rooted tree with (1) leaves labeled by data (genes, phenotypes, or other taxonomica), and (2) all interior edges (those not connected to a leaf) labeled with positive real numbers (and represented as the length of the edge). These trees can be readily illustrated as planar graphs, but the embedding of the tree into the plane is *not* part of the data.

Let $\mathcal{T}_n$ denote the following **tree space** of [37]: the space of rooted metric trees with n labeled (non-root) leaves. For simplicity, assume that the label on the root is 0 and labels on the leaves are $1, 2, 3, \ldots, n$. This space has a natural cube complex structure as follows. All non-leaf vertices of the graphs are assumed to be essential in that the number of edges attached is greater than two.

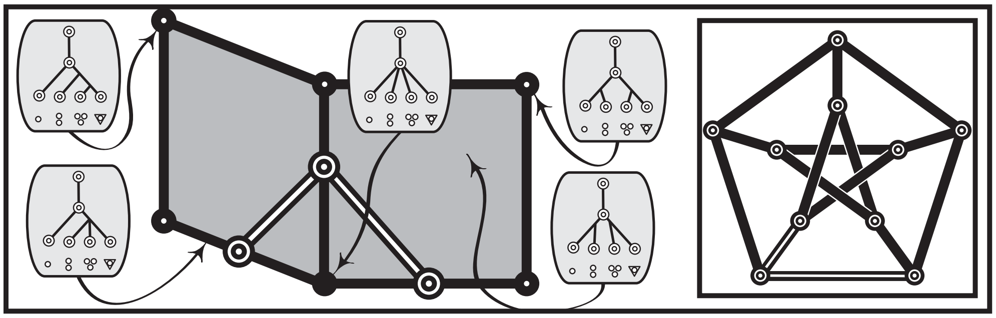

Two such trees are *equivalent* if there is a mapping between them that preserves the graph structure (takes vertices to vertices and edges to edges), the leaf labels, and the interior edge lengths; they are *isomorphic* if there is a mapping between them that preserves the graph structure and the labels. A fixed isomorphism class of phylogenetic trees is parameterized by the interior edge lengths: this yields an open cube of dimension equal to the number of interior edges.

To compactify this to a closed cube, reparameterize all interior edge lengths to the open interval $(0, 1)$ and add the boundary faces. Those faces with one or more 0 factors may be thought of as having the corresponding interior edges collapsed ("zero length edges"). These edge-collapse faces are identifiable as trees of a different isomorphism class. By gluing together all such cubes according to isomorphism class, a finite cube complex $\mathcal{T}_n$ is obtained.

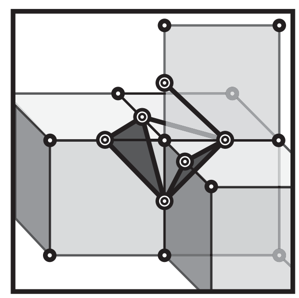

As topological spaces, $\mathcal{T}_n$ may seem too simple: it is contractible to an *origin*, the radial graph with no interior edges, by shrinking each interior edge continuously to length zero. Despite this simplicity, the manner in which the various cubes are assembled is topologically (and geometrically) of interest. The best way to analyze this assembly is through a simplical model called a **link**. Let X denote a cell complex built from simplices, cubes, or other simple polyhedral basis cells, and let $v \in X^{(0)}$ be a 0-cell. The link of v in X, $\ell k_X(v)$, is the simplicial complex whose k-simplices consists of $(k+1)$-dimensional cells attached to v, with faces inherited from faces of cells in X.

Example 2.5 (Tree spaces) ⊙

The tree space $\mathcal{T}_3$ is a "letter Y." The space $\mathcal{T}_4$ consists of 15 2-dimensional squares (corresponding to the 15 possible binary rooted trees with 4 labeled leaves) glued together so as to have a common corner (corresponding to the radial rooted tree with 4 leaves) and a link equal to the 10-vertex, 15-edged **petersen graph**. As one might guess, links of tree spaces $\mathcal{T}_n$ grow quickly in size and complexity with n; their topology, however, is classifiable.

Theorem 2.6 ([37]). *The link of the origin in $\mathcal{T}_n$ is a flag complex that has the homotopy type of $(n-1)!$ copies of $\mathbb{S}^{n-3}$ glued together at a single point.* ⊙

There are numerous reasons why scientists and engineers want to work with *spaces* of objects rather than merely the objects themselves. The tree spaces $\mathcal{T}_n$ (though arising in other contexts in topology and algebraic geometry as certain moduli spaces) were generated in response to challenges in statistics associated to genetic data. Experimental data is used to generate phylogenetic trees, and this data does not spring forth fulled formed from the mind of the researcher: the data are often noisy and not perfectly repeatable. A collection of numerical or vector-valued data can be averaged, but what does it mean to average a sequence of phylogenetic trees in a scientifically meaningful manner? This requires a geometry (and thus a topology) on the space of all possible phylogenetic trees.

2.8 Strategy complexes and uncertainty

Many of the complexes of this chapter approximate manifolds or provide models of data; however, complexes also have broad applicability in engineering systems as spaces that collate states. Consider a **transition graph**, a directed graph X whose vertices V represent states of an abstract system, and whose edges E represent deterministic actions. For concreteness, consider the example [109] of a transportation network, where vertices are locations and edges represent flights, freeways, subways, trains, or other discrete enter-exit modes of transport. Other examples include motion-planning in robotics with discrete states (robots move from landmark to landmark) or manufacturing processes (states are subassemblies; edges are assembly steps).

Planning on the transition graph is straightforward: find a directed path from initial to goal vertices, then execute the appropriate actions sequentially. However, stepping onto a flight bound for Chicago at noon does not guarantee arrival, either at Chicago or at noon. Worse still are the kinds of uncertainties precipitated by an adversary. Ill weather and air travel does not *exactly* fit that scheme, but other situations do exhibit adversarial forms of uncertainty for which, if something *may* go amiss, it *must* be mitigated.

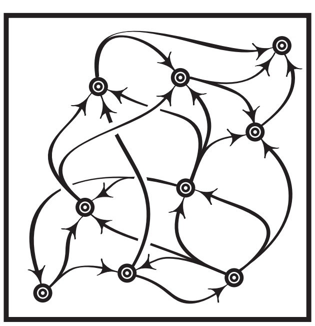

Erdmann [108, 109] models adversarial uncertainty in the transition graph via *branched* edges. Actions are deterministically chosen, but once chosen, the outcome is one of several possibilities, out of the control of the chooser. These possibilities are *not necessarily* chosen probabilistically; an adversary may determine any outcome among the terminal edges. The state chosen by the adversary is known to the game-player, but only *after* the nondeterministic action has been executed. In such a setting, the notion of a guaranteed strategy for reaching a goal state despite adversarial meddling seems difficult if not impossible.

The question of path-planning within a nondeterministic transition graph is one amenable to simplicial data structures. Given a nondeterministic transition graph X, define the **strategy complex**, Δ_X, as the abstract simplicial complex defined on the set of actions $E(X)$, in which a k-simplex is a collection of $k+1$ disjoint actions whose unions (including all branches of any chosen nondeterministic action) contain *no directed cycles*. This relation is closed under restriction (a smaller collection of such edges clearly also contains no directed cycle), and is sensible, as a directed cycle means that a skilled adversary can cycle the player through states *ad infinitum*. The intuition is that Δ_X consists of *trap-free* strategies: each simplex of Δ_X is a strategy for progressing to one or more states.

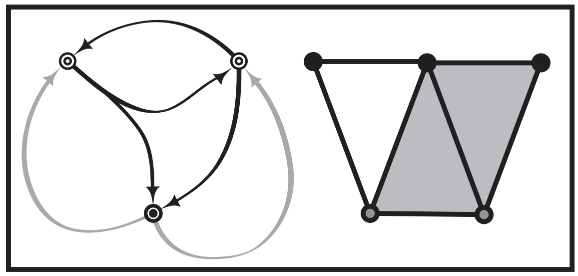

For $g \in V(X)$ a goal state, say that X has a *complete* strategy for attaining g if there exists a mapping from $V(X)$ to $E(X)$ which associates to v an action based at v, so that one attains g with the execution of at most N actions for N sufficiently large. Note: the goal is certain, but the path is not. The obstruction to attainability lies in the strategy **loopback** complex: denote by $\Delta_{X\leftarrow g}$ the strategy complex of X augmented by deterministic arrows from g to each $v \in V(X) - g$.

Theorem 2.7 ([108, 109]). *A nondeterministic transition graph X possesses a complete strategy for attaining a goal $g \in V(X)$ if and only if*

$$\Delta_{X\leftarrow g} \simeq \mathbb{S}^{\#V-2}.$$

Otherwise, it is contractible.

The proofs of this and related theorems use nerves extensively. Additional results include realizability of simplicial complexes as strategy complexes, incorporation of stochastic nondeterminism, and decompositions of strategy complexes by means of deterministic subsystems.

2.9 Decision tasks and consensus

A related use of simplicial structures assists in multi-agent decision-making. Consider the problems of consensus or distribution among collaborative agents, a common scenario in biology (animal flocking), sociology (voting, beliefs), engineering systems (synchronized network clocks), and robotics (cooperative action). The problem of collaborative consensus has a large literature: mathematical approaches depend sensitively on the modeling assumptions. Most work on distributed consensus phenomena [227, 287] focuses on graph-theoretical methods. Hidden beneath these 1-dimensional models is a rich higher-dimensional topological structure underlying influence, consensus, and division.

Example 2.8 (Allocation) ⊚

Consider the problem of N transmitters attempting to broadcast without interference over a frequency spectrum with C distinct, non-interfering channels. How may these channels be assigned? What if certain transmitters have frequency constraints? What does the space of solutions to this problem look like?

Build a simplicial complex on a vertex set of choices: vertices correspond to transmitter n choosing channel c, for each allowable assignment (n, c). A k-simplex consists of $k+1$ transmitters having made compatible (distinct channel) choices,

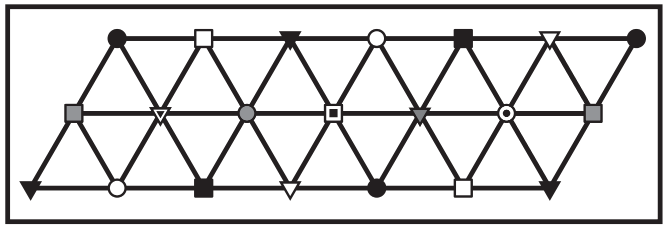

and the other transmitters having unassigned frequencies. This is closed under restriction, since letting one or more assigned transmitter release a frequency leaves the remaining transmitters still compatible. The resulting simplicial complex is rightly viewed as a *configuration space* of solutions to the allocation problem. For example, three transmitters choosing from four channels gives rise to a simplicial 2-torus $\mathbb{T}^2$. ⊙

Work of Herlihy and Shavit [178] pioneered a topological approach to consensus/decision problems in the context of distributed, fault-tolerant computation. In this setting, it is presumed that there are N independent processors P_i who share a common read/write memory element for communication purposes. For simplicity and concreteness, assume that there are $M > 1$ possible labels or states and that each processor has an initial preference. The problem is whether there exists an algorithm for the processors, via read/write communication to the common memory, to come to deterministic finite-time consensus. The interesting catch is that all computations are *asynchronous* (each processor works at its own speed) and *faulty* (sometimes, a processor crashes and can engage no further in communication). The problem in reasoning about such systems is that processors which take a very long time to reach decisions are not readily distinguished from processors which have failed and will never respond. A *fault tolerant* distributed algorithm is one which will terminate in finite time even in the event that some (but not all) processors fail irrevocably.

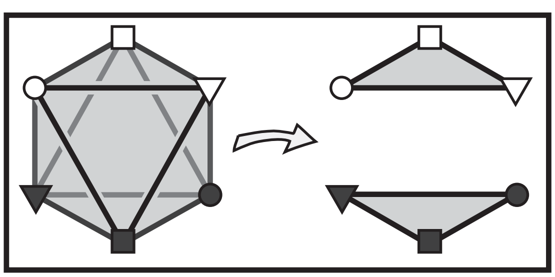

The approached used in [178] begins with a simplicial model of the problem. Given N processors with M possible initial states, define the **input complex** $\mathcal{I}$ to be the abstract simplicial complex whose k-simplices are labelings of the states of $k+1$ distinct processors. The **output complex** $\mathcal{O}$ is the abstract simplicial complex whose k-simplices are labelings of the states of $k+1$ distinct processors which are consistent with a desired output. The intuition behind a k-simplex is that $k+1$ of the processors exhibit a legal system state and the $N-(k+1)$ remaining processors have failed to report. Because of this, the topology of the complexes (encoded by the simplex faces) correlates with failure of processors. A **decision task** consists of the input and output complexes, along with a set of constraints. For example, if all N processors begin in consensus with the same label, then the putative consensus algorithm must of necessity output that precise label as the consensus state. Because of the relationship between faces and failures, it is possible to reduce questions of existence of distributed fault-tolerant consensus algorithms to the existence of certain maps from the input complex to the output complex.

Example 2.9 (Consensus) ⊚

In the case of N processors trying to come to binary consensus (*i.e.*, two labels, zero and one), the input complex $\mathcal{I} \simeq \mathbb{S}^{N-1}$ is a simplicial sphere and the output complex $\mathcal{O} = \Delta^{N-1} \sqcup \Delta^{N-1}$ is a disjoint pair of simplices: all zeros and all ones. The decision task is to come to consensus, with the proviso that initial consensus terminates the program immediately. Thus, any wait-free fault-tolerant protocol would have to induce a surjective continuous map from $\mathcal{I}$ to $\mathcal{O}$. This is impossible, since $\mathcal{I}$ is connected and $\mathcal{O}$ is not. Thus, there is no distributed asynchronous fault-tolerant solution to the problem of binary consensus. That in itself is no surprise; yet this approach has resolved more complex consensus problems for which other solutions were unknown. These deeper examples rely not on common-sense notions of connectivity but rather on holes of higher order – the homology of Chapter 4. ⊚

One of the key theorems [178] gives *necessary and sufficient* conditions for consensus stated in terms of existence of maps (with details concerning simplicial subdivisions and colorings). The salient feature of the result is that the topology of the decision task can determine whether an algorithm *exists*. This type of implicit inference is a hallmark of topological methods.

2.10 Discretized graph configuration spaces

The strategy complexes of §2.8 and the allocation, input, and output complexes of §2.9 are all viewable as a simplicial sort of configuration space. The types of configuration spaces used in robotics (from §1.2 and §1.5) are neither simplicial nor cellular complexes. One can, however, approximate such spaces with cell complexes. This is most easily done on the already-discrete structure of a graph.

Consider a finite graph X, now thought of not as an abstract system of states, but rather as a physical, geometric highway of paths. Keeping track of vehicles or robots (or *tokens*) on X leads, as in §1.5, to configuration spaces. The configuration space of n labeled points on X, $\mathcal{C}^n(X)$, is relevant to motion planning in robotics when the automated vehicles are constrained by tracks, guidewires, or optical paths [149]; however, $\mathcal{C}^n(X)$ is in general neither a manifold nor a simplicial/cell complex. The following constructions [2] are approximations to these spaces by cubical complexes.

Define the **discretized configuration space** of X as:

$$\mathcal{D}^n(X) := (X \times \cdots \times X) - \tilde{\Delta}, \tag{2.2}$$

where $\tilde{\Delta}$ denotes the set of all open cells in X^n whose closures intersect the topological diagonal Δ. Equivalently, $\mathcal{D}^n(X)$ is the set of configurations for which, given any two tokens on X and any path in X connecting them, the path contains at least one entire (open) edge. Thus, instead of restricting tokens to be at least some intrinsic distance ϵ apart, one now restricts tokens on X to be at least one full edge apart. Note that $\mathcal{D}^n(X)$ is a subcomplex of the cubical complex X^n and a subset of $\mathcal{C}^n(X)$ (it does not contain partial cells that arise when cutting along the diagonal), and is, in fact, the largest subcomplex of X^n that does not intersect Δ.

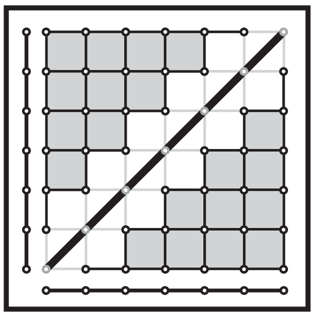

With this natural cell structure, one can think of the 0-cells of $\mathcal{D}^n(X)$ as *discretized* configurations – arrangements of labeled tokens at the vertices of the graph. The 1-cells of $\mathcal{D}^n(X)$ indicate which discrete configurations can be connected by moving one token along an edge of X. Each 2-cell in $\mathcal{D}^n(X)$ represents two physically independent edges: one can move a pair of tokens independently along disjoint edges. A k-cell in $\mathcal{D}^n(X)$ likewise represents the ability to move k tokens along k closure-disjoint edges in X in an independent manner.

Example 2.10 (Digital microfluidics) ⦾

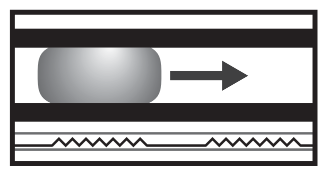

An excellent physical instantiation of the discretized configuration space appears in work on **digital microfluidics** [111], in which small (∼1mm diameter) droplets of fluid can be quickly and accurately manipulated in an inert oil suspension between two plates. The plates are embedded with a grid of wires, providing discrete localization of droplets via **electrowetting** – a process that exploits current-induced dynamic surface tension effects to propel and position a droplet. Applying an appropriate current drives the droplet a discrete distance along the wire grid. The goal of this and related microfluidics research is to create an efficient *lab on a chip* in which droplets of various chemicals or liquid suspensions can be positioned, mixed, and then directed to the appropriate outputs. Using the grid, one can manipulate many droplets in parallel. The discretized configuration space of droplets on the graph of the electrical grid captures the topology of the multi-droplet coordination problem. Note, however, that collisions are not always to be avoided – chemical reactions depend on controlled collisions. ⦾

Example 2.11 (A nonplanar graph) ⦾

Discretization can untangle a complex-looking configuration space. One interesting example is the configuration space of two labeled points on a complete connected graph of five vertices, K_5. The discretized configuration space, $\mathcal{D}^2(K_5)$, is a 2-dimensional cube complex.

Fix a single vertex $v \in \mathcal{D}^2(K_5)$: this corresponds to a pair of distinct labeled tokens on vertices of K_5. From this state v emanate six edges in $\mathcal{D}^2(K_5)$: each token can move to any of the three open vertices of K_5. For each edge e incident to v, the token not implicated by e can move to two of the three available vertices without interfering with e; thus, e is incident to exactly two 2-cells of $\mathcal{D}^2(K_5)$. There are a total of six such 2-cells incident to v, arranged cyclically. Thanks to the symmetry of K_5, the complex looks the same at each vertex, and $\mathcal{D}^2(K_5)$ is thus a topological 2-manifold: a compact surface. With the tools of the next chapter, it will be shown that this surface is of genus six. ⊙

The discretized configuration space is an accurate cubical model of the configuration space.

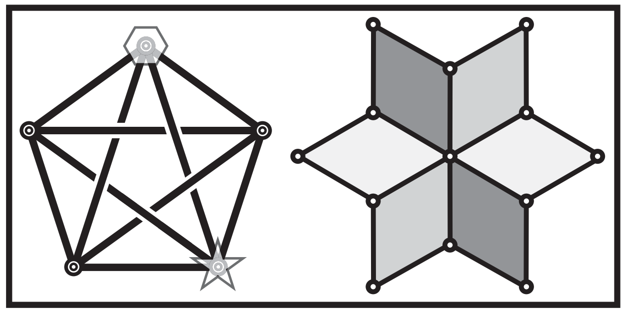

Theorem 2.12 ([2, 245]). *Let X be a finite topological graph (with no edges connecting a vertex to itself) and let $\tilde{X}$ be the subdivision of X that inserts vertices so as to split each edge of X into $n-1$ edges. Then $\mathcal{D}^n(\tilde{X})$ is homotopic to $\mathcal{C}^n(X)$.*

2.11 State complexes and reconfiguration

Discretized configuration spaces of graphs are generalizable to a very broad class of systems – including metamorphic robots, protein chains, and more – which are reconfigurable based on local rules. The resulting configuration spaces, called **state complexes**, are cube complexes with interesting topological and geometric properties.

Fix a graph X and an alphabet $\mathcal{A}$, used to label the vertices V, each such labeling $u\colon V \to \mathcal{A}$ comprising a state u. A **reconfigurable system** is a collection of states $\{u_\alpha\}$, closed under the actions of a fixed set of local reconfigurations, or **generators**. Each generator ϕ consists of a support subgraph $\mathrm{supp}_\phi \subset X$ and an *unordered* pair of **local states**, labelings of the vertex set of supp_ϕ by elements of $\mathcal{A}$. A generator is **admissible** at a given state u if one of the generator's local states matches the restriction of u to supp_ϕ. The **trace** of a generator is the (nonempty) subset of the support $\mathrm{trace}_\phi \subset \mathrm{supp}_\phi$ on which the local states differ.

A simple example comes from discretized configuration spaces of a graph from §2.10. An alphabet $\mathcal{A} = \{0, 1, \ldots, N\}$ labels the vertices of X as empty (0) or occupied with one of N (labeled) tokens. Local reconfigurations are supported on the closure of an edge in X; each of the N generators ϕ has $\mathrm{supp}_\phi = \mathrm{trace}_\phi$ equal to a single closed edge with local state a labeling of one vertex with 0 and the other $0 < i \leq N$ for some i, the other local state reversing these labelings. Applying a sequence of admissible generators is identical to moving labeled agents discretely along X. This system has a well-defined notion of a configuration space, $\mathcal{D}^N(X)$, in which k-dimensional cubes correspond to k local moves that can be executed *in parallel*: this generalizes to any reconfigurable system as follows.

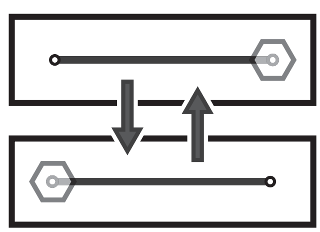

A collection of generators $\{\phi_\alpha\}$ admissible at a given state u is said to **commute** if their supports and traces are pairwise non-intersecting:

$$\text{trace}_{\phi_\alpha} \cap \text{supp}_{\phi_\beta} = \varnothing \qquad \forall \alpha \neq \beta.$$

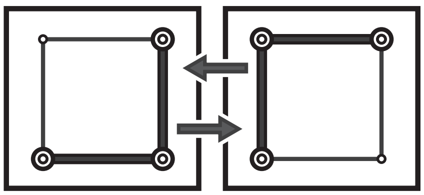

This means that the changes implicated by a generator (its trace) do not impact the applicability of the other generators (their supports): *commutativity encodes physical independence.* Define the **state complex** of a reconfigurable system to be the cube complex $\mathcal{S}$ with an abstract k-cube for each collection of k admissible commuting generators. That is, the 0-skeleton consists of the states in the reconfigurable system; the 1-skeleton is the graph whose edges pair two states which differ by the local states of a single generator; 2-cells correspond to commutative pairs of generators acting on a state, *etc.* For reconfigurable systems that admit only finitely many admissible generators at any given state, the state complex $\mathcal{S}$ is a locally compact cubical complex.

Example 2.13 (Articulated planar robot arm) ⊙

A reconfigurable system models the position of an articulated robotic arm with one end fixed at the origin and which can (1) rotate at the end and (2) "flip" a 2-segment corner from up-right to right-up and vice versa. This arm is *positive* in the sense that it may extend up and to the right only. A grammatical model is simplest: states are length n words in two letters. Specifically, let X be a linear graph on n vertices with an alphabet $\mathcal{A} = \{U, R\}$ encoding *up* and *right* respectively. Generators are of two types: (1) exchange the subwords UR and RU along an edge; (2) change the terminal letter in the word. Supports (and traces) are (1) a closed edge, and (2) the terminal vertex, respectively. Generators commute when the associated subwords are disjoint. The state complex in the case $n = 5$ has cubes of dimension at most three. In this case also, although the transition graph (the 1-skeleton of $\mathcal{S}$) for this system is complicated, the state complex itself is contractible. ⊙

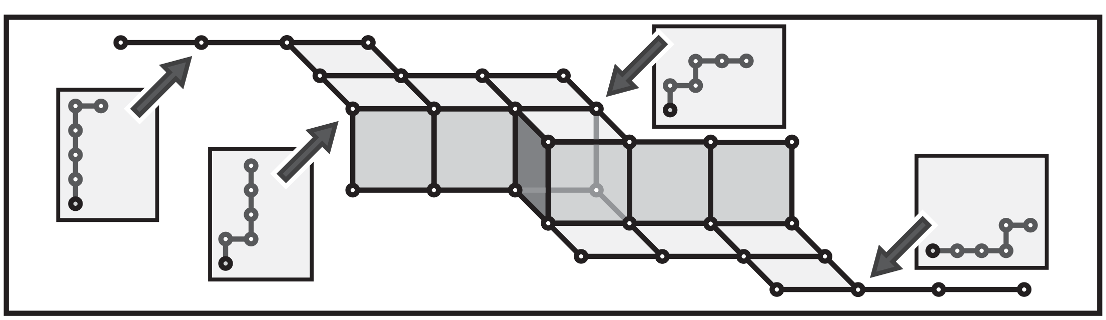

Example 2.14 (Robot path coordination) ⊙

The following example is inspired by robot coordination problems. Consider a collection of n graphs $(X_i)_1^n$, each embedded in the plane of a common workspace (with intersections permitted). On each X_i, a robot R_i with some particular fixed size/shape is free to translate along edges of X_i. The **coordination space** of this system is defined to be the space of all configurations in $\prod_i X_i$ for which there are no collisions – the geometric robots R_i have no overlaps in the workspace.

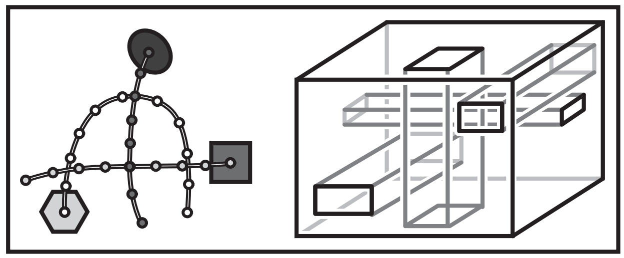

To model this with a reconfigurable system, let the underlying graph be the disjoint union $X = \sqcup_i X_i$. The generators for this system are as follows. For each edge $\alpha \in E(X_i)$, there is exactly one generator ϕ_α. The trace is the (closure of the) edge itself, $\mathrm{trace}(\phi_\alpha) = \alpha$, and its generator corresponds to sliding the robot R_i from one end of the edge to the other. The support, $\mathrm{supp}(\phi_\alpha)$ consists of this trace union any other edges β in X_j ($j \neq i$) for which the robot R_j sliding along the edge β can collide with R_i as it slides along α. The alphabet is $\mathcal{A} = \{0, 1\}$ and the local states for ϕ_α have zeros at all vertices of all edges in the support, except for a single 1 at the boundary vertices of α. Any state of this reconfigurable system has all vertices labeled with zeros except for one vertex per X_i with a label 1. The resulting state complex is a cubical complex that approximates the coordination space. In the case where $X_i = X$ is the same for all i and the robots R_i are sufficiently small, this reconfigurable system has $\mathcal{S} = \mathcal{D}^n(X)$. ⊙

State complexes encompass a wide variety of different systems and can appear to be full of holes, as seen in Examples 2.11 and 2.14 Nevertheless, the types of holes that arise are unusually restricted:

Theorem 2.15 ([153]). *Every state complex $\mathcal{S}$ is* **aspherical***: any map $f : \mathbb{S}^n \to \mathcal{S}$ for $n > 1$ is homotopic to a constant map $\mathbb{S}^n \to \star$.*

Note the recurring theme: state complexes, strategy complexes, tree spaces, and more all seem to be characterized or clarified by the absence of presence of *spheres*, leading one to wonder if these are not the primary building-blocks of topology.

Notes

1. A cell complex structure is a type of finite approximation to a space and, as such, is central to the ability to do computation. However, the type of cell structure matters. An n-dimensional sphere $\mathbb{S}^n$ has a CW structure with exactly two cells (dimensions 0 and n). As a regular cell complex, $2n$ cells suffice. As a cubical or simplicial complex, the minimal total number of cells is exponential in n.
2. One virtue of topological methods is their coordinate-free nature: it is not necessary to know exact locations. If, however, one does possess strong geometric data in the form of coordinates, there are numerous ways to improve on the Vietoris-Rips construction.

One popular approach is the **alpha complex** of Edelsbrunner [104] which restricts a Čech complex by a Delaunay triangulation to produce a topologically accurate simplicial model with computationally small footprint.

3. Theorem 2.3 is of interest primarily for its explicit computation of density bounds – it was previously known [204] that a sufficiently dense sampling recovers the topology of the submanifold.
4. The Nerve Lemma is ubiquitous in topology and its applications. It is often attributed to Leray [206], whose version, concerning sheaves, is much deeper than that stated. It echoes antecedents in the (independent) work of Čech [63] and Alexandrov [8]. The version stated is for open covers, but covers by compact contractible sets also suffice. It is often stated that $\mathcal{U}$ be comprised of convex sets, as these and all nonempty intersections thereof are contractible. It will sometimes be mentioned that the Nerve Lemma breaks down if the sets involved are not contractible. This statement is accurate but should propel the reader to more sophisticated methods (such as the **Leray spectral sequence**) rather than to despair.
5. The survey article [242] gives a broad treatment of applications of geometric and topological combinatorics to a broad class of problems in phylogenetics. This is a very active branch of applied topology and geometry.
6. Many results about the topology of state complexes are highly dependent on their geometric features [153]. There is a precise sense in which these complexes are devoid of positive curvature, and it is this property that most influences their global topology.
7. State complexes are *undirected* versions of the **high-dimensional automata** of Pratt [244]. They may also be viewed as weaker versions of the very interesting class of combinatorial cell complexes known as **hom complexes**. The monograph of Kozlov [198] has a wealth of information on this latter class of spaces.