Written draft of February 10, 2012, based on a lecture given January 5, 2011 at the
AMS Short Course on Computational Topology, JMM2011, New Orleans.

On the Topology of Discrete Planning with Uncertainty

Michael Erdmann

ABSTRACT. This chapter explores the topology of planning with uncertainty in discrete spaces. The chapter defines the *strategy complex* of a finite discrete graph as the collection of all plans for accomplishing all tasks specified by goal states in the graph. Transitions in the graph may be nondeterministic or stochastic. One key result is that a system can attain any state in its graph despite control uncertainty if and only if its strategy complex is homotopic to a sphere of dimension two less than the number of states in the graph.

1. Planning with Uncertainty in Robotics

The goal of Robotics is to animate the inanimate, so as to endow machines with the ability to act purposefully in the world. Roboticists, working in the subfield of planning, create software by which robots reason about future outcomes of potential actions. Using such planning software, robots combine individual actions into collections that together accomplish particular tasks in the world [**29, 30**].

Two fundamental and intertwined issues confound this seemingly straightforward approach. One is world complexity, the other is uncertainty.

1.1. Discrete Modeling. Modeling a seemingly continuous world with finite discrete symbols is a well-known problem at the heart of algorithmic computability [**32, 49**]. Practical robotics largely avoids the uncomputability questions by focusing on fixed levels of granularity deemed to be appropriate for the robot tasks at hand. The ensuing models may take the form of *a priori* shape families for describing the objects that a robot might encounter [**33, 16**], commensurate geometric representations of sensor data (e.g., visual, tactile, proprioceptive, auditory, laser) [**45, 34, 13, 38, 5, 31**], and partitions of configuration or control space into regions within which the dynamics of interaction are invariant [**50, 9**]. For instance, in [**12**] two robot palms manipulate objects based on prior geometric models of the objects and their frictional contact mechanics. A robot planner subdivides the configuration space of the palms and the object being manipulated into volumes within which the relative sliding motions at the contacts have invariant sign. These

2010 *Mathematics Subject Classification.* Primary 68T37, 68T40; Secondary 55U05, 55U10.
Key words and phrases. topology, graph, complex, strategy, robotics, planning, uncertainty.
This work was sponsored by DARPA under contract HR0011-07-1-0002. This work does not necessarily reflect the position or the policy of the Government. No official endorsement should be inferred.

volumes form the states in a discrete graph whose edges represent connectivity as a function of changing palm orientations. The planner constructs a strategy for reorienting the object by searching this graph for a path leading from the object's initial configuration to some desired final configuration.

1.2. Uncertainty. Uncertainty arises in modeling, control, and sensing. Models are inaccurate, control is errorful, and sensors are noisy. Despite uncertainty, roboticists seek to create planners that allow robots to operate purposefully and successfully in the world. The preimage methodology of [**35**] describes a general approach for planning in the presence of control and sensing uncertainty, generalized to model uncertainty in [**11**]. Specialized to physical systems, the planners often take the form of discrete graph searches. The states in these graphs need not simply be the configurations of the physical system. Instead, the graph states represent the information available to the robot at any given instant during plan execution [**30, 6**]. For instance, a common task in manufacturing systems is to reduce the entropy of small parts. These parts arrive in large numbers, jumbled together. An automatic system must orient and localize each part, so a robot can then, with little or no sensing, pick up the part and assemble it onto some product. The SONY SMART system [**43**] is a wonderful real world example. The robot systems described in [**15, 46, 24**] show how to construct planners for similar tasks, using the mechanics of the problem to reduce uncertainty. Each discrete state within the graphs for these planners is in fact a collection of underlying contact states, describing the extent to which the system has localized a part at runtime. Thus sensing uncertainty contributes to the definition of state.

As the previous discussion suggests, uncertainty and granularity are intertwined. A coarse world model produces relative certainty at the expense of expressive power. A coarse controller reduces search branching factors at the expense of local motion precision. A coarse sensor reduces hardware requirements at the expense of instantaneous localization. Tradeoffs between different levels of granularity and uncertainty are possible. Sequences of accurate motions may be traded for careful sensing in orienting parts, for instance. An open question is fine-tuning such tradeoffs so as to optimize system capabilities. It is a problem merely to describe these tradeoffs precisely.

1.3. Planners. There are two basic modes by which planners generate plans. In *backchaining*, a planner starts from a desired goal (possibly a set of configurations in some state space). The planner determines one or more actions (e.g., robot motions) that achieve the goal directly. The preconditions to those actions (e.g., the initial configurations of motions leading to the goal) combined with the original goal, then define a new subgoal. The planner repeats this process until it produces a subgoal that includes the current configuration of the system or until it determines that no such subgoal exists. The set of actions produced by this process constitutes a plan for attaining the original goal from the current configuration, and in fact, from any configuration satisfying the preconditions of any of the actions.

In *forward-chaining*, the planner starts from the current configuration. The planner determines the outcomes of all possible actions whose preconditions are currently satisfied. Then the planner repeats this process starting from each of the outcomes just determined, and so forth, until it produces a frontier of outcomes that satisfies the goal conditions or determines that the goal cannot be attained.

There exist numerous variations of these basic modes, for instance, simultaneously forward-chaining and backchaining [**4, 3, 30**].

1.4. Collections of Plans. Whatever the precise planning process, the result of planning is a plan for attaining some particular goal, perhaps from some particular initial configuration. Often when planning with robots, such unique plans are all one needs. In manufacturing, for instance, one may only need to know how to assemble, not disassemble. Yet, in principle, one could run the planner for all possible goals that are describable, not just those one explicitly needs. Doing so would reveal global system capabilities.

Plans fit together like jigsaw pieces. One can remove an action from a particular plan and obtain a new plan, with perhaps a slightly different goal. Sometimes one can add an action to a given plan, perhaps as a redundant backup, without changing the plan's outcomes. In other cases, adding an action might change the possible outcomes, possibly creating an infinite loop. By studying this jigsaw puzzle one gains insight into the dependence of a system's capabilities on design choices.

Moving up a level, one may thus view any one system as a point in a larger design space. If one understands the capabilities of a system and how those capabilities change when one alters the underlying system model, available actions, or sensing capabilities, one can begin to address the granularity tradeoffs mentioned earlier. For instance, in a directed graph one can readily decide whether the graph is strongly connected [**1**], modifying it accordingly if one desires different connectivity. No such straightforward tools exist currently for describing the capabilities of uncertain systems. This chapter proposes methods to help fill that deficit.

1.5. Topology of Plans. The aim of this chapter is to explore the topology of the collection of all plans that exist for an uncertain system and in so doing to characterize the system's capabilities. The many robotics results Ghrist found via algebraic topology [**10, 22, 20, 21**] motivate this exploration.

Our exploration focuses on systems that may be modeled as finite discrete spaces. As discussed, the states in such discrete spaces may represent fairly complicated system properties, so the tools presented here should have broad applicability. The core of this work has appeared previously as a robotics paper [**14**]. This chapter expands the topological perspective, generalizing several of the earlier results.

Beyond robotics applications, the research presented here is inspired topologically by the study of the collection of all partial orders on n items [**8, 26, 27**]. Two types of uncertain actions appear in this chapter, nondeterministic and stochastic. In the nondeterministic setting, as the earlier description of backchaining suggests, one may view a single plan as a particular partial order on a system's state space. When executing the plan, the system will visit states in some order consistent with this partial order. There will never be any cycling and the system will eventually wind up at the goal.

The collection of all plans is therefore related to the collection of all partial orders on the state space. A difference between our work and previous work is that the primitive motions in our collection of plans are nondeterministic actions rather than directed edges, producing slightly more complicated primitive partial orders (called *atoms* in the language of partial orders) than the single comparators $\{x > y\}$ that directed edges produce. By allowing these additional types of atoms, one can generate homotopy types of any finite simplicial complex, not just the spheres and points that are possible when atoms are single comparators [**26**]. From that perspective, nondeterministic planning is a natural "physical" realization of simplicial complexes, and, indeed, is forced upon us as soon as motions in a graph become uncertain or adversarial. In the stochastic setting, one cannot necessarily think of a plan as defining a partial order in the manner just described. The system may cycle between states. However, as long as cycling is transient, one may again view a plan as defining a stochastic ordering based on reachable states and convergence times.

The connection between planning and simplicial complexes holds categorically as well. The complexes defined in this chapter may be viewed, homeomorphically, as the classifying spaces of various planning categories, related to the forward-chaining and backchaining planners described earlier.

1.6. Chapter Outline. Section 2 introduces nondeterministic graphs and their strategy complexes. Section 3 introduces loopback complexes, using these to characterize goal attainability. Section 4 introduces stochastic graphs and their strategy complexes. Section 5 characterizes full controllability by the existence of a certain sphere, homotopically. Section 6 introduces source and dual complexes, indicating how these are useful for design and in assessing adversarial capabilities. Section 7 shows how a strategy complex factors into a part modeling full controllability on subgraphs and a part modeling obstructions to controllability. Section 8 examines the topology of links of actions. Section 9 develops a topological test to decide whether a set of actions is essential for accomplishing a goal. Section 10 uses decision trees to reveal further structure in loopback complexes. Section 11 examines the categorical foundations of strategy complexes and source complexes. Section 12 ends the chapter with a brief discussion.

2. Nondeterministic Graphs and Strategy Complexes

2.1. Nondeterministic Graphs. We model systems with uncertainty using discrete states and discrete actions with multiple outcomes. As suggested by the earlier robotics examples, a state may encapsulate not just the configuration of the robot but also the information known to the robot at runtime, such as that provided by sensors. Actions at a state represent the control choices available to the robot. The outcomes of an action describe the various state changes possible upon execution of that action. Following [**14**], we make the following definitions:

Definition 2.1. A *nondeterministic graph* $G = (V, \mathfrak{A})$ is a set of *states* V and a collection of *(nondeterministic) actions* $\mathfrak{A}$. V is also known as G's *state space*. Each $A \in \mathfrak{A}$ consists of a *source* state v and a nonempty set T of *targets*, with $v \in V$ and $T \subseteq V$. We may write action A as $v \to T$. If T consists of a single state, A is also said to be *deterministic*. In that case, with $T = \{u\}$, we may write A more simply as $v \to u$. All graphs in this chapter are finite.

Interpretation: Action A may be executed whenever the system is at state v. When action A is executed, the system moves from state v to one of the target states in T. If T contains multiple targets, the precise target attained is not known to the system before executing A, but is known after. Different execution instances of action A could attain different target states within T. (For instance, nature might choose a different target.) In order to model worst-case behaviors, we may imagine a potentially malevolent *adversary* who chooses the precise target attained.

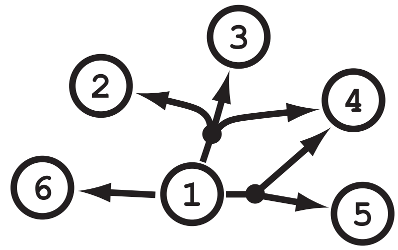

FIGURE 1. Three actions with source state 1. One action is a deterministic transition, with target state 6. The other two actions are both nondeterministic, one with targets $\{2, 3, 4\}$, the other with targets $\{4, 5\}$.

REMARKS 2.2. (1) Distinct actions may have the same source state and overlapping or identical target sets. (2) In particular, distinct actions may have the same written representation $v \to T$. (3) A nondeterministic graph in which each action is deterministic and no two actions have the same written representation is equivalent to a standard directed graph. (4) We allow the *null graph*, $(\emptyset, \emptyset)$.

EXAMPLE 2.3. In figures, we will draw an action $v \to T$ as a possibly bifurcated directed arrow from v to T. For example, Figure 1 shows three actions, each with source state 1, with different target sets. The written representations of these actions are $1 \to 6$, $1 \to \{2, 3, 4\}$, and $1 \to \{4, 5\}$.

2.2. Simplicial Complexes. Throughout this chapter we use the following definition of simplicial complex:

DEFINITION 2.4. An *(abstract) simplicial complex* Σ *with underlying vertex set* X is a collection of finite subsets of X, such that if σ is in Σ then so is every subset of σ (including the empty set $\emptyset$). The elements of Σ are *simplices*. We refer both to the elements of a simplex and to singleton simplices as *vertices*. Although traditionally many authors require that every element of X appear as a vertex in Σ, we do *not* impose this requirement. This broader definition is useful for modeling systems in which underlying states may or may not satisfy some monotone Boolean property. The set of vertices that actually appear in Σ is denoted by $\Sigma^{(0)}$, called the *zero-skeleton of* Σ. The *dimension* of a simplex is one less than its cardinality. All simplicial complexes and underlying vertex sets in this chapter are finite.

DEFINITION 2.5. The simplicial complex consisting solely of the empty simplex is the *empty complex*. The simplicial complex consisting of no simplices is the *void complex* [**27**].

REMARK 2.6. Every nonvoid finite abstract simplicial complex has a geometric realization in some finite-dimensional Euclidean space with relative topology the same as its weak/polytope topology [**37**]. Thus we may view any such complex as a topological space. The empty complex corresponds to the empty space. We also think of it as $\mathbb{S}^{-1}$, the sphere of dimension -1. The void complex does not seem to have such a nice topological interpretation, but is nonetheless convenient combinatorially. Viewing simplicial complexes as monotone Boolean functions which are TRUE for subsets of the underlying vertex set outside the complex, the void complex represents the constant function TRUE. The void complex is considered to be contractible. This is consistent with viewing the void complex as a collapse of any complex $\{\emptyset, \{v\}\}$ that represents a single point v.

DEFINITION 2.7. Suppose Σ and Γ are simplicial complexes with disjoint underlying vertex sets. The *simplicial join* [**47**] of Σ and Γ is the simplicial complex

$$\Sigma * \Gamma \;=\; \{\sigma \cup \gamma \mid \sigma \in \Sigma \text{ and } \gamma \in \Gamma\}.$$

The underlying vertex set of $\Sigma * \Gamma$ is the union of the underlying (disjoint) vertex sets of Σ and Γ.

2.3. Strategy Complexes Arising From Nondeterministic Graphs. We will define a simplicial complex for modeling the space of all plans on a nondeterministic graph. As the discussion of Section 1 suggests, we view a plan as a type of control law, that specifies what action the system should execute when it finds itself in a given state. Executing the plan should move the system from its current state to some desired state or set of states. In a nondeterministic graph, the result is a partial order on the system's state space. For this reason, we draw many of our techniques from [**8, 26, 27**].

We make one variation to the previous plan structure: We allow a plan to specify multiple possible actions at a given state. One may view multiple actions as additional permissible nondeterminism. At runtime, the system can execute any of the actions available at its current state or even leave the choice to an adversary. To emphasize this distinction from traditional plans and control laws, we generally speak of *strategies*. We capture the essence of a strategy via the following definitions:

DEFINITION 2.8. Suppose $G = (V, \mathfrak{A})$ is a nondeterministic graph and $\mathcal{A} \subseteq \mathfrak{A}$ is some set of actions. We say $\mathcal{A}$ *contains a circuit* if $\mathcal{A}$ contains a sequence of actions $v_1 \to T_1, \ldots, v_k \to T_k$, such that $v_{i+1} \in T_i$, for $i = 1, \ldots, k$, with $k \geq 1$ and $k+1$ meaning 1. We say $\mathcal{A}$ *converges* or *is convergent* if $\mathcal{A}$ does not contain a circuit.

If $\mathcal{A}$ contains a circuit, then an adversary could select action transitions to keep the system looping forever within the directed cycle $\{v_1, \ldots, v_k\}$, so we would not want to view $\mathcal{A}$ as a strategy. If $\mathcal{A}$ converges, then we may view $\mathcal{A}$ as a strategy.

DEFINITION 2.9. Suppose $G = (V, \mathfrak{A})$ is a nondeterministic graph, with $V \neq \emptyset$. The *strategy complex* of G, denoted Δ_G, is the simplicial complex with underlying vertex set $\mathfrak{A}$ whose simplices are all the convergent subsets $\mathcal{A}$ of $\mathfrak{A}$. Every simplex of Δ_G is a *(nondeterministic) strategy*. If $V = \emptyset$, we let Δ_G be the void complex.

REMARKS 2.10. (1) If V is nonempty, then Δ_G always contains the empty simplex. Intuitively, the empty simplex represents the strategy "DO NOT MOVE". (2) A nondeterministic action $v \to T$ with a self-loop, meaning $v \in T$, cannot appear in any strategy/simplex. (3) As outlined earlier, we view each strategy as a type of control law. In particular, to say that a system *executes strategy* σ means the following: Suppose the current state of the system is v. Strategy σ may contain zero, one, or several actions with source v. The system *stops* moving precisely when σ contains no action with source v. Otherwise, the system *must* execute some action $v \to T \in \sigma$. If there are several such actions, the strategy leaves open the method for choosing between those actions. From a worst-case perspective, an adversary may make the choice. Upon execution of action $v \to T$, the system finds itself at one of the targets $t \in T$. The process repeats, with t the system's new current state.

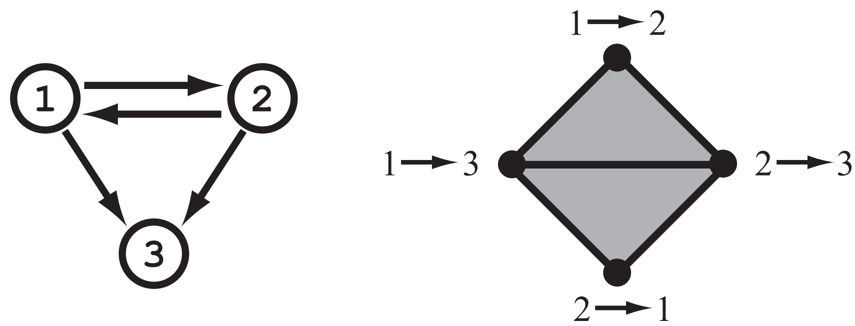

FIGURE 2. A directed graph and its associated strategy complex.

EXAMPLE 2.11. Figure 2 shows a standard directed graph along with its strategy complex. There are four directed edges in the graph, each representing a deterministic action. The strategy complex therefore could be as large as the complex generated by a tetrahedron. However, the two actions $\mathtt{1} \rightarrow \mathtt{2}$ and $\mathtt{2} \rightarrow \mathtt{1}$ together contain a circuit. Consequently, the strategy complex is actually generated by two triangles touching at an edge.

We may interpret each of the simplices in the resulting strategy complex as a strategy for accomplishing some goal, much like a traditional control law. For instance, the edge $\{\mathtt{1}\rightarrow\mathtt{2},\ \mathtt{2}\rightarrow\mathtt{3}\}$ represents the control law:

- When at state 1, execute the action $\mathtt{1}\rightarrow\mathtt{2}$.
- When at state 2, execute the action $\mathtt{2}\rightarrow\mathtt{3}$.
- Otherwise, stop moving.

Together, the two actions $\mathtt{1} \rightarrow \mathtt{2}$ and $\mathtt{2} \rightarrow \mathtt{3}$ constitute a strategy for attaining state 3 from anywhere in the graph.

The triangle $\{\mathtt{1} \rightarrow \mathtt{2},\ \mathtt{1} \rightarrow \mathtt{3},\ \mathtt{2} \rightarrow \mathtt{3}\}$ represents another strategy for attaining state 3 from anywhere in the graph. It happens to have two actions with source 1, indicating a purposeful disinterest: whichever of these two actions executes at runtime, the system will ultimately converge to state 3.

The edge common to both triangles, $\{\mathtt{1} \rightarrow \mathtt{3},\ \mathtt{2} \rightarrow \mathtt{3}\}$, represents the strategy one would obtain by traditional backchaining from state 3. This strategy moves to state 3 as directly as possible.

The edge $\{\mathtt{1}\rightarrow\mathtt{2},\ \mathtt{1}\rightarrow\mathtt{3}\}$ represents a strategy for attaining the goal set $\{\mathtt{2}, \mathtt{3}\}$. Effectively, this strategy says: "Move *away* from state 1; I do not care whereto."

EXAMPLE 2.12. In contrast, Figure 3 again shows three states, with the same *possible* transitions as in Example 2.11. However, in this example there are actually only two actions, $\mathtt{1}\rightarrow\{\mathtt{2},\mathtt{3}\}$ and $\mathtt{2}\rightarrow\{\mathtt{1},\mathtt{3}\}$, each of which is nondeterministic with two possible targets. Together, these actions contain a circuit; an adversary could force the system into an infinite loop between states 1 and 2. Consequently, the strategy complex consists of two isolated vertices, one for each action.

NOTATION 2.13. For any $m \geq -1$, $\mathbb{S}^m$ denotes the sphere of dimension m.

EXAMPLE 2.14. Figure 4 shows two strongly connected directed graphs on three states, along with their strategy complexes. The graphs are not isomorphic, but their strategy complexes are both homotopic to $\mathbb{S}^1$, the circle.

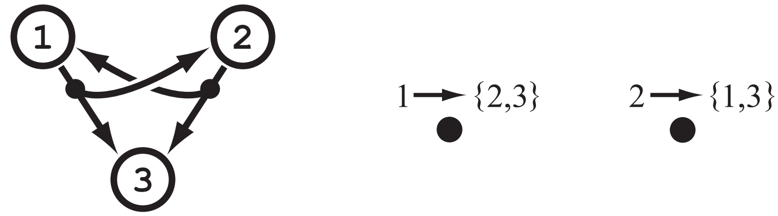

FIGURE 3. Left: A nondeterministic graph with two actions that together could produce a directed cycle. Right: The graph's strategy complex; it consists of two isolated vertices, one for each action.

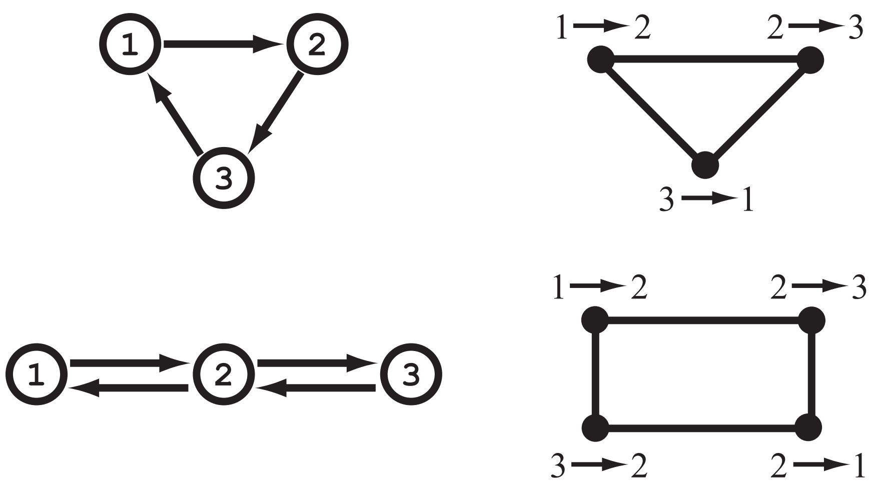

FIGURE 4. The left column shows two different strongly connected graphs on three states. The right column shows their respective strategy complexes.

Indeed, by [**26**], the strategy complex of any directed graph is homotopic either to a sphere or to a point. If the graph has n states and can be written as the disjoint union of k strongly connected subgraphs, then the strategy complex is homotopic to $\mathbb{S}^{n-k-1}$. Otherwise, the strategy complex is homotopic to a point. This result will re-appear in more general form for uncertain graphs, as Theorem 7.9.

Observe further that the graph in the lower left of Figure 4 may be viewed as the overlapping union of two strongly connected graphs on two states (the two graphs touch at state 2). Each of the subgraphs therefore has $\mathbb{S}^0$ as a strategy complex. One $\mathbb{S}^0$ is formed from the two actions $1 \rightarrow 2$ and $2 \rightarrow 1$, the other from the two actions $2 \rightarrow 3$ and $3 \rightarrow 2$. Mirroring this graph decomposition, observe that the join $\mathbb{S}^0 * \mathbb{S}^0$ is homotopic to $\mathbb{S}^1$ in general and in fact isomorphic in this case to the strategy complex of the overall graph.

Our research generalizes from directed graphs to nondeterministic (and stochastic) graphs. Doing so leads to a much larger class of strategy complexes than

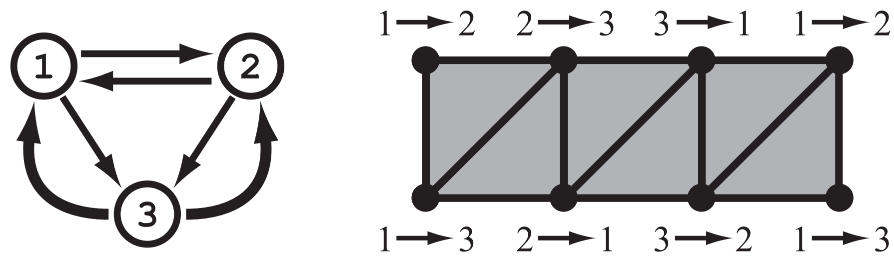

FIGURE 5. Left panel: The loopback graph $G_{\leftarrow 3}$ formed from the graph of Figure 2 by adding loopback actions at state 3 (indicated by thick arrows). Right panel: The associated loopback complex $\Delta_{G_{\leftarrow 3}}$. It is a polygonal cylinder, homotopic to $\mathbb{S}^1$.

just spheres and points, as the next theorem shows. One conclusion is that precise control is very much a special case in motion planning and that planning for uncertain systems is both topologically interesting and physically natural.

There is another observation. Directed graphs represent locally certain connectivity. Strongly connected directed graphs represent globally certain connectivity. That global property is reflected by the spherical nature of the graph's strategy complex. The connection between spheres and globally certain connectivity turns out to be significant as well for graphs in which local motions are uncertain. Much of the remainder of this chapter explores that connection.

NOTATION 2.15. Let Γ and Σ be simplicial complexes. (a) We write $\Gamma \cong \Sigma$ to mean that Γ and Σ are isomorphic, disregarding underlying vertex sets. (b) We let $\mathrm{sd}(\Sigma)$ denote the *first barycentric subdivision* of Σ. See [**44, 37, 42**].

THEOREM 2.16. *For any finite simplicial complex* Σ, *there exists a nondeterministic graph* G *such that* $\mathrm{sd}(\Sigma) \cong \Delta_G$.

PROOF. We give the basic construction and point to [**14**] for further details. Let $G = (V, \mathfrak{A})$, with V consisting of $(\mathrm{sd}(\Sigma))^{(0)}$ plus one additional state, and $\mathfrak{A}$ containing exactly one action $v \to T_v$ for each $v \in (\mathrm{sd}(\Sigma))^{(0)}$. The target set T_v consists of all states of V not adjacent to or equal to v in $\mathrm{sd}(\Sigma)$. □

3. Topological Characterization of Goal Attainability

The question of whether a nondeterministic graph contains a strategy for attaining some particular goal state may be rephrased as the problem of deciding whether the strategy complex associated with a variation of the graph is homotopic to a sphere or to a point. We first illustrate this property with two examples, then state the property as a theorem.

Consider again the graph of Figure 2. We seek to construct a topological space whose homotopy type tells us whether the graph contains a strategy for attaining state 3 from anywhere in the graph. (Of course, we know the graph contains such a strategy, by inspection, but we want a topological characterization.) Imagine that we add to the graph two deterministic actions at state 3, transitioning back to the other two states. We call these actions *loopback actions*, we call the resulting graph a *loopback graph*, and we call the associated strategy complex a *loopback complex*. See Figure 5. Observe that the loopback complex is homotopic to $\mathbb{S}^1$.

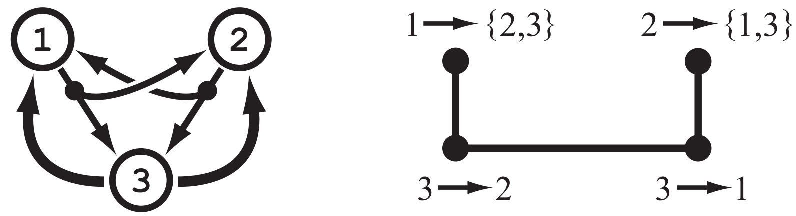

FIGURE 6. Left panel: The loopback graph $G_{\leftarrow 3}$ formed from the graph of Figure 3 by adding loopback actions at state 3. Right panel: The associated loopback complex $\Delta_{G_{\leftarrow 3}}$. It is contractible.

In contrast, suppose we ask whether the graph of Figure 3 contains a strategy for attaining state 3 from anywhere in the graph. (Again, we know it does not, by inspection, but again we seek a topological characterization.) As before, we add loopback actions at state 3 and compute the associated loopback complex, as shown in Figure 6. This time the loopback complex is homotopic to a point.

It turns out that loopback complexes are always homotopic either to a point or to $\mathbb{S}^{n-2}$, with n being the number of states in the graph. Moreover, the complex is a sphere precisely when the graph contains a strategy for moving from anywhere in the graph to the state at which we have added loopback actions. We now make this statement precise with the following definitions and theorem.

DEFINITION 3.1. Let $G = (V, \mathfrak{A})$ be a nondeterministic graph and s a desired *stop state* in V. A *complete strategy for attaining s in G* is a convergent set $\sigma \subseteq \mathfrak{A}$, such that σ contains at least one action with source v for every v in $V \setminus \{s\}$. If such a σ exists, we also say *G contains a complete strategy for attaining s.* (Later, we will use the same terminology for stochastic graphs.)

REMARK 3.2. In the previous definition, σ cannot contain any action with source s, as otherwise σ would not be convergent.

DEFINITION 3.3. With notation as above, define $G_{\leftarrow s}$ to be the nondeterministic graph constructed from G by first removing all actions of G that have source s, then adding all possible *loopback actions* at state s, that is, all deterministic actions $s \to v$, with $v \in V \setminus \{s\}$. Call $G_{\leftarrow s}$ the *loopback graph formed from G and s.* Define $\Delta_{G_{\leftarrow s}}$ to be the strategy complex associated with $G_{\leftarrow s}$ and call it the *loopback complex formed from G and s.* (Later, we will use the same terminology for stochastic graphs.)

REMARK 3.4. In the previous definition, we could simply have *added* loopback actions at s, without first *removing* any existing actions. There would be no difference in the homotopy type of the resulting complex, as Example 3.10 will show.

NOTATION 3.5. The rest of this chapter employs the following notation:

- $\mathbf{x} \in \mathbb{R}^n$ means the point $(x_1, \ldots, x_n)$ in n-dimensional Euclidean space.
- $X \simeq Y$ means that X and Y are homotopic as topological spaces (this is the same as saying that X and Y have the same homotopy type). The notation makes sense for simplicial complexes by Remark 2.6. Later, by Section 8.2, the notation will make sense for partially ordered sets.

- $\diagup^n_=$ denotes the *diagonal* in $\mathbb{R}^n$ and $\odot^n_{\neq}$ denotes its complement. So $\diagup^n_= \;= \{\mathbf{x} \in \mathbb{R}^n \mid x_1 = \cdots = x_n\}$, $\odot^n_{\neq} = \{\mathbf{x} \in \mathbb{R}^n \mid x_i \neq x_j,\text{ for some } i, j\}$.

Theorem 3.6. *Let G be a nondeterministic graph with state space V, $s \in V$, and $n = |V|$. If G contains a complete strategy for attaining s, then $\Delta_{G_{\leftarrow s}} \simeq \mathbb{S}^{n-2}$. Otherwise, $\Delta_{G_{\leftarrow s}}$ is contractible.*

Proof. This proof is motivated by the techniques in [**8, 26**]. The proof given here appears in similar though not quite identical form in [**14**].

If $n = 1$, then $\Delta_{G_{\leftarrow s}}$ is necessarily the empty complex, which we view as $\mathbb{S}^{-1}$. The empty simplex is a complete strategy for attaining the only state there is in the graph. This shows the theorem holds for $n = 1$. So we may assume that $V = \{1, \ldots, n\}$, with $s = n > 1$.

I. Suppose σ is a complete strategy for attaining s in G. Let $\mathfrak{A}$ be the actions of $G_{\leftarrow s}$. For each action $A \in \mathfrak{A}$, with $A = i \rightarrow T$, define the following open polyhedral cone, which we refer to as a *nondeterministic covering set*:

$$U_A \;=\; \left\{ \mathbf{x} \in \mathbb{R}^n \;\middle|\; x_i > \max_{j \in T} x_j \right\}.$$

(Looking ahead, nondeterministic covering sets constitute a special case of the covering sets to appear in Def. 5.2.)

The nerve of the cover $\{U_A\}_{A \in \mathfrak{A}}$ conveys information. In particular, a set of actions $\{A_1, \ldots, A_k\}$ is convergent if and only if $U_{A_1} \cap \cdots \cap U_{A_k}$ is not empty. When nonempty, the intersection is contractible. By the Nerve Lemma [**7, 25**], $\Delta_{G_{\leftarrow s}}$ therefore has the homotopy type of $\bigcup_{A \in \mathfrak{A}} U_A$. We will show that $\bigcup_{A \in \mathfrak{A}} U_A = \odot^n_{\neq}$. Consequently, $\Delta_{G_{\leftarrow s}}$ is homotopic to $\mathbb{S}^{n-2}$.

If $\mathbf{x}$ is a point on the diagonal $\diagup^n_=$, then $\mathbf{x}$ cannot lie in any nondeterministic covering set U_A, by construction. If $\mathbf{x} \in \mathbb{R}^n$ with $x_n > x_i$ for some index i, then $\mathbf{x}$ lies in the nondeterministic covering set associated with the loopback action $n \rightarrow i$.

Otherwise, $\mathbf{x} \in \mathbb{R}^n$ with $x_i > x_n$ for some index i. Suppose that $\mathbf{x}$ does not lie in any U_A. Since σ is a complete strategy for attaining n, σ contains an action B with source i. Write this action as $i \rightarrow T$, for some set of targets T. Action B is an element of $\mathfrak{A}$, so $\mathbf{x}$ does not lie in U_B. Consequently, for some target $j \in T$, $x_i \leq x_j$, implying that $x_j > x_n$. One may now repeat the argument with index j. Continuing in this manner, one obtains an unbounded sequence of actions $v_1 \rightarrow T_1, v_2 \rightarrow T_2, \ldots$, all in σ, such that $v_{k+1} \in T_k$ for all $k = 1, 2, \ldots$. Since G is finite, this means σ contains a circuit, establishing a contradiction.

II. Suppose G contains no complete strategy for attaining s. Let Σ be the subcomplex of $\Delta_{G_{\leftarrow s}}$ consisting of all simplices that contain no loopback actions and let σ_0 be a maximal simplex of Σ. Suppose $\sigma_0 \neq \emptyset$ and consider the collection (not a simplicial complex) $\Gamma = \{\gamma \in \Delta_{G_{\leftarrow s}} \mid \sigma_0 \subseteq \gamma\}$.

Since σ_0 is not a complete strategy for attaining $s = n$, there must be some state $i \neq n$ such that σ_0 contains no action with source i. Suppose γ is in Γ and does not contain the loopback action $n \rightarrow i$. By maximality of σ_0 in Σ, γ also does not contain an action with source i. Consequently, $\gamma \cup \{n \rightarrow i\}$ is convergent and thus in Γ. Suppose γ is in Γ and does already contain the loopback $n \rightarrow i$. Then removing that loopback produces as well an element of Γ. By Lemma 7.6 of [**8**], this

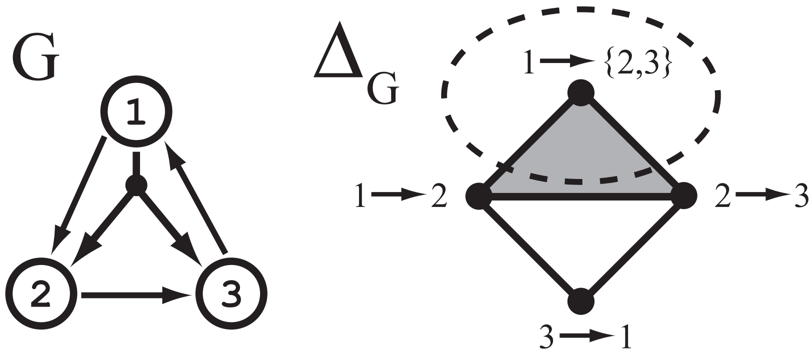

FIGURE 7. Action $1 \to 2$ is more precise than action $1 \to \{2, 3\}$, in G. Consequently, all simplices containing the action $1 \to \{2, 3\}$ may collapse away without changing the homotopy type of Δ_G.

means the complex $\Delta_{G_{\leftarrow s}}$ collapses to the complex $\Delta_{G_{\leftarrow s}} \setminus \Gamma$, preserving homotopy type. Repeating this process, one may collapse away all nonempty simplices of Σ along with their supersets in $\Delta_{G_{\leftarrow s}}$, leaving only the loopback actions. All the loopback actions together are convergent. So, we have shown how to collapse the complex $\Delta_{G_{\leftarrow s}}$ to a single nonempty simplex, which in turn collapses to a point. $\square$

REMARKS 3.7. (1) The contradiction argument in part I of the proof is much like planning, now from an adversary's perspective. (2) Part II of the proof actually establishes that $\Delta_{G_{\leftarrow s}}$ is collapsible when G fails to contain a complete strategy for attaining s. Later, Corollary 10.7 will establish the yet stronger property that $\Delta_{G_{\leftarrow s}}$ is nonevasive. (3) Allowing $\sigma_0 = \emptyset$ in part II would be fine though less explicit.

COROLLARY 3.8. *With notation as above, G contains a complete strategy for attaining s if and only if $\Delta_{G_{\leftarrow s}}$ contains an odd number of simplices (counting $\emptyset$).*

PROOF. The reduced Euler characteristic is 0 for points and ± 1 for spheres. $\square$

COROLLARY 3.9. *Let G be a nondeterministic graph with state space V. Let $s \in V$. The number of complete strategies for attaining s in G is either zero or odd.*

EXAMPLE 3.10. Understanding the homotopy types of strategy complexes in general is an open question. Globally, we have seen the significance of spheres of a certain dimension. Locally, homotopy collapse ignores imprecise actions in favor of more precise actions, as follows: Whenever a nondeterministic graph contains two actions with the same source state and comparable target sets, then all simplices containing the less precise action may collapse away. Imagine actions $v \to T$ and $v \to S$ with $S \subseteq T$. If $v \to T \in \sigma \in \Delta_G$, then $\sigma \cup \{v \to S\} \in \Delta_G$. So all simplices containing $v \to T$ may collapse away without changing homotopy type. See Figure 7.

4. Stochastic Graphs and Strategy Complexes

EXAMPLE 4.1. Consider again the graph of Figure 3. Now suppose that the transitions of each action are not so uncertain as to be nondeterministic, but instead have associated probabilities, as indicated in Figure 8. The probabilities mean that

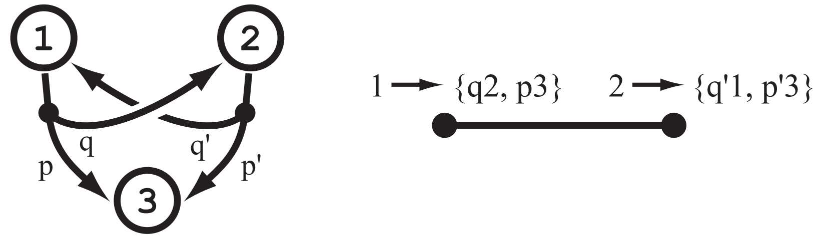

FIGURE 8. The graph on the left contains two stochastic actions. Although the actions together may cause the system to cycle between states 1 and 2, the cycling will be transient. As a result, the graph's strategy complex on the right contains not only the vertices representing the individual stochastic actions, but the edge between them. Compare with Figure 3.

during each execution of an action, a particular transition occurs with the indicated probability, independent of the past. Although such uncertainty in the actions might cause the system to cycle between states 1 and 2 for a while, the probability that the system would cycle forever is zero. With every action execution there is some minimum nonzero probability that the system will exit the cycle and move to state 3. Consequently, we should consider the set consisting of both actions to be a simplex in the graph's strategy complex, as shown in Figure 8.

In order to make this intuition precise, we need to generalize the notion of circuit given in Def. 2.8 from the nondeterministic setting to the stochastic setting. The earlier definition models an adversary who selects action transitions in such a way that the system finds itself stuck in some set of states, moving endlessly between those states. In the nondeterministic setting, the adversary can choose the transitions so as to create a cyclic path, but that is almost incidental; the key idea is that the system is stuck in a subspace. That idea generalizes readily to the stochastic setting: instead of a cyclic path, one obtains a *recurrent class* [**17, 28**] in a Markov chain created by the adversary. We formalize these concepts as follows:

DEFINITION 4.2. A *stochastic action* A consists of a *source* state v and a nonempty set T of *target* states, along with a strictly positive probability distribution $p : T \to (0, 1]$. We may write action A as $v \to pT$. If $T = \{u_1, \ldots, u_k\}$ and $p_i = p(u_i)$, for $i = 1, \ldots, k$, then we may also write A as $v \to \{\, p_1 u_1, \ldots, p_k u_k\}$. In this representation, $p_i > 0$, for $i = 1, \ldots, k$, and $\sum_{i=1}^{k} p_i = 1$. If T consists of a single state, then A is deterministic. In that case, with $T = \{u\}$, we may write A in several different ways, including $v \to u$, $v \to \{1u\}$, and $v \to \{u\}$.

INTERPRETATION: As in the nondeterministic case, action A may be executed whenever the system is at state v. When action A is executed, the system moves from state v to one of the targets u in T, selected from all the targets with probability $p(u)$. This process is Markovian, that is, independent of how or when the system arrived at state v. If T contains multiple targets, the precise target attained is not known to the system before executing A, but is known after.

DEFINITION 4.3. A *stochastic graph* $G = (V, \mathfrak{A})$ is a set of *states* V and a collection of *actions* $\mathfrak{A}$, such that each action of $\mathfrak{A}$ is either nondeterministic or stochastic, with its source and targets in V. We may refer to V as G's *state space*.

(**NB:** A stochastic graph may contain *both* nondeterministic and stochastic actions. Of course, this also includes the special case of deterministic actions.)

In the remaining definitions of this section, let $G = (V, \mathfrak{A})$ be a stochastic graph.

DEFINITION 4.4. Define $\mathrm{src} : \mathfrak{A} \to V$ so $\mathrm{src}(A)$ is the source of A. Extend to sets of actions. If $\mathcal{A} \subseteq \mathfrak{A}$, we say $\mathrm{src}(\mathcal{A})$ is the *start region* of $\mathcal{A}$.

DEFINITION 4.5. Let $W \subseteq V$ and $A \in \mathfrak{A}$. Action A *moves off* W if $\mathrm{src}(A) \in W$ and one of the following is true: (i) A is stochastic with at least one of its targets in $V \setminus W$, or (ii) A is nondeterministic with all of its targets in $V \setminus W$.

DEFINITION 4.6. Let $\mathcal{A} \subseteq \mathfrak{A}$ be some set of actions in G. We say $\mathcal{A}$ *contains a stochastic circuit* if, for some nonempty subset $\mathcal{B}$ of $\mathcal{A}$, no action of $\mathcal{B}$ moves off $\mathrm{src}(\mathcal{B})$. We say $\mathcal{A}$ *converges stochastically* or *is stochastically convergent* if $\mathcal{A}$ does not contain a stochastic circuit.

REMARK 4.7. Suppose $\mathcal{A}$ contains a stochastic circuit. Then an adversary could select some nonempty subset of actions $\mathcal{B}$ in $\mathcal{A}$ and some nonempty subset of states W in V, such that: (i) $W = \mathrm{src}(\mathcal{B})$, (ii) $\mathcal{B}$ contains exactly one action with source w for every $w \in W$, and (iii) no B in $\mathcal{B}$ moves off W. Now consider an action B in $\mathcal{B}$. If B is stochastic, then every target of B lies in W. If B is nondeterministic, then the adversary could further select one target of B lying in W. The complete selection process just described amounts to the construction of a Markov chain on state space W. Since the chain is finite, it must contain a recurrent class [**17, 28**]. This means that there is a nonempty subset R of W such that the probability of the chain eventually moving from any given state of R to any other state of R is 1, while the probability of ever leaving R is 0. Restricting the Markov chain from W to R defines a new Markov chain. This new chain has state space R and is irreducible; it is the stochastic analogue of any irreducible directed cycle appearing via Def. 2.8 for the nondeterministic setting. (With this generalization in mind, we usually omit the explicit "stochastic" designation in the terms of Def. 4.6.)

We may now define a strategy complex much as we did earlier for nondeterministic graphs, but now allowing both stochastic and nondeterministic actions:

DEFINITION 4.8. If $V \neq \emptyset$, the *strategy complex* Δ_G of G is the simplicial complex whose underlying vertex set is $\mathfrak{A}$ and whose simplices are all the stochastically convergent subsets $\mathcal{A}$ of $\mathfrak{A}$. Every simplex of Δ_G is a *(stochastic) strategy.* If $V = \emptyset$, we let Δ_G be the void complex.

REMARKS 4.9. (1) Remarks 2.10 carry over to the stochastic setting with small changes. For instance, a stochastic action $v \to_p T$ with a self-loop may appear in some simplices, so long as $|T| > 1$. (2) Theorem 3.6 holds as well in the stochastic setting. The proof carries over with small changes. See for example the covering sets of Section 5. The decision trees of Section 10 will provide yet a different type of proof.

REMARK 4.10. The Markov chain construction of Remark 4.7 suggests an alternate definition of stochastic convergence. Intuitively, no matter what Markov chain the adversary constructs, all states at which the strategy specifies actions should be transient states of the Markov chain, not recurrent states. The actions of a strategy should eventually move the system to some set of states at which the strategy specifies no actions. The next two definitions and lemma rephrase this intuition algebraically.

DEFINITION 4.11. Define the "adversity" function $\mathrm{adv} : \mathfrak{A} \times \mathbb{R}^{|V|} \to \mathbb{R}$ by

$$\mathrm{adv}(A, \{x_v\}_{v \in V}) \;=\; \begin{cases} \displaystyle\sum_{u \in T} p(u) x_u, & \text{if } A \;=\; w \to p\,T; \\ \displaystyle\max_{u \in T} x_u, & \text{if } A \;=\; w \to T; \end{cases} \qquad \text{(for some } w\text{)}.$$

INTERPRETATION: Given a family $\{x_v\}_{v \in V}$ of real numbers indexed by V and given an action A, the expression $\mathrm{adv}(A, \{x_v\}_{v \in V})$ computes an expectation based on A's targets when A is stochastic and a maximization based on A's targets when A is nondeterministic. This definition will help us model expected outcomes arising from worst-case adversarial choices. Notation: We frequently drop the index set V to write $\mathrm{adv}(A, \{x_v\})$, or use vector notation to write $\mathrm{adv}(A, \mathbf{x})$.

DEFINITION 4.12. If $\mathcal{A}$ is a collection of actions and w a state, let $\mathcal{A}|w$ denote all actions in $\mathcal{A}$ that have source w.

LEMMA 4.13. *Let $G = (V, \mathfrak{A})$ be a stochastic graph with $V \neq \emptyset$ and let $\mathcal{A} \subseteq \mathfrak{A}$. Then $\mathcal{A}$ is stochastically convergent if and only if the following system of equations in the real variables $\{x_v\}_{v \in V}$ has a unique finite solution, identically zero:*

$$x_w = \max_{A \in \mathcal{A}|w} \mathrm{adv}(A, \{x_v\}), \quad \textit{for all } w \in V. \tag{4.1}$$

(We take any maximization over the empty set to be 0.*)*

PROOF. Follows by standard techniques for Markov chains [**17, 28**]. □

REMARK 4.14. Suppose $V \neq \emptyset$ and $\mathcal{A}$ is stochastically convergent. Then $\mathrm{src}(\mathcal{A})$ must be a proper subset of V, so System (4.1) has at least one explicit equation of the form $x_w = 0$. We may view strategy $\mathcal{A}$ as attaining the goal set consisting of all states w for which (4.1) contains the explicit equation $x_w = 0$, that is, all states in $V \setminus \mathrm{src}(\mathcal{A})$.

REMARK 4.15. Now imagine that we associate to each action A of $\mathfrak{A}$ a nonnegative *action transition time* δ_A. Let $\mathcal{A} \subseteq \mathfrak{A}$. Then the following system of equations in the real variables $\{t_v\}_{v \in V}$ again has a unique finite solution if and only if $\mathcal{A}$ is stochastically convergent, in which case each of the t_v is nonnegative:

$$t_w = \max_{A \in \mathcal{A}|w} (\mathrm{adv}(A, \{t_v\}) + \delta_A), \quad \text{for all } w \in V. \tag{4.2}$$

We may interpret the unique finite solution $\{t_v\}_{v \in V}$, when it exists, as *worst-case expected convergence times*. Intuitively, an adversary can choose actions and nondeterministic transitions from $\mathcal{A}$ in such a way that the expected time for the system to enter the set of states $V \setminus \mathrm{src}(\mathcal{A})$, when started at state w, can be as great as t_w.

DEFINITION 4.16. Given a stochastically convergent set of actions $\mathcal{A}$, let $t_{\max}(\mathcal{A})$ be the maximum t_w obtained as a solution to System (4.2). For any nonnegative $\mathcal{T}$, let $\Delta_G^{\mathcal{T}}$ be the subcomplex of Δ_G consisting of all simplices σ for which $t_{\max}(\sigma) \leq \mathcal{T}$.

5. Topological Characterization of Full Controllability

Definition 5.1. (a) Let $G = (V, \mathfrak{A})$ be a stochastic graph and suppose I and S are nonempty subsets of V. We say that a simplex σ of Δ_G is a *stochastic strategy for attaining S from I* if, with probability 1, the system eventually stops at *some* state of S whenever it starts at *any* state of I and executes strategy σ. In the case of singleton sets $I = \{i\}$ and $S = \{s\}$, we may simply say that σ is a *stochastic strategy for attaining s from i.*

(b) A nonempty set of states S is *certainly attainable (in G)* if there is some stochastic strategy $\sigma \in \Delta_G$ for attaining S from all of V.

(c) A graph G is *fully controllable* if, for any initial state i and any stop state s in G's state space, G contains a stochastic strategy for attaining s from i.

This section characterizes full controllability by the condition that Δ_G be homotopic to $\mathbb{S}^{n-2}$. This condition and its proof are a generalization of Theorem 3.6 and its proof. We will employ a generalization of nondeterministic covering sets and show how transitivity of actions translates to unions of open sets even in the stochastic setting. Later (see Remark 6.7) we will see the basis for a more combinatorial proof.

In order to simplify the discussion, throughout the rest of this section we assume that G is a stochastic graph with actions $\mathfrak{A}$ and state space $V = \{1, \dots, n\}$, $n \geq 1$.

Definition 5.2. Let $\{\delta_A\}_{A \in \mathfrak{A}}$ be nonnegative action transition times, associated to the actions of G. For each action A in $\mathfrak{A}$, define the *covering set* of A to be the following open subset of $\mathbb{R}^n$:

$$U_{A,\delta_A} \;=\; \{\mathbf{x} \in \mathbb{R}^n \mid x_i > \mathrm{adv}(A, \mathbf{x}) + \delta_A\}, \quad \text{with } i = \mathrm{src}(A).$$

For the special case in which $\delta_A = 0$, we may write U_A in place of U_{A,δ_A}. This notation is consistent with the notation for nondeterministic covering sets that appeared earlier in the proof of Theorem 3.6.

Remark 5.3. Suppose we write stochastic action A as $i \to_p T$, with $\emptyset \neq T \subseteq V$, as per Def. 4.2. States are now integers. Then, letting $p_j = p(j)$ for $j \in T$,

$$U_{A,\delta_A} \;=\; \Big\{\mathbf{x} \in \mathbb{R}^n \;\Big|\; x_i > \sum_{j \in T} p_j x_j + \delta_A\Big\}.$$

Similarly, if we write nondeterministic action A as $i \to T$, then

$$U_{A,\delta_A} \;=\; \Big\{\mathbf{x} \in \mathbb{R}^n \;\Big|\; x_i > \max_{j \in T} x_j + \delta_A\Big\}.$$

In particular, for stochastic A with $\delta_A = 0$, the set U_A is an open homogeneous halfspace, whose defining hyperplane normal is determined by A's transition probabilities. This hyperplane includes the diagonal $\not\equiv^n$. For nondeterministic A with $\delta_A = 0$, the set U_A is the intersection of several such open homogeneous halfspaces, as specified by A's possible source-to-target transitions, just as in the proof of Theorem 3.6.

Lemma 5.4. *Suppose G contains a stochastic strategy σ for attaining state k from state ℓ. Let t_ℓ be the worst-case expected convergence time starting from state ℓ, obtained by solving System* (4.2) *with $\mathcal{A} = \sigma$. Then*

$$\{\mathbf{x} \in \mathbb{R}^n \mid x_\ell > x_k + t_\ell\} \;\subseteq\; \bigcup_{A \in \mathfrak{A}} U_{A,\delta_A}.$$

PROOF. We can assume without loss of generality that $k = n$ and that σ is a complete strategy for attaining n. If $\ell = n$, then there is nothing to prove, since the set on the left is empty, so we may assume that $1 \leq \ell < n$.

Let $\{t_i\}_{i=1}^n$ be the solution to System (4.2) with $\mathcal{A} = \sigma$.

Suppose there is some $\mathbf{x}^* \in \mathbb{R}^n$ such that $x^*_\ell > x^*_n + t_\ell$, but $\mathbf{x}^*$ lies in no U_{A,δ_A}. Since the sets U_{A,δ_A} are invariant with respect to translation along the diagonal $\nearrow^n_=$, we may assume that $x^*_n = 0$.

When System (4.2) has a unique finite solution, one can obtain that solution by iteration, that is by computing

$$t_w^{(m+1)} = \max_{A \in \mathcal{A}|w} (\mathrm{adv}(A, \{t_v^{(m)}\}) + \delta_A), \quad \text{for all } w \in V, \text{ for } m = 0, 1, 2, \ldots,$$

starting from any finite initial seed values for $\{t_v^{(0)}\}_{v \in V}$. As $m \to \infty$, the iteration will converge to the solution of (4.2).

In the present case, if we set $t_i^{(0)} = x_i^*$, then in the limit, by induction on m, we see that $t_i \geq x_i^*$, for all $1 \leq i \leq n$. Key in the induction is the observation that, for each action A of σ with source i, we have the inequality $\mathrm{adv}(A, \mathbf{x}^*) + \delta_A \geq x_i^*$, by the contrary assumption on $\mathbf{x}^*$. Thus $t_\ell \geq x^*_\ell > x^*_n + t_\ell = t_\ell$, a contradiction. □

LEMMA 5.5. *A nonempty set of actions $\mathcal{A}$ is stochastically convergent if and only if*

$$\bigcap_{A \in \mathcal{A}} U_A \neq \emptyset.$$

PROOF. I. Suppose $\bigcap_{A \in \mathcal{A}} U_A \neq \emptyset$, and choose $\mathbf{x}^*$ to lie in the intersection.

For each $A \in \mathcal{A}$, define $\delta_A = x^*_{\mathrm{src}(A)} - \mathrm{adv}(A, \mathbf{x}^*)$. By definition of U_A, $\delta_A > 0$. Writing $\mathbf{t}$ for $\{t_j\}_{j=1}^n$, consider the following variant of System (4.2):

$$t_i = \max_{A \in \mathcal{A}|i} (\mathrm{adv}(A, \mathbf{t}) + \delta_A) + b_i, \quad \text{for all } i \in V, \tag{5.1}$$

with $b_i = 0$ for all $i \in \mathrm{src}(\mathcal{A})$ and $b_i = x_i^*$ for all $i \notin \mathrm{src}(\mathcal{A})$.

By construction of the $\{\delta_A\}_{A \in \mathcal{A}}$, this system has at least one finite solution, given by $t_i = x_i^*$ for all i in V.

If $\mathcal{A}$ contains a stochastic circuit, then by Remark 4.7, we can construct from $\mathcal{A}$ an irreducible Markov chain on some nonempty subset R of V. Without loss of generality, $R = \{1, \ldots, k\}$, for some $1 \leq k \leq n$. Let (p_{ij}) denote the stochastic matrix of this Markov chain. Combining with System (5.1), we obtain

$$x_i^* \geq \sum_{j=1}^{k} p_{ij} x_j^* + \delta_i, \quad \text{i} = 1, \ldots, \text{k},$$

where $\delta_i = \delta_A$, A being the particular action used to construct the transition(s) at state i as per Remark 4.7. Since (p_{ij}) is a stochastic matrix, this is only possible if $\delta_i \leq 0$ for at least one i, establishing a contradiction.

II. Suppose $\mathcal{A}$ is stochastically convergent. For each $A \in \mathcal{A}$, let $\delta_A = 1$. System (4.2) has a unique finite solution, call it $\mathbf{t}^*$. Pick some arbitrary $B \in \mathcal{A}$. Suppose B has source i. Then

$$t_i^* \; = \; \max_{A \in \mathcal{A}|i} (\mathrm{adv}(A, \mathbf{t}^*) + 1) \; \geq \; \mathrm{adv}(B, \mathbf{t}^*) + 1 \; > \; \mathrm{adv}(B, \mathbf{t}^*),$$

implying that $\mathbf{t}^* \in U_B$. So the intersection of all the U_A, with $A \in \mathcal{A}$, contains $\mathbf{t}^*$ and thus is not empty. □

Corollary 5.6. *Let $G = (V, \mathfrak{A})$ be a stochastic graph, with $V \neq \emptyset$. Then*

$$\Delta_G \simeq \bigcup_{A \in \mathfrak{A}} U_A.$$

Proof. If V consists of a single state, then Δ_G is the empty complex, which corresponds to the empty space. Any covering set is empty as well. So the corollary holds.

If V contains multiple states, then each covering set U_A is open and convex, so intersections of such covering sets are contractible when nonempty. Together, the Nerve Lemma and Lemma 5.5 establish the corollary. □

Theorem 5.7. *Let $G = (V, \mathfrak{A})$ be a stochastic graph, with $V \neq \emptyset$. G is fully controllable if and only if $\Delta_G \simeq \mathbb{S}^{n-2}$, with $n = |V|$.*

Proof. If V consists of a single state, then G is fully controllable and Δ_G is $\mathbb{S}^{-1}$, so the theorem holds. So assume that $n > 1$ in what follows.

I. Suppose G is fully controllable. By Lemma 5.4, with all $\delta_A = 0$, the union of all the covering sets U_A is $\bigodot^n_{\neq}$. Combining with Corollary 5.6, we see that

$$\Delta_G \;\simeq\; \bigcup_{A \in \mathfrak{A}} U_A \;=\; \bigodot^n_{\neq} \;\simeq\; \mathbb{S}^{n-2}.$$

II. Suppose $\Delta_G \simeq \mathbb{S}^{n-2}$. If G is not fully controllable, then there must be some state $s \in V$ such that G does not contain a complete strategy for attaining s. Let $G_{+s} = (V, \mathfrak{A}^+)$ be the graph obtained from G by *adding* all possible loopbacks at s, that is, all actions $s \to v$, with $v \in V \setminus \{s\}$. We see

$$\mathbb{S}^{n-2} \;\simeq\; \Delta_G \;\simeq\; \bigcup_{A \in \mathfrak{A}} U_A \;\subseteq\; \bigcup_{A \in \mathfrak{A}^+} U_A \;\simeq\; \Delta_{G_{+s}}.$$

Since each covering set U_A is homogeneous and invariant with respect to translation along the diagonal $\not\nearrow^n_{=}$, and since no proper subset of $\mathbb{S}^{n-2}$ is homotopic to $\mathbb{S}^{n-2}$, it must be that the covering sets U_A arising from G cover all of $\bigodot^n_{\neq}$. The subset relation above therefore implies that $\Delta_{G_{+s}} \simeq \mathbb{S}^{n-2}$. On the other hand, a collapsibility argument nearly identical to that appearing in the proof of Theorem 3.6 shows that $\Delta_{G_{+s}}$ is contractible, establishing a contradiction. □

Remark 5.8. A similar result follows from Lemma 5.4 for time-bounded strategies: For any $\mathcal{T} \geq 0$, G is fully controllable using only strategies whose worst-case expected convergence times are bounded by $\mathcal{T}$ if and only if $\Delta^{\mathcal{T}}_G \simeq \mathbb{S}^{n-2}$. See [**14**].

Example 5.9. Figure 9 shows a three-state graph in which every action is uncertain: two actions are stochastic, one is nondeterministic. The strategy complex is homotopic to $\mathbb{S}^1$, suggesting the graph is fully controllable. Indeed, for every state there is a strategy for moving to any other state. For instance, the actions at states `1` and `3` together will with certainty move the system to state `2`. What is uncertain is the precise time this will take (one may of course compute a worst-case expected convergence time) and the precise route taken. The system may move directly to state `2` or it may cycle for a while between states `1` and `3`. We see in this example how strategy complexes and the topological characterization of Theorem 5.7 have abstracted away detailed trajectory information, while preserving a description of

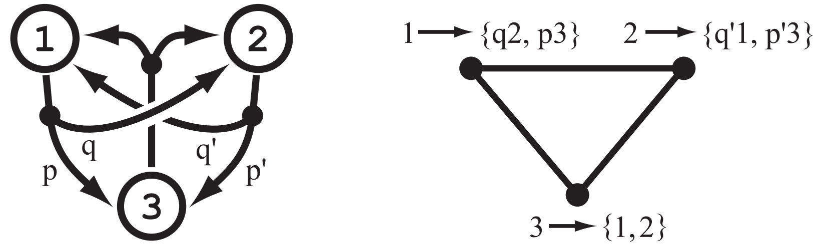

FIGURE 9. The graph on the left has a stochastic action at each of states `1` and `2` and a nondeterministic action at state `3`. Its strategy complex on the right is the boundary of a triangle.

the system's overall capabilities. — We mention in passing that changing either or both of the stochastic actions to nondeterministic actions would dissipate full controllability. For instance, if all three actions were nondeterministic, then the strategy complex would consist of three isolated vertices.

6. Source and Dual Complexes

Strategy complexes model control laws for accomplishing tasks specified by goal states in nondeterministic or stochastic graphs. Each simplex (aka strategy) is a collection of actions. The source of an action describes the conditions under which the action is applicable (modeled as a state in the graph). The targets of the action describe the possible outcomes (again modeled as states in the graph). The emphasis in strategies is on actions, yet often one cares primarily about the high-level capability of accomplishing some task, that is, moving from some set of initial states to some set of final states. This section shows how to compress the strategy complex into a smaller complex, called the *source complex*, modeling the start regions of all strategies available to the system. Moreover, this compression preserves homotopy type. Construction of the source complex leads very naturally to a dual complex that models the potentially unattainable goals. These two complexes provide a basis for analyzing and designing systems with control uncertainty.

6.1. Modeling System Capabilities. Let G be a stochastic graph with state space V. We may view any $\sigma \in \Delta_G$ as a strategy for attaining the goal set $V \setminus \mathrm{src}(\sigma)$, as suggested by Remarks 2.10(3) and 4.14: If $v \in \mathrm{src}(\sigma)$, then there is at least one action in σ with source v, possibly several. The system must execute one such action when it is at state v. If $v \notin \mathrm{src}(\sigma)$, then the system does not move when it is at state v. In short, the system moves so long as its current state lies in $\mathrm{src}(\sigma)$, and stops otherwise. Since σ is convergent, with probability 1 the system will eventually find itself in $V \setminus \mathrm{src}(\sigma)$. (Of course, in some instances one may be able to make a more precise prediction as to where the system will stop, but from a global perspective, the outcome $V \setminus \mathrm{src}(\sigma)$ is a general bound.) We model this abstraction with the following definition:

DEFINITION 6.1. The *source complex* $\overline{\Delta}_G$ of a stochastic graph $G = (V, \mathfrak{A})$ is the simplicial complex whose underlying vertex set is V and whose simplices are the start regions of all strategies in G:

$$\overline{\Delta}_G \;=\; \{\mathrm{src}(\sigma) \mid \sigma \in \Delta_G\}.$$

LEMMA 6.2. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and let W be a nonempty simplex of $\overline{\Delta}_G$. Then the following subcomplex of Δ_G is contractible:*

$$\Sigma_W \;=\; \{\tau \in \Delta_G \mid \mathrm{src}(\tau) \subseteq W\}.$$

PROOF. Let $\sigma \in \Delta_G$ such that $\mathrm{src}(\sigma) = W$. Since σ is convergent, there is some action $A \in \sigma$ that moves off W. Now let $\tau \in \Sigma_W$ and suppose $A \notin \tau$. If $\tau \cup \{A\}$ were to contain a circuit, then A could not move off W. So we see that $\tau \cup \{A\} \in \Sigma_W$, establishing that Σ_W is a cone with apex A. □

LEMMA 6.3. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and suppose W is a nonempty subset of V such that every proper subset of W is a simplex of $\overline{\Delta}_G$. Then $W \in \overline{\Delta}_G$ if and only if some action of G moves off W.*

PROOF. I. Suppose $W \in \overline{\Delta}_G$. Some action of G moves off W, by the first part of the proof of Lemma 6.2.

II. Let A be an action of G that moves off W and let $w = \mathrm{src}(A)$. By assumption, there exists $\tau \in \Delta_G$ such that $\mathrm{src}(\tau) = W \setminus \{w\}$. Arguing as in the last part of the proof of Lemma 6.2, we see that $\tau \cup \{A\} \in \Delta_G$, establishing $W \in \overline{\Delta}_G$. □

REMARKS 6.4. (1) Lemma 6.3 is backchaining topologized, abstracting a planning method known as DYNAMIC PROGRAMMING [**4**]. The connection to backchaining will appear more explicitly in the proof of Lemma 8.6.

(2) One could modify the hypotheses of the lemma, explicitly requiring merely that subsets of size $|W| - 1$ lie in $\overline{\Delta}_G$, rather than all proper subsets of W. The formulation given emphasizes the intuition that a goal is certainly attainable precisely when every subspace of the goal's complement has an exit.

THEOREM 6.5. *For any stochastic graph G, $\Delta_G \simeq \overline{\Delta}_G$.*

PROOF. If G is null, both complexes are void. Otherwise, the theorem follows from Lemma 6.2 and the Quillen Fiber Lemma [**41, 7, 8, 47**]. It is also a corollary to upcoming Theorem 8.10, so we omit further details here. □

COROLLARY 6.6. *A stochastic graph G with nonempty state space V is fully controllable if and only if $\overline{\Delta}_G$ is the boundary complex of the full simplex on V.*

REMARK 6.7. Here we view Corollary 6.6 as a consequence of Theorems 5.7 and 6.5. Alternatively, we could view Theorem 5.7 as a consequence of Corollary 6.6 and Theorem 6.5, thus giving us a different, combinatorial, proof of our main controllability theorem. Indeed, Corollary 6.6 is almost self-evident from the definition of source complex and Remark 6.9 below. To prove it fully, one would also need to make a backchaining argument showing how to combine individual strategies for attaining a particular state into a complete strategy for attaining that state.

DEFINITION 6.8. The *dual complex* of a stochastic graph $G = (V, \mathfrak{A})$ is the combinatorial Alexander dual of the source complex:

$$\overline{\Delta}_G^{*} \;=\; \bigl\{V \setminus W \;\bigm|\; W \subseteq V \text{ and } W \notin \overline{\Delta}_G\bigr\}.$$

(The underlying vertex set of $\overline{\Delta}_G^{*}$ is again V.)

Observe that $\emptyset$ is always a simplex of $\overline{\Delta}_G^{*}$, since V is never a simplex of $\overline{\Delta}_G$.

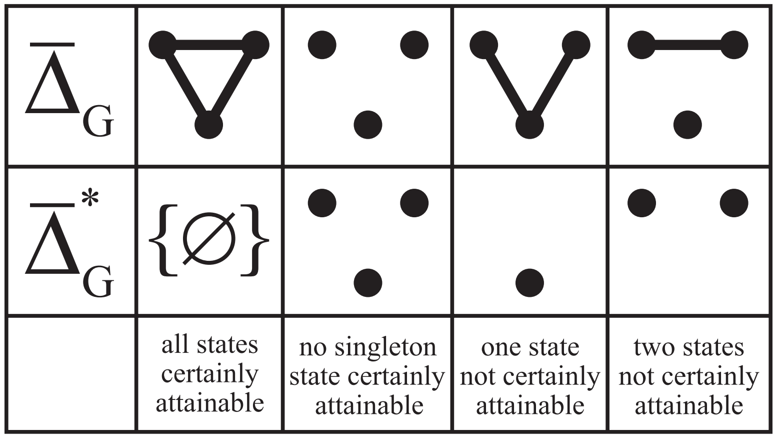

FIGURE 10. Possible source and dual complexes for a graph with three states and at least one convergent action at each state.

REMARK 6.9. The source complex $\overline{\Delta}_G$ of a graph G is the collection of all start regions of convergent sets of actions of G. The complements (relative to V) of these start regions are all the certainly attainable goals (see again Def. 5.1). The simplices of $\overline{\Delta}_G^*$ describe all *potentially unattainable* goals, that is, all sets of states that are not certainly attainable (from everywhere in the graph).

EXAMPLE 6.10. Figure 10 shows the relationship between $\overline{\Delta}_G$ and $\overline{\Delta}_G^*$, along with their meanings, for graphs on three states, assuming that, for every state, some action moves off that state. Given such actions, there are exactly four source complexes possible, ignoring state permutations, since there are three possible edges that may or may not be present in the source complex. Observe that there could be many different graphs that give rise to these complexes, but the details of these graphs are irrelevant at the level of understanding global capabilities. That observation is one interpretation of Theorem 6.5.

EXAMPLE 6.11. The source and dual complexes for the loopback graph of Figure 5 appear in the first column of complexes in Figure 10. The graph is fully controllable. The example of Figure 9 similarly maps to this same column in Figure 10. In contrast, the loopback graph of Figure 6 maps to the third column of complexes in Figure 10. Its source complex is in fact 1 —— 3 —— 2, with dual complex given by the singleton vertex representing state 3. This means that all goals are certainly attainable (in the loopback graph) except for state 3 alone, as we have seen in a variety of ways elsewhere.

The next theorem establishes arbitrary finite complexes as source complexes.

THEOREM 6.12. *For any finite simplicial complex Σ, there exists a nondeterministic graph G such that $\Sigma = \overline{\Delta}_G$ (disregarding underlying vertex sets).*

PROOF. We give the basic construction and point to [**14**] for further details. Let $G = (V, \mathfrak{A})$, with V consisting of $\Sigma^{(0)}$ plus one additional state, and let $\mathfrak{A}$ consist of all actions $x \to V \setminus X$, with $x \in X$ and X a maximal simplex of Σ. $\square$

6.2. System Design. The minimal nonfaces of $\overline{\Delta}_G$ are useful indicators of how a system loses full controllability, as Lemma 6.3 suggests. By a *minimal nonface* of $\overline{\Delta}_G$ we mean a set of states that is not a simplex of $\overline{\Delta}_G$ but all of whose proper subsets are simplices of $\overline{\Delta}_G$. An adversary can prevent the system from leaving any minimal nonface W of $\overline{\Delta}_G$, since no action moves off W. In particular, any stochastic action with source in W has all its targets in W, while any nondeterministic action with source in W has at least one target in W. For such actions, an adversary can, in effect, delete all the targets outside W (if even there are any). Doing so would produce a new graph with state space W. That graph would be fully controllable. To see this, observe that all proper subsets of W lie in $\overline{\Delta}_G$ since W is a minimal nonface. So the original system can certainly attain any goal of the form $\{w\} \cup (V \backslash W)$, with $w \in W$. Relative to the adversary preventing exit from W, this says the system has full controllability within W. This *relative controllability* is consistent with the simplices of $\overline{\Delta}_G$ inside W forming a sphere of dimension $|W| - 2$.

From a design perspective, one can treat the minimal nonfaces of an existing system as hints for improving system capabilities. The key is to fill in nonfaces by adding or modifying actions. For instance, in the example of Figure 6, since $\{\mathtt{1}, \mathtt{2}\}$ is the only minimal nonface of $\overline{\Delta}_{G_{\leftarrow 3}}$, adding *any* action with source $\mathtt{1}$ or $\mathtt{2}$ that moves off $\{\mathtt{1}, \mathtt{2}\}$, will establish full controllability. There are many different possibilities, including deterministic action $\mathtt{1} \rightarrow \mathtt{3}$ and stochastic action $\mathtt{1} \rightarrow \{\mathtt{p}_1\mathtt{1}, \mathtt{p}_2\mathtt{2}, \mathtt{p}_3\mathtt{3}\}$.

See [**14**] for further discussion of design.

These ideas extend to improving the performance of a system by considering the complexes $\Delta_G^{\mathcal{T}}$, along with their source and dual variants, for various times $\mathcal{T}$.

6.3. Inferring Adversarial Capabilities. The source complex also allows one to infer fairly high-level adversarial capabilities, again by looking for missing simplices. A simple observation is that

$$\overline{\Delta}_G \ = \ \Gamma * \Sigma_1 * \cdots * \Sigma_k,$$

where Γ is generated by a full simplex consisting of all states in the graph G that do not lie in any minimal nonface of $\overline{\Delta}_G$, while the $\{\Sigma_i\}$ are defined as follows: Define an equivalence relation on the states outside $\Gamma^{(0)}$ as the transitive closure of a simple relation in which two states are related whenever they lie in a common minimal nonface of $\overline{\Delta}_G$. Then Σ_i consists of all simplices of $\overline{\Delta}_G$ that lie within the i^{th} equivalence class of this equivalence relation.

This decomposition of $\overline{\Delta}_G$ reveals some time-varying adversarial capabilities. At any state that lies within multiple minimal nonfaces of $\overline{\Delta}_G$, an adversary may select within which of these minimal nonfaces to keep the system. Consequently, an adversary has some control over an *impatient* system, meaning a system that keeps trying to escape a minimal nonface, for instance by moving to every state of the minimal nonface and by eventually executing all actions available to it. If the system starts within a particular $\Sigma_i^{(0)}$ and is impatient, then the adversary can eventually force the system to reach any particular state within $\Sigma_i^{(0)}$. (Adversarially chosen transitions between different equivalence classes may also be possible, but that information is not directly knowable from the source complex.)

Finally, one can infer some adversarial capabilities more abstractly from homology and cohomology representatives, again by finding minimal nonfaces. Here are some sample results (see [**37, 25**] for background and notation):

LEMMA 6.13. *Suppose* $0 \neq [\alpha] \in \tilde{H}_p(\overline{\Delta}_G; \mathbb{Z})$, *with* $G = (V, \mathfrak{A})$ *a stochastic graph,* $p \geq 0$, *and* α *a simplicial cycle. Write* $\alpha = \sum_i n_i \sigma_i$, *with* $\{\sigma_i\}$ *a basis of elementary p-chains for* $\overline{\Delta}_G$ *and each* n_i *an integer. Then:*

(a) *For every* v *in* V, *there is some* σ_i *with* $n_i \neq 0$ *such that* $\sigma_i \cup \{v\} \notin \overline{\Delta}_G$.
(b) *No action moves off the set of states* $\bigcup_{n_i \neq 0} \sigma_i$ *(called the* support *of* α*).*

PROOF. (a) follows from the definition of nontrivial reduced homology: if no σ_i with the claimed property existed, then α would be homologous to zero.

(b) follows from (a) and Lemma 6.3. □

COROLLARY 6.14. *With notation as above, given any initial state of the system, either an adversary can force the system into the support of* α *and keep it there or the system has no action for moving off its initial state.*

LEMMA 6.15. *Suppose* $[\alpha^*] \in \tilde{H}^p(\overline{\Delta}_G; \mathbb{Z})$, *with* $G = (V, \mathfrak{A})$ *a stochastic graph,* $p \geq 0$, *and* α^* *a simplicial cocycle (possibly cobounding). Write* $\alpha^* = \sum_i n_i \sigma_i^*$, *with* $\{\sigma_i\}$ *a basis of elementary p-chains for* $\overline{\Delta}_G$ *and each* n_i *an integer.*

Let $v \in V \setminus \bigcup_{n_i \neq 0} \sigma_i$. *Then, for every* σ_i *with* $n_i \neq 0$, $\sigma_i \cup \{v\} \notin \overline{\Delta}_G$.

PROOF. Write $\tau_i = \sigma_i \cup \{v\}$ and suppose $\tau_i \in \overline{\Delta}_G$ for some i such that $n_i \neq 0$. With δ as coboundary operator and ∂ as boundary operator, bearing in mind that v lies outside the support of α^*, we calculate: $< \delta\alpha^*, \tau_i > = < \alpha^*, \partial\tau_i > = \pm n_i \neq 0$. That is a contradiction, since α^* is a cocycle. □

COROLLARY 6.16. *With notation as above, given any initial state of the system outside the support of* α^*, *either the system has no action for moving off its initial state or an adversary can force the system into some subspace of* any *of the defining simplices of* α^*. *Thereafter, the adversary can hold the system to the union of that subspace and the system's initial state.*

REMARKS 6.17. Lemmas 6.13 and 6.15 infer missing simplices of $\overline{\Delta}_G$ from simplicial cycles and cocycles. The two lemmas differ primarily in quantification. Lemma 6.13 and its corollary assert that an adversary can force a system to *some* defining simplex of a nonbounding cycle, whereas Lemma 6.15 and its corollary make a similar assertion for *any* defining simplex of a (possibly cobounding) cocycle. Lemma 6.15 is a generalization of the observation that the coboundary operator implicitly reveals missing simplices. (The two lemmas and their corollaries also impose slightly different conditions on the initial state of the system and describe different subspaces within which the adversary can hold the system.)

EXAMPLE 6.18. Suppose $\overline{\Delta}_G$ is the triangulation of the torus shown in Figure 11. Each minimal nonface of $\overline{\Delta}_G$ is a triangle and a generator of some one-dimensional homology. An adversary can hold the system to any such triangle once the system is at a state defining the triangle. For instance, the oriented 1-cycle $\alpha = [a, b] + [b, c] + [c, a]$ is a generator of some homology and constitutes a minimal nonface of $\overline{\Delta}_G$, so certainly the adversary can hold the system to the support of α. Corollary 6.14 makes the stronger point that the adversary can force the system to this 1-cycle *from any* state, should the system move from that state. Figure 11 also depicts the 1-cocycle α^*, whose support consists of all states except state c. Lemma 6.15 asserts that state c does not form a triangle with any of the defining edges of α^*. Consequently, the adversary can force the system from state c *to any* of α^*'s defining edges, should the system move from state c.

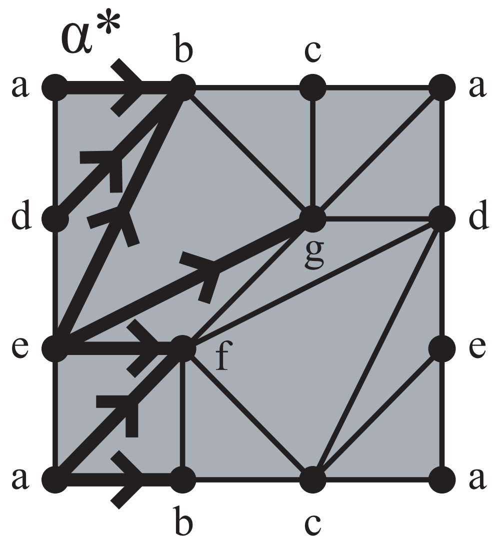

FIGURE 11. A particular triangulation of the torus. The thickened oriented edges depict a 1-cocycle α^* dual to the oriented 1-cycle $\alpha = [a, b] + [b, c] + [c, a]$.

The triangulation of Figure 11 is in fact very bad for the system. All states are equivalent by the equivalence relation defined at the beginning of this subsection, meaning an adversary can eventually force an impatient system anywhere.

7. Unions and Quotients

In this section we explore how the relationship between subgraphs and their encompassing graphs carries over to strategy complexes.

DEFINITION 7.1. A *stochastic subgraph* $H = (W, \mathfrak{B})$ of a stochastic graph $G = (V, \mathfrak{A})$ is a stochastic graph in its own right such that $W \subseteq V$ and $\mathfrak{B} \subseteq \mathfrak{A}$.

7.1. Graph Unions.

DEFINITION 7.2. If $G_1 = (V_1, \mathfrak{A}_1)$ and $G_2 = (V_2, \mathfrak{A}_2)$ are stochastic graphs with nonempty (possibly overlapping) state spaces, define their union $G_1 \cup G_2$ to be the stochastic graph $(V_1 \cup V_2, \mathfrak{A}_1 \sqcup \mathfrak{A}_2)$. Here "$\sqcup$" means "disjoint union"; we treat actions in $\mathfrak{A}_1$ and $\mathfrak{A}_2$ as distinct even if their written representations are identical.

LEMMA 7.3. *With notation and hypotheses as above, write $G = G_1 \cup G_2$.*

(a) *Always, $\Delta_G \subseteq \Delta_{G_1} * \Delta_{G_2}$.*
(b) *If either of the following conditions is satisfied, then $\Delta_G = \Delta_{G_1} * \Delta_{G_2}$:*
   (i) *At least one of G_1 and G_2 has no actions with sources in $V_1 \cap V_2$;*
   (ii) *$|V_1 \cap V_2| \leq 1$.*

PROOF. (a) The simplicial join $\Delta_{G_1} * \Delta_{G_2}$ is sensible since G combines actions of G_1 and G_2 with a disjoint union. The join is nonvoid since neither graph is null.

Finally, any convergent subset σ of $\mathfrak{A}_1 \sqcup \mathfrak{A}_2$ can be written as $(\sigma \cap \mathfrak{A}_1) \sqcup (\sigma \cap \mathfrak{A}_2)$, each term of which is a convergent set of actions in one of the original graphs.

(b)(i). Suppose $\mathrm{src}(\mathfrak{A}_1) \subseteq V_1 \setminus V_2$. Choose $\sigma_i \in \Delta_{G_i}$, $i = 1, 2$, and let $\sigma = \sigma_1 \sqcup \sigma_2$. If σ contains a circuit, then there is some nonempty set of actions $\tau \subseteq \sigma$ such that no action of τ moves off $\mathrm{src}(\tau)$. Since actions of G_2 have no targets in $V_1 \setminus V_2$ and since $\mathrm{src}(\tau \cap \mathfrak{A}_1)$ is disjoint from V_2, either $\tau \cap \mathfrak{A}_2$ is empty or it too defines a circuit. Since $\tau \cap \mathfrak{A}_2 \subseteq \sigma_2$, the second possibility cannot occur. The first possibility implies that $\tau \subseteq \sigma_1$, which means τ could not have defined a circuit.

(ii). If $V_1 \cap V_2 = \emptyset$, then $\Delta_G = \Delta_{G_1} * \Delta_{G_2}$ by (i). Otherwise, suppose τ is some nonempty subset of $\sigma_1 \sqcup \sigma_2$, with $\sigma_i \in \Delta_{G_i}$, $i = 1, 2$, such that no action of τ moves off $\mathrm{src}(\tau)$. One may choose τ so no two actions have the same source. Specializing (i) to a subgraph of G_1 and a subgraph of G_2, whose collections of actions are $\tau \cap \mathfrak{A}_1$ and $\tau \cap \mathfrak{A}_2$, respectively, produces a contradiction. □

7.2. Collapsing Fully Controllable Subgraphs.

DEFINITION 7.4. Suppose $G = (V, \mathfrak{A})$ is a stochastic graph with $V \neq \emptyset$. Let $\sim$ be an equivalence relation on V. Define the quotient graph $G/\!\sim\, = (V/\!\sim, \mathfrak{A}/\!\sim)$ as follows: The state space $V/\!\sim$ consists of one representative state for every equivalence class of $\sim$. The set of actions $\mathfrak{A}/\!\sim$ is nearly identical to the set $\mathfrak{A}$; the difference is we relabel source and target states by their representatives in $V/\!\sim$.

Three comments: (a) State relabeling may identify targets. In the case of stochastic actions, we sum the corresponding transition probabilities when such identifications occur. (b) Distinct actions of G may appear identical in $G/\!\sim$. We treat them as distinct. (c) Convergent actions of $\mathfrak{A}$ may become self-loops in $\mathfrak{A}/\!\sim$.

DEFINITION 7.5. Suppose $G = (V, \mathfrak{A})$ is a stochastic graph and W is a nonempty subset of V. Let $\sim$ be the relation in which two states are equivalent if they are identical or if they both lie in W. Write G/W for the quotient graph $G/\!\sim$.

LEMMA 7.6. *Suppose $H = (W, \mathfrak{B})$ is a fully controllable stochastic subgraph of stochastic graph $G = (V, \mathfrak{A})$, with $W \neq \emptyset$. Then*

$$\Delta_G \;\simeq\; \Delta_H * \Delta_{G/W}.$$

PROOF. This lemma is a special case of upcoming Lemma 8.12. □

DEFINITION 7.7. Suppose $G = (V, \mathfrak{A})$ is a stochastic graph. Let $\leftrightarrow$ be the equivalence relation in which two states u and v are equivalent if there exists some fully controllable stochastic subgraph $H = (W, \mathfrak{B})$ of G such that u and v both lie in W. (H may depend on u and v.)

REMARK 7.8. Relation $\leftrightarrow$ is not the same as attainability, by Figure 12.

THEOREM 7.9. *Let G be a non-null stochastic graph. Then*

$$\Delta_G \;\simeq\; \mathbb{S}^{n-k-1} \,*\, \Delta_{G/\leftrightarrow},$$

with n the size of G's state space and k the number of $\leftrightarrow$ equivalence classes.

PROOF. Apply Lemma 7.6 repeatedly, once for each $\leftrightarrow$ equivalence class (the order does not matter), using the fact that $\mathbb{S}^i * \mathbb{S}^j \simeq \mathbb{S}^{i+j+1}$. □

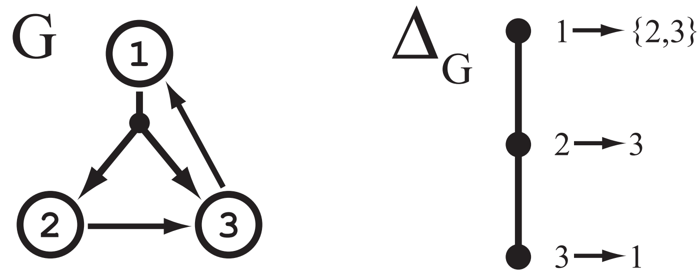

FIGURE 12. States 1 and 3 are certainly attainable from each other, but they do not together lie in a fully controllable subgraph. The nondeterminism of the action at state 1 may or may not force the system to pass through state 2 on its way to state 3; the system can neither definitely avoid state 2 nor be certain of attaining it.

REMARK 7.10. For deterministic directed graphs, $\leftrightarrow$ is the same as strong connectivity. In that case, $\Delta_{G/\leftrightarrow}$ is generated by a full simplex consisting of all directed edges between the strong components of G. If there are no such edges, then $\Delta_{G/\leftrightarrow}$ is the empty complex and $\Delta_G \simeq \mathbb{S}^{n-k-1}$. Otherwise, Δ_G is contractible. This is Hultman's result [**26**], mentioned earlier in Example 2.14. For nondeterministic graphs, $\Delta_{G/\leftrightarrow}$ can be much more general, as we have seen. Via the methods of Section 6, the complex $\Delta_{G/\leftrightarrow}$ reveals system limitations and adversarial power.

8. Topology of Prescribed Motions

In some situations, when creating a plan to accomplish a task, some actions may be prescribed, perhaps by earlier choices or by external constraints. For example, construction on a highway may force a local detour, or a broken finger may require a robot to perform an assembly from some specific angle. The link of the prescribed actions in the original strategy complex describes all the strategies consistent with the prescription. This section explores links in strategy complexes. We will see that many of our earlier ideas generalize. Indeed, we deferred proofs for some of the earlier results since they really are corollaries to this section. The next section will further use these results to characterize essential actions topologically.

8.1. Links and Sources. Following [**27**], we define deletion and link more generally than is customary, allowing deletions and links with respect to sets of elements that might not even be in the complex's underlying vertex set:

DEFINITION 8.1. Given a simplicial complex Σ and some set $\mathcal{E}$, we define the *deletion* $\mathrm{dl}(\Sigma, \mathcal{E})$ and the *link* $\mathrm{lk}(\Sigma, \mathcal{E})$ to be the following subcomplexes of Σ:

$$\begin{aligned} \mathrm{dl}(\Sigma, \mathcal{E}) &= \{\tau \in \Sigma \mid \tau \cap \mathcal{E} = \emptyset\}, \\ \mathrm{lk}(\Sigma, \mathcal{E}) &= \{\tau \in \Sigma \mid \tau \cap \mathcal{E} = \emptyset \text{ and } \tau \cup \mathcal{E} \in \Sigma\}. \end{aligned}$$

REMARKS 8.2. (1) If the underlying vertex set of Σ is X, then the underlying vertex set of both $\mathrm{dl}(\Sigma, \mathcal{E})$ and $\mathrm{lk}(\Sigma, \mathcal{E})$ is generally understood to be $X \setminus \mathcal{E}$.

(2) If $\mathcal{E}$ is not a simplex of Σ, then $\mathrm{lk}(\Sigma, \mathcal{E})$ is the void complex.

First, we generalize the definition of "moves off". To set the stage, suppose $G = (V, \mathfrak{A})$ is a stochastic graph, $\sigma \in \Delta_G$, and $W \subseteq V$. Suppose A is an action of G whose source lies in W. We are interested in the behavior of the system starting from $\mathrm{src}(A)$, given that the system first executes action A and then executes strategy σ, with the provision that execution of σ stops if ever the system re-enters W. Since $\sigma \cup \{A\}$ need not lie in Δ_G, we refer to this behavior by saying that the system *moves according to* $[A; \sigma \dashv W]$.

Definition 8.3. Let $G = (V, \mathfrak{A})$ be a stochastic graph, $\sigma \in \Delta_G$, $W \subseteq V$, and $A \in \mathfrak{A}$. Action A *moves off* W *subject to* σ if $A \notin \sigma$, $\mathrm{src}(A) \in W$, and the worst-case probability of returning to W from $\mathrm{src}(A)$ is strictly less than 1 whenever the system moves according to $[A; \sigma \dashv W]$.

Remarks 8.4. (1) If $\sigma = \emptyset$, Def. 8.3 is equivalent to Def. 4.5.

(2) Algebraically, we may express the probability condition of Def. 8.3 by requiring that $\mathrm{adv}(A, \{q_v\}) < 1$, where $\{q_v\}_{v \in V}$ is the solution to the following system of equations (the solution exists and is unique since σ is convergent):

$$q_w = \begin{cases} \max\limits_{B \in \sigma|w} \mathrm{adv}(B, \{q_v\}), & \text{if } w \in \mathrm{src}(\sigma) \setminus W; \\ 1, & \text{if } w \in W; \\ 0, & \text{otherwise.} \end{cases} \tag{8.1}$$

(The maximization at w is taken over $\sigma|w$, that is, all actions of σ with source w.)

Next, we generalize the definition of source complex:

Definition 8.5. Suppose $G = (V, \mathfrak{A})$ is a stochastic graph and $\sigma \in \Delta_G$. The *source link of* σ *in* Δ_G is the subcomplex of $\overline{\Delta}_G$ given by:

$$\overline{\mathrm{Lk}}(\Delta_G, \sigma) \;=\; \{\mathrm{src}(\tau) \mid \tau \in \mathrm{lk}(\Delta_G, \sigma)\}.$$

Lemma 6.3 generalizes as follows:

Lemma 8.6. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and $\sigma \in \Delta_G$. Suppose W is a nonempty subset of V such that every proper subset of W is a simplex of $\overline{\mathrm{Lk}}(\Delta_G, \sigma)$. Then $W \in \overline{\mathrm{Lk}}(\Delta_G, \sigma)$ if and only if some action of G moves off W subject to σ.*

Proof. I. Suppose $W \in \overline{\mathrm{Lk}}(\Delta_G, \sigma)$. Then there exists $\tau \in \Delta_G$ such that $\tau \cap \sigma = \emptyset$, $\tau \cup \sigma \in \Delta_G$, and $\mathrm{src}(\tau) = W$. We will now backchain using actions of τ and σ to define a sequence of triples $(W_1, A_1, v_1), \ldots, (W_k, A_k, v_k)$ such that $W \cup \mathrm{src}(\sigma) = \{v_1, \ldots, v_k\}$ and action A_j has source v_j and moves off W_j (in the standard sense of Def. 4.5):

- Let $W_1 = W \cup \mathrm{src}(\sigma)$ and $k = |W_1|$. (So $W_1 = \mathrm{src}(\tau) \cup \mathrm{src}(\sigma)$.)
- For $j = 1, \ldots, k$: Choose $v_j \in W_j$ so that all actions of $\tau \cup \sigma$ with source v_j move off W_j. Such a v_j exists since $\tau \cup \sigma$ is convergent and since $W_j \subseteq \mathrm{src}(\tau) \cup \mathrm{src}(\sigma)$. If $v_j \in W$, let A_j be an action of τ with source v_j. Otherwise, let A_j be an action of σ with source v_j. Finally, define $W_{j+1} = W_j \setminus \{v_j\}$. (Observe: $W_{k+1} = \emptyset$.)

Choose i to be the smallest index in $\{1, \ldots, k\}$ such that $A_i \in \tau$ (well-defined since $W \neq \emptyset$). This is the action A we seek. To verify: Action A is in τ, so not in σ. Action A moves off W in the standard sense, by minimality of i. We need to show that $\mathrm{adv}(A, \{q_v\}) < 1$, with $\{q_v\}_{v \in V}$ being the solution to System (8.1).

Suppose $i = 1$. Then A moves off $W \cup \mathrm{src}(\sigma)$. If A is nondeterministic, then all of A's targets lie outside $W \cup \mathrm{src}(\sigma)$, so $\mathrm{adv}(A, \{q_v\}) = 0 < 1$. If A is stochastic, then at least one target u of A lies outside $W \cup \mathrm{src}(\sigma)$, meaning $q_u = 0$, so $\mathrm{adv}(A, \{q_v\}) < 1$.

Suppose $i > 1$. By Def. 4.11, the previous conclusion holds more generally whenever $0 \leq q_u < 1$ for all relevant targets u (meaning all targets when A is nondeterministic; at least one target when A is stochastic). Recall that *all* actions of σ with source v_j move off W_j, for $j = 1, \ldots, i-1$. Inductively we therefore see that $q_{v_j} < 1$ for all $j = 1, \ldots, i-1$, and so $\mathrm{adv}(A, \{q_v\}) < 1$.

II. Let A be an action of G that moves off W subject to σ and let $w = \mathrm{src}(A)$. By assumption, there exists $\tau \in \Delta_G$ such that $\tau \cap \sigma = \emptyset$, $\tau \cup \sigma \in \Delta_G$, and $\mathrm{src}(\tau) = W \setminus \{w\}$. Write $\tau' = \tau \cup \{A\}$. Then $\mathrm{src}(\tau') = W$. Since $A \notin \sigma$, $\tau' \cap \sigma = \emptyset$. To establish that $W \in \overline{\mathrm{Lk}}(\Delta_G, \sigma)$, we therefore should show that $\tau' \cup \sigma \in \Delta_G$.

Suppose otherwise. Then, for some $\gamma \subseteq \tau' \cup \sigma$, with $A \in \gamma$, no action of γ moves off $\mathrm{src}(\gamma)$, in the standard sense of Def. 4.5. Let $\gamma' = \{B \in \gamma \mid \mathrm{src}(B) \notin W\} \subseteq \sigma$. Either γ' is empty or an adversary could select nondeterministic transitions so that, with probability 1, execution of γ' eventually moves the system into W whenever it starts in $\mathrm{src}(\gamma')$. That means $q_v = 1$ in (8.1) for every $v \in \mathrm{src}(\gamma)$, implying $\mathrm{adv}(A, \{q_v\}) = 1$, a contradiction. ☐

Finally, we generalize Lemma 6.2:

LEMMA 8.7. *Let $G = (V, \mathfrak{A})$ be a stochastic graph, $\sigma \in \Delta_G$, and $W \in \overline{\mathrm{Lk}}(\Delta_G, \sigma)$, with $W \neq \emptyset$. Then the following subcomplex of $\mathrm{lk}(\Delta_G, \sigma)$ is contractible:*

$$\Sigma_{W,\sigma} \;=\; \{\tau \in \mathrm{lk}(\Delta_G, \sigma) \mid \mathrm{src}(\tau) \subseteq W\}.$$

PROOF. Let A move off W subject to σ, as by Lemma 8.6. We claim that $\Sigma_{W,\sigma}$ is a cone with apex A. To see this, suppose $\tau \in \Sigma_{W,\sigma}$ and $A \notin \tau$. Write $\tau' = \tau \cup \{A\}$. We know that $\tau \cap \sigma = \emptyset$, $\tau \cup \sigma \in \Delta_G$, and $\mathrm{src}(\tau) \subseteq W$. We also know that $A \notin \sigma$ and $\mathrm{src}(A) \in W$. So we should show that $\tau' \cup \sigma \in \Delta_G$. The argument is identical to that given in the last paragraph of the proof of Lemma 8.6. ☐

8.2. Quillen Fiber Lemma. Before generalizing the key results of Sections 6 and 7, we review the Quillen Fiber Lemma [**40, 41**] in a form for our proofs. See [**7, 8, 47**] for further details regarding partially ordered sets (posets), order complexes, and the Quillen Fiber Lemma.

Every simplicial complex Σ defines a poset $\mathfrak{F}(\Sigma)$, called the *face poset* of Σ. The elements of $\mathfrak{F}(\Sigma)$ are the nonempty simplices of Σ, partially ordered by set inclusion. Conversely, any poset P defines a simplicial complex $\Sigma(P)$, called the *order complex* of P. The simplices of $\Sigma(P)$ are given by the finite chains $p_1 < \cdots < p_k$ in P. The order complex of a face poset, $\Sigma(\mathfrak{F}(\Sigma))$, is isomorphic to $\mathrm{sd}(\Sigma)$, showing that Σ and $\Sigma(\mathfrak{F}(\Sigma))$ are homeomorphic. One may therefore speak of the topology of a poset, implicitly meaning the topology of its order complex.

DEFINITION 8.8. If Q is a poset, let $Q_{\leq q}$ denote the set $\{q' \in Q \mid q' \leq_Q q\}$, with $\leq_Q$ being the partial order on Q (set inclusion in the case of face posets derived from simplicial complexes).

THEOREM 8.9 (Quillen Fiber Lemma). *Suppose $f : P \to Q$ is an order-preserving map between two posets. If $f^{-1}(Q_{\leq q})$ is contractible for all $q \in Q$, then P and Q are homotopy equivalent.*

Quillen's original version (as Theorem A in [**40**]) makes a stronger category-theoretic assertion than the combinatorial version we have reproduced here. Quillen's version for posets [**41**] is also stronger, asserting that f actually induces a homotopy equivalence. In this chapter, we only require the version stated above.

8.3. Two Link-Based Homotopy Equivalences. We generalize Theorem 6.5 and Lemma 7.6 to links.

THEOREM 8.10. *For any stochastic graph G and any $\sigma \in \Delta_G$,*

$$\mathrm{lk}(\Delta_G, \sigma) \;\simeq\; \overline{\mathrm{Lk}}(\Delta_G, \sigma).$$

PROOF. Let $P = \mathfrak{F}(\mathrm{lk}(\Delta_G, \sigma))$ and $Q = \mathfrak{F}(\overline{\mathrm{Lk}}(\Delta_G, \sigma))$ be the face posets of the two complexes. Define $f : P \to Q$ by $f(\tau) = \mathrm{src}(\tau)$. For any $W \in Q$, $f^{-1}(Q_{\leq W})$ is the face poset of $\Sigma_{W,\sigma}$, which is contractible by Lemma 8.7. The Quillen Fiber Lemma completes the argument. □

REMARK 8.11. For non-null G, Theorem 6.5 follows as a corollary, with $\sigma = \emptyset$.

In what follows, let tilde accents (as in $\tilde{\sigma}$) designate the image of G's actions in G/W as per Defs. 7.4 and 7.5.

LEMMA 8.12. *Suppose $H = (W, \mathfrak{B})$ is a fully controllable stochastic subgraph of stochastic graph $G = (V, \mathfrak{A})$, with $W \neq \emptyset$. Let $\sigma \in \Delta_G$ with $W \cap \mathrm{src}(\sigma) = \emptyset$. Then*

$$\mathrm{lk}(\Delta_G, \sigma) \;\simeq\; \Delta_H * \mathrm{lk}(\Delta_{G/W}, \tilde{\sigma}).$$

REMARK 8.13. Lemma 7.6 follows as a corollary, with $\sigma = \emptyset$.

PROOF OF LEMMA 8.12. Observe that $\tilde{\sigma} \in \Delta_{G/W}$, since $\sigma \in \Delta_G$ and σ contains no actions at states in W, so the lemma's statement makes sense. Moreover, by Theorems 6.5 and 8.10, we only need to prove that

$$\overline{\mathrm{Lk}}(\Delta_G, \sigma) \;\simeq\; \overline{\Delta}_H * \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma}).$$

Let $P = \mathfrak{F}(\overline{\Delta}_H * \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma}))$ and $Q = \mathfrak{F}(\overline{\mathrm{Lk}}(\Delta_G, \sigma))$ be the associated face posets. The elements of P are the nonempty simplices of $\overline{\Delta}_H * \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma})$. We may write any such simplex uniquely as $X \cup Y$, with $X \in \overline{\Delta}_H$ and $Y \in \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma})$, not both X and Y empty.

Let $\diamond$ be the state in G/W representing W identified to a point and define $f : P \to Q$ by

$$f\,(X \cup Y) = \begin{cases} X \cup Y, & \text{if } \diamond \notin Y; \\ W \cup Y \setminus \{\diamond\}, & \text{if } \diamond \in Y; \end{cases}$$

with $X \in \overline{\Delta}_H$ and $Y \in \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma})$. Observe that $X \subseteq W$ and $Y \subseteq (V \setminus W) \cup \{\diamond\}$.

Establishing the following conditions will complete the proof, by the Quillen Fiber Lemma:

(i) f is well-defined, meaning that $f(X \cup Y)$ really is a simplex of $\overline{\mathrm{Lk}}(\Delta_G, \sigma)$;
(ii) f is order-preserving;
(iii) the fibers $f^{-1}(Q_{\leq q})$ are contractible.

There are several cases to verify. We will prove the most interesting cases, leaving the remainder for the reader.

(i). We assume here that $\diamond \in Y$, leaving to the reader the case for which $\diamond \notin Y$.

We know there exists a set of actions $\tau \subseteq \mathfrak{A}$ such that $\tilde{\tau} \cap \tilde{\sigma} = \emptyset$, $\tilde{\tau} \cup \tilde{\sigma} \in \Delta_{G/W}$, and $\mathrm{src}(\tilde{\tau}) = Y$. (Recall again that, for example, $\tilde{\tau}$ is the image of G's actions τ

in the quotient graph G/W.) We can assume that $\tilde{\tau}$ contains exactly one action $\tilde{A}$ with source $\diamond$. Let $w \in W$ be the source of the corresponding action A of τ. Since H is fully controllable, there exists $\gamma \in \Delta_H$ such that $\mathrm{src}(\gamma) = W \setminus \{w\}$. Since $W \cap \mathrm{src}(\sigma) = \emptyset$, $(\gamma \cup \tau) \cap \sigma = \emptyset$. We know $\mathrm{src}(\gamma \cup \tau) = (W \setminus \{w\}) \cup (Y \setminus \{\diamond\}) \cup \{w\} = W \cup Y \setminus \{\diamond\}$. So it remains to establish that $\gamma \cup \tau \cup \sigma \in \Delta_G$.

The set of actions $(\tau \setminus \{A\}) \cup \sigma$ is convergent since the set $(\tilde{\tau} \setminus \{\tilde{A}\}) \cup \tilde{\sigma}$ is convergent and since the only difference, if any, between these two sets of actions is in the labeling of targets at which neither set has sources. Since γ has no sources or targets outside W and $(\tau \setminus \{A\}) \cup \sigma$ has no sources in W, Lemma 7.3(b)(i) tells us $\gamma \cup (\tau \setminus \{A\}) \cup \sigma$ is convergent. Any circuit contained in $\gamma \cup \tau \cup \sigma$ would therefore necessarily involve the action A. Mapping that circuit to the quotient graph G/W then tells us $\tilde{\tau} \cup \tilde{\sigma}$ must contain a circuit, which is impossible.

(ii). Easy.

(iii). Every $q \in Q$ is a nonempty simplex of $\overline{\mathrm{Lk}}(\Delta_G, \sigma)$, so we will write U in place of q, with $U \subset V$. Let such a nonempty U be given. We need to show that $f^{-1}(Q_{\leq U})$ is contractible. We assume here that $U \cap W = W$, leaving to the reader the case in which $U \cap W$ is a proper subset of W.

To establish contractibility, we will show that $f^{-1}(Q_{\leq U})$ is the face poset of a cone with apex $\diamond$. Observe that $f^{-1}(Q_{\leq U})$ is indeed the face poset of some simplicial complex, since f is order-preserving.

Let $X \cup Y \in f^{-1}(Q_{\leq U})$, with $X \in \overline{\Delta}_H$ and $Y \in \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma})$. Suppose $\diamond \notin Y$. We need to show that $X \cup Y \cup \{\diamond\}$ is an element of $f^{-1}(Q_{\leq U})$. Observe:

$$\begin{aligned} f\,(X \cup Y) &= X \cup Y \subseteq U, && \text{since } X \cup Y \in f^{-1}(Q_{\leq U}); \\ f\,(X \cup Y \cup \{\diamond\}) &= W \cup Y \subseteq U, && \text{since } W \subseteq U. \end{aligned}$$

To complete the proof, we will show that $Y \cup \{\diamond\} \in \overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma})$.

Suppose otherwise. Then there must be some set of states $Y' \subseteq Y$, such that every proper subset of $Y' \cup \{\diamond\}$ is a simplex of $\overline{\mathrm{Lk}}(\Delta_{G/W}, \tilde{\sigma})$ but $Y' \cup \{\diamond\}$ is not. By Lemma 8.6, no action of G/W moves off $Y' \cup \{\diamond\}$ subject to $\tilde{\sigma}$. Consequently, no action of G moves off $Y' \cup W$ subject to σ. Lemma 8.6 therefore implies that $Y' \cup W \notin \overline{\mathrm{Lk}}(\Delta_G, \sigma)$. That establishes a contradiction, since $Y' \cup W \subseteq Y \cup W \subseteq U \in \overline{\mathrm{Lk}}(\Delta_G, \sigma)$.

The same argument shows that $\{\diamond\}$ is itself in $f^{-1}(Q_{\leq U})$. □

9. Essential Actions

A graph may contain redundant actions. For instance, in the graph of Figure 7, the action $1 \to \{2, 3\}$ is completely unessential. We can remove the action without changing the overall capabilities of the system: with or without the action, the system can move from any state to any other state. The link of this action in the graph's strategy complex is the edge appearing in the right panel of Figure 13. It is contractible. That observation is generally true, as the next lemma shows.

Definition 9.1. Given stochastic graph $G = (V, \mathfrak{A})$ and actions $\mathcal{A} \subseteq \mathfrak{A}$, let $G \backslash \mathcal{A}$ denote the graph $(V, \mathfrak{A} \backslash \mathcal{A})$ formed from G by removing the actions $\mathcal{A}$.

Remark 9.2. Viewing $G \backslash \mathcal{A}$ as a subgraph of G, we may think of $\Delta_{G \backslash \mathcal{A}}$ as a subcomplex of Δ_G. In fact, $\Delta_{G \backslash \mathcal{A}} = \mathrm{dl}(\Delta_G, \mathcal{A})$.

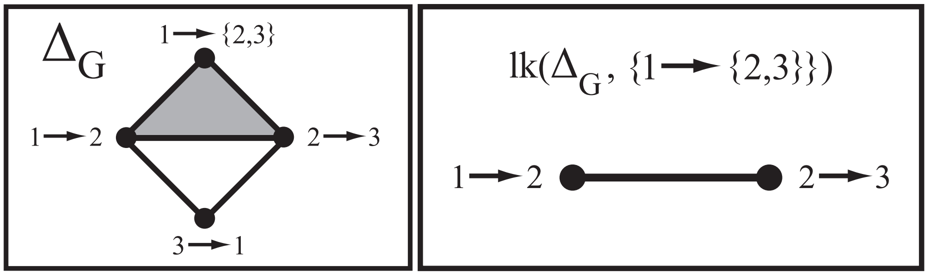

FIGURE 13. Left panel: The strategy complex Δ_G of Figure 7. Right panel: The link of the action $\texttt{1}\rightarrow\{\texttt{2},\texttt{3}\}$ in Δ_G. The link is contractible since $\texttt{1}\rightarrow\{\texttt{2},\texttt{3}\}$ is not essential for full controllability.

LEMMA 9.3. *Let $G = (V, \mathfrak{A})$ be a fully controllable stochastic graph. Suppose $\mathcal{A}$ is a nonempty set of actions in $\mathfrak{A}$ such that $G\backslash\mathcal{A}$ is also fully controllable. Then $\mathrm{lk}(\Delta_G, \mathcal{A})$ is contractible.*

PROOF. If $\mathcal{A}$ is not a simplex in Δ_G, then $\mathrm{lk}(\Delta_G, \mathcal{A})$ is the void complex, which is considered to be contractible. So we may assume that $\mathcal{A} \in \Delta_G$.

The covering set homotopy equivalence of Corollary 5.6 carries over to links of strategies as follows:

$$\mathrm{lk}(\Delta_G, \mathcal{A}) \;\simeq\; \left(\bigcup_{A\in\mathfrak{A}\backslash\mathcal{A}} U_A\right) \bigcap \left(\bigcap_{A\in\mathcal{A}} U_A\right).$$

Since $G\backslash\mathcal{A}$ is fully controllable, $\Delta_{G\backslash\mathcal{A}} \simeq \mathbb{S}^{n-2}$, with $n = |V|$. As a result, the union of all the covering sets U_A, with $A \in \mathfrak{A}\backslash\mathcal{A}$, is all of $\bigodot^n_{\neq}$. So

$$\mathrm{lk}(\Delta_G, \mathcal{A}) \;\simeq\; \bigcap_{A\in\mathcal{A}} U_A.$$

That last intersection is nonempty and convex, hence contractible. □

DEFINITION 9.4. Let $G = (V, \mathfrak{A})$ be a stochastic graph and let $s \in V$. Suppose G contains a complete strategy for attaining s. Let $\mathcal{A}$ be a set of actions in $\mathfrak{A}$. We say *$\mathcal{A}$ is essential for attaining s in G* if $G\backslash\mathcal{A}$ does not contain a complete strategy for attaining s.

THEOREM 9.5. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and let $s \in V$. Suppose G contains a complete strategy for attaining s. Let $\mathcal{A}$ be a nonempty subset of $\mathfrak{A}$.*

(a) *If $\mathcal{A}$ is not essential for attaining s in G, then $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A})$ is contractible.*
(b) *If $\mathcal{A}$ is essential, but no proper subset of $\mathcal{A}$ is essential, for attaining s in G, then $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A}) \simeq \mathbb{S}^{n-3}$, with $n = |V|$.*

PROOF SKETCH. Part (a) follows from Lemma 9.3. For part (b), we first observe that $\mathbb{S}^{n-3}$ is sensible, since $\mathcal{A}$'s existence means n is at least 2. We next outline the key steps of a proof, leaving the details to the reader:

(1) One may replace any stochastic action A of $\mathfrak{A} \setminus \mathcal{A}$ with a collection of deterministic actions, one for each stochastic transition of A, without changing the homotopy types of these complexes: Δ_G, $\Delta_{G_{\leftarrow s}}$, $\mathrm{lk}(\Delta_G, \mathcal{A})$, $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A})$. See Remark 9.7 below for more detail.

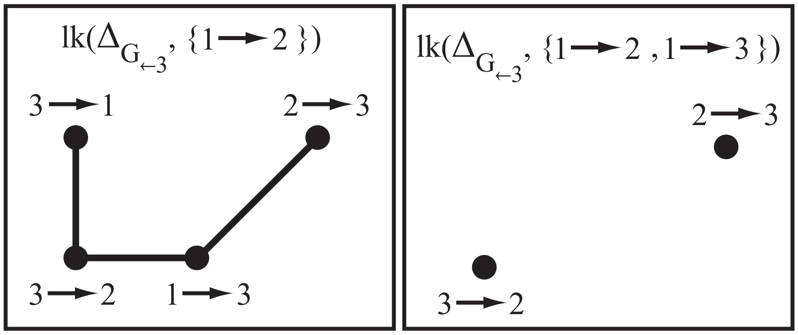

FIGURE 14. The link shown in the left panel is contractible since action $\texttt{1} \rightarrow \texttt{2}$ is not essential for attaining state $\texttt{3}$ in the graph of Figure 2. The link shown in the right panel is homotopic to $\mathbb{S}^0$ since together the actions $\texttt{1} \rightarrow \texttt{2}$ and $\texttt{1} \rightarrow \texttt{3}$ are essential for attaining state $\texttt{3}$ while neither is individually. See also Figure 5.

(2) Let G_{ndet} be G with all stochastic actions of $\mathfrak{A} \setminus \mathcal{A}$ replaced as in Step (1). Define W to be all states w in V for which G_{ndet} contains a strategy that attains s from w without requiring any of the actions $\mathcal{A}$. Observe $s \in W$.
(3) The loopback graph formed from G_{ndet} and s contains a fully controllable subgraph with state space W, disjoint from $\mathrm{src}(\mathcal{A})$. Prereasoning some of Step (4) shows that $\mathcal{A}$ is convergent. Lemma 8.12 then factors $\mathbb{S}^{|W|-2}$ out of the link of $\mathcal{A}$ in the loopback complex formed from G_{ndet} and s.
(4) Consider G again. Steps (1)–(3), along with the hypotheses regarding $\mathcal{A}$, allow us to assume without loss of generality that every action in $\mathcal{A}$ moves off $V \setminus \{s\}$ and that no action in $\mathfrak{A} \setminus \mathcal{A}$ moves off $V \setminus \{s\}$.
(5) Let $H = (V', \mathfrak{A}')$, with $V' = V \setminus \{s\}$ and $\mathfrak{A}' = \{A \in \mathfrak{A} \setminus \mathcal{A} \mid \mathrm{src}(A) \in V'\}$. This construction makes sense for all stochastic actions of $\mathfrak{A} \setminus \mathcal{A}$ by Step (4). There could, however, be nondeterministic actions in $\mathfrak{A} \setminus \mathcal{A}$ that have transitions both to s and to one or more states in $V \setminus \{s\}$. In constructing $\mathfrak{A}'$, remove from any such action the transition to s.
(6) Observe that for any $v \in \mathrm{src}(\mathcal{A})$, H contains a complete strategy for attaining state v, by the hypotheses regarding $\mathcal{A}$.
(7) Let H_+ designate H with all possible loopbacks added at every state of $\mathrm{src}(\mathcal{A})$. Using Theorems 6.5 and 8.10, along with the Quillen Fiber Lemma, one sees that $\Delta_{H_+} \simeq \mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A})$.
(8) Since H_+ is fully controllable and contains $n-1$ states, Theorem 5.7 establishes that $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A}) \simeq \mathbb{S}^{n-3}$.

□

EXAMPLE 9.6. The three-state graph of Figure 2 contains a complete strategy for attaining state $\texttt{3}$. The associated loopback graph and complex appear in Figure 5. The action $\texttt{1} \rightarrow \texttt{2}$ is not essential for attaining state $\texttt{3}$. Topologically, the link $\mathrm{lk}(\Delta_{G_{\leftarrow 3}}, \{\texttt{1} \rightarrow \texttt{2}\})$ is contractible, as shown in the left panel of Figure 14. Action $\texttt{1} \rightarrow \texttt{3}$ is not essential by itself either. However, the two actions $\texttt{1} \rightarrow \texttt{2}$ and $\texttt{1} \rightarrow \texttt{3}$ together clearly are essential for attaining state $\texttt{3}$. Topologically, the right panel of Figure 14 shows that $\mathrm{lk}(\Delta_{G_{\leftarrow 3}}, \{\texttt{1} \rightarrow \texttt{2}, \texttt{1} \rightarrow \texttt{3}\})$ is indeed homotopic to $\mathbb{S}^0$.

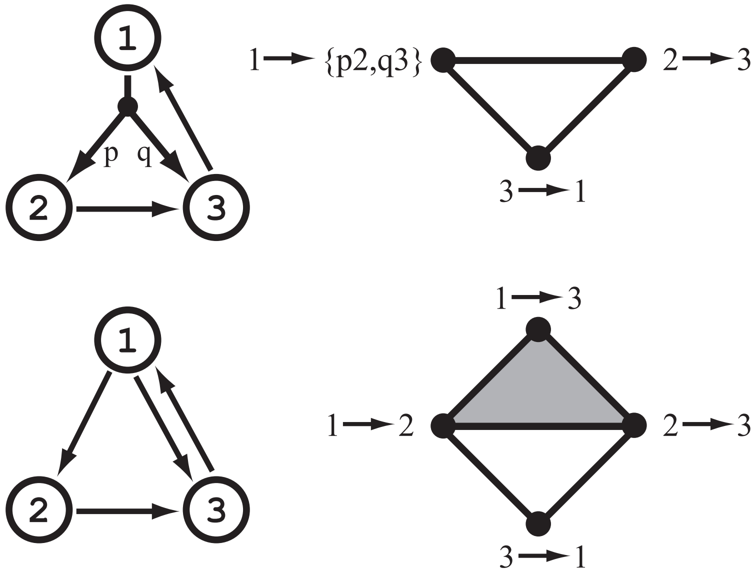

FIGURE 15. The graph in the top left panel contains an action with stochastic transitions. These transitions appear as separate deterministic actions in the graph shown in the bottom left panel. The strategy complexes for the two graphs, shown in the right two panels, have the same homotopy type.

REMARK 9.7. In Step (1) of the proof of Lemma 9.5, we observed that one may replace any stochastic action in a graph G by a collection of deterministic actions, one such action for each of the original stochastic transitions, without changing the homotopy type of Δ_G. In fact, the corresponding source complexes are identical, as a straightforward application of Lemma 6.3 shows. See Figure 15 for an example. The argument generalizes with the aid of Lemma 8.6, showing that replacement of stochastic transitions by deterministic actions does not change the homotopy type of $\mathrm{lk}(\Delta_G, \mathcal{A})$, so long as one does not make such replacements for any actions of $\mathcal{A}$. The same reasoning applies to $\Delta_{G_{\leftarrow s}}$ and $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A})$.

EXAMPLE 9.8. Lemma 9.3 is phrased in terms of a fully controllable graph, whereas Lemma 9.5 is specialized to loopback graphs. One wonders whether part (b) of Lemma 9.5 holds as well for fully controllable graphs. In fact, it need not hold when stochastic actions appear in the essential set $\mathcal{A}$, as Figure 16 shows.

10. Decision Trees

This section re-examines the structure of loopback complexes. We have seen that strategy complexes and source complexes can have the topology of any finite simplicial complex (Theorems 2.16 and 6.12), but loopback complexes are either spheres or points, homotopically (Theorem 3.6). Decision trees allow us to be more specific. We are motivated in this exploration by the extensive results Jonsson obtained for directed graph complexes via decision trees [**27**] as well as the connections Forman established between decision trees and discrete Morse theory [**18, 19**].

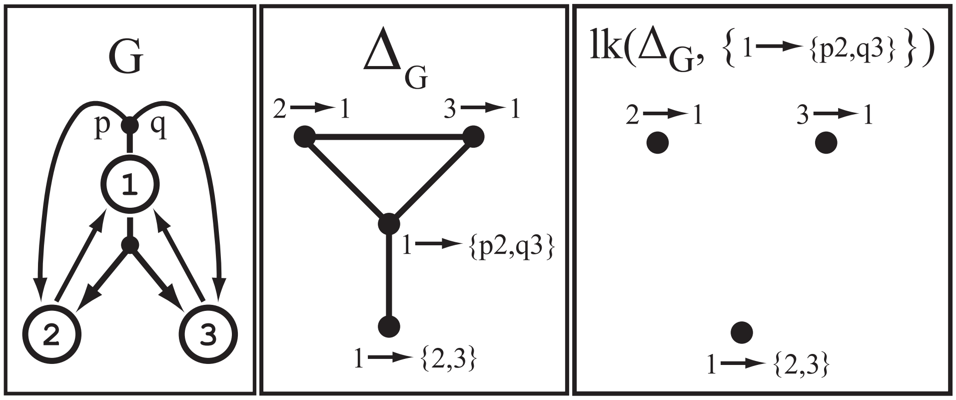

Figure 16. The graph on the left is fully controllable. Stochastic action $1 \rightarrow \{\mathtt{p2}, \mathtt{q3}\}$ is essential for full controllability. However, the link of that action in Δ_G is not homotopic to $\mathbb{S}^0$.

Intuitively, a decision tree is a variant of the "20 Questions" game, for determining whether an unknown set lies in some known collection of sets. Let Σ (perhaps a simplicial complex) be such a collection, drawn from an underlying vertex set X. We can view any subset σ of X as a bit vector over X. Suppose we know Σ exactly and someone has a secret σ that may or may not lie in Σ. We may ask whether individual bits in σ's bit vector are on or off. Any question we ask may depend on the answers to earlier questions. Our goal is to ask as few questions as possible in order to decide whether $\sigma \in \Sigma$. For example, suppose $x_0 \in X$ is some specific point, and suppose Σ consists of all subsets of X that do not contain x_0. Then the answer to one question, "Is x_0 in σ?", is sufficient to establish whether $\sigma \in \Sigma$. (We do not need to figure out what σ is exactly, merely whether it lies in Σ.)

In the worst case, one may need to ask $|X|$-many questions. Such sets σ are called *evasive* (relative to Σ and the questions being asked). Simplicial complexes for which one can structure the questions in such a way that no simplex is evasive are called *nonevasive*. This is a strong property. For finite simplicial complexes it is well-known [**7**] that the following proper inclusions hold:

$$\text{Cones} \subset \begin{matrix}\text{Nonevasive}\\ \text{Complexes}\end{matrix} \subset \begin{matrix}\text{Collapsible}\\ \text{Complexes}\end{matrix} \subset \begin{matrix}\text{Contractible}\\ \text{Complexes}\end{matrix}.$$

This section shows that a contractible loopback complex is in fact nonevasive. Similarly, for any loopback complex $\Delta_{G_{\leftarrow s}}$ homotopic to a sphere, one can pose the membership questions in an order such that exactly one simplex is evasive, corresponding to a complete strategy for attaining goal state s.

We now define decision trees recursively, much as one would in a functional programming language such as SML [**39**]. A decision tree is an object containing some data, with the object assuming one of two forms. One form of object, which we designate Node below, contains a simplicial complex, a vertex, and two subtrees. By containing two subtrees, a Node spawns two structural recursions. The other form of object, which we designate Leaf below, stops such recursions. A Leaf contains only a simplicial complex. There are further restrictions on the data in each object, made explicit in the next definition.

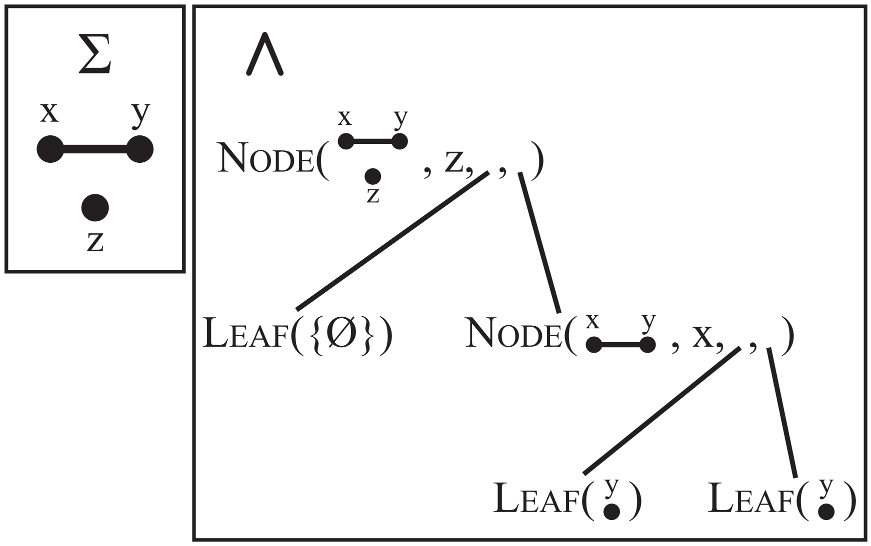

FIGURE 17. Left panel: A simplicial complex Σ generated by an edge $\{x, y\}$ and a point z. Right panel: One possible decision tree $\wedge$ for Σ. (Notation: Formally, a NODE contains its subtrees. Pictorially, it is convenient to draw edges to the subtrees. For ease of viewing, the figure depicts nonempty complexes geometrically rather than algebraically.)

DEFINITION 10.1. Suppose Σ is a finite simplicial complex with finite underlying vertex set X. A *decision tree for* Σ *(with underlying vertex set* X*)* is defined recursively as follows:

- Suppose Σ is one of the following complexes: $\emptyset$, $\{\emptyset\}$, $\{\emptyset, \{x\}\}$ (void complex, empty complex, single point x, respectively). Then LEAF(Σ) is a decision tree for Σ.
- Suppose $x \in X$. Define the underlying vertex set for both $\mathrm{lk}(\Sigma, \{x\})$ and $\mathrm{dl}(\Sigma, \{x\})$ to be $X \setminus \{x\}$. Suppose $\wedge_{\mathrm{lk}}$ is a decision tree for $\mathrm{lk}(\Sigma, \{x\})$ and $\wedge_{\mathrm{dl}}$ is a decision tree for $\mathrm{dl}(\Sigma, \{x\})$. Then NODE($\Sigma, x, \wedge_{\mathrm{lk}}, \wedge_{\mathrm{dl}}$) is a decision tree for Σ.
- Nothing else is a decision tree for Σ.

REMARKS 10.2. (1) One may think of NODE and LEAF as functions that take data and produce a decision tree, or, equivalently, as the actual objects containing that data. (2) Figure 17 shows one possible decision tree for a simplicial complex generated by an edge and a point. (3) Decision trees need not be unique. For instance, the complex consisting of a single point x with underlying vertex set $\{x\}$ gives rise to these two possible decision trees:

$$\text{LEAF}(\{\emptyset, \{x\}\}) \qquad \text{NODE}(\{\emptyset, \{x\}\}, x, \text{LEAF}(\{\emptyset\}), \text{LEAF}(\{\emptyset\})).$$

DEFINITION 10.3. A decision tree $\wedge$ for a finite simplicial complex Σ and a simplex σ of Σ together define a *descent path* through the tree, as follows: The path starts at $\wedge$. Suppose $\wedge = \text{NODE}(\Sigma, x, \wedge_{\mathrm{lk}}, \wedge_{\mathrm{dl}})$. If $x \in \sigma$, then the path continues via tree $\wedge_{\mathrm{lk}}$ and simplex $\sigma \setminus \{x\}$; otherwise, the path continues via tree

$\wedge_{\mathrm{dl}}$ and simplex σ. The path stops upon encountering a LEAF. The simplex σ is *evasive with respect to* $\wedge$ if the LEAF attained contains the empty complex, and *nonevasive with respect to* $\wedge$ otherwise. A simplicial complex Σ is *nonevasive* if there is some decision tree $\wedge$ for Σ such that every simplex of Σ is nonevasive with respect $\wedge$. See [**18, 27**] for further details.

REMARK 10.4. The void complex is nonevasive, as is the complex representing any full nonempty simplex. The empty simplex is evasive with respect to any decision tree for the empty complex.

The following lemma is immediate from the definitions:

LEMMA 10.5. *Suppose* $\wedge = \mathrm{NODE}(\Sigma, x, \wedge_{\mathrm{lk}}, \wedge_{\mathrm{dl}})$ *is a decision tree for finite simplicial complex* Σ.

(a) *If* $\wedge_{\mathrm{lk}}$ *and* $\wedge_{\mathrm{dl}}$ *establish that* $\mathrm{lk}(\Sigma, \{x\})$ *and* $\mathrm{dl}(\Sigma, \{x\})$*, respectively, are nonevasive, then* $\wedge$ *establishes that* Σ *is nonevasive.*
(b) *If exactly one simplex* τ *of* $\mathrm{lk}(\Sigma, \{x\})$ *is evasive with respect to* $\wedge_{\mathrm{lk}}$ *and every simplex of* $\mathrm{dl}(\Sigma, \{x\})$ *is nonevasive with respect to* $\wedge_{\mathrm{dl}}$*, then exactly one simplex* σ *of* Σ *is evasive with respect to* $\wedge$*, given by* $\sigma = \tau \cup \{x\}$.
(c) *If every simplex of* $\mathrm{lk}(\Sigma, \{x\})$ *is nonevasive with respect to* $\wedge_{\mathrm{lk}}$ *and exactly one simplex* τ *of* $\mathrm{dl}(\Sigma, \{x\})$ *is evasive with respect to* $\wedge_{\mathrm{dl}}$*, then exactly one simplex* σ *of* Σ *is evasive with respect to* $\wedge$*, given by* $\sigma = \tau$.

The following technical lemma regarding strategy complexes is also immediate:

LEMMA 10.6. *Let* $G = (V, \mathfrak{A})$ *be a stochastic graph,* $v \in V$*, and* $\tau \in \Delta_G$*. Suppose* $\mathcal{B}$ *is some collection of actions, all with source* v*, such that* $\tau \cup \{B\} \in \Delta_G$ *for each single action* $B \in \mathcal{B}$*. Then* $\tau \cup \mathcal{B} \in \Delta_G$.

We now strengthen the contractibility claim of Theorem 3.6:

COROLLARY 10.7. *Let* $G = (V, \mathfrak{A})$ *be a stochastic graph and* $s \in V$*. Suppose* G *does not contain a complete strategy for attaining* s*. Then* $\Delta_{G_{\leftarrow s}}$ *is nonevasive.*

PROOF. Let $\mathcal{A}_{\not s}$ denote the set of all actions of $G_{\leftarrow s}$ except for those with source s. Define a partial decision tree for $\Delta_{G_{\leftarrow s}}$ by arranging the actions $\mathcal{A}_{\not s}$ in any order and recursively constructing $\mathrm{NODE}(\Sigma, A, \wedge_{\mathrm{lk}}, \wedge_{\mathrm{dl}})$, with A ranging over $\mathcal{A}_{\not s}$, starting from $\Sigma = \Delta_{G_{\leftarrow s}}$, until all actions of $\mathcal{A}_{\not s}$ have been used along every possible descent path. Consider the (yet to be instantiated) subtrees $\{\wedge_i\}$ at the frontier of this partial decision tree. The "lk" branches on the descent path from the overall decision tree for $\Delta_{G_{\leftarrow s}}$ to subtree $\wedge_i$ define a subset $\mathcal{A}_i$ of $\mathcal{A}_{\not s}$. Subtree $\wedge_i$ is a decision tree for a subcomplex Σ_i of $\Delta_{G_{\leftarrow s}}$, consisting of all simplices σ such that $\sigma \cap \mathcal{A}_{\not s} = \emptyset$ and $\sigma \cup \mathcal{A}_i \in \Delta_{G_{\leftarrow s}}$. Subcomplex Σ_i could be void and thus nonevasive. Otherwise, its zero-skeleton is some subset of the loopback actions. That subset is nonempty. (To see this, observe that $\mathcal{A}_i$ must be a simplex of $\Delta_{G_{\leftarrow s}}$ when Σ_i is nonvoid, but it cannot be a complete strategy for attaining s. As in the proof Theorem 3.6, $\mathcal{A}_i$ can therefore join with at least one loopback action.) By Lemma 10.6, Σ_i must in fact represent a full nonempty simplex, so is nonevasive. Lemma 10.5(a) then implies that $\Delta_{G_{\leftarrow s}}$ is nonevasive. $\square$

The next lemma will be a useful stepping stone to two results. By a *minimal strategy* σ we mean a strategy with exactly one action at every state in $\mathrm{src}(\sigma)$.

Lemma 10.8. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and let $s \in V$. Suppose σ_0 is a minimal complete strategy for attaining s in G. Let $\mathcal{A}$ be some nonempty subset of $\mathfrak{A} \setminus \sigma_0$. Then $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A})$ is nonevasive.*

Proof. By Lemma 10.5(a), along with Remarks 8.2(2), 9.2, and 10.4, the proof reduces to the case in which $\mathcal{A}$ is convergent and consists of *all* actions in $\mathfrak{A} \setminus \sigma_0$ that do not have source s.

Let $\tau_0 = \{B \in \sigma_0 \mid \mathrm{src}(B) \notin \mathrm{src}(\mathcal{A})\}$ and write $\tau_0 = \{B_1, \ldots, B_k\}$, for some $k \geq 0$ ($k = 0$ means $\tau_0 = \emptyset$). Again by Lemma 10.5(a), the proof further reduces to showing that the complexes $\mathrm{dl}(\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A} \cup \{B_1, \ldots, B_{j-1}\}), \{B_j\})$, for $j = 1, \ldots, k$, and the complex $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A} \cup \tau_0)$ are all nonevasive. We will show that the last of these complexes, call it Σ, is nonevasive, leaving the first k complexes to the reader.

For every $v \in V \setminus \{s\}$, now let B_v designate the action of σ_0 with source v. By convergence of σ_0, the following system of equations has a unique finite solution:

$$\begin{aligned} t_v &= \mathrm{adv}(B_v, \{t_u\}) + 1, \quad \text{for } v \in V \setminus \{s\}; \\ t_s &= 0. \end{aligned} \tag{10.1}$$

Observe that if B_v is nondeterministic, then $t_v > t_u$ for all targets u of B_v. If B_v is stochastic, then $t_v > t_u$ for at least one target u.

Let $t_0 = \min\{\, t_v \mid v \in \mathrm{src}(\mathcal{A})\,\}$, let v_0 be a state in $\mathrm{src}(\mathcal{A})$ at which the minimum t_0 is attained, and let C_0 be the action of σ_0 with source v_0. (At least one v_0 exists since $\mathcal{A} \neq \emptyset$. If more than one v_0 is possible, pick any one.)

We will show that Σ, if not void, is a cone with apex C_0. Thus Σ is nonevasive. Let $\sigma \in \Sigma$ with $C_0 \notin \sigma$. Define $\sigma' = \sigma \cup \{C_0\}$. By construction, $\sigma \cap (\mathcal{A} \cup \tau_0) = \emptyset$ and $\sigma \cup \mathcal{A} \cup \tau_0 \in \Delta_{G_{\leftarrow s}}$. By definition, $C_0 \notin \mathcal{A}$. Since $\mathrm{src}(C_0) \in \mathrm{src}(\mathcal{A})$, $C_0 \notin \tau_0$. So $\sigma' \cap (\mathcal{A} \cup \tau_0) = \emptyset$. We need to establish that $\sigma' \cup \mathcal{A} \cup \tau_0 \in \Delta_{G_{\leftarrow s}}$, for then $\sigma' \in \Sigma$.

Suppose $\sigma' \cup \mathcal{A} \cup \tau_0$ is not convergent. Then there is some nonempty set of actions $\gamma \subseteq \sigma' \cup \mathcal{A} \cup \tau_0$ such that no action of γ moves off $\mathrm{src}(\gamma)$. Necessarily, $C_0 \in \gamma$. Some target u of C_0 with $t_0 > t_u$ lies in $\mathrm{src}(\gamma)$. Let $v_1 = u$, $t_1 = t_u$, and let C_1 be an action of γ with source u. Suppose $u \neq s$. By definition of t_0, this means C_1 must be an action of σ_0. Now repeat the construction. We obtain a sequence of distinct states $v_0, v_1, \ldots, v_i$, all in $\mathrm{src}(\gamma)$. By finiteness, we must eventually find that $v_i = s$. So $s \in \mathrm{src}(\gamma)$. On the other hand, observe that $\mathrm{src}(\mathcal{A} \cup \tau_0) = V \setminus \{s\}$, implying $\mathrm{src}(\sigma \cup \mathcal{A} \cup \tau_0) = V$. That is a contradiction, since $\sigma \cup \mathcal{A} \cup \tau_0$ is convergent. □

Remark 10.9. The method of decision trees is a combinatorial approach for inferring topology, seemingly different from the covering set approach we used earlier. A connection to covering sets appears via System (10.1).

We can now strengthen as well Lemma 9.5(a):

Corollary 10.10. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and let $s \in V$. Suppose G contains a complete strategy for attaining s. Let $\mathcal{A}$ be a nonempty subset of $\mathfrak{A}$ that is not essential for attaining s in G. Then $\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \mathcal{A})$ is nonevasive.*

The following result is useful in a decision-tree-theoretic proof that loopback complexes are spheres when complete strategies exist:

Corollary 10.11. *Let $G = (V, \mathfrak{A})$ be a stochastic graph and let $s \in V$. Suppose σ_0 is a minimal complete strategy for attaining s in G. There exists a decision tree for $\Delta_{G_{\leftarrow s}}$ with exactly one evasive simplex, given by σ_0.*

Proof. By using Lemma 10.5(c) and specializing Lemma 10.8 to cases in which $\mathcal{A}$ consists of a single action, we reduce to the case in which $\mathfrak{A} = \sigma_0$. Writing $\sigma_0 = \{A_1, \ldots, A_{n-1}\}$, with $n = |V|$, we may construct a partial decision tree for $\Delta_{G_{\leftarrow s}}$ whose frontier consists of decision trees for the following n complexes:

$$\begin{aligned} \Sigma_i &= \mathrm{dl}(\mathrm{lk}(\Delta_{G_{\leftarrow s}}, \{A_1, \ldots, A_{i-1}\}), A_i), \quad \text{for } i = 1, \ldots, n-1; \\ \Sigma_n &= \mathrm{lk}(\Delta_{G_{\leftarrow s}}, \{A_1, \ldots, A_{n-1}\}). \end{aligned}$$

For $i = 1, \ldots, n-1$, complex Σ_i is a cone with apex given by the loopback action $s \rightarrow \mathrm{src}(A_i)$. Complex Σ_n is the empty complex. Repeated application of Lemma 10.5(b) finishes the proof. □

We may now obtain the sphere result of Theorem 3.6 using decision trees:

Corollary 10.12. *Let $G = (V, \mathfrak{A})$ be a stochastic graph, $s \in V$, and $n = |V|$. If G contains a complete strategy for attaining s, then $\Delta_{G_{\leftarrow s}} \simeq \mathbb{S}^{n-2}$.*

Proof. We may assume that $n > 1$, as otherwise the claim is trivially true. Now recall the following general result from discrete Morse theory [**18**]:

Suppose $\wedge$ is a decision tree for a simplicial complex Σ, such that the empty simplex is nonevasive with respect to $\wedge$. Then Σ is homotopy equivalent to a CW complex consisting of exactly one p-cell for every p-dimensional simplex of Σ that is evasive with respect to $\wedge$, along with one additional 0-cell. In particular, if Σ is nonempty and admits a decision tree with exactly one evasive simplex, then the associated CW complex consists of a 0-cell and a k-cell, with k the dimension of the evasive simplex. The complex Σ is therefore homotopic to a sphere of dimension k.

In our case, Corollary 10.11 produces exactly one evasive simplex. That simplex has dimension $n-2$. □

11. Category Connections

This section explores connections between strategy complexes and category theory. For an introduction to category theory, see [**2**]. For more advanced treatments, see [**36, 23, 48**]. Then recall that the *nerve* of a small category is the simplicial set whose simplices are the diagrams of composable morphisms and that the *classifying space* of a small category is the geometric realization of its nerve. A strategy complex looks almost like the nerve of some category, particularly since actions look like morphisms. Moreover, any finite simplicial complex is a small category, via its face poset. In that setting, categorical nerve amounts to barycentric subdivision.

These observations suggest viewing strategy and source complexes as homeomorphic to the classifying spaces of planning processes, each described by a poset. Informally, "Plans are the nerve of planning." This section explores the foundations for that statement. We focus on nondeterministic graphs and give a sampling of key results. The constructions extend to stochastic graphs, with some technical modifications to account for probabilities and quantification differences.

11.1. Each Graph as a Category. We may view a given nondeterministic graph $G = (V, \mathfrak{A})$ as a category in several distinct but related ways:

(1) We *view G directly as a category* as follows: The objects are the individual states v of V plus every subset T of V that is the target set of some action in $\mathfrak{A}$. The morphisms are the actions $\mathfrak{A}$ plus all required identities. There are no compositions

except between an action $v \to T$ of $\mathfrak{A}$ and identities at v and T. In particular, here we view a source state v and a singleton target set $\{v\}$ as different objects.

(2) We *view G as a category with subsets* as follows: Much as in (1), but with additional objects and morphisms. The objects now are the individual states v plus *all* nonempty subsets S of target sets T. The morphisms now are *all* labeled arrows $v \xrightarrow{A} S$, with $A \in \mathfrak{A}$, $A = v \to T$, and $\emptyset \neq S \subseteq T$, plus all required identities.

(3) We *view G as a category with supersets* as follows: Much as in (2), except that we include supersets S of T instead of subsets.

Intuitively, one might interpret the original G in category (2) as an upper bound on an adversary's actual choices. One might interpret the original G in category (3) as a benign instantiation of more evil adversaries.

11.2. All Graphs as a Category. We may view all nondeterministic graphs as a category, again in several distinct but related ways. The objects in all cases are the graphs themselves. We include the null graph; it is an initial object in each category. A morphism $f : G \to H$ between two nondeterministic graphs is simultaneously a function on states and on actions that induces a functor from G to H, viewed as categories via Section 11.1. Composition of morphisms is composition of the underlying functions. To be explicit, we write out the analogue for case (2):

(2) In the *category of nondeterministic graphs with cycle-preserving morphisms*, a morphism $f : G \to H$, with $G = (V, \mathfrak{A})$ and $H = (W, \mathfrak{B})$, is simultaneously a function $f : V \to W$ and a function $f : \mathfrak{A} \to \mathfrak{B}$ such that: if $v \to S \in \mathfrak{A}$ and $f(v \to S) = w \to T \in \mathfrak{B}$, then $f(v) = w$ and $f(S) \subseteq T$. One can verify that identities and composition are well-defined in this category. The name for this category comes from Lemma 11.1 below.

(3) The *category of nondeterministic graphs with goal-preserving morphisms* is much like (2), except that we replace subset with superset. The name for this category comes from Lemma 11.2 below.

Lemma 11.1. *If $f : G \to H$ is a cycle-preserving morphism, then $f(\mathcal{A})$ contains a circuit in H whenever $\mathcal{A}$ contains a circuit in G.*

Proof. Suppose $\mathcal{A}$ contains a sequence of actions $v_1 \to S_1, \ldots, v_k \to S_k$, such that $v_{i+1} \in S_i$, for $i = 1, \ldots, k$, with $k \geq 1$ and $k+1$ meaning 1. For each i, $f(v_i \to S_i) = w_i \to T_i$ (some action of H), with $f(v_i) = w_i$ and $f(S_i) \subseteq T_i$. So $w_{i+1} \in T_i$, telling us $f(\mathcal{A})$ contains a circuit in H. $\Box$

Lemma 11.2. *If $f : G \to H$ is a goal-preserving morphism, with $G = (V, \mathfrak{A})$ and $H = (W, \mathfrak{B})$, and if Z is certainly attainable in G, then $f(Z) \cup (W \setminus f(V))$ is certainly attainable in H.*

Proof. We know $V \setminus Z \in \overline{\Delta}_G$. We need to show that $f(V) \setminus f(Z) \in \overline{\Delta}_H$. Using Lemma 6.3, it is enough to show that for every nonempty subset Y of $f(V) \setminus f(Z)$, some action of H moves off Y. Pick any such Y and let $X = V \setminus f^{-1}(f(V) \setminus Y)$. Observe that $\emptyset \neq X \subseteq V \setminus Z$, so $X \in \overline{\Delta}_G$. By Lemma 6.3, some action $v \to S$ in $\mathfrak{A}$ moves off X. In particular, $v \in X$ and $S \subseteq V \setminus X$. Since f is a goal-preserving morphism, $f(v \to S)$ is an action $w \to T$ in $\mathfrak{B}$, with $w = f(v) \in f(X) \subseteq Y$ and $f(S) \supseteq T$. Thus $T \subseteq f(V \setminus X) \subseteq f(V) \setminus Y$. So $f(v \to S)$ moves off Y. $\Box$

11.3. Functoriality of Strategy Complexes. Let JOIN^{+} be the following category: The objects are finite join semi-lattices with two restrictions: (1) Each semi-lattice must contain a distinguished top element $\hat{1}$. (2) A semi-lattice with more than one element must contain a distinguished bottom element $\hat{0}$. Morphisms are functions between semi-lattices viewed as sets, that further respect the join structure, send $\hat{1}$ to $\hat{1}$, and send $\hat{0}$ to $\hat{0} \neq \hat{1}$ whenever $\hat{0} \neq \hat{1}$ exists in the domain.

Comment: Let J, K be objects of JOIN$^+$, with $|J| = 1$ and $|K| > 1$. Then JOIN$^+$ contains identity morphisms at J and K, a unique morphism $J \to K$, but *no* morphism $K \to J$.

Definition 11.3. Given a nondeterministic graph G, we may view its strategy complex Δ_G as an object $L(G)$ in JOIN$^+$ by adjoining a top element $\hat{1}$. Intuitively, $\hat{1}$ represents circuits. The join operation is defined by saying that $\tau \vee \sigma = \tau \cup \sigma$ whenever τ, σ, and $\tau \cup \sigma$ are all simplices of Δ_G, and is $\hat{1}$ otherwise.

Comment: If G is null, then $L(G)$ consists just of $\hat{1}$. Otherwise, $L(G)$ also contains $\hat{0} = \emptyset \in \Delta_G$, distinct from $\hat{1}$.

Lemma 11.4. *The assignment operator L from Def. 11.3 is a functor from the category of nondeterministic graphs with cycle-preserving morphisms to* JOIN$^+$.

Proof. To define the functor, let $L(G)$ be as in Def. 11.3, then extend to morphisms as follows: If $f : G \to H$ is a cycle-preserving morphism, define the morphism $L(f) : L(G) \to L(H)$ in JOIN$^+$ as follows:

$$\begin{aligned} L(f)(\hat{1}) &= \hat{1}; \\ L(f)(\tau) &= f(\tau), \quad \text{if } \tau \in \Delta_G \text{ and } f(\tau) \in \Delta_H; \\ L(f)(\tau) &= \hat{1}, \quad \text{otherwise.} \end{aligned}$$

We now assume that all graphs are non-null; the null cases are straightforward.

To verify that $L(f)$ is a morphism, first observe that $L(f)$ sends $\hat{1}$ to $\hat{1}$ and $\hat{0}$ to $\hat{0}$ (since $\hat{0} = \emptyset$). Moreover, if x and y are elements in $L(G)$, with at least one of them being $\hat{1}$, then $L(f)(x \vee y) = \hat{1} = L(f)(x) \vee L(f)(y)$. So suppose $\tau, \sigma \in \Delta_G$. If $\tau \cup \sigma \notin \Delta_G$, then $f(\tau \cup \sigma) \notin \Delta_H$, by Lemma 11.1. So $L(f)(\tau \vee \sigma) = L(f)(\hat{1}) = \hat{1} = L(f)(\tau) \vee L(f)(\sigma)$. The last equality holds either because one of the terms is $\hat{1}$ already or because $f(\tau) \cup f(\sigma) = f(\tau \cup \sigma) \notin \Delta_H$. If $\tau \cup \sigma \in \Delta_G$ but $f(\tau \cup \sigma) \notin \Delta_H$, then $L(f)(\tau \vee \sigma) = L(f)(\tau \cup \sigma) = \hat{1} = L(f)(\tau) \vee L(f)(\sigma)$, with the last equality holding for the same reasons as before. Finally, if $\tau \cup \sigma \in \Delta_G$ and $f(\tau \cup \sigma) \in \Delta_H$, then $L(f)(\tau \vee \sigma) = f(\tau \cup \sigma) = f(\tau) \cup f(\sigma) = L(f)(\tau) \vee L(f)(\sigma)$.

To verify that L is a functor:

Identities. If $G \xrightarrow{i} G$ is an identity morphism, then $L(G) \xrightarrow{L(i)} L(G)$ is as well.

Composition. Suppose $G \xrightarrow{g} H \xrightarrow{h} K$. Consider $L(G) \xrightarrow{L(g)} L(H) \xrightarrow{L(h)} L(K)$. $L(h \circ g)$ and $L(h) \circ L(g)$ both send $\hat{1}$ to $\hat{1}$ and $\hat{0}$ to $\hat{0}$. Let $\tau \in \Delta_G$. If $g(\tau) \notin \Delta_H$, then $h(g(\tau)) \notin \Delta_K$ by Lemma 11.1. Thus $L(h \circ g)(\tau) = \hat{1}$ and $(L(h) \circ L(g))(\tau) = L(h)(L(g)(\tau)) = L(h)(\hat{1}) = \hat{1}$. If $g(\tau) \in \Delta_H$ but $h(g(\tau)) \notin \Delta_K$, then $L(h \circ g)(\tau) = \hat{1}$ and $(L(h) \circ L(g))(\tau) = L(h)(g(\tau)) = \hat{1}$. Finally, if $g(\tau) \in \Delta_H$ and $h(g(\tau)) \in \Delta_K$, then $L(h \circ g)(\tau) = h(g(\tau)) = (L(h) \circ L(g))(\tau)$, by definition of $\circ$. □

Remark 11.5. The join structure shows how Δ_G is homeomorphic to a planner's classifying space. Imagine a forward-chaining planner that takes unions of convergent sets of actions, retaining only those unions that remain convergent. This planner defines a category whose objects are the nonempty convergent sets of actions. The morphisms are induced by other sets of convergent actions, which may be adjoined without creating a circuit. The categorical nerve of this category is $\mathrm{sd}(\Delta_G)$.

11.4. Functoriality of Source Complexes. It is harder to see source complexes as functors. The reason is that source complexes discard information about precise paths and targets attained, retaining only the start regions of strategies. For functoriality, some contextual information appears to be necessary. This subsection explores one possible approach.

Definition 11.6. Let $G = (V, \mathfrak{A})$ be a nondeterministic graph. If $\mathcal{A}$ is any set of actions in $\mathfrak{A}$, let $\mathrm{trg}(\mathcal{A})$ denote all the targets of those actions. Formally, $\mathrm{trg}(\mathcal{A}) = \bigcup_{v \to T \in \mathcal{A}} T$.

Definition 11.7. Let $G = (V, \mathfrak{A})$ be a nondeterministic graph. Define the *poset of local strategies of* G, denoted $P(G)$, to consist of all pairs (W, X) such that: (i) $X \subseteq W \subseteq V$, and (ii) there exists a $\sigma \in \Delta_G$ with $\mathrm{src}(\sigma) = W \setminus X$ and $\mathrm{trg}(\sigma) \subseteq W$. Define a partial order on $P(G)$ by $(U, Y) \leq (W, X)$ if and only if $X \subseteq Y \subseteq U \subseteq W$.

Remarks 11.8. (1) If X is a nonempty proper subset of W, then (W, X) is in $P(G)$ precisely when G contains a nonempty strategy for attaining goal set X from start region $W \setminus X$, moving wholly within *ambient subspace* W. (2) An element of $P(G)$ "less" than another has a reduced ambient subspace and/or a looser goal set. (3) If G is the null graph, then $P(G)$ is the empty poset. Otherwise, $P(G)$ includes, for each subset W of G's state space, the trivial element (W, W).

Lemma 11.9. *The assignment operator P from Def. 11.7 is a functor from the category of nondeterministic graphs with goal-preserving morphisms to the category of finite posets.*

Proof. To define the functor, let $P(G)$ be as in Def. 11.7, then extend to morphisms as follows: If $f : G \to H$ is a goal-preserving morphism, define the poset morphism $P(f) : P(G) \to P(H)$ by $P(f)(W, X) = (f(W), f(X))$. One should verify that: (a) $P(f)(W, X) \in P(H)$ whenever $(W, X) \in P(G)$, (b) $P(f)$ is a poset morphism, and (c) P preserves identities and composition.

We will prove (a) and leave verification of (b) and (c) to the reader.

Let $(W, X) \in P(G)$. So there exists $\sigma \in \Delta_G$ such that $\mathrm{src}(\sigma) = W \setminus X$ and $\mathrm{trg}(\sigma) \subseteq W$. Consider any $v \to S \in \sigma$. Let $u \to T = f(v \to S)$. Since f is goal-preserving, $T \subseteq f(S) \subseteq f(W)$. So, we see that $\mathrm{trg}(f(\sigma)) \subseteq f(W)$.

Now let U be any nonempty subset of $f(W) \setminus f(X)$. We will show that some action of $f(\sigma)$ moves off this set. As a result, by applying Lemma 6.3 repeatedly, we see that there is some $\tau \in \Delta_H$ for which $\mathrm{src}(\tau) = f(W) \setminus f(X)$ and $\mathrm{trg}(\tau) \subseteq f(W)$.

Let $Y = W \setminus f^{-1}(f(W) \setminus U)$. Observe $\emptyset \neq Y \subseteq W \setminus X$. So σ contains an action $v \to S$ that moves off Y. Let $u \to T = f(v \to S)$. Then $u = f(v) \in f(Y) \subseteq U$ and $T \subseteq f(S) \subseteq f(W \setminus Y) \subseteq f(W) \setminus U$. This says the action $u \to T$ moves off U. $\square$

Definition 11.10. Given a nondeterministic graph $G = (V, \mathfrak{A})$ and a nonempty subset W of V, let $G|W = (W, \mathfrak{A}_W)$ with $\mathfrak{A}_W = \{v \to T \in \mathfrak{A} \mid v \in W \text{ and } T \subseteq W\}$. So $G|W$ consists of state space W and all actions of G whose motions lie within W. Also let $P|W = \{(W, X) \in P(G)\}$ (W fixed, X varying), inheriting $P(G)$'s partial order. So $P|W$ models the certainly attainable goals of $G|W$, suggesting the next lemma.

Lemma 11.11. *Let $G = (V, \mathfrak{A})$ be a nondeterministic graph. Suppose $\emptyset \neq W \subseteq V$. Then $P|W$ and $\overline{\Delta}_{G|W}$ are isomorphic when viewed as partially ordered sets.*

Proof. Define $P|W \xrightarrow{f} \overline{\Delta}_{G|W}$ and $\overline{\Delta}_{G|W} \xrightarrow{g} P|W$ by $f(W, X) = W \setminus X$ and $g(Y) = (W, W \setminus Y)$. The reader may verify that f and g are well-defined and that they are poset maps. They are inverses. □

11.5. A Planner's Classifying Space. Lemma 11.11 suggests how one may view $\overline{\Delta}_{G|W}$ as homeomorphic to a planner's classifying space. In particular, $P|W$ defines a poset category which one may interpret as a planner's search space. The category's objects are the elements of $P|W$. The category's morphisms are arrows of the form $(W, Y) \to (W, X)$, exactly one such arrow for each comparison $(W, Y) \leq (W, X)$ in $P|W$. This subsection elaborates the interpretation of $P|W$ as a planner.

Lemma 11.12. *Let $G = (V, \mathfrak{A})$ be a nondeterministic graph. Suppose $\emptyset \neq X \subseteq W \subseteq V$. Then $(W, X) \in P(G)$ if and only if there exists a sequence of triples $(X_0, A_0, v_0), \ldots, (X_{k-1}, A_{k-1}, v_{k-1})$, with $k \geq 0$, such that:*

(i) $W = X_0 \supset X_1 \supset \cdots \supset X_{k-1} \supset X_k \stackrel{\text{def}}{=} X$;
(ii) $X_j \setminus X_{j+1} = \{v_j\}$, *for* $j = 0, \ldots, k-1$;
(iii) A_j *is an action of* $G|W$ *with source* v_j *that moves off* $W \setminus X_{j+1}$, *for* $j = 0, \ldots, k-1$.

Proof. The lemma holds when $W = X$, by letting $k = 0$, so we may assume that X is a (nonempty) proper subset of W.

I. Suppose $(W, X) \in P(G)$. Let $\sigma \in \Delta_G$ be a minimal strategy satisfying $\operatorname{src}(\sigma) = W \setminus X$ and $\operatorname{trg}(\sigma) \subseteq W$. In particular, all actions of σ are actions of $G|W$. Now backchain from X, starting by letting $k = |W \setminus X|$ and $X_k = X$. Inductively, suppose $X_{j+1} \supset \cdots \supset X_k$ (and the corresponding triples) have been defined, for some $j \in \{0, \ldots, k-1\}$. Since σ is convergent, some action A_j of σ must move off $W \setminus X_{j+1}$. Let $v_j = \operatorname{src}(A_j)$ and $X_j = X_{j+1} \cup \{v_j\}$. Decrement j by one. Repeat this process until reaching (X_0, A_0, v_0), at which point $X_0 = W$.

II. Suppose the specified sequence of triples exists. Then the set of actions $\sigma = \{A_0, \ldots, A_{k-1}\}$ is convergent. Moreover, $\operatorname{src}(\sigma) = W \setminus X$ and $\operatorname{trg}(\sigma) \subseteq W$. So $(W, X) \in P(G)$. □

Remark 11.13. Viewing $P|W$ as a poset category, Lemma 11.12 says that for every object (W, X) in this category, there is a diagram of composable morphisms

$$(W, W) = (W, X_0) \to (W, X_1) \to \cdots \to (W, X_k) = (W, X), \tag{11.1}$$

with the $\{X_i\}$ satisfying the conditions stated in the lemma. This type of diagram defines a forward-chaining planner, as described next.

The planner starts its search at (W, W) and thereafter regards (W, W) as *visited*. The planner iterates as follows:

(a) Suppose the planner has visited $(W, Z) \in P|W$.
(b) Suppose there is some state $z \in Z$ and some action A of $G|W$ such that A has source z and A moves off $W \setminus Z'$, with $Z' = Z \setminus \{z\}$.
(c) This means (W, Z') is an object in $P|W$ and thus $(W, Z) \to (W, Z')$ is a morphism in $P|W$. If (W, Z') has not yet been visited, the planner *traverses the arrow* and visits (W, Z').

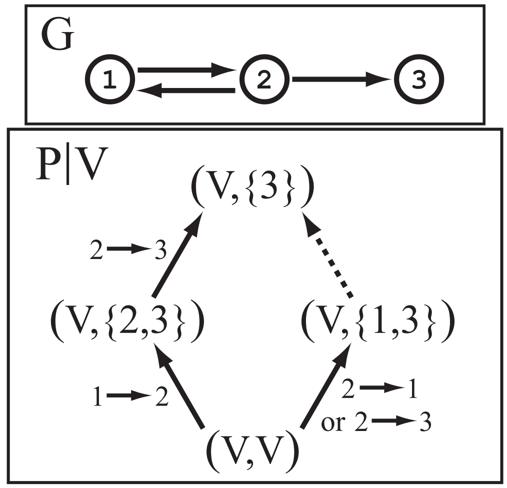

FIGURE 18. Top panel: A graph with state space $V = \{1, 2, 3\}$. Bottom panel: The category $P|V$, drawn with its four generating morphisms. (Not shown are four identity morphisms and the morphism $(V, V) \to (V, \{\mathtt{3}\})$ arising from composition.) Three of the four morphisms (solid arrows) are labeled by the actions that the forward-chaining planner of Remark 11.13 might produce. The fourth morphism (dashed arrow) amounts to a strategy switch; the planner would not traverse this arrow merely by forward-chaining, since no action of G with source `1` moves off $\{\mathtt{1}, \mathtt{2}\}$.

The planner loops over all possible choices of z and A at all visited objects, visiting more objects in the process, until it has visited every possible object in $P|W$. Some further comments:

(1) The conclusion that $(W, Z') \in P(G)$ in step (c) is independent of whatever strategy establishes that $(W, Z) \in P(G)$. This observation means that the planner needs to keep track merely of the objects (W, X) it encounters not of the actions leading to those objects.
(2) If desired, one may also view the planner as an output device. As such, the planner reports each arrow $(W, Z) \to (W, Z')$ that it traverses, labeled with the action A found in step (b). Implicitly, this linear output defines a strategy for attaining X from W for each object (W, X) visited by the planner. For any given (W, X), one may recover that strategy by scanning the output in reverse. (Of course, backchaining directly in G would be more efficient.)
(3) The category $P|W$ contains non-identity morphisms beyond the arrows traversed by the basic forward-chaining search just described, either because an arrow points to a previously visited object or because no action can label the arrow. See Figure 18. Viewed as planning operations, these additional morphisms constitute *strategy switches*: If the planner discovers both (W, X) and (W, Y) via forward-chaining, with $X \subset Y$, then the planner may, at least implicitly, traverse the arrow $(W, Y) \to (W, X)$,

thereby instantiating a preference for strategies with tighter goals over those with looser goals. (Morphism composition is a particular instance.)

Summary: The poset category $P|W$ defines the search space of a forward-chaining planner augmented with strategy switches that improve goal attainment. The objects (W, X) of $P|W$ constitute the planner's search states; the morphisms $(W, Y) \to (W, X)$ constitute the planner's possible state transitions.

Corollary 11.14. *Let $G = (V, \mathfrak{A})$ be a nondeterministic graph. Suppose $\emptyset \neq W \subseteq V$. Let $P^{\circ}|W$ be the poset subcategory of $P|W$ formed by removing the object (W, W). Then the categorical nerve of $P^{\circ}|W$ is isomorphic to $\mathrm{sd}(\overline{\Delta}_{G|W})$. Consequently, $\overline{\Delta}_G$ is homeomorphic to the classifying space of a category which one may interpret as a planner for finding all nontrivial certainly attainable goals in G.*

Proof. By Lemma 11.11, $P|W$ and $\overline{\Delta}_{G|W}$ are isomorphic posets. The bottom elements, (W, W) and $\emptyset$, respectively, correspond via this isomorphism. As a result, $P^{\circ}|W$ (which is $P|W \setminus \{(W, W)\}$) and $\overline{\Delta}_{G|W} \setminus \{\emptyset\}$ are isomorphic poset categories, implying that the categorical nerve of $P^{\circ}|W$ is isomorphic to $\mathrm{sd}(\overline{\Delta}_{G|W})$. Using $W = V$ in Remark 11.13, then dropping the planner's trivial initial state (V, V), establishes the lemma's planning assertion. □

Remark 11.15. The diagram of morphisms (11.1) gives rise to a $(k-1)$-simplex in $P^{\circ}|W$'s categorical nerve, of the form $(W, X_1) \to \cdots \to (W, X_k)$, with $k \geq 1$.

The following lemma provides an alternative perspective:

Lemma 11.16. *Let $G = (V, \mathfrak{A})$ be a nondeterministic graph. Suppose $\emptyset \neq W \subseteq V$. Then $(W, X_1) < (W, X_2) < \cdots < (W, X_k)$ is a chain in $P(G)$ if and only if there exist minimal strategies $\sigma_1 \subset \sigma_2 \subset \cdots \subset \sigma_k$ in $\Delta_{G|W}$, with all inclusions proper, such that $\mathrm{src}(\sigma_j) = W \setminus X_j$, with $X_j \subseteq W$, for $j = 1, \ldots, k$.*

Proof. I. Suppose $(W, X_1) < (W, X_2) < \cdots < (W, X_k)$ is a chain in $P(G)$. Let $\sigma_k \in \Delta_G$ be a minimal strategy satisfying $\mathrm{src}(\sigma_k) = W \setminus X_k$ and $\mathrm{trg}(\sigma_k) \subseteq W$. So $\sigma_k \in \Delta_{G|W}$. For $j = 1, \ldots, k-1$, let $\sigma_j = \{A \in \sigma_k \mid \mathrm{src}(A) \in W \setminus X_j\}$. Observe that $\mathrm{src}(\sigma_j) = W \setminus X_j$ since $X_k \subset X_j$. Moreover, σ_j is minimal since σ_k is minimal. By definition of $P(G)$, $X_j \subseteq W$. Finally, $\sigma_1 \subset \sigma_2 \subset \cdots \subset \sigma_k$ since $X_1 \supset X_2 \supset \cdots \supset X_k$, with all inclusions proper.

II. Suppose the specified strategies exist. Then $(W, X_j) \in P(G)$, for $j = 1, \ldots, k$. Since the strategies are minimal, proper inclusion of strategies implies proper inclusion of start regions. Consequently, $(W, X_1) < \cdots < (W, X_k)$. □

Remark 11.17. Consider a planner that constructs all possible strategies of $\Delta_{G|W}$ in the manner described by Remark 11.5. Suppose further that the planner reports the status of its search not by outputting full strategies but simply their start regions, along with corresponding arrows. Passing to complements in W, Lemma 11.16 tells us that this planner will traverse exactly all the arrows of the category $P|W$ (some arrows perhaps more than once). Omitting again the trivial object of $P|W$, one may therefore view $\overline{\Delta}_{G|W}$ as homeomorphic to the classifying space of a category derived from a planner. The difference between this planner and that described in Remark 11.13 lies in each planner's search space. The planner of Remark 11.13 never needs to remember actions; its search space truly is $P|W$. In contrast, the planner suggested by Lemma 11.16 operates in the join semi-lattice $L(G|W)$, then produces $P|W$ as a trace.

12. Discussion

12.1. Uncertainty, Geometry, and Topology. This research shows how the following are equivalent topologically: (a) the convergent sets of motions in finite graphs with control uncertainty, (b) finite simplicial complexes, (c) certain families of polyhedral cones in $\mathbb{R}^n$. Similar results were known for braid arrangements in $\mathbb{R}^n$ via earlier work on directed graph complexes and partially ordered sets [**7, 8, 26, 47, 27**]. Those ideas extend readily from directed graphs to graphs with nondeterministic transitions. The ideas extend as well to graphs with stochastic transitions, by allowing the hyperplanes comprising the $\mathbb{R}^n$ arrangements to rotate more freely about the line $\{x_1 = \cdots = x_n\}$. The usefulness of these hyperplanes is their ability to cast geometrically the expected convergence time equations of adversarially chosen Markov chains. That geometry provides a stepping stone to topology via the Nerve Lemma.

The significance of the step to topology is in showing how to infer global system capabilities from local graph connectivity. Traditionally, in order for a system with uncertainty to know that it can attain a goal, it tries to exhibit a strategy for doing so. Exhibiting a strategy entails combining uncertain actions in a manner that converges to the desired goal. Our theorems provide an alternative: instead of creating a specific strategy, one merely needs to show that the available actions cover a sphere. Exhibiting a specific strategy is one way to cover a sphere, but not the only way. For instance, imagine a collection of strategies $\{\sigma_v\}$, such that σ_v converges to goal s in time $\mathcal{T}$ when started from state v. The strategies need not be consistent with each other; their union may contain a circuit or simply take too long. Nonetheless, we know there is some strategy $\sigma \subseteq \bigcup_v \sigma_v$ that will converge to goal s in time $\mathcal{T}$ from all relevant v, simply because the union of the original strategies along with loopbacks from s must cover a sphere of the correct dimension. Section 6.2 suggests related applications in system design.

12.2. Higher-Order Interactions. The compression of strategy complexes to source complexes discards detailed action information while preserving knowledge of the system's global capabilities. This permits higher-level reasoning about interactions with an adversary. In effect, the start regions of strategies now become much like actions, facilitating reasoning about time-varying goals and tactics. This perspective holds as well when there are prescribed motions, via the source complexes of links. An interesting direction for future exploration is to vary these links and see how the complexes vary. From a practical perspective, this may be useful in disaster preparation.

12.3. Computational Complexity. Our recent robotics paper [**14**] provides explicit algorithms for computing many of the structures presented in this chapter, such as the source complex and the fully controllable subgraphs of any stochastic graph. That paper also discusses several approaches for computing the strategy complex of a stochastic graph. The subroutine used within many of these algorithms is a form of backchaining, much as it appeared in several proofs throughout this chapter. The backchaining algorithm starts with a desired goal set S. It searches for some action A that moves off the complement $V \setminus S$, then enlarges the goal to a new subgoal S' by adding the source of A to S. The algorithm repeats this process, now with S' in place of S, enlarging to a new S', and so forth. If and when an enlarged subgoal S' engulfs all of V, the algorithm terminates successfully; the actions it found constitute a strategy for attaining the original goal S. Otherwise, for some S', the algorithm is unable to find an action that moves off $V \setminus S'$, meaning that no strategy exists for attaining the original goal S from all of V. The worst-case runtime complexity of this backchaining algorithm is $O(|V|^2|\mathfrak{A}|)$, with $G = (V, \mathfrak{A})$ being the underlying stochastic graph. Depending on the particular graph G and the particular goal S, faster versions may be possible.

While backchaining has low polynomial time complexity, computing the source or strategy complex of a graph may require exponential time. At first glance, one suspects a representational defect, since the number of simplices in a simplicial complex can be exponential in the size of the complex's zero-skeleton. However, the underlying reasons appear to be intrinsic to the questions we are investigating rather than merely an artifact of the methods. Observe that the maximal simplices of a graph's source complex correspond to the minimal certainly attainable goals in the graph (as complements with respect to the state space V). It turns out that the problem of finding the size of the smallest certainly attainable goal in a nondeterministic graph is *NP*-complete [**14**]. Consequently, it is unlikely that faster than exponential time algorithms exist for determining the global capabilities of an uncertain system (as via a source complex), in the worst case.

From a robotics perspective, discovering that a problem is *NP*-complete is good news. The problem is nontrivial enough to be interesting yet probably not so complicated as to be intractable in all settings. Indeed, in applications one may be fortunate to have small upper bounds on the sizes of the minimal certainly attainable goals. Backchaining then allows one to construct the maximal simplices of the source complex reasonably quickly. The maximal simplices fully determine the complex. Other practical efficiency improvements are possible. For instance, in some cases the dual complex has a compact representation, as for a fully controllable graph. More generally, one may find the fully controllable subgraphs of a stochastic graph in fairly low polynomial time [**14**]. The results of Section 7 then simplify the source and strategy complexes by collapsing each such subgraph to a single state.

12.4. Complex Structure and Future Work. Loopback complexes are highly specialized structures, as analysis via decision trees shows, yet source and strategy complexes can be fairly arbitrary. Full controllability appears as homotopy equivalence to a sphere of a particular dimension. Dually, subcomplexes arising from certain subgraphs in a fully controllable graph must be cones, as we saw in several proofs. Yet, we do not understand in and of itself what contractibility of a strategy complex implies. This gap, between specific structure on the one hand and almost arbitrary structure on the other, suggests a spectrum of potential results. Future research should further classify the homotopy types of uncertain systems in order to facilitate effective design. A robot encountering a novel scenario should be able to abstract from it a topological hash, index into a table of applicable strategies, then select one such strategy optimized with respect to attendant objectives, much like a robot hand grasping an object today will select from a collection of forces in the grasp's null space. The research discussed in this chapter provides one step in that direction.

Acknowledgments

The author is grateful to Ben Mann, Rob Ghrist, and the entire SToMP group for making this work possible, and to Afra Zomorodian for the opportunity to present the results at the Computational Topology Short Course during JMM2011.

References

1. A. V. Aho, J. E. Hopcroft, and J. D. Ullman, *The Design and Analysis of Computer Algorithms*, Addison-Wesley, Reading, MA, 1974.
2. S. Awodey, *Category Theory*, second ed., Oxford University Press, Oxford, 2010.
3. A. Barr, P. Cohen, and E. Feigenbaum (eds.), *The Handbook of Artificial Intelligence*, William Kaufmann, Los Altos, California, 1981–1989.
4. R. Bellman, *Dynamic Programming*, Princeton University Press, Princeton, 1957.
5. T. Bergquist, C. Schenck, U. Ohiri, J. Sinapov, S. Griffith, and A. Stoytchev, *Interactive object recognition using proprioceptive feedback*, Proc. IEEE IROS Workshop: Semantic Perception for Mobile Manipulation, 2009.
6. D. P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, Englewood Cliffs, N.J., 1987.
7. A. Björner, *Topological methods*, Handbook of Combinatorics (R. L. Graham, M. Grötschel, and L. Lovász, eds.), vol. II, Elsevier, Amsterdam, 1995, pp. 1819–1872.
8. A. Björner and V. Welker, *Complexes of directed graphs*, SIAM J. Discrete Math **12(4)** (1999), 413–424.
9. M. Brady, J. M. Hollerbach, T. L. Johnson, T. Lozano-Pérez, and M. T. Mason, *Robot Motion: Planning and Control*, MIT Press, Cambridge, MA, 1982.
10. V. de Silva and R. Ghrist, *Coordinate-free coverage in sensor networks with controlled boundaries via homology*, Intl. J. Robotics Research **25(12)** (2006), 1205–1222.
11. B. R. Donald, *Error Detection and Recovery in Robotics*, Lecture Notes in Computer Science, No. 336. Springer Verlag, Berlin, 1989.
12. M. A. Erdmann, *An exploration of nonprehensile two-palm manipulation*, Intl. J. Robotics Research **17(5)** (1998), 485–503.
13. ———, *Shape recovery from passive locally dense tactile data*, Robotics: The Algorithmic Perspective (The Third Workshop on the Algorithmic Foundations of Robotics) (P. K. Agarwal, L. E. Kavraki, and M. T. Mason, eds.), A K Peters, Natick, MA, 1998, pp. 119–132.
14. ———, *On the topology of discrete strategies*, Intl. J. Robotics Research **29(7)** (2010), 855–896.
15. M. A. Erdmann and M. T. Mason, *An exploration of sensorless manipulation*, IEEE J. Robotics and Automation **4(4)** (1988), 369–379.
16. R. S. Fearing, *Tactile sensing for shape interpretation*, Dexterous Robot Hands (S. T. Venkataraman and T. Iberall, eds.), Springer Verlag, New York, 1990, pp. 209–238.
17. W. Feller, *An Introduction to Probability Theory and Its Applications*, (revised printing) third ed., vol. 1, John Wiley & Sons, New York, 1968.
18. R. Forman, *Morse theory and evasiveness*, Combinatorica **20(4)** (2000), 489–504.
19. ———, *A user's guide to discrete Morse theory*, Séminaire Lotharingien de Combinatoire **48** (2002), B48c.
20. R. Ghrist and S. LaValle, *Nonpositive curvature and Pareto-optimal coordination of robots*, SIAM J. Control & Optimization **45(5)** (2006), 1697–1713.
21. R. Ghrist, J. O'Kane, and S. LaValle, *Computing Pareto optimal coordinations on roadmaps*, Intl. J. Robotics Research **24(11)** (2005), 997–1010.
22. R. Ghrist and V. Peterson, *The geometry and topology of reconfiguration*, Advances in Applied Mathematics **38(3)** (2007), 302–323.
23. P. G. Goerss and J. F. Jardine, *Simplicial Homotopy Theory*, Birkhäuser Verlag, Basel, 1999.
24. K. Y. Goldberg, *Orienting polygonal parts without sensors*, Algorithmica **10(2–4)** (1993), 201–225.
25. A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2002.
26. A. Hultman, *Directed subgraph complexes*, Elec. J. Combinatorics **11(1)** (2004), R75.
27. J. Jonsson, *Simplicial complexes of graphs*, Ph.D. thesis, Department of Mathematics, KTH, Stockholm, Sweden, 2005.

28. S. Karlin and H. M. Taylor, *A Second Course in Stochastic Processes*, Academic Press, New York, 1981.
29. J.-C. Latombe, *Robot Motion Planning*, Kluwer Academic Publishers, Boston, 1991.
30. S. M. LaValle, *Planning Algorithms*, Cambridge University Press, New York, 2006.
31. M. Levoy, *The digital Michelangelo project*, Proc. Second Intl. Conf. on 3-D Digital Imaging and Modeling, 1999, pp. 2–11.
32. H. R. Lewis and C. H. Papadimitriou, *Elements of the Theory of Computation*, Prentice-Hall, Englewood Cliffs, New Jersey, 1981.
33. T. Lozano-Pérez, *The design of a mechanical assembly system*, Tech. Report AI-TR-397, S.M. thesis, MIT, Cambridge, MA, 1976.
34. T. Lozano-Pérez, J. L. Jones, E. Mazer, and P. A. O'Donnell, *HANDEY: A Robot Task Planner*, MIT Press, Cambridge, MA, 1992.
35. T. Lozano-Pérez, M. T. Mason, and R. H. Taylor, *Automatic synthesis of fine-motion strategies for robots*, Intl. J. Robotics Research **3(1)** (1984), 3–24.
36. J. P. May, *Simplicial Objects in Algebraic Topology*, University of Chicago Press, Chicago, 1967.
37. J. R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley, Menlo Park, CA, 1984.
38. D. K. Pai, *Multisensory interaction: Real and virtual*, Robotics Research: The Eleventh International Symposium (P. Dario and R. Chatila, eds.), Springer Verlag, Berlin, 2005, pp. 489–498.
39. L. C. Paulson, *ML for the Working Programmer*, second ed., Cambridge University Press, Cambridge, 1996.
40. D. Quillen, *Higher algebraic K-theory: I*, Lecture Notes in Mathematics, vol. 341, Springer Verlag, Berlin, 1973, pp. 85–147.
41. ———, *Homotopy properties of the poset of nontrivial p-subgroups of a group*, Advances in Math. **28(2)** (1978), 101–128.
42. J. J. Rotman, *An Introduction to Algebraic Topology*, Springer Verlag, Berlin, 1988.
43. M. Shirai and A. Saito, *Parts supply in SONY's general-purpose assembly system SMART*, Japanese Journal of Advanced Automation Technology **1** (1989), 108–111.
44. E. H. Spanier, *Algebraic Topology*, McGraw-Hill, San Francisco, 1966.
45. S. A. Stansfield, *Robotic grasping of unknown objects: A knowledge-based approach*, Intl. J. Robotics Research **10(4)** (1991), 314–326.
46. R. H. Taylor, M. T. Mason, and K. Y. Goldberg, *Sensor-based manipulation planning as a game with nature*, Robotics Research: The Fourth International Symposium (R. Bolles and B. Roth, eds.), MIT Press, Cambridge, MA, 1988, pp. 421–429.
47. M. L. Wachs, *Poset Topology: Tools and Applications*, IAS/Park City Mathematics Institute, Summer 2004.
48. G. Warner, *Topics in Topology and Homotopy Theory*, University of Washington, Seattle, https://digital.lib.washington.edu/researchworks/handle/1773/2641, 1999.
49. S. Weinberger, *Computers, Rigidity, and Moduli: The Large-Scale Fractal Geometry of Riemannian Moduli Space*, Princeton University Press, Princeton, 2005.
50. D. E. Whitney, *Quasi-static assembly of compliantly supported rigid parts*, Journal of Dynamic Systems, Measurement, and Control **104(1)** (1982), 65–77.

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213
E-mail address: me@cs.cmu.edu