

On the Topology of Discrete Strategies*

Michael Erdmann
School of Computer Science
Carnegie Mellon University
me@cs.cmu.edu

December 2009

Abstract

This paper explores a topological perspective of planning in the presence of uncertainty, focusing on tasks specified by goal states in discrete spaces. The paper introduces *strategy complexes*. A strategy complex is the collection of all plans for attaining all goals in a given space. Plans are like jigsaw pieces. Understanding how the pieces fit together in a strategy complex reveals structure. That structure characterizes the inherent capabilities of an uncertain system. By adjusting the jigsaw pieces in a design loop, one can build systems with desired competencies.

The paper draws on representations from combinatorial topology, Markov chains, and polyhedral cones. Triangulating between these three perspectives produces a topological language for describing concisely the capabilities of uncertain systems, analogous to concepts of reachability and controllability in other disciplines. The major nouns in this language are topological spaces.

Three key theorems (numbered 1, 11, 20 in the paper) illustrate the sentences in this language: (a) Goal Attainability: There exists a strategy for attaining a particular goal from anywhere in a system if and only if the strategy complex of a slightly modified system is homotopic to a sphere. (b) Full Controllability: A system can move between any two states *despite control uncertainty* precisely when its strategy complex is homotopic to a sphere of dimension two less than the number of states. (c) General Structure: Any system's strategy complex is homotopic to the product of a spherical part, modeling full controllability on subspaces, and a general part, modeling adversarial capabilities.

The paper contains a number of additional results required as stepping stones, along with many examples. The paper provides algorithms for computing the key structures described. Finally, the paper shows that some interesting questions are hard. For instance, it is *NP*-complete to determine the most precisely attainable goal of a system with perfect sensing but uncertain control.

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 U.S. Government. No official endorsement should be inferred.

*A much abbreviated version of this paper appeared as [28].

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE DEC 2009	2. REPORT TYPE	3. DATES COVERED 00-00-2009 to 00-00-2009
4. TITLE AND SUBTITLE On the Topology of Discrete Strategies		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER	
	5e. TASK NUMBER	
	5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,School of Computer Science,Pittsburgh,PA,15213		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT This paper explores a topological perspective of planning in the presence of uncertainty focusing on tasks specified by goal states in discrete spaces. The paper introduces strategy complexes. A strategy complex is the collection of all plans for attaining all goals in a given space. Plans are like jigsaw pieces. Understanding how the pieces fit together in a strategy complex reveals structure. That structure characterizes the inherent capabilities of an uncertain system. By adjusting the jigsaw pieces in a design loop, one can build systems with desired competencies. The paper draws on representations from combinatorial topology, Markov chains, and polyhedral cones. Triangulating between these three perspectives produces a topological language for describing concisely the capabilities of uncertain systems, analogous to concepts of reachability and controllability in other disciplines. The major nouns in this language are topological spaces. Three key theorems (numbered 1, 11, 20 in the paper) illustrate the sentences in this language: (a) Goal Attainability: There exists a strategy for attaining a particular goal from anywhere in a system if and only if the strategy complex of a slightly modified system is homotopic to a sphere. (b) Full Controllability: A system can move between any two states despite control uncertainty precisely when its strategy complex is homotopic to a sphere of dimension two less than the number of states. (c) General Structure: Any system's strategy complex is homotopic to the product of a spherical part, modeling full controllability on subspaces, and a general part, modeling adversarial capabilities. The paper contains a number of additional results required as stepping stones, along with many examples. The paper provides algorithms for computing the key structures described. Finally, the paper shows that some interesting questions are hard. For instance it is NP-complete to determine the most precisely attainable goal of a system with perfect sensing but uncertain control.		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 81	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Contents

1	Introduction	4
1.1	Planning Manipulation Strategies with Uncertainty	4
1.2	Understanding System Capabilities	4
1.3	Topology	5
1.4	Result Flavor	5
1.5	Spheres as Topological Descriptors of Task Solvability	6
1.6	Contributions	6
1.7	An Underlying Motivation	7
1.8	Broader Context	7
1.9	Outline	8
2	An Example: Nondeterminism, Cycles, and Strategies	9
3	Nondeterministic Graphs and Strategy Complexes	11
3.1	Basic Definitions	11
3.2	Examples	12
4	Loopback Graphs and Complexes	15
5	Topological Tools and Homotopy Equivalence	17
5.1	Deformation Retractions	18
5.2	Collapsibility and Contractibility	18
5.3	The Nerve Lemma	19
5.4	The Quillen Fiber Lemma	19
5.5	Homotopy Interpretation	20
6	Stochastic Graphs and Strategy Complexes	20
6.1	Stochastic Actions and Graphs	20
6.2	Stochastic Acyclicity	21
6.3	Stochastic Strategy Complexes	23
7	Covering Sets	25
7.1	Homogeneous Covering Sets	25
7.2	Affine Covering Sets	25
7.3	Inferring Topology from Covering Sets	26
8	Controllability of Motions in Stochastic Graphs	30
8.1	Connectivity: Covers and Chains	31
8.2	Characterizing Controllability with Spheres	34
9	Topology as a Design Tool: An Example	36
9.1	How Many Design Scenarios?	36
9.2	Tuning Convergence Times and Designing System Capabilities	37

10 Duality	41
10.1 Start Region Contractibility	41
10.2 Source Complex	42
10.3 Contractibility Characterization of Goal Attainability	43
10.4 The Dual Complex	44
10.5 Duality in Design	44
11 Modularity	46
11.1 Graph Union	46
11.2 Testing Acyclicity	47
11.3 Simplification via Strongly Controllable Subspaces	48
11.4 An Example (Air Travel During Thunderstorm Season)	52
12 Algorithms	55
13 Realizability	59
14 Hardness	60
14.1 The Difficulty of Determining a System’s Precision	60
14.2 Small Realization is Uncomputable	61
14.3 Recognizing Repercussions is Uncomputable	62
15 Topological Thinking	63
15.1 Topology Precompiles an Existence Argument	63
15.2 Topological Analysis of Adversity	64
15.3 Topological Thinking in Partially Observable Spaces	66
15.3.1 Inferring Task Unsolvability From Duality	67
15.3.2 Hypothesis-Testing and Sphere Suspension	69
16 Conclusions	71
16.1 Summary	71
16.2 Other Results	71
16.3 Open Questions	71
17 Acknowledgments	72
List of Primary Symbols	73
List of Lemmas and Theorems	74
List of Algorithms	75
List of Key Definitions	75
List of Figures	76
References	77

1 Introduction

1.1 Planning Manipulation Strategies with Uncertainty

Two themes pervade the last four decades of research in robotic manipulation. The first is the difficulty of uncertainty, the second is a solution by discrete modeling. The jamming diagram appearing in Whitney’s classic paper on peg-in-hole assembly [77], grounded in extensive work at the C.S. Draper Laboratory [62] in the 1970s, illustrates both themes: (i) The precise location and motion of the peg are uncertain. (ii) The contact state of the peg naturally discretizes the assembly space. By analyzing this discretized space, [77] produces an automatic strategy for assembly that is robust in the presence of uncertainty.

As this early example demonstrates, discretization appears in manipulation research to simplify, but not artificially so. Instead, the mechanics of a problem often generate natural discrete states even when the initial rendering of the problem is continuous. Contact modes, describing which features of multiple objects are in contact and how are they sliding relative to each other, are particularly common methods of discretization. As further illustration, the cooperative manipulation strategies of [27] are based on discrete states representing regions of configuration space over which the frictional contact mechanics of the robot palms and the part are invariant. Discrete states may also capture higher-order information, perhaps modeling sensing uncertainty. For instance, in the part orienters of [29, 72, 37, 30], the discrete states considered by the motion planners were sets of underlying contact states of the parts being oriented. In this manner, sensing uncertainty itself contributes to the definition of state [5, 54].

Delineating relevant states is only part of discrete modeling. One must also describe transitions between states, with the aim of synthesizing robust manipulation strategies. Anyone who has ever programmed a robot manipulator quickly learns never to expect a particular grasping operation or a particular assembly trajectory to succeed. Instead, transitions between states are generally uncertain, due both to underlying control uncertainty and, in the discrete case, to approximation effects. As a result, the discrete modeling of a manipulation problem does not define a standard directed graph, but rather a nondeterministic or stochastic graph, in which commanded motions may have any one of several possible outcomes. Consequently, one must program robots by thinking in terms of sets of possible motions not individual trajectories. The preimage methodology of [56] emphasizes this point and shows how to generate robust strategies in a fairly general, continuous, setting. In the discrete setting, that methodology is akin to Bellman-style backchaining [4]. Recent manipulation results [42, 43] demonstrate the utility of these ideas in stochastic settings.

1.2 Understanding System Capabilities

The description of planning above is highly operational. One can almost turn a crank: Given a manipulation task, create a discrete representation of the task based on the task mechanics and uncertainties, then backchain from the task goal. Missing is an understanding of the full capabilities of such a discrete representation. Operationally, one can report whether a planner found a plan to perform some task, but one does not understand intrinsically when there should be a plan and what failure to find a plan really means. It would be good to understand what the remedy for a failure might be, for instance, more precise sensing, more precise control, a change in the task goal, or perhaps even a remodeling of the task description and discretization.

Our field does not even have an adequate language to describe the issues. In control theory, terms such as “controllable”, “unstable”, and “limit cycle” are common; in graph theory, “strongly connected”, “source” and “sink”; in Markov chains, “recurrent class”, “transient” and “absorbing”. In different languages and in different settings, these terms describe similar concepts, significant for characterizing system capabilities.

Such terms also offer standards by which to design systems. One learns how to create or destroy limit cycles in control theory, for example. One understands that a finite Markov chain cannot consist exclusively of transient states, whereas an infinite one can. All effective engineering fields have such concise descriptions of capabilities. For instance, to pick an example very different from the previous ones, in computer language theory, one learns that certain context-free grammars are easily compilable while others are not.

Planning in the presence of uncertainty has no such concise descriptions of capabilities and thus no concise design standards.

1.3 Topology

A thesis of this paper is that topology provides a language for concise descriptions of system capabilities. In order to develop that language, we follow a tack common in engineering disciplines: Given a method for solving any particular problem operationally, one should attempt to classify problems and solutions generally, for instance, by observing similarities in problems and exploiting commonalities of solution. In the case of planning for uncertain systems, a plan seldom exists in isolation; instead it is part of an ensemble of interconnected strategies for accomplishing an ensemble of interconnected goals in some system.

Topology is good at extracting structure from such an interconnected patchwork. That structure constitutes our language. Moreover, this language, as we will see, has two desirable traits:

1. The topological descriptors abstract away details of particular trajectories, focusing instead on overall system capabilities.
2. The topological descriptors are consistent with existing terminology for standard directed graphs and Markov chains, while generalizing both.

Our search specifically for a topological language to describe system capabilities was motivated by Robert Ghrist’s technology transfer between topology and robotics [13, 36, 34, 35], by Steven LaValle’s work on information spaces [38, 65, 74, 73], and by workshops on topology and robotics organized by Michael Farber at ETH Zürich in 2003 and 2006 [57].

1.4 Result Flavor

An analogy for understanding the relationship of topology to planning might be the relationship of linear algebra to vectors written with specific coordinates. Linear algebra provides abstract techniques for representing and manipulating vectors, independent of coordinates. Anything one can do with matrix and vector notation one can do as well by writing out arrays of coordinates and manipulating the coordinates directly. Indeed, ultimately computations in a computer must work with numbers in some coordinate system. However, the numbers are like trees that obscure the forest. By thinking instead at the abstract level of linear subspaces, kernels, eigenvectors, and so forth, one can recognize fundamental structure. How easy it is

to say that a high-dimensional positive definite system always decomposes into orthogonal one-dimensional systems; how cumbersome it would be to convey that truth by numbers or coordinates alone.

This perspective has precedent in other areas of computer science, such as the pioneering work of Herlihy on asynchronous computation. For instance, [40] shows that an asynchronous decision task has a wait-free protocol if and only if there exists a certain color-preserving continuous function between two chromatic simplicial complexes determined by the task. In other words, a computational problem is equivalent to a topological problem. The simplicial complexes reflect the structure of the input and output spaces of the task. The continuous function is a topological constraint between those spaces and thus a constraint on solvability of the decision task.

In topological robotics, Farber initiated a line of work to describe the topological complexity of motion planning problems on a space in terms of the cohomology of that space [31]. This complexity reflects the discontinuities inherent to any controller that maps start and goal configurations of a robot to trajectories between those configurations.

An early example demonstrating the applicability specifically of algebraic topology to robotics, and indeed manipulation, appears in the work of Hopcroft and Wilfong [41]. They use the Mayer-Vietoris sequence on the zeroth and first homology groups to study contact connectivity. A key result: Under certain conditions on the structure of configuration space, if there is a motion of two objects starting and ending with the objects in contact, then there exists a motion throughout which the objects remain in contact.

1.5 Spheres as Topological Descriptors of Task Solvability

The current paper focuses on tasks that may be specified by goal states in some nondeterministic or stochastic graph; the task is to attain some goal state starting from anywhere within the graph. A key result (Theorem 1) shows that such a task has a guaranteed solution if and only if a certain simplicial complex associated with the task is homotopic to a sphere of a certain dimension. This special result leads to a general graph controllability theorem (Theorem 11) that characterizes the ability of a system to achieve *any* goal despite control uncertainty, again in terms of the existence of a certain sphere. Spheres and contractible spaces are much like the linear subspaces and trivial kernels in our earlier analogy.

These results are motivated by similar results describing the structure of complete directed graphs [8] and strongly connected directed graphs [44]. Indeed, our proof techniques build on the foundations of those two papers. Of additional interest is the extensive analysis in [47].

It is worth noting that the domain of nondeterministic graphs is much richer than that of directed graphs. It turns out (Section 13) that every finite simplicial complex can be realized via some nondeterministic graph. This correspondence further underscores the natural connection of topology to planning in the presence of uncertainty.

1.6 Contributions

The primary contribution of this paper is the previously unseen structure it reveals in planning problems. The paper shows how the details of a nondeterministic or stochastic graph may be abstracted away, leaving a purely topological description of the task: One may reason about task solvability by thinking in terms of spheres and contractible spaces.

A second contribution is the introduction of particular tools, such as *strategy complexes*. A single strategy on a graph is an eventually convergent nondeterministic control law for moving within some portion of the graph. A strategy complex consists of all possible strategies on a graph. Strategy complexes can be useful for reasoning about alternate strategies that a system might need if a selected strategy for accomplishing a task fails unexpectedly.

Third, although largely beyond the scope of this paper, a stream of ancillary results flows from our characterization of task solvability, among them: (i) One can answer purely topologically the question whether a collection of actions is essential to solving a task. (ii) One can recast various questions about tasks into measurements by the Euler characteristic. (iii) One discovers that all strategies for accomplishing a goal in a nondeterministic graph must overlap a particular strategy found by backchaining. (iv) One observes that the number of strategies for attaining a single goal state is either zero or odd (this is the topological equivalent of Yogi Berra’s direction: “When you come to a fork in the road, take it.”).

We anticipate yet more structure will be discovered in planning problems via the lens of topology. Indeed, every key topological idea we have looked at thus far has had some significant meaning when recast in a planning context. The Combinatorial Alexander Dual discussed in Section 10 is a good example.

1.7 An Underlying Motivation

A longstanding motivation for our research has been a desire to understand system capabilities and the interplay of sensing and action. That motivation is rooted in a long line of work on motion planning with uncertainty from the 1980s and 1990s [9, 59, 60, 22, 10, 14, 12, 72, 15, 16, 53, 33, 55, 11] as well as related work on understanding information invariants, such as natural tradeoffs between sensing and action [70, 21, 18, 2, 19, 52, 26, 17, 45, 46]. The tools for characterizing system capabilities discussed in Sections 9 and 10 provide a concise topological language for making such comparisons.

1.8 Broader Context

Applying the ideas in this paper to robotics problems requires front-end work: one needs to cast a robot task as a motion planning problem within some graph whose transitions are (potentially) uncertain. Much of robotics research over the past three decades has been devoted precisely to such reductions. We will not review those techniques here, but point to two excellent books: [53] and [54]. For instance, in manipulation, one generally needs to describe the geometry and mechanics of contact, including dealing with issues such as friction, deformation, and so forth. In many instances, these analyses are fairly local. One then is left with the task of combining the local information to form global motion planners. The current paper presents a topological perspective of that second phase.

The topological perspective likely has broader applications than the motivating tasks of robotic manipulation. Section 11.4 provides a small example. We hope that the techniques in this paper will be useful for analyzing planning and decision-making problems in a wide variety of domains, namely those with two modular phases: a front-end discretization phase, followed by a planning phase. In particular, methods such as those discussed in Sections 9–15 should be useful not only for planning but for analyzing and fine-tuning the details of the front-end phase. Potentially relevant additional domains include decisional architecture design, temporal

reasoning, and hybrid control systems.

1.9 Outline

Our basic approach is to replace the combinatorial structure of a discrete uncertain graph with the topology of a collection of open sets in Euclidean space. This technique is well-established for standard directed graphs [8, 44]. We show how to generalize the approach to uncertain systems. In particular, convergence results for Markov chains appear as intersection properties of certain affine open cones.

The resulting topology characterizes system capabilities. We indicate how to use such characterizations for designing desired systems, for analyzing existing systems, and for general reasoning about motion with uncertainty.

- Section 1 discusses core approaches to planning manipulation strategies and how these have motivated this paper’s topological inquiry.
- Section 2 provides intuition for understanding nondeterministic graphs and their strategy complexes, via a series of examples.
- Section 3 formally introduces nondeterministic graphs and their strategy complexes.
- Section 4 introduces loopback graphs and complexes, using these to characterize task solvability by the existence of topological spheres.
- Section 5 provides background on the topological tools appearing in this paper.
- Section 6 adds stochastic actions to the mix, now permitting both nondeterminism and stochasticity in a graph.
- Section 7 develops a connection between the combinatorial structure of plans and certain conical open sets in Euclidean space, useful for measuring a strategy’s convergence times.
- Section 8 characterizes full controllability on a graph, despite control uncertainty, in terms of the existence of a sphere of dimension two less than the number of states in the graph.
- Section 9 explores methods for designing systems using the topological tools developed thus far.
- Section 10 dualizes the earlier spherical characterizations, leading both to contractibility characterizations and two new simplicial complexes: the source complex and its dual complex. The source complex consists of all start regions of convergent plans, the dual complex consists of all potentially unattainable goals.
- Section 11 discusses internal structure: how to combine graphs, how to simplify graphs. An implemented example illustrates the ideas.
- Section 12 contains algorithms to implement the core ideas of the paper.
- Section 13 proves two realizability theorems, establishing two different ways in which nondeterministic graphs are just like finite simplicial complexes.

- Section 14 shows that ascertaining an uncertain system’s precision is *NP*-complete. The section also shows that two natural questions one might ask of graphs are undecidable, based on the undecidability of recognizing contractibility.
- Section 15 discusses several different ways in which topological thinking illuminates planning problems.

2 An Example: Nondeterminism, Cycles, and Strategies




Figure 1: Direct access might be blocked.




Figure 2: Nondeterministic graph and associated strategy complex for the motions of Fig. 1.

Legend: Figures in this paper depict nondeterministic actions in two ways: In a graph, as directed edges tied together by a circular arc, along with a “dot” inside the arc. In a strategy complex, by multiple outcomes to the right of an “arrow” ($\rightarrow$).

Imagine an ambulance rescue in an old complicated city, perhaps after an earthquake. There are many opportunities for nondeterminism: Entries into the city might be blocked, maps might be wrong, navigation might lead to circular paths. Let us focus on the final step in which the ambulance medics must pass through a narrow opening to reach their patient, as in Fig. 1. Something might go wrong, a collapse of some sort, forcing the rescuers to either side. Let us suppose the rescuers may then take additional steps around buildings bounding the original narrow opening to reach their patient.

We can model this scenario using the graph in the left panel of Fig. 2. The action to move from A (ambulance) to X (patient) might nondeterministically lead to X but perhaps also to B or C, depending on whether and how a collapse occurs on the direct path from A to X. In the example, there then are deterministic actions from either B or C to X. Such a graph is essentially a compressed AND/OR graph [3].

We can now represent the *strategy* (or control law or plan) just described as a solid triangle (right panel of Fig. 2). The vertices of the triangle are the individual actions to be executed at any particular location during the rescue operation. So, for instance, the strategy says “When at location B, execute the action $B \rightarrow X$.”




Figure 3: Actions that move between the far locations B and C.

Perhaps it is also possible to move from B to C, as in Fig. 3, and vice-versa. We can augment our graph to include these actions (left panel of Fig. 4)




Figure 4: With additional (potentially cycle-inducing) actions, the strategy complex now contains two solid tetrahedra.

Once we have the additional actions $B \rightarrow C$ and $C \rightarrow B$, we have to be careful *not* to include both in a strategy. Otherwise, the rescuers (who could be robots, not humans) might cycle forever between B and C, never reaching their patient at X. Instead, we recognize that there now exist several distinct strategies for reaching X. These can be represented by the two solid tetrahedra in the right panel of Fig. 4. Observe that the two tetrahedra intersect in a triangle that is our original triangle from Fig. 2, but there is no simplex that simultaneously includes the two actions $B \rightarrow C$ and $C \rightarrow B$.

Each tetrahedron represents a strategy (or control law or plan) consisting of actions that may be executed without accidentally creating cycles in the graph. Each triangle or edge or vertex of one of these tetrahedra, formed from a subset of the actions comprising the tetrahedron, also represents some strategy (perhaps with a different goal).

The semantics of the top tetrahedron of Fig. 4 are:

- When at A, execute the action $A \rightarrow B, X, C$.
- When at B, execute either the action $B \rightarrow X$ or the action $B \rightarrow C$.
It does not matter which; pick one, perhaps nondeterministically.
- When at C, execute the action $C \rightarrow X$.

Nondeterminism appears in both the outcomes and choices of the strategy:

1. Nature acts as an adversary during execution of the action $A \rightarrow B, X, C$, making the outcome uncertain. This can be bad. Fortunately, in the example, the system can compensate, by executing additional actions.
2. The system has available multiple actions at location B. This is good; it provides redundancy, alternate paths. The system can either leave the choice of which action to execute at state B open or it can select a particular action. The tetrahedron *in toto* leaves the choice open, effectively increasing nondeterminism, perhaps handing the choice over to an adversary. Alternatively, if the system chooses a particular action at state B, it is effectively picking a particular triangular face of the tetrahedron as its true strategy.

Terminology: In the remainder of the paper, we will generally speak of a “strategy” rather than a “plan”, in order to avoid the suggestion that strategies are created or discovered by planners. They exist independently of being created or discovered.

3 Nondeterministic Graphs and Strategy Complexes

This section makes precise the intuition of the previous section, focusing on nondeterministic graphs; Section 6 introduces stochasticity.

3.1 Basic Definitions

Nondeterministic Graphs

A *nondeterministic graph* $G = (V, \mathcal{A})$ is a set of *states* V and a collection of (*nondeterministic*) *actions* $\mathcal{A}$. Each $A \in \mathcal{A}$ is a nonempty set of *directed edges* $\{(v, u_1), (v, u_2), \dots\}$, with v and all u_i in V . We refer to v as A ’s *source* and to each u_i as a (*nondeterministic*) *target* of A . If A has a single target, A is also said to be *deterministic*.

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_i . If A has multiple targets, the precise target attained is not known ahead of time. One can imagine an adversary choosing the target.

Distinct actions may have overlapping or identical edge sets. (For instance, two different global motion commands might have the same effect at a given state. We could introduce extra notation to label actions with names, thereby making explicit their individual identity, but the extra notation would be more cumbersome than informative.)

All graphs, sets of states, actions, and collections of actions in this paper are finite.

Remark: Nondeterministic graphs in which each action is deterministic and no two actions have the same edge set are equivalent to standard directed graphs.

Terminology: We speak of a graph “state”, reserving the term “vertex” for singleton simplices in simplicial complexes.

Acyclic Graphs and Subgraphs

Suppose $G = (V, \mathcal{A})$ is a nondeterministic graph. A *possible path of length k* in G is a sequence of states $v_0, v_1, \dots, v_k$ in V such that $(v_i, v_{i+1}) \in A_i$, with each A_i an action in $\mathcal{A}$, for $i = 0, \dots, k-1$. G is *acyclic* if none of its possible paths have $v_0 = v_k$ with $k \geq 1$.

A *nondeterministic subgraph* $H = (W, \mathcal{B})$ of G is a nondeterministic graph in its own right such that $W \subseteq V$ and $\mathcal{B} \subseteq \mathcal{A}$. In particular, any $\mathcal{B} \subseteq \mathcal{A}$ defines a nondeterministic subgraph $H_{\mathcal{B}} = (V, \mathcal{B})$ of G . We will say that a collection of actions $\mathcal{B} \subseteq \mathcal{A}$ is *acyclic* if its induced subgraph $H_{\mathcal{B}}$ is acyclic.

Simplicial Complexes

An (*abstract*) *simplicial complex* Σ is a collection of finite sets, such that if σ is in Σ then so is every subset of σ [61]. The elements of Σ are called *simplices*; the elements of a simplex and singleton simplices are both called *vertices*. Sometimes one refers to the *underlying vertex set* of a simplicial complex Σ , which contains all the vertices of Σ and may contain vertices that could be present in Σ even if they happen not to be. The *dimension* of a simplex is one less than the number of its elements. If $\tau \subseteq \sigma$, with $\tau, \sigma \in \Sigma$, then one says that τ is a *face* of σ .

We permit the empty simplex $\emptyset$, for combinatorial simplicity [8, 47]. The complex $\{\emptyset\}$, consisting solely of the empty simplex, is the *empty complex*. It is also the sphere of dimension -1 . The complex $\emptyset$, consisting of no simplices, is the *void complex*.

If Σ is a simplicial complex, if Σ' is some subcollection of Σ , and if Σ' is a simplicial complex in its own right, then one says that Σ' is a *subcomplex* of Σ .

All simplicial complexes in this paper are finite. Any nonvoid finite simplicial complex has a geometric realization in some Euclidean space, with relative topology the same as its polytope topology [61]. Thus we may view Σ as a topological space.

See [61, 6, 39] for a further introduction to topology and simplicial complexes. See Section 5 for a summary of topological tools used in this paper.

Strategy Complexes

Given a nondeterministic graph $G = (V, \mathcal{A})$ with $V \neq \emptyset$, let Δ_G be the simplicial complex whose simplices are the acyclic collections of actions $\mathcal{B} \subseteq \mathcal{A}$. If $V = \emptyset$, let $\Delta_G = \emptyset$. We refer to Δ_G as G 's *strategy complex* and to every simplex in Δ_G as a (*nondeterministic*) *strategy*.

Observe that no action of G with a self-loop can appear in any simplex of Δ_G .

3.2 Examples

In what follows, we abbreviate any action $\{(v, u_1), \dots, (v, u_k)\}$ by writing $v \rightarrow u_1, \dots, u_k$.

The graph on the left of Fig. 5 is a standard directed graph. Each edge of the directed graph is a possible “action” the system could perform, moving it from some state of the graph to some other state.

The strategy complex of this graph is shown in the right panel of Fig. 5. The biggest simplex one could possibly expect to see in the strategy complex would be a tetrahedron, consisting of all four actions present in the directed graph. However, two of the actions, namely $1 \rightarrow 2$ and $2 \rightarrow 1$, could give rise to a cycle in the graph, so no simplex of the strategy complex can contain both these actions. The complex is in fact generated by two triangles.




Figure 5: The graph on the left defines the strategy complex shown on the right.

The two triangles, as well as three of the five edges in the complex, constitute strategies for attaining state 3 in the graph. The central edge, consisting of the actions $\{1 \rightarrow 3; 2 \rightarrow 3\}$, is the strategy one would obtain by backchaining from state 3 in a traditional fashion.

Observe that a strategy complex may contain strategies for a variety of goals. For instance, the top left edge of the complex in Fig. 5, comprising the set of actions $\{1 \rightarrow 3; 1 \rightarrow 2\}$, is a strategy that simply says “move *away* from state 1.”




Figure 6: The graph on the left has two nondeterministic actions that could create a cycle, so the strategy complex on the right consists of two isolated vertices.

For contrast, consider the graph of Fig. 6. It contains two actions, one each at states 1 and 2. Each action has two nondeterministic outcomes. The two actions cannot appear together as a simplex since, depending on the actual nondeterministic transitions at runtime, these actions could cause cycling in the graph between states 1 and 2. As a result, the strategy complex consists of two isolated vertices, representing the two strategies “move away from state 1” and “move away from state 2.” In particular, *there are no strategies guaranteed to attain state 3* from the other two states.




Figure 7: A strongly connected directed graph and its strategy complex. Compare with Fig. 8.




Figure 8: Another strongly connected directed graph and its strategy complex. See also Fig. 7.

It is also instructive to compare the two graphs and strategy complexes shown in Figures 7 and 8. Both graphs are strongly connected directed graphs with three states. The graphs are not isomorphic or even homomorphic, but their strategy complexes are both topological circles.

Indeed, Hultman [44] proved: Any directed graph that can be written as the disjoint union of its strongly connected components generates a strategy complex (he used a different name) topologically similar to a sphere of dimension $n - k - 1$, where n is the number of states in the graph and k is the number of strongly connected components. All other directed graphs produce contractible strategy complexes (roughly meaning: they can be shrunk to a point; see Section 5 for a precise definition).

This is a first hint that topology is capturing some significant graph property independent of the detailed structure of a graph. In the case of directed graphs, the strategy complexes always look either like spheres or contractible spaces (homotopically; see Section 5).

Remark: The nondeterministic setting is considerably richer than the deterministic setting of directed graphs. As we will see in Section 13, nondeterministic graphs are able to generate a much larger collection of topological spaces via their strategy complexes, namely all spaces describable by finite simplicial complexes. This means nondeterministic graphs and finite simplicial complexes are essentially identical topological objects. Consequently, one *should* use topology to study strategies for solving tasks in uncertain discrete spaces.

4 Loopback Graphs and Complexes

Now let us modify the graph of Fig. 5. We will add artificial deterministic transitions from state 3 to each of states 1 and 2. We call these added transitions *loopback actions*. Think of these loopback actions as “topological electrodes” that will allow us to measure whether the graph contains a guaranteed strategy for attaining state 3 from all states.




Figure 9: A loopback graph and loopback complex associated with the graph of Fig. 5. The complex contains 6 vertices, 12 edges, and 6 triangles (shaded). The two triangular endcaps outlined in dashed red are *not* part of the complex, since each gives rise to a cycle in the graph. The complex is homotopic to S^1 , the circle.

The left panel of Fig. 9 shows the resulting *loopback graph*. Now imagine constructing the strategy complex associated with that graph, as shown in the right panel of the figure. We refer to it as a *loopback complex* of our original graph. The complex in this case looks roughly like a polygonal cylinder. The complex is homotopic to a circle, as represented by either open end of the cylinder.

“Homotopic” means, in this case, that the complex, viewed as a topological space, can be continuously deformed within itself into a subspace that is topologically a circle. Notice that one cannot continuously deform the complex (within itself) into a point. This is crucial. Homotopy type is an equivalence relation on topological spaces. Circles and points lie in different homotopy equivalence classes. (See Section 5 for precise definitions.)




Figure 10: A loopback graph and loopback complex associated with the graph of Fig. 6. The complex is contractible.

In contrast, suppose we add both possible loopback actions at state 3 to the graph of Fig. 6. The resulting graph and loopback complex are shown in Fig. 10. Now the loopback complex is homotopic to a point; one can continuously deform it within itself to a point.

The Punch Line: No matter how complicated the nondeterministic graph, if we add all loopback actions to it that transition from some state $\mathbf{s}$ to the remaining states, then the resulting loopback complex will *always* be homotopic either to a sphere or to a point. A sphere tells us that there is a strategy guaranteed to attain state $\mathbf{s}$ from all states in the graph; a point tells us that no such strategy exists. We are beginning to see a topological language by which to characterize system capabilities.

The following definitions and theorem make the previous observation precise. We provide a proof of the theorem following its statement, both to preserve continuity and to build intuition. The reader may nonetheless first wish to review the topological tools discussed in Section 5. The foundations of our proof appear in [8, 44]. Subsequently, we will generalize these techniques further, namely to stochastic graphs.

Definitions Let $G = (V, \mathcal{A})$ be a nondeterministic graph and suppose $s \in V$ is some desired *stop state*. We make the following definitions:

- G contains a *complete guaranteed strategy for attaining s* if there is some acyclic set of actions $\mathcal{B} \subseteq \mathcal{A}$ such that $\mathcal{B}$ contains at least one action with source v for every $v \in V \setminus \{s\}$.

Observe that $\mathcal{B}$ cannot contain any actions with source s . Moreover, any possible path in the graph $(V, \mathcal{B})$, that terminates at some $v_k \neq s$, may be extended to a longer possible path. Iterating, this process converges at s , since $\mathcal{B}$ is acyclic.

We say: $\mathcal{B}$ is a *complete guaranteed strategy for attaining s* .

- Define $G_{\leftarrow s}$ to be the nondeterministic graph identical to G except that all actions with source s have been discarded, replaced instead by $(|V|-1)$ -many *loopback* actions $\{(s, v)\}$, each consisting of a single edge from s to some v , with v ranging over $V \setminus \{s\}$.

$G_{\leftarrow s}$ is a *loopback graph* of G .

- Define $\Delta_{G_{\leftarrow s}}$ to be the strategy complex associated with $G_{\leftarrow s}$.

$\Delta_{G_{\leftarrow s}}$ is a *loopback complex* of G .

- $[n]$ is (standard) shorthand for the set $\{1, \dots, n\}$.

Theorem 1 (Goal Attainability) Let $G = (V, \mathcal{A})$ be a nondeterministic graph and $s \in V$.

If G contains a complete guaranteed strategy for attaining s , then $\Delta_{G_{\leftarrow s}}$ is homotopic to the sphere S^{n-2} , with $n = |V|$. Otherwise, $\Delta_{G_{\leftarrow s}}$ is contractible.

Proof. We may assume $V = [n]$ and $s = n$. The theorem is trivially true for $n = 1$, so suppose $n > 1$.

I. Let $\mathcal{B}$ be a complete guaranteed strategy for attaining s and let $\mathcal{A}'$ be the actions of $G_{\leftarrow s}$. For each $A \in \mathcal{A}'$, define the open polyhedral cone $U_A = \bigcap_{(i,j) \in A} \{\mathbf{x} \in \mathbf{R}^n \mid x_i > x_j\}$. Observe that a set of actions $\{A_1, \dots, A_k\}$ is acyclic if and only if $U_{A_1} \cap \dots \cap U_{A_k}$ is not empty. When nonempty, the intersection is contractible. By the Nerve Lemma (see Section 5.3), $\Delta_{G_{\leftarrow s}}$ therefore has the homotopy type of $\bigcup_{A \in \mathcal{A}'} U_A$. We claim that this union covers all of $\mathbf{R}^n$ except for the line on which all coordinates are equal. Thus it is homotopic to S^{n-2} .

To see coverage: Clearly no point with all coordinates equal can be in the union. The cones determined by the loopback actions cover all points $\mathbf{x} \in \mathbf{R}^n$ for which $x_n > x_i$, some i .

Suppose some $\mathbf{x}$ in $\mathbf{R}^n \setminus \{x_1 = \dots = x_n\}$ is not inside any U_A . Then $x_i \geq x_n$ for all i , with at least one $x_i > x_n$. Some action $B \in \mathcal{B}$ has that i as a source. $\mathcal{B} \subseteq \mathcal{A}'$, so $\mathbf{x} \notin U_B$, meaning there is some target j of B such that $x_j \geq x_i > x_n$. Repeating this argument with j , etc., produces an arbitrarily long and thus cyclic possible path in $H_{\mathcal{B}}$. Contradiction.

II. If G does not contain a complete guaranteed strategy for attaining s , then no simplex of $\Delta_{G \leftarrow s}$ contains actions at all states of $V \setminus \{s\}$. For every simplex $\sigma \in \Delta_{G \leftarrow s}$ there is therefore a unique nonempty maximal set τ_σ of loopback actions such that $\sigma \cup \tau_\sigma \in \Delta_{G \leftarrow s}$. A standard collapsing argument now shows that $\Delta_{G \leftarrow s}$ is contractible (see Section 5.2). $\square$

5 Topological Tools and Homotopy Equivalence

One of the most important functions of topology is to recognize equivalences between seemingly different objects. This section reviews some of the key topological equivalences and tools for establishing those equivalences. As suggested by Theorem 1, we will use these tools to characterize system capabilities by assigning topological structures to nondeterministic (and stochastic) graphs.

Two simplicial complexes, Γ and Σ , are said to be *isomorphic*, written $\Gamma \cong \Sigma$, if there is a bijective correspondence between the vertices of the two complexes given by some function f , such that $\{v_1, \dots, v_k\}$ is a simplex of Γ if and only if $\{f(v_1), \dots, f(v_k)\}$ is a simplex of Σ . In particular, f preserves simplex dimension.

Two topological spaces, X and Y , are said to be *homeomorphic*, written $X \approx Y$, if there exist two continuous functions, $f : X \rightarrow Y$ and $g : Y \rightarrow X$, such that $g \circ f = \text{id}_X$ and $f \circ g = \text{id}_Y$, where id_X is the identity function on X and id_Y is the identity function on Y .

Isomorphic simplicial complexes are homeomorphic when viewed as topological spaces.

$\approx$ is an equivalence relation in the category of topological spaces. It is perhaps the most familiar to the reader. Two homeomorphic spaces really are very much the same intuitively, differing “only” by some continuous transformation with a continuous inverse.

There is a weaker equivalence relation for topological spaces, based on homotopy. This relation is more akin to the notion of “morphing” similarity one sees in computer graphics. The equivalence classes one obtains under homotopy turn out to be useful descriptors for planning in the presence of uncertainty. We can think of these equivalence classes as constituting major nouns of the language we desired in Section 1. We give the definitions next and then some tools for establishing homotopy equivalence.

Suppose $f_0 : X \rightarrow Y$ and $f_1 : X \rightarrow Y$ are two continuous functions between topological spaces X and Y . One says that f_0 is *homotopic* to f_1 , written $f_0 \simeq f_1$, if there exists a continuous function $F : X \times [0, 1] \rightarrow Y$ such that $F(x, 0) = f_0(x)$ and $F(x, 1) = f_1(x)$ for all $x \in X$. (Here $[0, 1]$ is the unit interval of the real line.) One can think of $F(\cdot, t)$ as a continuous “morphing” of f_0 into f_1 as t varies from 0 to 1.

Two topological spaces, X and Y , are said to be *homotopy equivalent* (or to *have the same homotopy type* or to be *homotopic*), written $X \simeq Y$, if there exist two continuous functions, $f : X \rightarrow Y$ and $g : Y \rightarrow X$, such that $g \circ f \simeq \text{id}_X$ and $f \circ g \simeq \text{id}_Y$. Observe that the difference between “homeomorphic” and “homotopic” is the difference between compositions that are exactly equal to the identity versus merely homotopic to the identity. Thus homeomorphic spaces are also homotopic.

A topological space homotopic to a single point is said to be *contractible*.

It is often difficult to establish homotopy equivalence based on these definitions alone, but there are several key topological tools that one may use instead. We discuss these next. Once again, we also refer the reader to [61, 39, 6, 8] for more extensive treatments.

5.1 Deformation Retractions

Suppose A is a subspace of a topological space X . A (*strong*) *deformation retraction* of X onto A is a continuous function $F : X \times [0, 1] \rightarrow X$ such that:

- (a) $F(x, 0) = x$ for all $x \in X$,
- (b) $F(x, 1) \in A$ for all $x \in X$, and
- (c) $F(a, t) = a$ for all $a \in A$ and all $t \in [0, 1]$.

F establishes that $X \simeq A$ [61] and one says that A is a *deformation retract* of X . More generally, it is a fact that two spaces are homotopy equivalent precisely when each may be viewed as a deformation retract of some common encompassing space [39].

We encourage the reader to verify the following facts, as warmup for ideas to come:

- (1) Let X be all of n -dimensional Euclidean space except for the origin.
Then X is homotopy equivalent to S^{n-1} , the sphere of dimension $n - 1$.
- (2) Let $X = \mathbf{R}^n \setminus \{\mathbf{x} \in \mathbf{R}^n \mid x_1 = \cdots = x_n\}$, that is, n -dimensional Euclidean space with a line removed. Then X is homotopy equivalent to S^{n-2} .
(The proof of Theorem 1 uses this fact.)

5.2 Collapsibility and Contractibility

Suppose Σ is a simplicial complex and suppose τ and σ are simplices in Σ such that τ is a proper face of σ . If τ is a proper face of *no other* simplex in Σ , then one can remove both τ and σ from Σ to obtain a new complex $\Sigma' = \Sigma \setminus \{\tau, \sigma\}$ that has the same homotopy type as Σ . The process of constructing Σ' from Σ is called an *elementary collapse*. The reverse process is called an *elementary anti-collapse*. [6]

It is a fact that a finite simplicial complex is contractible if and only if there is a sequence of elementary collapses and elementary anti-collapses that transforms the complex into a single point. Finding such a sequence is an uncomputable problem (this goes back to a deep result that the word problem for groups is undecidable [63]). The special case in which only elementary collapses are needed is computable; one can try all possibilities. Such complexes are called *collapsible*. A classic example of a space that is contractible but not collapsible is the “House with Two Rooms” [39].

A finite simplicial complex Σ is a *cone* if there is some vertex v of Σ such that $\sigma \cup \{v\} \in \Sigma$ whenever $\sigma \in \Sigma$. In this case, v is an *apex* of the cone. A cone is a classic example of a collapsible complex.

The following very useful fact appears as Lemma 7.6 in [8] (the collapsibility argument in the proof of Theorem 1 uses this lemma):

Notation: $\sigma \pm v$ means $\sigma \setminus \{v\}$ if $v \in \sigma$ and $\sigma \cup \{v\}$ if $v \notin \sigma$.

Lemma 2 (Björner and Welker, [8]) *Suppose Σ' is a subcomplex of a finite simplicial complex Σ and suppose for some vertex v of Σ , the collection $\Sigma \setminus \Sigma'$ is closed under the map $\sigma \mapsto \sigma \pm v$. Then Σ collapses to Σ' (via a sequence of elementary collapses) and thus $\Sigma \simeq \Sigma'$.*

5.3 The Nerve Lemma

Suppose $\mathcal{U}$ is some collection of sets (not necessarily distinct). One may define a simplicial complex called the *nerve of $\mathcal{U}$* , written $\mathcal{N}(\mathcal{U})$, as follows: The simplices of $\mathcal{N}(\mathcal{U})$ are given by the empty simplex and all nonvoid finite subcollections $\{U_1, \dots, U_k\}$ of $\mathcal{U}$ such that the intersection $U_1 \cap \dots \cap U_k$ is not empty. See [61, 39].

Lemma 3 (Nerve Lemma, [39]) *Let X be a paracompact topological space. Suppose $\mathcal{U}$ is a collection of open subsets of X whose union covers X , such that the intersection of any nonvoid finite subcollection of sets in $\mathcal{U}$ is contractible whenever it is nonempty. Then $X \simeq \mathcal{N}(\mathcal{U})$.*

The reader can look up the term “paracompact”. See [50, 20, 39]. It is enough for our purposes to know that any topological subspace of Euclidean space is paracompact.

The Nerve Lemma infers global homotopy type from local contractibility. For instance, in Theorem 1, we associated to every action of a nondeterministic graph an open set. The Nerve Lemma allowed us to infer the overall topology of the graph’s strategy complex by the intersection properties of these open sets. In Section 7, we will associate open sets with actions in both nondeterministic and stochastic graphs so as to encode the actions’ convergence properties. Again, the Nerve Lemma will allow us to infer the topology of the graphs’ strategy complexes.

5.4 The Quillen Fiber Lemma

Another very useful tool is the Quillen Fiber Lemma. We will state it for partially ordered sets (known as *posets*). Every simplicial complex Σ defines a poset $\mathcal{F}(\Sigma)$, called the *face poset* of Σ . The elements of the face poset $\mathcal{F}(\Sigma)$ are the nonempty simplices of the complex Σ , partially ordered by set inclusion. One can also construct a simplicial complex $\Sigma(P)$, called the *order complex*, from any poset P . The simplices of $\Sigma(P)$ are given by the finite chains $p_1 < \dots < p_k$ in P .

It is a fact that Σ and $\Sigma(\mathcal{F}(\Sigma))$ are homeomorphic. Indeed, $\Sigma(\mathcal{F}(\Sigma))$ is the first *barycentric subdivision* of Σ , usually written as $\text{sd}(\Sigma)$. Abstractly, $\text{sd}(\Sigma)$ is a new complex, whose nonempty simplices are given by all sets of the form $\{\sigma_1, \sigma_2, \dots, \sigma_k\}$, with each σ_i a nonempty simplex of Σ , and with σ_i a proper face of σ_{i+1} , for $i = 1, \dots, k-1$ [69]. Geometrically, the first barycentric subdivision is a re-triangulation obtained by adding as vertices the centroids (called *barycenters*) of all simplices in Σ , then defining new simplices accordingly [71, 61].

Thus posets and simplicial complexes are essentially identical topological objects. For instance, one may speak of the topology of a poset, implicitly meaning the topology of its order complex. We will not elaborate on this connection further, but refer the reader to [8, 6, 75].

Notation: If Q is a poset, then $Q_{\leq q}$ denotes the set $\{q' \in Q \mid q' \leq_Q q\}$, where $\leq_Q$ is the partial order on Q (set inclusion in the case of face posets derived from simplicial complexes).

Lemma 4 (Quillen, [68]) Suppose $f : P \rightarrow Q$ 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.

5.5 Homotopy Interpretation

It is still a research question to determine fully what the homotopy type of a graph's strategy complex really means. Much of the rest of the paper will be devoted to understanding that meaning for some important special cases.




Figure 11: The deterministic action $1 \rightarrow 2$ is more precise than the nondeterministic action $1 \rightarrow 2, 3$. The strategy complex makes this explicit, by showing how simplices containing action $1 \rightarrow 2, 3$ can collapse away while preserving homotopy type.

There is however something very simple we can observe. Suppose $G = (V, \mathcal{A})$ is a nondeterministic graph. Suppose G contains two distinct actions $A \in \mathcal{A}$ and $B \in \mathcal{A}$ with identical source states such that one action is more precise than the other. Viewed as edge sets, suppose $A \subseteq B$. Intuitively, we would expect never to need action B . After all, anything we can be certain of doing with action B , we can also be certain of doing with action A .

Homotopy equivalence captures this observation via collapsibility. Figure Fig. 11 provides an example.

In particular, suppose $B \in \sigma \in \Delta_G$. Let $\tau = \sigma \pm A$. Then $B \in \tau \in \Delta_G$ as well. So Lemma 2 tells us that the collection $\Delta^- = \{\sigma \in \Delta_G \mid B \notin \sigma\}$ is homotopy equivalent to Δ_G . In other words, action B is irrelevant. Observe as well that $\Delta^- = \Delta_{G^-}$, where G^- is the nondeterministic graph identical to G except that action B has been removed.

6 Stochastic Graphs and Strategy Complexes

6.1 Stochastic Actions and Graphs

Stochastic Actions

Section 3 defined a nondeterministic action to be a nonempty set of directed edges with common source. We now define a *stochastic action* to be a nonempty set of directed edges with

common source in which each edge is labeled with a transition probability. Formally, we write a stochastic action A as a nonempty set of labeled pairs in the form $\{(v, p_1 u_1), (v, p_2 u_2), \dots\}$, with v and all u_i elements of some underlying state space V . As before, v is A 's *source* and each u_i is a (*stochastic*) *target* of A . Each label p_i is a *transition probability*.

The semantics of a stochastic action are Markovian: 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_i , selected from all of A 's targets with probability p_i . We require that each $p_i > 0$ and that $\sum_i p_i = 1$.

Stochastic Graphs

A *stochastic graph* $G = (V, \mathcal{A})$ is a set of *states* V and a collection of *actions* $\mathcal{A}$ whose sources and targets all lie in V . V is also known as G 's *state space*. An action may be either nondeterministic or stochastic. As before, distinct actions may have overlapping or identical edge sets or even identical sets of transition probabilities. And, again, all graphs, sets of states, actions, and collections of actions in this paper are finite.

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

Remarks:

- **To emphasize:** We allow both nondeterministic and stochastic actions in a stochastic graph.
- **Determinism:** We may view a deterministic action that transitions from state v to state u as a special case of either a nondeterministic action, $\{(v, u)\}$, or a stochastic action, $\{(v, 1u)\}$. In this paper, it does not matter which.
- **Notation:** In figures of graphs, we label the edges of a stochastic action with transition probabilities. In figures of strategy complexes, we indicate the vertex corresponding to a stochastic action $A = \{(v, p_1 u_1), \dots, (v, p_k u_k)\}$ with the label $v \rightarrow p_1 u_1, \dots, p_k u_k$. We sometimes use this notation in the text as well.

6.2 Stochastic Acyclicity




Figure 12: The graph on the left has two stochastic actions that together could create a cycle but must eventually converge to state 3. Consequently, the strategy complex on the right includes not only the individual actions but the full 1-simplex formed by these two actions.

The two nondeterministic actions of the graph in Fig. 6 did not appear together in a strategy simplex since together they could lead to infinite looping. Now imagine that the actions are merely stochastic, as in Fig. 12. Intuitively, we probably *do* want to permit both these actions together in a strategy simplex. Although such a strategy might lead to cycling between states 1 and 2, the cycling would be stochastic and therefore not last forever. Eventually the system would converge to state 3. Indeed, one may even compute the expected convergence times.

We therefore need to modify our definition of acyclicity to account for stochastic convergence. There are two natural approaches, both grounded in the idea that an adversary might try to choose actions and transitions in such a way as to keep the system cycling for as long as possible.

Definition: Given a collection of actions $\mathcal{B}$, let $V_{\mathcal{B}} = \{v \mid v \text{ is the source of some } B \in \mathcal{B}\}$. We refer to $V_{\mathcal{B}}$ as the *start region* of $\mathcal{B}$.

Markov Chain Perspective: Suppose $G = (V, \mathcal{A})$ is a stochastic graph (so it may contain either nondeterministic or stochastic actions or both). Let W be some subset of $V_{\mathcal{A}}$. Now imagine an adversary who constructs a Markov chain M as follows:

- For every $v \in W$, the adversary selects some action $A \in \mathcal{A}$ whose source is v .
 - If the action A is stochastic, that is, $A = \{(v, p_1 u_1), \dots, (v, p_k u_k)\}$, then the stochastic transitions of M at v are exactly those given by A .
 - If the action A is nondeterministic, that is, $A = \{(v, u_1), \dots, (v, u_k)\}$, then the adversary further selects one target u_i of A . There will be exactly one transition of M at v , given by an edge from v to u_i , occurring with probability 1.
- For every $v \in V \setminus W$, there is a single transition of M at v , consisting of a self-loop with probability 1.

We refer to such a construction as a *Markov chain M with support W* and say G *contains* (M, W) . Observe that in general G may contain many different (M, W) .

Convergence Time Perspective: Suppose $G = (V, \mathcal{A})$ is a stochastic graph. Associate to every action $A \in \mathcal{A}$ a nonnegative transition time δ_A . Consider the following system of equations:

$$t_v = \max \left(\begin{array}{cc} \max_{\substack{A \in \mathcal{A} \\ A = \{(v, u_j)\}}} \left(\max_j t_{u_j} + \delta_A \right), & \max_{\substack{A \in \mathcal{A} \\ A = \{(v, p_j u_j)\}}} \left(\sum_j p_j t_{u_j} + \delta_A \right) \end{array} \right), \text{ if } v \in V_{\mathcal{A}};$$

$$t_v = 0, \quad \text{if } v \notin V_{\mathcal{A}}.$$
(1)

The reader may recognize System (1) as describing a maximization over expected durations of random walks with variable step times [32], induced by Markov chains $(M, V_{\mathcal{A}})$ contained in G . It is a form of Bellman's equation, representing an optimization from an imagined adversary's perspective, who is trying to maximize the graph's convergence times.

The convergence times are 0 for all states at which there are no actions. Otherwise, the maximizations appearing in System (1) describe adversarial choices. At any state v , the adversary maximizes over all actions A with source v . If A is a nondeterministic action, the adversary performs an additional maximization over the convergence times of all targets of A ; if A is stochastic, there appears an expectation over the convergence times of the target states.

Reminder: A *recurrent class* of a Markov chain M is a set of states R of M such that the probability of reaching any state of R from any other state of R is 1, while the probability of reaching any state outside R is 0. The restriction of M to R induced by transitions at states of R therefore itself defines a Markov chain. If R consists of a single state, that state is called *absorbing*. [32, 48]

The following lemma establishes an equivalence between the Markov chain perspective and the convergence time perspective:

Lemma 5 (Stochastic Acyclicity) *Let $G = (V, \mathcal{A})$ be a stochastic graph with associated nonnegative action transition times $\{\delta_A\}_{A \in \mathcal{A}}$.*

System (1) has a unique finite solution if and only if the only recurrent classes of any Markov chain with support W contained in G are formed by the absorbing states $V \setminus W$.

Moreover, when the solution is unique and finite, it is nonnegative, that is, $t_v \geq 0$ for all $v \in V$.

Proof. We omit the details of the proof. The basic techniques are similar to those used in Markov Decision Processes. We point to [32, 48, 78]. $\square$

The previous lemma permits us to move back and forth between the Markov chain and convergence time perspectives. Often it is best to reason directly about the recurrent classes induced by a stochastic graph but easier to compute convergence times using System (1). In particular, when System (1) has a unique finite solution one can use the system in an iterative fashion to obtain that solution.

6.3 Stochastic Strategy Complexes

We now make some further definitions, leading to strategy complexes in the stochastic setting.

Stochastically Acyclic Collections of Actions

Suppose $G = (V, \mathcal{A})$ is a stochastic graph. G is *stochastically acyclic* if System (1) has a unique finite solution for some (and thus any) set of nonnegative action transition times $\{\delta_A\}_{A \in \mathcal{A}}$. Similarly, we say that a collection of actions $\mathcal{B} \subseteq \mathcal{A}$ is *stochastically acyclic* if the induced subgraph $H_{\mathcal{B}} = (V, \mathcal{B})$ is stochastically acyclic.

When $\mathcal{B}$ is stochastically acyclic, we may view $\mathcal{B}$ as a strategy for moving the system into the complement of $\mathcal{B}$'s start region $V_{\mathcal{B}}$: If the initial state of the system lies inside $V_{\mathcal{B}}$, then moving under actions of $\mathcal{B}$, the system will eventually stop at some state inside $V \setminus V_{\mathcal{B}}$. If the system initially starts in $V \setminus V_{\mathcal{B}}$, then it remains there. Given specific nonnegative action transition times $\{\delta_B\}_{B \in \mathcal{B}}$, we refer to the solution of System (1), written out for $H_{\mathcal{B}}$ with those transition times, as *the worst-case expected convergence times* of $\mathcal{B}$.

Stochastic Strategy Complexes

Given a stochastic graph $G = (V, \mathcal{A})$ with $V \neq \emptyset$, let Δ_G be the simplicial complex whose simplices are the stochastically acyclic collections of actions $\mathcal{B} \subseteq \mathcal{A}$. If $V = \emptyset$, let $\Delta_G = \emptyset$.

This definition is identical to the one we gave earlier for nondeterministic graphs, but our notion of “acyclic collections of actions” has now been enlarged to include stochastic actions. Again, we refer to Δ_G as G 's *strategy complex* and to every simplex in Δ_G as a (*stochastic*) *strategy*.

Time-Bounded Strategy Complexes

System (1) allows us to define a tower of strategy complexes for any graph $G = (V, \mathcal{A})$.

Suppose we associate nonnegative transition times $\{\delta_A\}_{A \in \mathcal{A}}$ to the actions of G . To every $\sigma \in \Delta_G$ we can then associate a *maximal worst-case expected convergence time*, $t_{\max}(\sigma)$, defined to be the maximum time t_v obtained as a solution to System (1) when written out for the graph (V, σ) .

Let $T \geq 0$ be given. Define $\Delta_G^T = \{\sigma \in \Delta_G \mid t_{\max}(\sigma) \leq T\}$. Then Δ_G^T is a subcomplex of Δ_G , representing all strategies whose maximal worst-case expected convergence times are no greater than T . (Exercise for the reader: Prove that removing actions from a simplex σ cannot raise $t_{\max}(\sigma)$, as is required to infer that Δ_G^T is a simplicial complex.)

7 Covering Sets

We wish to understand the topology of strategy complexes and thus the capabilities of nondeterministic and stochastic graphs. Let us adopt a technique that appeared in the proof of Theorem 1 and associate to each action an open set. The topology of the resulting collection of open sets will be the topology of a graph's strategy complex.

7.1 Homogeneous Covering Sets

Definition: Suppose $G = (V, \mathcal{A})$ is a stochastic graph with $V \neq \emptyset$. We can assume for simplicity that $V = [n]$, with $n \geq 1$. We associate to each $A \in \mathcal{A}$ an homogeneous open subset U_A of $\mathbf{R}^n$, which we refer to as a *covering set*:

- If the action A is stochastic, that is, $A = \{(i, p_j j)\}$ for some set of targets $\{j\}$, then

$$U_A = \left\{ \mathbf{x} \in \mathbf{R}^n \mid x_i > \sum_j p_j x_j \right\}.$$

- If the action A is nondeterministic, that is, $A = \{(i, j)\}$ for some set of targets $\{j\}$, then

$$U_A = \bigcap_j \{ \mathbf{x} \in \mathbf{R}^n \mid x_i > x_j \}.$$

In the stochastic case, the open set U_A is a halfspace of $\mathbf{R}^n$ whose defining hyperplane has a normal determined by action A 's transition probabilities. In the nondeterministic case, the open set U_A is the intersection of several such halfspaces, one for each possible nondeterministic target. Observe that the defining hyperplanes all contain the line $\{ \mathbf{x} \in \mathbf{R}^n \mid x_1 = \dots = x_n \}$.

Remark: The collection of all hyperplanes of the form $\{ \mathbf{x} \in \mathbf{R}^n \mid x_i = x_j \}$ is known classically as the *Type A braid arrangement in $\mathbf{R}^n$* [75]; it is very useful for studying the poset of all posets on n items [8]. It should therefore come as no surprise that the open sets $\{U_A\}_{A \in \mathcal{A}}$ will allow us to infer the topology of G 's strategy complex. In particular, for a nondeterministic graph G , we are simply looking at a subposet of that overall poset of posets on n items. For a stochastic graph G , we have effectively created a poset of *stochastic partial orders*, each of which we may think of as a collection of Markov chains all of whose states are either transient or trivially absorbing (that was the gist of Lemma 5).

7.2 Affine Covering Sets

Definition: In order to study the topology of time-bounded strategy complexes, it is useful to define affine covering sets. Suppose we associate nonnegative transition times $\{\delta_A\}_{A \in \mathcal{A}}$ to the actions of G . Then we can define for each $A \in \mathcal{A}$ an affine open subset U_A^+ of $\mathbf{R}^n$:

- If $A = \{(i, p_j j)\}$ is stochastic, then $U_A^+ = \left\{ \mathbf{x} \in \mathbf{R}^n \mid x_i > \sum_j p_j x_j + \delta_A \right\}$.
- If $A = \{(i, j)\}$ is nondeterministic, then $U_A^+ = \bigcap_j \{ \mathbf{x} \in \mathbf{R}^n \mid x_i > x_j + \delta_A \}$.

7.3 Inferring Topology from Covering Sets

The following lemma shows that Δ_G and $\mathcal{N}(\{U_A\}_{A \in \mathcal{A}})$ are isomorphic simplicial complexes (and therefore homeomorphic and homotopic). The lemma generalizes to stochastic graphs a statement that appeared for nondeterministic graphs early in the proof of Theorem 1.

Lemma 6 (Homogeneous Nerve) *Let $G = (V, \mathcal{A})$ be a stochastic graph with $V = [n]$, $n > 0$. Suppose $\emptyset \neq \mathcal{B} \subseteq \mathcal{A}$. Then:*

$$\bigcap_{B \in \mathcal{B}} U_B \neq \emptyset \quad \text{if and only if} \quad \mathcal{B} \text{ is stochastically acyclic.}$$

Proof. Recall that $H_{\mathcal{B}}$ means the subgraph $(V, \mathcal{B})$ of G .

I. Suppose $\bigcap_{B \in \mathcal{B}} U_B \neq \emptyset$.

Choose $\mathbf{x}^* \in \mathbf{R}^n$ so that $\mathbf{x}^* \in U_B$ for all $B \in \mathcal{B}$. Define δ_B for each $B \in \mathcal{B}$ as follows:

If $B = \{(i, p_j j)\}$ is stochastic, let $\delta_B = x_i^* - \sum_j p_j x_j^*$.

If $B = \{(i, j)\}$ is nondeterministic, let $\delta_B = x_i^* - \max_j(x_j^*)$.

Observe that each $\delta_B > 0$.

Now consider the system of equations:

$$t_i = \max \left(\begin{array}{cc} \max_{\substack{B \in \mathcal{B} \\ B = \{(i, j)\}}} \left(\max_j t_j + \delta_B \right), & \max_{\substack{B \in \mathcal{B} \\ B = \{(i, p_j j)\}}} \left(\sum_j p_j t_j + \delta_B \right) \end{array} \right), \text{ if } i \in V_{\mathcal{B}}; \quad (2)$$

$$t_i = x_i^*, \quad \text{if } i \notin V_{\mathcal{B}}.$$

By construction, System (2) has at least one finite solution, given by $\mathbf{x}^*$, that is, $t_i = x_i^*$, for all $i \in V$. (In fact, the solution holds for *all* actions $B \in \mathcal{B}$, that is, the maximum at state i in System (2) occurs for every action of $\mathcal{B}$ with source i .)

Now suppose that $H_{\mathcal{B}}$ contains a Markov chain (M, W) whose support W is a recurrent class. Let $P = (p_{ij})$ be the probability transition matrix of M restricted to W . P is a stochastic matrix in its own right since W is a recurrent class. The transition probabilities p_{ij} are determined from actions of $\mathcal{B}$ via the process outlined on page 22. In particular, for every $i \in W$, some action $B \in \mathcal{B}$ gives rise to M 's transitions at state i . Define $\delta_i = \delta_B$.

Combining this construction with System (2), we see that:

$$x_i^* \geq \sum_{j \in W} p_{ij} x_j^* + \delta_i, \quad \text{for all } i \in W.$$

That is only possible if $\delta_i \leq 0$ for some $i \in W$, since P is a stochastic matrix. Contradiction. So, by Lemma 5, $\mathcal{B}$ must be stochastically acyclic.

II. Suppose $\mathcal{B}$ is stochastically acyclic.

For each $B \in \mathcal{B}$, let $\delta_B > 0$ be arbitrary. Now write out System (1) for H_B :

$$t_i = \max \left(\begin{array}{cc} \max_{\substack{B \in \mathcal{B} \\ B = \{(i,j)\}}} \left(\max_j t_j + \delta_B \right), & \max_{\substack{B \in \mathcal{B} \\ B = \{(i,p_j j)\}}} \left(\sum_j p_j t_j + \delta_B \right) \end{array} \right), \text{ if } i \in V_B;$$

$$t_i = 0, \quad \text{if } i \notin V_B.$$

By assumption, this system has a unique finite solution, call it $\mathbf{t}^*$.

Now consider $B \in \mathcal{B}$:

If $B = \{(i, p_j j)\}$ is stochastic, then $t_i^* \geq \sum_j p_j t_j^* + \delta_B > \sum_j p_j t_j^*$.

If $B = \{(i, j)\}$ is nondeterministic, then $t_i^* \geq \max_j (t_j^*) + \delta_B > \max_j (t_j^*)$.

So $\mathbf{t}^* \in U_B$ for all $B \in \mathcal{B}$, establishing that $\bigcap_{B \in \mathcal{B}} U_B \neq \emptyset$. $\square$

In order to obtain the topology of time-bounded strategy complexes we need to be a little more careful. There may exist gaps in the affine covering sets near the line on which all coordinates are equal. We will therefore intersect the covering sets with the boundary of a polyhedral cylinder designed to measure convergence times.

Definition: Given real $r > 0$, let $C_r = \{\mathbf{x} \in \mathbf{R}^n \mid |x_i - x_j| < r \text{ for all } i, j \in [n]\}$ and let ∂C_r be the boundary of C_r .

Observe: ∂C_r is a polyhedral cylinder with axis given by the line $\{\mathbf{x} \in \mathbf{R}^n \mid x_1 = \dots = x_n\}$ and with r the cylinder's "radius".

Lemma 7 (Affine Nerve) *Let $G = (V, \mathcal{A})$ be a stochastic graph with associated nonnegative action transition times $\{\delta_A\}_{A \in \mathcal{A}}$. Assume $V = [n]$ and $n > 0$. Let $T \geq 0$ be given.*

There exists $\epsilon_T > 0$, such that for every $0 < \epsilon < \epsilon_T$:

For every $\emptyset \neq \mathcal{B} \subseteq \mathcal{A}$, $\mathcal{B}$ is stochastically acyclic with $t_{\max}(\mathcal{B}) \leq T$ if and only if

$$\bigcap_{B \in \mathcal{B}} U_B^+ \cap \partial C_{T+\epsilon} \neq \emptyset.$$

We omit the proof, except to note that the existence of ϵ_T in the proof depends on finiteness of G . We now move directly to the key theorem that describes the topology of strategy complexes in terms of the topology of covering sets in $\mathbf{R}^n$:

Theorem 8 (Cover Homotopy) *Let $G = (V, \mathcal{A})$ be a stochastic graph with $V = [n]$, $n > 0$.*

$$\text{Then} \quad \Delta_G \simeq \bigcup_{A \in \mathcal{A}} U_A.$$

Let the action transition times of G be given by nonnegative numbers $\{\delta_A\}_{A \in \mathcal{A}}$ and let $T \geq 0$.

$$\text{Then} \quad \Delta_G^T \simeq \bigcup_{A \in \mathcal{A}} U_A^+ \cap \partial C_{T+\epsilon},$$

with $0 < \epsilon < \epsilon_T$ and ϵ_T given as per Lemma 7.

Proof. First, observe that whenever a set of the form $U_{A_1} \cap \cdots \cap U_{A_k}$ is nonempty, then it is convex hence contractible. We therefore obtain

$$\Delta_G \cong \mathcal{N}(\{U_A\}_{A \in \mathcal{A}}) \simeq \bigcup_{A \in \mathcal{A}} U_A.$$

The isomorphism $\cong$ follows from Lemma 6 and the homotopy equivalence $\simeq$ follows from the Nerve Lemma.

Second, observe that whenever a set of the form $U_{A_1} \cap \cdots \cap U_{A_k} \cap \partial C_{T+\epsilon}$ is nonempty, then, while it may not be convex, it is the deformation retract of a convex set, so is contractible. Consequently, we may apply Lemma 7 and the Nerve Lemma to conclude

$$\Delta_G^T \cong \mathcal{N}\left(\left\{U_A^+ \cap \partial C_{T+\epsilon}\right\}_{A \in \mathcal{A}}\right) \simeq \bigcup_{A \in \mathcal{A}} U_A^+ \cap \partial C_{T+\epsilon}. \quad \square$$




Figure 13: The graph on the left has two stochastic actions (at states 1 and 2) and one nondeterministic action (at state 3). Its strategy complex, shown on the right, is the boundary of a triangle.

As an example, let G be the graph of Fig. 13. The graph has three states and three actions, one at each state. Two of the actions are stochastic and one is nondeterministic. Δ_G is the boundary complex of a triangle.

Fig. 14 shows the previous lemmas and theorem in action, depicting the covering sets $\{U_1^+, U_2^+, U_3^+\}$ associated with the three actions of G , assuming all action transition times




Figure 14: This figure shows a two-dimensional slice of the covering sets associated with the actions of the graph in Fig. 13, along with intersections of these covering sets. The slice describes the (x_1, x_2) -shape of the covering sets at $x_3 = 1$. The figure also shows ∂C_r (drawn in yellow) for various critical radii r (roughly), with the critical events circled. For reference, the graph's strategy complex appears in the upper left corner.

are 1 and assuming $p' > p > \frac{2}{3}$. The figure actually shows the (x_1, x_2) -slice of the cover $\{U_1^+, U_2^+, U_3^+\}$ at $x_3 = 1$. Since the covering sets are invariant with respect to a translation along the direction $(1, 1, 1)$, it is enough to look at such two-dimensional slices to determine the topologies of Δ_G and Δ_G^T .

The covering set U_3^+ associated with the action $3 \rightarrow 1, 2$ is the intersection of two halfspaces in $\mathbf{R}^3$. It therefore appears as the quadrant $\{(x_1, x_2) \mid x_1 < 0 \text{ and } x_2 < 0\}$ in Fig. 14. U_1^+ and U_2^+ each arise from stochastic actions and thus appear as affine halfspaces in Fig. 14.

The boundary of the cylinder C_r is drawn for several critical radii, namely those at which some simplex of Δ_G appears. The innermost cylinder has radius $r = 1 + \epsilon$. Since each action has transition time 1, this cylinder just touches each of the covering sets U_1^+ , U_2^+ , and U_3^+ . From a nerve perspective, this geometry produces the three vertices of the complex Δ_G . As the radius grows, the other simplices of Δ_G appear. The correspondence between the covering geometry and the complex topology is highlighted with a dashed double arrow for one critical event: The intersection $U_1^+ \cap U_3^+$ in the cover $\{U_1^+, U_2^+, U_3^+\}$ corresponds to the 1-simplex $\{1 \rightarrow p2, q3 ; 3 \rightarrow 1, 2\}$ in the complex Δ_G .

8 Controllability of Motions in Stochastic Graphs

Theorem 1 provides a topological test for guaranteed goal reachability in nondeterministic graphs. (A similar result holds in the stochastic setting.) We seek a more general result, along the lines of Hultman's topological characterization of strong connectivity in directed graphs. We have the tool of covering sets from Section 7, so let us use that tool to extract topology from strategies. First, here are some

Definitions: Throughout, let $G = (V, \mathcal{A})$ be a stochastic graph and let S be a nonempty subset of V (S stands for “stop states”).

- G is a stochastic strategy for attaining S if:

- (i) G is stochastically acyclic, and
- (ii) $V \setminus V_{\mathcal{A}} \subseteq S \subseteq V$.

This definition ensures that G contains actions at all states outside of S and that motions under those actions eventually converge (in a subset of S).

Observe that $V_{\mathcal{A}}$ cannot be all of V , since G is stochastically acyclic.

- G contains a complete stochastic strategy for attaining S (on the state space V) if there is some set of actions $\mathcal{B} \subseteq \mathcal{A}$ such that $H_{\mathcal{B}} = (V, \mathcal{B})$ is a stochastic strategy for attaining S . (This definition is consistent with the definition of complete guaranteed strategy from Section 4, p.16.)

- In this case:
- We refer to both $\mathcal{B}$ and $H_{\mathcal{B}}$ as being a *complete stochastic strategy for attaining S* (in the graph G or on the state space V).
 - We say S is a *stochastically attainable goal* (in G or within V).

- Suppose $\mathcal{I}$ is a subset of V ($\mathcal{I}$ stands for “initial states”). G contains a *stochastic strategy for attaining S from $\mathcal{I}$* if G contains a subgraph $H = (W, \mathcal{B})$ such that $\mathcal{I} \cup S \subseteq W$ and H is a stochastic strategy for attaining S . We refer to H as a *stochastic strategy for attaining S from $\mathcal{I}$* . Of course, we may also view H as a complete stochastic strategy for attaining S , now on the state space W .

This definition captures the idea that a strategy for attaining some set of states from some other set of states may require moving through some intermediate states, but not necessarily all of V .

We will presently focus on cases in which $\mathcal{I}$ and S are both singleton sets. In such cases, we speak of attaining one state from another.

8.1 Connectivity: Covers and Chains

The following two lemmas capture the idea that stochastically certain connectivity between two states in a stochastic (or nondeterministic) graph implies coverage of a particular halfspace in $\mathbf{R}^n$ by the graph’s covering sets. This fact appeared as a special case in the proof of Theorem 1. Subsequently, we will use these two lemmas to characterize graph controllability (Theorems 11 and 12).

Lemma 9 (Stochastic Connectivity) *Let $G = (V, \mathcal{A})$ be a stochastic graph with $V = [n]$, $n > 0$. Let $\ell, k \in V$.*

Suppose G contains a stochastic strategy for attaining state k from state ℓ . Then

$$\{\mathbf{x} \in \mathbf{R}^n \mid x_\ell > x_k\} \subseteq \bigcup_{A \in \mathcal{A}} U_A.$$

Proof. While we could give a direct proof, we may also view this lemma as a special case of Lemma 10, with all action transition times zero. That lemma appears next. $\square$

Lemma 10 (Time-Bounded Stochastic Connectivity) *Let $G = (V, \mathcal{A})$ be a stochastic graph with associated nonnegative action transition times $\{\delta_A\}_{A \in \mathcal{A}}$.*

Assume $V = [n]$ and $n > 0$. Let $\ell, k \in V$ be given.

Suppose G contains a stochastic strategy for attaining state k from state ℓ . Let t_ℓ be the worst-case expected convergence time to attain k from ℓ , as given by the solution of System (1) when written out for this stochastic strategy. Then

$$\{\mathbf{x} \in \mathbf{R}^n \mid x_\ell > x_k + t_\ell\} \subseteq \bigcup_{A \in \mathcal{A}} U_A^+.$$

Proof. Let $H = (W, \mathcal{B})$ be the subgraph of G constituting the given stochastic strategy for attaining k from ℓ . So $\ell, k \in W$. We can assume without loss of generality that $W = [k]$, so it is enough to show that

$$\{\mathbf{x} \in \mathbf{R}^k \mid x_\ell > x_k + t_\ell\} \subseteq \bigcup_{B \in \mathcal{B}} U_B^+,$$

with each set U_B^+ now a subset of $\mathbf{R}^k$.

If $\ell = k$, there is nothing to prove, since the set on the left is the empty set, which is a subset of all sets. So assume $\ell \neq k$.

Here is System (1) written out for H :

$$t_i = \max \left(\begin{array}{cc} \max_{\substack{B \in \mathcal{B} \\ B = \{(i,j)\}}} \left(\max_j t_j + \delta_B \right), & \max_{\substack{B \in \mathcal{B} \\ B = \{(i,p_j j)\}}} \left(\sum_j p_j t_j + \delta_B \right) \end{array} \right), \text{ for } 1 \leq i < k; \quad (3)$$

$$t_k = 0.$$

Since H is stochastically acyclic, this system has a unique finite solution with all $t_i \geq 0$.

Suppose there is some $\mathbf{x}^* \in \mathbf{R}^k$ such that $x_\ell^* > x_k^* + t_\ell$ but $\mathbf{x}^* \notin \bigcup_{B \in \mathcal{B}} U_B^+$. Since the sets U_B^+ are invariant with respect to translation along the line $\{\mathbf{x} \in \mathbf{R}^k \mid x_1 = \dots = x_k\}$, we can assume without loss of generality that $x_k^* = 0$.

Now define a Markov chain M contained in H , with support $[k-1]$:

- At state k , let M remain at k with probability 1.
- At state $i \in [k-1]$, pick some action $B \in \mathcal{B}$ whose source is i . Such an action must exist, since H is a complete stochastic strategy for attaining k on the state space W . Let $\delta_i = \delta_B$.
 - If B is stochastic, that is, $B = \{(i, p_j j)\}$ for some set of targets $\{j\}$, then let M 's transitions at i be exactly those of B .
Observe that $x_i^* \leq \sum_j p_j x_j^* + \delta_i$, since $\mathbf{x}^* \notin U_B^+$.
 - If B is nondeterministic, then there must be some target j of B such that $x_i^* \leq x_j^* + \delta_i$, again since $\mathbf{x}^* \notin U_B^+$. Let M move from i to j with probability 1.

Let $P = (p_{ij})$ be the probability transition matrix of M . (It is a $k \times k$ matrix.) Since H is stochastically acyclic, M cannot have any recurrent classes other than the absorbing state k , and therefore the following system has a unique finite solution (this reasoning is at the heart of Lemma 5; see also [48]):

$$\begin{aligned} x_i &= \sum_{j=1}^k p_{ij} x_j + \delta_i, & \text{for } 1 \leq i < k; \\ x_k &= 0. \end{aligned} \quad (4)$$

Moreover, one may obtain the solution to System (4) by iteration, starting from any initial seed for $x_1, \dots, x_{k-1}$. (This is a standard result from Markov chains; it follows in particular from the theorem on p. 389 of [32].)

We will now iterate from two different seeds, obtaining contradictory results:

Iteration Scheme #1:

- Initialize $x_i^{(0)} = t_i$, for $i = 1, \dots, k$, where $\{t_i\}_{i=1}^k$ is the solution to System (3).
- For $m = 0, 1, \dots$, iterate using the update rules:

$$\begin{aligned} x_i^{(m+1)} &= \sum_{j=1}^k p_{ij} x_j^{(m)} + \delta_i, & \text{for } 1 \leq i < k; \\ x_k^{(m+1)} &= 0. \end{aligned}$$

We claim that $x_i^{(m)} \leq t_i$ for all $i = 1, \dots, k$ and all $m = 0, 1, \dots$.

To see this:

The claim is certainly true for all i when $m = 0$ and also for $i = k$ for all m . Inductively, suppose the claim holds for some $m \geq 0$. Then, with $1 \leq i < k$:

$$x_i^{(m+1)} \leq \sum_{j=1}^k p_{ij} t_j + \delta_i \leq t_i.$$

The first inequality follows from the inductive hypothesis, the second from the fact that the right side of System (4) is a special case appearing in the maximizations of System (3).

Consequently,
$$x_i = \lim_{m \rightarrow \infty} x_i^{(m)} \leq t_i, \quad \text{for } i = 1, \dots, k,$$

where $\{x_i\}_{i=1}^k$ is the solution to System (4).

Iteration Scheme #2:

- Initialize $y_i^{(0)} = x_i^*$, for $i = 1, \dots, k$, where $\mathbf{x}^*$ is as supposed earlier.
- For $m = 0, 1, \dots$, iterate using the update rules:

$$\begin{aligned} y_i^{(m+1)} &= \sum_{j=1}^k p_{ij} y_j^{(m)} + \delta_i, & \text{for } 1 \leq i < k; \\ y_k^{(m+1)} &= 0. \end{aligned}$$

We claim that $y_i^{(m)} \geq x_i^*$ for all $i = 1, \dots, k$ and all $m = 0, 1, \dots$.

Verifying:

Again, the claim is true for all i when $m = 0$ and also for $i = k$ for all m .

Inductively, suppose the claim holds for some $m \geq 0$. Then, with $1 \leq i < k$:

$$y_i^{(m+1)} \geq \sum_{j=1}^k p_{ij} x_j^* + \delta_i \geq x_i^*.$$

The first inequality follows from the inductive hypothesis, the second from the construction of M .

Consequently, this time we see that

$$x_i = \lim_{m \rightarrow \infty} y_i^{(m)} \geq x_i^*, \quad \text{for } i = 1, \dots, k.$$

Combining the results of the two iteration schemes, we infer that $x_i^* \leq t_i$ for $i = 1, \dots, k$. In particular, for $i = \ell$, recalling the definition of $\mathbf{x}^*$, we see that

$$t_\ell = 0 + t_\ell = x_k^* + t_\ell < x_\ell^* \leq t_\ell.$$

That says $t_\ell < t_\ell$, a contradiction. $\square$

8.2 Characterizing Controllability with Spheres

The following two theorems characterize topologically the ability of a finite discrete system to reach any state from any other state despite control uncertainty.

Theorem 11 (Graph Controllability) *Let $G = (V, \mathcal{A})$ be a stochastic graph with $V \neq \emptyset$. The following two statements are equivalent:*

- (i) *For every pair of states $v, u \in V$, G contains a stochastic strategy for attaining u from v .*
- (ii) *$\Delta_G \simeq S^{n-2}$, with $n = |V|$.*

Clarification: In (i), the strategy may depend on v and u , that is, different pairs of states may give rise to different strategies.

Proof. If V consists of a single state, then the empty simplex is a stochastic strategy for attaining that state from itself. In fact, that is the only possible stochastically acyclic collection of actions of G ; any action of G must be a self-loop. Consequently, $\Delta_G = \{\emptyset\}$, which by convention is the sphere of dimension -1 , i.e., $n - 2$. So, henceforth we may assume that $V = [n]$, with $n \geq 2$.

I. Proof that (i) implies (ii):

We repeatedly use Lemma 9 to infer that $\bigcup_{A \in \mathcal{A}} U_A$ contains every point of $\mathbf{R}^n$ except the line on which all coordinates are equal. Thus $\Delta_G \simeq S^{n-2}$, by Theorem 8.

II. Proof that (ii) implies (i):

Suppose (i) is false. Then there must be some $s \in V$ such that G does not contain a complete stochastic strategy for attaining s (on the state space V). Define a new stochastic graph $G_{+s} = (V, \mathcal{A}')$, where $\mathcal{A}'$ is the union of $\mathcal{A}$ and all possible loopback actions at s . (This construction is similar to that appearing in the proof of Theorem 1, except that we have *added* loopbacks at s rather than merely replaced the existing actions at s with loopbacks.)

Since G_{+s} contains all the actions of G ,

$$\bigcup_{A \in \mathcal{A}} U_A \subseteq \bigcup_{A \in \mathcal{A}'} U_A.$$

By Theorem 8, $\bigcup_{A \in \mathcal{A}} U_A \simeq S^{n-2}$.

Using the fact that the covering sets U_A are homogeneous and invariant with respect to translation along the line $\{\mathbf{x} \in \mathbf{R}^n \mid x_1 = \dots = x_n\}$, as well as the fact that no proper subset of a sphere is homotopic to that same sphere, one sees that $\bigcup_{A \in \mathcal{A}'} U_A \simeq S^{n-2}$.

On the other hand, a collapsibility argument very similar to that used in the proof of Theorem 1 shows that $\Delta_{G_{+s}}$ must be contractible. By Theorem 8, $\bigcup_{A \in \mathcal{A}'} U_A$ has the same homotopy type as $\Delta_{G_{+s}}$. That says S^{n-2} is contractible, a contradiction. $\square$

Definition: A graph is *fully controllable* if, for any initial state and any stop state, the graph contains a stochastic strategy for attaining the stop state from the initial state.

Example: The strategy complex of the graph in Fig. 13 is homotopic to a circle, i.e., to S^{n-2} . Every action in the graph is uncertain, with either nondeterministic or stochastic outcomes. Nonetheless, Theorem 11 tells us that the graph is fully controllable.

Key Point: Despite significant control uncertainty, all states are precisely attainable.

The Power of Topology: One can easily verify the theorem's assertion by inspection for the example of Fig. 13. For instance, the system can be certain of attaining state 2 from state 1 via the strategy $\{1 \rightarrow p2, q3 ; 3 \rightarrow 1, 2\}$. What we do not know is the exact path the system will take: It may move directly to state 2 from state 1 or it may move through state 3. In fact, it may even cycle for a while between states 1 and 3 before moving to state 2. This is exactly one of the properties we sought in the Introduction (see again page 5): When planning in the presence of uncertainty, one should focus not on specific trajectories but entire classes of motions. Topology is doing this naturally for us.

Here is a time-bounded version of Theorem 11:

Theorem 12 (Time-Bounded Graph Controllability) *Let $G = (V, \mathcal{A})$ be a stochastic graph with $V \neq \emptyset$ and with associated nonnegative action transition times $\{\delta_A\}_{A \in \mathcal{A}}$. Let $T \geq 0$ be given. The following two statements are equivalent:*

- (i) *For every pair of states $v, u \in V$, G contains a stochastic strategy σ_{vu} for attaining u from v with maximal worst-case expected convergence time $t_{\max}(\sigma_{vu})$ no greater than T .*
- (ii) $\Delta_G^T \simeq S^{n-2}$, with $n = |V|$.

Proof. The proof is similar to that given for Theorem 11. One difference is that the covering sets U_A^+ are not homogeneous (they are still invariant with respect to translation along the line in $\mathbf{R}^n$ on which all coordinates are equal). This fact is one reason Lemma 7 uses the cylinder $\partial C_{T+\epsilon}$. $\square$

Remark: Theorem 1 in Section 4 established a two-world scenario for nondeterministic loopback graphs: $\Delta_{G_{\leftarrow s}}$ is homotopic either to S^{n-2} or to a point, depending on whether G contains a complete guaranteed strategy for attaining s , or not, respectively. The same result holds for stochastic loopback graphs, both in the general case and in the time-bounded case. In the time-bounded case, we associate transition time zero to each loopback action. Thus, $\Delta_{G_{\leftarrow s}}^T$ is homotopic either to S^{n-2} or to a point. If G contains a complete stochastic strategy for attaining s , all of whose worst-case expected convergence times are bounded by T , then $\Delta_{G_{\leftarrow s}}^T$ is homotopic to a sphere. Otherwise, $\Delta_{G_{\leftarrow s}}^T$ is homotopic to a point.

9 Topology as a Design Tool: An Example

We are beginning to understand how the topology of a strategy complex reflects a system's capabilities. Theorem 11 tells us that homotopy equivalence between the system's strategy complex and a sphere of dimension two less than the number of system states is equivalent to full controllability.

Moreover, while it remains a research question to understand the full implications of homotopy equivalence, we will soon see that the strategy complex precisely characterizes the stochastically attainable and (potentially) unattainable goals of a system. Of course, full controllability means that all goals are stochastically attainable. In the absence of full controllability, the strategy complex still informs us about system capabilities, in a manner to be explained in the next section.

This section illustrates how a strategy complex may be used as a design tool: We turn design knobs while watching how the strategy complex changes, freezing the design when the strategy complex exhibits a desired topology.




Figure 15: The parameter space for designing an action at state 1. The tuning parameter, drawn in green, consists of an interval and a point. The action is depicted, also in green, for four characteristic parameter values.

Let us revisit the three-state example of Fig. 13, but now consider the situation in which we are designing the actions and the control error. There is one action at each state. The action may be a deterministic transition to one of the other two states, or a stochastic motion to both those states, with tunable transition probabilities, or a nondeterministic motion to both those states. Fig. 15 shows the possible tuning parameter for the action at state 1.

9.1 How Many Design Scenarios?

Counting the design scenarios depends on one's perspective:

- **Degrees of Freedom:** The design problem is a three-degree-of-freedom problem. Each degree of freedom may be modeled as $[0, 1] \cup \{\infty\}$, where $[0, 1]$ represents the

stochastic continuum whose endpoints are the deterministic actions and $\triangleleft$ represents the nondeterministic action.

- **Characteristic Cases:** As the English description above suggests, there are four characteristic cases for each tuning parameter, so 64 characteristic cases overall. Ignoring symmetries, there are 16 cases.

Comment: *A priori*, the precise transition probabilities for a stochastic action could be significant. They certainly affect convergence times, and thus the complexes Δ_G^T . However, varying the probabilities does not affect the homotopy type of Δ_G , except possibly at an endpoint when some transition probability goes to zero.

- **Topologically:** There are 8 distinct strategy complexes possible, as we will see shortly. In fact, ignoring symmetries, there are only 4 topologically distinct cases. In other words, from the perspective of *overall* system capabilities, as measured by stochastically attainable goals, this design problem entails a choice between four different systems.

9.2 Tuning Convergence Times and Designing System Capabilities




Figure 16: One may design system capabilities by tuning the topology of Δ_G^T : As the design parameters change, so do the covering sets associated with the graph's actions. The homotopy type of the cover intersected with $\partial C_{T+\epsilon}$ coincides with the homotopy type of the simplicial complex Δ_G^T . This figure shows a snapshot of that process. As did Fig. 14, the figure shows a two-dimensional slice of the three-dimensional covering sets. See text for further details.

Color Legend: Corresponding graph actions, covering sets, and simplices are color-coded. In case color is not visible, placements are as in Figures 13 and 14.

Probabilities: The actions at states 1 and 2 are stochastic. To reduce clutter, the figure only shows the probabilities of moving from state 1 to 2 and vice-versa.

Fig. 16 shows a snapshot of the tuning process. The upper right frame shows the graph and its actions for the current choice of tuning parameters, the big frame on the left shows the (x_1, x_2) -slice at $x_3 = 1$ of the actions' covering sets. As in Section 7, each action's covering set gives rise to a vertex of the simplicial complex Δ_G^T , shown in the lower right frame of the figure. Also drawn over the covering sets is a convergence time cylinder $\partial C_{T+\epsilon}$ for some desired convergence time T . Key intersection points of the covering sets within this desired time are highlighted. They give rise to the edges in the complex Δ_G^T ; the thickness of an edge is roughly proportional to the difference between the maximal convergence time of the edge (viewed as a stochastic strategy) and the maximal time T permitted.




Figure 17: A slice of the design space from the perspective of the resulting stochastic graphs. A_i means the action at state i . In the slice shown, A_3 is nondeterministic while the parameters for actions A_1 and A_2 vary over their full ranges. For the stochastic range of each parameter, the figure depicts a representative action. To reduce clutter, the figure only shows the probability of moving from state 1 to 2 or vice-versa. See Fig. 18 for the associated covering sets and Fig. 19 for the associated simplicial complexes.

Figures 17–19 depict a two-dimensional slice of the three-degree-of-freedom design space from three perspectives. In this slice, the tuning parameter for the action at state 3 is fixed to be nondeterministic, while the tuning parameters for the actions at states 1 and 2 vary. Fig. 17 shows how the graphs vary as the tuning parameters vary, Fig. 18 shows how the covering sets vary, and Fig. 19 shows how the simplicial complexes Δ_G vary (we now ignore the precise convergence times, focusing on Δ_G not Δ_G^T). Since there are three actions without self-loops, Δ_G always contains all three actions as vertices. Δ_G cannot be the full triangle on those three vertices, since such a triangle would not be stochastically acyclic. Consequently, there should




Figure 18: A slice of the design space from the perspective of the resulting covering sets. The covering sets shown correspond to the graphs of Fig. 17, assuming all actions have unit transition time.




Figure 19: A slice of the design space from the perspective of the resulting simplicial complexes Δ_G . The complexes shown correspond to the covering sets of Fig. 18.

			
SmN, SSN , pSN, mNS, SNS, mmS, SmS, mSS, SSS, pSS, NSS, mpS, SpS, ppS, NpS, SNp, Smp, SSP, pSp, Spp, ppp, mmm, Smm, mSm, SSm, pSm, NSm	NNN	pmN, pNS, NNS, pmS, NmS, pNp, pmp, pmm, Nmm, SNN, SpN, ppN, mNm, SNm, Spm, ppm, Npm, mpm, mmN, mSN, NSN, mNp, mmp, mSp, NSp, mpp, Npp	mNN, mpN, NpN, NmN, NNp, Nmp, pNN, pNm, NNm

Figure 20: There are four topologically distinct simplicial complexes attainable by varying the design parameters of a three-state graph with one (acyclic) action at each state. These complexes describe the inherent capabilities of the system under varying degrees of control uncertainty. The columns classify the 64 characteristic design cases in terms of the complexes they generate. The leftmost column corresponds to full controllability. The example of Fig. 13 is circled.

Legend: Each characteristic case is summarized by a 3-letter code, with the i^{th} letter describing the action at state i as follows:

- | | |
|---------------------|--|
| S: stochastic | m: deterministic transition to state $i - 1$ |
| N: nondeterministic | p: deterministic transition to state $i + 1$
(with wraparound at 3/1) |

be as many different complexes in the full design space as there are ways to form edges from three vertices, namely eight ways. Indeed, inspection of Fig. 19 reveals eight unique complexes; the full design space repeats these. Ignoring symmetries, there are in fact only four distinct complexes, shown in the top row of Fig. 20.

One may now classify each of the 64 characteristic cases of the design space by the type of its associated strategy complex Δ_G , as shown in the columns of Fig. 20.

Section 10 discusses in more detail what each of the four complexes means. We already know that the complex defining the leftmost column constitutes full controllability. We thus see twenty-seven different characteristic implementations that attain full controllability. The example of Fig. 13 is circled.

Design Implications: From a design perspective, one may now use other constraints to decide which of the twenty-seven implementations might be desirable for full controllability. Moreover, one may tune the convergence times using the idea described in Fig. 16. This process is much like choosing forces within the nullspace of a grasp to satisfy some design criterion while maintaining a specific equilibrium grasp on an object [51].

10 Duality

In order to understand better how the complex Δ_G describes system capabilities, we need to develop a perspective dual to the spherical perspectives of Theorems 1 and 11. In the process, we will see how backchaining and contractibility of certain subcomplexes are manifestations of the same idea. Finally, we will see how to match complexes with design criteria.

10.1 Start Region Contractibility

Consider the start region V_σ of some simplex σ of stochastically acyclic actions in a graph G . As we will see in this section, the subcomplex induced by all actions of the graph with sources in that region is contractible. Intuitively, contractibility is consistent with viewing σ as a collapsing of V_σ , moving the system off those states.

Definitions Suppose $G = (V, \mathcal{A})$ is a stochastic graph and $W \subseteq V$. Define the following:

- $\mathcal{A}|W$ is the collection of all actions of G whose sources lie in W .
- $W_{\mathcal{A}}$ is the union of W and all targets of actions in $\mathcal{A}|W$.
- $G|W = (W_{\mathcal{A}}, \mathcal{A}|W)$. Intuition: $G|W$ is a subgraph of G , induced by all the actions of G whose sources lie in W but whose targets may lie outside of W .
- A moves off W if $A \in \mathcal{A}|W$ 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 $V \setminus W$.

Lemma 13 (Contractibility of Start Regions) *Let $G = (V, \mathcal{A})$ be a stochastic graph. Suppose $W = V_\sigma$ for some $\sigma \in \Delta_G$. Then $\Delta_{G|W}$ is contractible.*

Proof. If $W = \emptyset$, then $\Delta_{G|W}$ is the void complex $\emptyset$, which is considered contractible [47] (do not confuse it with the empty complex $\{\emptyset\}$).

If $W \neq \emptyset$, then there is some action $A \in \mathcal{A}|W$ that moves off W . To see this, suppose otherwise. Imagine constructing a Markov chain M with support W from the graph (V, σ) via the process of page 22. If no action A moves off W , then an adversary can ensure that M has no transitions to states outside of W . This means M must contain a recurrent class within W , since the chain is finite [32]. That contradicts the stochastic acyclicity of σ (recall Lemma 5).

Now suppose $\tau \in \Delta_{G|W}$. Consider $\tau' = \tau \cup \{A\}$. If τ' is not stochastically acyclic, then τ' must give rise to some Markov chain (M, W') whose support W' is a recurrent class. Since

τ is stochastically acyclic, M must include transitions induced by A . Since A moves off W it moves off W' , contradicting the assumed recurrent nature of W' .

Consequently, τ' is stochastically acyclic, establishing that $\Delta_{G|W}$ is a cone with apex A , hence contractible. $\square$

Remark: A similar result holds for the time-bounded case. The proof is more involved. The problem in the time-bounded case is that the action A appearing in the proof above need no longer be a cone apex for the complex $\Delta_{G|W}^T$. The reason is that some worst-case expected convergence time of τ' may exceed the desired bound T , even though τ and A separately satisfy the bound. One must therefore argue differently. In particular, one can return to the covering set approach and see that the cover corresponding to $\Delta_{G|W}^T$ must be contractible.

10.2 Source Complex

This section associates a new simplicial complex to every graph, called the *source complex*, in effect compressing the graph's strategy complex. The main result is that compression preserves homotopy type. This result explains a slight sleight of hand in the discussion of Section 9. Even though strategy complexes reside in a space of actions, we managed to view the complexes as sitting on the original graph. Doing so was very natural in the example of Section 9 since the graph contained exactly one action at each state. The current section establishes that one may always view the strategy complex as residing back on the graph, via the compression to source complexes.

Definition: The *source complex* $\overline{\Delta}_G$ of a stochastic graph $G = (V, \mathcal{A})$ is the collection $\{V_\sigma \mid \sigma \in \Delta_G\}$. The underlying vertex set of $\overline{\Delta}_G$ is V .

It is easy to check that $\overline{\Delta}_G$ really is a simplicial complex.

$\overline{\Delta}_G$ describes all possible start regions of strategies definable by actions of G . The complements of those start regions, that is, sets of the form $V \setminus V_\sigma$, describe all stochastically attainable goals. Observe that V is never a simplex of $\overline{\Delta}_G$, since any purported strategy with actions at all states of a graph would actually cycle forever.

Theorem 14 (Compression Preserves Homotopy Type) *Let G be a stochastic graph.*

$$\Delta_G \simeq \overline{\Delta}_G.$$

Proof. Let $P = \mathcal{F}(\Delta_G)$ and $Q = \mathcal{F}(\overline{\Delta}_G)$ be the face posets of the two complexes.

Define $f : P \rightarrow Q$ by $f(\sigma) = V_\sigma$. Suppose $W \in Q$. Then $f^{-1}(Q_{\leq W})$ is the face poset of $\Delta_{G|W}$, which is contractible by Lemma 13. The desired result now follows from the Quillen Fiber Lemma. $\square$

Remark: Again, a similar result holds for the time-bounded case.

The following lemma is a useful tool:

Lemma 15 (Source Complex Membership) *Let $G = (V, \mathcal{A})$ be a stochastic graph. Suppose W is a nonempty proper subset of V such that every proper subset of W is a simplex in $\overline{\Delta}_G$. Then $W \in \overline{\Delta}_G$ if and only if G contains an action that moves off W .*

Proof. One direction follows from the proof of Lemma 13.

For the other direction, suppose A moves off W . Let v be the source of A . By assumption, there is some $\sigma \in \Delta_G$ such that $V_\sigma = W \setminus \{v\}$. Consider $\tau = \sigma \cup \{A\}$. τ is stochastically acyclic since σ is stochastically acyclic and since any transitions at v induced by A in a Markov chain derived from τ have probability less than 1 of remaining in W . $V_\tau = W$, so $W \in \overline{\Delta}_G$. $\square$

10.3 Contractibility Characterization of Goal Attainability

We now obtain a result dual to Theorem 1. The proof of the following theorem makes explicit the connection between backchaining and contractibility.

Definition: Suppose Σ is a simplicial complex with underlying vertex set V . A *minimal nonface* of Σ is a set $W \subseteq V$ such that W is not a simplex of Σ but every proper subset of W is a simplex of Σ .

Theorem 16 (Contractibility Characterization of Goal Attainability)

Suppose $G = (V, \mathcal{A})$ is a stochastic graph and $\emptyset \neq S \subseteq V$. S is a stochastically attainable goal in G if and only if $\Delta_{G|W}$ is contractible for every $W \subseteq V \setminus S$.

Proof. **I.** Suppose G contains a complete stochastic strategy for attaining $S \subseteq V$.

Then there is some $\sigma \in \Delta_G$ such that $V_\sigma = V \setminus S$. By Lemma 13, and the simplicial nature of σ , we see that $\Delta_{G|W}$ is contractible for every $W \subseteq V_\sigma$.

II. Suppose $\Delta_{G|W}$ is contractible for every $W \subseteq V \setminus S$.

We now construct $\sigma \in \Delta_G$ by backchaining:

1. Initialize $S^* := S$ and $\sigma := \emptyset$.
2. If $S^* = V$, then done; return σ .
3. Otherwise, as shown below, G contains an action A that moves off $V \setminus S^*$.
Let v be the source of A . Update $S^* := S^* \cup \{v\}$ and $\sigma := \sigma \cup \{A\}$.
4. Repeat Steps 2 and 3 until done.

It is easy to see that the σ returned in Step 2 really is stochastically acyclic. $V_\sigma = V \setminus S$, so σ is a complete stochastic strategy for attaining S , as desired.

To complete the proof, we need to establish the existence of action A in Step 3. Suppose no such A exists. By the proof of Lemma 13, this means $V \setminus S^* \notin \overline{\Delta}_G$. Choose nonempty $W \subseteq V \setminus S^*$ such that $W \notin \overline{\Delta}_G$ but every proper subset of W is a simplex of $\overline{\Delta}_G$. In other words, W is a minimal nonface of $\overline{\Delta}_G$. Observe that W is also a minimal nonface of $\overline{\Delta}_{G|W}$.

A priori, the underlying vertex set of $\overline{\Delta}_{G|W}$ is the state space $W_{\mathcal{A}}$ of $G|W$, but no state outside W can be a vertex of $\overline{\Delta}_{G|W}$ since $G|W$ contains no actions at such states. Consequently, W being a minimal nonface, $\overline{\Delta}_{G|W}$ is the boundary complex of the full simplex on W . Thus $\overline{\Delta}_{G|W} \simeq S^{|W|-2}$. On the other hand, Theorem 14 implies $\overline{\Delta}_{G|W} \simeq \Delta_{G|W}$, which is contractible. Contradiction. $\square$

10.4 The Dual Complex

In algebraic topology, geometric duality becomes algebraic duality. The theorems are often formulated for nonempty proper subsets of spheres. Given the importance of spheres when reasoning about strategy complexes, one imagines that duality might help illuminate the homotopy type of a strategy complex. There is a simple combinatorial description for simplicial complexes of a particular duality known as *Alexander Duality* [7]. Formally, given a simplicial complex Σ with underlying vertex set V , its *Combinatorial Alexander Dual* is the complex $\Sigma^* = \{\sigma \subseteq V \mid V \setminus \sigma \notin \Sigma\}$. Observe that $(\Sigma^*)^* = \Sigma$.

The Alexander dual of a source complex has an important natural meaning: it describes all goals that are *not* stochastically attainable (in the sense of page 30), as follows.

Definition: The *dual complex* $\bar{\Delta}_G^*$ of a stochastic graph $G = (V, \mathcal{A})$ is the collection $\{V^* \subseteq V \mid V \setminus V^* \notin \bar{\Delta}_G\}$. The underlying vertex set of $\bar{\Delta}_G^*$ is again V .

It is again easy to verify that $\bar{\Delta}_G^*$ is a simplicial complex. Moreover, it follows from the definitions that $V^* \in \bar{\Delta}_G^*$ if and only if there is no strategy $\sigma \in \Delta_G$ such that $V \setminus V_\sigma \subseteq V^*$. In other words, we may think of $\bar{\Delta}_G^*$ as encoding all the *potentially unattainable* goals, meaning there is no stochastic strategy for attaining such a goal from everywhere in the graph.

The reader may be wondering why we have not also defined a simplicial complex to define the stochastically attainable goals. The answer is that simplicial complexes must satisfy some kind of monotone property: subsets of simplices must also be simplices. Failure to attain a goal is a monotone property; if we cannot be certain of attaining a goal of some size when planning with uncertainty, then we also cannot be certain of attaining any subset of that goal. In contrast, attainability is not similarly monotone; just because we can attain some goal does not mean we can attain a more precisely defined goal.

In short, start regions and potentially unattainable goals define simplicial complexes via the source and dual complexes. The stochastically attainable goals are given implicitly, as the complements of the start regions.

10.5 Duality in Design

We now understand better the topological information encoded in a graph's strategy complex. Earlier (Fig. 11) we saw that homotopy equivalence naturally favors precise actions over imprecise actions. We also realized (page 35) that topology abstracts away from particular trajectories to broad classes of motions. The source and dual complexes make this explicit. Moreover, the state space V of a graph provides a good reference frame onto which homotopy equivalences should map. One can compare different system capabilities of different graphs with underlying state space V by comparing their source (and/or dual) complexes.

Returning to the design problem of Section 9, Fig. 21 now shows both the source and dual complexes corresponding to Fig. 20. Moreover, the figure makes sense even if there are multiple actions or design parameters at each state.

Recall the various semantics we have been attaching to complexes. Consider the first column of complexes in Fig. 21. Previously we observed that this column corresponds to full controllability since the strategy complex is homotopic to a circle. That statement still holds,

$\overline{\Delta}_G$				
$\overline{\Delta}_G^*$	$\{\emptyset\}$			
Meaning	All goals stochastically attainable	No singleton state is stochastically attainable	One state is not stochastically attainable	Two states are not stochastically attainable

Figure 21: A comparison of the source and dual complexes for the design problem of Fig. 20.

only now we may refer to the source complex, since we now know that the strategy and source complexes have the same homotopy type. The dual complex in this case is the empty complex, meaning that there are no potentially unattainable goals. That statement is consistent with full controllability: all goals are stochastically attainable.

Remark: The reader may observe that this reasoning provides the basis for a purely combinatorial proof of Theorem 11, made possible by Theorem 14 and the semantics of source and dual complexes.

In some cases, one may not want full controllability. For instance, imagine that we are designing the hallways and door controllers in a bank. Suppose, for simplicity, there are three states: OUTSIDE, LOBBY, and VAULT. We might want to ensure that OUTSIDE and LOBBY are reachable from everywhere by everyone, but that VAULT is reachable only by designated people. This suggests a system whose capabilities switch between those given by the first and third columns of complexes in Fig. 21, depending on the people walking in the hallways. We could go back to Fig. 20 to select a particular implementation of these switchable capabilities that also avoids passing by the VAULT accidentally. For instance, we might use a switchable deterministic action in the LOBBY whose actual direction is under the control of the bank's security guards. Finally, we might also go back to Fig. 16 to fine-tune convergence times.

11 Modularity

This section discusses combination and simplification of graphs.

First, some topology. The (*topological*) *join* $X * Y$ of two topological spaces X and Y is the quotient space obtained from $X \times Y \times [0, 1]$ by identifying each set $\{x\} \times Y \times \{0\}$ to a point and each set $X \times \{y\} \times \{1\}$ to a point, for all $x \in X$ and $y \in Y$ [61]. Geometrically, one may think of $X * Y$ as the union of all line segments joining points in X to points in Y [39]. For instance, the topological join of two finite disjoint edges is a tetrahedron.

In the case of simplicial complexes, the join assumes a simple combinatorial form:

Definition: Suppose Σ and Γ are two simplicial complexes with disjoint underlying vertex sets. Their (*simplicial*) *join* is the simplicial complex defined as

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

For instance, if Σ and Γ represent two disjoint line segments, say

$$\begin{aligned} \Sigma &= \{\emptyset, \{p_1\}, \{p_2\}, \{p_1, p_2\}\}, \\ \Gamma &= \{\emptyset, \{q_1\}, \{q_2\}, \{q_1, q_2\}\}, \end{aligned}$$

then

$$\begin{aligned} \Sigma * \Gamma &= \{\emptyset, \{p_1\}, \{p_2\}, \{q_1\}, \{q_2\}, \\ &\quad \{p_1, p_2\}, \{q_1, q_2\}, \{p_1, q_1\}, \{p_1, q_2\}, \{p_2, q_1\}, \{p_2, q_2\}, \\ &\quad \{p_1, p_2, q_1\}, \{p_1, p_2, q_2\}, \{p_1, q_1, q_2\}, \{p_2, q_1, q_2\}, \{p_1, p_2, q_1, q_2\}\}, \end{aligned}$$

which represents a tetrahedron.

Observe: $\Sigma * \{\emptyset\} = \{\emptyset\} * \Sigma = \Sigma$ and $\Sigma * \emptyset = \emptyset * \Sigma = \emptyset$, for all simplicial complexes Σ .

Fact: Homotopy equivalence ($\simeq$) commutes with the join operator ($*$).

11.1 Graph Union

Suppose $G_1 = (V_1, \mathcal{A}_1)$ and $G_2 = (V_2, \mathcal{A}_2)$ are two stochastic graphs. (The state spaces V_1 and V_2 are allowed to overlap.)

Consider the stochastic graph $G = (V, \mathcal{A})$ representing the union of G_1 and G_2 . It is defined by $V = V_1 \cup V_2$ and $\mathcal{A} = \mathcal{A}_1 \sqcup \mathcal{A}_2$. The operator “ $\sqcup$ ” means “disjoint union”, that is, we treat actions in $\mathcal{A}_1$ and $\mathcal{A}_2$ as distinct even if they happen to have identical edge sets. (Any resulting redundancy washes away under homotopy equivalence, as we saw in Section 5.5 and Fig. 11.) We write $G = G_1 \cup G_2$ for shorthand.

We wish to understand the relationship between Δ_G , Δ_{G_1} , and Δ_{G_2} . Let us assume that V_1 and V_2 are both nonempty. Treating $\mathcal{A}_1$ and $\mathcal{A}_2$ as disjoint means Δ_{G_1} and Δ_{G_2} have disjoint underlying vertex sets, so we can form their join. Observe that $\Delta_G \subseteq \Delta_{G_1} * \Delta_{G_2}$, in fact,

$$\Delta_G = \{\tau \in \Delta_{G_1} * \Delta_{G_2} \mid \tau \text{ is stochastically acyclic}\}.$$

The next lemma describes some common cases in which there is equality of complexes. Case (a) covers completely disjoint graphs, case (b) covers graphs that touch at a single state, and case (c) covers graphs in which one graph is essentially feeding into the other.

Lemma 17 (Join Sufficiency) *Suppose $G = G_1 \cup G_2$, with notation as above. Assume that V_1 and V_2 are both nonempty. In each of the following cases, $\Delta_G = \Delta_{G_1} * \Delta_{G_2}$:*

- (a) $V_1 \cap V_2 = \emptyset$,
- (b) $V_1 \cap V_2$ contains a single state,
- (c) G_1 (or G_2) has no actions with sources in $V_1 \cap V_2$.

Proof. (a) Clear from the definitions. (Also, (a) follows from (c).)

(b) If $\sigma \cup \gamma$ is not stochastically acyclic although $\sigma \in \Delta_{G_1}$ and $\gamma \in \Delta_{G_2}$, then $\sigma \cup \gamma$ must induce a Markov chain (M, W) whose support W is a recurrent class containing the common state of V_1 and V_2 . The transitions of M at that state are induced by a single action A , with either $A \in \sigma$ or $A \in \gamma$. In the first case, all of A 's targets lie in V_1 and so σ alone could generate (M, W) , contradicting the stochastic acyclicity of σ . Similarly for the second case.

(c) Suppose the sources of all of G_1 's actions lie in $V_1 \setminus V_2$. If $\sigma \cup \gamma$ is not stochastically acyclic although $\sigma \in \Delta_{G_1}$ and $\gamma \in \Delta_{G_2}$, then $\sigma \cup \gamma$ must induce a Markov chain (M, W) whose support W is a recurrent class overlapping both $V_1 \setminus V_2$ and V_2 . Any action with source in V_2 must be an action of γ and thus has no targets in $V_1 \setminus V_2$, contradicting the assumption that W is a recurrent class of M . $\square$

11.2 Testing Acyclicity

Suppose G_1 and G_2 do not satisfy the conditions of Lemma 17. The following options for computing Δ_G exist:

1. One possibility is to compute $\overline{\Delta}_G$ rather than Δ_G . Conceptually, $\overline{\Delta}_G$ is easier to compute than Δ_G . The two complexes have the same homotopy type (Theorem 14). Moreover, $\overline{\Delta}_G$ is more explicitly useful in characterizing system capabilities, as we saw in Section 10. Computing $\overline{\Delta}_G$ amounts to repeated backchaining (see Section 12).
2. Another possibility is to work with the definitions directly. For instance, in deciding whether a potential simplex $\sigma \cup \gamma$ really is stochastically acyclic, one may attempt to solve System (1) (written out for $\sigma \cup \gamma$) using strictly positive action transition times $\{\delta_A\}$. The solution will diverge precisely when $\sigma \cup \gamma$ can induce a Markov chain whose support is a recurrent class.
3. Finally, a third possibility is to break the computations into simpler pieces using the tools described next.

Lemma 18 (Combining Actions) *Let $G = (V, \mathcal{A})$ be a stochastic graph and $\sigma \in \Delta_G$. Suppose there are actions $A_1, \dots, A_k \in \mathcal{A}$, all with the same source, such that $\sigma \cup \{A_i\} \in \Delta_G$ for each individual action A_i , $i = 1, \dots, k$. Then $\sigma \cup \{A_1, \dots, A_k\} \in \Delta_G$.*

Proof. Let v be the source of the actions $A_1, \dots, A_k$.

Any Markov chain (M, W) induced by $\sigma \cup \{A_1, \dots, A_k\}$ whose support W is a recurrent class must have transitions at v induced by some A_i . But then $\sigma \cup \{A_i\}$ could induce the same Markov chain. Contradiction. $\square$

Now suppose $\sigma \in \Delta_G$ and A is an action of G . To avoid trivialities, assume A is not a self-looping (non)deterministic action. Let v_A be A 's source. One can decide whether $\sigma \cup \{A\} \in \Delta_G$ as follows:

- By Lemma 18, one may assume without loss of generality that $v_A \notin V_\sigma$ (recall V_σ is the set of all sources of actions in σ).
- By Lemma 17(c), if v_A is not the target of some action of σ , or if no target of A lies in V_σ , then $\sigma \cup \{A\} \in \Delta_G$. (This step is not necessary but provides a convenient filter.)
- Now consider the following system of equations:

$$\begin{aligned}
 q_v &= \max \left(\begin{array}{cc} \max_{\substack{B \in \sigma \\ B = \{(v, u_j)\}}} \left(\max_j q_{u_j} \right), & \max_{\substack{B \in \sigma \\ B = \{(v, p_j u_j)\}}} \left(\sum_j p_j q_{u_j} \right) \end{array} \right), \text{ if } v \in V_\sigma; \\
 q_v &= 1, & \text{ if } v = v_A; \\
 q_v &= 0, & \text{ otherwise.}
 \end{aligned} \tag{5}$$

This system has a unique finite solution $\{q_v\}_{v \in V}$, since σ is stochastically acyclic. One can compute the solution using iteration. For each $v \in V$, q_v is the maximum probability under all Markov chains induced by σ that the system, when started at state v , will reach v_A , the source of action A .

Consequently:

- (i) If A is stochastic then $\sigma \cup \{A\} \in \Delta_G$ if and only if $q_u < 1$ for *some* target u of A .
- (ii) If A is nondeterministic, then $\sigma \cup \{A\} \in \Delta_G$ if and only if $q_u < 1$ for *all* targets u of A .

11.3 Simplification via Strongly Controllable Subspaces

For regular directed graphs, one defines an equivalence relation by saying that two states are equivalent if each is reachable from the other. The resulting equivalence classes are called *strongly connected components*. One may form a quotient graph by collapsing each strongly connected component to a single point. Ignoring induced self-loops, the result is a directed acyclic graph.

This section pursues such simplification for stochastic graphs. We will discover some enrichments not present in regular directed graphs.




Figure 22: States 1 and 3 are each certainly attainable from the other, but an excursion to state 2 cannot be either excluded or forced by the system when leaving state 1. (An adversary may have control over the precise path.)

An Equivalence Relation

The first subtlety concerns connectivity. It is not enough to define two states as equivalent if each is reachable from the other. For instance, in Fig. 22, there exist strategies for attaining state 3 from state 1 and vice-versa. The problem is that the strategy to attain 3 from 1 might or might not pass through state 2. That would be fine if the graph contained strategies for purposefully attaining state 2 from 1 and 3, but it does not.

We therefore need a stronger requirement than pairwise reachability. Here it is:

Definitions Suppose $G = (V, \mathcal{A})$ is a stochastic graph.

Define a binary relation $\leftrightarrow$ on $V \times V$ as follows:

- $v \leftrightarrow w$ if G contains a subgraph $H = (W, \mathcal{B})$ such that:
 - (a) $v, w \in W$, and
 - (b) $\Delta_H \simeq S^{|W|-2}$.

It is easy to see that $\leftrightarrow$ defines an equivalence relation on V . We refer to an equivalence class of $\leftrightarrow$ as a *strongly controllable subspace* of G .

For each equivalence class W there is some subgraph $H = (W, \mathcal{B})$ whose strategy complex is homotopy equivalent to a sphere of dimension $|W|-2$. The proof of Theorem 11 ensures that H contains strategies for attaining any state in W , *without leaving* W .

- If $\sim$ is any equivalence relation on V , let $G/\sim$ be the quotient graph obtained by collapsing states to their equivalence classes. Thus $G/\leftrightarrow$ collapses the strongly controllable subspaces of G .

A special case: Let $H = (W, \mathcal{B})$ be a subgraph of G with $W \neq \emptyset$. Define the equivalence $\sim$ as follows: Each state of $V \setminus W$ is equivalent only to itself, while all states in W are equivalent. This means W collapses to a point, call it $\diamond$. We designate the resulting quotient graph by G/W . $G/W = (V', \mathcal{A}')$, with $V' = V \setminus W \cup \{\diamond\}$ and $\mathcal{A}'$ essentially identical to $\mathcal{A}$ except that any states of W appearing as sources or targets in $\mathcal{A}$ have been remapped to $\diamond$ in $\mathcal{A}'$.

Side issues: Distinct actions of G may effectively become the same action in $G/\sim$. Nonetheless, we treat them as distinct. In other words, there is a bijective correspondence between $\mathcal{A}$ and $\mathcal{A}'$. Once again, any redundancy will be washed away by homotopy equivalence

once we pass to strategy complexes. The reader may also notice that for a given action, different targets in G may become the same target in $G/\sim$. This poses no real issue. Nondeterministic actions are sets of edges, so the redundancy disappears automatically. For stochastic actions, one can add up the probabilities of different edges with the same target to form a single edge. Finally, some actions that are stochastically acyclic in G may very well become self-looping (non)deterministic actions in $G/\sim$. Such actions disappear when we pass to strategy complexes.

Now imagine that $H = (W, \mathcal{B})$ is a subgraph of G satisfying $\Delta_H \simeq S^{|W|-2}$ (W need not be an equivalence class of G under $\leftrightarrow$, but must certainly be a subset of such an equivalence class). Then the system has full controllability within H , by Theorem 11. Consequently, H and G/W are *almost* separate graphs that touch at a single state, as in Lemma 17(b). We might therefore expect to see equality between complexes as in that lemma. This analogy is not quite correct; the single common state is $\diamond$ in G/W , but it might be all of W in H . As a result, equality becomes homotopy equivalence, as the following lemma and its proof demonstrate.

Lemma 19 (Factoring Controllable Subgraphs) *Let $G = (V, \mathcal{A})$ be a stochastic graph and $H = (W, \mathcal{B})$ a subgraph satisfying $\Delta_H \simeq S^{|W|-2}$, with $W \neq \emptyset$. Then*

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

Proof. By Theorem 14, we need merely prove $\overline{\Delta}_G \simeq \overline{\Delta}_H * \overline{\Delta}_{G/W}$.

Let $P = \mathcal{F}(\overline{\Delta}_H * \overline{\Delta}_{G/W})$ and $Q = \mathcal{F}(\overline{\Delta}_G)$ be the associated face posets. Every p in P is a nonempty simplex of $\overline{\Delta}_H * \overline{\Delta}_{G/W}$, so we will simply write $X \cup Y$ for elements of P .

Define $f : P \rightarrow Q$ 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 $Y \in \overline{\Delta}_{G/W}$. Observe that $X \subseteq W$ and $Y \subseteq V \setminus W \cup \{\diamond\}$.

Recall: $\diamond$ is the state of G/W that represents W collapsed to a point.

[Comment: Observe the power of posets: f maps a singleton set $\{\diamond\}$ to all of W .]

The Quillen Fiber Lemma will give us the desired result, if we can satisfy these preconditions:

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

Establishing these three conditions is a bit tedious. There are several cases. We will prove the most interesting cases and leave the rest to the reader. We emphasize that one really must prove (i); full controllability within H is needed to establish that f is well-defined.

(i). We assume $\diamond \in Y$. We leave to the reader the case in which $\diamond \notin Y$.

Let $X \in \overline{\Delta}_H$ and $Y \in \overline{\Delta}_{G/W}$ be given, with $\diamond \in Y$. This means there exists a stochastically acyclic set of actions $\tau' \in \Delta_{G/W}$ such that $V_{\tau'} = Y$. Now let τ be the actions of G that generate τ' (recall: the correspondence between actions in G and G/W is bijective). Since $\diamond \in Y$, τ' includes an action A' with source $\diamond$. We may assume there is exactly one such action in τ' . Let $w_0 \in W$ be the source of the corresponding action A of τ . Since $\Delta_H \simeq S^{|W|-2}$, the proof of Theorem 11 implies a stochastically acyclic set of actions $\sigma \in \Delta_H$ such that $V_\sigma = W \setminus \{w_0\}$. It is not hard to see that $\sigma \cup \tau$ is stochastically acyclic since τ' is (Lemma 17(c) makes part of the argument, establishing that $\sigma \cup \tau \setminus \{A\}$ is stochastically acyclic). Finally,

$$V_{\sigma \cup \tau} = W \setminus \{w_0\} \cup Y \setminus \{\diamond\} \cup \{w_0\} = W \cup Y \setminus \{\diamond\},$$

meaning $f(X \cup Y) \in \overline{\Delta}_G$.

(ii). Easy.

(iii). Every $q \in Q$ is a nonempty simplex of $\overline{\Delta}_G$, so we will write $\overline{V}$ in place of q , with $\overline{V} \subset V$. Let $\overline{V}$ be given. We need to show that $f^{-1}(Q_{\leq \overline{V}})$ is contractible.

We deal here with the case in which $\overline{V} \cap W = W$. We leave to the reader the case in which $\overline{V} \cap W$ is a proper subset of W .

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

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

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

There remains to show that $Y \cup \{\diamond\} \in \overline{\Delta}_{G/W}$. Suppose that is false. Then there must be some set of states $Y' \subseteq Y$, such that every proper subset of $Y' \cup \{\diamond\}$ is a simplex of $\overline{\Delta}_{G/W}$ but $Y' \cup \{\diamond\}$ is not. By Lemma 15, no action of G/W moves off $Y' \cup \{\diamond\}$, which means no action of G moves off $Y' \cup W$. On the other hand, $Y' \cup W \subseteq Y \cup W \subseteq \overline{V}$, implying $Y' \cup W \in \overline{\Delta}_G$. Now Lemma 15 reveals a contradiction.

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

Theorem 20 (Controllability Structure) *Let G be a stochastic graph. Then*

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

where n is the size of G 's state space and k is the number of equivalence classes induced by $\leftrightarrow$.

Proof. Use Lemma 19 repeatedly on the equivalence classes determined by $\leftrightarrow$. Bear in mind that $S^i * S^j \simeq S^{i+j+1}$ [6]. $\square$

Theorem 20 generalizes a result of Hultman's for directed graphs [44]. For directed graphs, it is not hard to see that $\Delta_{G/\leftrightarrow}$ is either the empty complex or homotopic to a point. So for directed graphs, Δ_G always looks like a sphere or a point, homotopically. For stochastic

(in particular, nondeterministic) graphs, the domain is considerably richer (see Section 13 for instance); $\Delta_{G/\leftrightarrow}$ may have arbitrary homotopy type (with n and k varying) within the range of finite simplicial complexes.

We may think of S^{n-k-1} as measuring controllability in the graph G , whereas $\Delta_{G/\leftrightarrow}$ captures the adversary's ability to constrain the system. The full meaning of the factor $\Delta_{G/\leftrightarrow}$ is still a research question. We learned a lot in Section 10; the following example offers further insight.

11.4 An Example (Air Travel During Thunderstorm Season)




Figure 23: This graph contains three types of actions: deterministic (green), stochastic (red), and nondeterministic (blue). The graph might represent some poor passengers' potential flight paths during a day of thunderstorms.

Consider the graph of Fig. 23, which might represent possible air travel routes between cities in the United States during some thunderstorm-infested July afternoon. The graph includes deterministic, stochastic, and nondeterministic actions. Perhaps deterministic actions represent flights that are certainly possible, stochastic actions represent flights whose destinations are stochastically determined by emerging thunderstorms, and nondeterministic actions represent sets of possible flights, to any one of which a hopeful passenger will be assigned in an otherwise unpredictable fashion (nondeterminism is Murphy's Law in action).

Of interest is the set of cities reachable from anywhere in the system and whether a passenger might become trapped in an endless loop trying to reach some particular city.

The equivalence relation $\leftrightarrow$ partitions the state space into five equivalence classes. Three




Figure 24: The two nontrivial equivalence classes of the $\leftrightarrow$ relation for the graph of Fig. 23, along with the actions establishing full controllability within each equivalence class.

of these are the singleton sets $\{3\}$, $\{6\}$, and $\{8\}$. The two nontrivial equivalence classes are $\{1, 2\}$ and $\{4, 5, 7, 9, 10\}$, as shown in Fig. 24.

This means that passengers wishing to travel between cities 1 and 2 are in luck. They can go directly from one city to the next. Passengers traveling between cities in the set $\{4, 5, 7, 9, 10\}$ are also fairly lucky. They will certainly get to their destinations, though the exact routing may not always be certain in advance.

Unfortunately, there is a convention in city 4, towards which passengers from all over the country are traveling. Many will not make it today, as we shall see.

The left panel of Fig. 25 shows the quotient graph $G/\leftrightarrow$ for the graph G in Fig. 23. Each equivalence class is represented by one of its states. The source complex $\overline{\Delta}_{G/\leftrightarrow}$ of this graph appears in Fig. 26. The underlying vertex set of $\overline{\Delta}_{G/\leftrightarrow}$ is the state space of $G/\leftrightarrow$, represented by the cities $\{1, 3, 4, 6, 8\}$. (Keep in mind that, for instance, state 4 really represents the entire equivalence class of cities $\{4, 5, 7, 9, 10\}$. Theorem 20 assures us that the capabilities of the graph overall are captured by this reduced quotient representation.)

The complements of the maximal simplices of $\overline{\Delta}_{G/\leftrightarrow}$ describe the most precisely attainable goals, as implied by the results of Section 10. Since $\{1, 3, 4, 8\}$ is a simplex in $\overline{\Delta}_{G/\leftrightarrow}$, this means city 6 is attainable from anywhere within the graph $G/\leftrightarrow$ and thus from anywhere in the graph G . The right panel of Fig. 25 exhibits a strategy (in the quotient graph) for attaining city 6. Too bad the convention is not happening in that city.

Particularly informative are the minimal nonfaces of $\overline{\Delta}_{G/\leftrightarrow}$. The complements of these sets are maximal simplices in $\overline{\Delta}_{G/\leftrightarrow}^*$, meaning they are the maximal goals not certain to be attainable from everywhere in the graph $G/\leftrightarrow$. $\overline{\Delta}_{G/\leftrightarrow}$ contains two minimal nonfaces, namely $\{4, 6\}$ and $\{1, 3, 6, 8\}$. The complement of $\{1, 3, 6, 8\}$ with respect to the underlying vertex set of $\overline{\Delta}_{G/\leftrightarrow}$ is $\{4\}$. Consequently, city 4, the location of the convention, is *not* stochastically attainable from everywhere in the system.




Figure 25: The left panel shows the quotient graph $G/\leftrightarrow$ obtained by collapsing the strongly controllable subspaces of the graph G in Fig. 23 to single states (not shown are self-looping actions). The right panel shows a particular strategy for attaining state 6 in this quotient graph.




Figure 26: Source complex $\bar{\Delta}_{G/\leftrightarrow}$ of the quotient graph $G/\leftrightarrow$ of Fig. 25.




Figure 27: The graph $(G/\leftrightarrow)|W$ appears in the left panel, with $W = \{1, 3, 6, 8\}$ being a minimal nonface of $\overline{\Delta}_{G/\leftrightarrow}$ (see Fig. 26). An adversary can prevent attainment of state 4 by selecting transitions as in the right panel. The system then has full controllability within W relative to this adversarial choice.

The left panel of Fig. 27 shows the induced subgraph formed by flights available at the equivalence classes of cities appearing in the minimal nonface $\{1, 3, 6, 8\}$. Sure enough, adversarially, city 4 may become unreachable, as the right panel of Fig. 27 demonstrates.

Remark: The reader may have observed that in the right graph of Fig. 27, although city 4 becomes unreachable, the system can move to any city within the minimal nonface $\{1, 3, 6, 8\}$. This is in fact generally true: if some set W is a nonempty minimal nonface of $\overline{\Delta}_H$ for some graph H , then the system has full controllability within W *relative* to the assumption that it does not exit from W . In other words, it might happen that the system leaves W , but if not, for instance if an adversary makes nondeterministic choices to prevent motions out of W , then the system can move to any state within W from any other state in W . This is consistent with the observation we made in the proof of Theorem 16 that $\overline{\Delta}_{G|W} \simeq S^{|W|-2}$ for any such W .

12 Algorithms

This section provides algorithms for backchaining from a goal set and for computing a graph's strongly controllable subspaces. With these tools one can then reduce a graph to its potentially simpler quotient graph, as well as compute the source complex of the graph or of its quotient graph. One then obtains a characterization of the underlying system's capabilities, as discussed in the previous sections. We have implemented all these algorithms. Sample runs produced the figures in Section 11.4.

We do not provide an explicit algorithm for computing the strategy complex of a graph. Our homotopy results show that one does not need the strategy complex *per se* in order

to understand a system's capabilities. Moreover, the algorithm BACKCHAIN given next will produce a strategy for attaining any particular goal set S , or determine that no strategy exists. However, there certainly may exist strategies in a strategy complex other than those computed by backchaining, as Fig. 5 suggests. These extra strategies may be useful practically as backup strategies. The procedure outlined in Section 11.2, particularly starting with "Now suppose ..." on page 48, allows one to construct such strategies individually or as part of the full strategy complex.

Algorithm 1 BACKCHAIN(G, S)

Input: A stochastic graph $G = (V, \mathcal{A})$ and a nonempty subset S of V .

Output: A (possibly empty) set $\{v_1, \dots, v_k\}$ of states in $V \setminus S$ and a corresponding set $\{\sigma_1, \dots, \sigma_k\}$ of collections of actions in $\mathcal{A}$, such that every action in σ_i has source v_i and moves off the set $\{v_i, \dots, v_k\}$, for $i = 1, \dots, k$.
Moreover, if v is a state in $V \setminus S$ for which G contains a stochastic strategy for attaining S from v , then v will be one of the states v_i returned.

Procedure:

1. Let $W_0 := S$.
2. For $i = 1, 2, \dots$ until done do:
 - (a) Let v_i be any state in $V \setminus W_{i-1}$ such that some action with source v_i moves off $V \setminus W_{i-1}$. If no such v_i exists, then done.
 - (b) Otherwise, let σ_i be all actions in $\mathcal{A}$ with source v_i that move off $V \setminus W_{i-1}$.
 - (c) Let $W_i := W_{i-1} \cup \{v_i\}$.
3. Return $\{v_1, \dots, v_k\}$ and $\{\sigma_1, \dots, \sigma_k\}$, with k the maximum index for which v_i is defined. (If there is no such k , then return $\emptyset$ and $\emptyset$.)

Remarks: When Algorithm 1 is done, the collection of actions $\sigma = \bigcup_{i=1}^k \sigma_i$ is stochastically acyclic, constituting a strategy for attaining $V \setminus \{v_1, \dots, v_k\}$. In particular, if $W_k = V$ when Step 2 is done, then σ is a complete stochastic strategy for attaining S on the state space V . Conversely, if S is stochastically attainable in G , then W_k will be all of V , as the proof of Theorem 16 shows.

Aside: Applying Algorithm 1 to the graph of Fig. 5 with $S = \{3\}$ returns one or the other of the two triangles in the strategy complex shown in the figure, rather than merely the central edge. This is because Algorithm 1 grows the sets W_i one state at a time. Traditionally, one might add multiple states at once. Doing so would then produce the central edge as the backchained strategy in Fig. 5.

Given a goal S , the next algorithm uses backchaining to iteratively winnow a graph's state space V down to *the* maximal subspace W within which S is stochastically attainable. This procedure will be a useful step in determining a graph's strongly controllable subspaces.

(The reader may wish to verify that *the* is correct, that there really is a single maximal subspace W within which S is stochastically attainable.)

Algorithm 2 STRATEGY(G, S)

Input: A stochastic graph $G = (V, \mathcal{A})$ and a nonempty subset S of V .

Output: A stochastically acyclic subgraph $H = (W, \mathcal{B})$ of G such that $S \subseteq W$ and $V_{\mathcal{B}} = W \setminus S$, with W maximal among all such W .

Procedure:

1. Let $\{v_1, \dots, v_k\}, \{\sigma_1, \dots, \sigma_k\}$ be the results returned by BACKCHAIN(G, S).
2. Let $W := \{v_1, \dots, v_k\} \cup S$ and $\sigma := \bigcup_{i=1}^k \sigma_i$.
3. If $W = V$, return the graph (W, σ) .
4. Otherwise, recursively call STRATEGY($(W, \mathcal{C}), S$), where $\mathcal{C}$ consists of all actions in $\mathcal{A}$ whose sources *and* targets lie in W .

Given a stochastic graph $G = (V, \mathcal{A})$, the next algorithm computes a regular directed graph (V, E) such that (v, u) is an edge in E if and only if G contains a stochastic strategy for attaining u from v , with v distinct from u . (See again the definitions on page 31.)

Algorithm 3 REACHABLE(G)

Input: A stochastic graph $G = (V, \mathcal{A})$, with $V \neq \emptyset$.

Output: A directed graph $D = (V, E)$, with specifications as above.

Procedure:

1. For each $u \in V$, let $H_u := (W_u, \mathcal{B}_u)$ be the result returned by STRATEGY($G, \{u\}$).
2. Let $E := \{(v, u) \in V \times V \mid v \in W_u, v \neq u\}$.
3. Return (V, E) .

The next algorithm computes the strongly controllable subspaces of G .

Algorithm 4 SUBSPACES(G)

Input: A stochastic graph $G = (V, \mathcal{A})$, with $V \neq \emptyset$.

Output: The set $\{W_1, \dots, W_m\}$ of equivalence classes of $\leftrightarrow$ for G .

Procedure:

1. Let $\{V_1, \dots, V_\ell\}$ be the strongly connected components of the directed graph returned by REACHABLE(G). See [1] for an algorithm to compute these components.
2. If $\ell = 1$, return $\{V_1\}$.
3. Otherwise, for each $i = 1, \dots, \ell$:
 - (a) Let $H_i := (V_i, \mathcal{C}_i)$, with $\mathcal{C}_i$ all actions of $\mathcal{A}$ whose sources and targets lie in V_i .
 - (b) Let $\mathcal{V}_i$ be the result of calling SUBSPACES(H_i) recursively.

Return $\bigcup_{i=1}^{\ell} \mathcal{V}_i$.

Remarks:

- **Correctness.** It is not difficult to prove that the previous algorithms correctly compute output as specified.
- **Runtime.** The previous algorithms all run in almost-reasonable polynomial time. For instance, the slowest, SUBSPACES(G), has time-complexity $O(|V|^5|\mathcal{A}|)$, with $G = (V, \mathcal{A})$. No doubt, faster implementations exist.
- **Quotient Graphs.** Computing the quotient graph $G/\leftrightarrow$ from a graph G entails calling SUBSPACES(G), then relabeling sources and targets of actions. Specifically, if SUBSPACES(G) returns $\{W_1, \dots, W_m\}$, then one relabels every state w in W_i by some representative state for W_i , with $i = 1, \dots, m$.
- **Source Complexes.** Computing the source complex $\overline{\Delta}_G$ of a graph G is conceptually straightforward: For every possible goal set S , one calls BACKCHAIN(G, S). To see this:

As we observed earlier, if the set of states $\{v_1, \dots, v_k\}$ returned by BACKCHAIN(G, S) is equal to $V \setminus S$, then S is a stochastically attainable goal. Always, $\{v_1, \dots, v_k\}$ is a simplex of $\overline{\Delta}_G$. Conversely, any simplex of $\overline{\Delta}_G$ must show up as the start region of some minimal complete stochastic strategy for attaining some goal S in G . By itself, without other actions of G , that strategy necessarily looks like a backchained strategy, as can be seen via recursive application of Lemma 15. Consequently, given other actions of G , backchaining may produce a different strategy, but it will produce some complete stochastic strategy for attaining S in G .

This naïve algorithm has running time exponential in the size of V . One can tinker with the procedure to make it slightly more efficient, but the exponential nature is likely to be fundamental, as the first hardness result of Section 14 indicates.

13 Realizability

We have previously indicated that finite simplicial complexes and stochastic graphs are essentially equivalent objects from a topological perspective. In fact, this equivalence depends on nondeterminism not stochasticity. The underlying reason is that finite simplicial complexes are purely combinatorial and thus are equivalent to finite posets [6, 75]. Nondeterministic graphs alone, without stochastic actions, are able to capture all finite posets.

This section makes precise these comments with two realizability theorems. The first theorem says that for any finite simplicial complex Σ , there is some strategy complex isomorphic to the first barycentric subdivision of Σ . The second theorem says that Σ may actually be realized exactly as a source complex of some graph. In both cases, the generating graphs are nondeterministic.

Reminder: $\text{sd}(\Sigma)$ means the first barycentric subdivision of Σ (Section 5.4).

Theorem 21 (Realization by Strategy Complexes) *Let Σ be a finite simplicial complex. There exists a nondeterministic graph G such that $\text{sd}(\Sigma) \cong \Delta_G$.*

Proof. If $\Sigma = \emptyset$, then $\text{sd}(\Sigma) = \emptyset$, so we may let $G = (\emptyset, \emptyset)$.

Otherwise, define $G = (V, \mathcal{A})$ as follows: V consists of all vertices in the complex $\text{sd}(\Sigma)$ along with one new vertex, $\perp$. $\mathcal{A}$ contains exactly one nondeterministic action A_v at every $v \in V$ other than at $\perp$. $A_v = \{(v, u) \mid u \in V \text{ and } \{v, u\} \text{ is not a simplex of } \text{sd}(\Sigma)\}$. In other words, A_v has transitions to all vertices that are not adjacent or equal to v in the barycentric subdivision of Σ . Observe that every action A_v contains at least one edge, namely a transition from v to $\perp$, so this definition is well-formed. Moreover, no action contains a self-loop.

The map $v \mapsto A_v$ is a bijective correspondence between the vertices of $\text{sd}(\Sigma)$ and those of Δ_G . In order to establish that $\text{sd}(\Sigma) \cong \Delta_G$, we need to show that this correspondence preserves simplices. Definitionally, both $\text{sd}(\Sigma)$ and Δ_G contain the empty simplex.

(Aside: If $\Sigma = \{\emptyset\}$, then $G = (\{\perp\}, \emptyset)$ and $\Delta_G = \{\emptyset\} = \text{sd}(\Sigma)$.)

(a) Suppose $\emptyset \neq \sigma \in \text{sd}(\Sigma)$. $\sigma = \{v_1, \dots, v_k\}$, with $v_i \in V \setminus \{\perp\}$. Let A_i be the action of G defined at v_i , for $i = 1, \dots, k$. No v_i can appear as the target of any A_j , since $\{v_i, v_j\}$ is a simplex of $\text{sd}(\Sigma)$. That means $\{A_1, \dots, A_k\}$ is a simplex of Δ_G .

(b) Suppose $\emptyset \neq \tau \in \Delta_G$. $\tau = \{A_1, \dots, A_k\}$, with each $A_i \in \mathcal{A}$. Let v_i be the source of A_i , for $i = 1, \dots, k$. Each v_i is the barycenter of some nonempty simplex $\sigma_i \in \Sigma$.

Consider the set $\{v_i, v_j\}$ formed by any two distinct such vertices. This set must be a simplex of $\text{sd}(\Sigma)$ as otherwise $\{A_i, A_j\}$ could generate a cycle in G . Consequently, either σ_i is a proper face of σ_j or vice-versa. So, without loss of generality, $\emptyset \subset \sigma_1 \subset \sigma_2 \subset \dots \subset \sigma_k$, all inclusions being proper. In turn, that means $\{v_1, \dots, v_k\}$ is a simplex of $\text{sd}(\Sigma)$. $\square$

Theorem 22 (Realization as Source Complexes) *Let Σ be a finite simplicial complex. There exists a nondeterministic graph G such that $\Sigma = \overline{\Delta}_G$.*

Proof. If $\Sigma = \emptyset$, let $G = (\emptyset, \emptyset)$. Then $\overline{\Delta}_G = \emptyset$.

Otherwise, let V be the vertices of Σ . If Σ is the complex of the full simplex on V , enlarge V by one new state, $\perp$. Let $\overline{\Sigma}$ be the set of maximal simplices of Σ . Observe: By construction of V , $V \notin \overline{\Sigma}$.

Construct $G = (V, \mathcal{A})$ as follows: For every $X \in \overline{\Sigma}$ and every $v \in X$, define the nondeterministic action $A_{v,X} = \{(v, u) \mid u \in V \setminus X\}$. Each such action contains at least one edge, since X cannot be all of V . Let $\mathcal{A}$ consist of all such actions.

(Aside: If $\Sigma = \{\emptyset\}$, then $G = (\{\perp\}, \emptyset)$ and $\overline{\Delta}_G = \{\emptyset\}$.)

(a) Suppose $\emptyset \neq X \in \Sigma$. We can assume without loss of generality that $X \in \overline{\Sigma}$. Consider the set of actions $\sigma = \{A_{v,X} \mid v \in X\}$. No source of an action in σ is the target of an action in σ , so $\sigma \in \Delta_G$. Thus $X = V_\sigma \in \overline{\Delta}_G$.

(b) Suppose $\emptyset \neq X \in \overline{\Delta}_G$. Let $\sigma \in \Delta_G$ such that $V_\sigma = X$. By Lemma 15, σ contains an action that moves off X . The action must be of the form $A_{v,Y}$, for some $Y \in \overline{\Sigma}$, with $v \in Y$. By the definition of “moves off”, $v \in X$ and all targets of $A_{v,Y}$ lie outside X . By construction, the targets of $A_{v,Y}$ are all the states of $V \setminus Y$. That means $X \subseteq Y$. So X is a simplex of Σ , since Y is.

Both Σ and $\overline{\Delta}_G$ contain the empty simplex. $\square$

Remark: The underlying vertex sets of Σ and $\overline{\Delta}_G$ may differ slightly, but the simplices in the two complexes are identical.

14 Hardness

This section presents some simple hardness results. The first result indicates that the exponential computation of $\overline{\Delta}_G$ described in Section 12 is fundamental, assuming NP and P are different complexity classes. The other two results suggest that homotopy equivalence of strategy complexes is either a stronger topological condition than one needs for understanding system capabilities or that homotopy equivalence encodes some interesting properties about uncertain systems beyond those one usually asks at the fixed graph level. It remains a research question to know which of these perspectives is correct.

14.1 The Difficulty of Determining a System’s Precision

The first result has practical implications. Given an uncertain system, it may be very difficult to determine just how precise the system can be.

Lemma 23 (Precision is Hard) *Let $G = (V, \mathcal{A})$ be a stochastic graph. Determining the size of the smallest stochastically attainable goal is NP-complete.*

Proof. The proof is by a reduction from INDEPENDENT SET. The underlying ideas are related to the realizability ideas from Section 13.

We define the following two problems:

INDEPSET: Given an undirected graph (V, E) and an integer ℓ , is there a set $I \subseteq V$ of size ℓ such that no two states in I constitute an edge in E ?

PRECISION: Given a stochastic graph $G = (V, \mathcal{A})$ and an integer k , is there some set $S \subseteq V$ of size k such that G contains a complete stochastic strategy for attaining S ?

INDEPSET is a known *NP*-complete problem [49, 66].

I. PRECISION lies in *NP*:

Given S , one can verify in polynomial time that S has size k and that G contains a complete stochastic strategy for attaining S , using Algorithm 1.

II. PRECISION is *NP*-hard:

Suppose (V, E) is an undirected graph and ℓ an integer. Let $n = |V|$.

Define a nondeterministic graph $G = (V \cup \{\perp\}, \mathcal{A})$ as follows: $\perp$ is a new state. For every $v \in V$, $\mathcal{A}$ contains a nondeterministic action $A_v = \{(v, u) \mid u = \perp \text{ or } (v, u) \text{ is an edge in } E\}$. So A_v transitions to every state adjacent to v in the original undirected graph, as well as to $\perp$ (ensuring A_v is well-formed).

(a) Suppose I is an independent set of size ℓ in (V, E) .

Consider the set of actions $\sigma = \{A_i \mid i \in I\}$. No source of an action in σ is the target of an action in σ , by definition of I . So $\sigma \in \Delta_G$. That means G contains a complete strategy for attaining $\{\perp\} \cup V \setminus I$, which is a set of size $n + 1 - \ell$.

(b) Suppose G contains a complete strategy for attaining set S of size k .

S necessarily contains $\perp$. There is some $\sigma \in \Delta_G$ such that $V_\sigma = V \setminus S$ and $|V_\sigma| = n + 1 - k$. V_σ is an independent set in (V, E) . To see this, suppose otherwise. Then σ must contain actions A_v and A_u such that (v, u) is an edge in E . That would mean σ could generate a cycle in G , a contradiction.

To summarize: (V, E) contains an independent set of size ℓ if and only if G contains a stochastically attainable goal of size k with $k + \ell = n + 1$. One may construct G from (V, E) in polynomial time. Thus PRECISION is *NP*-hard. $\square$

14.2 Small Realization is Uncomputable

There are classic deep results telling us it is algorithmically impossible to recognize the homeomorphism or homotopy type of even a finite general simplicial complex [76]. These results have their roots in work on the algorithmic impossibility of deciding whether a word in a finitely presented group is trivial [63]. One particular implication is that it is algorithmically uncomputable to decide whether a finite simplicial complex is contractible (this follows from

[64]). We present two consequences for nondeterministic (and thus stochastic) graphs.

The following lemma says: Finding the smallest graph whose strategy complex is homotopic to a given complex is an uncomputable problem. Compare this with the realizability results of Section 13.

Lemma 24 (Compact Realization) *The following question is undecidable:*

Given a finite simplicial complex Σ and a nonnegative integer m , is there a stochastic graph G with m actions such that $\Delta_G \simeq \Sigma$?

Proof. Observe: If G is a stochastic graph containing a single action A , then either $\Delta_G = \{\emptyset\}$ or $\Delta_G = \{\emptyset, \{A\}\}$, depending on whether A contains a (non)deterministic self-loop or not, respectively.

We now reduce from the uncomputability of deciding contractibility:

Suppose Σ is a finite simplicial complex. We can check whether it is trivial (void or empty) easily. Assuming it is not trivial, we ask whether there is a stochastic graph G containing a single action such that $\Delta_G \simeq \Sigma$. The answer is “yes” if and only if Σ is contractible. $\square$

14.3 Recognizing Repercussions is Uncomputable

Suppose we are given a stochastic graph $G = (V, \mathcal{A})$ that somehow does not satisfy some criterion we desire. We wish to understand whether adding a particular new action A , with source and targets in V , will satisfy our criterion. We can certainly compute the strategy or source complexes of the resulting graph, to see whether any larger start regions or more precise goals have been created. What we cannot do, in general, is decide whether the homotopy type of the strategy (or source) complex has changed, as the following lemma indicates.

Lemma 25 (Detecting Change) *The following question is undecidable:*

Given a stochastic graph and a new action, does the graph’s strategy complex change homotopy type when one adds the action to the graph?

Proof. Suppose Σ is a finite simplicial complex. Again, we can check whether it is trivial (void or empty) easily. Assuming it is not trivial, we construct a graph G much as in the proof of Theorem 22, except that we always add *two* new states, $\perp$ and $\top$. So the graph looks like $G = (V, \mathcal{A})$, with V the union of $\{\perp, \top\}$ and the vertices of Σ . The set of actions $\mathcal{A}$ is computed as in the proof of Theorem 22. Observe that there are no actions at either $\perp$ or $\top$. As before, one obtains exact equality: $\Sigma = \overline{\Delta}_G$.

Now consider adding to G a new action A , representing a deterministic transition from $\top$ to $\perp$. Call the resulting graph $\widehat{G}$. Then $\Delta_{\widehat{G}}$ is a cone with apex A , hence contractible.

- In summary:
- $\Sigma = \overline{\Delta}_G \simeq \Delta_G$ (by proof of Theorem 22 and by Theorem 14).
 - $\Delta_G \simeq \Delta_{\widehat{G}}$ if and only if Δ_G is contractible (by construction).

So Σ is contractible if and only if $\Delta_G \simeq \Delta_{\widehat{G}}$.

Recognizing contractibility is undecidable. $\square$

15 Topological Thinking

This section discusses some scenarios in which topological thinking reveals the essence of a problem and its solution.

15.1 Topology Precompiles an Existence Argument

Suppose $G = (V, \mathcal{A})$ is a stochastic graph with associated nonnegative action transition times $\{\delta_A\}_{A \in \mathcal{A}}$. Let S be a desired stop set (goal states) within V and let T be some nonnegative time.

Suppose for every u in $V \setminus S$, someone has produced a strategy σ_u by which the system will converge to S when started at u , with worst-case expected-convergence time for u no greater than T . (In other words, $0 \leq t_u \leq T$ in the solution of System (1) on p. 22, written out for σ_u . We make no assertions concerning σ_u 's times t_v for states v other than u .)

We ask the question: Is there a complete stochastic strategy σ that converges to S from all states of V with all worst-case expected-convergence times no greater than T ?

Intuitively, we certainly expect the answer to be “yes”. After all, one can imagine a meta-strategy that determines the initial state u , then executes strategy σ_u . However, this dispatch-based strategy is not a strategy in our sense. Our strategies have the flavor of feedback-based control laws, that is, they are mappings from states to (sets of) actions. The question therefore is whether we can flatten the dispatch-based strategy to be a true strategy. Observe the difficulties:

Let v and u be distinct states in $V \setminus S$.

Issue 1: It could very well be that σ_u contains no actions at state v .

Issue 2: Even if σ_u contains one or more actions at state v , it could be that $t_v > T$ when we solve System (1) written out for σ_u .

Issue 3: The two strategies σ_u and σ_v may contain inconsistent actions, meaning that $\sigma_u \cup \sigma_v$ might not be stochastically acyclic. For instance, one strategy might move left through a hallway and the other right.

One can make a backchaining argument that some strategy σ exists for attaining S from everywhere in V . Arguing that among all such strategies there is at least one for which all worst-case expected convergence times are bounded by T is possible, but tedious.

There is a short topological argument:

- Write $V = [n]$ and $S = \{n\}$, without loss of generality.
- Consider the loopback graph $G_{\leftarrow n} = (V, \mathcal{A}')$, as in the proof of Theorem 1. Associate transition time 0 to every loopback action.
- By Lemma 10, the affine cover $\bigcup_{A \in \mathcal{A}'} U_A^+$ contains all sets $\{\mathbf{x} \in \mathbf{R}^n \mid x_i > x_n + T\}$ and all sets $\{\mathbf{x} \in \mathbf{R}^n \mid x_n > x_j\}$. Thus it contains all sets $\{\mathbf{x} \in \mathbf{R}^n \mid x_i > x_j + T\}$, with i, j in V . Using Theorem 8 and the structure of $\partial C_{T+\epsilon}$,

$$\Delta_{G_{\leftarrow n}}^T \simeq \bigcup_{A \in \mathcal{A}'} U_A^+ \cap \partial C_{T+\epsilon} \simeq S^{n-2}.$$

- By the remark at the end of Section 8 (p. 35), Δ_G^T contains a complete stochastic strategy for attaining S . In particular, all worst-case expected convergence times of this strategy are bounded by T , as desired.

The intuitive summary:

- Strategy existence depends on cover connectivity in $\mathbf{R}^n$.
- Adding strategies, even partial or pairwise inconsistent strategies, can only improve cover connectivity.

Topology has precompiled an existence argument: Bellman's Principle of Optimality shows up as the union of open sets.

15.2 Topological Analysis of Adversity

In looking at examples, the reader may have begun to wonder where adversity sits: In the stochastic nature of actions or in the nondeterministic nature of actions? Topology helps us understand an answer to that question.

The definition of *moves off* on page 41 makes explicit a difference between stochastic and nondeterministic actions: A stochastic action is almost the same as a set of deterministic actions. The following definition and lemma make this idea precise:

Definition: If $G = (V, \mathcal{A})$ is a stochastic graph, let $\det(G) = (V, \det(\mathcal{A}))$ designate the nondeterministic graph in which every stochastic action of G has been replaced by a set of deterministic actions. Specifically, let us define $\det(\mathcal{A})$ as follows:

- $\det(\mathcal{A})$ contains every deterministic or nondeterministic action of $\mathcal{A}$.
- For every stochastic action $A \in \mathcal{A}$ of the form $A = \{(v, p_1 u_1), \dots, (v, p_k u_k)\}$, $\det(\mathcal{A})$ contains k deterministic actions, each consisting of a single edge (v, u_i) , with $i = 1, \dots, k$. (As usual, we permit redundancies, that is, distinct actions with identical edge sets.)

Lemma 26 (Stochastic Determinism) *Let $G = (V, \mathcal{A})$ be a stochastic graph. Then*

$$\Delta_G \simeq \Delta_{\det(G)} \quad \text{and} \quad \overline{\Delta}_G = \overline{\Delta}_{\det(G)}.$$

Proof. The first assertion follows from the second, by Theorem 14.

One can establish the second assertion by using Algorithm 1 and Lemma 15. $\square$

Conclusion: From the perspective of homotopy type or from the perspective of understanding a system's stochastically attainable goals, one may study $\det(G)$ in place of G . Fig. 28 provides an example, demonstrating homotopy equivalence of Δ_G and $\Delta_{\det(G)}$. Of course, convergence times in the two graphs may very well be different, so one must be careful not to oversimplify.




Figure 28: The top row shows a pure stochastic graph and its strategy complex. The bottom row shows the deterministic rendition of this graph plus the resulting strategy complex. Both complexes are homotopic to S^1 .

Pursuing this further, we can ask where true adversity resides. The realizability theorems of Section 13 all used nondeterministic graphs, as did the *NP*-completeness result of Lemma 23. Suppose G is a *pure* stochastic graph, in the sense that it contains only deterministic and stochastic actions but no multi-edged nondeterministic actions. Then Lemma 26 and Theorem 20 allow us to conclude that Δ_G is homotopic either to a sphere or to a point. Pure stochastic graphs simply cannot attain the richness of strategy complexes that nondeterministic graphs attain. Said differently, true adversity comes from nondeterministic adversaries. Stochastic adversaries can slow a system down, but do not have the ability to switch between limit cycles in the way nondeterministic adversaries sometimes can (limit cycles are the minimal nonfaces of $\overline{\Delta}_G$; see also the remark on page 55).

Some further remarks:

Suppose again that G is a pure stochastic graph. When Δ_G and $\Delta_{\det(G)}$ are homotopic to a sphere, that sphere has dimension $n - k - 1$, where n is the number of states in G and k is the number of strongly connected components of $\det(G)$. In particular, if $\det(G)$ can be written as the disjoint union of its strongly connected components, then Δ_G and $\Delta_{\det(G)}$ will be homotopic to a sphere; otherwise, they will be contractible, as follows from [44] or Theorem 20. Observe that k is no greater than the number of strongly controllable subspaces of G and in fact can be less, but must be the same when Δ_G and $\Delta_{\det(G)}$ are homotopic to a sphere.

When passing to quotient graphs, let us discard any actions that become cyclic (deterministically self-looping). Then the graph $\det(G)/\leftrightarrow$ is a directed acyclic graph. All the actions of this graph together are acyclic and thus the complex $\Delta_{\det(G)/\leftrightarrow}$ consists of a single simplex (all acyclic actions of $\det(G)/\leftrightarrow$) and its faces. It is empty precisely when $\det(G)$ can be written as the disjoint union of its strongly connected components.

The graph $G/\leftrightarrow$ is similar, yet slightly different. As we observed already, the state space of $G/\leftrightarrow$ may be larger than that of $\det(G)/\leftrightarrow$. With regard to strategies, as was the case for $\Delta_{\det(G)/\leftrightarrow}$, $\Delta_{G/\leftrightarrow}$ consists of a single simplex (all acyclic actions, now of $G/\leftrightarrow$) and its faces. A difference lies in certainty. $G/\leftrightarrow$ represents a collection of Markov chains, with all states either transient or absorbing. Whereas in $\det(G)/\leftrightarrow$ the system may be able to move purposefully from some state to some other state, in $G/\leftrightarrow$ it may only be able to do so with some probability; the adversary may have some stochastic choices by which to influence the outcome of motions. This is a simple Markov Decision Process [67].

15.3 Topological Thinking in Partially Observable Spaces

This section shows by example how the topological characterization of task solvability may help one decide whether a task has a guaranteed solution. Throughout this section all graphs are nondeterministic.

Implicitly in the analyses thus far, the system has always known its current state. If in fact the system's sensing function is imperfect, then one may redefine the system's state space so as to obtain perfect sensing in the redefined state space. The construction is straightforward [5, 54], albeit often with added complexity.

The nondeterministic version of this construction leads to a derived nondeterministic graph, sometimes called a *knowledge space* [23, 25]. A sensorless version of knowledge space appeared as the search space for the part orienter of [29].

Intuitively, the knowledge space of a graph is a new graph whose states are sets of potential robot locations consistent with the robot's sensing and action history at runtime. For completeness, we give a general definition here.

Definitions Suppose $G = (V, \mathcal{A})$ is a nondeterministic graph.

- For the purposes of the next few definitions, we assume that every action in $\mathcal{A}$ is labeled with a *name*, which we refer to notationally as A . Names of actions at a given state are unique. However, different actions at different states might have the same name. (Think of a global control with possibly different effects at different states.)

We write A_v to mean the action in $\mathcal{A}$ whose name is A and whose source is v , whenever such an action exists.

- A *sensor on V* is a covering Ξ of V by nonempty subsets of V , each of which is called a *sensory interpretation set*. Sensory interpretation sets may overlap.

Intuition: Whenever the system is at some state v , a sensor returns some value whose interpretation is a set $I \in \Xi$ such that $v \in I$. (We could model the sensor via some intermediate sensory space, but we will simply map directly to interpretations of the sensor. More general formulations of sensing appear, for instance, in [5, 54].) Sensing in our case is nondeterministic, meaning that the sensor could return any $I \in \Xi$ with $v \in I$. Observe that $\Xi = \{V\}$ is equivalent to no sensing.

- Suppose $X \subseteq V$ and A is an action name. The *forward projection of X under A* , written $F_A(X)$, is the set of all possible locations the system might move to after executing some action named A , given that the system starts in X . In order for this definition to be sensible, the name A must actually mean something at every state in X . Otherwise, by convention, $F_A(X)$ is not defined. Formally, if A_v exists for every $v \in X$, then

$$F_A(X) = \{u \in V \mid (v, u) \in A_v \text{ for some } v \in X\}.$$

- Once the system moves from X under A , it can sense again. The sensor may return any sensory interpretation set $I \in \Xi$ for which $F_A(X) \cap I$ is not empty. If the system knows it is in set X before moving under A , then after moving and sensing a particular I , the system knows it is somewhere in the set $F_A(X) \cap I$.
- Given G and Ξ as above, define a new nondeterministic graph $K = (V^+, \mathcal{A}^+)$ as follows:
 - V^+ consists of all nonempty subsets of V . (Depending on Ξ and $\mathcal{A}$, one may not really need to construct all of V^+ , of course.)
 - For every $X \in V^+$, and for every action name A such that A_v exists for all $v \in X$, $\mathcal{A}^+$ contains a nondeterministic action with source X given by:

$$\{(X, Y) \mid \emptyset \neq Y = F_A(X) \cap I \text{ for some } I \in \Xi\}.$$

K is the *knowledge space* associated with G and Ξ .

Remark: One can define stochastic analogues of K for pure stochastic graphs G . This leads to the so-called *belief states*, of great interest in Partially Observable Markov Decision Processes [54]. Some caution is advised, since problems quickly become intractable [67, 58].

15.3.1 Inferring Task Unsolvability From Duality

Recall Theorem 16. Failure of the theorem's contractibility condition is evidence of a *potentially inescapable cycle*, appearing as a minimal nonface of $\overline{\Delta}_G$, indicating that the graph contains no complete guaranteed strategy for attaining set S . We saw an example in Fig. 27.

Now consider the graph of Fig. 29, which might model a robot moving in a building. There are two corridors (states 1 and 2). In any one corridor the robot can move RIGHT or LEFT. Atriums connect the corridors at either ends. The task is to reach one particular atrium (state 3). Entry into the corridors from the other atrium (state 4) is imprecise. The gray triangle




Figure 29: A graph, an imperfect sensor, and a potentially inescapable cycle between knowledge states $\{4\}$ and $\{1, 2\}$ (indicated by the solid arrows in the right panel).




Figure 30: With perfect sensing, the strategy complex of the state graph of Fig. 29 consists of a solid triangle joined to an edge of a solid tetrahedron. With sensing ambiguity at states 1 and 2, this strategy complex collapses, inducing a non-contractible subcomplex of strategies in knowledge space that signals a potentially inescapable cycle.

in the strategy complex of Fig. 30 constitutes a strategy for accomplishing this task, assuming perfect sensing. The strategy is to ENTER from state 4, move RIGHT from 1, LEFT from 2.

Now imagine a robot controller unable to distinguish the two corridors (states 1 and 2) based on sensing alone. The task no longer has a guaranteed solution. One could see this in a variety of ways, for instance, by explicitly constructing the robot’s knowledge space, part of which is shown in the right panel of Fig. 29. Let us take a related but more topological perspective.

One does not need to construct the knowledge space directly. Instead, one can reason about strategies. Consider the strategy complex of the original graph (left panel, Fig. 30). As we will see, this complex induces a non-contractible subcomplex of strategies in knowledge space (right panel, Fig. 30) that violates the contractibility condition of Theorem 16.

Let us focus on the actions at two key knowledge states, namely $\{4\}$, representing certainty that the system is at state 4, and $\{1, 2\}$, representing uncertainty as to whether the system is at state 1 or 2. The tetrahedron of the original complex describes the strategies possible with perfect sensing at the graph states 1 and 2. The tetrahedron collapses to an edge under sensing ambiguity. This edge represents the two actions, RIGHT and LEFT, possible at state $\{1, 2\}$ in knowledge space. In the original complex, only a portion of the tetrahedron joins with the action ENTER; action RIGHT at state 2 and action LEFT at state 1 do not join with ENTER. In the knowledge space complex, the edge $\{\text{RIGHT}, \text{LEFT}\}$ consequently *cannot join* with the action ENTER.

We have thus exhibited a non-contractible complex describing the strategies available at a subset of knowledge space, with that subset lying in the complement of the goal (state 3). This means, *no matter what the surrounding knowledge space might look like*, there can exist no complete guaranteed strategy for attaining state 3 in the presence of control uncertainty at state 4 and sensing confusion between states 1 and 2.

15.3.2 Hypothesis-Testing and Sphere Suspension

Definition: A *suspension* of a complex is another complex formed by joining each simplex of the given complex with each of two new vertices [61, 39]. For instance, the complex in Fig. 4 is a suspension of the complex in Fig. 2. The key property relevant to us currently is that the suspension of a sphere of any dimension is another sphere, of one higher dimension. For instance, (a globe of) the Earth is a suspension of the Equator by the North and South Poles.

Fig. 31 shows a variant of the example of Fig. 29. Once again there is some control uncertainty, once again the sensor cannot distinguish certain corridors. The system cannot move reliably to state 3 using pure feedback control (that is, using sensing alone, without history). For some sensor values, neither of the actions RIGHT or LEFT will make progress toward the goal 3 at all graph states consistent with the sensor value.

Fortunately, this time the confusable corridors lie in different “wings” of the building. The sensor stratifies the graph into two subgraphs, both containing the goal 3. Within each subgraph the sensor is effectively perfect and each subgraph contains a strategy guaranteed to attain the goal from anywhere within that subgraph using sensor-indexable controls that are well-defined across subgraphs.

Whenever a sensor stratifies a graph into subgraphs in this manner, there exists a hypothesis-testing strategy for attaining the goal from anywhere in the overall graph. Hypothesis-testing means: The system assumes it is in one of the subgraphs; it commands




Figure 31: An imperfect sensor stratifies a graph into two subgraphs (the “a wing” and the “b wing”), over each of which sensing is effectively perfect.

actions and interprets sensor readings as if it really were in that subgraph, but also verifies consistency between predicted motions and observed sensor readings. If an inconsistency occurs, the hypothesis of being in that subgraph has been falsified and the system moves on to another hypothesis. Intuitively, this strategy eventually converges at the goal.

Hypothesis-testing is a strategy in knowledge space, but one does not need to construct knowledge space. In general, knowledge space may contain additional, possibly shorter, strategies. Fortuitously, for the example of Fig.31, hypothesis-testing is effectively the only strategy.

There is a short topological argument that hypothesis-testing converges:

Hypothesis-testing amounts to repeated sphere suspension.

Further details: The graph H_i describing motions under the i^{th} hypothesis is equivalent to the subgraph G_i being hypothesized, except that some actions may move nondeterministically to a new state $\perp_i$, signaling falsification of the hypothesis. We also add an “action” from $\perp_i$ to the goal, as explained below. Without loss of generality, every action of G_i , and thus H_i , contains a *nondeterministic* transition to the goal. (Adding these transitions does not affect the existence of strategies for attaining the goal, but focuses on the topologically significant simplices in the complexes.)

The relevant loopback complex Γ_{H_i} of H_i is then a suspension, formed by joining the relevant loopback complex Γ_{G_i} of G_i with the loopback action $\text{goal} \rightarrow \perp_i$ and the action $\perp_i \rightarrow \text{goal}$. We now use Theorem 1, once in each direction: Since G_i contains a strategy guaranteed to attain the goal from any state in G_i , Γ_{G_i} is homotopic to a sphere, of the correct dimension. Since Γ_{H_i} is a suspension of Γ_{G_i} , it too is a sphere, again of the correct dimension, proving that hypothesis-testing converges.

What is the mysterious action $\perp_i \rightarrow \text{goal}$? It is the inductive-hypothesis that there exists a strategy for attaining the goal once the i^{th} graph-hypothesis has been falsified! (The base case is similar, except that there is no need for a state $\perp_i$. The final graph-hypothesis is certain.)

Remarks: (1) One can make a very similar argument directly at the graph level, further underscoring the parallel between spheres and task solvability. That parallel shows as well that backchaining in nondeterministic graphs is much like repeated sphere suspension. (2) Hypothesis-testing is related to randomization [23, 24, 25].

16 Conclusions

16.1 Summary

The controllability structure Theorem 20 on page 51 is the most concise description of system capabilities we have currently. Implicitly, this theorem contains as special cases the full controllability Theorem 11 and the goal attainability Theorem 1. These three theorems characterize controllability of uncertain systems by the existence of spheres of a specific dimension. Theorem 16 provides a dual statement via contractibility of subcomplexes. That theorem helps one understand an adversary’s capabilities in terms of forced limit cycles.

From a practical perspective, Theorem 14 on page 42 is the most useful. It suggests that when considering homotopy equivalences, one should take the homotopy simplification of a strategy complex down to the state space, and not further. That rendition provides a natural mechanism by which to compare the capabilities of different systems with a common state space, as illustrated by the design example of Sections 9 and 10.

16.2 Other Results

We have omitted many interesting results from this paper. The Contributions section (§1.6) mentioned several. For instance, the number of nonempty simplices in the loopback complex associated with a particular goal is even if the goal is stochastically attainable and odd if it is not. This is a consequence of the Euler characteristic being even for spheres and odd for contractible spaces.

There also exist generalizations of the results presented, by reduction to them. For instance, perhaps one is interested in whether a goal is attainable from some initial subset of the state space rather than from the entire state space. One can reduce this problem to a full attainability question by modifying the graph slightly. In the nondeterministic case this is easy to see: one merely adds a nondeterministic *hyperjump* at each state, that jumps to every state in the initial region. These hyperjumps will never be needed by an actual strategy in the original system but serve as a convenient reduction tool.

Similarly, perhaps one is merely interested in reaching a goal with some minimum probability, but not necessarily for certain. Again, one can add hyperjumps at some states, assigning strictly positive transition times to these hyperjumps and zero transition times to all other actions. One thus reduces the original problem to the problem of attaining a goal with certainty, but now with a bound on the worst-case expected convergence times.

16.3 Open Questions

Our original goal was to develop a topological language for describing system capabilities. We have come a long way. Strategy and source complexes abstract away the detailed connectivity between states in a graph while retaining information about overall goal attainability.

Consequently, one can assign topological spaces to graphs, much like labels, that succinctly summarize the graphs’ capabilities.

One wonders whether this new language holds more expressive power than we have used thus far. Our key theorems are not yet completely “coordinate-free”: The number of states shows up in the dimensionality of the spheres that assure controllability. Likewise, we find ourselves working with source complexes on specific state spaces, rather than more abstractly.

Hidden inside this observation is another: Our contractibility results tend to have the flavor of collapsibility (see Section 5.2). This begs the question whether the more general (and more difficult) form of contractibility is relevant, and if so, what it is telling us.

Related to this general interpretation of contractibility is the question of whether our results generalize to continuous state spaces. As we observed in the Introduction, from a practical perspective one tends to work in discrete spaces. Nonetheless, the question is interesting intellectually. For instance, imagine a sequence of finite discrete graphs that approximate ever more finely a continuous space. An amusing issue here is that finite-dimensional spheres are homotopically different from points, whereas the infinite-dimensional sphere is homotopically equivalent to a point.

17 Acknowledgments

Many thanks to Ben Mann for his leadership, energy, and vision, which made this research possible, enjoyable, and addictive.

This research was part of a multi-year multi-institutional effort involving topologists and roboticists working on a variety of problems at the interface of algebraic topology and autonomous systems. Many thanks to the entire “SToMP” group for their feedback and encouragement. Particular thanks to Robert Ghrist and Steven LaValle for their infectious enthusiasm and indefatigable leadership. Thanks to Robert Ghrist and Shmuel Weinberger for answering my topology questions so quickly and precisely. Thanks to Matt Mason for advice, as always, on a large number of topics.

List of Primary Symbols

Symbol	Meaning	Page Reference
G, H	nondeterministic or stochastic graph	11, 21
V	set of all states in a graph (state space)	11, 21
W	some set of states in a graph	12, 21
A, B	nondeterministic or stochastic action in a graph	11, 20
$\mathcal{A}, \mathcal{B}, \mathcal{C}$	collection of nondeterministic or stochastic actions	11, 21
$H_{\mathcal{B}}$	subgraph of G induced by actions $\mathcal{B}$	12, 23
$G W$	subgraph of G induced by actions with sources in W	41
$\det(\cdot)$	replacement of stochastic actions by sets of deterministic actions	64
$\mathbf{R}^n$	n -dimensional Euclidean space	18
$S^1; S^{n-2}$	circle; sphere of dimension $n-2$	15; 16
Σ, Γ	simplicial complex	12, 17
σ, τ, γ	simplex in a simplicial complex	12, 46
$\mathcal{N}(\cdot)$	nerve of a collection of sets	19
$\mathcal{F}(\cdot)$	face poset of a simplicial complex	19
$\text{sd}(\cdot)$	barycentric subdivision of a simplicial complex	19
Δ_G	strategy complex of a graph	12, 24
$\overline{\Delta}_G$	source complex of a graph	42
$\overline{\Delta}_G^*$	dual complex of a graph	44
$t_{\max}(\cdot)$	maximal worst-case expected convergence time of a strategy	24
Δ_G^T	subcomplex of Δ_G containing time-bounded strategies	24
$V_A, V_B, V_\sigma, V_\tau$	start region of a collection of actions	22
U_A, U_B, U_A^+, U_B^+	covering sets: open subsets of $\mathbf{R}^n$ associated with actions A and B	25
$[n]$	shorthand for $\{1, \dots, n\}$	16
s	desired stop state in a graph (e.g., a task goal)	16
S	set of desired stop states in a graph (e.g., alternative task goals)	30
$G_{\leftarrow s}$	loopback graph	16
$\Delta_{G_{\leftarrow s}}$	loopback complex	16
(M, W)	Markov chain M with support W	22
$*$	join operator (for spaces and simplicial complexes)	46
$\cong$	isomorphic (equivalence relation on simplicial complexes)	17
$\approx$	homeomorphic (equivalence relation on topological spaces)	17
$\simeq$	homotopic (equivalence relation on functions and spaces)	17
$\leftrightarrow$	strongly controllable (equivalence relation on states)	49
$G/\leftrightarrow$	quotient graph under strong controllability	49, 54
G/W	quotient graph formed by collapsing W to a point	49
$C_r, \partial C_r, \partial C_{T+\epsilon}$	polyhedral cylinders measuring convergence times	27

List of Lemmas and Theorems

#	Description	Page
*1	characterization of task solvability by spheres	16
2	Collapsibility Tool (from the literature)	19
3	Nerve Lemma (from the literature)	19
4	Quillen Fiber Lemma (from the literature)	20
5	equivalence of convergence time and Markov chain perspectives	23
6	characterization of stochastically acyclic actions	26
7	characterization of time-bounded stochastically acyclic actions	27
8	the homotopy types of strategy complexes are given by covering sets	28
9	stochastically certain connectivity implies a homogeneous halfspace	31
10	time-bounded stochastic connectivity implies an affine halfspace	31
*11	characterization of graph controllability by spheres	34
12	characterization of time-bounded graph controllability by spheres	35
13	contractibility of start regions	41
*14	strategy and source complexes are homotopic	42
15	characterization of source complexes	42
*16	characterization of task solvability by contractibility	43
17	some conditions under which graph union implies strategy join	47
18	simplices combine if they differ by actions with the same source	47
19	controllable subspaces factor strategy complexes	50
*20	structure theorem for strategy complexes	51
21	realizability of simplicial complexes by strategy complexes	59
22	realizability of simplicial complexes as source complexes	59
23	determining the precision of a system is <i>NP</i> -complete	60
24	small realizability is uncomputable	62
25	detecting changes in homotopy type is uncomputable	62
26	making stochastic actions deterministic preserves homotopy type	65

* — indicates a key theorem of the paper

List of Algorithms

#	Call Format	Description	Page
1	BACKCHAIN(G, S)	Finds actions to move off the complement of S , and repeats.	56
2	STRATEGY(G, S)	Finds a subgraph of G with maximal state space within which S is stochastically attainable.	57
3	REACHABLE(G)	Constructs the directed graph representing stochastically certain strategies for moving between states of G .	57
4	SUBSPACES(G)	Constructs the strongly controllable subspaces of G .	58
–	–	Description of an algorithm for computing quotient graphs.	58
–	–	Description of an algorithm for computing source complexes.	58
–	–	An approach for computing strategy complexes.	47

List of Key Definitions

Term	Page Reference
contractible	18
collapsible	18
cone	18
acyclic	12
stochastically acyclic	23
(nondeterministic) strategy	12
(stochastic) strategy	24
complete guaranteed strategy	16
complete stochastic strategy, stochastically attainable goal, etc.	30
System (1)	22
worst-case expected convergence times	24
an action moves off a set	41
fully controllable	35
strongly controllable subspace	49
pure stochastic graph	65
knowledge space	67

List of Figures

Figure	Description	Page
1	nondeterministic motion	9
2	nondeterministic graph; its strategy complex	9
3	cyclic motions	10
4	graph with a cycle; its strategy complex	10
5	directed graph with three states; its strategy complex	13
6	graph with two cycle-inducing nondeterministic actions; its complex	13
7	a strongly connected graph with three states; its complex	14
8	another strongly connected graph with three states; its complex	14
9	loopback graph and complex for the graph of Fig. 5	15
10	loopback graph and complex for the graph of Fig. 6	15
11	homotopy equivalence favors more precise actions	20
12	stochastic graph with two eventually-convergent actions; its complex	21
13	stochastic graph with three uncertain actions; its complex	28
14	covering sets for the graph of Fig. 13	29
15	action tuning parameters	36
16	snapshot of designing a system by considering covering sets	37
17	design space from graph perspective	38
18	design space from covering set perspective	39
19	design space from strategy complex perspective	39
20	the design space summarized by four simplicial complexes	40
21	source and dual complexes describing the design space	45
22	attainability versus strong controllability	49
23	a sample stochastic graph, perhaps modeling air travel	52
24	nontrivial strongly controllable subspaces of the previous graph	53
25	quotient graph and a complete strategy for attaining one state	54
26	source complex of the quotient graph	54
27	interpretation of a minimal nonface	55
28	deterministic instantiation of stochastic actions preserves homotopy type	65
29	imperfect sensing causes cycling	68
30	effect of imperfect sensing on strategy complex	68
31	an imperfect sensor stratifies a graph perfectly	70

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. Akella and M. T. Mason. Posing polygonal objects in the plane by pushing. *Proc. IEEE Intl. Conference on Robotics and Automation*, pages 2255–2262, 1992.
- [3] A. Barr, P. Cohen, and E. Feigenbaum, editors. *The Handbook of Artificial Intelligence*. William Kaufmann, Inc., Los Altos, California, 1981–1989.
- [4] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, N.J., 1957.
- [5] D. P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, Englewood Cliffs, N.J., 1987.
- [6] A. Björner. Topological methods. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume II, pages 1819–1872. North-Holland, Amsterdam, 1995.
- [7] A. Björner and M. Tancer. Note: Combinatorial Alexander duality — A short and elementary proof. *Discrete and Computational Geometry*, 42(4):586–593, 2009.
- [8] A. Björner and V. Welker. Complexes of directed graphs. *SIAM J. Discrete Math*, 12(4):413–424, 1999.
- [9] M. Brady, J. M. Hollerbach, T. Johnson, T. Lozano-Pérez, and M. T. Mason. *Robot Motion: Planning and Control*. MIT Press, Cambridge, MA, 1982.
- [10] R. C. Brost. Automatic grasp planning in the presence of uncertainty. *Intl. J. Robotics Research*, 7(1):3–17, 1988.
- [11] R. C. Brost and K. Y. Goldberg. A complete algorithm for designing planar fixtures using modular components. *IEEE Trans. on Robotics and Automation*, 12(1):31–46, 1996.
- [12] J. F. Canny. *The Complexity of Robot Motion Planning*. MIT Press, Cambridge, MA, 1988.
- [13] V. de Silva and R. Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *Intl. J. Robotics Research*, 25(12):1205–1222, 2006.
- [14] B. R. Donald. Robot motion planning with uncertainty. *Artificial Intelligence*, 37(1–3):223–271, 1988.
- [15] B. R. Donald. *Error Detection and Recovery in Robotics*. Lecture Notes in Computer Science, No. 336. Springer-Verlag, Berlin, 1989.
- [16] B. R. Donald. Planning multi-step error detection and recovery strategies. *Intl. J. Robotics Research*, 9(1):3–60, 1990.
- [17] B. R. Donald. On information invariants in robotics. *Artificial Intelligence*, 72:217–304, 1995.

- [18] B. R. Donald and J. Jennings. Constructive recognizability for task-directed robot programming. *Robotics and Autonomous Systems*, 9:41–74, 1992.
- [19] B. R. Donald, J. Jennings, and D. Rus. Towards a theory of information invariants for cooperating autonomous mobile robots. *Sixth International Symposium on Robotics Research*, pages 29–48, 1993.
- [20] J. Dugundji. *Topology*. Allyn and Bacon, Boston, 1966.
- [21] H. F. Durrant-Whyte. Sensor models and multisensor integration. *Intl. J. Robotics Research*, 7(6):97–113, 1988.
- [22] M. A. Erdmann. Using backprojections for fine motion planning with uncertainty. *Intl. J. Robotics Research*, 5(1):19–45, 1986.
- [23] M. A. Erdmann. *On Probabilistic Strategies for Robot Tasks*. PhD thesis, EECS, MIT, Cambridge, MA, 1989. Also available as Technical Report AI-TR-1155.
- [24] M. A. Erdmann. Randomization in robot tasks. *Intl. J. Robotics Research*, 11(5):399–436, 1992.
- [25] M. A. Erdmann. Randomization for robot tasks: Using dynamic programming in the space of knowledge states. *Algorithmica*, 10(2–4):248–291, 1993.
- [26] M. A. Erdmann. Understanding action and sensing by designing action-based sensors. *Intl. J. Robotics Research*, 14(5):483–509, 1995.
- [27] M. A. Erdmann. An exploration of nonprehensile two-palm manipulation. *Intl. J. Robotics Research*, 17(5):485–503, 1998.
- [28] M. A. Erdmann. On the topology of plans. *Eighth Intl. Workshop on the Algorithmic Foundations of Robotics*, Guanajuato, Mexico, December 2008.
- [29] M. A. Erdmann and M. T. Mason. An exploration of sensorless manipulation. *IEEE J. Robotics and Automation*, 4(4):369–379, 1988.
- [30] M. A. Erdmann, M. T. Mason, and G. Vaněček, Jr. Mechanical parts orienting: The case of a polyhedron on a table. *Algorithmica*, 10(2–4):226–247, 1993.
- [31] M. Farber. Topological complexity of motion planning. *Discrete and Computational Geometry*, 29(2):211–221, 2003.
- [32] W. Feller. *An Introduction to Probability Theory and Its Applications. Volume I*. John Wiley and Sons, New York, third edition, 1968.
- [33] A. Fox and S. Hutchinson. Exploiting visual constraints in the synthesis of uncertainty-tolerant motion plans. *IEEE Trans. on Robotics and Automation*, 11(1):56–71, 1995.
- [34] R. Ghrist and S. LaValle. Nonpositive curvature and Pareto-optimal coordination of robots. *SIAM J. Control & Opt*, 45(5):1697–1713, 2006.

- [35] R. Ghrist, J. O’Kane, and S. LaValle. Computing Pareto optimal coordinations on roadmaps. *Intl. J. Robotics Research*, 24(11):997–1010, 2005.
- [36] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Adv. Appl. Math*, 38:302–323, 2007.
- [37] K. Y. Goldberg. Orienting polygonal parts without sensors. *Algorithmica*, 10(2–4):201–225, 1993.
- [38] L. Guilamo, B. Tovar, and S. LaValle. Pursuit-evasion in an unknown environment using gap navigation trees. *Proc. IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, pages 3456–3462, 2004.
- [39] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
- [40] M. Herlihy and N. Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, 1999.
- [41] J. Hopcroft and G. Wilfong. Motion of objects in contact. *Intl. J. Robotics Research*, 4(4):32–46, 1986.
- [42] K. Hsiao, L. Kaelbling, and T. Lozano-Pérez. Grasping POMDPs. *Proc. IEEE Intl. Conference on Robotics and Automation*, pages 4685–4692, 2007.
- [43] K. Hsiao, L. Kaelbling, and T. Lozano-Pérez. Robust belief-based execution of manipulation programs. *Eighth Intl. Workshop on the Algorithmic Foundations of Robotics*, Guanajuato, Mexico, December 2008.
- [44] A. Hultman. Directed subgraph complexes. *Elec. J. Combinatorics*, 11(1):R75, 2004.
- [45] Y.-B. Jia and M. A. Erdmann. Geometric sensing of known planar shapes. *Intl. J. Robotics Research*, 15(4):365–392, 1996.
- [46] Y.-B. Jia and M. A. Erdmann. Pose and motion from contact. *Intl. J. Robotics Research*, 18(5):466–490, 1999.
- [47] J. Jonsson. *Simplicial Complexes of Graphs*. PhD thesis, Department of Mathematics, KTH, Stockholm, Sweden, 2005.
- [48] S. Karlin and H. M. Taylor. *A Second Course in Stochastic Processes*. Academic Press, New York, 1981.
- [49] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972.
- [50] J. L. Kelley. *General Topology*. Springer Verlag, New York, 1955.
- [51] J. Kerr and B. Roth. Analysis of multifingered hands. *Intl. J. Robotics Research*, 4(4):3–17, 1986.

- [52] D. E. Koditschek. An approach to autonomous robot assembly. *Robotica*, 12(2):137–155, 1994.
- [53] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.
- [54] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, New York, 2006.
- [55] A. Lazanas and J.-C. Latombe. Motion planning with uncertainty: A landmark approach. *Artificial Intelligence*, 76(1–2):285–317, 1995.
- [56] T. Lozano-Pérez, M. Mason, and R. Taylor. Automatic synthesis of fine-motion strategies for robots. *Intl. J. Robotics Research*, 3(1):3–24, 1984.
- [57] D. Mackenzie. ROBOTICS: Topologists and Roboticists Explore an ‘Inchoate World’. *Science*, 301(5634):756, 8 August 2003.
- [58] O. Madani, S. Hanks, and A. Condon. On the undecidability of probabilistic planning and related stochastic optimization problems. *Artificial Intelligence*, 147(1–2):5–34, 2003.
- [59] M. T. Mason. Automatic planning of fine motions: Correctness and completeness. *Proc. IEEE Intl. Conference on Robotics and Automation*, pages 492–503, 1984.
- [60] M. T. Mason. Mechanics and planning of manipulator pushing operations. *Intl. J. Robotics Research*, 5(3):53–71, 1986.
- [61] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Menlo Park, 1984.
- [62] J. Nevins, D. Whitney, S. Drake, D. Killoran, M. Lynch, D. Seltzer, S. Simunovic, R. Spencer, P. Watson, and A. Woodin. Exploratory research in industrial modular assembly. Technical Report R-921, The Charles Stark Draper Laboratory, 1975.
- [63] P. S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *Trudy Mat. Inst. Steklov*, 44:1–143, 1955.
- [64] S. P. Novikov. Appendix to: The problem of the algorithmic discrimination of the standard three-dimensional sphere, by I. A. Volodin, V. E. Kuznetsov, and A. T. Fomenko. *Uspekhi Mat. Nauk*, 29(5):71–168, 1974.
- [65] J. M. O’Kane and S. M. LaValle. Comparing the power of robots. *Intl. J. Robotics Research*, 27(1):5–23, 2008.
- [66] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Englewood Cliffs, N.J., 1982.
- [67] C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987.
- [68] D. Quillen. Homotopy properties of the poset of non-trivial p -subgroups of a group. *Advances in Math.*, 28:101–128, 1978.
- [69] J. J. Rotman. *An Introduction to Algebraic Topology*. Springer Verlag, New York, 1988.

- [70] R. C. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *Intl. J. Robotics Research*, 5(4):56–68, 1986.
- [71] E. H. Spanier. *Algebraic Topology*. McGraw-Hill, San Francisco, 1966.
- [72] R. H. Taylor, M. T. Mason, and K. Y. Goldberg. Sensor-based manipulation planning as a game with nature. In R. Bolles and B. Roth, editors, *Robotics Research*, volume IV, pages 421–429. MIT Press, Cambridge, MA, 1988.
- [73] B. Tovar, R. Murrieta, and S. LaValle. Distance-optimal navigation in an unknown environment without sensing distances. *IEEE Transactions on Robotics*, 23(3):506–518, 2007.
- [74] B. Tovar, A. Yershova, J. O’Kane, and S. LaValle. Information spaces for mobile robots. *Proc. Intl. Workshop on Robot Motion and Control*, 2005.
- [75] M. L. Wachs. Poset topology: Tools and applications. Technical report, IAS/Park City Mathematics Institute, Summer 2004.
- [76] S. Weinberger. *Computers, Rigidity, and Moduli: The Large-Scale Fractal Geometry of Riemannian Moduli Space*. M. B. Porter Lectures. Princeton University Press, Princeton, 2004.
- [77] D. E. Whitney. Quasi-static assembly of compliantly supported rigid parts. *Journal of Dynamic Systems, Measurement, and Control*, 104(1):65–77, 1982.
- [78] Wikipedia, The Free Encyclopedia. *Markov Decision Process*.
http://en.wikipedia.org/wiki/Markov_decision_process.