Deception, Delay, and Detection of Strategies

Michael Erdmann*
Carnegie Mellon University
June 27, 2019

Abstract

Homology generators in a relation offer individuals the ability to delay identification, by guiding the order via which the individuals reveal their attributes [6]. This perspective applies as well to the identification of goal-attaining strategies in systems with errorful control, since the strategy complex of a fully controllable nondeterministic or stochastic graph is homotopic to a sphere. Specifically, such a graph contains for each state v a maximal strategy σ_v that converges to state v from all other states in the graph and whose identity may be shrouded in the following sense: One may reveal certain actions of σ_v in a particular order so that the full strategy becomes known only after at least $n-1$ of these actions have been revealed, with none of the actions revealed definitively inferable from those previously revealed. Here n is the number of states in the graph. Moreover, the strategy contains at least $(n-1)!$ such *informative action release sequences*, each of length at least $n-1$.

The earlier work described above sketched a proof that *every* maximal strategy in a *pure nondeterministic* or *pure stochastic* graph contains *at least one* informative action release sequence of length at least $n-1$. The primary purpose of the current report is to fill in the details of that sketch. To build intuition, the report first discusses several simpler examples. These examples suggest an underlying structure for hiding capabilities or bluffing capabilities, as well as for detecting such deceit.

Contents

*This report is based upon work supported in part by the National Science Foundation under award number IIS-1409003. Any opinions, findings and conclusions or recommendations expressed in this report are those of the author and do not necessarily reflect the views of the Government or the National Science Foundation.

1 Introductory Examples

1.1 Paths and Constituent Transitions

Figure 1 shows four islands connected by bridges, as might be found in one of the great oceanic cities of the world. One of the bridges allows traffic in two directions, the others are one-way bridges. Of interest are the possible paths a bus of tourists or the motorcade of a prominent dignitary might take from the HOTEL ISLAND to the PALACE ISLAND, via one or two intermediary islands (the LEFT ISLAND and/or the RIGHT ISLAND).

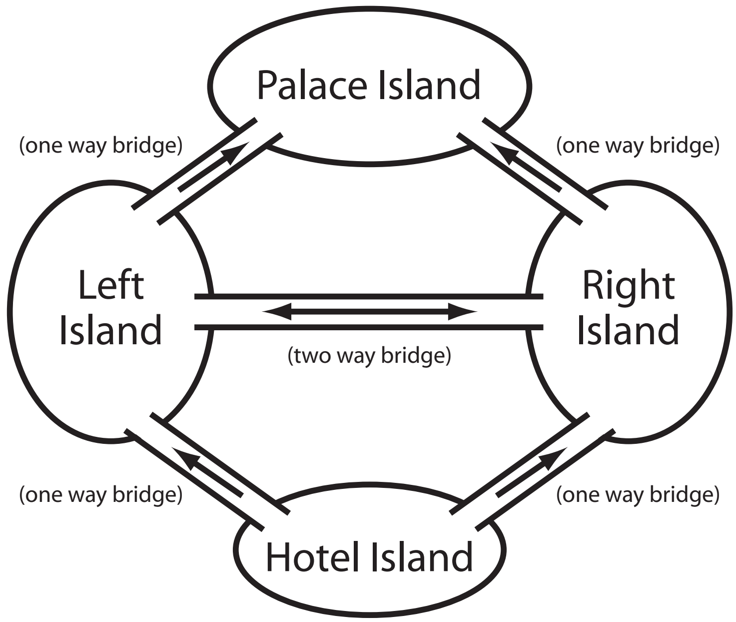

Figure 1: Four islands along with some directional bridges connecting the islands.

Since the bridge between LEFT ISLAND and RIGHT ISLAND is bidirectional, there are infinitely many such paths, parameterized by the number of times the bus or motorcade cycles over the two-way bridge. For the purposes of this report, we will disallow such infinite cycling. One can imagine different restrictions. In this first example, we impose the restriction that a vehicle may traverse the two-way bridge at most once in each of its possible directions (perhaps for legal or monetary reasons). In a more general setting, we would disallow traversing any bridge in the same direction more than once. Later, in Section 1.2, we will discuss a different example with a different restriction that prevents infinite cycling.

We define a *permissible path* to be any path that a vehicle might take from HOTEL ISLAND to PALACE ISLAND, subject to the "no directional transition twice" restriction. The next page enumerates all permissible paths; there are six. For clarity, we abbreviate each island name to its first letter and give paths the names $\pi_1, \ldots, \pi_6$. Throughout this subsection, we consider only these six paths, each of which starts at HOTEL ISLAND and ends at PALACE ISLAND.

π_1: H → L → P
π_2: H → R → P
π_3: H → L → R → P
π_4: H → R → L → P
π_5: H → L → R → L → P
π_6: H → R → L → R → P

Figure 2 describes the islands and bridges of Figure 1 as a directed graph, and the six permissible paths as a relation. The relation has a row for each permissible path and a column for each directed edge in the graph, that is, for each directional transition across a bridge. Since a permissible path may traverse any bridge direction at most once, each permissible path defines a *set* of directed edges, modeling all directional bridge transitions in the path. The relation therefore contains a nonblank entry • for a given path π_i and a given directed edge A → B if and only if path π_i includes transition A → B.

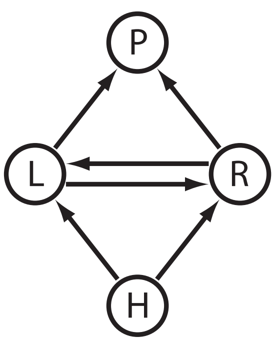

	H → L	H → R	L → P	R → P	L → R	R → L
π_1	•		•			
π_2		•		•		
π_3	•			•	•	
π_4		•	•			•
π_5	•		•		•	•
π_6		•		•	•	•

Figure 2: **Left Panel:** Directed graph representing the sketch of Figure 1. (Island names appear as first letter abbreviations.) **Right Panel:** A relation describing all paths leading from the Hotel Island to the Palace Island, while traversing any bridge direction at most once.

Identifying Paths from Transitions at Execution Time

Suppose an observer is watching a bus drive from the Hotel Island to the Palace Island. At what point during the trip can the observer identify uniquely the specific path followed by the bus, assuming the bus is traversing one of the permissible paths $\pi_1, \pi_2, \pi_3, \pi_4, \pi_5$, or π_6?

- Certainly, once the bus arrives at its destination, Palace Island, the observer can identify the path uniquely, since at that point the observer knows that he/she has seen the entire path.

- The observer cannot identify any path uniquely after observing only the first bridge transition. For instance, after observing transition H → L, the possible paths consistent with this observation are π_1, π_3, and π_5. Similarly, after observing transition H → R, the possible paths consistent with the observation are π_2, π_4, and π_6.

- Consider paths π_1 and path π_2, each of which consists of two transitions. By the previous point, the observer must see the entire path in order to identify either of these paths uniquely.

- The first two transitions of paths π_3 and π_5 are the same, namely $\mathtt{H} \rightarrow \mathtt{L}$ and $\mathtt{L} \rightarrow \mathtt{R}$. Consequently, upon observing these transitions, the observer cannot identify a path uniquely; the path could be either π_3 or π_5. Path π_3 consists of three transitions. Thus, if the actual path is π_3, the observer must see the entire path before identifying the path as π_3. A similar argument holds for path π_4.

- Paths π_5 and π_6 contain four transitions. For each of these two paths, the observer only needs to see the first three transitions in order to identify the path; no other permissible path shares those same three transitions with the path being observed.

In summary: Paths π_1, π_2, π_3, and π_4 can only be identified uniquely after seeing all their transitions, assuming one observes transitions in consecutive order. Paths π_5 and π_6 can be identified uniquely after seeing the first three of their four transitions, again assuming one observes transitions in consecutive order.

Identifying Paths from Transitions in Arbitrary Order

Previously we assumed that the observer was observing consecutive motions of a bus. Suppose now that the observer merely learns of particular transitions made by the bus, *without* any explicit ordering in time. For instance, perhaps the observer is listening to stories told by tourists on the bus after their trip, from which the observer attempts to reconstruct the path taken. Or perhaps the observations are coming from many trips taken over the course of several days by a bus following a particular fixed bus route each day. Or perhaps the observer overhears the bus driver commenting on particular bridges he will encounter on his next trip, from which the observer is trying to predict the path yet to be taken.

We now ask: What *set* of transitions allows an observer to identify a path uniquely?

- Recall that path π_1 consists of the set of transitions $\{\mathtt{H} \rightarrow \mathtt{L},\ \mathtt{L} \rightarrow \mathtt{P}\}$. Previously, when observing transitions in consecutive order, seeing both these transitions identified path π_1 uniquely. That is no longer true when transitions may be observed nonconsecutively. The reason is that path π_5 contains these same transitions, plus others. In fact, it is *no longer possible* to identify path π_1 uniquely. Similarly, it is no longer possible to identify path π_2 uniquely.

- If the observer learns that a path contains the transitions $\mathtt{H} \rightarrow \mathtt{L}$ and $\mathtt{R} \rightarrow \mathtt{P}$, then the observer can infer that the path must also contain the transition $\mathtt{L} \rightarrow \mathtt{R}$ and must in fact be path π_3. Whereas previously an observer needed to see the entire path π_3 in order to identify it uniquely, now a pair of nonconsecutive transitions identifies the path. In effect, continuity of paths allows the observer to infer an unobserved transition. A similar argument holds for path π_4. Of course, if a story teller wishes to draw out identification of the path, he/she might simply talk about the sights seen during the bus ride in consecutive order, thus preventing such a leap of inference.

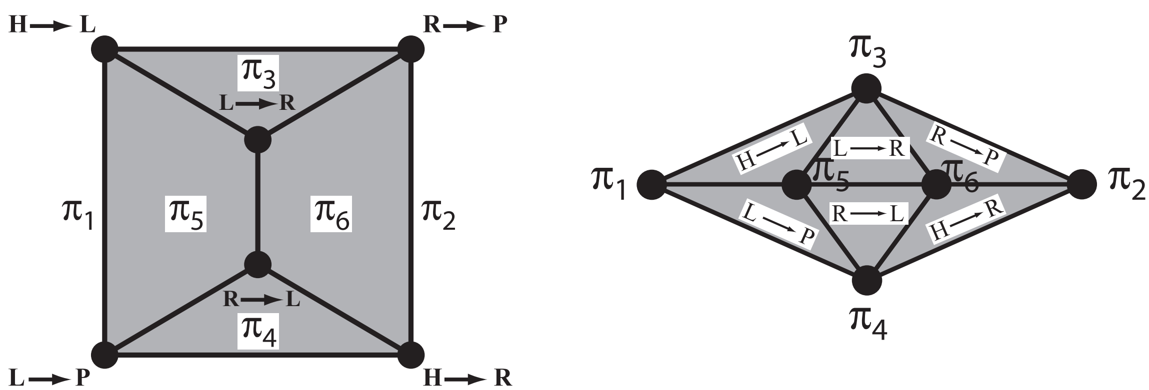

Figure 3: Two simplicial complexes derived from the relation of Figure 2. The two complexes are Dowker dual [6] to each other with respect to that relation. **Left Panel:** The underlying vertex set of this complex is the collection of directed edges in the graph of Figure 2. The paths $\pi_1, \pi_2, \pi_3, \pi_4, \pi_5, \pi_6$ generate simplices as indicated by the path labels. (The two quadrilaterals are actually solid tetrahedra, flattened for ease of viewing in the figure.) **Right Panel:** The underlying vertex set of this complex is the collection of permissible paths in the graph of Figure 2. Each maximal simplex in the complex is a triangle, reflecting the fact that each possible directed edge in the graph appears in three permissible paths. Each triangle is labeled with that directed edge.

- If the observer learns that a path contains the transitions $\mathtt{H} \to \mathtt{L}$ and $\mathtt{R} \to \mathtt{L}$, then the observer can actually infer two unobserved transitions, namely $\mathtt{L} \to \mathtt{R}$ and $\mathtt{L} \to \mathtt{P}$, thereby concluding that the path is π_5. A similar inference is possible for path π_6. The observer is in effect taking advantage both of path continuity and knowledge of the path's destination. Again, a story teller could draw out identification of path π_5 slightly by reporting transitions in consecutive order.

Figure 3 encodes these conclusions geometrically, using simplicial complexes [11, 13, 6]. As in the relation of Figure 2, we now view each path as a *set* of directed edges. These sets constitute the generating simplices of the left simplicial complex shown in Figure 3. The vertices in this complex are the directed edges of the graph of Figure 2. Path π_1 generates a one-dimensional simplex (edge) in the complex. This simplex is a subset of the three-dimensional simplex (tetrahedron) generated by path π_5, modeling the earlier conclusion that one cannot identify path π_1 uniquely when observing transitions in arbitrary order. Observe that the set $\{\mathtt{H} \to \mathtt{L}, \mathtt{R} \to \mathtt{P}\}$ is a free face* in the complex and is not itself a path. This geometry models the inference and identification of path π_3 discussed previously. Similarly, the set $\{\mathtt{H} \to \mathtt{L}, \mathtt{R} \to \mathtt{L}\}$ forms a free face in the complex, modeling the inferences and identification of path π_5 discussed above. (The set $\{\mathtt{H} \to \mathtt{L}, \mathtt{R} \to \mathtt{L}\}$ is an undrawn "diagonal" of the tetrahedron labeled π_5.)

The right simplicial complex of Figure 3 contains the same information as the left complex, but in dual form. The duality is with respect to the relation of Figure 2. (Details of such "Dowker duality" are discussed further in [6].)

*Simplicial complexes in this report are *abstract*, i.e., collections of sets and all their subsets. A simplex is a *free face* of an abstract simplicial complex if it is a proper subset of exactly one maximal simplex in the complex.

The vertices in the right complex are the permissible paths of the graph of Figure 2. The generating simplices of the complex are given by the columns of the relation of Figure 2. In other words, each generating simplex consists of all the paths that share a given directed edge. Thus the right complex tells us how to interpret observations of transitions as intersections of generating simplices. For instance, if we know that a path contains the transitions $\mathtt{H} \rightarrow \mathtt{L}$ and $\mathtt{L} \rightarrow \mathtt{P}$, then we can intersect the triangle labeled with $\mathtt{H} \rightarrow \mathtt{L}$ and the triangle labeled with $\mathtt{L} \rightarrow \mathtt{P}$ to see that the possible paths are π_1 and π_5. (This geometric intersection is exactly the intersection of the two columns indexed by $\mathtt{H} \rightarrow \mathtt{L}$ and $\mathtt{L} \rightarrow \mathtt{P}$ in the relation of Figure 2.)

Considering such intersections, our earlier observations plus some others are immediate:

- π_1 and π_2 are not uniquely identifiable.
- Observing $\mathtt{H} \rightarrow \mathtt{L}$ and $\mathtt{R} \rightarrow \mathtt{P}$ identifies path π_3.
- Observing $\mathtt{H} \rightarrow \mathtt{R}$ and $\mathtt{L} \rightarrow \mathtt{P}$ identifies path π_4.
- Observing $\mathtt{H} \rightarrow \mathtt{L}$ and $\mathtt{R} \rightarrow \mathtt{L}$ identifies path π_5 (as does observing $\mathtt{L} \rightarrow \mathtt{P}$ and $\mathtt{L} \rightarrow \mathtt{R}$).
- Observing $\mathtt{H} \rightarrow \mathtt{R}$ and $\mathtt{L} \rightarrow \mathtt{R}$ identifies path π_6 (as does observing $\mathtt{R} \rightarrow \mathtt{P}$ and $\mathtt{R} \rightarrow \mathtt{L}$).

(These are the smallest sets of identifying observations for each path; there exist larger sets of observations as well.)

1.2 Strategies and Underlying Capabilities

Figure 4 shows a river with two islands, a fishing area upstream of the islands, and a marina downstream from the islands. The two islands create three passages within the river that boats may traverse, either upstream or downstream, as they move between the fishing area and the marina. The three passages produce currents of different strengths. One of these currents is so strong that only boats with powerful motors are able to traverse the current going upstream. Fish in the fishing area like to gather near the start of that strong current. Consequently, boats with powerful motors have an advantage reaching nice fish over boats with weaker motors. On the other hand, revealing that one has a powerful motor leads to envy and other competitions. As a result, skippers tend to underplay the power of their motors.

We will examine the possible strategies for reaching the fishing area (along with strategies for reaching the marina). We will further examine the extent to which someone can reveal portions of a strategy without revealing the entire strategy. Conversely, we will examine the extent to which an observer can infer that a boat has a powerful motor even when the observer never sees the boat traversing upstream over the strong current.

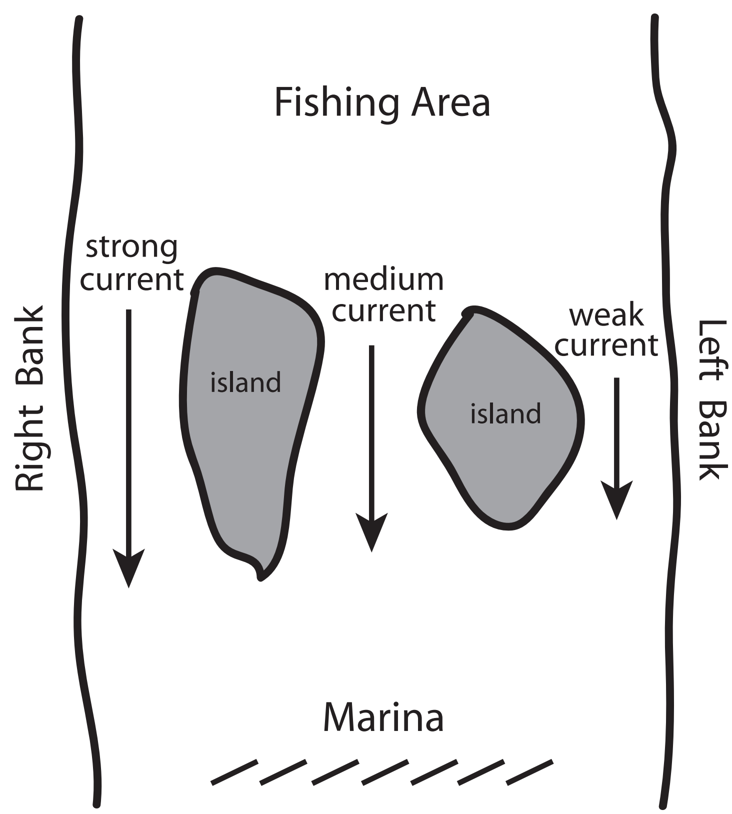

Figure 4: A river along with islands that create three passages and consequent currents of different strengths, with a fishing area upstream and a marina downstream.

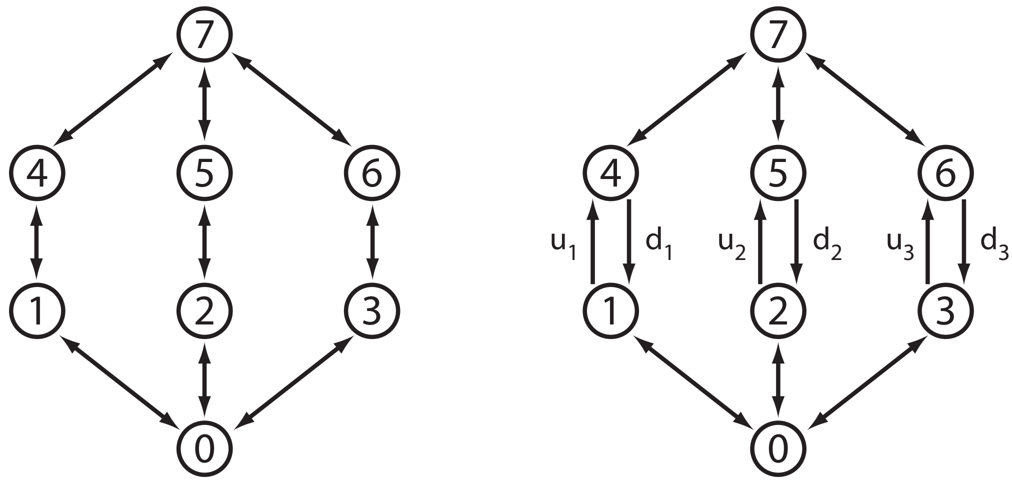

Figure 5: **Left Panel:** A directed graph that describes the possible transitions a boat might make while traversing the river of Figure 4 between the marina and the fishing area. (State #0 represents the marina, while state #7 represents the fishing area.)
Right Panel: The same graph, now with the transitions that move upstream or downstream through the passages given explicit names.

Figure 5 describes the river setting of Figure 4 using a directed graph, much as we did in the earlier example of Section 1.1. We will be interested primarily in the transitions upstream and downstream through the passages beside the islands, so we give those directed edges explicit names: $\mathtt{u}_1$, $\mathtt{d}_1$, $\mathtt{u}_2$, $\mathtt{d}_2$, $\mathtt{u}_3$, $\mathtt{d}_3$, as shown in the right panel of the figure.

In the example of Section 1.1, we focused on paths. In the current example, we adopt a slightly different perspective. We are interested in a generalization of what is frequently called a *control law*, namely a mapping from states to commanded motions. The generalization is that of a *strategy*, to be reviewed in Section 2.1. A strategy is a mapping from states to *sets* of possible motions, in this case sets of directed edges. The semantics are as follows: When a boat is at a particular location, a strategy specifies a set of directed edges leading from that location to some neighboring locations. The boat must move along some one of those directions, with the particular direction determined possibly by circumstance rather than chosen by the skipper. (The strategy specifies a set since sometimes the precise direction is not so important as is a general direction. For instance, a boat with a powerful motor that is currently at the marina might be instructed to move toward any of the three passages in the river. A corresponding strategy would therefore include the set of transitions $\{0 \to 1,\, 0 \to 2,\, 0 \to 3\}$.) If the set specified for a particular location is empty, then the boat must stop if it is at that location.

In the example of Section 1.1, we prevented infinite cycling by disallowing any path that traversed any directed edge more than once. With strategies, it is more natural to disallow any strategy whose motion sets might cause the system to revisit a state.

There is a well-developed theory for strategies in graphs with directed edges [1, 8, 9] as well as in graphs with nondeterministic and/or stochastic transitions [4, 5]. One can model the collection of all strategies as a simplicial complex, similar to the constructions of Section 1.1. In a directed graph, a strategy is a set of directed edges that produces no cycle(s) in the graph. A graph's strategies constitute the simplices of a simplicial complex whose underlying vertex set consists of the graph's directed edges. For a strongly connected directed graph, this simplicial complex has the homotopy type of a sphere, namely $\mathbb{S}^{n-2}$, with n the number of states in the graph [8]. For the graph of Figure 5, the simplicial complex is therefore homotopic to $\mathbb{S}^6$.

Since a sphere has homology, our prior work on privacy [6] offers some lower bounds on how long a skipper may delay identification of a strategy or a boat's final destination, relative to all possible strategies in the graph of Figure 5. However, rather than explore the entire space of strategies, we will focus in this example on some simpler scenarios.

Strategies for Attaining the Fishing Area

For the moment, let us consider only all maximal strategies that ultimately attain the fishing area from anywhere in the graph. (By a *maximal strategy* we mean here a cycle-free set of directed edges in the graph of Figure 5 that is maximal among all such sets.)

- Here is one such strategy, consisting of all possible upstream motions in the graph of Figure 5:

  $$\sigma_{123} \;=\; \{0 \to 1,\; 0 \to 2,\; 0 \to 3,\; 1 \to 4,\; 2 \to 5,\; 3 \to 6,\; 4 \to 7,\; 5 \to 7,\; 6 \to 7\}.$$

  (This strategy contains the three upstream transitions $\mathtt{u}_1$, $\mathtt{u}_2$, and $\mathtt{u}_3$, with $\mathtt{u}_i = i \to i{+}3$.)

- The strategy σ_{123} only makes sense for a boat with a powerful enough motor to traverse the strong current of Figure 4. A boat without such a powerful motor might instead use the following strategy:

  $$\sigma_{23} \;=\; \{1 \to 0,\; 0 \to 2,\; 0 \to 3,\; 4 \to 1,\; 2 \to 5,\; 3 \to 6,\; 4 \to 7,\; 5 \to 7,\; 6 \to 7\}.$$

  (This strategy contains downstream transition $\mathtt{d}_1$ and upstream transitions $\mathtt{u}_2$ and $\mathtt{u}_3$.)

  Strategy σ_{23} is very similar to strategy σ_{123}, but in place of the upstream transitions $0 \to 1$ and $1 \to 4$, the strategy contains the downstream transitions $1 \to 0$ and $4 \to 1$. As a result, if necessary, the boat will first return to the marina via the leftmost passage of Figure 4, then move up to the fishing area via either of the other two passages.

Permissible Strategies: Strategy σ_{23} specifies two transitions at state #4, namely $4 \to 7$ and $4 \to 1$. A boat moving under strategy σ_{23} may therefore reach the fishing area from state #4 either by moving directly to the fishing area or by moving first downstream to the marina then upstream via one of the other passages. Intuitively, this bifurcation arises because there are two arcs between any two points on a circle. Some maximal strategy must contain both.

While generally useful, motion multiplicity may merely add bookkeeping clutter, so we restrict it: To start, we define a *permissible strategy* to be a maximal strategy that (i) attains the fishing area from anywhere in the graph and (ii) specifies a unique motion at each state in the set $\{1, 2, 3\}$. We also stipulate that whenever a strategy specifies a passage transition and some other motion at a boat's current location, then the boat will move through the passage.

We may now model permissible strategies via the relation of Figure 6. The relation contains a row for each permissible strategy and a column for each passage transition. An entry in the

	$\mathtt{u}_1$	$\mathtt{d}_1$	$\mathtt{u}_2$	$\mathtt{d}_2$	$\mathtt{u}_3$	$\mathtt{d}_3$
σ_{123}	•		•		•	
σ_{12}	•		•			•
σ_{23}		•	•		•	
σ_{13}	•			•	•	
σ_1	•			•		•
σ_2		•	•			•
σ_3		•		•	•	

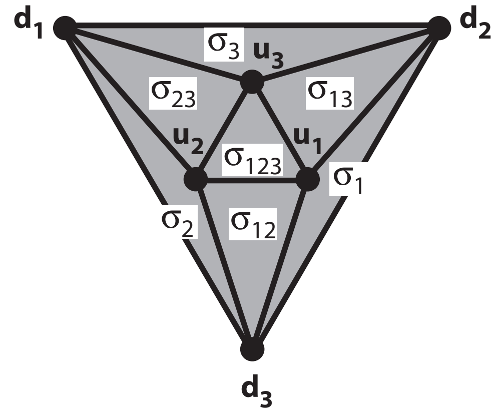

Figure 6: **Left Panel:** Relation describing permissible strategies for attaining state #7 in Figure 5. (State #7 is the fishing area in Figure 4. See page 8 for the meaning of *permissible.*) For each permissible strategy, the relation shows only the subset of upstream and downstream transitions $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{u}_3, \mathtt{d}_3\}$ contained in the strategy. These transitions describe motions through the three passages of the river. They fully determine the strategy, given that it is permissible. **Right Panel:** A simplicial complex derived from the relation, with underlying vertex set being the upstream and downstream transitions $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{u}_3, \mathtt{d}_3\}$. Each maximal simplex is labeled with its strategy name, as specified by its row in the relation.

relation is nonblank if and only if the given strategy contains the given transition. Every maximal strategy in the graph of Figure 5 must contain exactly one transition from each of the three sets $\{\mathtt{u}_i, \mathtt{d}_i\}$, $i = 1, 2, 3$. Furthermore, among the *permissible* strategies, each strategy is uniquely characterized by the three passage transitions it contains. (Of course, it is impossible for a permissible strategy to contain all three downstream transitions $\mathtt{d}_1$, $\mathtt{d}_2$, $\mathtt{d}_3$, since then the strategy would not be guaranteed to attain the fishing area from the marina.) Figure 6 further depicts a simplicial complex generated by the strategies of the relation. The underlying vertex set of this complex is $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{u}_3, \mathtt{d}_3\}$, comprising the six passage transitions in the river.

Each of the free faces in the complex of Figure 6 consists of a pair of downstream transitions, e.g., $\{\mathtt{d}_1, \mathtt{d}_2\}$, suggesting inference of an upstream transition, e.g., $\mathtt{u}_3$. Indeed, if an observer learns that a permissible strategy specifies downstream motion through two passages, then the observer can infer that the strategy must specify an upstream motion through the remaining passage (since the fishing area is given as destination). Consequently, observing two downstream transitions in a permissible strategy identifies the strategy uniquely. There are no other free faces in the complex. Consequently, observing any other proper subset of a permissible strategy's passage transitions does not identify that strategy uniquely.

Comment: On a given fishing expedition, an observer may only see a boat move through a single passage, but over the course of several days the observer may see the boat take different routes. Or perhaps a crewmember speaks of the transitions specified by a strategy. Assuming the skipper's strategy is constant, the observer may be able to eventually infer the overall permissible strategy, much like an observer could infer a bus route in Section 1.1, after observing different bridge crossings on different days or by listening to tourist stories.

Inferring Motor Strength: How might an observer of a boat's transitions infer that the boat has a strong motor? Directly observing the upstream transition u_1 is one way, of course. Additionally, if the observer learns that a strategy specifies downstream transitions $\mathtt{d}_2$ and $\mathtt{d}_3$ *and* if the observer knows that these are part of a strategy to reach the fishing area, then the observer can infer that the strategy must contain the upstream transition $\mathtt{u}_1$, implying that the boat has a powerful motor. (We assume that each boat only follows strategies it can execute.)

Four of the permissible strategies in Figure 6, namely σ_{123}, σ_{12}, σ_{13}, and σ_1, presuppose a powerful motor. If a skipper is following one of the strategies σ_{123}, σ_{12}, or σ_{13}, then the skipper can carefully reveal up to two different passage transitions while still hiding the motor's power. In contrast, for strategy σ_1, the skipper can reveal at most one passage transition; revealing a second transition necessarily exposes or implies the motor's power.

Strategies for Attaining the Fishing Area and Strategies for Attaining the Marina

Let us augment our collection of *permissible strategies*, in order to model boat excursions that are *either* outbound to the fishing area *or* returning to the marina. We now permit any maximal strategy for attaining the fishing area that specifies a unique motion at each state in the set $\{1, 2, 3\}$, *plus* any maximal strategy for attaining the marina that specifies a unique motion at each state in the set $\{4, 5, 6\}$. (We retain the stipulation regarding passage transitions.)

Focusing on the subcollection of strategies for attaining the marina, we may again construct a simplicial complex whose underlying vertex set is $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{u}_3, \mathtt{d}_3\}$, much as in Figure 6, now with the upstream and downstream transitions interchanged. We thus have two simplicial complexes, one for the fishing-attaining strategies, the other for the marina-attaining strategies. We wish to combine these complexes. In order to not confuse simplices, we conify each complex with a vertex identifying the complex. We then glue the resulting two complexes together at common boundary locations. The final simplicial complex thus obtained is homotopic to $\mathbb{S}^2$.

In order to visualize this construction more easily, let us simplify the problem, by removing one of the islands, as in Figure 7. Now there are only two passages, one with a strong current requiring a powerful motor for the upstream direction, the other with a mild current, traversable by all boats in both directions. (The new graph's state and transition names are consistent with those of Figure 5. State #0 is the marina and state #7 is the fishing area.)

There are now three permissible strategies that attain the fishing area from anywhere in the graph. We name them σ_{12}, σ_1, and σ_2. Similarly, there are three permissible strategies that attain the marina from anywhere in the graph. We name them τ_{12}, τ_1, and τ_2. Figure 8 describes these six strategies via a relation. Strategy names index the rows of the relation. The upstream and downstream transitions, $\mathtt{u}_1$, $\mathtt{d}_1$, $\mathtt{u}_2$, and $\mathtt{d}_2$, plus two additional attributes, $\mathtt{f}$ and $\mathtt{m}$, index the columns of the relation. Previously, when we were considering only permissible strategies for attaining the fishing area, the component upstream and downstream transitions of that strategy fully determined the strategy (given that the strategy was permissible). Now, knowing a strategy's upstream and downstream transitions *may not* fully determine the strategy. However, each permissible strategy in the set $\{\sigma_{12}, \sigma_1, \sigma_2, \tau_{12}, \tau_1, \tau_2\}$ is fully determined by its upstream and downstream transitions *and* by its destination. The attributes $\mathtt{f}$ and $\mathtt{m}$ model this destination. Attribute $\mathtt{f}$ means a strategy's destination is the fishing area and attribute $\mathtt{m}$ means the strategy's destination is the marina. For each permissible strategy, the relation lists the strategy's upstream and downstream transitions along with its destination.

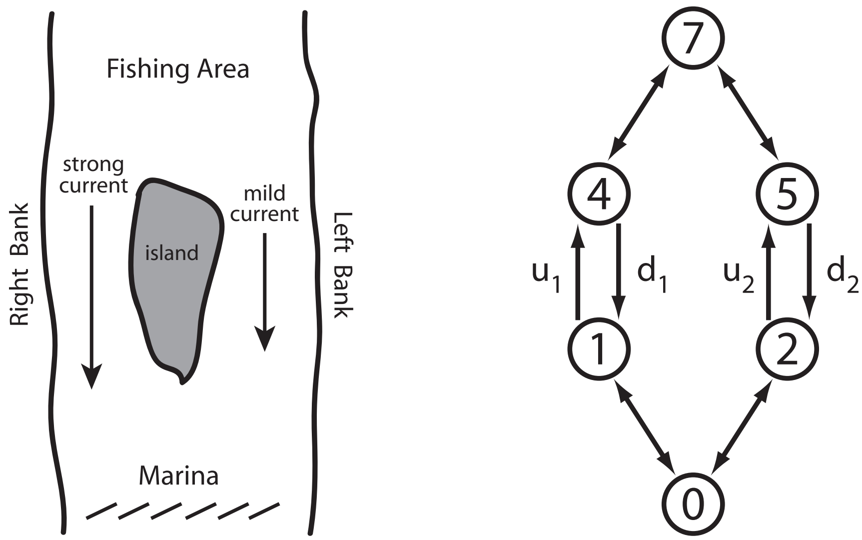

Figure 7: **Left Panel:** Simplified version of the river from Figure 4, in which there now are only two passages and consequent currents. **Right Panel:** Corresponding simplified graph, again highlighting the upstream and downstream transitions through the passages.

	$\mathtt{u}_1$	$\mathtt{d}_1$	$\mathtt{u}_2$	$\mathtt{d}_2$	f	m
σ_{12}	•		•		•	
σ_1	•			•	•	
σ_2		•	•		•	
τ_1		•	•			•
τ_2	•			•		•
τ_{12}		•		•		•

Figure 8: Relation describing permissible strategies σ_{12}, σ_1, σ_2 for attaining state #7 (fishing area) and permissible strategies τ_{12}, τ_1, τ_2 for attaining state #0 (marina) in Figure 7. The relation uses the upstream and downstream transitions for each strategy as attributes. Without knowing a strategy's destination, these transitions are not necessarily enough to uniquely determine the strategy. Consequently, the relation includes two additional attributes, to indicate the destination, as either the fishing area (attribute `f`) or the marina (attribute `m`).

Comment: One may readily observe a boat traversing a passage, but what does it mean to observe attribute `f` or `m`? One possibility is that a skipper announces a boat's destination. Another possibility is that the "observation" of `f` or `m` is actually an inference made from other observations. For instance, if an observer learns that the skipper has made preparations for both transitions $4 \to 7$ and $5 \to 7$, then the observer may conclude that the boat's destination is the fishing area. Or perhaps someone observes a boat departing the marina, for instance by making the transition $0 \to 2$. Then the observer knows that the boat's destination cannot be the marina, and so has "observed `f`" (assuming the strategy being observed is permissible).

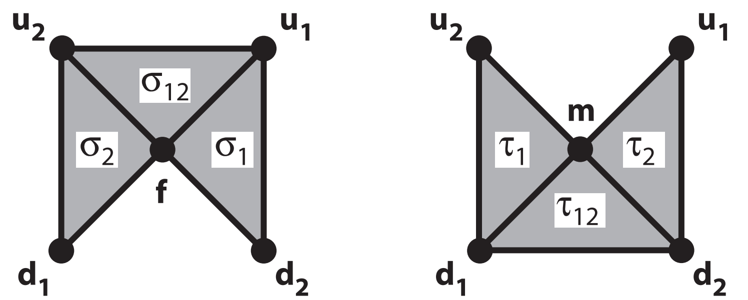

Figure 9: Two simplicial complexes derived from the relation of Figure 8, with underlying vertex set being the attributes $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{f}, \mathtt{m}\}$. The left complex shows the permissible strategies σ_{12}, σ_1, and σ_2 for attaining the fishing area. The right complex shows the permissible strategies τ_{12}, τ_1, and τ_2 for attaining the marina. See Figure 10 as well.

We will now construct a simplicial complex to represent the relation of Figure 8, much as we constructed the complex in Figure 6. Let us proceed in steps. First, we construct a simplicial complex representing the permissible strategies for attaining the fishing area and a separate simplicial complex representing the permissible strategies for attaining the marina. In both cases, we let the underlying vertex set be $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{f}, \mathtt{m}\}$. See Figure 9. Observe that attribute $\mathtt{f}$ is a cone apex for the complex generated by the fishing-attaining strategies, while attribute $\mathtt{m}$ is a cone apex for the complex generated by the marina-attaining strategies.[†] Finally, we glue these two simplicial complexes together along shared simplices, obtaining the complex of Figure 10. The homotopy type of this complex is $\mathbb{S}^1$, suggesting some ability to delay identification of strategies [6]. Observe in particular that every upstream or downstream transition appears in three permissible strategies.

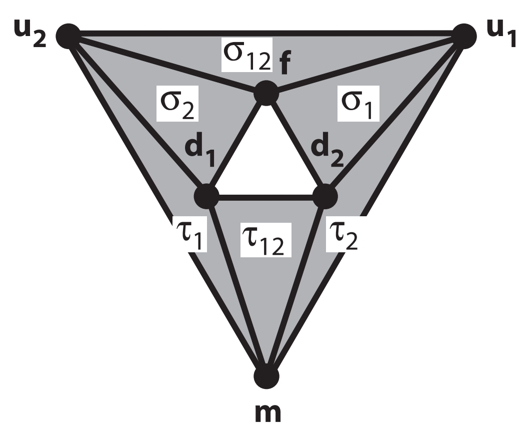

Figure 10: Simplicial complex derived from the relation of Figure 8, with underlying vertex set being the attributes $\{\mathtt{u}_1, \mathtt{d}_1, \mathtt{u}_2, \mathtt{d}_2, \mathtt{f}, \mathtt{m}\}$. (This complex is obtained from the two complexes shown in Figure 9 by gluing those complexes together along common vertices and edges.) Each maximal simplex is labeled with its strategy name, as specified in the relation.

[†]A *cone apex* for a finite simplicial complex is a vertex contained in every maximal simplex of the complex.

Considering the free and nonfree faces of the simplicial complex in Figure 10, or directly from the relation of Figure 8, we may make the following inferences:

- Observing upstream transitions $\mathtt{u}_1$ and $\mathtt{u}_2$ in a boat's strategy implies attribute $\mathtt{f}$. In other words, one may infer that the boat's destination is the fishing area and that the encompassing permissible strategy is σ_{12}. This inference reflects the physical reality that a boat cannot have the marina as destination if its strategy always entails moving upstream, via the two passages on both sides of the island.
- Similarly, observation of $\mathtt{d}_1$ and $\mathtt{d}_2$ implies $\mathtt{m}$ and identifies strategy τ_{12}.
- Observing an upstream transition in one passage and a downstream transition in the other passage (without knowing which occurred first) leaves the boat's destination ambiguous. This ambiguity reflects the physical reality that the boat could have rounded either the upstream end or the downstream end of the island.
- Knowing the boat's destination and observing a passage transition *away* from that destination implies the other passage transition and thus the overall permissible strategy. For instance, knowing that the boat is heading to the fishing area ($\mathtt{f}$) and observing the boat move downstream along the strong current ($\mathtt{d}_1$) means the boat's strategy must also specify a motion upstream over the mild current ($\mathtt{u}_2$). The geometry of the river forces this implication and thus identifies the strategy σ_2 (assuming strategies are permissible).
- Knowing the boat's destination and observing a passage transition *toward* that destination leaves open the directionality of the transition through the other passage. For instance, knowing that the boat is heading to the marina ($\mathtt{m}$) and observing the boat move downstream along the mild current ($\mathtt{d}_2$) does not nail down whether the strategy specifies an upstream or a downstream transition through the passage with the strong current. There remains an ambiguity as to whether the encompassing permissible strategy is τ_{12} or τ_2.

Deception and Detection

Skippers of boats with motors capable of moving upstream over the strong current tend to hide their strength for when it is really needed, such as a competition to snag nice fish at the head of the strong current. They may hide their strength either by never exercising it or by moving upstream over the strong current only under cover of fog or darkness.

Let us suppose that this deception is so pervasive that, for all intents and purposes, the upstream transition $\mathtt{u}_1$ is *never* observable. Two questions emerge:

1. How does the unobservability of $\mathtt{u}_1$ change the relation of Figure 8 and the complex of Figure 10?
2. How can one distinguish between a boat that is truly incapable of moving upstream over the strong current and a boat whose skipper is merely hiding that capability?

An initial answer to question #2 is that one cannot distinguish the two types of boats *if* the powerful boat *always* acts like the weaker boat. However, if the powerful boat does sometimes exercise its capabilities, then other observations *may* imply the boat's power.

	$\mathtt{d}_1$	$\mathtt{u}_2$	$\mathtt{d}_2$	$\mathtt{f}$	$\mathtt{m}$
σ_{12}		•		•	
σ_1			•	•	
σ_2	•	•		•	
τ_1	•	•			•
τ_2			•		•
τ_{12}	•		•		•

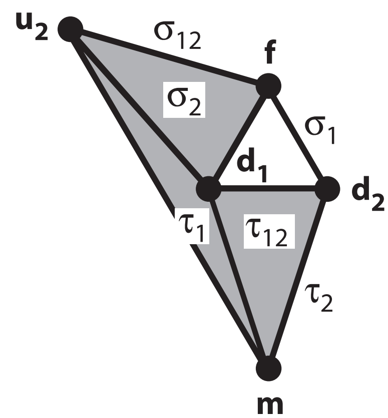

Figure 11: Here is a version of the relation from Figure 8 and the simplicial complex from Figure 10 in which the attribute $\mathtt{u}_1$ exists but is not observable. Each strategy in the relation generates a simplex in the complex, as indicated by the labels. Some strategies now appear as edges rather than triangles, due to the unobservable transition $\mathtt{u}_1$ in those strategies. Although this transition is unobservable, one can infer its existence indirectly if one observes attributes $\mathtt{d}_2$ and $\mathtt{f}$. This is because the edge $\{\mathtt{d}_2, \mathtt{f}\}$ is not a subsimplex of some larger strategy.

Figure 11 answers question #1. The relation one obtains when $\mathtt{u}_1$ is unobservable is the same as the original relation of Figure 8, except that the column indexed by the upstream transition $\mathtt{u}_1$ disappears. The resulting simplicial complex is now the deletion $\mathrm{dl}(\Gamma, \mathtt{u}_1)$, with Γ the complex of Figure 10. (Formally, $\mathrm{dl}(\Gamma, \mathtt{u}_1) = \{\gamma \in \Gamma \mid \mathtt{u}_1 \notin \gamma\}$. In other words, the resulting complex contains all simplices of the original complex except those that included $\mathtt{u}_1$.)

In the new complex, some strategies that appeared as triangles originally now appear as edges. These are the strategies σ_{12}, σ_1, and τ_2, namely all strategies that include the now unobservable upstream transition $\mathtt{u}_1$.

The observable portions of two of those strategies, namely σ_{12} and τ_2, are subsets of other strategies. For instance, the observable portion of σ_{12} is a subset of strategy σ_2. This means: If one observes the upstream transition $\mathtt{u}_2$ and if one knows that the boat's destination is the fishing area (attribute $\mathtt{f}$), then there remains an ambiguity regarding the strategy's specified transition over the strong current; it could be either upstream (as in σ_{12}) or downstream (as in σ_2). Consequently, if a skipper with a powerful motor always follows strategy σ_{12}, but traverses the strong current only during fog and the mild current during clear weather, then no one will know of the boat's power. (This assumes that the traversal times are unmeasured or constant, e.g., the skipper avoids traversing the mild current excessively quickly with the strong motor.)

In contrast, strategy σ_1 generates a maximal simplex even in the new complex and is thus uniquely identifiable from its observable attributes. This means: If one observes the downstream transition $\mathtt{d}_2$ but knows that the overall destination of the boat is the fishing area (attribute $\mathtt{f}$), then one can conclude that the strategy must be σ_1 and that the boat will traverse the strong current upstream. In other words, even though $\mathtt{u}_1$ is unobservable, one can infer its existence and conclude that the boat has a powerful motor.

	d_1	u_2	d_2	f	m
σ_2	•	•		•	
τ_1	•	•			•
τ_{12}	•		•		•

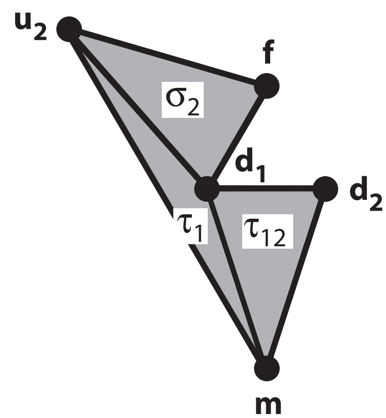

Figure 12: **Left Panel:** A relation describing the permissible strategies for a boat that is incapable of moving upstream over the strong current of Figure 7. This incapacity amounts to removing the directed edge u_1 from the graph of Figure 7 *as well as* removing all strategies from the relation of Figure 8 that contain transition u_1. **Right Panel:** The resulting simplicial complex, with maximal simplices labeled by strategy names.

It is instructive to construct the strategy relation and attendant simplicial complex for a boat that truly is incapable of traversing the strong current upstream. These appear in Figure 12. Not only does the transition u_1 now disappear, but so do all strategies that relied on that transition. (Again, we assume that the boat only follows strategies which it is capable of executing.)

The space of strategies is much smaller now. The resulting simplicial complex looks similar to that in which u_1 was merely unobservable, but there is a key difference: The simplex $\{f, d_2\}$ is gone. (In fact, all three strategies that once contained u_1 are now gone, but only one of those, namely σ_1, formed a maximal simplex previously in the complex of Figure 11.) The new complex tells us that it is *inconsistent* for a boat with a weak motor to announce its destination as the fishing area (attribute f) but to traverse the mild current downstream (attribute d_2).

Report [6] described how inconsistent observations sometimes suggest or identify unmodeled properties. Here, observing the inconsistency $\{f, d_2\}$ might suggest that the actual strategies for the boat being observed are not those of Figure 12 but those of Figure 11.

We may therefore interpret the difference between the two relations and complexes of Figures 11 and 12 as follows:

- The relations and complexes tell an *actor* what behavior to *avoid*, in order to *be successful* at deception.
- The relations and complexes tell an *observer* how to *look for* inconsistencies in behavior, in order to *detect* deception.

Inferences, Free Faces, Geometry

The simplicial complex of Figure 12 is a cone with apex $\mathtt{d}_1$. This geometry arises because *every* permissible strategy in the relation of that figure contains the downstream transition $\mathtt{d}_1$. Said differently, if a boat is incapable of moving upstream through the passage with the strong current, then there is no choice but to include the downstream transition in every permissible strategy for that boat. (Recall that a permissible strategy consists of a *maximal* set of motions that a boat *might* make to attain its destination without cycling.) One can thus infer the transition $\mathtt{d}_1$ "for free", i.e., without observing or learning of the motion directly, assuming one knows that the boat has a weak motor.

Let us therefore remove vertex $\mathtt{d}_1$ from the complex of Figure 12 to obtain the following:

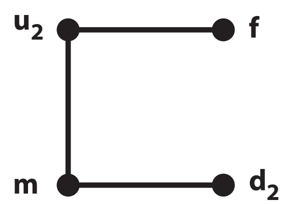

Figure 13: In the complex of Figure 12, the downstream transition $\mathtt{d}_1$ is a cone apex. Removing $\mathtt{d}_1$ produces the simplicial complex shown here. The free vertices of this complex highlight the implications $\mathtt{f} \implies \mathtt{u}_2$ and $\mathtt{d}_2 \implies \mathtt{m}$ that occur when transition $\mathtt{u}_1$ does not exist, as can also be seen from the relation of Figure 12.

This simplified complex highlights some inferences that are possible when observing boats known to have weak motors (as usual, we also assume that the strategies under consideration are permissible strategies and that each boat only follows strategies it can execute):

- $\mathtt{f} \implies \mathtt{u}_2$: If one knows that a boat is heading to the fishing area, then one can infer that the boat must move upstream when traversing the mild current.
- $\mathtt{d}_2 \implies \mathtt{m}$: If one observes a boat moving downstream over the mild current, then the boat must be heading to the marina.

These conclusions reflect the geometry of the fishing area and marina relative to the passages. Said differently, the geometry of the simplicial complex models the geometry of the river in such a way that one can draw conclusions about boat motions from the free faces in the complex. We saw this property as well when modeling paths in the example of Section 1.1.

1.3 Hidden State

The examples and discussion of this introduction have assumed that the underlying state is known to the observer. For example, the analysis of Section 1.2 assumed a state space defined by the river geometry, along with strategies whose motions were not allowed to revisit states.

In reality, a skipper might have additional *hidden state*. The restriction on revisiting states would then apply to the composite of observable and hidden states but not necessarily to the observable states alone. For instance, a skipper might be willing to revisit some parts of the river once or twice, relative to some hidden internal counter (perhaps fuel consumption). A skipper could then selectively hide some motions and reveal other motions, in order to bluff the capability of a strong motor. (An observer unaware of the skipper's hidden state might infer $\mathtt{u}_1$ from observations $\{\mathtt{f}, \mathtt{d}_2\}$ even though the boat only has a weak motor and has made surreptitious cyclic motions not involving $\mathtt{u}_1$ in order to give the appearance of a strong motor.)

One possible approach for dealing with such hidden state is to construct many possible models of that state, then hope to observe inconsistencies in behavior to rule out or imply some of these models. We leave such higher-order deception and detection for future work.

2 Review of Prior Work and Notation

This section briefly reviews key concepts and notation regarding strategy complexes and relations. Detailed discussions of strategies and strategy complexes appear in [4, 5]. The connection of strategies to relations appears in [6]. Background material on topology may be found in [11, 13], on posets in [15], and on privacy in [14, 2, 12, 3].

Assumption: All graphs, relations, and simplicial complexes in this report are finite.

2.1 Graphs and Strategies

This subsection reviews material on strategies, taken fairly directly from [5], with some descriptions verbatim.

Nondeterministic and Stochastic Graphs: We are interested in finite graphs, viewed as state spaces with errorful transitions. We model any such graph G as a pair $(V, \mathfrak{A})$, consisting of a finite set of *states* V and a finite collection of *actions* $\mathfrak{A}$. Each action $a \in \mathfrak{A}$ consists of a *source* state v and a nonempty set T of *targets*, with $v \in V$ and $\emptyset \neq T \subseteq V$.

If T consists of a single state t, we say that the action is *deterministic*. We may write a deterministic action a as $v \rightarrow t$, just like a directed edge in a directed graph. If T contains more than one state, there are two possibilities: The action is either *nondeterministic* or *stochastic*. We discuss each of these possibilities next. We may also regard a deterministic action as a special instance of a nondeterministic action and/or as a special instance of a stochastic action.

We write a nondeterministic action a as $v \rightarrow T$. The semantics of such an action are as follows: Action a may be executed whenever the system is at state v. When action a is executed, the system moves from state v to one of the target states in T. If $|T| > 1$, then the precise target attained is not predictable in advance, but is known after execution completes. Different execution instances of action a could attain different target states within T. An adversary might be choosing the target attained.

We write a stochastic action a as $v \rightarrow_p T$, with p a strictly positive probability distribution $p : T \rightarrow (0, 1]$, such that $\sum_{t \in T} p(t) = 1$. The semantics of a stochastic action are very similar to those of a nondeterministic action, except that the target state attained at execution time is now determined stochastically, according to the probability distribution p, rather than nondeterministically. Different execution instances of action a are assumed to be independent of each other.

We say that a graph $(V, \mathfrak{A})$ is a *pure nondeterministic* graph if all the actions in $\mathfrak{A}$ are either deterministic or nondeterministic. We say that a graph $(V, \mathfrak{A})$ is a *pure stochastic* graph if all the actions in $\mathfrak{A}$ are either deterministic or stochastic. This report is primarily interested in such pure graphs. However, see [5] for a discussion of graphs with a mix of deterministic, nondeterministic, and stochastic actions. See also Section 6.

Suppose action a has source v and targets T. We refer to each possible transition $v \rightarrow t$, with $t \in T$, as an *action edge (of action a)*.

Comment: We permit multiple actions to be distinct yet have the same source v and the same target set T (and the same probability distribution p if the actions are stochastic). Such duplication flexibility is useful, for instance when forming quotient graphs (see page 20).

Strategies and Strategy Complexes: We next define *strategies* and *strategy complexes* via a series of intermediate concepts. Intuitively, a strategy is a generalization of a control law, now viewed as a mapping from states to sets of actions.

Let $G = (V, \mathfrak{A})$ be a graph as on page 18, and suppose $\mathcal{A} \subseteq \mathfrak{A}$. We view $\mathcal{A}$ as a generalized control law as follows: Suppose the system is currently at state v. The set $\mathcal{A}$ may contain zero, one, or several actions with source v. The system *stops* moving precisely when $\mathcal{A}$ contains no action with source v. Otherwise, the system *must* execute some action $a \in \mathcal{A}$ with source v. If there are several such actions, any one of the actions might execute, determined nondeterministically. (Worst-case, an adversary might make the choice. In Section 1.2, perhaps sometimes a boat's skipper could.) Upon execution of action a, the system finds itself at one of the targets t of action a. The process then repeats, with t the system's new current state.

We are interested in only those control laws that eventually stop at some state or states. Intuitively, for pure nondeterministic graphs, this means that executing any of the actions contained in $\mathcal{A}$ will never cause the system to cycle (i.e., revisit a previously encountered state). For pure stochastic graphs, it means that no subset of the actions contained in $\mathcal{A}$ forms a recurrent Markov chain. We model these requirements with the following definitions, again taken fairly directly from [5]:

Let $G = (V, \mathfrak{A})$ be a graph as on page 18.

- With $a \in \mathfrak{A}$, let $\mathrm{src}(a)$ denote the source of action a. When $\mathcal{A} \subseteq \mathfrak{A}$, define $\mathcal{A}$'s *source set* $\mathrm{src}(\mathcal{A})$ as the set of all the individual actions' sources: $\mathrm{src}(\mathcal{A}) = \{\mathrm{src}(a) \mid a \in \mathcal{A}\}$.

- With $a \in \mathfrak{A}$, let $\mathrm{trg}(a)$ denote the set of targets of action a.

- Let $W \subseteq V$ and $a \in \mathfrak{A}$. We say that *action a moves off W (in G)* if $\mathrm{src}(a) \in W$ and one (or both) of the following is true: (i) action a is nondeterministic with all of its targets in $V \setminus W$, or (ii) action a is stochastic with at least one of its targets in $V \setminus W$. (The two requirements are identical when a is deterministic.)

- Let $\mathcal{A} \subseteq \mathfrak{A}$. We say $\mathcal{A}$ *contains a circuit* if, for some nonempty subset $\mathcal{B}$ of $\mathcal{A}$, no action of $\mathcal{B}$ moves off $\mathrm{src}(\mathcal{B})$. We say $\mathcal{A}$ *converges* or *is convergent* if $\mathcal{A}$ does not contain a circuit. Comment: If $\mathcal{A}$ is convergent and $V \neq \emptyset$, then necessarily $\mathrm{src}(\mathcal{A}) \neq V$.

- If $V \neq \emptyset$, then the *strategy complex Δ_G of G* is the simplicial complex whose underlying vertex set is $\mathfrak{A}$ and whose simplices are all the convergent subsets $\mathcal{A}$ of $\mathfrak{A}$. Every simplex of Δ_G is called a *strategy*. The empty simplex $\emptyset$ is one such strategy, modeling no motion. It appears in Δ_G whenever $V \neq \emptyset$. If $V = \emptyset$, then one would let Δ_G be the void complex, containing no simplices. (This report will always require $V \neq \emptyset$.)

- If σ is a strategy in Δ_G, then $V \setminus \mathrm{src}(\sigma)$ is called the *goal* or *goal set* of σ. The goal set consists of all states in the graph at which σ does not specify a motion. We may say that σ *converges to* $V \setminus \mathrm{src}(\sigma)$. If the goal set is a singleton $\{v\}$, we may refer directly to state v as σ's *goal*.

- Suppose $V \neq \emptyset$. We say that G is *fully controllable* if every nonempty subset of V is the goal set of some strategy in Δ_G. Observe that G is fully controllable if and only if every singleton state is the goal set of at least one *maximal* strategy (simplex) in Δ_G.

- Topologically, G is fully controllable if and only if Δ_G is homotopic to the sphere $\mathbb{S}^{n-2}$, with $n = |V|$. See [4, 5].

We will soon model strategy complexes via relations. The maximal simplices of a strategy complex will index the rows of the relation and the graph's actions will index the columns. Of interest will be how to reveal the constituent actions of a maximal strategy in such a way as to delay identification of the strategy for as long as possible.

We end this subsection with two definitions that will be useful later in the report:

Subgraphs: Suppose $G = (V, \mathfrak{A})$ is a graph as on page 18. By a *subgraph H of G* we mean a graph $H = (W, \mathfrak{B})$ in its own right such that $W \subseteq V$ and $\mathfrak{B} \subseteq \mathfrak{A}$.

Quotient graphs: Suppose $G = (V, \mathfrak{A})$ is a graph as on page 18 and suppose $\emptyset \neq W \subseteq V$. We define the *quotient graph* $G/W = (V', \mathfrak{A}')$ as follows:

- The state space is $V' = (V \setminus W) \cup \{\diamond\}$.

  Here $\diamond$ is a new state. It represents the set of states W all identified to one state.[‡]

- The actions $\mathfrak{A}'$ are in one-to-one correspondence with the actions $\mathfrak{A}$, but source and target states of actions in $\mathfrak{A}'$ are relabeled to match the new state space. Specifically, any source or target in $V \setminus W$ remains unchanged, while any source or target in W becomes $\diamond$.

  The relabeling of targets may identify some or all of the targets of an action $a \in \mathfrak{A}$. For a stochastic action of the form $v \rightarrow p\,T$, with $T \cap W \neq \emptyset$, one therefore sums the transition probabilities of the targets in $T \cap W$ in order to determine the transition probability to state $\diamond$ of the relabeled action $a' \in \mathfrak{A}'$.

  Comment: "one-to-one correspondence" means that distinct actions of $\mathfrak{A}$ remain distinct in $\mathfrak{A}'$ even if their sources become the same and their target sets become the same (and even if their probability distributions become the same, in the stochastic case).

The following facts are easy to establish:

1. A convergent set of actions in G may contain a circuit once one views the actions in G/W. (In particular, individual actions may become nonconvergent.) However, the set of actions remains convergent if its source set does not overlap W.

2. Any convergent set of actions in G/W will remain convergent if one views the actions back in their original form in G.

3. If G is fully controllable, then so is G/W.

Generalization: Suppose $W_1, \ldots, W_k$ are nonempty pairwise disjoint subsets of V, with $k \geq 1$. Let $\diamond_1, \ldots, \diamond_k$ be new and distinct states. The definition of quotient graph given above generalizes to this setting: For each $i = 1, \ldots, k$, we identify all states of G that lie in W_i to the single state $\diamond_i$. We denote the resulting quotient graph by $G/\{W_1, \ldots, W_k\}$.

[‡]In this report, *identify* typically means *determine identity of*, but sometimes, as here, it means *treat as same*.

2.2 Relations and Dowker Complexes

This subsection reviews material on relations, taken fairly directly from [6], with some descriptions verbatim.

Let X and Y be nonempty finite discrete spaces. A *relation* R on $X \times Y$ is a set of ordered pairs constituting a subset of the cross product $X \times Y$. We frequently view R as a matrix of blank and nonblank entries, with X indexing rows and Y indexing columns. We often refer to elements of X as *individuals* and to elements of Y as *attributes*.

For each $x \in X$, we let Y_x be the set of attributes that individual x has. Formally, $Y_x = \{\, y \in Y \mid (x, y) \in R\}$. We may view Y_x as the row of R indexed by x (or more precisely, as all the attributes with nonblank entries in the row indexed by x). Similarly, for each $y \in Y$, we let X_y be the set of individuals who have attribute y, that is, $X_y = \{x \in X \mid (x, y) \in R\}$. We may view X_y as the column of R indexed by y (or more precisely, as all the individuals with nonblank entries in the column indexed by y).

Given a relation R, we define two simplicial complexes, Φ_R and Ψ_R, as follows:

Φ_R is called the *Dowker attribute complex*. It has underlying vertex set Y and is generated by the rows of R. Ψ_R is called the *Dowker association complex*. It has underlying vertex set X and is generated by the columns of R. Thus:

$$\Phi_R = \bigcup_{x \in X} <Y_x> \qquad \text{and} \qquad \Psi_R = \bigcup_{y \in Y} <X_y>.$$

(The symbol $<\sigma>$ means the *simplicial complex generated by* σ, that is, the collection of all subsets of σ, including the empty simplex $\emptyset$ and σ itself.)

We define two *interpretation maps*, ϕ_R and ψ_R, as follows:

$$\begin{aligned} \phi_R(\sigma) &= \bigcap_{x \in \sigma} Y_x, \qquad \text{for any } \sigma \subseteq X, \\ \psi_R(\gamma) &= \bigcap_{y \in \gamma} X_y, \qquad \text{for any } \gamma \subseteq Y. \end{aligned}$$

Thus $\phi_R(\sigma)$ consists of all attributes shared by at least all the individuals in σ, while $\psi_R(\gamma)$ consists of all individuals who each have at least all the attributes in γ.

Observe that $\phi_R(\emptyset) = Y$ and $\psi_R(\emptyset) = X$.

One may regard ϕ_R both as an interpretation map as well as a test for membership in the Dowker complex Ψ_R. Specifically, for all $\sigma \subseteq X$, $\sigma \in \Psi_R$ if and only if $\phi_R(\sigma) \neq \emptyset$. Moreover, if $\emptyset \neq \sigma \in \Psi_R$, then $\emptyset \neq \phi_R(\sigma) \in \Phi_R$.

Similarly, for all $\gamma \subseteq Y$, $\gamma \in \Phi_R$ if and only if $\psi_R(\gamma) \neq \emptyset$. And $\emptyset \neq \psi_R(\gamma) \in \Psi_R$ whenever $\emptyset \neq \gamma \in \Phi_R$.

We say that individual $x \in X$ is *identifiable via R* whenever $\psi_R(Y_x) = \{x\}$. In other words, an individual x is identifiable when no other individual's attributes include all of x's attributes.

The following facts are useful to remember [6]:

1. Each of ϕ_R and ψ_R is inclusion-reversing.
2. For all $\gamma \subseteq Y,\ \ \gamma \subseteq (\phi_R \circ \psi_R)(\gamma)$. Similarly for $\psi_R \circ \phi_R$.
3. If γ is a maximal simplex of Φ_R, then $(\phi_R \circ \psi_R)(\gamma) = \gamma$. Similarly for $\psi_R \circ \phi_R$.
4. Each of the compositions $\phi_R \circ \psi_R$ and $\psi_R \circ \phi_R$ is idempotent.
5. $\phi_R \circ \psi_R \circ \phi_R = \phi_R$. Similarly with the roles of ϕ_R and ψ_R interchanged.

The maps ϕ_R and ψ_R define homotopy equivalences between the two Dowker complexes [6]. In particular, the compositions $\phi_R \circ \psi_R$ and $\psi_R \circ \phi_R$ are homotopy equivalent to the identity maps on their respective simplicial complexes. These equivalences allow one to construct a poset whose elements may be viewed as pairs of sets (σ, γ) satisfying $\emptyset \neq \sigma = \psi_R(\gamma) \in \Psi_R$ and $\emptyset \neq \gamma = \phi_R(\sigma) \in \Phi_R$. This poset has an encompassing lattice structure and is amenable to topological analysis: When the poset has high-dimensional homology, one can be assured that it contains long chains. We will not need the details of that poset construction in this report. Instead, we jump directly to one additional definition that we will need:

Informative Attribute Release Sequences: An *informative attribute release sequence (for relation R)*, abbreviated as *iars*, is a nonempty set of attributes in Y released in a particular sequential order

$$y_1, y_2, \ldots, y_k, \quad \text{with } k \geq 1,$$

satisfying

$$y_i \notin (\phi_R \circ \psi_R)(\{y_1, \ldots, y_{i-1}\}), \quad \text{for all } 1 \leq i \leq k.$$

In order to understand this last condition, recall from [6] that $(\phi_R \circ \psi_R)(\gamma) \setminus \gamma$, with $\gamma \in \Phi_R$, is the set of all attributes inferable from γ. For instance, one may have directly observed attributes γ for some unknown individual known to be modeled by relation R. Then the set of attributes $(\phi_R \circ \psi_R)(\gamma) \setminus \gamma$ is inferable without direct observation. Thus the condition above requires that no attribute y_i in the sequence be inferable from the attributes $y_1, \ldots, y_{i-1}$ released before y_i. In particular, y_1 must not be inferable "for free", i.e., without observation. (The cone apex $\mathtt{d}_1$ of the complex in Figure 12 on page 15 was inferable for free and thus would never appear in an informative attribute release sequence for the relation in that figure.)

Comment: An informative attribute release sequence $y_1, y_2, \ldots, y_k$ might not form a simplex in Φ_R, but any proper prefix of the sequence will. It may at first seem counterintuitive to have a nonsimplex be informative, but the inconsistency one obtains with the last attribute released may provide information in some relation containing R, as discussed in [6].

Of interest in some privacy settings is how long one can delay identifying an individual while revealing information: Given an individual x, how large can one make k in defining an informative attribute release sequence $y_1, y_2, \ldots, y_k$ for which $\psi_R(\{y_1, \ldots, y_k\}) = \{x\}$? Topology offers lower bounds [6]. In this report, we will consider that question with strategies in place of individuals and actions in place of attributes. We will argue from first principles.

Relations from Complexes: Suppose Γ is a nonvoid simplicial complex with underlying vertex set $Y \neq \emptyset$. We can define a relation R on $\mathfrak{M} \times Y$, with $\mathfrak{M}$ the maximal simplices of Γ:

$$R = \{(\gamma, y) \mid y \in \gamma \in \mathfrak{M}\}.$$

One readily sees that $\Phi_R = \Gamma$. (See footnote § for a side comment.)

Observe that every maximal simplex of Γ, i.e., every $\gamma \in \mathfrak{M}$, is identifiable via relation R.

Action Relations: Suppose $G = (V, \mathfrak{A})$ is a graph as on page 18, now with both $V \neq \emptyset$ and $\mathfrak{A} \neq \emptyset$. We may substitute Δ_G for Γ in the previous construction to obtain a relation on $\mathfrak{M} \times \mathfrak{A}$, with $\mathfrak{M}$ the maximal simplices of Δ_G. We refer to such a relation as an *action relation*, or more specifically, as *graph G's action relation*. In an action relation, maximal strategies[¶] play the role of individuals while actions play the role of attributes. The strategy relations of Section 1.2 were of this form, though there we were only considering subcomplexes of the full strategy complex.

In this context, informative *attribute* release sequences become informative *action* release sequences. We may thus ask the question:

How many actions can one reveal informatively before one has identified a maximal strategy?

Caution: The order via which actions are revealed in an informative action release sequence need not correspond to the order in which actions might be executed at runtime.

Terminology: Let G be a graph as on page 18. We will make statements of the form "maximal strategy σ in G contains informative action release sequence $a_1, a_2, \ldots, a_k$". This statement means that the following three conditions hold:

(a) σ is a maximal simplex in Δ_G.

(b) $\{a_1, a_2, \ldots, a_k\} \subseteq \sigma$.

(c) $a_1, a_2, \ldots, a_k$ is an informative attribute release sequence for G's action relation.

In particular, the order of the actions $a_1, a_2, \ldots, a_k$ is significant. If we view the same set of actions in a different order we obtain a different sequence. Consequently, a statement of the form "σ contains $k!$ different informative action release sequences" has meaning even when $|\sigma| = k$. (Caution: A permutation of an informative attribute release sequence need not itself be an informative attribute release sequence. See page 29.)

[§] Void vs. Empty: The *void complex* is the simplicial complex $\Gamma = \emptyset$, containing no simplices. In contrast, the *empty complex* is the simplicial complex $\Gamma = \{\emptyset\}$, containing a single simplex, namely the *empty simplex* $\emptyset$. One may construct R from the empty complex, assuming $Y \neq \emptyset$. In that case, $\mathfrak{M} = \{\emptyset\}$, $R = \emptyset$, and $\Phi_R = \{\emptyset\} = \Gamma$.

[¶] The term *maximal strategy (in G or Δ_G)* is synonymous with *maximal simplex (in Δ_G)*.

2.3 Minimal Nonfaces

Suppose Γ is a simplicial complex with underlying vertex set Y. A *minimal nonface* of (or in) Γ is a subset of Y that is not itself a simplex in Γ but all of whose proper subsets are simplices in Γ.

If we arrange the vertices of a minimal nonface in any order, we obtain an informative attribute release sequence. That fact is the content of our first lemma:

Lemma 1 (Minimal Nonfaces as Informative Attribute Release Sequences)**.** *Let R be a relation on $X \times Y$, with both X and Y nonempty. Suppose κ is a minimal nonface of Φ_R. Then any ordering of the attributes in κ is an informative attribute release sequence for R.*

Proof. Suppose the lemma fails for some minimal nonface κ of Φ_R. Necessarily, $\kappa \neq \emptyset$. Let $k = |\kappa|$. Then, for some ordering $y_1, y_2, \ldots, y_k$ of the attributes in κ, y_k must be implied by $y_1, \ldots, y_{k-1}$ (if $k = 1$, this means y_1 is inferable "for free"). Since κ is a minimal nonface of Φ_R, $\{y_1, \ldots, y_{k-1}\}$ is a simplex in Φ_R. Thus $\{y_1, \ldots, y_{k-1}\} \subseteq (\phi_R \circ \psi_R)(\{y_1, \ldots, y_{k-1}\}) \in \Phi_R$. By supposition, $y_k \in (\phi_R \circ \psi_R)(\{y_1, \ldots, y_{k-1}\})$. Consequently, $\kappa \subseteq (\phi_R \circ \psi_R)(\{y_1, \ldots, y_{k-1}\}) \in \Phi_R$, contradicting the assumption that κ is a nonface of Φ_R. □

Suppose all individuals in a relation R are identifiable. Then all rows of R are distinct and each row forms a maximal simplex in the attribute complex Φ_R. Suppose an observer has observed attributes η for some unknown individual x known to be modeled by relation R. Even if η is a proper subset of Y_x, it is possible that η identifies x, meaning $\psi_R(\eta) = \{x\}$. In that case, Y_x is the only maximal simplex containing η. Conversely, if the observed attributes η do not identify individual x, then η must be contained in some maximal simplex besides Y_x. Thus there exists an attribute y that is not one of x's attributes but that is consistent with all the observed attributes η of x, meaning $\eta \cup \{y\} \in \Phi_R$. The following lemma characterizes this situation more generally for a simplicial complex, in terms of minimal nonfaces.

Lemma 2 (Minimal Nonfaces between a Maximal Simplex and a Separate Vertex)**.** *Suppose Γ is a simplicial complex with underlying vertex set Y. Let γ be a maximal simplex of Γ and let $y \in Y$ such that $y \notin \gamma$. Define*

$$\mathcal{K} \;=\; \{\kappa \subseteq \gamma \cup \{y\} \mid \kappa \text{ is a minimal nonface of } \Gamma\}.$$

Suppose $\eta \subseteq \gamma$. Let $\eta' = \eta \cup \{y\}$.

Then $\eta' \in \Gamma$ if and only if $\kappa \setminus \eta' \neq \emptyset$ for every $\kappa \in \mathcal{K}$.

Comments: (i) $\mathcal{K} \neq \emptyset$, since γ is a maximal simplex in Γ and $y \notin \gamma$. (ii) If $\{y\} \notin \Gamma$, then $\mathcal{K}$ consists solely of $\{y\}$ and $\{y\} \setminus \eta' = \emptyset$, no matter what η is. Indeed, no η' can be in Γ. (iii) If $\{y\} \in \Gamma$, then γ cannot be the empty simplex. Every $\kappa \in \mathcal{K}$ now contains at least two vertices, namely y and some element of γ. Therefore the lemma's assertion for $\eta = \emptyset$ is clear. (iv) More generally, the lemma says: Even though vertex y cannot enlarge simplex γ, it may be able to enlarge a face η of γ. Such enlargement is possible precisely when the enlarged set contains no minimal nonfaces of the type described by $\mathcal{K}$.

Proof. Suppose $\eta' \in \Gamma$. If for some $\kappa \in \mathcal{K}$, $\kappa \setminus \eta' = \emptyset$, then η' would contain κ as a minimal nonface, a contradiction. Now suppose $\eta' \notin \Gamma$. Then η' must contain some minimal nonface, necessarily a set κ in $\mathcal{K}$ since $\eta' \subseteq \gamma \cup \{y\}$. Thus $\kappa \setminus \eta' = \emptyset$. □

Minimal Nonfaces in a Strategy Complex: Specializing to minimal nonfaces of a strategy complex yields additional results, as discussed below.

Lemma 3 (Minimal Nonfaces in Strategy Complexes). *Let $G = (V, \mathfrak{A})$ be a graph as on page 18, with $V \neq \emptyset$. Suppose κ is a minimal nonface of Δ_G. Then the actions in κ all have distinct sources and no action in κ moves off* $\mathrm{src}(\kappa)$ *in G.*

Proof. Let $k = |\kappa|$. Since $V \neq \emptyset$, $k > 0$.

Write $\kappa = \{a_1, \ldots, a_k\}$. For each $i = 1, \ldots, k$, define $\kappa_i = \kappa \setminus \{a_i\}$, that is, remove one action from κ. Then $\kappa_i \in \Delta_G$, for $i = 1, \ldots, k$. Thus, for every $\emptyset \neq \tau \subseteq \kappa_i$, some action in τ must move off $\mathrm{src}(\tau)$. On the other hand, $\kappa \notin \Delta_G$, so for some $\emptyset \neq \xi \subseteq \kappa$, no action in ξ moves off $\mathrm{src}(\xi)$. Consequently $\xi = \kappa$, establishing the second assertion of the lemma.

The first assertion is trivial if $k = 1$, so assume $k > 1$ and suppose $\mathrm{src}(a_1) = \mathrm{src}(a_2)$. Then $\mathrm{src}(\kappa) = \mathrm{src}(\kappa_1) = \mathrm{src}(\kappa_2)$. Some action in κ_1 moves off $\mathrm{src}(\kappa_1) = \mathrm{src}(\kappa)$. That contradicts the previous paragraph, thereby establishing the first assertion of the lemma. □

Interpreting Minimal Nonfaces in Strategy Complexes: Let us examine the meaning of minimal nonfaces for the two types of pure graphs discussed in this report. Assume the notation of Lemma 3 and its proof.

- Suppose G is a pure nondeterministic graph. Inductively, Lemma 3 produces a cycle of actions $a_1, \ldots, a_k$, such that $\mathrm{src}(a_{i+1}) \in \mathrm{trg}(a_i)$, for $i = 1, \ldots, k$ (here indices wrap around, so that a_{k+1} again means a_1). Moreover, for each action a_i, *exactly one* of the action's targets lies in $\mathrm{src}(\kappa)$; any additional targets lie outside $\mathrm{src}(\kappa)$. (Otherwise, one could create a shorter cycle and thus a proper subset of κ would be a nonface of Δ_G.)

- Suppose G is a pure stochastic graph. In the definition of "moves off" from page 19, the quantification over targets is different for stochastic actions than for nondeterministic actions. Consequently, Lemma 3 now implies that *all* targets of every action in κ must lie within $\mathrm{src}(\kappa)$. One may therefore create a subgraph H of G defined by $H = (\mathrm{src}(\kappa), \kappa)$. One sees that κ is also a minimal nonface in Δ_H, that Δ_H is the boundary complex[‖] on the set κ, and that H is a fully controllable pure stochastic graph. In fact, H defines an irreducible Markov chain [7, 10, 5].

[‖]The *boundary complex on the set Z* (with Z finite) is the simplicial complex whose underlying vertex set is Z and whose simplices are all the proper subsets of Z.

2.4 Sample Graphs, Relations, and Informative Action Release Sequences

This subsection provides examples of graphs, action relations, and strategy complexes, along with discussion of the extent to which strategy or goal identification may be delayed. (Actions here are deterministic or nondeterministic. Stochastic actions appear in Sections 5 and 6.)

2.4.1 A Directed Cycle Graph

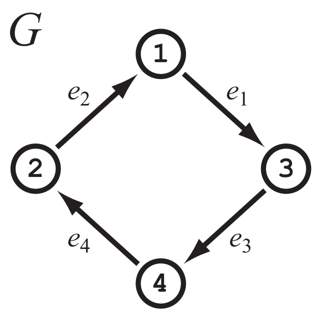

A	e_1	e_2	e_3	e_4	Goal
σ_1		•	•	•	1
σ_2	•		•	•	2
σ_3	•	•		•	3
σ_4	•	•	•		4

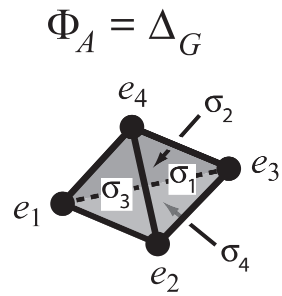

Figure 14: **Left Panel:** A deterministic graph with four states and four actions that form a directed cycle. **Middle Panel:** The graph's action relation A, along with each maximal strategy's goal. **Right Panel:** The Dowker attribute complex Φ_A (which is necessarily the same as the strategy complex Δ_G). It is a hollow tetrahedron. Vertices are actions and triangles are maximal strategies, as indicated by the labels.

As a first example, consider the directed graph G in the left panel of Figure 14. The graph contains four states and four directed edges. (The directed edges represent deterministic actions.) These edges form a directed cycle. Any proper subset of the four directed edges does not form a cycle. Consequently, any set of three directed edges forms a strategy, in fact a maximal strategy, that converges to one of the states in the graph, from any other state in the graph. For instance, the strategy σ_4, consisting of the set $\{e_1, e_2, e_3\}$ of directed edges, converges to state #4, for any initial starting state of the system.

Comment: Any subset of a maximal strategy is also a strategy, since it too will be acyclic. For instance, the set of directed edges $\{e_1, e_4\}$ is a strategy that stops at either state #2 or state #3. (The precise stopping point depends on the starting point during a particular execution of the strategy.) The set $\{e_1, e_4\}$ is a strategy but it it is *not* a *maximal* strategy.

The middle panel of Figure 14 shows graph G's action relation, describing each maximal strategy by its constituent actions. For each state v, there is a maximal strategy σ_v converging to that state from anywhere else in the graph. Therefore, G is fully controllable. The strategy complex Δ_G is in fact generated by four such maximal strategies, each consisting of three directed edges. Consequently, the strategy complex is a hollow tetrahedron, as shown in the right panel of the figure. In particular, the strategy complex contains a single minimal nonface, namely the set $\{e_1, e_2, e_3, e_4\}$, consisting of all four directed edges (actions) in the graph.

There are no free faces in the strategy complex, so it is impossible to infer any actions of a strategy from any actions revealed — "attribute privacy is preserved" [6]. Thus it is impossible to identify a maximal strategy uniquely if one knows only a proper subset of its actions. Each maximal strategy consists of three actions and has no free faces. Consequently, each maximal strategy contains 3! different informative action release sequences that identify the strategy, and each such sequence has length 3. For instance, the six sequences for strategy σ_4 are:

$$e_1, e_2, e_3 \qquad e_2, e_3, e_1 \qquad e_3, e_1, e_2$$
$$e_3, e_2, e_1 \qquad e_2, e_1, e_3 \qquad e_1, e_3, e_2.$$

Ability to Delay Strategy Identification: Let $G = (V, \mathfrak{A})$ be a fully controllable graph with $n = |V| > 1$. The following property holds [6]: For every state $v \in V$, there is some maximal strategy $\sigma_v \in \Delta_G$ such that σ_v has goal v and contains at least $(n-1)!$ different informative action release sequences of length at least $n-1$ each. For each such sequence $a_1, \ldots, a_\ell$, this means the following: An observer cannot infer (via G's action relation) that σ_v contains action a_i merely from knowing that σ_v contains the set $\{a_1, \ldots, a_{i-1}\}$ of actions appearing earlier in the sequence. In particular, an observer cannot identify σ_v uniquely before seeing all actions in the sequence $a_1, \ldots, a_\ell$. Moreover, $\ell \geq n-1$.

Comment: In the example of Figure 14, the six informative action release sequences of length 3 within each maximal strategy were permutations of the strategy's three constituent actions. In general, the $(n-1)!$ different sequences need not be permutations of each other.

2.4.2 A Graph with a Subspace Cycle

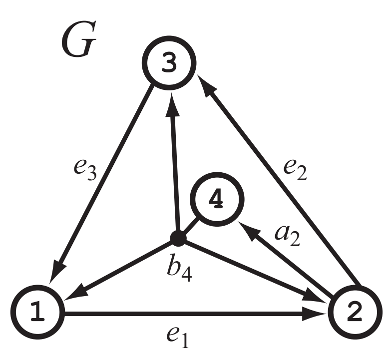

Figure 15: A pure nondeterministic graph with four states, $1, 2, 3, 4$, four deterministic actions, e_1, e_2, e_3, a_2, and one nondeterministic action, b_4.

As a second example, let us consider a graph with a directed cycle merely on a proper subset of the state space, as shown in Figure 15. The graph again consists of four states. The set of deterministic actions $\{e_1, e_2, e_3\}$ forms a directed cycle on the set of states $\{1, 2, 3\}$. In addition, there is a deterministic action a_2 that moves off this cycle space, specifically from state #2 to state #4. Finally, there is a nondeterministic action b_4 that moves from state #4 back to the set of target states $\{1, 2, 3\}$. (To say that the action is nondeterministic means that the precise target state attained cannot be predicted in advance, not even stochastically.)

The actions e_1, e_2, and e_3 form a directed cycle. Any two of these actions form a convergent strategy. The remaining two actions, a_2 and b_4, taken together, could cause the system to cycle

A	e_1	e_2	e_3	a_2	b_4	Goal
σ_1		•	•		•	1
σ_2	•		•		•	2
σ_3	•	•			•	3
σ_4	•		•	•		4
σ_{14}		•	•	•		$\{1,4\}$
σ_{34}	•	•		•		$\{3,4\}$

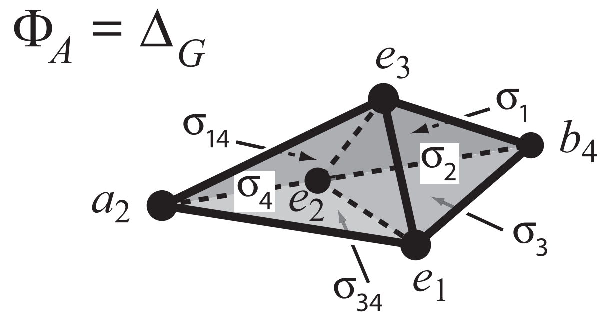

Figure 16: **Left Panel:** The action relation A for the graph G of Figure 15, along with each maximal strategy's goal (or goal set). **Right Panel:** The Dowker attribute complex Φ_A (which is necessarily the same as the strategy complex Δ_G). The complex consists of two party hats glued together, forming an $\mathbb{S}^2$ hole. Vertices are actions and triangles are maximal strategies, as indicated by the labels.

between states #2 and #4. Any one of these actions is convergent by itself. Therefore, the two sets of actions $\{e_1, e_2, e_3\}$ and $\{a_2, b_4\}$ each form a minimal nonface in the strategy complex Δ_G. In fact, these are the only minimal nonfaces in the strategy complex. They are independent of each other. Consequently, G's strategy complex is the simplicial join of the boundary of the triangle $\{e_1, e_2, e_3\}$ and the boundary of the edge $\{a_2, b_4\}$. In other words, Δ_G is a *suspension* [11, 15] of a triangle boundary. Figure 16 depicts this complex along with G's action relation. Observe that the complex is homotopic to $\mathbb{S}^2$, consistent with G being fully controllable.

As in the example of Section 2.4.1, Δ_G contains no free faces. So, again, it is impossible to identify a maximal strategy uniquely from a proper subset of its constituent actions.

One salient difference between this example and the previous one is that some maximal strategies now have goal sets with more than one state in them. For instance, strategy σ_{34}, consisting of actions $\{e_1, e_2, a_2\}$ has goal set $\{3, 4\}$. This multi-state goal arises because the actions e_1 and e_2, taken together, converge to state #3 *assuming* the system state lies within the subset of states $\{1, 2, 3\}$. However, if the starting state happens to be state #4, then the system will simply remain at that state. Consequently, the strategy $\{e_1, e_2\}$ has goal set $\{3, 4\}$. That strategy is not itself maximal. One can augment it either with action b_4, in which case the resulting maximal strategy would be σ_3, converging to state #3. Or one can augment $\{e_1, e_2\}$ with action a_2 to produce σ_{34}. Adding action a_2 introduces some nondeterminism at state #2, but nothing that changes the overall goal set; it remains $\{3, 4\}$.

Informative Action Release Sequences: As in the example of Section 2.4.1, the longest informative action release sequences within each maximal strategy in Δ_G are simply permutations of the strategy's constituent actions. Whenever a maximal strategy has no free faces, one can release its actions in any order without definitively identifying the strategy before all actions have been released.

In order to understand the more general picture, suppose we simply interchange the roles of strategies and actions in this example. The "individuals" are now e_1, e_2, e_3, a_2, b_4 and the "attributes" are σ_1, σ_2, σ_3, σ_4, σ_{14}, σ_{34}. We are thus interested in the Dowker association complex Ψ_A of the original action relation. That complex appears in Figure 17. (Comment: We are not asserting that this complex is the strategy complex of a fully controllable graph, merely using the complex to illustrate a point. One can however construct fully controllable graphs with strategy complexes that make the same underlying point: permutations of informative attribute release sequences need not themselves be informative attribute release sequences.)

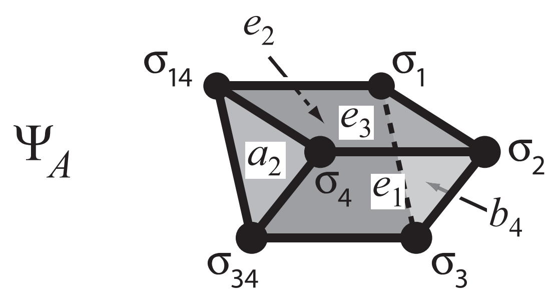

Figure 17: The Dowker association complex Ψ_A for the action relation of Figure 16, drawn as a hollow cylinder with a triangular cross-section and two triangular endcaps. The quadrilaterals drawn in the figure are actually solid tetrahedra, flattened in the figure for ease of viewing.

Each maximal simplex of Ψ_A continues to offer (at least) 3! different informative release sequences of length 3 each, with elements in each sequence drawn from the simplex's vertices. Two of the maximal simplices (namely, the "endcaps" in the figure) are solid triangles with no free faces, so their sequences are again simply permutations of each other. Three of the maximal simplices are solid tetrahedra. The undrawn "diagonals" of these tetrahedra are free faces, so releasing their endpoints would completely identify the tetrahedron. For instance, releasing vertices σ_1 and σ_4 identifies the tetrahedron labeled e_3. Consequently, one cannot simply choose arbitrary sequences of length 3 and expect them to be informative. Nonetheless, each tetrahedron does contain 16 informative release sequences of length 3. Here are the sequences for the tetrahedron labeled e_3:

$$
\begin{array}{llll}
\sigma_1, \sigma_2, \sigma_4 & \sigma_2, \sigma_1, \sigma_4 & \sigma_1, \sigma_2, \sigma_{14} & \sigma_2, \sigma_1, \sigma_{14} \\
\sigma_4, \sigma_{14}, \sigma_1 & \sigma_{14}, \sigma_4, \sigma_1 & \sigma_4, \sigma_{14}, \sigma_2 & \sigma_{14}, \sigma_4, \sigma_2 \\
\sigma_1, \sigma_{14}, \sigma_4 & \sigma_{14}, \sigma_1, \sigma_4 & \sigma_1, \sigma_{14}, \sigma_2 & \sigma_{14}, \sigma_1, \sigma_2 \\
\sigma_2, \sigma_4, \sigma_{14} & \sigma_4, \sigma_2, \sigma_{14} & \sigma_2, \sigma_4, \sigma_1 & \sigma_4, \sigma_2, \sigma_1.
\end{array}
$$

Incorporating Additional Constraints: Suppose, in some context, the system only executes strategies that converge to singleton goals. From an inference perspective, the action relation A and strategy complex Δ_G of Figure 16 would be misleading. To understand the possible inferences, one should consider a relation $A^{(1)}$ that models all the maximal strategies with singleton goal sets, and only those, as shown in Figure 18.

With this added information, it is no longer true that one can find 6 different informative action release sequences of length 3 within every maximal strategy. For instance, action a_2

$A^{(1)}$	e_1	e_2	e_3	a_2	b_4	Goal
σ_1		•	•		•	1
σ_2	•		•		•	2
σ_3	•	•			•	3
σ_4	•		•	•		4

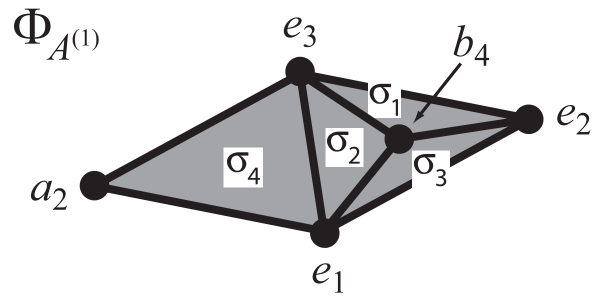

Figure 18: **Left Panel:** Modified relation from Figure 16, containing only those maximal strategies that each have a singleton state as goal set. **Right Panel:** The corresponding Dowker attribute complex $\Phi_{A^{(1)}}$, along with labels for actions and maximal strategies.

identifies strategy σ_4. Consequently, as soon as one releases that action, the other two actions in σ_4, if not previously released, would be implied. As a result, there are only two informative action release sequences of length 3 for identifying strategy σ_4, namely e_1, e_3, a_2 and e_3, e_1, a_2.

Similarly, action e_2 implies action b_4, again limiting the ordering of any sequences containing both those actions. Strategy σ_1 now contains only three, rather than six, informative action release sequences of length 3, namely:

$$e_3, b_4, e_2 \qquad b_4, e_3, e_2 \qquad b_4, e_2, e_3.$$

Ability to Delay Goal Identification: Suppose $G = (V, \mathfrak{A})$ is a fully controllable graph with $n = |V| > 1$. The following property holds [6]: For every state $v \in V$, there is some maximal strategy $\tau_v \in \Delta_G$ such that τ_v has goal v and contains an informative action release sequence whose sequential release leaves the goal ambiguous at least until all actions in the sequence have been revealed. Moreover, the sequence reduces the goal ambiguity by at most one state with each action revealed, so the sequence has length at least $n - 1$.

One sees this property in the complex of Figure 18 since: (i) every maximal simplex contains a vertex shared by three strategies with different goals and (ii) the vertex lies within one of the simplex's edges that is shared by two strategies with different goals.

2.4.3 An Augmented Cycle Graph

Let us augment the graph of Figure 14 with two nondeterministic actions, as shown in Figure 19. The actions are $a_1 = 1 \to \{2, 3\}$ and $a_2 = 2 \to \{3, 4\}$.

The new graph $\overline{G}$ has a strategy complex $\Delta_{\overline{G}}$ described by the action relation $\overline{A}$ shown in Figure 19. The complex is a partially puffed up version of the hollow tetrahedron from Figure 14, now consisting of three solid tetrahedra and two solid triangles glued together to enclose an $\mathbb{S}^2$ hole. Figure 20 shows the 1-skeleton of this complex.

The maximal strategies from the original strategy complex Δ_G are still present in $\Delta_{\overline{G}}$. Two of these strategies now lie within larger maximal simplices. For instance, strategy σ_2 for attaining goal state #2 in the original graph G consisted of the actions $\{e_1, e_3, e_4\}$, whereas now the corresponding maximal simplex $\overline{\sigma}_2$ consists of the actions $\{e_1, e_3, e_4, a_1\}$. The original

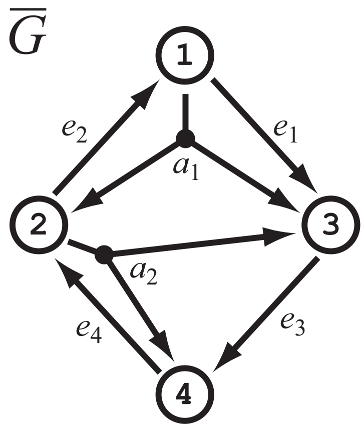

$\overline{A}$	e_1	e_2	e_3	e_4	a_1	a_2	Goal
$\overline{\sigma}_1$		•	•	•			1
$\overline{\sigma}_2$	•		•	•	•		2
$\overline{\sigma}_3$	•	•		•			3
$\overline{\sigma}_4$	•	•	•			•	4
$\overline{\sigma}_5$	•		•		•	•	4

Figure 19: **Left Panel:** A pure nondeterministic graph with four states, $1, 2, 3, 4$, four deterministic actions, e_1, e_2, e_3, e_4, and two nondeterministic actions, a_1, a_2. The graph here is the graph from Figure 14 augmented with two nondeterministic actions.

Right Panel: The graph's action relation $\overline{A}$, along with each maximal strategy's goal.

strategy σ_2 always executed action e_1 when the system was at state #1, thus transitioning to state #3. In the new graph $\overline{G}$, the new $\overline{\sigma}_2$ might execute either action e_1 or action a_1 at state #1, selected nondeterministically (possibly by an adversary). Action a_1 will transition either to state #2 (the goal) or to state #3. (If an adversary controls the outcome of action a_1, then the adversary might choose to make action a_1 mimic action e_1, in which case the old σ_2 and the new $\overline{\sigma}_2$ would behave equivalently.)

Actions e_2 and e_4 both appear in both the original strategies σ_1 and σ_3 of Δ_G. In the new graph $\overline{G}$, the set $\{e_2, a_1\}$ contains a circuit, as does the set $\{e_4, a_2\}$. Consequently, one cannot augment the strategies σ_1 and σ_3 with either of the actions a_1 or a_2. These strategies remain unchanged as one passes from G to $\overline{G}$, that is, $\overline{\sigma}_1 = \sigma_1$ and $\overline{\sigma}_3 = \sigma_3$.

There are five maximal strategies in the new graph, whereas there were four previously. New strategy $\overline{\sigma}_5$ has the same goal state, namely state #4, as does strategy $\overline{\sigma}_4$, but arrives there with different actions, trading off action e_2 for action a_1. The minimal nonface $\{e_2, a_1\}$ of $\Delta_{\overline{G}}$ hints at this possible tradeoff.

The ability to delay strategy identification mentioned on page 27, as well as our analysis of the original graph G, ensures that each of $\overline{\sigma}_1, \overline{\sigma}_2, \overline{\sigma}_3, \overline{\sigma}_4$ contains (at least) 6 informative action release sequences of length (at least) 3 each. What can we say about strategy $\overline{\sigma}_5$?

Releasing both of the two nondeterministic actions a_1 and a_2 identifies $\overline{\sigma}_5$. Releasing either one of these actions implies both deterministic actions in $\overline{\sigma}_5$. Consequently, one obtains the longest possible informative action release sequences within $\overline{\sigma}_5$ by first revealing the two deterministic actions (in either order) and then the two nondeterministic actions (in either order). Here are the four possible longest informative action release sequences within $\overline{\sigma}_5$:

$$e_1, e_3, a_1, a_2 \qquad e_1, e_3, a_2, a_1 \qquad e_3, e_1, a_1, a_2 \qquad e_3, e_1, a_2, a_1.$$

(The previous reasoning can be generalized and formalized using lattice representations of links, as discussed in [6], but we will not develop or use that machinery in this report.)

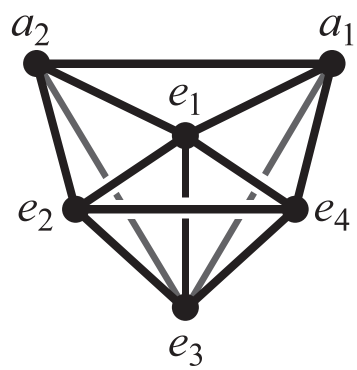

Figure 20: The 1-skeleton of the strategy complex $\Delta_{\overline{G}}$, with $\overline{G}$ the graph of Figure 19. The complex consists of three solid tetrahedra and two solid triangles, enclosing a hollow tetrahedron. The hollow tetrahedron has vertices $\{e_1, e_2, e_3, e_4\}$, just as in Figure 14.

Informative Action Release Sequences for Maximal Strategies in Pure Graphs: Suppose G is a fully controllable graph containing n states, with $n > 1$. Suppose further that G is either pure nondeterministic or pure stochastic. The remainder of this report will provide proofs that *every* maximal strategy in G contains at least one informative action release sequence of length at least $n - 1$. The proofs will be different for the two types of graphs. This property need not hold for graphs containing a mix of nondeterministic and stochastic actions. See [6] and Section 6 for counterexamples.

3 Basic Tools

Graphs and strategies were defined in Section 2.1. We wish to prove the following theorem:

Theorem 4 (Informative Action Release Sequences for Maximal Strategies)**.**
Let $G = (V, \mathfrak{A})$ be a fully controllable graph with $n = |V| > 1$. Suppose that G is either pure nondeterministic or pure stochastic. Then every maximal strategy in Δ_G contains an informative action release sequence of length at least $n - 1$.

Section 4 will prove Theorem 4 for the pure nondeterministic case, while Section 5 will prove Theorem 4 for the pure stochastic case.

This section builds some tools that are useful for both settings. Throughout this section the graph G may in fact contain a mix of deterministic, nondeterministic, and stochastic actions.

Terminology and Notation:

- We frequently abbreviate *informative action release sequence* as *iars*, for both the singular and plural cases.
- If G is a graph, the phrase *$a_1, \ldots, a_k$ is an iars for G* means that $a_1, \ldots, a_k$ is an informative action release sequence for G's action relation as defined on page 23.
- Suppose $G = (V, \mathfrak{A})$ is a graph and $G' = (V', \mathfrak{A}')$ is some quotient graph of G. Recall from page 20 that the sets of actions $\mathfrak{A}$ and $\mathfrak{A}'$ are in one-to-one correspondence. Corresponding actions $a \in \mathfrak{A}$ and $a' \in \mathfrak{A}'$ differ only in that the source and/or target states of action a' may have changed from those of action a in order to reflect the quotient graph's state identifications. One may therefore view any action of G' in the form a', with a being the unique action of G corresponding to a'. We will use this prime notation from now on.

Our first tool allows us to combine two strategies when one of the strategies comes from a subgraph and the other comes from a quotient graph formed by collapsing that subgraph to a single state.

Lemma 5 (Combining Quotient and Subgraph Strategies)**.** *Let $G = (V, \mathfrak{A})$ be a graph and let $H = (W, \mathfrak{B})$ be a subgraph of G. (As always, both V and W are assumed to be nonempty.)*

If $\sigma' \in \Delta_{G/W}$ and $\gamma \in \Delta_H$, then $\sigma \cup \gamma \in \Delta_G$.

(Prime notation indicates corresponding actions in a graph and a quotient, as discussed.)

Comment: Since all sources and targets of actions in γ lie within W, no such action is convergent in the quotient graph G/W. Therefore $\sigma \cap \gamma = \emptyset$.

Proof. Suppose $\emptyset \neq \tau \subseteq \sigma \cup \gamma$. We need to show that some action of τ moves off $\mathrm{src}(\tau)$ in G.

If $\mathrm{src}(\tau \backslash \gamma) \cap W = \emptyset$, then this assertion follows from Lemma 7.3(b)(i) in [5].

Otherwise, let $\kappa = \tau \backslash \gamma$. Then $\emptyset \neq \kappa' \subseteq \sigma' \in \Delta_{G/W}$, so some action b' in κ' moves off $\mathrm{src}(\kappa')$ in G/W. If b' is nondeterministic, then all its targets lie in the set $\big((V \backslash W) \cup \{\diamond\}\big) \setminus \mathrm{src}(\kappa')$. If b' is stochastic, then at least one of its targets lies in that set. (As usual, $\diamond$ represents W identified to a single state in G/W.) Since κ contains an action with source in W, κ' contains an action with source $\diamond$. Thus $\big((V \backslash W) \cup \{\diamond\}\big) \setminus \mathrm{src}(\kappa') \;=\; V \setminus \big(W \cup \mathrm{src}(\kappa')\big) \;=\; V \setminus \big(W \cup \mathrm{src}(\kappa)\big) \;\subseteq\; V \setminus \mathrm{src}(\tau)$. That means action b lies in τ and moves off $\mathrm{src}(\tau)$ in G. □

The next lemma ensures that unquotienting a simplicial informative action release sequence (iars) again produces an informative action release sequence, when the quotienting is over a proper subspace that is fully controllable.

Lemma 6 (Lifting Quotient IARS). *Let $G = (V, \mathfrak{A})$ be a graph and let $H = (W, \mathfrak{B})$ be a fully controllable subgraph of G, with $\emptyset \neq W \subsetneq V$.*

Suppose $a'_1, \ldots, a'_k$ is an iars for G/W, with $\{a'_1, \ldots, a'_k\} \in \Delta_{G/W}$ and $k \geq 1$.

Then $a_1, \ldots, a_k$ is an iars for G.

Proof. Let A be G's action relation. We need to show that

$$a_i \notin (\phi_A \circ \psi_A)(\{a_1, \ldots, a_{i-1}\}), \quad \text{for } i = 1, \ldots, k.$$

Suppose this assertion is false for some i. Then a_i is contained in every maximal simplex of Δ_G that contains $\{a_1, \ldots, a_{i-1}\}$.

Let $\mathfrak{A}'$ be the actions of the quotient graph G/W. Since $a'_1, \ldots, a'_k$ is an iars for G/W and $\{a'_1, \ldots, a'_k\} \in \Delta_{G/W}$, there exist actions τ', with $\emptyset \neq \tau' \subseteq \mathfrak{A}' \setminus \{a'_1, \ldots, a'_i\}$, such that

$$\{a'_1, \ldots, a'_{i-1}\} \cup \tau' \in \Delta_{G/W} \qquad \text{but} \qquad \{a'_1, \ldots, a'_i\} \cup \tau' \notin \Delta_{G/W}.$$

Let η' be a minimal nonface of $\Delta_{G/W}$ contained in $\{a'_1, \ldots, a'_i\} \cup \tau'$. (Since $\{a'_i\} \in \Delta_{G/W}$, η' contains at least two actions, one of them being a'_i.) By Lemma 3 on page 25, no action of η' moves off $\mathrm{src}(\eta')$ in G/W.

We consider two cases below, deriving a contradiction for each.

By Lemma 3, there are no further cases.

State $\diamond$ represents W identified to a single state in the quotient graph G/W, as per page 20.

I: <u>No action in η' has source $\diamond$:</u>

Then $\mathrm{src}(\eta) = \mathrm{src}(\eta')$, so no action of η moves off $\mathrm{src}(\eta)$ in G.

On the other hand, $\{a_1, \ldots, a_{i-1}\} \cup \tau \in \Delta_G$, by Fact 2 on page 20.
So $\{a_1, \ldots, a_i\} \cup \tau \in \Delta_G$, by the falsity assumption above. Since $\eta \subseteq \{a_1, \ldots, a_i\} \cup \tau$, that means $\emptyset \neq \eta \in \Delta_G$ and some action of η must move off $\mathrm{src}(\eta)$ in G, a contradiction.

II: <u>Exactly one action in η' has source $\diamond$:</u>

Suppose the source of the corresponding action in G is w. Then $w \in W$. Let $\kappa = \eta \cup \gamma$, with $\gamma \in \Delta_H$ a strategy that attains w from anywhere in W using actions of H. Then $\mathrm{src}(\kappa) = (\mathrm{src}(\eta') \setminus \{\diamond\}) \cup W$. Since no action of η' moves off $\mathrm{src}(\eta')$ in G/W and since γ has all its sources and targets in W, no action of κ moves off $\mathrm{src}(\kappa)$ in G.

Since $\{a'_1, \ldots, a'_{i-1}\} \cup \tau' \in \Delta_{G/W}$ and $\gamma \in \Delta_H$, Lemma 5 on page 33 implies that $\{a_1, \ldots, a_{i-1}\} \cup \tau \cup \gamma \in \Delta_G$. By the falsity assumption, $\{a_1, \ldots, a_i\} \cup \tau \cup \gamma \in \Delta_G$. Now $\kappa \subseteq \{a_1, \ldots, a_i\} \cup \tau \cup \gamma$, so $\emptyset \neq \kappa \in \Delta_G$, and some action of κ must move off $\mathrm{src}(\kappa)$ in G, again a contradiction. □

Combining Informative Action Release Sequences: The next lemma shows how one may combine an iars in a graph with an iars from a subgraph. The subsequent corollary leverages this result with those discussed earlier, showing how one may combine an iars from a quotient graph with an iars from a fully controllable subgraph. That combinability forms a stepping stone in several proofs during the rest of the report.

Lemma 7 (Combining Graph and Subgraph Informative Action Release Sequences)**.** *Let $G = (V, \mathfrak{A})$ be a graph and let $H = (W, \mathfrak{B})$ be a subgraph of G (with both V and W nonempty). Suppose $a_1, \ldots, a_k$, with $k \geq 1$, is an iars for G, such that:*

(i) $a_i \in \mathfrak{A} \setminus \mathfrak{B}$, for $i = 1, \ldots, k$, and

(ii) $\{a_1, \ldots, a_k\} \cup \tau \in \Delta_G$, for every $\tau \in \Delta_H$.

Suppose $b_1, \ldots, b_\ell$ is an iars for H, with $\{b_1, \ldots, b_\ell\} \in \Delta_H$ and $\ell \geq 1$.
Then $a_1, \ldots, a_k, b_1, \ldots, b_\ell$ is an iars for G, with $\{a_1, \ldots, a_k, b_1, \ldots, b_\ell\} \in \Delta_G$.

Comment: The lemma also holds when $k = 0$, meaning every iars for H is also an iars for G.

Proof. Suppose the iars part of the assertion is false. Let A be G's action relation.

Then, for some $i \in \{0, 1, \ldots, \ell - 1\}$, $b_{i+1} \in (\phi_A \circ \psi_A)(\{a_1, \ldots, a_k, b_1, \ldots, b_i\})$.
(When $i = 0$, this notation means $b_1 \in (\phi_A \circ \psi_A)(\{a_1, \ldots, a_k\})$.)

Consequently, every maximal simplex of Δ_G containing $\{a_1, \ldots, a_k, b_1, \ldots, b_i\}$ also contains b_{i+1}.

Since $b_1, \ldots, b_\ell$ is an iars for H's action relation, there exists a maximal simplex $\tau \in \Delta_H$ such that $\{b_1, \ldots, b_i\} \subseteq \tau$ but $\{b_1, \ldots, b_{i+1}\} \not\subseteq \tau$.

By assumption, $\{a_1, \ldots, a_k\} \cup \tau \in \Delta_G$. Consequently, $\{a_1, \ldots, a_k, b_{i+1}\} \cup \tau \in \Delta_G$. Thus $\tau \cup \{b_{i+1}\} = \mathfrak{B} \cap \big(\{a_1, \ldots, a_k, b_{i+1}\} \cup \tau\big) \in \Delta_H$, contradicting the maximality of τ in Δ_H. ☐

Corollary 8 (Lifting and Combining Informative Action Release Sequences)**.** *Let $G = (V, \mathfrak{A})$ be a graph and let $H = (W, \mathfrak{B})$ be a fully controllable subgraph of G with $\emptyset \neq W \subsetneq V$.*

Suppose $a'_1, \ldots, a'_k$ is an iars for G/W, with $\{a'_1, \ldots, a'_k\} \in \Delta_{G/W}$ and $k \geq 1$.

Suppose further that $b_1, \ldots, b_\ell$ is an iars for H, with $\{b_1, \ldots, b_\ell\} \in \Delta_H$ and $\ell \geq 1$.

Then $a_1, \ldots, a_k, b_1, \ldots, b_\ell$ is an iars for G, with $\{a_1, \ldots, a_k, b_1, \ldots, b_\ell\} \in \Delta_G$.

Proof. By Lemma 6, $a_1, \ldots, a_k$ is an iars for G. By Lemma 5, $\{a_1, \ldots, a_k\} \cup \tau \in \Delta_G$, for every $\tau \in \Delta_H$. Since actions of H become self-loops in G/W, $a_i \notin \mathfrak{B}$, for $i = 1, \ldots, k$. The desired result therefore follows from Lemma 7. ☐

Comment: The corollary also holds if one of k or ℓ is 0.

4 The Nondeterministic Setting

The aim of this section is to prove Theorem 4 from page 33 for the case in which the graph G is pure nondeterministic. Throughout, this section assumes that all graphs are pure nondeterministic, meaning each action is either deterministic or nondeterministic (but not stochastic). First, we need some additional definitions and results.

4.1 Hierarchical Cyclic Graphs

We start with a recursive definition:

Definition 9 (Hierarchical Cyclic Graph). *A pure nondeterministic graph $G = (V, \mathfrak{A})$ is a hierarchical cyclic graph if one of conditions (i) or (ii) holds:*

(i) $|V| = 1$ *and* $\mathfrak{A} = \emptyset$.

(ii) There exist $V_1, \ldots, V_k$, $\mathfrak{A}_1, \ldots, \mathfrak{A}_k$, $a_1, \ldots, a_k$, *with* $k > 1$, *such that:*

- *(a)* $V_1, \ldots, V_k$ *are nonempty pairwise disjoint subsets of* V *and* $V = \bigcup_{i=1}^{k} V_i$.
- *(b)* $\mathfrak{A}_i$ *consists of all actions in* $\mathfrak{A}$ *whose sources and targets lie in* V_i, *for* $i = 1, \ldots, k$.
- *(c)* $(V_i, \mathfrak{A}_i)$ *is a hierarchical cyclic graph, for* $i = 1, \ldots, k$.
- *(d)* $\mathfrak{A} = \{a_1, \ldots, a_k\} \cup \bigcup_{i=1}^{k} \mathfrak{A}_i$.
- *(e) For* $i = 1, \ldots, k$, $\mathrm{src}(a_i) \in V_i$ *and* $\mathrm{trg}(a_i) \subseteq V_{i+1}$ *(here indices wrap around, so* V_{k+1} *again means* V_1*).*

The decomposition above need not be unique. We implicitly assume a specific decomposition when stating that a graph is hierarchical cyclic. We refer to it as the tree decomposition of G.

A graph of type (i) is a leaf and a graph of type (ii) is a node.

When G is a node, we refer to the subgraphs $(V_1, \mathfrak{A}_1), \ldots, (V_k, \mathfrak{A}_k)$ *in G's tree decomposition as the children of G. Each subgraph* $(V_i, \mathfrak{A}_i)$ *is itself either a leaf or a node, with parent* $(V, \mathfrak{A})$. *When* $(V_i, \mathfrak{A}_i)$ *is a node, we may then speak of its children, and so forth. Transitively, we may therefore speak of all the nodes and leaves within G (that includes* $(V, \mathfrak{A})$*). Finally, we may speak of the root of the tree decomposition of G, meaning the node or leaf* $(V, \mathfrak{A})$, *i.e., G itself.*

For a graph of type (ii), the actions $a_1, \ldots, a_k$ *are the (top-level) cycle actions of G. Similarly, if N is any node within G, the cycle actions of N are the top-level cycle actions of N when N is viewed as a hierarchical cyclic graph in its own right.*

Comments and Observations:

- Given a hierarchical cyclic graph G of type (ii) as above, we can form the quotient graph $G/\{V_1, \ldots, V_k\}$ (see again page 20). This quotient graph has state space $\{\diamond_1, \ldots, \diamond_k\}$, where $\diamond_i$ represents all of V_i identified to a single state, for $i = 1, \ldots, k$.

  All actions in each $\mathfrak{A}_i$ become nonconvergent in $G/\{V_1, \ldots, V_k\}$ (actions in $\mathfrak{A}_i$ become self-loops on state $\diamond_i$), so we may ignore them. In contrast, each action a_i turns into a *deterministic transition* a'_i from state $\diamond_i$ to state $\diamond_{i+1}$.

We may therefore view the quotient graph $G/\{V_1, \ldots, V_k\}$ as the cycle graph

$$\diamond_1 \xrightarrow{a'_1} \diamond_2 \xrightarrow{a'_2} \cdots \xrightarrow{a'_{k-1}} \diamond_k \xrightarrow{a'_k} \diamond_1$$

(the first and last states in the diagram above are the same state, namely $\diamond_1$).

- More generally, suppose G is a hierarchical cyclic graph and $(W, \mathfrak{B})$ is some node that appears within the tree decomposition of G. We can form the quotient graph G/W. The quotienting identifies all of W to a single state $\diamond$. The actions $\mathfrak{B}$ become self-loops on state $\diamond$. Technically, G/W includes these self-loops, but there is no harm ignoring them, thereby allowing us to view G/W as a hierarchical cyclic graph. If $(W, \mathfrak{B})$ is G itself, then we may view G/W as the leaf $(\{\diamond\}, \emptyset)$. Otherwise, the tree decomposition of G/W is largely unchanged from that of G, except that one node, along with the subtree rooted at that node, has now become a leaf, and any actions of G with source or target states in W have had those states relabeled as $\diamond$. The only actions that become nonconvergent (by creating self-loops) are those in $\mathfrak{B}$, which we now ignore and discard.

- Conversely, suppose $(\{s\}, \emptyset)$ is a leaf that appears in the tree decomposition of a hierarchical cyclic graph $G = (V, \mathfrak{A})$. Suppose $H = (W, \mathfrak{B})$ is another hierarchical cyclic graph, with states and actions distinct from those of G.

  We can replace the leaf $(\{s\}, \emptyset)$ with node H, to form a new hierarchical cyclic graph $\overline{G} = (\overline{V}, \overline{\mathfrak{A}} \cup \mathfrak{B})$.

  Here $\overline{V} = (V \setminus \{s\}) \cup W$. In forming $\overline{\mathfrak{A}}$ from $\mathfrak{A}$, we have some choices:

  Suppose $a = v \to T$ is an action in $\mathfrak{A}$. We create a corresponding action $\overline{a} \in \overline{\mathfrak{A}}$ as follows:

  - If $v = s$, we let $\overline{v}$ be *any* state in W and define $\overline{a} = \overline{v} \to T$.
  - If $s \in T$, we let S be *any* nonempty subset of W and then define $\overline{a} = v \to \overline{T}$, with $\overline{T} = (T \setminus \{s\}) \cup S$.
  - In all other cases, $\overline{a} = a$.

- A special case of the previous construction is to replace a single state s in a hierarchical cyclic graph with a deterministic cycle on some new set of states, while adjusting all other actions of the encompassing graph accordingly. Actions of the encompassing graph that used to start at s now start at an arbitrary state of the cycle. Actions that used to have a transition to s now might transition to one or more states comprising the cycle.

- Every hierarchical cyclic graph is fully controllable and each of its actions is convergent.

- Conversely, the lemma below shows that every fully controllable pure nondeterministic graph contains a hierarchical cyclic subgraph with the same state space. (There may be more than one such subgraph.)

Lemma 10 (Hierarchical Cyclic Subgraphs)**.** *Let $G = (V, \mathfrak{A})$, with $V \neq \emptyset$, be a fully controllable pure nondeterministic graph. Then G contains a hierarchical cyclic subgraph $H = (V, \mathfrak{B})$.*

Proof. By strong induction on $|V|$. The base case $|V| = 1$ is clear, so suppose $|V| > 1$. For every state v in V one can find a nonlooping *deterministic* action with target v (since G is fully controllable and pure nondeterministic). Backchaining such actions produces a deterministic cycle $\mathcal{C}$ on some subspace W of V (possibly all of V), containing at least two states.

Consider $G/W = (V', \mathfrak{A}')$. Here $V' = (V \setminus W) \cup \{\diamond\}$, with $\diamond$ representing W. G/W is fully controllable (by Fact 3 on page 20) and pure nondeterministic, with $0 < |V'| < |V|$, so the induction hypothesis applies. We therefore obtain a hierarchical cyclic subgraph $H' = (V', \mathfrak{B}')$ of G/W.

We may now replace leaf $(\{\diamond\}, \emptyset)$ in H' with cycle $\mathcal{C}$ on state space W. When adjusting the encompassing actions $\mathfrak{B}'$, we choose sources and targets so as to undo any relabeling of states that occurred in forming G/W. These adjustments produce a hierarchical cyclic subgraph $H = (V, \mathfrak{B})$ of G. ☐

4.2 Core Cycle Actions, Leaf Covers, Disruptive Sets of Actions

Suppose $H = (W, \mathfrak{B})$ is a hierarchical cyclic graph with $|W| > 1$. Each state t of W appears as a leaf $(\{t\}, \emptyset)$ in the tree decomposition of H and has some parent node $N = (U, \mathfrak{E})$. Some action $c_t \in \mathfrak{E}$, necessarily a cycle action of N, must be deterministic with *target* t. We refer to c_t as *t's core cycle action.* This action is determined uniquely by t and the tree decomposition of H. (H may contain multiple deterministic actions with target t, but one and only one of those actions will be a cycle action in the parent node of $(\{t\}, \emptyset)$.)

With that construct in mind, we now make a series of definitions and observations.

Definition 11 (Core Cycle Actions)**.** *Let $H = (W, \mathfrak{B})$ be a hierarchical cyclic graph. The set $\mathfrak{C}_H$ of* <u>core cycle actions of H</u> *is*

$$\mathfrak{C}_H = \begin{cases} \{\, c_t \mid t \in W \,\}, & \text{if } |W| > 1 \quad \textit{(with } c_t \textit{ as defined above);} \\ \emptyset, & \textit{otherwise.} \end{cases}$$

Definition 12 (Leaf Covers)**.** *Let N be a node in a hierarchical cyclic graph H. We say that N* <u>covers only leaves in H</u> *whenever each of N's children is a leaf in H's tree decomposition.*

Definition 13 (Disruptive Sets of Actions)**.** *Let $H = (W, \mathfrak{B})$ be a hierarchical cyclic graph and suppose $\mathfrak{D} \subseteq \mathfrak{B}$. We say that $\mathfrak{D}$ is* <u>disruptive (in H)</u> *whenever the following condition is satisfied:*

For every node that covers only leaves in H,
at least two of the node's cycle actions are missing from $\mathfrak{D}$.

Observations:

- When a hierarchical cyclic graph $H = (W, \mathfrak{B})$ contains at least two states, $|\mathfrak{C}_H| = |W|$.
- A node covers only leaves in H if and only if all the node's cycle actions lie in $\mathfrak{C}_H$.
- The empty set of actions is always disruptive, even when H is a leaf.

4.3 Cycle-Breaking Strategies

Sets of actions that do not contain any node's full set of cycle actions are convergent and may be arranged informatively, as the following definition and lemmas make precise.

Definition 14 (Cycle-Breaking)**.** *Suppose $H = (W, \mathfrak{B})$ is a hierarchical cyclic graph. A set of actions $\tau \subseteq \mathfrak{B}$ is* cycle-breaking (in H) *if, for each node N in the tree decomposition of H, τ does not contain all of N's cycle actions.*

Lemma 15 (Cycle-Breaking is Convergent)**.** *Suppose τ is a cycle-breaking set of actions in a hierarchical cyclic graph H. Then $\tau \in \Delta_H$.*

Proof. By structural induction on the tree decomposition of H. The lemma holds if H is a leaf, since only $\tau = \emptyset$ is possible. Otherwise, suppose the children of H are $(W_1, \mathfrak{B}_1), \ldots, (W_k, \mathfrak{B}_k)$. Inductively, the lemma holds for the set of actions $\tau \cap \mathfrak{B}_i$ in the hierarchical cyclic graph $(W_i, \mathfrak{B}_i)$, for $i = 1, \ldots, k$. Let σ consist of the top-level cycle actions of H that are in τ. Since $H/\{W_1, \ldots, W_k\}$ is a directed cycle graph (see top of page 37) and since τ is cycle-breaking, $\sigma' \in \Delta_{H/\{W_1,\ldots,W_k\}}$. Thus, by repeated application of Lemma 5 on page 33, $\tau \in \Delta_H$. □

Caution: Not all strategies in a hierarchical cyclic graph need be cycle-breaking (see page 54).

Lemma 16 (Cycle-Breaking is Informative)**.** *Suppose τ is a nonempty cycle-breaking set of actions in a hierarchical cyclic graph H. Then some ordering of all the actions in τ is an informative action release sequence for H.*

Proof. The proof will associate to each leaf and node of H an informative action release sequence, with the sequence associated to the root of H comprising all of τ.

For the purposes of this proof, it will be convenient to consider the empty sequence of actions as an informative action release sequence. Since τ is nonempty, the final sequence produced below will be nonempty, satisfying the standard requirement of page 22 that informative attribute release sequences be nonempty.

Base Case: Associate to each leaf of H the empty sequence.

Inductive Step: Consider a node N of H and assume each child C of N has an associated informative action release sequence consisting of all the actions of τ that appear in the graph C. View N as a hierarchical cyclic subgraph in its own right, and form the quotient graph N' obtained by identifying each child to a singleton state. The cycle actions of N create a deterministic directed cycle in N'. This cycle forms a minimal nonface in $\Delta_{N'}$. By Lemma 1 on page 24, any sequential ordering of the directed edges comprising this cycle forms an iars for N', any proper subset of which is convergent. Let $\{a_1, \ldots, a_\ell\}$ be the set of N's cycle actions in τ, this being $\emptyset$ with $\ell = 0$ when none of N's cycle actions lie in τ. Since τ is cycle-breaking, the reasoning just given implies $\{a'_1, \ldots, a'_\ell\} \in \Delta_{N'}$ and $a'_1, \ldots, a'_\ell$ is an iars for N'.

Since the children of N are fully controllable subgraphs of N, repeated application of Corollary 8 on page 35 shows that $a_1, \ldots, a_\ell, c_1, \ldots, c_m$ is an iars for N, with $c_1, \ldots, c_m$ being some concatenation of all the informative action release sequences associated to N's children. Associate $a_1, \ldots, a_\ell, c_1, \ldots, c_m$ to N. Observe that this iars consists of all actions of τ that appear in the graph N. Associated to H itself therefore is an iars consisting of all of τ. □

4.4 Markings

Let $H = (W, \mathfrak{B})$ be a hierarchical cyclic graph. We will view each node of H as being either *marked* or *unmarked*. Each node is unmarked initially. Later, we will define an algorithm that *marks* nodes according to some criteria. Once marked, a node remains marked.

In order to consider marking a node N, we first require that each child of N be either a leaf or an already marked node. The collection of marked nodes at any instant therefore defines a set $\mathcal{M}$ of *maximal marked nodes*, consisting of those nodes that are marked but have no marked parent. After some nodes have been marked, $\mathcal{M} = \{(W_1, \mathfrak{B}_1), \ldots, (W_\ell, \mathfrak{B}_\ell)\}$, for some $\ell \geq 1$, with the sets $W_1, \ldots, W_\ell$ nonempty and pairwise disjoint. We may therefore form the quotient graph $H/\{W_1, \ldots, W_\ell\}$, which we abbreviate as $H/\mathcal{M}$. We view $H/\mathcal{M}$ as a hierarchical cyclic graph, much as on page 37, by discarding any actions that have become self-loops.

The process will be iterative, adding an additional node to the collection of marked nodes with each step. We abbreviate the notation by writing $\mathcal{M}^{(j)}$ to mean the maximal marked nodes at the j^{th} step, with $j \geq 1$, and by writing $H^{(j)}$ to mean $H/\mathcal{M}^{(j)}$. We also define $\mathcal{M}^{(0)} = \emptyset$ and $H^{(0)} = H$.

Observation: Any node covering only leaves in $H^{(j)}$ corresponds to a node in H that is not yet marked but that could be marked at the $(j+1)^{st}$ step, and vice-versa.

4.5 Forward Projections

Strategies in a pure nondeterministic graph define partial orders. One may view those partial orders as forward projections of possible system states.

Definition 17 (A Strategy's Partial Order)**.** *Let $G = (V, \mathfrak{A})$ be a pure nondeterministic graph. If $\sigma \in \Delta_G$, then* <u>σ induces a partial order $\geq_\sigma$ on V</u> *as follows:*

For each $w, v \in V$, $w \geq_\sigma v$ if and only if either $w = v$ or there exist actions $a_1, \ldots, a_k \in \sigma$, with $k \geq 1$, such that:

(i) $\mathrm{src}(a_1) = w$,

(ii) $\mathrm{src}(a_{i+1}) \in \mathrm{trg}(a_i)$, *for $i = 1, \ldots, k-1$, and*

(iii) $v \in \mathrm{trg}(a_k)$.

Thus, $w \geq_\sigma v$ if and only if w is v or the system might move from w to v while executing strategy σ (in the diagram below, $a_i \in \sigma$, $v_i = \mathrm{src}(a_i)$, and $v_{i+1} \in \mathrm{trg}(a_i)$, for $i = 1, \ldots, k$):

$$w = v_1 \xrightarrow{a_1} v_2 \xrightarrow{a_2} \cdots\ v_k \xrightarrow{a_k} v_{k+1} = v.$$

The partial order $\geq_\sigma$ is well-defined since σ cannot create any cycles.

Definition 18 (Forward Projection)**.** *Suppose $G = (V, \mathfrak{A})$ is a pure nondeterministic graph. Let $\sigma \in \Delta_G$ and $\emptyset \neq W \subseteq V$. The* <u>forward projection of W under σ</u> *is the set*

$$\mathcal{F}_\sigma(W) \;=\; \{v \in V \mid w \geq_\sigma v,\ \textit{for some } w \in W\}.$$

In other words, $\mathcal{F}_\sigma(W)$ consists of all states that the system might pass through or stop at, assuming the system starts at some state in W and moves according to strategy σ. (In some papers, *forward projection* refers only to the states the system might stop at. Here, *forward projection* includes all states through which the system might move, including starting states.)

Lemma 19 (Disjoint Forward Projections — Core Cycle Actions).

Let $H = (W, \mathfrak{B})$ be a hierarchical cyclic graph and suppose $\tau \in \Delta_H$.

Define $\tau_+ = \mathfrak{C}_H \cap \tau$ and $\tau_- = \mathfrak{C}_H \setminus \tau$.

For each $c \in \tau_-$, let $J_c = \mathcal{F}_{\tau_+}(\{t\})$, with t being the unique target of action c. Then:

(a) The sets in the family $\{J_c\}_{c \in \tau_-}$ are pairwise disjoint.

(b) Suppose further that τ is disruptive. Let $c \in \tau_-$. Write $c = w \to t$. Then $w \notin J_c$.

In words: We split the core cycle actions $\mathfrak{C}_H$ of H into two sets, those that lie in the strategy τ and those that do not. The first set is itself a strategy, so we can consider forward projections under that strategy. For each core cycle action that is not in τ, we consider the forward projection of that action's target state. The lemma asserts that the resulting forward projections are pairwise disjoint. Moreover, if τ is disruptive, then each such forward projection does not loop back far enough to include the source state of its generating core cycle action.

These properties will help us later to construct minimal nonfaces from which we can then extract an informative action release sequence that is sufficiently long to establish Theorem 4.

Proof. (a) Let $\geq$ be the partial order induced by τ_+ on W. Suppose $v \in J_c \cap J_d$, with $c, d \in \tau_-$. Write $c = w \to t$ and $d = u \to s$. Then $t \geq v$ and $s \geq v$. Since $\tau_+ \subseteq \mathfrak{C}_H$, backchaining from v produces a unique backwards path of action edges in τ_+, with each edge actually being a deterministic action. (The path could be degenerate, consisting of no edges, just the state v.) That backwards path eventually encounters both t and s, establishing that t and s are comparable. For example, $s \to \cdots \to t \to \cdots \to v$ would establish $s \geq t$. Since core cycle actions $w \to t$ and $u \to s$ are missing from τ_+, this is only possible if $s = t$, meaning $c = d$.

(b) Suppose $w \in J_c$, with $c \in \tau_-$ and $c = w \to t$. Arguing as in (a), we now obtain a cycle:

$$w \to t \to w_1 \to \cdots \to w_k = w, \quad \text{with } k \geq 1.$$

All but one of the actions comprising this cycle lie in τ_+.

The exception is $w \to t$, which lies in τ_-.

All the actions comprising the cycle lie in $\mathfrak{C}_H$. Consider any action $u \to s$ of $\mathfrak{C}_H$. The depth** of the leaf $(\{u\}, \emptyset)$ in the tree decomposition of H must be greater than or equal to the depth of the leaf $(\{s\}, \emptyset)$. Consequently, all the states in the cycle appear in H as leaves at the same depth and with the same parent node. The cycle must therefore consist of that parent node's cycle actions and the parent node cannot contain any other children. So, the parent node covers only leaves. Since τ is disruptive, at least two of the node's cycle actions lie in τ_-, not just one, establishing a contradiction. ☐

**Here, the depth of a node or leaf in a tree is defined recursively as follows:
The depth of the tree's root is 0. The depth of a child is one more than the depth of its parent.

4.6 Quotienting until Disruption

The proof path now is to iteratively mark and quotient by nodes that prevent a strategy from being disruptive. Concurrently, one assembles several sets of actions that satisfy a property similar to the disjointness of forward projections described in Lemma 19.

Construction 20 (Acyclic Dissection). *Let $H = (W, \mathfrak{B})$ be a hierarchical cyclic graph. Suppose $\tau \subseteq \mathfrak{B}$. An* acyclic dissection $(\tau_\circ, \tau_+, \tau_-, \xi)$ of τ in H *is defined iteratively as follows:*

1. Initialize $\tau^{(0)} = \tau$ and $\kappa^{(0)} = \emptyset$. Assume all nodes in the tree decomposition of H are unmarked and initialize $H^{(0)} = H$, as per Section 4.4.

   Set Done to `true` if τ is disruptive in H and to `false` otherwise.

2. While not Done, run the following loop, starting from $j = 0$:

   (a) At this stage, $\tau^{(j)}$ consists of actions in $H^{(j)}$ and is not disruptive in $H^{(j)}$. Let N be some unmarked node in H such that the corresponding quotient node N' in $H^{(j)}$ covers only leaves and at most one of the cycle actions in N' is absent from $\tau^{(j)}$.

   (b) Suppose N' has k cycle actions $\{c'_1, \ldots, c'_k\}$. Discard one of these, so that the rest all lie in $\tau^{(j)}$. Without loss of generality, assume one may discard c'_k. Now let

   $$\kappa^{(j+1)} \;=\; \kappa^{(j)} \cup \{c_1, \ldots, c_{k-1}\}.$$

   Inductively: $\kappa^{(j+1)}$ consists of (unquotiented) actions in H. In fact, $\kappa^{(j+1)} \subseteq \tau$.

   (c) Mark node N, then let $H^{(j+1)}$ be the quotient graph formed from the resulting maximal marked nodes, as per page 40, again viewed as a hierarchical cyclic graph.

   (d) Suppose $H^{(j+1)} = (W', \mathfrak{D}')$. Let $\tau^{(j+1)} = \{a' \in \mathfrak{D}' \mid a \in \tau\}$. So $\tau^{(j+1)}$ is nearly the same as τ', except that $\tau^{(j+1)}$ ignores any action of τ whose source and targets all lie within any one maximal marked node of H. (Prime notation indicates the correspondence between an action in H and its relabeled form in a quotient graph.)

   (e) If $\tau^{(j+1)}$ is disruptive in $H^{(j+1)}$, set Done to `true`. The loop ends. Otherwise, the loop continues, with $j+1$ in place of j.

3. If τ was already disruptive in H, let $H^* = H$, $\tau^* = \tau$, and $\tau_\circ = \emptyset$. Otherwise, let $H^* = H^{(j+1)}$, $\tau^* = \tau^{(j+1)}$, and $\tau_\circ = \kappa^{(j+1)}$, with $j+1$ as above when the loop ends. In either case, τ^* is disruptive in H^*. Finally, let $\mathfrak{C} = \{c \in \mathfrak{B} \mid c' \in \mathfrak{C}_{H^*}\}$. In other words, $\mathfrak{C}$ is the set of actions in H that become core cycle actions in the quotient graph H^*.

4. Define ξ as follows (H contains a marked node if and only if the loop of step 2 was run):

   Start with $\xi = \emptyset$. Then, for each *unmarked* node N in H, let $\mathcal{C}_N$ be N's cycle actions. If $\mathcal{C}_N \cap \tau$ is a proper subset of $\mathcal{C}_N$, add all of $\mathcal{C}_N \cap \tau$ to ξ. Otherwise, select an action c in $\mathcal{C}_N \setminus \mathfrak{C}$. Add the actions $\mathcal{C}_N \setminus \{c\}$ to ξ. (Why does c exist? If not, let N' be the node in H^* corresponding to N. It is well-defined since N is unmarked. Then N' would cover only leaves in H^* and thus τ^* would not be disruptive in H^*, a contradiction.)

5. Step 3 defined $\tau_\circ$. Now define $\tau_+ = \mathfrak{C} \cap \xi$, and $\tau_- = \mathfrak{C} \setminus \xi$.

Lemma 21. *Construction 20 produces an acyclic dissection* $(\tau_\circ, \tau_+, \tau_-, \xi)$ *of* τ *such that:*

(i) $\tau_+ \subseteq \xi$ *and* $\tau_- \subseteq \mathfrak{B}$,

(ii) $\tau_\circ \cup \xi \subseteq \tau$ *and* $\tau_\circ \cap \xi = \emptyset$,

(iii) $\tau_\circ \cup \xi$ *is cycle-breaking in* H, *and*

(iv) $\tau_- \cap \tau = \emptyset$.

Proof. The loop in step 2 of Construction 20 runs at most a finite number of times, since the graph H is finite. As a result, an acyclic dissection $(\tau_\circ, \tau_+, \tau_-, \xi)$ of τ is well-defined by step 5.

Assertions (i), (ii), and (iii) are clear from the construction.

To establish assertion (iv), suppose $a \in \tau_- \cap \tau$. Let prime notation denote quotienting from H to H^*, with H^* as defined in step 3 of the construction, and assume the rest of the notation from the construction.

Then $a \in \mathfrak{C}$, $a \in \tau$, and $a \notin \xi$.

So $a' \in \mathfrak{C}_{H^*}$, implying $a \in \mathcal{C}_N$, with $\mathcal{C}_N$ the cycle actions of some unmarked node N in H.

If $\mathcal{C}_N \cap \tau$ is proper subset of $\mathcal{C}_N$, then $a \in \xi$, by step 4 of the construction, producing a contradiction.

So $\mathcal{C}_N \cap \tau = \mathcal{C}_N$. Let c be the action removed in step 4 of the construction. So $c \notin \mathfrak{C}$. Since $a \notin \xi$, $a \in \mathcal{C}_N$, and $\mathcal{C}_N \setminus \{c\} \subseteq \xi$, it must be that $a = c$, but that contradicts $a \in \mathfrak{C}$. □

And here is a generalization of Lemma 19:

Lemma 22 (Disjoint Forward Projections)**.** *Suppose* $H = (W, \mathfrak{B})$ *is a hierarchical cyclic graph and* $\tau \subseteq \mathfrak{B}$. *Construct* $(\tau_\circ, \tau_+, \tau_-, \xi)$ *from* τ *as per Construction 20. Let* $\eta = \tau_\circ \cup \tau_+$.

Given $c \in \tau_-$, *write* $c = w \to T$ *and define*

$$J_c = \mathcal{F}_\eta(T).$$

(The definition is sensible since η *is cycle-breaking in* H *and so* $\eta \in \Delta_H$.*)*

Then:

(a) The sets in the family $\{J_c\}_{c \in \tau_-}$ *are pairwise disjoint.*

(b) For each $c \in \tau_-$, $\mathrm{src}(c) \notin J_c$.

Proof. Let prime notation denote quotienting from H to H^*, where sensible, with H^* as in step 3 of the construction. Write $H^* = (W', \mathfrak{D}')$, viewed with a tree decomposition derived from that of H.

Since ξ arises only from cycle actions of unmarked nodes, each action in ξ' is a well-defined convergent action in $\mathfrak{D}'$. By construction, ξ' is cycle-breaking in H^*, so $\xi' \in \Delta_{H^*}$. Since $\xi' \subseteq \tau^*$, ξ' is disruptive in H^*. Consequently, Lemma 19 applies to the graph H^* and the disruptive strategy ξ'.

In H, we have $\tau_+ = \mathfrak{C} \cap \xi$ and $\tau_- = \mathfrak{C} \setminus \xi$. Therefore, each action $c \in \tau_+$ corresponds to an action $c' \in \tau'_+ = \mathfrak{C}_{H^*} \cap \xi'$ in H^*, and each action $c \in \tau_-$ corresponds to an action $c' \in \tau'_- = \mathfrak{C}_{H^*} \setminus \xi'$. (Recall the comment about "one-to-one correspondence" on page 20.)

Let $\geq$ be the partial order induced by η on W and let $\geq^*$ be the partial order induced by τ'_+ on W'. Suppose $v \geq w$, with $v, w \in W$. Then $v' \geq^* w'$, with $v', w' \in W'$ being the state relabelings of v and w, respectively. (Why? If there is a path of action edges from v to w with the actions drawn from η, then there is a path of action edges from v' to w' with the actions drawn from η'. Some of the action edges between states in W may become self-loops when sources and targets are relabeled as states in W'. Indeed, $v' = w'$ is possible even if $v \neq w$. Any such self-loops could only come from actions in $\tau_\circ'$. Conversely, all actions in $\tau_\circ'$ are self-loops. One discards those actions in forming H^*, leaving only τ_+' from η'. Thus there is a path of action edges from v' to w' with the actions drawn from τ_+'.)

It follows that $v \in J_c$ implies $v' \in J_{c'}$, with $J_{c'}$ defined for H^* and ξ' as in Lemma 19, now using c' in place of c, τ_+' in place of τ_+, and τ_-' in place of τ_-. (To see this, write $c = w \to T$. The set of targets T becomes a single state $t' \in W'$, since $c' \in \mathfrak{C}_{H^*}$. Write $c' = w' \to t'$. If $v \in J_c$, then $t \geq v$ for some $t \in T$, so $t' \geq^* v'$, and thus $v' \in J_{c'}$.)

Consequently, Lemma 19 establishes the claims of the current lemma. □

4.7 Alternate Development: Quotienting until Disruption

This subsection restates Construction 20 recursively without mentioning markings, then provides induction proofs of the corresponding lemmas. The key steps are the same as before. The rest of Section 4 will prove Theorem 4 for pure nondeterministic graphs using the earlier iterative construction, side-stepping any issue of strategy maximality in quotient graphs. Section 5 will engage that issue when proving Theorem 4 for pure stochastic graphs.

Construction 23 (Alternate Construction: Acyclic Dissection)**.** *Let $H = (W, \mathfrak{B})$ be a hierarchical cyclic graph. Suppose $\tau \subseteq \mathfrak{B}$. An* acyclic dissection $(\tau_\circ, \tau_+, \tau_-, \xi)$ of τ in H *is defined recursively as follows:*

I. Suppose τ is disruptive in H:

  1. Define ξ as follows, starting from $\xi = \emptyset$:
     For each node N in H, let $\mathcal{C}_N$ be N's cycle actions. If $\mathcal{C}_N \cap \tau$ is a proper subset of $\mathcal{C}_N$, add all of $\mathcal{C}_N \cap \tau$ to ξ. Otherwise, there is at least one action c in $\mathcal{C}_N \setminus \mathfrak{C}_H$. (If not, then N would cover only leaves in H and thus τ would not be disruptive.) Pick one such action c and add the remaining actions $\mathcal{C}_N \setminus \{c\}$ to ξ.
  2. Let $\tau_\circ = \emptyset$, $\tau_+ = \mathfrak{C}_H \cap \xi$, and $\tau_- = \mathfrak{C}_H \setminus \xi$.

II. Suppose τ is not disruptive in H:

  1. Let $N = (U, \mathfrak{E})$ be a node in H that covers only leaves and at most one of whose cycle actions is absent from τ. The actions $\mathfrak{E}$ are necessarily N's cycle actions. Discard one of those actions, so the rest all lie in τ. Denote that resulting set by $\mathcal{C}$.
  2. Let $H^* = (W', \mathfrak{D}')$ be the hierarchical cyclic graph formed from the quotient graph H/U by discarding self-loops. Let $(\tau_\star', \tau_+', \tau_-', \xi')$ be a recursively constructed acyclic dissection of $\tau' \cap \mathfrak{D}'$ in H^*. (As usual, prime notation describes the correspondence between actions of H and actions of H/U.)
  3. Now define the sets of actions τ_+, τ_-, and ξ by unquotienting, that is, by direct correspondence from the sets of actions τ_+', τ_-', and ξ', respectively. Finally, let $\tau_\circ = \tau_\star \cup \mathcal{C}$, with $\tau_\star$ formed from $\tau_\star'$ by unquotienting.

Proof of Lemma 21, assuming alternate acyclic dissection given by Construction 23:

Proof. The construction terminates because H is finite and each recursive invocation of the construction replaces a node with a leaf.

The proof of the specific assertions is by induction, with Case I of the construction defining the base case and Case II defining the inductive step:

I: In the base case, assertions (i), (ii), and (iii) are immediate from the construction. Assertion (iv) follows as it did in the earlier proof of Lemma 21, but now working directly with H rather than needing to form a quotient.

II: Inductively, we assume assertions (i)–(iv) hold for an acyclic dissection $(\tau'_\star, \tau'_+, \tau'_-, \xi')$ of $\tau' \cap \mathfrak{D}'$ in H^*, using the notation from the construction. Then:

- (i) $\tau'_+ \subseteq \xi'$, so $\tau_+ \subseteq \xi$ and $\tau'_- \subseteq \mathfrak{D}'$, so $\tau_- \subseteq \mathfrak{D} \subseteq \mathfrak{B}$.
- (ii) $\tau'_\star \cup \xi' \subseteq \tau' \cap \mathfrak{D}'$, so $\tau_\circ \cup \xi = \tau_\star \cup \mathcal{C} \cup \xi \subseteq \tau$, since $\mathcal{C} \subseteq \tau$.
  $\tau'_\star \cap \xi' = \emptyset$, so $\tau_\star \cap \xi = \emptyset$. Since all actions in $\mathfrak{C}$ become self-loops and are discarded when forming H^* from H, $\mathcal{C} \cap \xi = \emptyset$, and so $\tau_\circ \cap \xi = \emptyset$.
- (iii) $\tau'_\star \cup \xi'$ is cycle-breaking in H^*, so $\tau_\circ \cup \xi = \tau_\star \cup \mathcal{C} \cup \xi$ is cycle-breaking in H, since $\mathcal{C}$ consists of a proper subset of one node's cycle actions and since $\tau_\star \cup \xi$ does not include any of that node's actions.
- (iv) $\tau'_- \cap \tau' \cap \mathfrak{D}' = \emptyset$, so $\tau_- \cap \tau \cap \mathfrak{D} = \emptyset$. Since $\tau_- \subseteq \mathfrak{D}$, it follows that $\tau_- \cap \tau = \emptyset$. □

Proof of Lemma 22, assuming alternate acyclic dissection given by Construction 23:

Proof. Again by induction:

I: The base case follows from Lemma 19, with ξ in place of τ, since ξ is a disruptive strategy by Lemmas 21 and 15.

II: Inductively, the argument is much the same as in the earlier proof of this lemma. One assumes the assertions hold for the hierarchical cyclic quotient graph H^*. In moving back to H, one state of H^* turns back into a cycle of states, with all but one of the cycle actions added to $\tau_\star$ to form $\tau_\circ$. The other sets of actions in the dissection do not change as one moves back from H^* to H, except for relabelings of sources and targets. Consequently, execution paths of $\tau_\circ \cup \tau_+$ (within H) imply execution paths of $\tau'_\star \cup \tau'_+$ (within H^*), thereby establishing the lemma's assertions for H. □

4.8 Acyclic Dissection Sizes

This subsection measures the size of the set $\tau_\circ \cup \tau_+ \cup \tau_-$ in an acyclic dissection. When reading the lemma below, recall that Construction 20 marks nodes in H.

Lemma 24 (Subgraph Sizes)**.** *Suppose $H = (W, \mathfrak{B})$ is a hierarchical cyclic graph and $\tau \subseteq \mathfrak{B}$. Let $H^* = (W', \mathfrak{D}')$ and $\tau_\circ$ be derived from τ as per step 3 in Construction 20 on page 42.*

For each $u \in W'$, define $(W_u, \mathfrak{B}_u)$ as follows: If $u \in W$, let $(W_u, \mathfrak{B}_u)$ be the leaf $(\{u\}, \emptyset)$ of H. If $u \notin W$, let $(W_u, \mathfrak{B}_u)$ be the maximal marked node of H for which u represents W_u.

Then $|\tau_\circ \cap \mathfrak{B}_u| = |W_u| - 1$, for each $u \in W'$.

Proof. The proof is by induction on the iteration count j in the loop of Construction 20, now using $H^{(j)}$ in place of H^*, $\kappa^{(j)}$ in place of $\tau_\circ$, and with the collection of marked nodes dependent on j. The base case, $j = 0$, corresponds to all u being in W, for which the lemma's assertion is clear. Inductively, suppose the lemma's assertion is true for $H^{(j)}$.

In forming $H^{(j+1)}$ from $H^{(j)}$, one marks an unmarked node $N = (V, \mathfrak{A})$ of H whose corresponding node N' in $H^{(j)}$ covers only leaves. So N is now a maximal marked node. Let $W^{(j)}$ be the states of $H^{(j)}$, $W^{(j+1)}$ the states of $H^{(j+1)}$, and V' the states of N'. Then $W^{(j+1)} = \big(W^{(j)} \setminus V'\big) \,\cup\, \{\diamond\}$, with $\diamond$ representing the states V' identified to a singleton.

Let $\{c_1, \ldots, c_k\}$ and $\{(V_1, \mathfrak{A}_1), \ldots, (V_k, \mathfrak{A}_k)\}$ be the cycle actions and children of N in H, respectively. Without loss of generality, $\kappa^{(j+1)} = \kappa^{(j)} \cup \{c_1, \ldots, c_{k-1}\}$. One has $k = |V'| > 1$.

Case I: Suppose $u \in W^{(j+1)} \cap W^{(j)}$. Then the definition of $(W_u, \mathfrak{B}_u)$ is the same via $H^{(j+1)}$ as via $H^{(j)}$. So $|\kappa^{(j)} \cap \mathfrak{B}_u| = |W_u| - 1$. Since $(W_u, \mathfrak{B}_u)$ is either a leaf or a marked node of H at the j^{th} iteration of the loop in Construction 20, $\mathfrak{B}_u$ contains none of node N's cycle actions. Thus $\kappa^{(j+1)} \cap \mathfrak{B}_u = \kappa^{(j)} \cap \mathfrak{B}_u$ and so $|\kappa^{(j+1)} \cap \mathfrak{B}_u| = |W_u| - 1$, inductively.

Case II: Suppose $u = \diamond$. Then the definition of $(W_\diamond, \mathfrak{B}_\diamond)$ via $H^{(j+1)}$ is N, so $W_\diamond = V$ and $\mathfrak{B}_\diamond = \mathfrak{A}$. The states V' of N' in $H^{(j)}$ are in one-to-one correspondence with the children $\{(V_i, \mathfrak{A}_i)\}$ of N. Inductively, $|\kappa^{(j)} \cap \mathfrak{A}_i| = |V_i| - 1$. Again, $\kappa^{(j+1)} \cap \mathfrak{A}_i = \kappa^{(j)} \cap \mathfrak{A}_i$. Moreover,

$$W_\diamond \;=\; V_1 \cup \cdots \cup V_k \qquad \text{and} \qquad B_\diamond \;=\; \{c_1, \ldots, c_k\} \cup \mathfrak{A}_1 \cup \cdots \cup \mathfrak{A}_k.$$

By reasoning about markings, one further knows that $\kappa^{(j)} \cap \{c_1, \ldots, c_k\} = \emptyset$. Therefore

$$\kappa^{(j+1)} \cap B_\diamond \;=\; \{c_1, \ldots, c_{k-1}\} \;\cup\; \bigcup_{i=1}^{k} \big(\kappa^{(j)} \cap \mathfrak{A}_i\big)$$

and

$$\begin{aligned} |\kappa^{(j+1)} \cap B_\diamond| \;&=\; (k-1) \;+\; \sum_{i=1}^{k} |\kappa^{(j)} \cap \mathfrak{A}_i| \\ &=\; (k-1) \;+\; \sum_{i=1}^{k} \big(|V_i| - 1\big) \\ &=\; -1 \;+\; \sum_{i=1}^{k} |V_i| \\ &=\; |W_\diamond| - 1. \end{aligned}$$

□

Corollary 25 (Dissection Sizes). *Suppose $H = (W, \mathfrak{B})$ is a hierarchical cyclic graph and $\tau \subseteq \mathfrak{B}$. Let $\tau_\circ$, τ_+, τ_-, and H^* be derived from τ as per Construction 20 on page 42.*

Let $n = |W|$ and $m = |\tau_\circ \cup \tau_+ \cup \tau_-|$.

If H^* is a leaf, then $m = n - 1$. If H^* is a node, then $m = n$.

Proof. Suppose H^* is a leaf. Then $H^* = (\{u\}, \emptyset)$, for some u, and $\tau_+ = \tau_- = \emptyset$. Using the notation of Lemma 24, $(W_u, \mathfrak{B}_u)$ must be all of H. The lemma then implies that $m = |\tau_\circ| = |\tau_\circ \cap \mathfrak{B}| = |\tau_\circ \cap \mathfrak{B}_u| = |W_u| - 1 = |W| - 1 = n - 1$, as claimed.

Suppose H^* is a node. Let W' be the states of H^*. For each $u \in W'$, let $(W_u, \mathfrak{B}_u)$ be defined as in Lemma 24. Since $\tau_\circ$ is formed from cycle actions in marked nodes of H, $\tau_\circ = \bigcup_{u \in W'} (\tau_\circ \cap \mathfrak{B}_u)$. We also know that $W = \bigcup_{u \in W'} W_u$. Thus, by the lemma,

$$\begin{aligned} |\tau_\circ| &= \sum_{u \in W'} |\tau_\circ \cap \mathfrak{B}_u| \\ &= \sum_{u \in W'} (|W_u| - 1) \\ &= |W| - |W'| \\ &= n - |W'|. \end{aligned}$$

Let $\mathfrak{C}$ and $\mathfrak{C}_{H^*}$ be as in Construction 20. Then $\tau_+ \cup \tau_- = \mathfrak{C}$, so $|\tau_+ \cup \tau_-| = |\mathfrak{C}| = |\mathfrak{C}_{H^*}| = |W'|$. Consequently, $m = |\tau_\circ| + |\tau_+ \cup \tau_-| = n$, as claimed. □

4.9 Informative Action Release Sequences for Maximal Strategies

This subsection assembles the previous results to prove Theorem 4 for pure nondeterministic graphs.

Lemma 26 (Minimal Nonfaces Overlapping Forward Projections). *Let $G = (V, \mathfrak{A})$ be a fully controllable pure nondeterministic graph with $V \neq \emptyset$ and suppose $H = (V, \mathfrak{B})$ is a hierarchical cyclic subgraph of G. (Recall that H exists, by Lemma 10 on page 37.)*

Let σ be a maximal strategy in Δ_G, define $\tau = \sigma \cap \mathfrak{B}$, let $(\tau_\circ, \tau_+, \tau_-, \xi)$ be an acyclic dissection of τ obtained from Construction 20, and set $\eta = \tau_\circ \cup \tau_+$.

Suppose $c \in \tau_-$. Write $c = w \to T$ and define

$$\begin{aligned} \mathcal{K}_c &= \{\kappa \subseteq \sigma \cup \{c\} \mid \kappa \textit{ is a minimal nonface of } \Delta_G\}, \\ J_c &= \mathcal{F}_\eta(T), \\ \sigma_c &= \{a \in \sigma \mid \mathrm{src}(a) \in J_c \textit{ and } \mathrm{trg}(a) \not\subseteq J_c\}. \end{aligned}$$

Then $\mathcal{K}_c \neq \emptyset$ and $\kappa \cap \sigma_c \neq \emptyset$ for every $\kappa \in \mathcal{K}_c$.

Comment: The lemma tells us that every action $c \in \tau_-$ is part of a minimal nonface of Δ_G whose remaining actions lie in σ, and that at least one of those actions has its source, but not all its targets, in the forward projection of c's targets. Here the forward projection is based on those actions of σ that lie in the acyclic dissection sets $\tau_\circ$ and τ_+. Intuitively, it is useful to think of H as a single node defining a directed cycle, with the projection of σ onto that cycle being disruptive. Disruption means that the cycle splits into at least two pairwise disjoint directed arcs, as follows: The cycle edges present in σ constitute τ_+, the cycle edges missing from σ constitute τ_-, and $\tau_\circ$ is empty in this simple scenario. There is one directed arc for each action $c \in \tau_-$, starting at c's target. Each arc is formed from contiguous action edges of τ_+. Each arc has a forward projection flow defined on it by the directionality of those action edges. A directed arc ends when it encounters the source of another action in τ_-. An arc may be degenerate, consisting of a single state. The lemma says that, for each missing cycle edge c, there is some action $a \in \sigma$ whose source lies in an arc that starts at c's target, such that at least one of a's targets lies outside this arc and such that a and c appear together in a minimal nonface of Δ_G. As we will see shortly, the "a or c?" choice is therefore informative.

Proof. By Lemma 21(iv) on page 43, $c \notin \sigma$. So, since σ is maximal in Δ_G, $\mathcal{K}_c \neq \emptyset$.

Let $\kappa \in \mathcal{K}_c$ be given. Define $\gamma = \{a \in \kappa \mid \mathrm{src}(a) \in J_c\}$. Since κ is a minimal nonface of Δ_G, no action of κ moves off $\mathrm{src}(\kappa)$, by Lemma 3 on page 25. Since $c \in \kappa$ and $\{c\} \in \Delta_G$, $\gamma \neq \emptyset$. By Lemma 22(b) on page 43, $\mathrm{src}(c) \notin J_c$, so $c \in \kappa \setminus \gamma$. Thus $\emptyset \neq \gamma \subsetneq \kappa$ and $\gamma \subseteq \sigma$.

Now suppose the lemma's second assertion is false for this κ. Then every action in γ has all its targets in J_c. Pick some $a \in \gamma$. Since $a \in \kappa$, there exists $b \in \kappa$ such that $\mathrm{src}(b) \in \mathrm{trg}(a) \subseteq J_c$. We see therefore that $b \in \gamma$ and that no action of γ moves off $\mathrm{src}(\gamma)$. Consequently, $\gamma \notin \Delta_G$, which contradicts κ being a minimal nonface of Δ_G. □

Imagine revealing actions of some secret maximal strategy $\sigma \in \Delta_G$ to an observer who knows G but initially merely that σ is maximal in Δ_G. Suppose c is an action in τ_-, as previously defined. So $c \notin \sigma$ and $\sigma \cup \{c\} \notin \Delta_G$. Let σ_c be as before. The next corollary says that so long as one has not explicitly revealed any actions of σ_c, the observer cannot exclude the possibility that one is revealing actions of some maximal strategy other than σ, some strategy that does include action c. Moreover, there exists some unrevealed and unimplied action in σ_c that one may yet release informatively. (The explicitly revealed actions may imply some actions in σ_c, but so long as none of the explicitly revealed actions themselves lie in σ_c, these assertions hold.)

Corollary 27 (Informative Actions in Forward Projections)**.** *Let the hypotheses and notation be as in Lemma 26. In particular, σ is maximal in Δ_G and $c \in \tau_-$.*

Suppose $\gamma \subseteq \sigma$ such that $\gamma \cap \sigma_c = \emptyset$.

Let A be G's action relation and define $\overline{\gamma} = (\phi_A \circ \psi_A)(\gamma)$. Then:

(i) $\overline{\gamma} \cup \{c\} \in \Delta_G$.

(ii) $\mathcal{K}_c \neq \emptyset$ *and* $(\kappa \cap \sigma_c) \setminus \overline{\gamma} \neq \emptyset$ *for every* $\kappa \in \mathcal{K}_c$.

Proof. We may prove (i) by establishing that $\psi_A(\overline{\gamma} \cup \{c\}) \neq \emptyset$. By reasoning similar to that on page 122 in [6], $\psi_A(\overline{\gamma} \cup \{c\}) = \psi_A(\gamma \cup \{c\})$, so it is enough to show that $\gamma \cup \{c\} \in \Delta_G$. Suppose this is false. Then there exists a minimal nonface κ of Δ_G such that $\kappa \subseteq \gamma \cup \{c\} \subseteq \sigma \cup \{c\}$, so $\kappa \in \mathcal{K}_c$. By Lemma 26, $\kappa \cap \sigma_c \neq \emptyset$. That establishes a contradiction to $\gamma \cap \sigma_c = \emptyset$ and $c \notin \sigma_c$.

Turning to (ii), $\mathcal{K}_c \neq \emptyset$ by Lemma 26. Suppose now that $\gamma = \sigma \setminus \sigma_c$. Establishing the second part of (ii) for this particular γ will establish it for all hypothesized γ, by monotonicity of closure operators. By (i), $\overline{\gamma} \cup \{c\} \in \Delta_G$. Since σ is maximal in Δ_G, $\overline{\gamma} \subseteq \sigma$. Since $c \notin \sigma$ and by Lemma 2 on page 24, $\kappa \setminus (\overline{\gamma} \cup \{c\}) \neq \emptyset$ for every $\kappa \in \mathcal{K}_c$. Since $\sigma \setminus \sigma_c = \gamma \subseteq \overline{\gamma} \subseteq \sigma$ and $\kappa \subseteq \sigma \cup \{c\}$, $\kappa \setminus (\overline{\gamma} \cup \{c\}) = \big(\kappa \cap \sigma_c\big) \setminus \overline{\gamma}$, completing the proof. □

The following theorem has as corollary Theorem 4 of page 33 for pure nondeterministic graphs:

Theorem 28 (Informative Action Release Sequences : Pure Nondeterministic Graphs)**.**
Let $G = (V, \mathfrak{A})$ be a fully controllable pure nondeterministic graph with $n = |V| > 1$ and suppose $H = (V, \mathfrak{B})$ is a hierarchical cyclic subgraph of G.

Suppose σ is a maximal strategy in Δ_G. Set $\tau = \sigma \cap \mathfrak{B}$, then define H^* by step 3 of Construction 20 on page 42.

I. If H^* is a leaf, then σ contains an informative action release sequence for G of length at least $n - 1$.

II. If H^* is a node, then σ contains an informative action release sequence for G of length at least n.

Proof. Throughout the proof we assume notation as given in Construction 20 and Lemma 26. Observe that $\sigma \neq \emptyset$, since G is fully controllable with $n > 1$.

I. Suppose H^* is a leaf. Then $\tau_+ = \tau_- = \emptyset$. By Lemma 21 on page 43, $\tau_\circ \subseteq \sigma$ and $\tau_\circ$ is cycle-breaking in H; by Corollary 25 on page 47, $|\tau_\circ| = n - 1$; and by Lemma 16 on page 39, one may find an ordering of the actions in $\tau_\circ$ such that they form an informative action release sequence for H. This sequence is also informative for G by the comment after Lemma 7 on page 35.

II. Suppose H^* is a node. As in part I, one may find an ordering of the actions in $\tau_\circ \cup \tau_+$ such that they form an informative action release sequence for G ($\tau_\circ \cup \tau_+ \neq \emptyset$, by maximality of σ and full controllability of G). Write this sequence as $a_1, \ldots, a_\ell$. It is contained in σ.

By Corollary 25, $|\tau_\circ \cup \tau_+ \cup \tau_-| = n$. Since τ^* is disruptive in H^*, $\tau_- \neq \emptyset$. Of course, one cannot release the actions in τ_-, since they are not in σ. Instead, as we will see shortly, for each $c \in \tau_-$ one may release some action of σ_c informatively, thereby completing the proof.

First, observe that $\big(\tau_\circ \cup \tau_+\big) \cap \sigma_c = \emptyset$, for every $c \in \tau_-$. To see this, write $\eta = \tau_\circ \cup \tau_+$ and suppose $a \in \eta$ and $\mathrm{src}(a) \in J_c$ for some $c \in \tau_-$. Write $c = w \to T$. Let $\geq$ be the partial order induced on V by η. Then $t \geq \mathrm{src}(a)$ for some $t \in T$. Since $a \in \eta$, $\mathrm{src}(a) \geq s$ for every $s \in \mathrm{trg}(a)$. So $t \geq s$ for every $s \in \mathrm{trg}(a)$, meaning $\mathrm{trg}(a) \subseteq J_c$. Consequently, $a \notin \sigma_c$.

Inductively, suppose we have released, for some sequence of distinct actions $c_1, ..., c_k$ in τ_-, with $k \geq 0$, a corresponding sequence of distinct actions $b_1, ..., b_k$ in σ, such that $b_i \in \sigma_{c_i}$, for $i = 1, \ldots, k$, and such that the overall sequence $a_1, \ldots, a_\ell, b_1, \ldots, b_k$ is an iars for G. If $k = |\tau_-|$, we are done. Otherwise, we need to show how to extend this sequence.

Let $\gamma = \{a_1, \ldots, a_\ell, b_1, \ldots, b_k\}$ and $\overline{\gamma} = (\phi_A \circ \psi_A)(\gamma)$, with A being G's action relation. Pick some $c \in \tau_- \setminus \{c_1, \ldots, c_k\}$. We already observed that $\big(\tau_\circ \cup \tau_+\big) \cap \sigma_c = \emptyset$. By construction, $\mathrm{src}(b_i) \in J_{c_i}$, for $i = 1, \ldots, k$. By part (a) of Lemma 22 on page 43, $J_{c_i} \cap J_c = \emptyset$, for $i = 1, \ldots, k$. Consequently, $\gamma \cap \sigma_c = \emptyset$. By Corollary 27, there exist $\kappa \in \mathcal{K}_c$ and $b \in \big(\kappa \cap \sigma_c\big) \setminus \overline{\gamma}$, so b may be released informatively. Let $c_{k+1} = c$ and $b_{k+1} = b$. □

4.10 Examples for Pure Nondeterministic Graphs

This subsection shows how the proof of Theorem 28 produces informative action release sequences for various pure nondeterministic graphs and strategies.

4.10.1 A Hierarchical Pure Nondeterministic Graph

The first example considers the pure nondeterministic graph of Figure 15 on page 27. The graph may be viewed directly as a hierarchical cyclic graph, as indicated by Figure 21.

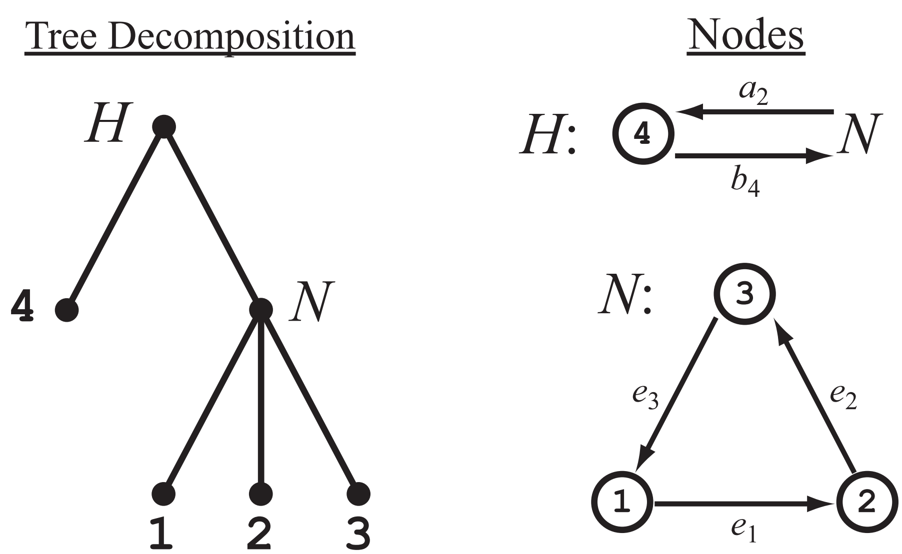

Figure 21: A view of graph G from Figure 15 directly as a hierarchical cyclic graph H. The left panel shows the tree decomposition of H. The right panel shows each node's constituent parts. The root H contains two children, a leaf modeling state #4 and a node N, along with two cycle actions, b_4 and a_2. Although action b_4 is nondeterministic with multiple targets inside node N, for simplicity the figure merely depicts an arrow pointing from state #4 to node N. Node N contains three leaves as children, modeling the set of states $\{1, 2, 3\}$, along with three cycle actions, e_1, e_2, and e_3.

Let us consider two maximal strategies and see how our constructions generate informative action release sequences using the hierarchical cyclic graph H. Since H and G have the same actions, τ in Construction 20 on page 42 is the maximal strategy under consideration.

$\underline{\tau = \{e_2, e_3, b_4\}}$ (This strategy converges to state #1.)

- Not used by the construction, but just for reference: τ is cycle-breaking in H.
- τ is not disruptive in H, so we run the loop of step 2 in Construction 20:
  1. First we mark node N, defining $\kappa^{(1)} = \{e_2, e_3\}$.
  2. Then we mark node H, defining $\kappa^{(2)} = \{e_2, e_3, b_4\}$.

- At step 3, H^* is a leaf. So $\tau_\circ = \kappa^{(2)} = \tau$ and $\tau_+ = \tau_- = \emptyset$.
- The proof of Lemma 16 on page 39 now produces either the sequence b_4, e_2, e_3 or the sequence b_4, e_3, e_2 as an informative action release sequence.

$\underline{\tau = \{e_1, e_2, a_2\}}$ (This strategy converges to the set of states $\{3, 4\}$.)

- τ is cycle-breaking in H.
- τ is not disruptive in H, so we run the loop of step 2 in the construction:
  1. First we mark node N, defining $\kappa^{(1)} = \{e_1, e_2\}$.
  2. Then we mark node H, defining $\kappa^{(2)} = \{e_1, e_2, a_2\}$.
- At step 3, H^* is a leaf. So $\tau_\circ$ is again all of τ and $\tau_+ = \tau_- = \emptyset$.
- Again, one may release the actions of $\tau_\circ$ informatively, as per the proof of Lemma 16, for instance as the sequence a_2, e_1, e_2.

Comments: (i) G's action relation in Figure 16 on page 28 shows that no maximal strategy is disruptive in H, so Construction 20 will always run the loop of step 2. (ii) The construction will always assemble the entire strategy as an iars. In fact, as Figure 16 shows, the strategy complex Δ_G is a triangulation of $\mathbb{S}^2$, and in particular has no free faces. Consequently, any ordering of the actions in a maximal strategy will be an informative action release sequence for G.

4.10.2 A Pure Nondeterministic Graph with Several Nondeterministic Actions

Let us add some nondeterministic actions to the previous graph, as shown in Figure 22.

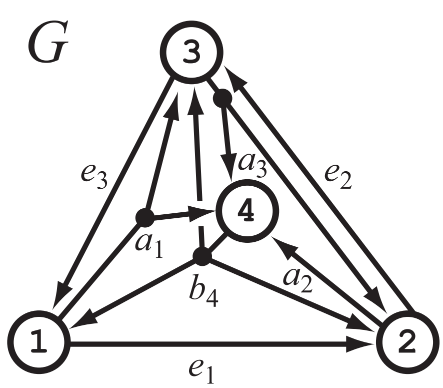

A	e_1	e_2	e_3	a_1	a_2	a_3	b_4	Goal
σ_1		•	•				•	1
σ_2	•		•				•	2
σ_3	•	•					•	3
σ_4	•		•		•	•		4
σ_5	•			•	•	•		4
σ_{14}		•	•		•			$\{1, 4\}$
σ_{34}	•	•		•	•			$\{3, 4\}$

Figure 22: **Left Panel:** A pure nondeterministic graph G with four states, $1, 2, 3, 4$, four deterministic actions, e_1, e_2, e_3, a_2, and three nondeterministic actions, a_1, a_3, b_4.
Right Panel: G's action relation and goal sets.
(This figure is a copy of Figures 47 and 48 in [6].)

The earlier hierarchical cyclic graph H of Figure 21 is a subgraph of the new G, on the same state space (but with fewer actions), so we can use the same H as before to construct informative action release sequences for maximal strategies, now in the new G. Almost every maximal strategy in the new Δ_G is either identical to or a proper superset of a maximal strategy in the old Δ_G. Intersecting one of these strategies with the actions of H, as Theorem 28 requires, therefore produces the same constructions as before.

There is one exception: The new Δ_G contains a maximal strategy, namely σ_5, that does not restrict to a maximal strategy in the old Δ_G. Let us look at that strategy more carefully:

$\underline{\sigma = \sigma_5 = \{e_1, a_1, a_2, a_3\}}$ (This strategy converges to state #4.)

- τ is the intersection of σ with the actions of H, so $\tau = \{e_1, a_2\}$.
- τ is cycle-breaking in H.
- Now τ *is* disruptive in H, so Construction 20 *does not* run the loop of step 2, but skips directly to step 3.
- At step 3, H^* is all of H, so $\tau_\circ = \emptyset$.
- $\mathfrak{C}$ consists of all the core cycle actions of H, so $\mathfrak{C} = \{e_1, e_2, e_3, a_2\}$.
- The construction of ξ in step 4 incorporates all of τ, starting from $\xi = \emptyset$, as follows:
  1. For node N, step 4 adds action e_1 to ξ.
  2. For node H, step 4 adds action a_2 to ξ.
- Thus $\tau_+ = \mathfrak{C} \cap \xi = \{e_1, a_2\}$ and $\tau_- = \mathfrak{C} \setminus \xi = \{e_2, e_3\}$.
- The actions of τ_+ may be released informatively in depth order, as the sequence a_2, e_1.
- For each action in τ_-, one finds an action in σ as per the proof of Theorem 28:
  1. For action $e_2 \in \tau_-$, action $a_3 \in \sigma$ lies "downstream" from e_2, forms a minimal nonface with e_2, and is not implied by $\{a_2, e_1\}$.
  2. For action $e_3 \in \tau_-$, action $a_1 \in \sigma$ lies "downstream" from e_3, forms a minimal nonface with e_3, and is not implied by $\{a_2, e_1, a_3\}$.

  (The term "downstream" refers to the partial order determined by $\eta = \tau_\circ \cup \tau_+$. Since $\tau_\circ = \emptyset$, that simply means τ_+ here. Specifically, the phrase "action b lies downstream from action a" means that b's source lies in the forward projection of a's targets under η, that is, "$t \geq_\eta \mathrm{src}(b)$, for some $t \in \mathrm{trg}(a)$".

  Moreover, this and subsequent examples, following the proof of Theorem 28, further choose b so that *not* all of b's targets lie within the forward projection of a's targets.)

Consequently, one may arrange all four actions of σ $(= \sigma_5)$ into an informative action release sequence for G. This is consistent with Theorem 28, since H^* is a node in the construction. For instance, the sequence a_2, e_1, a_3, a_1 is an iars. There are other orderings that will also produce iars of length 4, but not all will do so. For instance, releasing action a_1 as the first action in a sequence would limit the length of that sequence as an iars to 2. See [6] for further discussion of this example.

4.10.3 A Directed Graph with Several Cycles, Represented Hierarchically

Consider the directed graph G of Figure 23. All the actions in this graph are deterministic. The graph has several directed cycles in it, giving us the opportunity to explore more than one hierarchical decomposition for G. The figure also shows a maximal strategy σ in Δ_G. We will focus on this one strategy, using two different hierarchical cyclic subgraphs of G to construct informative action release sequences for G in two different ways, such that each sequence consists of actions contained in σ. For reference, G's full action relation appears in Figure 24.

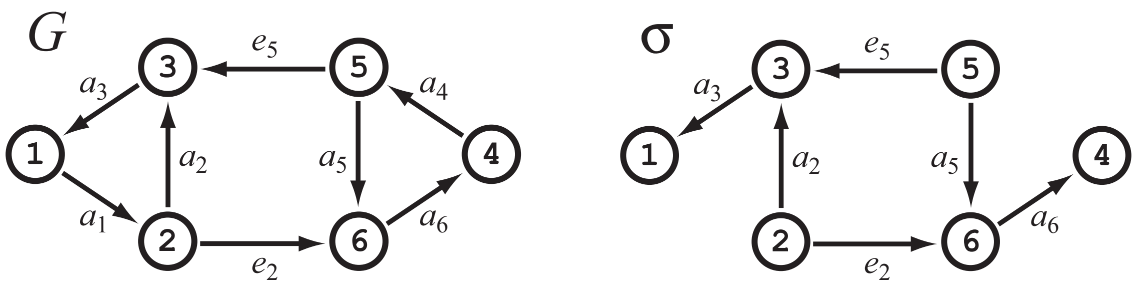

Figure 23: **Left Panel:** A directed graph G, consisting of six states and eight directed edges. **Right Panel:** A maximal strategy $\sigma \in \Delta_G$, depicted by its directed edges.

A	a_1	a_2	a_3	a_4	a_5	a_6	e_2	e_5	Goal
σ		•	•		•	•	•	•	$\{1,4\}$
	•		•		•	•	•	•	4
	•	•			•	•	•	•	$\{3,4\}$
		•	•	•		•	•	•	1
	•		•	•		•	•		5
	•		•	•		•		•	2
	•	•		•		•	•	•	3
		•	•	•	•		•	•	$\{1,6\}$
	•		•	•	•		•	•	6
	•	•		•	•		•	•	$\{3,6\}$

Figure 24: Action relation and goal sets for the graph of Figure 23. The row corresponding to maximal strategy σ is labeled. This strategy has a multi-state goal, namely $\{1,4\}$.

A Multi-Node Hierarchical Decomposition: The decomposition H shown in Figure 25 models G directly as a hierarchical cyclic graph, meaning H and G contain the same states and actions. In this decomposition, the smaller two cycles of G define two nodes. Each of these nodes contains only leaves, comprising the state spaces $\{1,2,3\}$ and $\{4,5,6\}$, respectively. The root of the tree has these two nodes as children, connected by a two-cycle.

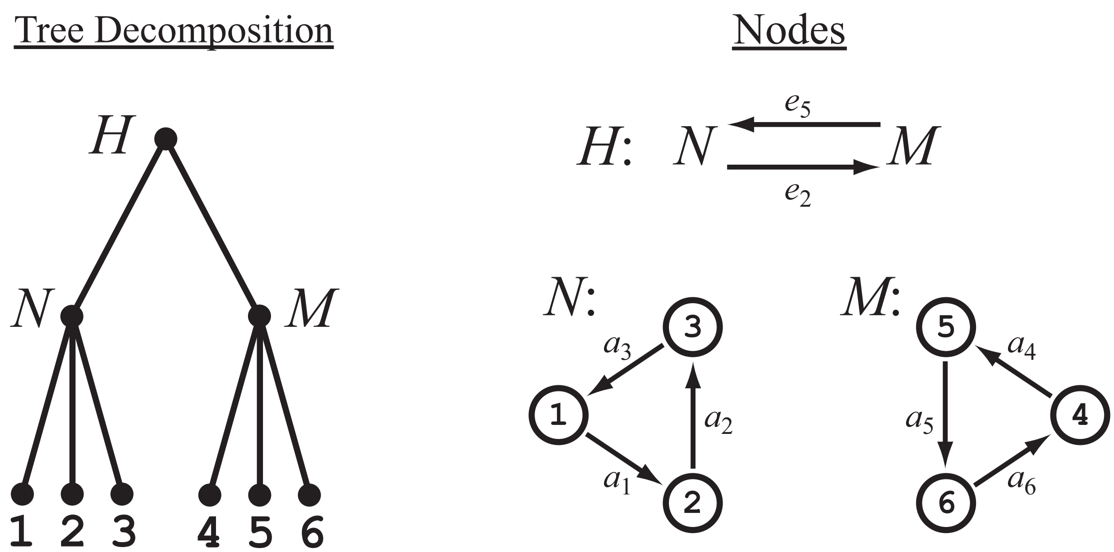

Figure 25: A view of graph G from Figure 23 directly as a hierarchical cyclic graph H. The left panel shows the tree decomposition of H. The right panel shows each node's constituent parts. The root of H contains two nodes as children, along with two cycle actions. The child nodes N and M each contain three leaves as children along with three cycle actions.

$\sigma = \{e_2, e_5, a_2, a_3, a_5, a_6\}$

- Since $H = G$, also $\tau = \sigma$.
- τ is not cycle-breaking in H, since it contains both cycle actions of H's root node.
- τ is not disruptive in H, so Construction 20 runs the loop of step 2. The construction may mark nodes N and M in either order. Here we start with N.
  1. Mark node N, defining $\kappa^{(1)} = \{a_2, a_3\}$.
  2. Mark node M, defining $\kappa^{(2)} = \{a_2, a_3, a_5, a_6\}$.
  3. Mark node H. Since τ contains both of H's cycle actions, the construction could add either action to $\kappa^{(2)}$ in defining $\kappa^{(3)}$. Here we add e_2, so $\kappa^{(3)} = \{a_2, a_3, a_5, a_6, e_2\}$.
- At step 3, H^* is a leaf. So $\tau_\circ = \{a_2, a_3, a_5, a_6, e_2\}$ and $\tau_+ = \tau_- = \emptyset$.
- One may release the actions of $\tau_\circ$ informatively in depth order, as per the proof of Lemma 16 on page 39, for instance as the sequence e_2, a_2, a_3, a_5, a_6.
- Observe that $\tau_\circ$ is almost all of σ, excluding only action e_5. The construction discarded that one action when forming $\kappa^{(3)}$.

A Flat Decomposition: Figure 26 shows another hierarchical cyclic subgraph H of G, on the same state space but with fewer actions. In this subgraph, the Hamiltonian cycle of G defines a single node, necessarily the root of H, with all six states as leaves. Two of G's actions do not appear in H.

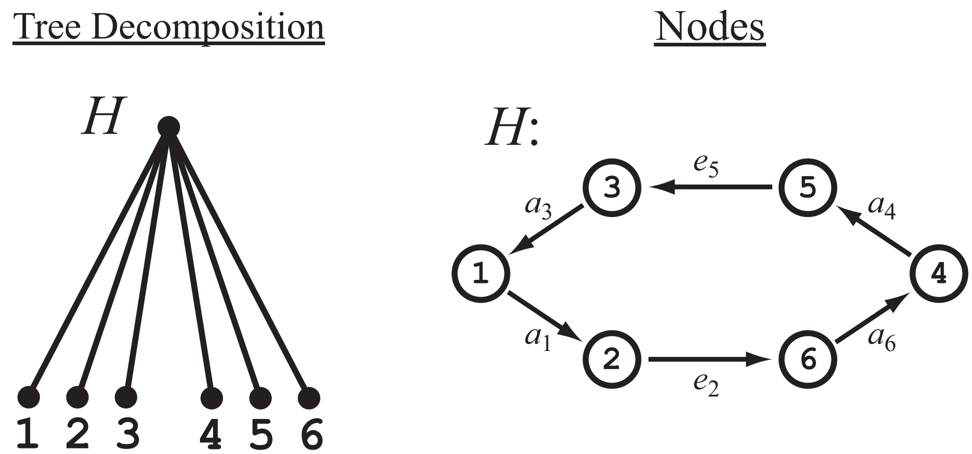

Figure 26: A hierarchical cyclic subgraph H of graph G from Figure 23, with the same states. The left panel shows the tree decomposition of H. The right panel shows constituent parts. The root of H contains six leaves, connected by six actions forming a directed cycle.

$\sigma = \{e_2, e_5, a_2, a_3, a_5, a_6\}$

- τ is the intersection of σ with H's actions, so $\tau = \{e_2, e_5, a_3, a_6\}$.
- τ is both cycle-breaking and disruptive in H.
- Since τ is disruptive, $H^* = H$ and $\tau_\circ = \emptyset$ in step 3 of Construction 20.
- $\mathfrak{C}$ consists of all the core cycle actions of H, which means all the actions of H since H defines a Hamiltonian cycle. So $\mathfrak{C} = \{a_1, e_2, a_6, a_4, e_5, a_3\}$.
- The construction of ξ incorporates all of τ, since H consists of a single unmarked node. Thus $\tau_+ = \mathfrak{C} \cap \xi = \tau = \{e_2, e_5, a_3, a_6\}$ and $\tau_- = \mathfrak{C} \setminus \xi = \{a_1, a_4\}$.
- The actions of τ_+ may be released informatively in any order, for instance as the sequence e_2, e_5, a_3, a_6.
- For each action in τ_-, one finds an action in σ as per the proof of Theorem 28:

  (Again, "downstream" refers to the partial order determined by τ_+.)

  1. For action $a_1 \in \tau_-$, action $a_2 \in \sigma$ lies "downstream" from a_1, participates in the minimal nonface $\{a_1, a_2, a_3\}$ with a_1, and is not implied by $\{e_2, e_5, a_3, a_6\}$.
  2. For action $a_4 \in \tau_-$, action $a_5 \in \sigma$ lies "downstream" from a_4, participates in the minimal nonface $\{a_4, a_5, a_6\}$ with a_4, and is not implied by $\{e_2, e_5, a_3, a_6, a_2\}$.
- Consequently, all actions of σ may be arranged into the informative action release sequence $e_2, e_5, a_3, a_6, a_2, a_5$.

Comment: We have seen the following: (i) With H as in Figure 25, Construction 20 produces an informative action release sequence for G consisting of 5 actions in σ. (ii) With H as in Figure 26, the construction produces an informative action release sequence consisting of all 6 actions in σ. These sequence lengths match the assertions of Theorem 28 on page 49.

4.10.4 A Directed Graph with a Disruptive but not Cycle-Breaking Strategy

This example will illustrate an instance in which τ contains all the cycle actions in a node during step 4 of Construction 20. Figure 27 depicts a graph G and a maximal strategy $\sigma \in \Delta_G$. Figure 28 displays G's action relation. Figure 29 shows a hierarchical cyclic subgraph H of G, on the same state space (but with fewer actions). (Other such subgraphs exist, of course.)

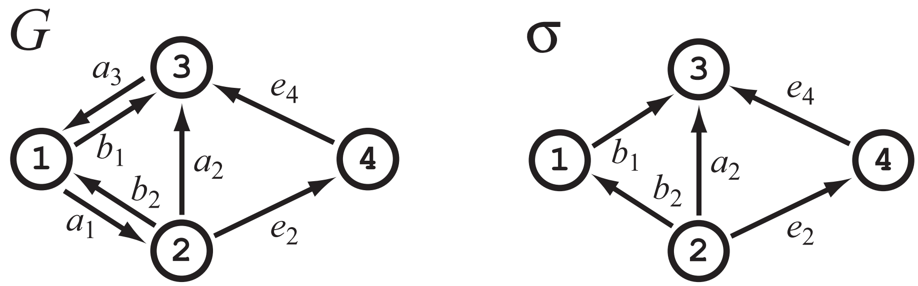

Figure 27: **Left Panel:** A directed graph G, consisting of four states and seven directed edges. **Right Panel:** A maximal strategy $\sigma \in \Delta_G$, depicted by its directed edges.

A	a_1	a_2	a_3	b_1	b_2	e_2	e_4	Goal
		•	•		•	•	•	1
	•		•			•		4
	•		•				•	2
	•	•		•		•	•	3
σ		•		•	•	•	•	3

Figure 28: Action relation and goals for the graph of Figure 27. The row corresponding to maximal strategy σ is labeled. This strategy converges to state #3.

$\sigma = \{e_2, e_4, a_2, b_1, b_2\}$

- τ is the intersection of σ with H's actions, so $\tau = \{e_2, e_4, a_2\}$.
- τ is not cycle-breaking in H, since it contains both cycle actions of H's root node.
- τ is disruptive in H, since it contains only one of node N's three cycle actions.
- Since τ is disruptive, $H^* = H$ and $\tau_\circ = \emptyset$ in step 3 of Construction 20.
- $\mathfrak{C}$ consists of all the core cycle actions of H, so $\mathfrak{C} = \{a_1, a_2, a_3, e_2\}$.

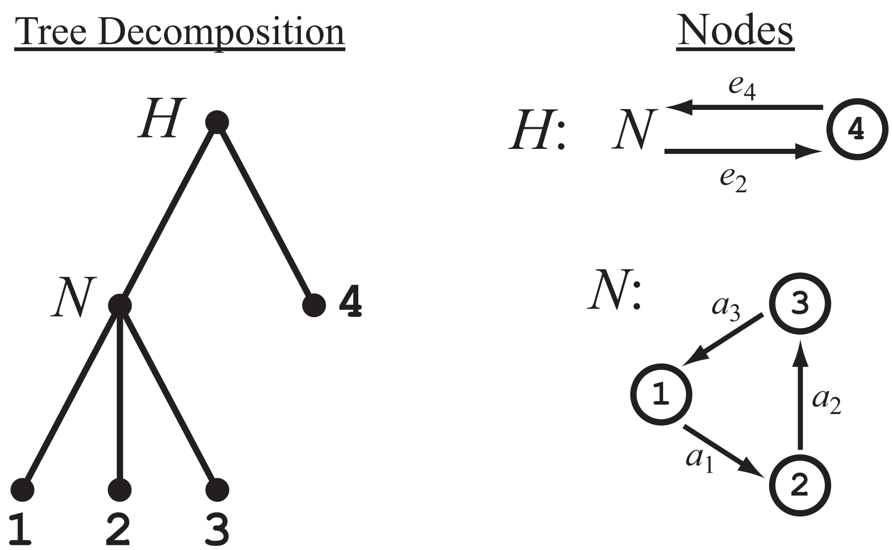

Figure 29: A hierarchical cyclic subgraph H of graph G from Figure 27, on the same same state space (but with fewer actions). The left panel shows the tree decomposition of H. The right panel shows each node's constituent parts. The structure is very similar to that of Figure 21.

- The construction of $\xi = \{a_2, e_2\}$ in step 4 occurs as follows, starting from $\xi = \emptyset$:
  1. For node N, $\mathcal{C}_N = \{a_1, a_2, a_3\}$ and $\mathcal{C}_N \cap \tau = \{a_2\}$, so one adds action a_2 to ξ.
  2. For node H, $\mathcal{C}_H = \{e_2, e_4\}$ and $\mathcal{C}_H \cap \tau = \mathcal{C}_H$, so one must discard some action of $\mathcal{C}_H$ that is not in $\mathfrak{C}$. That action is e_4. One adds action e_2 to ξ.
- Consequently, $\tau_+ = \mathfrak{C} \cap \xi = \{a_2, e_2\}$ and $\tau_- = \mathfrak{C} \setminus \xi = \{a_1, a_3\}$.
- The actions of τ_+ may be released informatively in depth order, so as the sequence e_2, a_2.
- For each action in τ_-, one finds an action in σ as per the proof of Theorem 28:

  (Once again, "downstream" refers to the partial order determined by τ_+.)
  1. For action $a_1 \in \tau_-$, action $b_2 \in \sigma$ lies "downstream" from a_1, forms a minimal nonface with a_1, and is not implied by $\{e_2, a_2\}$.
  2. For action $a_3 \in \tau_-$, action $b_1 \in \sigma$ lies "downstream" from a_3, forms a minimal nonface with a_3, and is not implied by $\{e_2, a_2, b_2\}$.
- Therefore e_2, a_2, b_2, b_1 is an informative action release sequence for G, contained in σ.

  (The order matters: Revealing either b_1 or b_2 at the beginning of the sequence would narrow the set of maximal strategies consistent with the revealed action to two instantly. One then could reveal only one more action informatively before identifying σ, that action being the other "b_i" action not yet revealed. Indeed, revealing actions b_1 *and* b_2 in either order identifies the maximal strategy to be σ. Revealing action a_2 implies action e_2, while revealing e_2 at the beginning does not imply a_2. Although not part of the construction, observe that revealing b_1 would in and of itself declare the goal to be state #3.)

5 The Stochastic Setting

The aim of this section is to prove Theorem 4 from page 33 for the case in which the graph G is pure stochastic. Throughout, this section assumes that all graphs are pure stochastic, meaning each action is either deterministic or stochastic (but not nondeterministic).

Caution: Even though all actions in a pure stochastic graph are deterministic or stochastic, there may still be a component of nondeterminism in a strategy: When multiple actions have the same source state, any one of those actions might execute from that state, with the choice potentially made by an adversary. (See again the discussion of generalized control laws on page 19, as well as the definitions of "moves off", "contains a circuit", and "strategy complex".)

5.1 Expanding Fully Controllable Subgraphs via Minimal Nonfaces

As mentioned on page 25, a minimal nonface of Δ_G in a pure stochastic graph G defines an irreducible Markov chain and thus a fully controllable subgraph of G. If G is itself fully controllable, one may construct such a minimal nonface κ for each maximal strategy $\sigma \in \Delta_G$, for instance by considering some action at a goal state of σ. The actions of the minimal nonface κ that lie within σ then form an informative action release sequence z for G, contained in σ.

One may expand the state space $\mathrm{src}(\kappa)$ covered by this minimal nonface by considering some action outside σ, in a manner to be discussed. This process yields a new minimal nonface and thus additional actions of σ with which to enlarge the informative action release sequence z. Repeating this process one may eventually encounter a situation in which there are no further useful actions outside σ. Instead, one forms a quotient graph by identifying all the states covered thus far. Recursively, one obtains an informative action release sequence within this quotient graph. Patching the two sequences together gives an overall informative action release sequence contained in σ of length one less than the number of states in G's state space.

The following construction and subsequent results describe this process formally:

Construction 29 (Minimal Nonface Expansion)**.** *Let $G = (V, \mathfrak{A})$ be a fully controllable pure stochastic graph with $n = |V| > 1$ and suppose σ is a maximal strategy in Δ_G.*

Construct a collection $\{b_1, \ldots, b_k\}$ of convergent actions in $\mathfrak{A}$, a collection $\{\kappa_1, \ldots, \kappa_k\}$ of minimal nonfaces of Δ_G, a collection $\{\mathcal{A}_1, \ldots, \mathcal{A}_k\}$ of subsets of $\mathfrak{A}$, and a collection $\{W_1, \ldots, W_k\}$ of subsets of V, with $k \geq 1$, as follows:

1. Let $g \in V \setminus \mathrm{src}(\sigma)$. Choose $b_1 \in \mathfrak{A}$ so that $\{b_1\} \in \Delta_G$ and $\mathrm{src}(b_1) = g$. Such an action exists since G is fully controllable and $|V| > 1$.

2. Since σ is maximal, $\sigma \cup \{b_1\} \notin \Delta_G$, so there exists a minimal nonface κ_1 of Δ_G such that $b_1 \in \kappa_1 \subseteq \sigma \cup \{b_1\}$. (For later reference, observe also that $|\kappa_1| > 1$.)

3. Let $\mathcal{A}_1 = \kappa_1$ and $W_1 = \mathrm{src}(\kappa_1)$.

4. Set Done to `false`. While not Done, run the following loop, starting from $i = 1$:
   (a) Consider the quotient graph G/W_i and let prime notation refer to the correspondence between actions in G and G/W_i, as per the discussion on page 20.

(b) Define $\xi_i = \{a \in \sigma \mid \mathrm{src}(a) \in V \setminus W_i\}$. So $\xi_i \subseteq \sigma$, $\xi_i \in \Delta_G$, and $\mathrm{src}(\xi_i) \cap W_i = \emptyset$.

(c) By Fact 1 on page 20, $\xi_i' \in \Delta_{G/W_i}$, so extend ξ_i' to a maximal simplex $\tau_i' \in \Delta_{G/W_i}$.

(d) If $\tau_i \subseteq \sigma$, then set k to the current value of i and Done to `true`. The loop ends.

Otherwise:
- Let $b_{i+1} \in \tau_i \setminus \sigma$.
- As in step 2, there exists a minimal nonface κ_{i+1} of Δ_G such that $b_{i+1} \in \kappa_{i+1} \subseteq \sigma \cup \{b_{i+1}\}$. (Again, $|\kappa_{i+1}| > 1$.)
- Let $\mathcal{A}_{i+1} = \mathcal{A}_i \cup \kappa_{i+1}$ and $W_{i+1} = W_i \cup \mathrm{src}(\kappa_{i+1})$.
- The loop continues, with $i + 1$ in place of i.

Lemma 30 (Expansive Subspaces)**.** *Let the hypotheses and notation be as in Construction 29. Then $W_i \subsetneq W_{i+1}$, for all i such that W_i and W_{i+1} are well-defined.*
(Consequently, the loop of step 4 in the construction ends, that is, k is well-defined finite.)

Proof. It is enough to show that $\mathrm{src}(\kappa_{i+1}) \not\subseteq W_i$. Suppose otherwise. Since $b_{i+1} \in \kappa_{i+1}$, that would mean $\mathrm{src}(b_{i+1}) \in W_i$ and $\mathrm{trg}(b_{i+1}) \subseteq W_i$. (The inclusion holds because κ_{i+1} is a minimal nonface in Δ_G, so no action of κ_{i+1} moves off $\mathrm{src}(\kappa_{i+1})$, and because G is pure stochastic.) Thus b_{i+1}' would become self-looping in G/W_i, contradicting $b_{i+1}' \in \tau_i' \in \Delta_{G/W_i}$. □

Lemma 31 (Fully Controllable Expansion)**.** *Let the hypotheses and notation be as in Construction 29. Then $(W_i, \mathcal{A}_i)$ is a fully controllable pure stochastic graph, for $i = 1, \ldots, k$.*

Proof. Observe that $k \geq 1$, since $|V| > 1$.

Base Case: $(W_1, \mathcal{A}_1) = (\mathrm{src}(\kappa_1), \kappa_1)$. Since κ_1 is a minimal nonface in the strategy complex of a pure stochastic graph, $(\mathrm{src}(\kappa_1), \kappa_1)$ is a fully controllable pure stochastic graph.

Inductive Step: As in the base case, $(\mathrm{src}(\kappa_{i+1}), \kappa_{i+1})$ is a fully controllable pure stochastic graph. Inductively, $(W_i, \mathcal{A}_i)$ is a fully controllable pure stochastic graph. Showing that $W_i \cap \mathrm{src}(\kappa_{i+1}) \neq \emptyset$ would therefore establish full controllability of the pure stochastic graph $(W_{i+1}, \mathcal{A}_{i+1})$. Suppose this intersection is empty. Then κ_{i+1}' is a minimal nonface in Δ_{G/W_i}. On the other hand, $\kappa_{i+1} \setminus \{b_{i+1}\} \subseteq \xi_i$, so $\kappa_{i+1}' \subseteq \tau_i' \in \Delta_{G/W_i}$, producing a contradiction. □

Lemma 32 (Distinct Actions)**.** *Let the hypotheses and notation be as in Construction 29. Then $|\{b_1, \ldots, b_k\}| = k$, that is, the actions $b_1, \ldots, b_k$ are distinct.*

Proof. Suppose $1 \leq j \leq i < k$. Then $\mathrm{src}(b_j) \in W_j$ and $\mathrm{trg}(b_j) \subseteq W_j$. Since $W_j \subseteq W_i$, action b_j' is self-looping in G/W_i and thus b_j cannot be a candidate for b_{i+1}. □

Lemma 33 (Expansive Sets of Actions)**.** *Let hypotheses and notation be as in Construction 29. Suppose $1 < i \leq k$. Let $\ell_i = |W_i \setminus W_{i-1}|$. (By Lemma 30, $\ell_i > 0$.)*

Then there exist actions $\mathcal{E}_i \subseteq \kappa_i \setminus (\mathcal{A}_{i-1} \cup \{b_i\})$ such that $|\mathcal{E}_i| = \ell_i$ and at most one action in $\mathcal{E}_i$ has its source in W_{i-1}. (We refer to $\mathcal{E}_i$ as an expansive set *of actions.)*

Moreover, suppose for all $\mathcal{E} \subseteq \kappa_i \setminus (\mathcal{A}_{i-1} \cup \{b_i\})$ with $|\mathcal{E}| = \ell_i$, $\mathrm{src}(\mathcal{E}) \cap W_{i-1} \neq \emptyset$. Then $\mathrm{src}(b_i) \notin W_{i-1}$ and one may choose $\mathcal{E}_i$ to contain an action e such that $\mathrm{src}(e) \in W_{i-1}$ and such that the probability of reaching $\mathrm{src}(b_i)$ from $\mathrm{src}(e)$ under actions of $\mathcal{E}_i$ is nonzero.

Comments: (a) Let $\mathcal{A}_0 = \emptyset$, $W_0 = \{\mathrm{src}(b_1)\}$, and $\mathcal{E}_1 = \kappa_1 \backslash \{b_1\}$. Then the lemma holds for $i = 1$, with $\mathrm{src}(\mathcal{E}_1) \cap W_0 = \emptyset$. (b) For $i = 1, \ldots, k$, $\mathcal{E}_i \subseteq \mathcal{A}_i \cap \sigma$, since $\kappa_i \subseteq \mathcal{A}_i$ and $\kappa_i \backslash \{b_i\} \subseteq \sigma$.

Proof. Assume $1 < i \leq k$. Readily, $W_i \setminus W_{i-1} = \mathrm{src}(\kappa_i) \setminus W_{i-1}$ and $W_{i-1} = \mathrm{src}(\mathcal{A}_{i-1})$. Thus $W_i \setminus W_{i-1} = \mathrm{src}(\kappa_i) \setminus \mathrm{src}(\mathcal{A}_{i-1}) \subseteq \mathrm{src}\left(\kappa_i \setminus \mathcal{A}_{i-1}\right)$, meaning each state in $W_i \setminus W_{i-1}$ is the source of some action in κ_i that is not also an action in $\mathcal{A}_{i-1}$. If in fact every state in $W_i \setminus W_{i-1}$ is the source of some action in κ_i that is neither an action in $\mathcal{A}_{i-1}$ nor the action b_i, then we may construct $\mathcal{E}_i \subseteq \kappa_i \backslash \left(\mathcal{A}_{i-1} \cup \{b_i\}\right)$ such that $|\mathcal{E}_i| = \ell_i$ and $\mathrm{src}(\mathcal{E}_i) \cap W_{i-1} = \emptyset$.

Otherwise, since all actions in a minimal nonface have distinct sources, it is only possible to find $\ell_i - 1$ actions in $\kappa_i \setminus \left(\mathcal{A}_{i-1} \cup \{b_i\}\right)$ whose sources lie outside W_{i-1}. Moreover, $\mathrm{src}(b_i) \in W_i \setminus W_{i-1}$. By the proof of Lemma 31, $\mathrm{src}(\kappa_i) \cap W_{i-1} \neq \emptyset$, meaning κ_i contains at least one action with source in W_{i-1}. We now show by backchaining from $\mathrm{src}(b_i)$ how to select one such action e so that $\mathcal{E}_i$ may consist of action e and the $\ell_i - 1$ actions just mentioned.

To reduce index clutter, we fix i and make the following definitions for the rest of the proof:

$$\begin{aligned}
\mathcal{A} &= \mathcal{A}_{i-1} \quad \text{and} \quad W = \mathrm{src}(\mathcal{A}), \\
b &= b_i, \\
\kappa &= \kappa_i, \\
\mathcal{E}^- &= \left\{ a \in \kappa \backslash \left(\mathcal{A} \cup \{b\}\right) \mid \mathrm{src}(a) \in \mathrm{src}(\kappa) \backslash W \right\}.
\end{aligned}$$

(By assumption for this case, $|\mathcal{E}^-| = \ell_i - 1$ and $\mathrm{src}(b) \in \mathrm{src}(\kappa) \backslash W$.)

We now define a backchaining algorithm, with a loop index j, for constructing sets of actions $\emptyset \neq \tau^{(0)} \subsetneq \cdots \subsetneq \tau^{(j)} \subsetneq \cdots$. Inductively, each iteration assumes that (i) $b \in \tau^{(j)} \subsetneq \kappa$, (ii) $\tau^{(j)} \subseteq \mathcal{E}^- \cup \{b\}$, and (iii) for each $s \in \mathrm{src}(\tau^{(j)})$, there exists a sequence of zero or more action edges leading from s to $\mathrm{src}(b)$, with the edges coming from actions in $\tau^{(j)} \setminus \{b\}$.

We initialize the loop with $\tau^{(0)} = \{b\}$. The loop will end by defining an action e such that we may let $\mathcal{E}_i = \mathcal{E}^- \cup \{e\}$, establishing the lemma. The loop starts from $j = 0$:

(a) Since κ is a minimal nonface in Δ_G, with G pure stochastic, $(\mathrm{src}(\kappa), \kappa)$ is a fully controllable graph in its own right and $\emptyset \neq \mathrm{src}(\tau^{(j)}) \subsetneq \mathrm{src}(\kappa)$, by Lemma 3 on page 25. Thus some action $a^{(j)} \in \kappa$ moves off $\mathrm{src}(\kappa) \setminus \mathrm{src}(\tau^{(j)})$ in this graph.

(b) If we can pick $a^{(j)}$ so that $\mathrm{src}(a^{(j)}) \in W$, then we do so and in that case we let $e = a^{(j)}$. Either way, we define $\tau^{(j+1)} = \tau^{(j)} \cup \{a^{(j)}\}$. Condition (iii) above is satisfied by $\tau^{(j+1)}$ since it is satisfied by $\tau^{(j)}$ and $\mathrm{trg}(a^{(j)}) \cap \mathrm{src}(\tau^{(j)}) \neq \emptyset$.

(c) If step (b) defined action e, then the loop ends.

Otherwise, necessarily $a^{(j)} \in \mathcal{E}^- \setminus \tau^{(j)}$. Thus, in this case, $\tau^{(j+1)}$ also satisfies conditions (ii) and (i) above, since in particular some action of κ has source in W but no action of $\tau^{(j+1)}$ does. The loop continues, with $j + 1$ in place of j.

By finiteness, the loop must eventually end, for some j. The probability of reaching $\mathrm{src}(b)$ from $\mathrm{src}(e)$ under actions of $\tau^{(j+1)} \backslash \{b\}$ is nonzero by condition (iii), so the same will be true under actions of $\mathcal{E}_i = \mathcal{E}^- \cup \{e\} \subseteq \kappa$. Moreover, $e \in \kappa \backslash \left(\mathcal{A} \cup \{b\}\right)$ with $\mathrm{src}(e) \in W$, since $e \in \kappa$ and $\emptyset \neq \mathrm{trg}(e) \cap \mathrm{src}(\tau^{(j)}) \subseteq \mathrm{src}(\kappa) \backslash W$, whereas $\mathrm{trg}(a) \subseteq W$, for all $a \in \mathcal{A}$, and $\mathrm{src}(b) \notin W$. $\square$

For the remainder of Section 5: Assume the hypotheses and notation of Construction 29 starting on page 58. Let $\mathcal{A}_0 = \emptyset$ and $W_0 = \{\mathrm{src}(b_1)\}$. Define $H_i = (W_i, \mathcal{A}_i)$, for $i = 0, 1, \ldots, k$, with $k \geq 1$. Each H_i is a fully controllable pure stochastic graph, by Lemma 31. Also, for $i = 1, \ldots, k$, H_{i-1} is a subgraph of H_i, with $\emptyset \neq W_{i-1} \subsetneq W_i$ and $\mathcal{A}_{i-1} \subsetneq \mathcal{A}_i$, by Lemmas 30 and 32, and since $|\kappa_1| > 1$. Let $\mathcal{E}_1 = \kappa_1 \setminus \{b_1\}$. For $i = 2, \ldots, k$, define $\mathcal{E}_i$ via Lemma 33, choosing $\mathcal{E}_i$ so that $\mathrm{src}(\mathcal{E}_i) \cap W_{i-1} = \emptyset$ whenever possible.

Corollary 34 (Expansion Independence)**.** *Let hypotheses and notation be as above. Suppose* $1 \leq i \leq k$. *Then* $\mathcal{E}_i \cup \tau \in \Delta_{H_i}$, *for every* $\tau \in \Delta_{H_{i-1}}$.

Proof. If $\mathrm{src}(\mathcal{E}_i) \subseteq W_i \setminus W_{i-1}$, then the lemma's assertion follows from Lemma 7.3(b)(i) in [5].

Otherwise, $i > 1$. Let $e \in \mathcal{E}_i$ be as per Lemma 33. There exists a sequence of action edges

$$\mathrm{src}(e) = v_1 \xrightarrow{a_1 = e} v_2 \xrightarrow{a_2} \cdots\ v_m \xrightarrow{a_m} v_{m+1} = \mathrm{src}(b_i),$$

for some $m \geq 1$, with $a_j \in \mathcal{E}_i$, $v_j = \mathrm{src}(a_j)$, and $v_{j+1} \in \mathrm{trg}(a_j)$, for all $j = 1, \ldots, m$. Moreover, $v_1 \in W_{i-1}$, $v_j \notin W_{i-1}$, for $j = 2, \ldots, m$, and $v_{m+1} = \mathrm{src}(b_i) \notin W_{i-1} \cup \mathrm{src}(\mathcal{E}_i)$.

Suppose $\mathcal{E}_i \cup \tau \notin \Delta_{H_i}$, for some $\tau \in \Delta_{H_{i-1}}$. Let ϱ be a minimal nonface of Δ_{H_i}, with $\emptyset \neq \varrho \subseteq \mathcal{E}_i \cup \tau$. Some or all of the actions $\{a_1, \ldots, a_m\}$ lie in ϱ. Certainly $e \in \varrho$, again by Lemma 7.3(b)(i) in [5]. Since ϱ is a minimal nonface, e is the only action of ϱ with source $\mathrm{src}(e)$. Since $\tau \in \Delta_{H_{i-1}}$, $\mathrm{src}(\varrho \setminus \mathcal{E}_i) \subseteq W_{i-1}$. The actions in $\mathcal{E}_i$ all have distinct sources. Thus no action in ϱ other than a_j (if a_j is even in ϱ) can have source v_j, for $j = 2, \ldots, m$.

Consequently, there is a nonzero probability that the system will transition to and stop at a state outside $\mathrm{src}(\varrho)$ when started at $\mathrm{src}(e)$, while moving under actions of ϱ. Some action of ϱ therefore moves off $\mathrm{src}(\varrho)$, which is a contradiction. □

Corollary 35 (Cardinality of Expansive Actions)**.** *Let hypotheses and notation be as above. Then* $\left|\bigcup_{i=1}^{k} \mathcal{E}_i\right| = |W_k| - 1$.

Proof. By Construction 29, Lemma 33, and subsequent comments,

$$\left|\bigcup_{i=1}^{k} \mathcal{E}_i\right| \;=\; \sum_{i=1}^{k} |\mathcal{E}_i| \;=\; \sum_{i=1}^{k} \ell_i \;=\; \sum_{i=1}^{k} |W_i \setminus W_{i-1}| \;=\; \sum_{i=1}^{k} \left(|W_i| - |W_{i-1}|\right) \;=\; |W_k| - 1.$$

□

5.2 Informative Action Release Sequences from Expansive Sets of Actions

This subsection shows how the constructions of the previous subsection produce informative action release sequences. Some notational abbreviations will be useful:

Notation and Terminology:

1. Rather than merely write sequences of actions, $b_1, \ldots, b_m$, we may write sequences of sets of actions $\mathcal{B}_1, \ldots, \mathcal{B}_m$, assuming the sets $\mathcal{B}_1, \ldots, \mathcal{B}_m$ are nonempty and pairwise disjoint.

   The meaning of a set $\mathcal{B}_i$ of actions is to indicate a multiplicity of sequences of actions, one for each possible permutation of the actions in the set $\mathcal{B}_i$. The sequence of sets $\mathcal{B}_1, \ldots, \mathcal{B}_m$ represents all possible orderings of the actions $\cup_{i=1}^{m} \mathcal{B}_i$ consistent with the top-level ordering $\mathcal{B}_1, \ldots, \mathcal{B}_m$. Here, "consistent" means actions in $\mathcal{B}_i$ must appear before actions in $\mathcal{B}_j$ whenever $i < j$, but the ordering is otherwise unconstrained.

For example, the sequence of sets $\{a, b\}, \{c\}, \{d, e, f\}$ represents 12 sequences of actions:

$$\begin{array}{cccccc} a,b,c,d,e,f & a,b,c,e,f,d & a,b,c,f,d,e & a,b,c,f,e,d & a,b,c,e,d,f & a,b,c,d,f,e \\ b,a,c,d,e,f & b,a,c,e,f,d & b,a,c,f,d,e & b,a,c,f,e,d & b,a,c,e,d,f & b,a,c,d,f,e \end{array}$$

2. We say that a sequence $\mathcal{B}_1, \ldots, \mathcal{B}_m$ of sets of actions is *informative for G* if each of the sequences of actions it represents is an informative action release sequence for graph G.

3. In place of a singleton set, we may also simply write the action it contains. For instance, we could write the top-level sequence in the example above as $\{a, b\}, c, \{d, e, f\}$.

Lemma 36 (Expanding Informative Actions). *Suppose $G = (V, \mathfrak{A})$ is a fully controllable pure stochastic graph with $n = |V| > 1$. Let σ be a maximal strategy in Δ_G.*

From G and σ construct $H_1, \ldots, H_k$ and $\mathcal{E}_1, \ldots, \mathcal{E}_k$, with $k \geq 1$, as per Construction 29 on page 58, Lemma 33 on page 59, and the definitions and notation of page 61.

Then, for each i, with $1 \leq i \leq k$, the sequence $\mathcal{E}_i, \mathcal{E}_{i-1}, \ldots, \mathcal{E}_1$ is informative for H_i.

Proof. By induction on i. Let i, with $1 \leq i \leq k$, be given.

The set $\mathcal{E}_i$ is a nonempty proper subset of a minimal nonface of Δ_G and thus of Δ_{H_i}. By Lemma 1 on page 24, every ordering of actions in $\mathcal{E}_i$ is an informative action release sequence for H_i. Moreover, $\mathcal{E}_i \in \Delta_{H_i}$.

If $i = 1$, these observations establish the base case.

If $i > 1$, then inductively $\mathcal{E}_{i-1}, \ldots, \mathcal{E}_1$ is informative for H_{i-1} and $\mathcal{E}_{i-1} \cup \cdots \cup \mathcal{E}_1 \in \Delta_{H_{i-1}}$.

By Corollary 34 on page 61, $\mathcal{E}_i \cup \tau \in \Delta_{H_i}$, for every $\tau \in \Delta_{H_{i-1}}$. By construction, no action in $\mathcal{E}_i$ is an action in the graph H_{i-1}. Therefore, by Lemma 7 on page 35, $\mathcal{E}_i, \mathcal{E}_{i-1}, \ldots, \mathcal{E}_1$ is informative for H_i. Moreover, $\mathcal{E}_i \cup \cdots \cup \mathcal{E}_1 \in \Delta_{H_i}$, since, for instance, $\mathcal{E}_i \cup \cdots \cup \mathcal{E}_1 \subseteq \sigma$. □

Corollary 37 (Expanding Informative Actions in G). *Let the hypotheses and notation be as for Lemma 36. For each i with $1 \leq i \leq k$, the sequence $\mathcal{E}_i, \mathcal{E}_{i-1}, \ldots, \mathcal{E}_1$ is informative for G.*

Proof. By the previous lemma, $\mathcal{E}_i, \mathcal{E}_{i-1}, \ldots, \mathcal{E}_1$ is informative for H_i. The comment after the statement of Lemma 7 on page 35 establishes the corollary. □

5.3 An Informative Action Release Sequence from a Quotient

The loop in Construction 29 may end in step 4(d) (on page 59) with $\tau_k = \emptyset$. This will occur if and only if $W_k = V$. In that case, Corollary 37 (above) and Corollary 35 (on page 61) imply that the sequence $\mathcal{E}_k, \ldots, \mathcal{E}_1$ provides an informative action release sequence for G of length $n - 1$, with all actions of the sequence contained in σ, and with $n = |V| > 1$.

Otherwise, the following lemma ensures that one may add a prefix of informative actions to that sequence whenever one can find an informative sequence in the quotient graph G/W_k.

Lemma 38 (Informative Actions from Quotient)**.** *Suppose $G = (V, \mathfrak{A})$ is a fully controllable pure stochastic graph with $n = |V| > 1$. Let σ be a maximal strategy in Δ_G.*

From G and σ construct k, W_k, τ_k, H_k, and $\mathcal{E}_1, \ldots, \mathcal{E}_k$, as per Construction 29 on page 58, Lemma 33 on page 59, and the definitions and notation of page 61. (Recall that $k \geq 1$.)

Suppose further that $a'_1, \ldots, a'_\ell$ is an informative action release sequence for G/W_k, with $\ell \geq 1$ and $\{a'_1, \ldots, a'_\ell\} \subseteq \tau'_k \in \Delta_{G/W_k}$.

Then $a_1, \ldots, a_\ell, \mathcal{E}_k, \ldots, \mathcal{E}_1$ is informative for G, with all actions contained in σ.

Proof. By construction, $\{a_1, \ldots, a_\ell\} \subseteq \tau_k \subseteq \sigma$ and $\cup_{i=1}^{k} \mathcal{E}_i \subseteq \sigma$.

By Lemma 31 on page 59, H_k is a fully controllable subgraph of G. Also, $\emptyset \neq W_k \subsetneq V$.

By Lemma 36 on page 62, $\mathcal{E}_k, \ldots, \mathcal{E}_1$ is informative for H_k. Any informative sequence of actions formed from $\mathcal{E}_k, \ldots, \mathcal{E}_1$ is a subset of σ, therefore convergent in both G and H_k.

Corollary 8 on page 35 therefore establishes the desired result. □

The following theorem instantiates Theorem 4 of page 33 for pure stochastic graphs:

Theorem 39 (Informative Action Release Sequences : Pure Stochastic Graphs)**.**

Let $G = (V, \mathfrak{A})$ be a fully controllable pure stochastic graph with $n = |V| > 1$.

Suppose σ is a maximal strategy in Δ_G.

Then σ contains an informative action release sequence for G of length at least $n - 1$.

Proof. By induction on n.

Base Case: $n = 2$.

In this case, Δ_G consists of two (nonempty) maximal strategies, one for each state in V. (The strategy for state v consists of all actions with source v that are not deterministic self-loops. The strategy converges to the other state.) Any single action in one of these strategies constitutes an informative action release sequence for G and is contained in the given strategy.

Inductive Step: $n > 2$.

From G and σ construct k, W_k, τ_k, and $\mathcal{E}_1, \ldots, \mathcal{E}_k$, using Construction 29 on page 58, Lemma 33 on page 59, and subsequent comments. Recall that $k \geq 1$.

As discussed on page 62, if $\tau_k = \emptyset$, then the sequence $\mathcal{E}_k, \ldots, \mathcal{E}_1$ provides an informative action release sequence for G of length $n - 1$, consisting of actions in σ.

Otherwise, let $\ell = |V \setminus W_k|$. Then $\ell > 0$. The quotient graph G/W_k is pure stochastic and fully controllable, by Fact 3 on page 20. It has state space $V' = (V \setminus W_k) \cup \{\diamond\}$, with $\diamond$ representing all of W_k identified to a single state.

Since the minimal nonface κ_1 in Construction 29 contains at least two actions, W_k contains at least two states. Therefore $2 \leq |V'| < n$. Inductively, the theorem holds for graph G/W_k and maximal strategy τ'_k, producing an informative action release sequence $a'_1, \ldots, a'_\ell$ for G/W_k with $\{a'_1, \ldots, a'_\ell\} \subseteq \tau'_k \in \Delta_{G/W_k}$. By Lemma 38, $a_1, \ldots, a_\ell, \mathcal{E}_k, \ldots, \mathcal{E}_1$ is informative for G, with all actions contained in σ. Any consequent informative action release sequence has length $\ell + |\cup_{i=1}^{k} \mathcal{E}_i| = |V \setminus W_k| + (|W_k| - 1) = n - 1$, by Corollary 35 on page 61. □

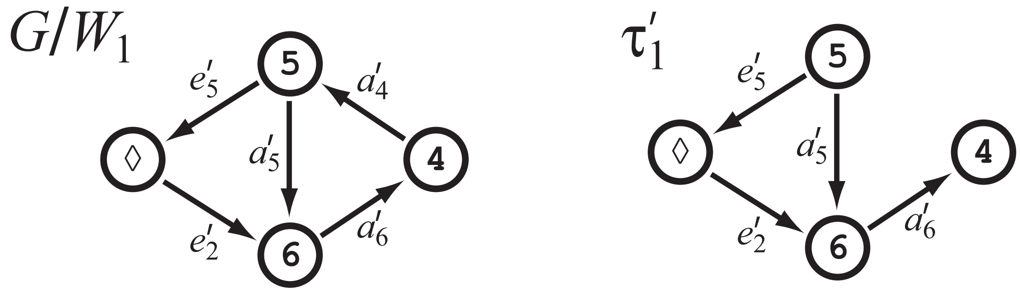

Figure 30: **Left Panel:** The quotient graph G/W_1, with G as in Figure 23 and $W_1 = \{1, 2, 3\}$. **Right Panel:** The maximal strategy τ_1', obtained in step 4(c) on page 59 during the first iteration of the loop of Construction 29, as applied to the graph G and strategy σ of Figure 23.

5.4 Examples for Pure Stochastic Graphs

This subsection shows how the proof of Theorem 39 produces informative action release sequences for some pure stochastic graphs and strategies. For clarity, figures discard self-loops.

5.4.1 A Directed Graph with Several Cycles

Consider again the strongly connected directed graph G and maximal strategy σ of Figure 23 on page 53. Earlier, we viewed G as a pure nondeterministic graph with different hierarchical cyclic subgraphs. Now, we view G as a pure stochastic graph and apply Construction 29 to obtain an informative action release sequence of length 5 for G, contained in σ.

For this example, it turns out that the loop of step 4 in the construction runs once, ending with $k = 1$, but without having covered the entire state space of the graph. Consequently, as indicated by Theorem 39's inductive proof, one needs to invoke the construction again, on a quotient graph. Again, the loop runs only once. In total, there are three invocations of the construction. The synopses below show how local variables in the construction are instantiated.

1. In the first invocation of Construction 29, the graph is G as in Figure 23 and the maximal strategy is $\sigma = \{e_2, e_5, a_2, a_3, a_5, a_6\}$. For g, one may use either state in σ's goal set $\{1, 4\}$. Using $g = 1$, one finds $b_1 = a_1$, yielding the minimal nonface $\kappa_1 = \{a_1, a_2, a_3\}$. Thus $W_1 = \{1, 2, 3\}$. The comments at the top of page 60 produce $\mathcal{E}_1 = \{a_2, a_3\}$.

   Running the loop of step 4 in the construction, with $i = 1$, one obtains $\xi_1 = \{e_5, a_5, a_6\}$. In G/W_1, ξ_1' has a single maximal extension, namely $\tau_1' = \{e_2', e_5', a_5', a_6'\}$. Figure 30 shows both G/W_1 and τ_1', with $\Diamond$ representing all of W_1 identified to a singleton.

   Since $\tau_1 \subseteq \sigma$, the loop ends with $k = 1$.

2. In the second invocation of Construction 29, the graph is G/W_1 and the maximal strategy is $\tau_1' = \{e_2', e_5', a_5', a_6'\}$. The strategy has goal state 4, so let $g = 4$. Therefore, in this invocation of the construction, $b_1' = a_4'$, yielding the minimal nonface $\kappa_1' = \{a_4', a_5', a_6'\}$. (We use single prime notation to indicate actions in G/W_1, including references to local variables within this invocation of Construction 29.)

We now write U_1 for $\mathrm{src}(\kappa_1')$, in order to avoid confusion with the earlier W_1. Thus $U_1 = \{4, 5, 6\}$. The comments at the top of page 60 produce $\mathcal{E}_1' = \{a_5', a_6'\}$.

(Below, we will now also use double prime notation, specifically to indicate actions in $(G/W_1)/U_1$, including references to local variables within Construction 29. We therefore write ξ_1' in place of ξ_1 in step 4(b) and ξ_1'' in place of ξ_1' in step 4(c).)

Running the loop of step 4, with $i = 1$, one obtains $\xi_1' = \{e_2'\}$. In $(G/W_1)/U_1$, ξ_1'' has a single maximal extension, namely itself. We refer to that extension as ρ_1'', in order to avoid confusion with the earlier τ_1. Figure 31 shows both $(G/W_1)/U_1$ and ρ_1'', with $\Diamond$ as before and $\Box$ representing all of U_1 identified to a singleton.

Since $\rho_1' \subseteq \tau_1'$, the loop ends with $k = 1$.

3. Since the graph $(G/W_1)/U_1$ has only two states, one could now simply refer to the base case in the proof of Theorem 39. However, we will invoke Construction 29 yet a third time, with graph $(G/W_1)/U_1$ and maximal strategy $\rho_1'' = \{e_2''\}$. This strategy has goal state $\Box$, thus yielding minimal nonface $\kappa_1'' = \{e_2'', e_5''\}$ with source set $\{\Diamond, \Box\}$. The comments at the top of page 60 produce the expansive set of actions $\mathcal{E}_1'' = \{e_2''\}$. The loop ends because the source set is the entire state space, as discussed at the beginning of Section 5.3 on page 62.

Finally, one assembles the various expansive sets in reverse order of the recursive invocations of Construction 29. This process produces the following sequence of sets of actions in G:

$$\{e_2\},\ \{a_5, a_6\},\ \{a_2, a_3\}.$$

That sequence of sets represents four informative action release sequences for G, each consisting of actions in σ:

$$e_2, a_5, a_6, a_2, a_3 \qquad e_2, a_5, a_6, a_3, a_2$$
$$e_2, a_6, a_5, a_2, a_3 \qquad e_2, a_6, a_5, a_3, a_2$$

(We know from Section 4.10.3 that σ also contains other informative action release sequences.)

Figure 31: **Left Panel:** The quotient graph $(G/W_1)/U_1$, with G/W_1 as in Figure 30 and $U_1 = \{4, 5, 6\}$. **Right Panel:** The maximal strategy ρ_1'', obtained in step 4(c) during the first iteration of the loop of Construction 29 as applied to the graph G/W_1 and the strategy τ_1' of Figure 30. (In order to avoid overloaded letters, while retaining indices as in Construction 29, this figure refers to U_1, ρ_1'', and uses double prime notation to indicate actions in $(G/W_1)/U_1$.)

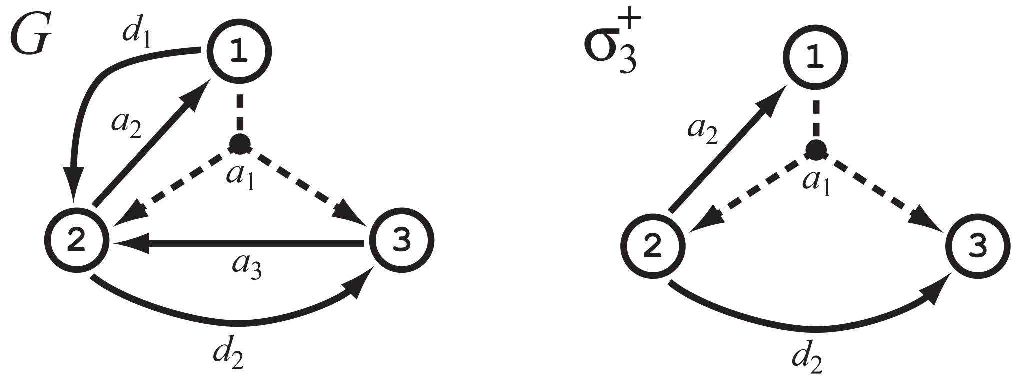

Figure 32: **Left Panel:** A pure stochastic graph G, consisting of three states, four deterministic actions, and one stochastic action. The stochastic action is $a_1 = 1 \to p\{2,3\}$; its action edges appear as dashed lines. The precise probability distribution p is not significant here, except to indicate that each of the transitions $1 \to 2$ and $1 \to 3$ has nonzero probability.
Right Panel: The maximal strategy $\sigma_3^+ \in \Delta_G$, depicted by its actions. See also Figure 33.

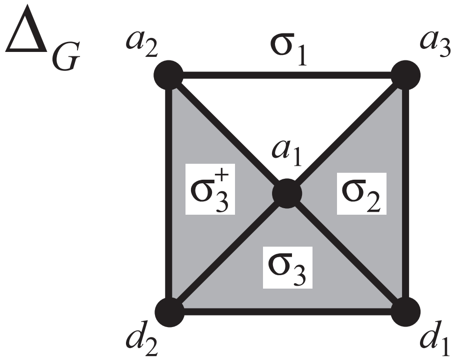

A	a_1	a_2	a_3	d_1	d_2	Goal
σ_1		•	•			1
σ_2	•		•	•		2
σ_3	•			•	•	3
σ_3^+	•	•			•	3

Figure 33: The strategy complex Δ_G of the graph G from Figure 32 appears on the left. Each maximal simplex is labeled with its strategy name, as specified by the relation on the right.

5.4.2 A Pure Stochastic Graph

Figure 32 depicts a fully controllable pure stochastic graph G, along with a maximal strategy that contains an action with stochastic transitions. The strategy complex Δ_G appears in Figure 33, along with G's action relation. The maximal strategy under consideration is σ_3^+.

Unlike in a pure nondeterministic graph, cycling is permitted in a pure stochastic graph, so long as the cycling is transient. The definition of "moves off" from page 19 captures this distinction. For instance, in the current example, $\{a_1, a_2\}$ is a convergent (nonmaximal) strategy, with goal state #3. The set of actions $\{a_1, a_2\}$ would *not* be convergent if action a_1 were nondeterministic, since then an adversary could force infinite cycling between states #1 and #2. However, a_1 is stochastic, so there is a nonzero probability that the system will exit such a cycle, transitioning to state #3 instead. The precise transition probabilities of action a_1 affect expected convergence times, as discussed in [5, 4], but not overall convergence.

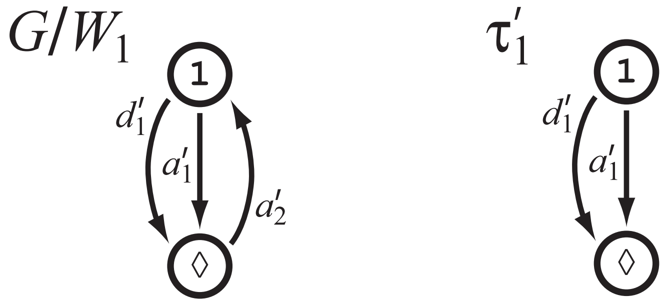

Figure 34: **Left Panel:** The quotient graph G/W_1, with G as in Figure 32 and $W_1 = \{2, 3\}$. **Right Panel:** The maximal strategy τ_1', obtained in step 4(c) during the first iteration of the loop of Construction 29, as applied to the graph G and strategy σ_3^+ of Figure 32.

Given G and $\sigma = \sigma_3^+ = \{a_1, a_2, d_2\}$ as in Figure 32, Construction 29 computes as follows:

1. There is one state outside $\mathrm{src}(\sigma) = \{1, 2\}$, so $g = 3$. Action b_1 in the construction should be an action with source g, meaning it is action a_3 of Figure 32.
2. One may then use minimal nonface $\kappa_1 = \{d_2, a_3\}$. (Another possibility is $\{a_1, a_2, a_3\}$.)
3. So $\mathcal{A}_1 = \{d_2, a_3\}$ and $W_1 = \{2, 3\}$.
4. Now the loop of the construction runs:

$i = 1$:
- (a) Figure 34 depicts graph G/W_1, omitting the self-looping actions at state $\Diamond$.
- (b) $\xi_1 = \{a \in \sigma \mid \mathrm{src}(a) \notin W_1\} = \{a_1\}$.
- (c) ξ_1' is not maximal in Δ_{G/W_1}. It has unique maximal extension $\tau_1' = \{a_1', d_1'\}$.
- (d) $\tau_1 \not\subseteq \sigma$. Since $\tau_1 \setminus \sigma = \{d_1\}$, $b_2 = d_1$. Thus $\kappa_2 = \{d_1, a_2\}$, $\mathcal{A}_2 = \mathcal{A}_1 \cup \kappa_2 = \{d_2, a_3, d_1, a_2\}$, and $W_2 = W_1 \cup \mathrm{src}(\kappa_2) = \{1, 2, 3\}$.

$i = 2$:
- (a) Graph $G/W_2 = (\{\Box\}, \emptyset)$, with $\Box$ representing all states of G identified to a singleton. So $\Delta_{G/W_2} = \{\emptyset\}$, the empty simplicial complex.
- (b) $\xi_2 = \emptyset$.
- (c) ξ_2' is maximal in Δ_{G/W_2}, so $\tau_2' = \emptyset$.
- (d) $\tau_2 \subseteq \sigma$, so the loop ends, with $k = 2$.

Lemma 33 and subsequent comments construct expansive sets $\mathcal{E}_1$ and $\mathcal{E}_2$ as follows (variable bindings for g, W_1, W_2, b_1, b_2, κ_1, and κ_2 are as above, actions d_2 and a_2 are as in Figure 32):

0. Let $W_0 = \{g\} = \{3\}$.
1. Since $W_1 \setminus W_0 = \{2\} = \mathrm{src}(\kappa_1 \setminus \{b_1\})$, $\mathcal{E}_1 = \kappa_1 \setminus \{b_1\} = \{d_2\}$. (See also page 60.)
2. While $W_2 \setminus W_1 = \{1\} \neq \{2\} = \mathrm{src}(\{a_2\}) = \mathrm{src}(\kappa_2 \setminus \{b_2\})$, $\mathrm{src}(b_2) \in \mathrm{trg}(a_2)$, so $\mathcal{E}_2 = \{a_2\}$.

By Corollary 37, the sequence $\mathcal{E}_2, \mathcal{E}_1$ is informative for G, with all actions contained in σ_3^+. We thus obtain the informative action release sequence a_2, d_2. Side note: This is not the only iars contained in σ_3^+. The longest such iars consists of all actions in σ_3^+, in the order a_1, d_2, a_2.

Interpretation: Figure 35 depicts $\mathcal{A}_2$ as a graph. The graph consists of two independent two-cycles. Relationally, we may think of these two two-cycles as two independent bits of information, forming a basis for informative action release sequences. In more general examples (see Section 5.4.3), there may be less independence. Consequently, Lemma 33 (page 59) constructs expansive sets of actions, which Lemma 36 (page 62) then arranges informatively.

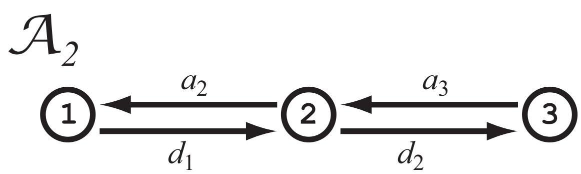

Figure 35: The set of actions $\mathcal{A}_2$ viewed as a graph. Construction 29 produces this set when applied to graph G and maximal strategy σ_3^+ of Figure 32 in the manner discussed on page 67. The actions of $\mathcal{A}_2$ contained in σ_3^+ form an informative action release sequence for G. (In fact, any ordering of any convergent set of actions in $\mathcal{A}_2$ is an iars contained in some strategy, by independence of the two two-cycles in $\mathcal{A}_2$.)

5.4.3 A Pure Stochastic Graph Highlighting Expansive Set Order

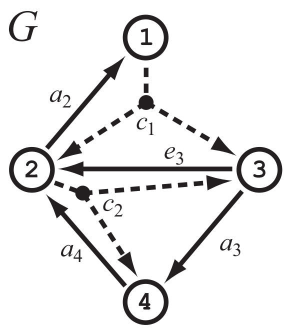

A	a_2	a_3	a_4	e_3	c_1	c_2	Goal
	•	•	•	•			1
		•	•	•	•		2
	•		•		•	•	3
		•		•	•	•	4
	•	•			•	•	4
σ	•	•		•		•	$\{1, 4\}$

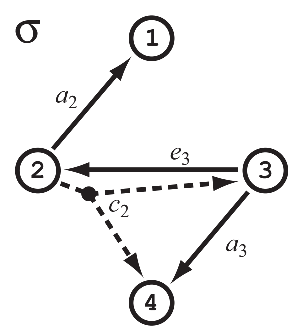

Figure 36: The left panel displays a pure stochastic graph G, consisting of four states, $1, 2, 3, 4$, four deterministic actions, a_2, a_3, a_4, e_3, and two stochastic actions, c_1, c_2. The center panel shows G's action relation. The right panel depicts maximal strategy $\sigma \in \Delta_G$, via its actions.

We now apply Construction 29 to the pure stochastic graph G and maximal strategy $\sigma \in \Delta_G$ of Figure 36. We may let $g = 1$. Then $b_1 = c_1$, $\kappa_1 = \{c_1, a_2, e_3\}$, and $W_1 = \{1, 2, 3\}$. Thus $\mathcal{E}_1 = \{a_2, e_3\}$. Now $\xi_1 = \emptyset$, so there is a choice in constructing τ_1'. If we choose $\tau_1' = \{a_4'\}$, then $\tau_1 \not\subseteq \sigma$ and so $b_2 = a_4$. There are two minimal nonfaces within $\sigma \cup \{a_4\}$. Let us use $\kappa_2 = \{a_4, c_2, a_3\}$. Thus $W_2 = \{1, 2, 3, 4\}$ and the loop ends with k=2. Constructing $\mathcal{E}_2$ involves a choice since actions c_2 and a_3 each have action a_4's source as a target. Let us pick $\mathcal{E}_2 = \{a_3\}$.

Corollary 37 on page 62 arranges the expansive sets of actions in the order $\mathcal{E}_2, \mathcal{E}_1$. Indeed, both a_3, a_2, e_3 and a_3, e_3, a_2 are informative action release sequences for G. Notice that action e_3 implies action a_3 in relation A. Consequently, the order $\mathcal{E}_1, \mathcal{E}_2$ would *not* be acceptable.

Comment: We can lengthen a_3, a_2, e_3 to the iars c_2, a_3, a_2, e_3. In fact, 12 of the 24 possible permutations of all the actions in σ constitute informative action release sequences for G.

6 Counterexamples for Mixed Graphs

The assertion of Theorem 4 on page 33 need not hold for graphs containing a mix of deterministic, nondeterministic, and stochastic actions. Of course, there are many settings in which the assertion does hold. For instance, if one can find a hierarchical cyclic subgraph with the same state space as the given graph, then one can again prove the theorem, even if the graph contains stochastic actions. All the proof needs is for the hierarchical cyclic subgraph to be composed only of deterministic and nondeterministic actions. Absent such structure, it is very easy to construct a counterexample involving a mix of deterministic, nondeterministic, and stochastic actions. With some added effort, one may also construct counterexamples involving maximal strategies that attain singleton goals. This section presents such counterexamples.

6.1 A Counterexample with a Large Goal Set

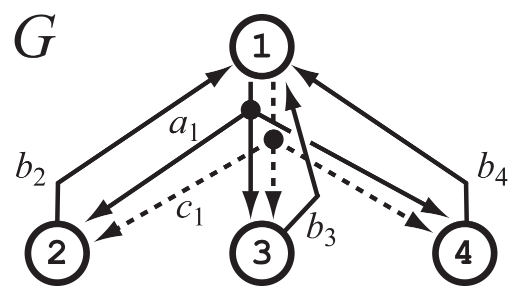

A	a_1	c_1	b_2	b_3	b_4	Goal
σ_1			•	•	•	1
σ_2		•		•	•	2
σ_3		•	•		•	3
σ_4		•	•	•		4
σ_{234}	•	•				$\{2,3,4\}$

Figure 37: **Left Panel:** A graph G with four states, $1, 2, 3, 4$, three deterministic actions, b_2, b_3, b_4, one nondeterministic action, a_1, and one stochastic action, c_1.
Right Panel: G's action relation and goal sets.

Consider the graph and action relation of Figure 37. The graph contains two actions that have identical action edges, but differ in that one action is stochastic and the other action is nondeterministic. The stochastic action is $c_1 = 1 \to p\{2,3,4\}$, for some probability distribution p ascribing nonzero probabilities to each target, and the nondeterministic action is $a_1 = 1 \to \{2,3,4\}$. Additionally, the graph contains three deterministic action, $b_i = i \to 1$, for $i = 2,3,4$.

The graph is fully controllable based just on the set of actions $\{c_1, b_2, b_3, b_4\}$, as one can see from the action relation or as follows: The system can attain state #1 from any other state by using strategy $\{b_2, b_3, b_4\}$. The system can reach a desired state in the set $\{2,3,4\}$ by repeatedly trying to do so using the stochastic action c_1, cycling back to state #1 if that action transitions to the wrong state. For instance, strategy $\{c_1, b_2, b_3\}$ will converge to state #4.

None of the actions $\{b_i\}$ can be in a strategy together with action a_1, but action c_1 can be. In fact, $\sigma_{234} = \{c_1, a_1\}$ is a maximal strategy. The longest informative action release sequence contained in σ_{234} is the strategy itself, revealed in the order c_1, a_1. That iars has length 2, which is less than the number 3 demanded by Theorem 4.

One may easily generalize this example to graphs with n states, for $n > 4$, such that some maximal strategy consists of only two actions and therefore has an iars of length at most 2.

Key to this example is the oddity of having two nearly identical actions, with the only difference being that one action is stochastic and the other is nondeterministic. The graph would *not* be fully controllable with just the nondeterministic action. The stochastic action is needed to “sample with replacement”, i.e., “try and try again, until success”.

From a worst-case perspective, the strategy consisting of the nearly identical stochastic and nondeterministic actions amounts to no more than the nondeterministic action itself. So, why even include the nondeterministic action in the graph?

The answer is that it is a choice a system may make. Executing the stochastic action may entail greater cost than executing the nondeterministic action, because the nondeterministic action relieves the system of guaranteeing stochastic behavior. That may be desirable in some settings. At first it seems hardly so, because the only goal set one can attain using any strategy containing the nondeterministic action is a very large set (consisting of $n-1$ states in the generalized version). However, not caring about precise transitions is sensible when the graph is part of a larger graph and it does not matter what state the system passes through as a subgoal while attaining some overall goal. The next subsection explores such graphs further.

6.2 A Counterexample with a Small Goal Set

Previously, we saw the basis for a family of counterexamples in which the graph has n states but contains a maximal strategy consisting of two actions with a goal set of size $n-1$. One might therefore hypothesize that Theorem 4 should merely assert the existence of an informative action release sequence of length $n-k$, with k being the size of the goal set. In fact, such a theorem would also be false, since one can construct counterexamples to Theorem 4 using strategies that have goal sets of size 1, as this subsection demonstrates.

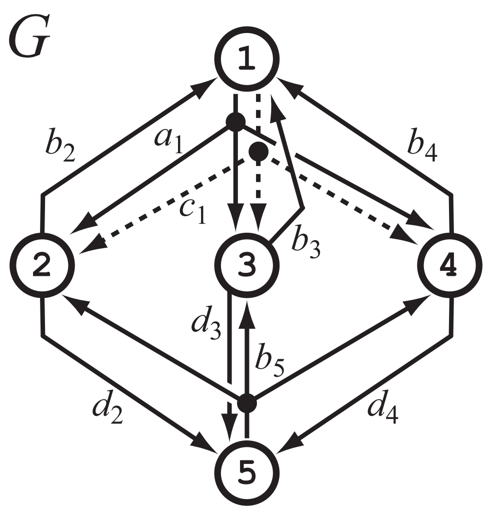

A	a_1	d_2	d_3	d_4	c_1	b_2	b_3	b_4	b_5	Goal
σ_1						•	•	•	•	1
σ_2					•		•	•	•	2
σ_3					•	•		•	•	3
σ_4					•	•	•		•	4
σ_{234}	•				•				•	$\{2,3,4\}$
σ_5	•	•	•	•	•					5
σ_{54}		•	•	•	•	•	•			5
σ_{53}		•	•	•	•	•		•		5
σ_{52}		•	•	•	•		•	•		5
σ_{15}		•	•	•		•	•	•		$\{1,5\}$

Figure 38: **Left Panel:** A graph with five states, $1,2,3,4,5$, six deterministic actions, d_2, d_3, d_4, b_2, b_3, b_4, two nondeterministic actions, a_1, b_5, and one stochastic action, c_1.

Right Panel: The graph’s action relation and goal sets.

(This figure is a copy, with minor notational changes, of Figures 67 and 68 in [6].)

In constructing such counterexamples, one may take the fully controllable graph on n states of the previous subsection and glue the set $\{2, \ldots, n\}$ to another graph, permitting direct motion from any state in $\{2, \ldots, n\}$ to a new state, $\#(n+1)$. In order to retain full controllability, the system also needs to be able to move back from state $\#(n+1)$. In this subsection, we use a single nondeterministic action. In the next subsection, will will use $n-1$ different nondeterministic actions. Numerous variations exist.

Figure 38 replicates a counterexample taken from [6], showing a graph and its action relation. Maximal strategy $\sigma_5 = \{a_1, d_2, d_3, d_4, c_1\}$ converges to singleton goal state #5. The strategy contains 5 actions, but the longest informative action release sequences contained in σ_5 have length 3, which is less than the 4 demanded by Theorem 4. The reason no longer iars exists is because any one of the "downward" actions in the set $\{d_2, d_3, d_4\}$ implies the other two. That fact is clear from the action relation, but can also be understood as follows: First, knowing that a strategy contains one of the actions d_2, d_3, or d_4 means the strategy cannot contain the nondeterministic action b_5. Second, for a maximal strategy, not containing b_5 means the strategy must contain the entire set $\{d_2, d_3, d_4\}$.

6.3 A Counterexample with a Small Goal Set and Nonequivalent Inferences

Finally, we construct a counterexample similar to that of Figure 38, but without requiring equivalence between the deterministic downward actions flowing into state $\#(n+1)$. Instead, any *two* of these downward actions will imply all the downward actions. One may achieve this inference by replacing the single nondeterministic action at state $\#(n+1)$ with $n-1$ different nondeterministic actions. Each of these actions now has a target set of size $n-2$ contained within the set $\{2, \ldots, n\}$. With this counterexample in mind, one may see yet another infinite family of counterexamples, parameterized now by the number of downward actions $\{d_i\}_{i=2}^{n}$ that may be released before all are implied (with n sufficiently large).

Figure 39 shows a graph G in three panels. Figure 40 shows G's action relation. There are six states. Four of the graph's deterministic actions, namely d_2, d_3, d_4, d_5, transition *to* state #6, while four nondeterministic actions, e_2, e_3, e_4, e_5, transition *away* from state #6. Each of those nondeterministic actions has a target set of size three that is a subset of $\{2, 3, 4, 5\}$. The following table shows which pairings of d_i and e_j actions create minimal nonfaces in Δ_G. One sees that any single action drawn from $\{d_2, d_3, d_4, d_5\}$ is potentially consistent with strategies not involving any other d_i action, but that any two of the $\{d_i\}$ actions imply them all. (Any two of the $\{d_i\}$ eliminate all $\{e_j\}$, so the given maximal strategy must contain all $\{d_i\}$.)

	e_5	e_4	e_3	e_2
d_2	•	•	•	
d_3	•	•		•
d_4	•		•	•
d_5		•	•	•

Maximal strategy $\sigma = \{a_1, d_2, d_3, d_4, d_5, c_1\}$ converges to goal state #6 and contains 6 actions. However, since at most two of the actions $\{d_i\}$ are informative, the longest informative action release sequences contained in σ have length 4. That is less than the 5 demanded by Theorem 4.

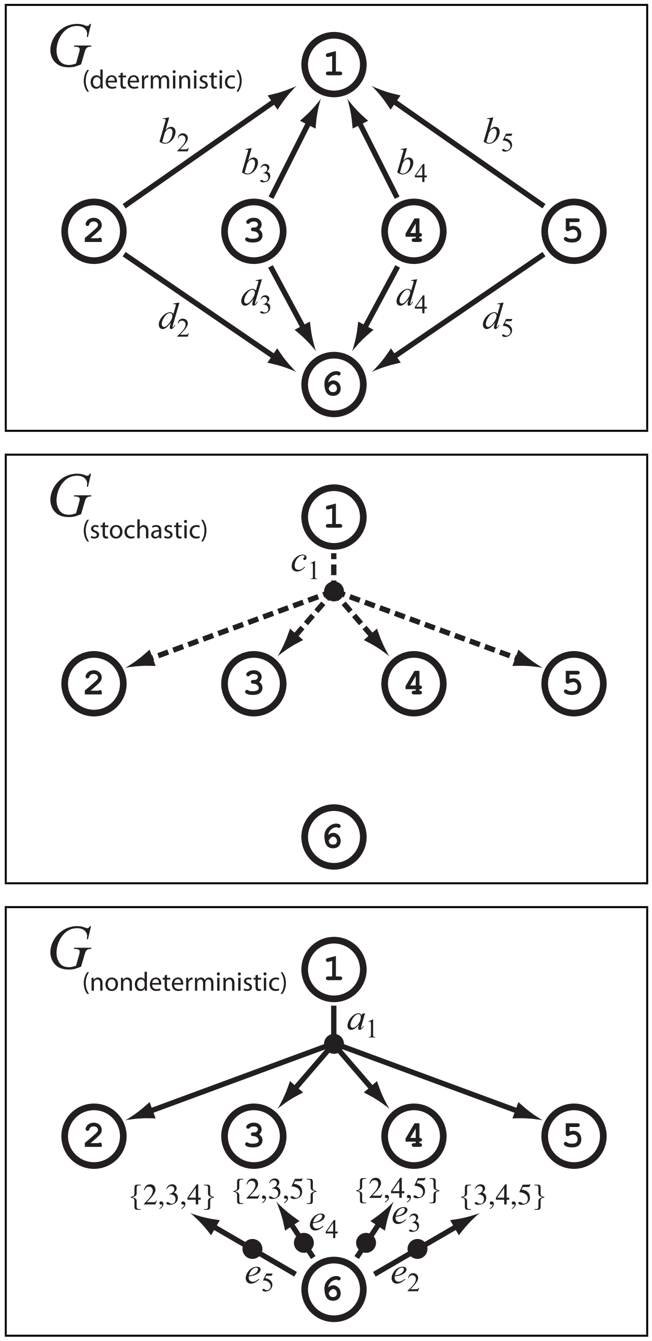

Figure 39: A graph G with eight deterministic actions, one stochastic action, and five nondeterministic actions, depicted in three panels. Figure 40 displays G's action relation.

Top: The top panel shows G's deterministic actions. Action b_i moves "back" from state $\#i$ to state $\#1$, while action d_i moves "down" from state $\#i$ to state $\#6$, for $i = 2, 3, 4, 5$.

Middle: The middle panel shows G's stochastic action c_1, with source state $\#1$ and target set $\{2, 3, 4, 5\}$.

Bottom: The bottom panel shows G's nondeterministic actions. Action a_1 has the same source and targets as the stochastic action, but is nondeterministic. The remaining four actions each have source state $\#6$, and some target set of size three in the set of states $\{2, 3, 4, 5\}$. Specifically, the target set of action e_i is $\{2, 3, 4, 5\}$ but with state $\#i$ "excised", for $i = 2, 3, 4, 5$. The figure displays these four actions in abbreviated form, with written target sets rather than all the arrows drawn.

A	a_1	d_2	d_3	d_4	d_5	c_1	b_2	b_3	b_4	b_5	e_5	e_4	e_3	e_2	Goal
							•	•	•	•	•	•	•	•	1
						•		•	•	•	•	•	•	•	2
						•	•		•	•	•	•	•	•	3
						•	•	•		•	•	•	•	•	4
						•	•	•	•		•	•	•	•	5
	•					•					•	•	•	•	$\{2,3,4,5\}$
					•		•	•	•	•	•				1
					•	•		•	•	•	•				2
					•	•	•		•	•	•				3
					•	•	•	•		•	•				4
	•				•	•					•				$\{2,3,4\}$
				•			•	•	•	•		•			1
				•		•		•	•	•		•			2
				•		•	•		•	•		•			3
				•		•	•	•	•			•			5
	•			•		•						•			$\{2,3,5\}$
			•				•	•	•	•			•		1
			•			•		•	•	•			•		2
			•			•	•	•		•			•		4
			•			•	•	•	•				•		5
	•		•			•							•		$\{2,4,5\}$
		•					•	•	•	•				•	1
		•				•	•		•	•				•	3
		•				•	•	•		•				•	4
		•				•	•	•	•					•	5
	•	•				•								•	$\{3,4,5\}$
		•	•	•	•		•	•	•	•					$\{1,6\}$
		•	•	•	•	•		•	•	•					6
		•	•	•	•	•	•		•	•					6
		•	•	•	•	•	•	•		•					6
		•	•	•	•	•	•	•	•						6
σ	•	•	•	•	•	•									6

Figure 40: Action relation and goal sets for the graph G of Figure 39. The maximal strategy in the bottommost row is labeled σ for reference in the text.

Acknowledgments

Many thanks to Rob Ghrist, Steve LaValle, Ben Mann, and Matt Mason, for advice and support related to this research and its earlier foundations within SToMP.

References

[1] A. Björner and V. Welker. Complexes of directed graphs. *SIAM J. Discrete Math*, 12(4):413–424, 1999.

[2] I. Dinur and K. Nissim. Revealing information while preserving privacy. In *Proceedings of the 22nd ACM Symposium on Principles of Database Systems*, pages 202–210, 2003.

[3] C. Dwork. A firm foundation for private data analysis. *Communications of the ACM*, 54(1):86–95, 2011.

[4] M. A. Erdmann. On the topology of discrete strategies. *International Journal of Robotics Research*, 29(7):855–896, 2010.

[5] M. A. Erdmann. On the topology of discrete planning with uncertainty. In A. Zomorodian, editor, *Advances in Applied and Computational Topology*, pages 147–194. AMS, 2012.

[6] M. A. Erdmann. Topology of privacy: Lattice structures and information bubbles for inference and obfuscation. `http://arxiv.org/abs/1712.04130`, 2017.

[7] W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, New York, third edition, 1968.

[8] A. Hultman. Directed subgraph complexes. *Elec. J. Combinatorics*, 11(1):R75, 2004.

[9] J. Jonsson. *Simplicial Complexes of Graphs.* PhD thesis, Department of Mathematics, KTH, Stockholm, Sweden, 2005.

[10] S. Karlin and H. M. Taylor. *A Second Course in Stochastic Processes.* Academic Press, New York, 1981.

[11] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, Menlo Park, CA, 1984.

[12] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 111–125, 2008.

[13] J. J. Rotman. *An Introduction to Algebraic Topology.* Springer-Verlag, New York, 1988.

[14] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.

[15] M. L. Wachs. *Poset Topology: Tools and Applications.* IAS/Park City Mathematics Institute, Summer 2004. Also available here: `http://arxiv.org/abs/math/0602226`.