A Appendix

A.1 Mathematical background on cube complexes and Gromov's Link Condition

Cube complexes and simplicial complexes are higher dimensional analogues of graphs that appear prominently in topology, geometric group theory, and combinatorics. Background on cube complexes can be found in [Sch19, Sag14, Wis12][5], while simplicial complexes are detailed in standard algebraic topology texts [EH10]. Here, we will only provide brief explanations in order to discuss Gromov's Link Condition.

Cube complexes. Informally, a cube complex is a space that can be constructed by gluing cubes together in a fashion not too dissimilar to a child's building blocks. An n–cube is modelled on

$$\{(x_1, \ldots, x_n) \in \mathbb{R}^n \ : 0 \leq x_i \leq 1 \text{ for all } i\}.$$

By restricting some co-ordinates to either 0 or 1, we can obtain lower dimensional subcubes. In particular, an n–cube has 2^n vertices and is bounded by $2n$ faces which are themselves $(n-1)$–cubes. A *cube complex* X is a union of cubes, where the intersection of every pair of distinct cubes is either empty, or a common subcube.

Simplicial complexes. Simplicial complexes are constructed in a similar manner to cube complexes, except that we use higher dimensional analogues of triangles or tetrahedra instead of cubes. An n–dimensional simplex (or n–simplex) is modelled on

$$\{(x_1, \ldots, x_{n+1}) \in \mathbb{R}^{n+1} \ : x_i \geq 0 \text{ for all } i, \ \sum_i x_i = 1\};$$

this has $n+1$ vertices and is bounded by $n+1$ faces which are themselves $(n-1)$–simplices. For $n = 0, 1, 2, 3$, an n–simplex is respectively a point, line segment, triangle, and tetrahedron. A *simplicial complex* K is an object that can be constructed by taking a graph and then inductively filling in simplices of progressively higher dimension; this graph is called the 1–*skeleton* of K. We require that every finite set of vertices in K form the vertices of (or *spans*) at most one simplex; thus simplices in K are uniquely determined by their vertices. (This rules out loops or multi-edges in the 1–skeleton.)

Links. The local geometry about a vertex v in a cube complex X is captured by a simplicial complex known as its *link* $lk(v)$. Intuitively, this is the intersection of a small sphere centred at v within X, and can be regarded as the space of possible directions emanating from v. Each edge in X emanating from v determines a vertex (0–simplex) in $lk(v)$. If two such edges bound a 'corner' of a square in X based at v, then there is an edge (1–simplex) connecting the associated vertices in $lk(v)$. More generally, each 'corner' of an n–cube incident to v gives rise to an $(n-1)$–simplex in $lk(v)$; moreover, the boundary faces of the simplex naturally correspond to the faces of the cube bounding the corner. Since the cube complexes we consider have cubes completely determined by their vertices, each simplex in $lk(v)$ is also completely determined by its vertices. Figure 8 illustrates four separate examples of links of vertices in cube complexes.

Gromov's Link Condition. Local curvature in a cube complex can be detected by examining the combinatorial structure of the links of its vertices. Specifically, Gromov's Link Condition gives a method for proving that a cube complex is *non-positively curved* (NPC)[6], where there is an absence of positive curvature. In the bottom-right example in Figure 8, where there is positive curvature, we observe a 'hollow' triangle in its link. In the other examples of Figure 8, where there is only negative or zero curvature, there are no such hollow triangles (or hollow simplices).

This absence of 'hollow' or 'empty' simplices is formalised by the *flag* property: a simplicial complex is *flag* if whenever a set of $n+1$ vertices spans a complete subgraph in the 1–skeleton, they must span an n–simplex. In particular, a flag simplicial complex is determined completely by its 1–skeleton. If v is a vertex in a cube complex X, then the flag condition on $lk(v)$ can be re-interpreted as a 'no empty corner' condition for the cube complex: whenever we see (what appears to be) the corner of an n–cube, then the whole n–cube actually exists.

[5] Much of the literature in geometric group theory focusses primarily on non-positively curved cube complexes, whereas in our study, the presense of positive curvature plays a crucial role.

[6] In the sense that geodesic triangles are no fatter than Euclidean triangles [BH99].

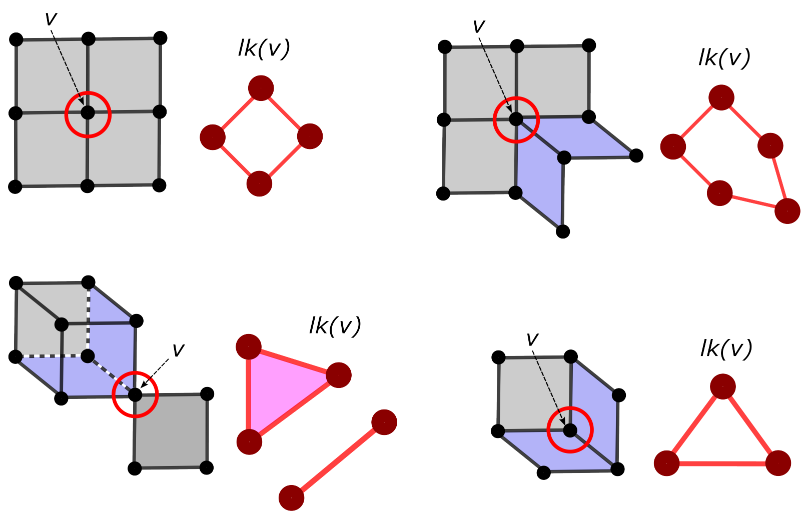

Figure 8: Four separate examples of links of vertices in cube complexes. In the bottom-right example, where there is positive curvature, $lk(v)$ is a 'hollow' triangle and is thus not a flag simplicial complex. For the other examples, $lk(v)$ is a flag complex and therefore, by Gromov's Link Condition, there is only negative or zero curvature. In the bottom-left example, the cube complex is a solid cube joined to a filled-in square at a common vertex v.

Theorem (Gromov's Link Condition [Gro87])**.** *A finite-dimensional cube complex X is non-positively curved if and only if the link of each vertex in X is a flag simplicial complex.* □

Thus, the local geometry of a cube complex is determined by the combinatorics of its links.

A.2 Proof of Theorem 5.2

Before giving our proof, we first classify low-dimensional simplices in $lk(v)$ for a vertex v in our modified state complex $\mathcal{S}'$. A 0–simplex in $lk(v)$ corresponds to an admissible move at v. However, a 1–simplex either represents a pair of commuting moves, or two moves in a common dance. A 2–simplex either represents three agents moving pairwise independently, or a dancing agent commuting with a moving agent. Finally, a 3–simplex represents either four agents moving pairwise independently, one dancing agent and two moving agents that pairwise commute, or a pair of commuting dancers.

Theorem (Gromov's Link Condition in the modified state complex)**.** *Let v be a vertex in the modified state complex $\mathcal{S}'$ of an agent-only gridworld. Then*

- *$lk(v)$ satisfies Gromov's Link Condition if and only if it has no empty 2–simplices nor 3–simplices, and*
- *if $lk(v)$ fails Gromov's Link Condition then there exist a pair of agents whose positions differ by either a knight move or a 2–step bishop move (as in Figure 7).*

Proof. If $lk(v)$ satisfies Gromov's Link Condition, then it has no empty simplices of any dimension, giving the forward implication. For the converse, assume that $lk(v)$ has no empty 2–simplices nor 3–simplices. Suppose there exist $n+1$ vertices spanning a complete subgraph of $lk(v)$, where $n \geq 4$. We want to show that these vertices span an n–simplex. By induction, we may assume that every subset of n vertices from this set spans an $(n-1)$–simplex. Since $n \geq 4$, every quartuple of vertices in this subgraph spans a 3–simplex. Therefore, appealing to our classification of low-dimensional simplices, every pair of moves or dances involved has disjoint supports. Thus, the desired n–simplex exists. Consequently, potential failures can only be caused by empty 2–simplices or 3–simplices.

Next, we want to determine when three pairwise adjacent vertices in $lk(v)$ span a 2–simplex. These vertices represent three admissible moves at v. Since they are pairwise adjacent, they either correspond to three agents each doing a Move, or to one agent dancing with another one moving. In the former case, the supports are pairwise disjoint and so these moves form a commuting set of generators. Therefore, the desired 2–simplex exists (indeed, in the absence of dancers, the situation is the same as the original Abrams, Ghrist & Peterson setup). For the latter case, suppose that the first agent is dancing while the second moves. Since the 0–simplices are pairwise adjacent, each of the two admissible moves within the dance has disjoint support with the second agent's move. Thus, the only way the support of the dance fails to be disjoint from that of the second agent's move is if the second agent can move into the diagonally opposite corner of the dance. Therefore, the only way an empty 2–simplex can arise is if the agents' positions differ by a 'knight move' (see Figure 7 for illustration).

It remains to determine when four pairwise adjacent vertices in $lk(v)$ span a 3–simplex. We may assume that each triple of vertices in this set spans a 2–simplex, for otherwise we can reduce to the previous case. Let us analyse each case by the number of involved agents. If there are four involved agents, then each 0–simplex corresponds to exactly one agent moving. Since no dances are involved, it immediately follows that the desired 3–simplex exists. If there are three involved agents, then one is dancing while the other two move. Since each triple of 0–simplices spans a 2–simplex, we deduce that each move has disjoint support with the dance. Therefore, the dance and the two moves form a commuting set, and so the 3–simplex exists. Finally, if there are two agents then they must both be dancers. By the assumption on 2–simplices, each admissible move within the dance of one agent has disjoint support from the dance of the other agent. Thus, the only way for the two dances to have overlapping supports is if their respective diagonally opposite corners land on the same cell. Therefore, the only way an empty 3–simplex can arise (assuming no empty 2–simplices) is if two agents' positions differ by a '2–step bishop move' (see Figure 7 for illustration). □

A.3 Python tool for constructing gridworlds and their state complexes

We developed a Python-based tool for constructing gridworlds with objects and agents. It includes a GUI application for the easy specification of gridworlds and a script which will produce plots and data of the resulting state complex. We ran all experiments on a Lenovo IdeaPad 510-15ISK laptop. [A link to the open-source code will be placed here in the final version, but details are redacted in this review version, in keeping with NeurIPS's double-blind review policy. A copy of our code is included in this NeurIPS submission.]

For the sake of generality and future-proofing of our software, we chose to construct the links in our implementation of checking Gromov's Link Condition in gridworlds, which is not necessary in-practice. Instead, in practical situations, one can directly check for supports of knight or two-step bishop moves between agents, which per Theorem 5.2 provides a computational short-cut for detecting failures in agent-only gridworlds. Another area of computational efficiency available in many rooms are in the symmetries of the room itself. For example, an evenly-sized square room can be cut into eighths (like a square pizza), where each eighth is geometrically identical to every other.

Users of the code will notice a small but important implementation detail in the code which we chose to omit the particulars of in this paper: in the code, we need to include labelled walls along the borders of our gridworlds. This is because we construct our gridworlds computationally as coordinate-free, abstract graphs. For Move, the lack of a coordinate system is not an issue – if an agent label sees a neighbouring vertex with an empty floor label, the support exists and the generator can be used. However, Push/Pull only allows objects to be pushed or pulled by the agent in a straight line within the gridworld. We ensure this straightness in the abstract graph by identifying a larger subgraph around the object and agent than is illustrated in Figure 2. Essentially, we incorporate three wildcard cells (cells of any labelling) adjacent to three labelled cells (`'agent'`, `'object'`, and `'floor'`), such that together they form a 2×3 grid.

A.4 Experiments in small rooms

Summary statistics for the 3×3 room with varying numbers of agents is shown in Table 1, showing the distribution in failures of Gromov's Link Condition across these conditions. The 2×3 room with two agents shows multiple instances of local positive curvature in the associated (modified) state complex. Figure 9 shows one such state where Gromov's Link Condition fails due to the agents being separated by a knight's move (see Theorem 5.2). At this state, there are actually two empty 2–simplices in its link – this is because the pattern appearing in the 5–cell subgrid with two agents

Table 1: Data of Gromov's Link Condition failure and commuting 4–cycles in the state complexes of a 3×3 room with varying numbers of agents and no objects. The percentage of NPC states (shown in brackets in the second column) is rounded to the nearest integer. The mean number of Gromov's Link Condition failures (shown in the penultimate column) is the mean number of failures over the total number of states, and is rounded to two decimal places.

				Gromov's Link Condition Failures		
Agents	States (% NPC)	Dances	Commuting moves	Total	Mean	Max
0	1 (100)	0	0	0	0	0
1	9 (100)	4	0	0	0	0
2	36 (78)	20	44	32	0.89	4
3	84 (62)	40	220	184	2.19	14
4	126 (65)	40	440	288	2.29	11
5	126 (68)	20	440	152	1.21	6
6	84 (86)	4	220	16	0.19	2
7	36 (100)	0	44	0	0	0
8	9 (100)	0	0	0	0	0
9	1 (100)	0	0	0	0	0

(as in Figure 7) arises in two different ways within the given state on the gridworld. The only other state where Gromov's Link Condition fails is a mirror image of the one shown.

Further small gridworlds and their respective state complexes are shown in Figures 10, 11, and 12.

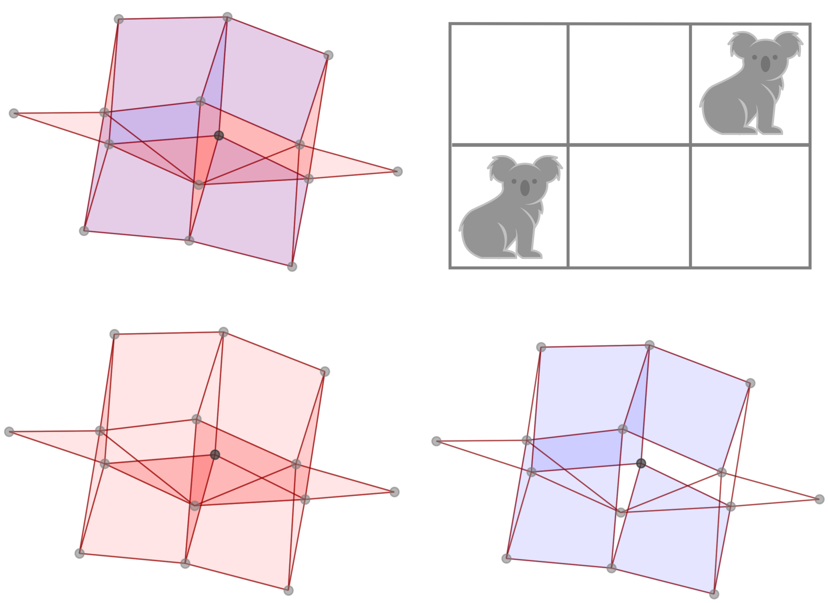

Figure 9: A 2×3 room with two agents (top right) and its state complex (top left), where dances are shaded blue and commuting moves are shaded red. The darker-shaded vertex represents the state of the gridworld shown. Also shown is the state complex with only commuting moves (bottom left) and only dances (bottom right).

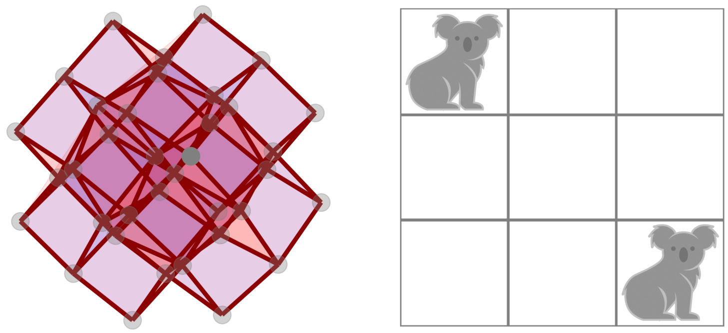

Figure 10: A 3×3 room with two agents (right) and its state complex (left), where dances are shaded blue and commuting moves are shaded red. The darker-shaded vertex represents the state of the gridworld shown. Naturally-occurring copies of this state complex can be found as sub-complexes in the state complex shown in Figure 12.

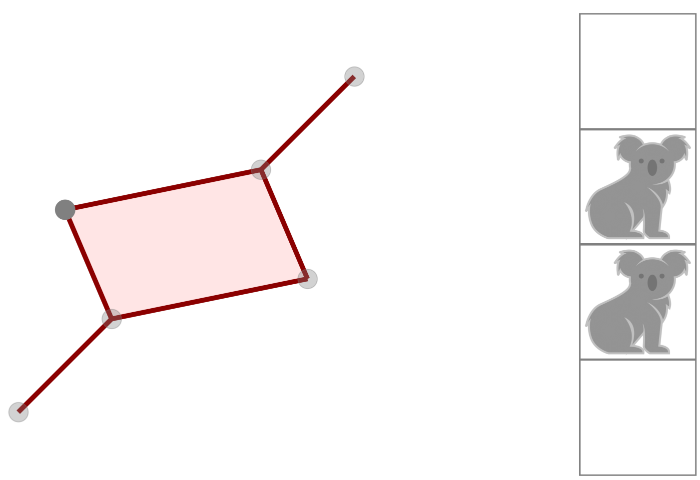

Figure 11: A 4×1 corridor with two agents (right) and its state complex (left). There are no dances and only one commuting move, shaded red. The darker-shaded vertex represents the state of the gridworld shown. Naturally-occurring copies of this state complex can be found as sub-complexes in the state complex shown in Figure 12.

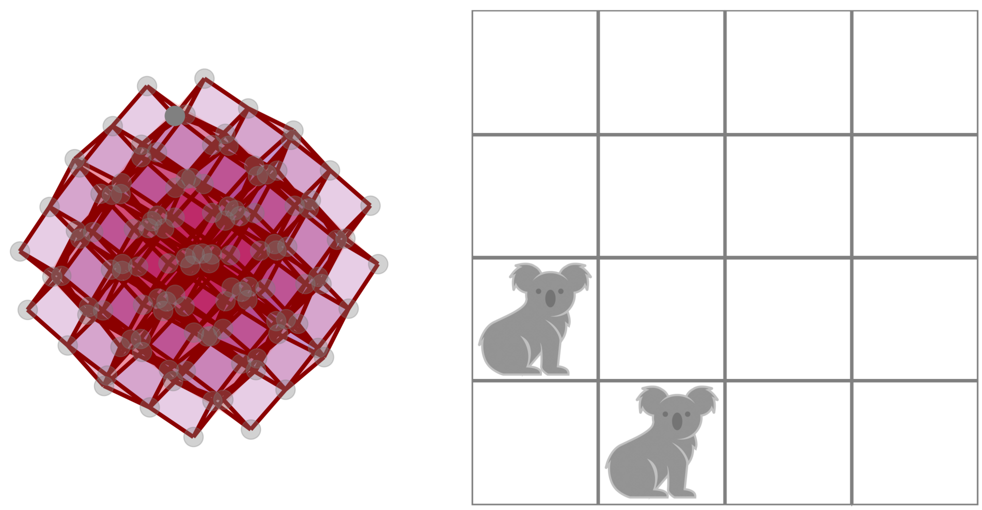

Figure 12: A 4×4 room with two agents (right) and its state complex (left), where dances are shaded blue and commuting moves are shaded red. The darker-shaded vertex represents the state of the gridworld shown. Embedded within this state complex are naturally-occurring copies of the state complex of the 4×1 corridor with two agents, shown in Figure 11. There are also naturally-occurring copies of state complex of the 3×3 room with two agents, shown in Figure 10.