Computational Complexity of k-Block Conjugacy

Tyler Schrock Rafael Frongillo
University of Colorado, Boulder

September 9, 2019

Abstract

We consider several computational problems related to conjugacy between subshifts of finite type, restricted to k-block codes: verifying a proposed k-block conjugacy, deciding if two shifts admit a k-block conjugacy, and reducing the representation size of a shift via a k-block conjugacy. We give a polynomial-time algorithm for verification, and show GI- and NP-hardness for deciding conjugacy and reducing representation size, respectively. Our approach focuses on 1-block conjugacies between vertex shifts, from which we generalize to k-block conjugacies and to edge shifts. We conclude with several open problems.

1 Introduction

One-dimensional subshifts of finite type (SFTs) are of fundamental importance in the study of symbolic dynamical systems. Despite their central role in symbolic dynamics, however, several basic questions about SFTs remain open, particularly with regard to computation. Most prominent is the conjugacy problem: whether it is possible to decide if two given SFTs are conjugate. In this work, we study restricted versions of the conjugacy problem, with an eye toward applications (algorithms to simplify representations of SFTs) as well as developing insights toward the full conjugacy problem. In particular, we address the computational complexity of deciding or verifying conjugacy when given a bound on the block size of the corresponding sliding block code. We focus on the case of vertex shifts; see below for other representations, notably edge shifts.

First consider the question of verification: given two vertex shifts and a proposed sliding block code, what is the computational complexity of verifying that the code induces a conjugacy? We give a polynomial-time algorithm, for both the irreducible and reducible cases; a polynomial-time algorithm in latter case was not known (§ 3). Second, the question of deciding k-block conjugacy: given two vertex shifts, what is the complexity of deciding if there exists a sliding block code, with block length at most k, that induces a conjugacy? By the first result on efficient verification, this problem is in NP; we show it to be GI-hard (at least as hard as the Graph Isomorphism problem) for all k (§ 4). Third, the question of reduction: given a vertex shift and integer ℓ, what is the complexity of deciding whether there exists a k-block conjugacy which reduces the number of vertices by ℓ? Extending a construction from previous work [1], we show that this problem, for $k = 1$, is NP-complete (§ 5).

It is interesting to contrast our results with those of previous work [1], on the special case of $k = 1$ with the restriction that the block code be a sequence of amalgamations. (Recall that any conjugacy can be expressed as a sequences of splittings followed by amalgamations; see § 2.) This previous work shows that the analogous version of our third problem, of reducing the number of vertices using only amalgamations, is NP-complete, but it does not address the verification problem; intuitively it seems plausible that verification would also be NP-hard. Returning to our setting, note that general 1-block codes need not be sequences of amalgamations (Figure 1). Thus, while it is unsurprising that the reduction problem remains NP-hard in our setting, it is perhaps surprising given that verification can be done in polynomial time, as a priori the number of splittings required could be super-polynomial.

Edge shifts have received more attention in the literature, perhaps because of their succinct representations as integer matrices. Precisely because of their succinct representations, the question of verification is somewhat nuanced: verifying a given sliding block code requires writing down the proposed code, which can be exponential in the description size of the original edge shifts, so while the runtime of our algorithm can be exponential in the description sizes of the shifts, it is still polynomial-time (§ 6). We also show GI-hardness for the corresponding conjugacy problems, and leave several open questions (§ 7).

2 Setting

We begin with basic graph-theoretic definitions and convention. A directed graph $G = (V, E)$ is a set of vertices V along with a set of edges $E \subseteq V \times V$. When multiple graphs are in play, we will write $G = (V_G, E_G)$ to clarify which graphs the vertices or edges correspond to. For a directed graph $G = (V, E)$ and a vertex $v \in V$, we define $N^+(v) = \{u \in V : (v, u) \in E\}$ and $N^-(v) = \{u \in V : (u, v) \in E\}$ to be the set of out-neighbors and in-neighbors of v, respectively. Unless specified otherwise, *a cycle of length* n will mean a sequence $v_1 v_2 \cdots v_n \in V$ such that $(v_i, v_{i+1}) \in E$ for $i \in \{1, \dots, n+1\}$ where $v_{n+1} := v_1$. That is, $v_1 v_2 \cdots v_n$ is a cycle in our terminology if the path $v_1 v_2 \cdots v_n v_1$ forms a cycle in G. We define $C_n(G)$ to be the set of cycles of length n in $G = (V, E)$. For example, if the path $v_1 v_2 v_3 v_1$ forms a cycle in G, then $v_1 v_2 v_3, v_2 v_3 v_1, v_3 v_1 v_2 \in C_3(G)$.

Let $\mathcal{A}$ be a finite set. The *full shift* $\mathcal{A}^{\mathbb{Z}}$ over alphabet $\mathcal{A}$ is the set $\{(x_i)_{i\in\mathbb{Z}} : x_i \in \mathcal{A} \text{ for all } i \in \mathbb{Z}\}$. An element of $\mathcal{A}^{\mathbb{Z}}$ is called a *point*. A *block* (or *word*) in $\mathcal{A}$ is a string $a_1 a_2 \cdots a_n$ of symbols from $\mathcal{A}$. We will use the term *infinite word* to describe strings in $\mathcal{A}$ which are infinite in exactly one direction. If $x = (x_i)_{i\in\mathbb{Z}} \in \mathcal{A}^{\mathbb{Z}}$, we use $x_{[a,b]}$ to denote the block $x_a x_{a+1} \cdots x_b$. Similarly, we use $x_{[a,\infty)}$ to denote the infinite word $x_a x_{a+1} \cdots$. Let $\mathcal{F}$ be a set of blocks over $\mathcal{A}$ called *forbidden blocks*. Then $X_{\mathcal{F}}$ is defined to be the subset of $\mathcal{A}^{\mathbb{Z}}$ where each $x \in X_{\mathcal{F}}$ contains none of the forbidden block in $\mathcal{F}$. A *shift space* (or *shift*) is a subset $X \subseteq \mathcal{A}^{\mathbb{Z}}$ such that $X = X_{\mathcal{F}}$ for some set of forbidden blocks $\mathcal{F}$. If there exists a finite set $\mathcal{F}$ such that $X = X_{\mathcal{F}}$, then X is called a *shift of finite type*.

Given a directed graph $G = (V, E)$ with labeled vertices (each distinct), we associate to it the shift space $X_G = \{(v_i)_{i\in\mathbb{Z}} : v_i \in V, (v_i, v_{i+1}) \in E \text{ for all } i \in \mathbb{Z}\}$, which is the collection of all bi-infinite walks on G. Note that X_G is a shift of finite type with $\mathcal{F} = \{v_i v_j : (v_i, v_j) \notin V\}$. Any shift space of this form is called a *vertex shift*. Similarly, given a directed multigraph $G = (V, E)$, i.e. where E is a multiset, and a labeling of the edges from $\mathcal{A}$, we define the *edge shift* X^e_G of labelings of bi-infinite walks on G. Again edge shifts are shifts of finite type with $\mathcal{F} = \{e_1 e_2 : e_1 \text{ does not terminate at the initial vertex of } e_2\}$.

A shift X is *irreducible* if for every pair of words w_1, w_2 in X, there is a word w_3 such that $w_1 w_3 w_2$ is a word in X, and X is *reducible* if it is not irreducible. In graph-theoretic terms, first consider any graph containing a vertex with either no out-neighbors or no in-neighbors. Such a vertex is called *stranded*. A graph (or multigraph) with no stranded vertices is called *essential*. A graph (or multigraph) with the property that for every pair of vertices u, v there is a path from u to v is called *strongly connected*. Finally, a vertex shift X_G (or edge shift X^e_G) is irreducible if G is essential and strongly connected.

Given a shift X with alphabet $\mathcal{A}_1$, we can transform X into a shift space over another alphabet $\mathcal{A}_2$ in the following way. Fix integers m, a with $-m \leq a$. Then letting $\mathcal{B}_n(X)$ denote the set of blocks of size n from the shift X and given a function $\Phi : \mathcal{B}_{m+a+1}(X) \to \mathcal{A}_2$, the corresponding *sliding block code* with memory m and anticipation a is the function Φ_∞ defined by $\Phi_\infty((x_i)_{i\in\mathbb{Z}}) = (\Phi(x_{[i-m,i+a]}))_{i\in\mathbb{Z}}$. That is, Φ_∞ looks at a block of size $m + a + 1$ through a window to determine a character from $\mathcal{A}_2$. Then the window is slid infinitely in both directions. Letting $k = m + a + 1$, we will call any sliding block code with window size k a *k-block code*. Given a sliding block code as $\Phi : \mathcal{A}_1^k \to \mathcal{A}_2$, we extend Φ to all finite and infinite words w of length at least k by $\Phi((w_i)_{i\in I}) = (\Phi(w_{[i-m,i+a]}))_{i-m,i+a\in I}$, where $I \subsetneq \mathbb{Z}$. That is, we extend Φ to words by sliding Φ over the entire word.

Let X be any shift space with alphabet $\mathcal{A}_1$. We define the kth higher block shift $X^{[k]}$ with alphabet $\mathcal{A}_2 = \mathcal{B}_k(X)$ by the image of X under $\beta_N : X \to (\mathcal{B}_k(X))^{\mathbb{Z}}$ where for any point $p \in X$, $\beta_N(p)_i = p_{[i,i+k-1]}$. If $X = X_G$ happens to be a vertex shift, we can construct the kth higher block shift in terms of the graph. For any directed graph G, construct the graph $G^{[k]}$ by $V_{G^{[k]}} = \{v_1 \cdots v_k : v_1 \cdots v_k \text{ is a path in } G\}$ and $E_{G^{[k]}} = \{(v_1 v_2 \cdots v_k, v_2 \cdots v_k v_{k+1}) : v_1 \cdots v_{k+1} \text{ is a path in } G\}$. Then $X_{G^{[k]}} = X_G^{[k]}$. When dealing with k-block codes, it is often useful to *pass to a higher block shift* by noting that there is a k-block conjugacy $\Phi_\infty : X \to Y$ if and only if there is a 1-block conjugacy $\Phi_\infty^{[k]} : X^{[k]} \to Y$ [6, Proposition 1.5.12].

Furthermore, any sliding block code $\Phi_\infty : X_G \to X_H$ between vertex shifts induces the function $\Phi_{\mathsf{c}} : \bigcup_{n=1}^{\infty} C_n(G) \to \bigcup_{n=1}^{\infty} C_n(H)$ defined as follows. Given a cycle c in G, there is a unique cycle d in H with $|c| = |d|$ such that $\Phi_\infty(c^\infty) = d^\infty$; we set $\Phi_{\mathsf{c}}(c) = d$. In the special case of a 1-block code, the block map $\Phi : \mathcal{A}_1 \to \mathcal{A}_2$ is simply a map between the alphabets. In this case, when X_G is a vertex shift, we have $\Phi_{\mathsf{c}}(v_1 \cdots v_n) = \Phi(v_1) \cdots \Phi(v_n)$.

Definition 1. Let X_G be a vertex shift. We say states $u, v \in V_G$ can be *amalgamated* if one the following conditions is met.

1. $N^+(u) = N^+(v)$ and $N^-(u) \cap N^-(v) = \emptyset$
2. $N^-(u) = N^-(v)$ and $N^+(u) \cap N^+(v) = \emptyset$

If u and v are amalgamated, they are replaced by the vertex uv which has $N^+(uv) = N^+(u) \cup N^+(v)$ and $N^-(uv) = N^-(u) \cup N^-(v)$.

Definition 2. Let X_G be a vertex shift. A vertex $v \in V_G$ can be split into two vertices v_1 and v_2 provided the edges of v_1, v_2 satisfy one of the following conditions.

1. $\{N^+(v_1), N^+(v_2)\}$ is a partition of $N^+(v)$ and $N^-(v_1) = N^-(v_2) = N^-(v)$.
2. $\{N^-(v_1), N^-(v_2)\}$ is a partition of $N^-(v)$ and $N^+(v_1) = N^+(v_2) = N^+(v)$.

The corresponding new graph is called a *state splitting* of v. Note that state splittings and amalgamations are inverse operations.

The definitions for edge shifts are similar. Since edges shifts are based on multigraphs, $N^-(v)$ and $N^+(v)$ are multisets. The definition of a state splitting is identical noting that the partition is a multiset partition. For amalgamations, two vertices u, v can be amalgamated if $N^-(u) = N^-(v)$ or $N^+(u) = N^+(v)$. In the case where $N^-(u) = N^-(v)$, u, v are replaced by a single vertex uv with $N^-(uv) = N^-(u) = N^-(v)$ and $N^+(uv) = N^+(u) \uplus N^+(v)$, where $N^+(u) \uplus N^+(v)$ is the multiset disjoint union.

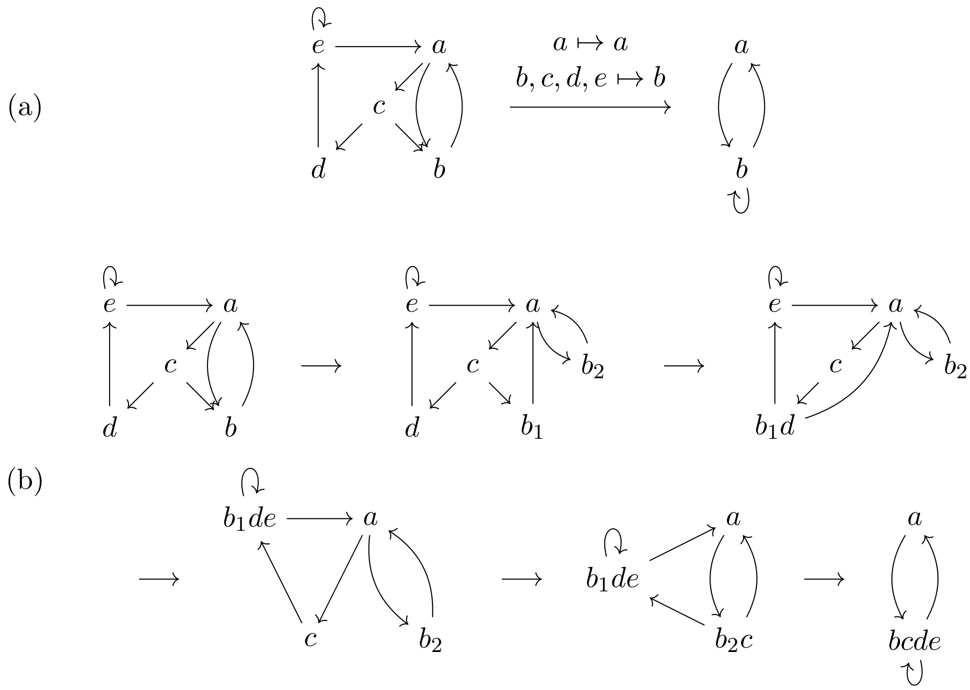

Figure 1: (a) A minimal example of two vertex shifts which are conjugate by a 1-block code but by not a sequence of amalgamations. (b) The conjugacy, demonstrated via a splitting followed by four amalgamations.

Theorem 3 ([9, 6])**.** *Let X_G, X_H be vertex shifts (or edge shifts). Then X_G and X_H are conjugate if and only if there is a sequence of state splittings followed by a sequence of amalgamations which transform G into H.*

In the case of a 1-block code $\Phi : V_G \to V_H$, we may view the block map as a partition of the vertices of G, where each element of the partition is converted to a vertex of H. In light of Theorem 3, it may be tempting to think that every 1-block code can be written as a sequence of amalgamations only, as intuitively splitting a vertex while requiring the vertices be re-amalgamated has no benefit. Yet this statement is not true; there are simple examples of two graphs admitting a 1-block conjugacy, where no pair of vertices can be amalgamated in either graph (Figure 1).

We conclude the background with a common way a sliding block code can fail to be injective. Given a k-block code $\Phi_\infty : X_G \to X_H$, if there exist distinct words w_2, w_2' such that $\Phi(w_1 w_2 w_3) = \Phi(w_1 w_2' w_3)$ with $|w_1| = |w_3| = k$, we say Φ_∞ *collapses a diamond.* As we now state, if a sliding block code is injective, it cannot collapse a diamond. (As we discuss in § 3.2, if Φ_{c} is injective, collapsing a diamond is actually the only way Φ_∞ can fail to be injective.) We prove the result for completeness; see, e.g., [6, Theorem 8.1.16] for a similar result in the irreducible case.

Lemma 4. *Let $\Phi_\infty : X_G \to X_H$ be a k-block code. If Φ collapses a diamond, then Φ_∞ is not injective.*

Proof. Suppose Φ collapses a diamond. That is, $\Phi(w_1 w_2 w_3) = \Phi(w_1 w_2' w_3)$ for some words w_1, w_3 of length k and distinct words w_2, w_2' in G. Consider any infinite word w_0 which can precede w_1 and any infinite word w_4 which can follow w_3. Then $\Phi_\infty(w_0w_1w_2w_3w_4) = \Phi_\infty(w_0w_1w_2'w_3w_4)$, so Φ_∞ is not injective. □

3 Verification: Testing a k-Block Map for Conjugacy

Given a pair of directed graphs G, H, and a proposed k-block map Φ, we wish to verify whether or not Φ induces a conjugacy between the vertex shifts X_G, X_H. We will focus in this section on the case $k = 1$, as the case $k > 1$ follows immediately by recoding to the kth higher block shift. When G and H are irreducible (strongly connected), this problem boils down to checking that the two graphs have the same number of cycles of each length up to some constant, and furthermore that Φ induces an injection on these cycles. While cycle counting can be done efficiently using powers of the adjacency matrices, the challenge remains of checking injectivity efficiently.

The reducible case, when G and H are not strongly connected, is much more complex. We give counter-examples to several statements which would have led to a straightforward algorithm wherein one subdivides the graphs into their irreducible components and uses the algorithm for the irreducible case on each, together with some other global checks. Instead, we give a more direct reduction to the irreducible case: we efficiently augment the graphs and block map with new vertices and edges, until the resulting graphs are irreducible, in such a way as to preserve conjugacy (or lack thereof).

3.1 Irreducible Case

As described above, we will focus first on 1-block codes. When G, H are irreducible, the following straightforward topological result allows us to restrict attention to the map induced on cycles between the graphs.

Proposition 5. *Suppose X, Y are compact metric spaces, $\psi : X \to Y$ is continuous, and $D \subseteq Y$ is a dense subset of Y. If ψ surjects onto D, then ψ surjects onto all of Y.*

Proof. Suppose X, Y are compact metric spaces, $\psi : X \to Y$ is continuous, $D \subseteq Y$ is a dense subset of Y, and $D \subseteq \psi(X) \subseteq Y$. Let $p \in Y$. Since D is dense, there is a sequence $\{p_n\}$ in D which converges to p. Since ψ surjects onto D, every p_n has a preimage in X. Pick $\gamma : D \to X$ such that $\psi \circ \gamma = \mathrm{id}_D$. Then $\{\gamma(p_n)\}$ is a sequence in X. Since X is compact, there is a subsequence $\{\gamma(p_{n_k})\}$ which converges to some $q \in X$. As limits commute with continuous functions, we have

$$\psi(q) = \psi(\lim_{n_k\to\infty} \gamma(p_{n_k})) = \lim_{n_k\to\infty} \psi(\gamma(p_{n_k})) = \lim_{n_k\to\infty} p_{n_k} = p.$$

Thus $\psi(q) = p$, so ψ is surjective on all of Y. □

We will applying Proposition 5 with D being the set of periodic points of X_H, which are in bijection with cycles of H. The following result, that Φ induces a 1-block conjugacy if and only if it induces a bijection on cycles, appears to be known; we give the proof for completeness.

Theorem 6. *Irreducible vertex shifts X_G, X_H are conjugate via a 1-block code if and only if there is a vertex map $\Phi : V_G \to V_H$ such that the induced map Φ_{c} is a bijection.*

Proof. If Φ_∞ is a conjugacy, then it is a bijection on periodic points; we conclude Φ_{c} is a bijection. For the converse, suppose Φ_∞ is not a conjugacy. We proceed in cases.

(Case 1) If Φ_∞ is not injective, there exist distinct points $p, q \in X_G$ such that $\Phi_\infty(p) = \Phi_\infty(q)$. (Case 1a) Suppose first that p, q disagree at $|V_G|^2 + 1$ consecutive indices, meaning the words $p_{[a,b]}, q_{[a,b]}$ disagree at every index for $a, b \in \mathbb{Z}$ with $b - a = |V_G|^2 + 1$. Consider all possible pairs of states in G; there are $|V_G|^2$ such pairs. Thus there exist distinct indices $c, d \in \{a, a+1 \ldots, b\}$ such that $(p_c, q_c) = (p_d, q_d)$. But then $\Phi_{\mathsf{c}}(p_{[c,d-1]}) = \Phi_{\mathsf{c}}(q_{[c,d-1]})$.

(Case 1b) Suppose instead that p, q do not disagree at $|V_G|^2 + 1$ consecutive indices: there exist indices a, b with $a < b - 1$ such that p, q agree at indices a and b, but p, q disagree at every index between a and b. Let w be any word connecting $p_b = q_b$ to $p_a = q_a$. Then $\Phi_{\mathsf{c}}(p_{[a,b]}w) = \Phi_{\mathsf{c}}(q_{[a,b]}w)$.

(Case 2) If Φ_∞ is not surjective, suppose for a contradiction that Φ_{c} is bijective. Then every periodic point in X_H is mapped to by Φ_∞. Since the periodic points are a dense subset of the compact metric space X_H, Proposition 5 contradicts the fact that Φ_∞ is not surjective. □

To verify that the cycle map Φ_{c} is bijective, we will test for injectivity explicitly, and rely on counting arguments to check surjectivity. For injectivity, it turns out that checking cycles up to length $|V_G|^2$ suffices.

Proposition 7. *Suppose $\Phi_\infty : X_G \to X_H$ is a 1-block code between irreducible vertex shifts. If Φ_{c} is injective on $\bigcup_{n=1}^{|V_G|^2} C_n(G)$, then Φ_{c} is injective.*

Proof. Let c, d be distinct cycles of size $|c| = |d| = k > |V_G|^2$. Proceeding by strong induction, suppose Φ_{c} is injective on all cycles of size less than k. There are $|V_G|^2$ possible pairs of states in G. Thus there exist distinct indices a, b such that $(c_a, d_a) = (c_b, d_b)$. That is, $c_{[a,b-1]}, d_{[a,b-1]}$ are cycles of the same length and $c_{[b,a-1]}, d_{[b,a-1]}$ are cycles of the same length. Since c, d were distinct, we can assume without loss of generality that $c_{[a,b-1]}, d_{[a,b-1]}$ are distinct. By the induction hypothesis, $\Phi_{\mathsf{c}}(c_{[a,b-1]}) \neq \Phi_{\mathsf{c}}(d_{[a,b-1]})$. Thus $\Phi_{\mathsf{c}}(c) \neq \Phi_{\mathsf{c}}(d)$. □

Proposition 7 suggests the naïve algorithm of checking all cycles up to length $|V_G|^2$ to verify injectivity of Φ_{c}. This algorithm is remarkably inefficient, however; letting $n = |V_G|$, there can be $\Omega(n^{n^2})$ cycles of length up to n^2, as is the case for the complete graph. Fortunately, these checks can be performed much more efficiently, by rephrasing them as a search problem in a graph built from pairs of vertices in G. This procedure is outlined in Algorithm 1.

Theorem 8. *Let X_G be a vertex shift and $A = \{1, 2, \ldots, m\}$. Then any given map $\Phi : V_G \to A$ induces a map $\Phi_{\mathsf{c}} : \bigcup\limits_n C_n(G) \to \bigcup\limits_n A^n$. Deciding if Φ_{c} is injective can be determined in $O(|V_G|^4)$ time.*

Proof. First we build the directed meta-graph $M = (V_M, E_M)$ where $V_M = \{(v_1, v_2) : v_1, v_2 \in V_G\}$ and $E_M = \{((v_1, v_2), (u_1, u_2)) : \Phi(v_1) = \Phi(u_1), \Phi(v_2) = \Phi(u_2), (v_1, u_1) \in E_G$, and $(v_2, u_2) \in E_G\}$. That is, M is a graph on pairs of vertices from G, with an edge connecting pairs P_1, P_2 if and only if (i) there is a pair of (possibly non-distinct) edges in G connecting the two vertices in P_1 to the vertices in P_2, and (ii) the induced map on words of length two (i.e., edges) maps the two edges together. M can be constructed in $O(|V_G|^4)$ time.

Given M, the map Φ_{c} is injective if an only if there is no cycle in M which passes through a vertex $(v_1, v_2) \in V_M$ with $v_1 \neq v_2$. Furthermore, such a cycle in M exists if and only if M has a strongly connected component containing an edge and a vertex (v_1, v_2) with $v_1 \neq v_2$. Tarjan's strongly connected components algorithm [8] now applies, in $O(|V_M| + |E_M|) = O(|V_G|^4)$ time. □

Putting the above results together with the higher-block codes gives the desired algorithm to verify k-block conjugacies; the full conjugacy algorithm for $k = 1$ is outlined in Algorithm 2.

Corollary 9. *Given a k-block code $\Phi_\infty : X_G^k \to X_H$ between irreducible vertex shifts, deciding if Φ_∞ is a conjugacy is in* P. *In particular, it can be determined in $O(|V_G|^{4k})$ time.*

Proof. Given G, H, we first pass to the kth higher block shift $X_{G^{[k]}}$ of X_G, recalling that $\Phi_\infty^{[k]}$ is a 1-block code and Φ_∞ is a conjugacy if and only if $\Phi_\infty^{[k]}$ is a conjugacy [6, Proposition 1.5.12]. We can construct $\Phi_\infty^{[k]} : X_{G^{[k]}} \to X_H$ in time $O(|V_{G^{[k]}}| + |E_{G^{[k]}}|) = O(|V_{G^{[k]}}|^3)$. Noting that $|V_{G^{[k]}}| \le |V_G|^k$, it thus suffices to show the case $k = 1$.

By Theorem 6, Φ_∞ is a conjugacy if and only if Φ_{c} is a bijection. As $k = 1$, Theorem 8 shows that injectivity of Φ_{c} can be determined in $O(|V_G|^4)$ time. To show Φ_{c} is surjective, it suffices to check that $|C_i(G)| = |C_i(H)|$ for all $i \in \mathbb{N}$. Letting $A(G), A(H)$ be the adjacency matrices of G, H, we note $|C_i(G)| = \operatorname{tr}(A(G)^i)$, so our desired check is equivalent to checking $\operatorname{tr}(A(G)^i) = \operatorname{tr}(A(H)^i)$ for all $i \in \mathbb{N}$ [6, Proposition 2.2.12]. In fact, it suffices to check up to $i = |V_G|$ [2, 5]. Calculating $\operatorname{tr}(A(G)^i), \operatorname{tr}(A(H)^i)$ for all $i \in \{1, \ldots, |V_G|\}$ can be done by repeated multiplication in $O(|V_G|^{1+\omega}) = O(|V_G|^4)$ time, where ω is the exponent of matrix multiplication. □

3.2 Reducible Case

Several useful statements about conjugacy between irreducible vertex shifts fail to hold in the reducible case. First, given a sliding block code $\Phi_\infty : X_G \to X_H$ between irreducible vertex shifts, it is known that if Φ_∞ is injective and G, H have the same topological entropy, then Φ_∞ is a conjugacy [6, Corollary 8.1.20]. (The topological entropy of a shift X is defined as $h(X) = \lim_{n\to\infty} \frac{1}{n} \log_2 |\mathcal{B}_n(X)|$.) If the shifts are reducible, however, Φ_∞ can satisfy these conditions but fail to be surjective (Figure 2a). Second, we have from Theorem 6 that if Φ_∞ is a 1-block code between irreducible vertex shifts, then Φ_{c} being a bijection implies Φ_∞ is bijective. In the reducible case, Φ_{c} can be bijective while Φ_∞ fails injectivity (Figure 2b) or surjectivity (Figure 2a).

As an even stronger test, one might guess for reducible vertex shifts that if $\Phi_\infty : X_G \to X_H$ is surjective and the induced maps between irreducible subgraphs are all conjugacies, then Φ_∞ is a conjugacy. If true, this statement would suggest applying the algorithm in Corollary 9 to each irreducible subgraph, at which point one would only need to test surjectivity. Yet this statement is also false; Φ_{c} being a bijection implies neither the injectivity nor the surjectivity of Φ_∞ (Figure 2). By extending the argument of Theorem 6, one can correct the statement by adding a check for diamonds: if Φ_∞ is surjective, the induced maps between irreducible subgraphs are all conjugacies, and Φ does not collapse a diamond, then Φ_∞ is a conjugacy. Unfortunately, while this revised statement does break the problem of verifying a proposed 1-block conjugacy into more manageable pieces, how to turn it into a decision procedure, let alone an efficient algorithm, is far from clear.

To verify a potential conjugacy between vertex shifts efficiently, we will instead apply a more direct reduction to the irreducible case. Given a 1-block code $\Phi_\infty : X_G \to X_H$ between reducible vertex shifts, we will extend G and H to irreducible graphs while preserving the conjugacy or non-conjugacy of Φ_∞. The key operation for this extension is the following procedure, which adds a new sink vertex to a sink component in such a way as to preserve conjugacy/non-conjugacy. We will then apply this procedure to every sink component, and in reverse to every source component, until we have enough structure to connect the new

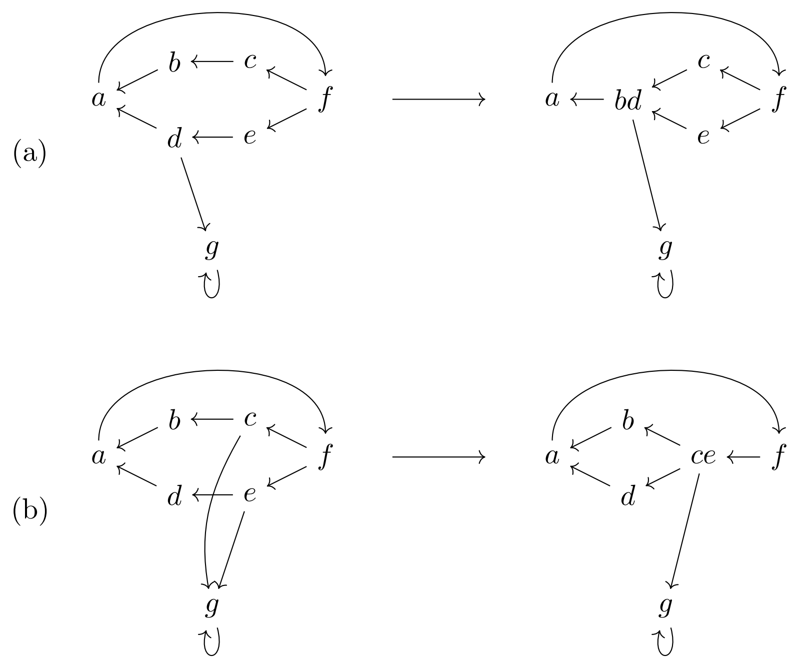

Figure 2: Counter-examples showing various statements which hold in the irreducible case fail in the reducible case. Note that all four shifts have the same topological entropy, $h(X) = \frac{1}{4}$. (a) A 1-block code between two reducible shifts which restricts to conjugacies between the irreducible components (and hence Φ_{c} is a bijection) but is not surjective. (b) A 1-block code between two reducible shifts which restricts to conjugacies between the irreducible components but is not injective.

sink verticies back to the new source vertices through a new vertex $*$, rendering both graphs irreducible.

Let T be a sink component of H and $T' = \Phi^{-1}(T)$ be the subgraph of G which maps to T under Φ_∞. The procedure is as follows:

1. Pick an arbitrary vertex v in T.
2. Pick an arbitrary cycle c in T ending at v of length $|c| \leq |T|$.
3. Add the vertex t along with the edges $(t,t),(v,t)$ to H. Call this new graph $\hat{H}$.
4. Select the vertices $v' \in \Phi^{-1}(v)$ which are followed by an infinite word w' such that $\Phi(v'w') = vc^\infty$. Call this set of vertices V'.
5. Add the vertex t' and the edges $\{(v',t') : v' \in V' \cup \{t'\}\}$ to G. Call this new graph $\hat{G}$.
6. Define $\hat{\Phi}_\infty : X_{\hat{G}} \to X_{\hat{H}}$ by $\hat{\Phi}(u) = \begin{cases} \Phi(u), & \text{if } u \neq t' \\ t, & \text{if } u = t' \end{cases}$.

Proposition 10. *Let $\Phi_\infty : X_G \to X_H$ be a 1-block code between reducible vertex shifts. Then $\hat{\Phi}_\infty : X_{\hat{G}} \to X_{\hat{H}}$ as described above is a conjugacy if and only if Φ_∞ is a conjugacy.*

Proof. Since X_G is a subshift of $X_{\hat{G}}$ (and similarly for H) and $\hat{\Phi}_\infty$ preserves Φ_∞, it immediately follows that Φ_∞ is a conjugacy whenever $\hat{\Phi}_\infty$ is. For the converse, suppose $\hat{\Phi}_\infty$ is not a conjugacy.

If $\hat{\Phi}_\infty$ is not injective, we have distinct points $p_1, p_2 \in X_{\hat{G}}$ such that $\hat{\Phi}_\infty(p_1) = \hat{\Phi}_\infty(p_2) = q$. If $q \in X_H$, then $p_1, p_2 \in X_G$ by the definition of $\hat{\Phi}_\infty$, so Φ_∞ is not injective. If $q \notin X_H$, then $q = wvt^\infty$. By the definition of $\hat{\Phi}_\infty$, we have $p_1 = w_1 v_1 t'^\infty, p_2 = w_2 v_2 t'^\infty$. By the construction of $N^-(t')$, there exist infinite words w_1', w_2' such that $v_1 w_1', v_2 w_2'$ are words in G and $\Phi(w_1') = c^\infty = \Phi(w_2')$. Thus $\Phi_\infty(w_1 v_1 w_1') = \Phi_\infty(w_2 v_2 w_2')$, and Φ_∞ is not injective.

If $\hat{\Phi}_\infty$ is not surjective, then there exists $p \in X_{\hat{H}}$ which is not mapped to. If $p \in X_H$, then Φ is not surjective. Otherwise, $p \notin X_H$, so $p = wvt^\infty$. But again noting the construction of $N^-(t)$, the point wvc^∞ is not mapped to. ◻

We now construct the final graphs G^*, H^* using the above procedure as well as one additional step below. Let $T_1, \ldots, T_m$ be the sink components of H, and $S_1, \ldots, S_\ell$ the source components, with $T_i' = \Phi^{-1}(T_i)$ and $S_i' = \Phi^{-1}(S_i)$ the corresponding inverse image subgraphs of G. We apply the above procedure iteratively to every sink component, and every source component by reversing the edge direction in G, H, applying the procedure, and reversing edges back. Let $\hat{G}, \hat{H}$ denote the graphs after applying the procedure to the m sinks and ℓ sources. Note that, by construction, each sink or source component in $\hat{H}$ contains a single state. Furthermore, the preimage of each of these states under the induced map $\hat{\Phi}$ in $\hat{G}$ also contains a single state. Denote the source states in $\hat{H}$ as $\{s_1, \ldots, s_\ell\}$ and the sink states as $\{t_1, \ldots, t_m\}$. In $\hat{G}$, denote the preimages as $s_i' = \Phi^{-1}(s_i), t_j' = \Phi^{-1}(t_j)$. We extend $\hat{H}, \hat{G}$ to the irreducible graphs H^*, G^* as follows:

1. Add a new vertex $*$ to both $\hat{H}$ and $\hat{G}$.
2. In H^*, set $N^-(*) = \{t_1, \ldots, t_m\}$, $N^+(*) = \{s_1, \ldots, s_\ell\}$. In G^*, set $N^-(*) = \{t_1', \ldots, t_m'\}$, $N^+(*) = \{s_1', \ldots, s_\ell'\}$.

3. Define $\Phi^*_\infty : X_{G^*} \to X_{H^*}$ by $\Phi^*(u) = \begin{cases} \hat{\Phi}(u), & \text{if } u \neq * \\ *, & \text{if } u = * \end{cases}$.

Proposition 11. *Let $\Phi_\infty : X_G \to X_H$ be a 1-block code between reducible vertex shifts. Then $\Phi^*_\infty : X_{G^*} \to X_{H^*}$ as described in the construction above is a conjugacy if and only if Φ_∞ is a conjugacy.*

Proof. By Proposition 10, $\hat{\Phi}_\infty : X_{\hat{G}} \to X_{\hat{H}}$ is a conjugacy if and only if Φ_∞ is a conjugacy, where as above the graphs $\hat{G}, \hat{H}$ immediately precede the addition of the vertex $*$. As in the proof of Proposition 10, $X_{\hat{G}}$ is a subshift of X_{G^*} (and similarly for $\hat{H}$) and Φ^*_∞ preserves Φ_∞, so $\hat{\Phi}_\infty$ is a conjugacy if Φ^*_∞ is. For the other direction, suppose Φ^*_∞ is not a conjugacy.

First suppose Φ^*_∞ is not injective. Since G^*, H^* are irreducible, we have cycles c, d in G^* from Theorem 6 such that $\Phi^*_{\mathsf{c}}(c) = \Phi^*_{\mathsf{c}}(d)$. Without loss of generality, c, d pass through $*$. Since Φ^* is bijective on $\{s_1, \ldots, s_\ell, t_1, \ldots, t_m\}$, we have $c = *s_i w_1 t_j, d = *s_i w_2 t_j$. But then $\hat{\Phi}$ collapses the diamond $(s_i w_1 t_j, s_i w_2 t_j)$, so by Lemma 4, $\hat{\Phi}_\infty$ is not injective.

Now suppose Φ^*_∞ is not surjective. Again by Theorem 6, we know there is a cycle c which is not in the image of Φ^*_{c}. Without loss of generality, we can assume $c = *s_i w t_j$. By the construction of $N^-(*), N^+(*)$, we conclude that $s_i w t_j$ is not in the image of $\hat{\Phi}$. Thus $s_i^\infty w t_j^\infty$ is a point in $X_{\hat{H}}$ which is not in the image of $\hat{\Phi}_\infty$. □

We now have that given reducible vertex shifts X_G, X_H and a proposed 1-block conjugacy between them, the shifts can be embedded into irreducible shifts such that the conjugacy or non-conjugacy is preserved. Next we show this embedding can be performed efficiently; the procedure described in the proof is outlined in Algorithm 5.

Theorem 12. *Given reducible vertex shifts X_G, X_H and a 1-block code as $\Phi : V_G \to V_H$, the graphs G^* and H^* can be constructed in $O(|V_G|^3)$ time.*

Proof. Let T be an arbitrary sink component in H and T' be the subgraph $\Phi^{-1}(T)$ of G. We will show the corresponding sink vertices t, t' can be added in $O(|V_{T'}|^3)$ time. Iterating over all sink components $T \in \mathcal{T}$ and source components $S \in \mathcal{S}$ will give an overall complexity of $O(\sum_{T' \in \mathcal{T}'} |V_{T'}|^3 + \sum_{S' \in \mathcal{S}'} |V_{S'}|^3) = O(|V_G|^3)$ time. (Adding the vertex $*$ takes linear time.)

Let v be an arbitrary vertex of T, and let c be the shortest cycle in T through v, which can be computed using breadth-first search in $O(|V_T| + |E_T|) = O(|V_H|^2) = O(|V_G|^2)$ time. Note that $|c| \leq |T|$, so we have completed steps 1 and 2. Step 3 is constant time. The only nontrivial step that remains is step 4, the computation of the set $V' \subseteq V_{T'}$, from which steps 5 and 6 follow trivially in linear time.

Let $C = (V_C, E_C)$ be the subgraph of T corresponding to c, and let $C' = (V_{C'}, E_{C'})$ be the subgraph of T' which maps onto C as follows: $V_{C'} = \Phi^{-1}(V_C)$, and $E_{C'} = \{(u', v') \in E_{T'} : (\Phi(u'), \Phi(v')) \in E_C\}$. The subgraph C' can be constructed in $O(|V_{T'}|^2)$ time. Note that infinite walks in C' starting from any $v' \in \Phi^{-1}(v)$ are precisely the walks in T' that map onto c^∞, and moreover, there is an infinite walk in C' starting from v' if and only if there is a path in C' from v' to a cycle in C'. We therefore define $V' \subseteq \Phi^{-1}(v) \subseteq V_{C'}$ to be the set of nodes v' such that there is a path in C' from v' to a cycle in C'. To compute V', we can simply run breadth-first search from each vertex in $\Phi^{-1}(v)$, in $O(|\Phi^{-1}(v)| \cdot (|V_{C'}| + |E_{C'}|)) = O(|V_{T'}|^3)$ time. □

We have now seen an efficient procedure to embed a pair of reducible graphs into a pair of irreducible graphs, such that the original pair admits a 1-block conjugacy if and only if the embedded pair does. Moreover, the embedded irreducible graphs have at most twice the number of vertices as the original graphs. With this procedure in hand, we can extend our verification algorithm to the reducible case.

Corollary 13. *Given vertex shifts X_G, X_H and a k-block code Φ_∞ as $\Phi : V_G^k \to V_H$, deciding if Φ_∞ is a conjugacy can be determined in $O(|V_G|^{4k})$ time.*

Proof. If G, H are irreducible, Corollary 9 applies immediately. For the reducible case, as in Corollary 9, by passing to the kth higher block shift it suffices to show the case $k = 1$. From Theorem 12, we can embed G, H into the irreducible shifts G^*, H^* in $O(|V_G|^3)$ time. Furthermore, $|V_{G^*}| < 2|V_G|$, so $|V_{G^*}| = O(|V_G|)$. Then by Corollary 9, we can verify if Φ^*_∞ (and hence Φ_∞) is a conjugacy in $O(|V_{G^*}|^4) = O(|V_G|^4)$. □

4 Deciding k-Block Conjugacy

We now turn to the question of deciding k-block conjugacy. Specifically, we wish to understand the complexity of the problem k-BC, which is to decide given directed graphs G, H whether the vertex shifts X_G, X_H are conjugate via a k-block code $\Phi_\infty : X_G \to X_H$. Note that the description size of Φ is polynomial in $|V_G|$ and $|V_H|$, and thus from Corollary 13 we know that a potential k-block conjugacy can be verified in polynomial time; hence, k-BC is in NP. We will show that k-BC is GI-hard for all k, where GI is the class of problems with a polynomial-time Turing reduction to the Graph Isomorphism problem [4]. (A graph isomorphism is bijection between the vertices of two graphs which preserves the edges/non-edge relation; the Graph Isomorphism problem is to decide if two given undirected graphs are isomorphic.)

Definition 14. Given directed graphs G, H, the *k-Block Conjugacy Problem*, denoted k-BC, is to decide if there is a k-block conjugacy $\Phi_\infty : X_G \to X_H$ between the vertex shifts X_G and X_H.

To begin, we give the straightforward result that the case $k = 1$ is GI-hard, essentially because 1-block conjugacies between vertex shifts for equal sized graphs must be isomorphisms.

Theorem 15. *The 1-Block Conjugacy Problem,* 1-BC*, is* GI*-hard.*

Proof. Given strongly connected graphs directed G, H with $|V_G| = |V_H|$, we show that the shifts X_G, X_H are conjugate via 1-block code if and only if the graphs are isomorphic (cf. [6, Ex. 2.2.14]). The result then follows as graph isomorphism between strongly connected directed graphs is GI-hard, by the usual reduction from the undirected case (replace each edge with two directed edges).

First suppose $\Psi : V_G \to V_H$ is a graph isomorphism. As $\Psi(v_1 v_2)$ is a legal word in X_H for all words of length 2, by definition of a graph isomorphism, we have that $\Psi_\infty : X_G \to X_H$ is a valid 1-block code. Letting $\Phi = \Psi^{-1} : V_H \to V_G$, we have $\Psi_\infty(\Phi_\infty((x_i)_{i\in\mathbb{Z}})) = (\Psi(\Phi(x_i)))_{i\in\mathbb{Z}} = (x_i)_{i\in\mathbb{Z}}$ for all $x \in X_H$, and $\Phi_\infty(\Psi_\infty((x_i)_{i\in\mathbb{Z}})) = (\Phi(\Psi(x_i)))_{i\in\mathbb{Z}} = (x_i)_{i\in\mathbb{Z}}$ for all $x \in X_G$. Thus, Φ_∞ is the 2-sided inverse of Ψ_∞, and Ψ_∞ is a 1-block conjugacy.

For the other direction, suppose $\Phi_\infty : X_G \to X_H$ is a 1-block conjugacy. Then $\{\Phi(v) : v \in V_G\}$ must be exactly the set of words of length 1 in X_H, i.e., the vertices of H. Since $|V_G| = |V_H|$, $\Phi : V_G \to V_H$ is a bijection. Also, for any edge $(v_1, v_2) \in E_G$, we have $\Phi(v_1 v_2) = \Phi(v_1)\Phi(v_2)$, so $(\Phi(v_1), \Phi(v_2)) \in E_H$ as Φ_∞ is a well-defined sliding block code. Even more, consider any pair v_3, v_4 of vertices in V_G such that $(v_3, v_4) \notin E_G$. Noting that $\Phi(v_3)\Phi(v_4)$ has a unique preimage as Φ is a bijection and Φ_∞ is surjective, we have $(\Phi(v_3), \Phi(v_4)) \notin E_H$.

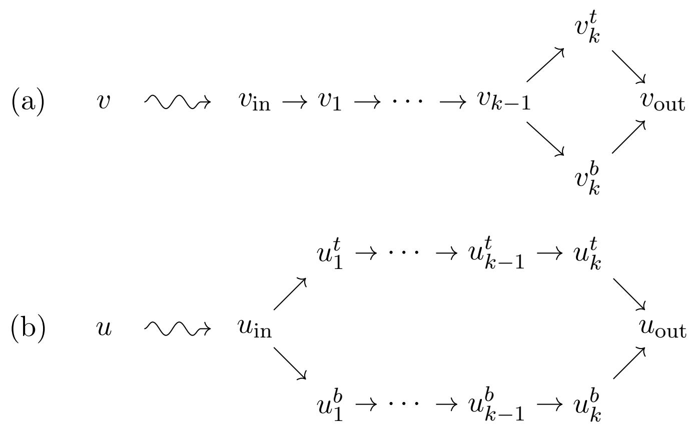

Figure 3: The vertex gadgets for (a) each vertex v in G, and (b) each vertex u in H.

Thus $\Phi : V_G \to V_H$ is a bijection on vertices which preserves the edge relationship; that is, Φ is a graph isomorphism. $\square$

Next, we will show k-BC is GI-hard for all k, by reduction to the 1-block case. Specifically, given directed graphs G, H, we will construct graphs G', H' such that there exists a 1-block conjugacy $\Phi_\infty : X_G \to X_H$ if and only if there exists a k-block conjugacy $\Phi'_\infty : X_{G'} \to X_{H'}$ exists. To form G', we replace every vertex $v \in V_G$ with a path $v_{\text{in}} v_1 v_2 \cdots v_{k-1}$ followed by the diamond with sides $v_{k-1} v_k^t v_{\text{out}}$ and $v_{k-1} v_k^b v_{\text{out}}$ (Figure 3a). To form H', we replace every vertex $u \in V_H$ with two parallel paths $u_{\text{in}} u_1^t u_2^t \cdots u_k^t u_{\text{out}}$ and $u_{\text{in}} u_1^b u_2^b \cdots u_k^b u_{\text{out}}$ (Figure 3b).

Lemma 16. *Given directed graphs G, H, let G', H' be constructed as above. If there exists a k-block conjugacy $\Phi'_\infty : X_{G'} \to X_{H'}$, then for all $v \in V_G$ there exists $u \in V_H$ such that $\Phi'(v_{in} v_1 \cdots v_{k-1}) = u_{in}$.*

Proof. Suppose for a contradiction that Φ'_∞ is a k-block code such that for some $v \in V_G$ we have $\Phi'(v_{\text{in}} v_1 \cdots v_{k-1}) \neq u_{\text{in}}$ for all $u \in V_H$. We break the argument into two cases.

First, suppose $\Phi'(v_{\text{in}} v_1 \cdots v_{k-1}) = u_i^t$. (The case u_i^b is identical.) Since the shift map commutes with sliding block codes, we must have $\Phi'(v_1 \cdots v_{k-1} v_k^t) = \Phi'(v_1 \cdots v_{k-1} v_k^b) = z$ where $z \in \{u_{i+1}^t, u_{\text{out}}\}$. Picking any edge $(v, \hat{v}) \in E_G$ and continuing to slide the block window, we must have $\Phi'(\hat{v}_{\text{in}} \hat{v}_1 \cdots \hat{v}_{k-1}) \in \{\hat{u}_i^t, \hat{u}_i^b\}$ for some $\hat{u} \in V_H$. Without loss of generality, assume $\Phi'(\hat{v}_{\text{in}} \hat{v}_1 \cdots \hat{v}_{k-1}) = \hat{u}_i^t$. Furthermore, since there is only one word in H' between u_i^t and $\hat{u}_i^t$ of proper length but two words in G' between $\hat{v}_{\text{in}}$ and v_{in}, we have

$$\Phi'(v_{\text{in}} \cdots v_k^t v_{\text{out}} \hat{v}_{\text{in}} \cdots \hat{v}_{k-1}) = u_i^t \cdots u_k^t u_{\text{out}} \hat{u}_{\text{in}} \hat{u}_1^t \cdots \hat{u}_i^t = \Phi'(v_{\text{in}} \cdots v_k^b v_{\text{out}} \hat{v}_{\text{in}} \cdots \hat{v}_{k-1}).$$

That is, Φ' collapses a diamond, so by Lemma 4, Φ' is not a conjugacy.

Second, suppose $\Phi'(v_{\text{in}} v_1 \cdots v_{k-1}) = u_{\text{out}}$. Pick any edge $(\hat{v}, v) \in E_G$. Then without loss of generality, $\Phi'(\hat{v}_{\text{out}} v_{\text{in}} \cdots v_{k-2}) = u_k^t$. Continuing to slide the block window, we have $\Phi'(\hat{v}_{\text{in}} \cdots \hat{v}_{k-1}) = \hat{u}_{\text{out}}$ for some $\hat{u} \in V_H$. Again, there are two words in G' between $\hat{v}_{\text{in}}$ and v_{in} but only one word in H' between $\hat{u}_{\text{out}}$ and u_{out} which passes through u_k^t. Thus, we have

$$\Phi'(\hat{v}_{\text{in}} \cdots \hat{v}_k^t \hat{v}_{\text{out}} v_{\text{in}} \cdots v_{k-1}) = \hat{u}_{\text{out}} u_{\text{in}} \cdots u_k^t u_{\text{out}} = \Phi'(\hat{v}_{\text{in}} \cdots \hat{v}_k^b \hat{v}_{\text{out}} v_{\text{in}} \cdots v_{k-1}),$$

so Φ' again collapses a diamond, and by Lemma 4, Φ' is not a conjugacy. $\square$

We now show that graphs G, H admit a 1-block conjugacy if and only if the graphs G', H' constructed as above admit a k-block conjugacy. To do this, we first introduce a natural operation on shift spaces, which "stretches" each point by a factor N. Given alphabet $\mathcal{A}$, and any point $p = \cdots v.v' \cdots \in \mathcal{A}^{\mathbb{Z}}$, we write $p^{\langle N \rangle} = \cdots v \cdots v.v' \cdots v' \cdots$ to be the point p with each symbol repeated N times. Given a shift X over alphabet $\mathcal{A}$, we define the shift space $X^{\langle N \rangle} = \{\sigma^i(p^{\langle N \rangle}) : p \in X, i \in \mathbb{Z}\}$ where σ is the shift map. In particular, $X^{\langle N \rangle}$ contains all shifts of the Nth expansion of points in X. While in general $X^{\langle N \rangle}$ is not a vertex shift when $N > 1$, it is still structured enough that the following lemma is immediate.

Lemma 17. *Given shifts X, Y, there exists a 1-block conjugacy $\Phi_\infty : X \to Y$ if and only if there exists a 1-block conjugacy $\Phi_\infty^{\langle N \rangle} : X^{\langle N \rangle} \to Y^{\langle N \rangle}$, where $\Phi = \Phi^{\langle N \rangle}$ as block maps.*

To make use of this definition and lemma, we will project points in $X_{G'}, X_{H'}$ to $X_G^{\langle k+2 \rangle}, X_H^{\langle k+2 \rangle}$ by simply erasing the subscript and superscript information. Formally, we define the 1-block map $\Psi^G : V_{G'} \to V_G$ by $\Psi^G(u) = v$ for $u \in \{v_{\text{in}}, v_1, \ldots, v_{k-1}, v_k^t, v_k^b, v_{\text{out}}\}$, and let $\pi^G = \Psi^G_\infty : X_{G'} \to X_G^{\langle k+2 \rangle}$. We define Ψ^H, π^H similarly. Letting $S_p^G := \left(\pi^G\right)^{-1}(p) \subseteq X_{G'}$, we have that $\{S_p^G : p \in X_G^{\langle k+2 \rangle}\}$ is a partition of the points in $X_{G'}$. (Similarly for S_q^H and $X_{H'}$.)

Theorem 18. *Given graphs G, H, construct G', H' as above. Then there exists a 1-block conjugacy $\Phi_\infty : X_G \to X_H$ if and only if there exists a k-block conjugacy $\Phi' : X_{G'} \to X_{H'}$.*

Proof. ($\Rightarrow$) Suppose there exists a 1-block conjugacy $\Phi_\infty : X_G \to X_H$. By Lemma 17, there is a 1-block conjugacy $\Phi_\infty^{\langle k+2 \rangle} : X_G^{\langle k+2 \rangle} \to X_H^{\langle k+2 \rangle}$. Define the k-block code $\Phi'_\infty : X_{G'} \to X_{H'}$ with no memory by

- $\Phi'(v_{\text{in}} \cdots) = \Phi(v)_{\text{in}}$
- $\Phi'(v_i \cdots v_k^t \cdots) = \Phi(v)_i^t$, $i \in \{1, \ldots, k\}$
- $\Phi'(v_i \cdots v_k^b \cdots) = \Phi(v)_i^b$, $i \in \{1, \ldots, k\}$
- $\Phi'(v_{\text{out}} \cdots) = \Phi(v)_{\text{out}}$

To show that Φ'_∞ is a bijection, we will show that for any $p \in X_G^{\langle k+2 \rangle}$ and $q \in X_H^{\langle k+2 \rangle}$ with $\Phi_\infty^{\langle k+2 \rangle}(p) = q$, the map $\Phi'_\infty : S_p^G \to S_q^H$ is a bijection. The result then follows because $\Phi_\infty^{\langle k+2 \rangle}$ is a bijection between $X_G^{\langle k+2 \rangle}$ and $X_H^{\langle k+2 \rangle}$, and the sets $\{S_p^G : p \in X_G^{\langle k+2 \rangle}\}, \{S_q^H : q \in X_H^{\langle k+2 \rangle}\}$ partition $X_{G'}, X_{H'}$.

We first claim that $\Phi'_\infty(S_p^G) \subseteq S_q^H$, which is to say, for every $p' \in X_{G'}$ such that $\pi^G(p') = p$, we have $\pi^H(\Phi'_\infty(p')) = q$. To see this, note that by construction of Φ', for all $p' \in X_{G'}$ and all $i \in \mathbb{Z}$, we have $\Psi^H(\Phi'_\infty(p')_i) = \Phi(\Psi^G(p'_i)) = \Phi^{\langle k+2 \rangle}(\Psi^G(p'_i))$. The condition $\pi^G(p') = \Psi^G_\infty(p') = p$ implies $\Psi^G(p'_i) = p_i$ for all $i \in \mathbb{Z}$. Combining the above with the observation that $\Phi^{\langle k+2 \rangle}(p_i) = q_i$ gives $\Psi^H(\Phi'_\infty(p')_i) = q_i$, which implies the claim.

To see that Φ'_∞ is injective on S_p^G, consider distinct points $p', p'' \in S_p^G$ which differ at index i. Since $\pi^G(p') = \pi^G(p'')$, we can assume without loss of generality that $p'_i = v_k^t$ and $p''_i = v_k^b$. Then

$$\Phi'(p'_{[i-k,i]}) = \Phi'(v_1 \cdots v_{k-1} v_k^t) = \Phi(v)_1^t \neq \Phi(v)_1^b = \Phi'(v_1 \cdots v_{k-1} v_k^b) = \Phi'(p''_{[i-k,i]}),$$

so $\Phi'_\infty(p') \neq \Phi'_\infty(p'')$.

($\Leftarrow$) Suppose there exists a k-block conjugacy $\Phi'_\infty : X_{G'} \to X_{H'}$. Without loss of generality, assume Φ'_∞ has no memory. By Lemma 16, for every $v \in V_G$, there exists $u \in V_H$ such that

$$
\begin{tikzcd}
X_{G'} \arrow[r, "\Phi'_\infty"] \arrow[d, "\pi^G=\Psi^G_\infty"'] & X_{H'} \arrow[d, "\pi^H=\Psi^H_\infty"] \\
X_G^{\langle k+2\rangle} \arrow[r, "\Phi_\infty^{\langle k+2\rangle}"] & X_H^{\langle k+2\rangle}
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
p' \arrow[r, mapsto, "\Phi'_\infty"] \arrow[d, mapsto, "\pi^G"'] & q' \arrow[d, mapsto, "\pi^H"] \\
p \arrow[r, mapsto, "\Phi_\infty^{\langle k+2\rangle}"] & q
\end{tikzcd}
$$

Figure 4: Given Φ' or Φ, one can construct the other such that this diagram commutes.

$\Phi'(v_{\text{in}}\cdots) = u_{\text{in}}$. Define the 1-block code $\Phi_\infty : X_G \to X_H$ by $\Phi(v) = \Psi^H(\Phi'(v_{\text{in}}v_1\cdots v_{k-1}))$. We claim Φ_∞ is a conjugacy. To show this, we instead will show $\Phi_\infty^{\langle k+2\rangle}$ defined by the same block map is a conjugacy. To see $\Phi_\infty^{\langle k+2\rangle}$ is surjective, consider any $q \in X_H^{\langle k+2\rangle}$. Picking any $q' \in S_q^H$, set $p' = {\Phi'}_\infty^{-1}(q')$ and $p = \pi^G(p')$. We will now show $\Phi_\infty^{\langle k+2\rangle}(p) = q$, so the diagram in Figure 4 is commutative and Φ_∞ is surjective. For all $i \in \mathbb{Z}$,

$$\Phi(p_i) = \Psi^H(\Phi'((p_i)_{\text{in}}\cdots(p_i)_{k-1})) = \Psi^H((q_i)_{\text{in}}) = q_i. \tag{1}$$

To see $\Phi_\infty^{\langle k+2\rangle}$ is injective, consider distinct $p_1, p_2 \in X_G^{\langle k+2\rangle}$. Then $S_{p_1}^G \neq S_{p_2}^G$. Since Φ'_∞ is a conjugacy, $\Phi'_\infty(S_{p_1}^G) \cap \Phi'_\infty(S_{p_2}^G) = \emptyset$. By Lemma 16, there exist $q_1, q_2 \in X_H^{\langle k+2\rangle}$ such that $S_{q_1}^H = \Phi'_\infty(S_{p_1}^G)$ and $S_{q_2}^H = \Phi'_\infty(S_{p_2}^G)$. By the construction of Φ (and shown in (1)), $\Phi_\infty^{\langle k+2\rangle}(p_1) = q_1 \neq q_2 = \Phi_\infty^{\langle k+2\rangle}(p_2)$, so $\Phi_\infty^{\langle k+2\rangle}$ is injective. $\square$

The construction in Theorem 18 gives a polynomial-time reduction from 1-BC to k-BC for all k. (The same construction could plausibly give a reduction from m-BC to ℓ-BC where $\ell = (m-1)(k+2)+k$, though if true the proof would be much more involved.) Combining this reduction with Theorem 15 therefore gives GI-hardness for all k.

Corollary 19. *k-BC is GI-hard for all k.*

5 Reducing Representation Size

Thus far we have addressed two problems. We first gave an efficient algorithm, given directed graphs G, H and k-block map Φ, to verify whether $\Phi_\infty : X_G \to X_H$ is a conjugacy. We then showed that the problem of deciding whether X_G and X_H are conjugate, given only G and H, is GI-hard. We now address a problem given only G and an integer ℓ: whether we can find a k-block code which reduces the size of G by ℓ vertices while preserving conjugacy.

Definition 20. Given a directed graph G and integer ℓ, the *k-Block Reduction Problem*, denoted k-BR, is to decide if there exists a directed graph H with $|V_H| = |V_G| - \ell$ such that the vertex shifts X_G and X_H are conjugate via a k-block code.

We will show this problem is NP-complete for the case $k = 1$, by modifying the hardness proof of the State Amalgamation Problem (SAP), which asks if ℓ consecutive amalgamations can be performed on a graph G [1]. The proof that SAP is NP-hard shows that the set of graphs satisfying a certain structure property is closed under the amalgamation operation. This structure is then leveraged to encode an NP-complete problem (Hitting Set). While 1-block codes are more general than sequences of amalgamations (Figure 1), we find that, surprisingly, the same set of graphs is also closed under 1-block conjugacy. In fact, the rest of the construction of [1] suffices as well, though much of the argument needs to be strengthened to the general 1-block case.

We begin by recalling the structure property.

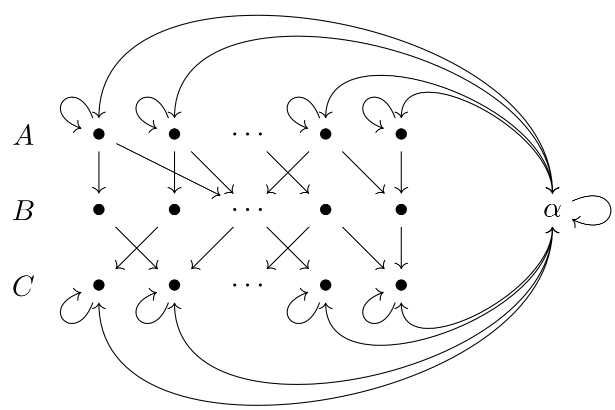

Figure 5: A graph which satisfies the structure property

Definition 21 ([1])**.** A directed graph G satisfies the *structure property* if it is essential and there exists a partition $\{\{\alpha\}, A, B, C\}$ of V_G such that the following four conditions hold.

1. $N^+(\alpha) = N^-(\alpha) = \{\alpha\} \cup A \cup C$.
2. For each $a \in A$, $N^-(a) = \{a, \alpha\}$ and $\{a, \alpha\} \subseteq N^+(a) \subseteq \{a, \alpha\} \cup B$.
3. For each $c \in C$, $N^+(c) = \{c, \alpha\}$ and $\{c, \alpha\} \subseteq N^-(c) \subseteq \{c, \alpha\} \cup B$.
4. For each $b \in B$, $N^-(b) \subseteq A$ and $N^+(b) \subseteq C$.

See Figure 5 for an example. We now show that the structure property is preserved under 1-block conjugacy.

Lemma 22. *Let G be a graph with the structure property having $\{\{\alpha\}, A, B, C\}$ as the partition of V_G, and let $\Phi_\infty : X_G \to X_H$ be a 1-block conjugacy. Then $\Phi(v_1) = \Phi(v_2)$ implies $v_1 = v_2$ or $v_1, v_2 \in B$, so H also satisfies the structure property with vertex partition $\{\{\Phi(\alpha)\}, \Phi(A), \Phi(B), \Phi(C)\}$.*

Proof. First note that if $\Phi : X_G \to X_H$ is a 1-block conjugacy from a graph G with vertex partition $\{\{\alpha\}, A, B, C\}$ such that $\Phi(v_1) = \Phi(v_2)$ implies $v_1 = v_2$ or $v_1, v_2 \in B$, then the fact that H satisfies the structure property with vertex partition $\{\{\Phi(\alpha)\}, \Phi(A), \Phi(B), \Phi(C)\}$ follows immediately. Now suppose for a contradiction that $v_1 \neq v_2 \in V_G$ and $\Phi(v_1) = \Phi(v_2)$; we proceed in cases.
Case 1: $v_1, v_2 \in \{\alpha\} \cup A \cup C$. Then $\Phi_\infty(v_1^\infty) = \Phi_\infty(v_2^\infty)$ and Φ_∞ is not a conjugacy.
Case 2: $v_1 = \alpha, v_2 \in B$. Let $a \in A, c \in C$ be such that av_2c is a word in X_G. (Such a, c exist as G is essential.) Then $\Phi_\infty((av_2c\alpha)^\infty) = \Phi_\infty((a\alpha c\alpha)^\infty)$ and Φ_∞ is not a conjugacy.
Case 3a: $v_1 \in A, v_2 \in B, (v_1, v_2) \notin E_G$. Let $a \in A$ be such that $(a, v_2) \in E_G$. Note that $a \neq v_1$. Consider the point

$$p = (\Phi(a)\Phi(v_2)\Phi(\alpha))^\infty = (\Phi(a)\Phi(v_1)\Phi(\alpha))^\infty$$

in X_H of period 3. Due to G having the structure property and our assumption that Φ_∞ is a 1-block conjugacy, the preimage of p must be defined by a 3-cycle whose vertices are contained in $\{\alpha\} \cup A \cup C$. In particular, the preimage must trace a self-loop, so we know $\Phi(a) = \Phi(\alpha)$ or $\Phi(a) = \Phi(v_1)$ or $\Phi(v_1) = \Phi(\alpha)$. Since we know Φ is injective on $\{\alpha\} \cup A \cup C$ by Case 1, none of these are possible.
Case 3b: $v_1 \in A, v_2 \in B, (v_1, v_2) \in E_G$. Let $c \in C$ be such that $(v_2, c) \in E_G$. Consider the point

$$p = (\Phi(v_2)\Phi(c)\Phi(\alpha))^\infty = (\Phi(v_1)\Phi(c)\Phi(\alpha))^\infty$$

in X_H of period 3. Again by the requirement that the preimage of p traces a self-loop, we know $\Phi(v_1) = \Phi(c)$ or $\Phi(v_1) = \Phi(\alpha)$ or $\Phi(\alpha) = \Phi(c)$. However, all of these situations violate the injectivity of Φ on $\{\alpha\} \cup A \cup C$.
Case 4: $v_1 \in C, v_2 \in B$. This is identical to Case 3 where the edges in the graph have been reversed. ☐

As in [1], we will need a "weight widget" which acts as a weighted switch, using the following notation. Let v be a vertex with $N^-(v) = D$ and $N^+(v) = E$. We will write $v : [D, E]$ in this situation, and as a slight abuse of notation, we will drop the curly brackets if E or D is a singleton and write $v : [u, E]$. Additionally, we extend this notation to sets S of vertices in the obvious way and write $S : [D, E]$.

Definition 23 ([1])**.** Let G satisfy the structure property with $V_G = \{\alpha\} \cup A \cup B \cup C$, and let $K > 0$ be a fixed even integer. Then for nonempty subsets $A_* \subseteq A, C_* \subseteq C$, the *weight widget* $w = \mathsf{weight}[A_*, C_*]$ is the following collection of vertices.

- $A_w = \{a_1^w, \ldots, a_{K/2}^w\}$
- $B_w = \{b_1^w \ldots, b_K^w\}$
- $C_w = \{c_1^w, \ldots, c_{K/2}^w\}$

where $A_w \cap A_* = \emptyset = C_w \cap C_*$, and for each $i \in \{1, \ldots, K/2\}$ we have

- $b_{2i-1} : [A_* \cup \{a_1^w, \ldots, a_{i-1}^w\}, c_i^w]$
- $b_{2i} : [a_i^w, C_* \cup \{c_1^w, \ldots, c_i^w\}]$.

Moreover, we require these to be the only images of A_w in B, i.e., $B \cap N^+(a_i^w) \subseteq B_w$ for all $a_i^w \in A_w$, and similarly for the preimage of C_w. For a given 1-block conjugacy Φ_∞, letting $S = \Phi^{-1}(\Phi(b_1^w)) \setminus B_w = \{b \in B : \Phi(b) = \Phi(b_1^w)\} \setminus B_w$, we say w is *activated* if $S : [A_*, C_*]$.

See Figure 6 for an example. The term "activate" comes from the following fact, which we show below in Lemma 25(1): if S is a singleton, then the construction of the weight widget allows the states $S \cup B_w$ to be amalgamated sequentially into a single state. For example, the vertex v in Figure 6 can activate the weight widget shown. The next two lemmas show that these amalgamations cannot be performed if the widget is not activated.

Lemma 24. *Let $w = \mathsf{weight}[A_*, C_*]$ be a weight widget in G. If $\Phi_\infty : X_G \to X_H$ is a 1-block conjugacy between graphs with the structure property, then for any $v \in V_H$, the statement $b_\ell^w \in \Phi^{-1}(v)$ for $\ell > 1$ implies $b_{\ell-1}^w \in \Phi^{-1}(v)$ or $|\Phi^{-1}(v)| = 1$.*

Proof. By contrapositive, suppose $|\Phi^{-1}(v)| > 1$ and there exists $b_\ell^w \in \Phi^{-1}(v)$ such that $b_{\ell-1}^w \notin \Phi^{-1}(v)$. Without loss of generality, let ℓ be the largest such subscript. We have two cases.

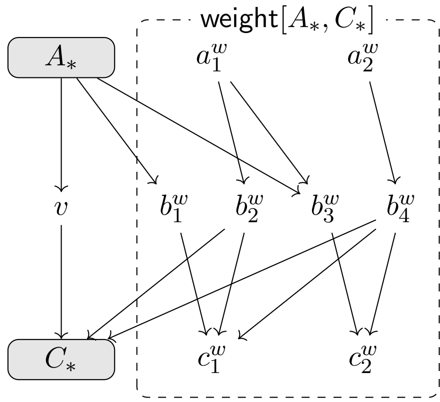

Figure 6: The weight widget $\mathsf{weight}[A_*, B_*]$ with $K = 4$.

Case 1: ℓ is even. We claim there must exist $a \in N^-(v) \setminus \{\Phi(a^w_{\ell/2}), \dots, \Phi(a^w_{K/2})\}$. To see this claim, first note the weight widget construction guarantees that for every $b \in B \setminus B_w$, we have $N^-(b) \cap A_w = \emptyset$. Also by the weight widget construction, for any $b^w_{2i+1} \in B_w$, we have $A_* \subseteq N^-(b^w_{2i+1}) \setminus A_w$. Noting that $\mathsf{weight}[A_*, B_*]$ is not defined when $A_* = \emptyset$ and G is not essential when $N^-(b) = \emptyset$, if we have $\Phi^{-1}(v) \setminus \{b^w_{2i} : b^w_{2i} \in B_w\} \neq \emptyset$, then we must have $N^-(v) \setminus \Phi(A_w) \neq \emptyset$. That is, the claim is satisfied in the case when $\Phi^{-1}(v) \not\subseteq \{b^w_{2i} : b^w_{2i} \in B_w\}$. Assuming now that $\Phi^{-1}(v) \subseteq \{b^w_{2i} : b^w_{2i} \in B_w\}$, we note our earlier assumption that $|\Phi^{-1}(v)| > 1$ guarantees there exists $b^w_{2j} \in \Phi^{-1}(v)$ for some index $2j \neq \ell$. By our other assumption that ℓ is the largest such subscript, we have $2j < \ell$. Then $\Phi(a^w_j) \in N^-(v)$, and the claim follows. Proceeding, we then have $\Phi(a)v\Phi(c^w_{\ell/2})$ is a word in X_H; however, $N^-(c^w_{\ell/2}) = \{b^w_{\ell-1}\} \cup \{b^w_\ell, b^w_{\ell+2}, \dots, b^w_K\}$ and for all $b^w_{2i} \in N^-(c^w_{\ell/2}) \setminus \{b^w_{\ell-1}\}, (a, b^w_{2i}) \notin E_G$. Since $b^w_{\ell-1} \notin \Phi^{-1}(v)$ by assumption, this word has no preimage in X_G, so Φ_∞ is not a conjugacy.
Case 2: ℓ is odd. Using an argument symmetric to the one in case 1, we get that there must exist $c \in N^+(v) \setminus \{\Phi(c^w_{(\ell+1)/2}), \dots, \Phi(c^w_{K/2})\}$. Then $\Phi(a^w_{(\ell-1)/2})v\Phi(c)$ is a word in X_H with no preimage, so Φ_∞ is not a conjugacy. $\square$

Lemma 25. *Suppose $w = \mathsf{weight}[A_*, C_*]$ is a weight widget in G. Then*

1. *Suppose $\Phi_\infty : X_G \to X_{G'}$ is a 1-block conjugacy such that $\Phi^{-1}(\Phi(b^w_i)) = \{b^w_i\}$ for all $b^w_i \in B_w$ and $V_{G'}$ contains $v : [A_*, B_*]$. Defining $\Phi'_\infty : X_G \to X_H$ by*

$$\Phi'(u) = \begin{cases} \Phi(u), & \text{if } u \notin \Phi^{-1}(v) \cup B_w \\ v, & \text{if } u \in \Phi^{-1}(v) \cup B_w \end{cases}$$

   where H is the minimal graph induced by G and Φ', then Φ' is a 1-block conjugacy with $|V_H| = |V_{G'}| - K$.

2. *If $w = \mathsf{weight}[A_*, C_*]$ is not activated and $\Phi_\infty : X_G \to X_H$ is a 1-block conjugacy, then $\Phi^{-1}(\Phi(b^w_i))$ is a singleton for every b^w_i with $i > 1$.*

Proof. (1) Consider the sequence of splittings followed by the sequence of amalgamations which transforms G into G'. Then note that by the construction of the weight widget, $\{v : [A_*, C_*]\} \cup B_w$ can be amalgamated sequentially for an additional K amalgamations.

(2) Suppose $b^w_1 \in \Phi^{-1}(v)$ for some $v \in V_H$ and consider $V = \Phi^{-1}(v) \setminus B_w$. By Lemma 24 it suffices to show $b^w_2 \notin \Phi^{-1}(v)$. By definition of w not being activated, we have two cases.

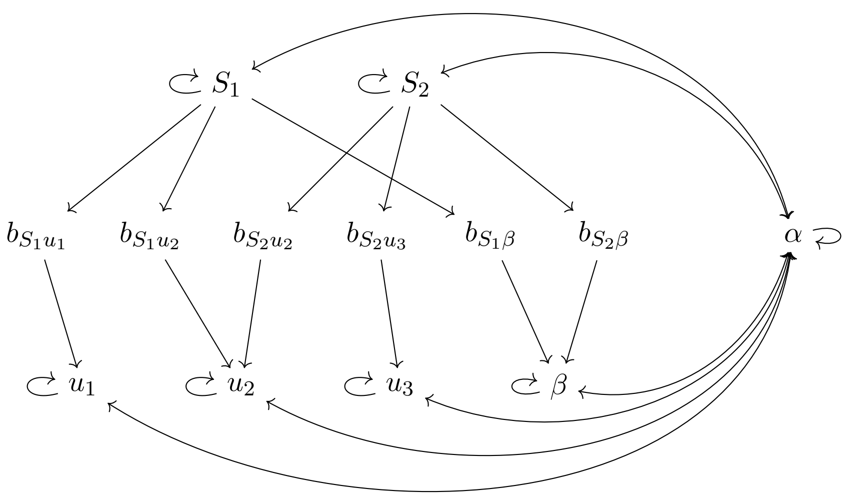

Figure 7: The graph constructed in Theorem 28 for the HittingSet instance with $\mathcal{S} = \{\{u_1, u_2\}, \{u_2, u_3\}\}$, without any weight widgets attached.

Case 1: $N^-(V) \neq A_*$. If there is some $a \in N^-(V) \setminus A_*$, then $\Phi(a)v\Phi(c_1^w)$ is a word in X_H. Since there is no state in G connecting a with c_1^w, the word has no preimage in X_G and Φ_∞ is not a conjugacy. Otherwise, there is some $a \in A_* \setminus N^-(V)$. By contrapositive, suppose $b_2^w \in \Phi^{-1}(v)$. Picking any $c \in C_*$, we have $\Phi(a)v\Phi(c)$ is a word in X_H. Since there is no state in $\Phi^{-1}(v)$ connecting a with c, the word has no preimage and Φ_∞ is not a conjugacy.
Case 2: $N^+(V) \neq C_*$. By contrapositive, suppose $b_2^w \in \Phi^{-1}(v)$. If there is some $c \in N^+(V) \setminus C_*$, then $\Phi(a_1^w)v\Phi(c)$ is a word in X_H. Since there is no state in G connecting a_1^w with c, the word has no preimage and Φ_∞ is not a conjugacy. Otherwise, there is some $c \in C_* \setminus N^+(V)$. Considering any $a \in N^-(V)$, we have $\Phi(a)v\Phi(c)$ is a word in X_H. Since there is no state in $\Phi^{-1}(v)$ connecting a with c, the word has no preimage and Φ_∞ is not a conjugacy. □

We now define the Hitting Set problem, which is NP-complete [3], and state a lemma which we will need in the proof.

Definition 26. Let $\mathcal{S} = \{S_1, \ldots, S_m\}$ be a collection of sets with $\bigcup_i S_i = U$. Given a subset $S \subseteq U$, we define its *hit set* as $\mathrm{hit}(S) = \{S_i : S \cap S_i \neq \emptyset\}$. Given $\mathcal{S}, U$, and an integer t, the *hitting set problem*, denoted HittingSet, is to decide whether there is a set H of cardinality t such that $\mathrm{hit}(H) = \mathcal{S}$. We will also overload this notation, and write $\mathrm{hit}(s)$ to mean $\mathrm{hit}(\{s\})$ for $s \in U$.

Lemma 27 ([1])**.** *Let $(\mathcal{S}, U, t)$ be an instance of* HittingSet*. Suppose for some $t \leq |\mathcal{S}|$ there is no H with $|H| \leq t$ and $hit(H) = \mathcal{S}$. Then for all $H \subseteq U$, $|hit(H)| - |H| < |\mathcal{S}| - t$.*

We now show that 1-BR is NP-complete, by reduction from HittingSet. Given the lemmas developed above, the result essentially follows from the argument in [1], with minor modifications for the 1-block case; for completeness, we give the full proof.

Theorem 28. 1-BR *is* NP*-complete.*

Proof. First we show 1-BR is in NP. Given a vertex shift X_G and $\Phi : V_G \to \{1, 2, \ldots, |V_G| - n\}$ from a proposed 1-block conjugacy Φ_∞, we construct the minimal image graph G' such that $\Phi_\infty : X_G \to X_{G'}$ is well-defined. In particular, $V_{G'} = \{\Phi^{-1}(u) : u \in V_G\}$ and $E_{G'} = \{(\Phi^{-1}(v_1), \Phi^{-1}(v_2)) : (v_1, v_2) \in E_G\}$. By Corollary 13, we can determine if Φ_∞ is a conjugacy in $O(|V_G|^4)$ time.

To show hardness, we reduce from HittingSet; let $\mathcal{S} = \{S_1, \ldots, S_m\}$ be the collection of sets and t the given integer. Defining $n = |U|$ for $U = \bigcup_i S_i$, we set the parameter $K = 5mn$ for the weight widgets. Then, as in [1], we build the following graph $G = (V_G, E_G)$ with the structure property $V_G = A \cup B \cup C \cup \{\alpha\}$.

1. Start with $A = \mathcal{S}, B = \emptyset, C = U \cup \{\beta\}$, where β is a new vertex.
2. For each $u \in A, v \in C$, add $b_{uv} : [u, v]$. That is, add the vertex $b_{S_i s}$ and path $S_i \to b_{S_i s} \to s$ for every (s, S_i) with $s \in S_i$ as well as the vertex $b_{S_i \beta}$ and path $S_i \to b_{S_i \beta} \to \beta$ for every $S_i \in \mathcal{S}$.
3. For each (s, S_i) with $s \in S_i$, add the weight widget $w = \mathsf{weight}[S_i, \{s, \beta\}] = (A_w, B_w, C_w)$. Note that A_w will added to A, B_w to B, and C_w to C.
4. For each $s \in U$, add the weight widget $\mathsf{weight}[\mathrm{hit}(s), \{s\}]$.
5. Finally, add the vertex α and the necessary edges for G to have the structure property, i.e., add the edges $\{(a, \alpha), (\alpha, a) : a \in A\} \cup \{(b, \alpha), (\alpha, b) : b \in B\} \cup \{(v, v) : v \in A \cup B \cup \{\alpha\}\}$.

Summarizing, if W is the collection of weight widgets added, $A = \mathcal{S} \cup \bigcup_{w \in W} A_w$, $B = \{b_{S_i s} : S_i \in \mathcal{S}, s \in S_i\} \cup \{b_{S_i \beta} : S_i \in \mathcal{S}\} \cup \bigcup_{w \in W} B_w$, and $C = U \cup \{\beta\} \cup \bigcup_{w \in W} C_w$. (See Figure 7 for an example with $\mathcal{S} = \{\{u_1, u_2\}, \{u_2, u_3\}\}$ where only steps (1), (2), and (5) have been performed.)

We will show there is a hitting set of size t if and only if there is a 1-block conjugacy $\Phi_\infty : X_G \to X_{G'}$ such that $|V_{G'}| \le |V_G| - (m + n - t)K$. The idea behind the reduction is that s can either choose to be in the hitting set by combining some $b_{S_i s}$ with the appropriate $b_{S_i \beta}$ to activate some of the $\mathsf{weight}[S_i, \{s, \beta\}]$, or choose not to be in the hitting set by combining all $b_{S_i s}$ for $S_i \in \mathrm{hit}(s)$ to activate $\mathsf{weight}[\mathrm{hit}(s), \{s\}]$. We will be able to activate $|\mathrm{hit}(H)| + |U \setminus H| = m + (n - t)$ weight widgets if there is a hitting set of size t and strictly fewer if no such set exists. By our choice of K, any reduction in the number of vertices not caused by activating weight widgets will be insignificant.

First, suppose there is a hitting set H for $\mathcal{S}$ of size t. We will give a sequence of $(m + n - t)K$ consecutive amalgamations, which together constitute a 1-block reducing the number of vertices by $(m + n - t)K$. For each $S_i \in \mathcal{S}$, pick some $s \in H$ such that $S_i \in \mathrm{hit}(s)$. After amalgamating $b_{S_i s}$ with $b_{S_i \beta}$, the weight widget $w = \mathsf{weight}[S_i, \{s, \beta\}]$ can be activated and B_w amalgamated sequentially. Doing this for each S_i gives a total of $m(K+1) \ge mK$ consecutive amalgamations. As the above amalgamations only affected the vertices in B associated with H, next consider any $s \in U \setminus H$. We can amalgamate the vertices $\{b_{S_i s} : S_i \in \mathrm{hit}(s)\}$ in any order to form $b_{\mathrm{hit}(s)s} : [\mathrm{hit}(s), \{s\}]$ which can then activate $\mathsf{weight}[\mathrm{hit}(s), \{s\}]$. Amalgamating all the vertices in this weight widgets give a total of at least K amalgamations for each vertex $s \in U \setminus H$. Thus we can perform $mK + (n - t)K = (m + n - t)K$ consecutive amalgamations, so there is a 1-block conjugacy $\Phi_\infty : X_G \to X_{G'}$ such that $|V_G| \ge |V_{G'}| - (m + n - t)K$.

Next suppose there is no hitting set H of size t. Let $\Phi : X_G \to X_{G'}$ be a 1-block conjugacy such that $N = |V_G| - |V_{G'}|$ is as large as possible. Define

$$\begin{aligned}\overline{H} &= \{s \in U : \mathsf{weight}[\mathrm{hit}(s), \{s\}] \text{ is activated}\},\\ F &= \{S_i : \mathsf{weight}[S_i, \{s, \beta\}] \text{ is activated for some } s \in S_i\},\\ H &= U \setminus \overline{H}.\end{aligned}$$

Note that there is a single path in G from S_i to s, through the vertex $b_{S_i s}$, which is required to activate both $\mathsf{weight}[\mathrm{hit}(s), \{s\}]$ and $\mathsf{weight}[S_i, \{s, \beta\}]$. Thus for every $b_{S_i s}$ we have that if $S_i \in F$, then $s \in H$. That is,

$$F \subseteq \{S_i : s \in H \text{ for some } b_{S_i s}\}. \tag{2}$$

We now count how much smaller $|V_{G'}|$ could be than $|V_G|$. By construction, each activated widget can lead to reducing the number of vertices by at most K. Let $B_{\text{non-weight}} = \{b_{S_i s} : S_i \in \mathcal{S}, s \in S_i\} \cup \{b_{S_i \beta} : S_i \in \mathcal{S}\}$ be the vertices in B not in weight widgets. By Lemma 25, if $u \in \Phi^{-1}(v)$ with $|\Phi^{-1}(v)| > 1$ for some u not in an activated widget, then $u \in B_{\text{non-weight}} \cup \bigcup_{w \in W} w_1$. Thus V_G can be reduced by at most $(|F| + |\overline{H}|)K + |B_{\text{non-weight}}| + |W|$. Since

$$|B_{\text{non-weight}}| + |W| = (mn + m) + (mn + n) < K,$$

we have

$$\begin{aligned} N &\le (|F| + |\overline{H}|)K + |B_{\text{non-weight}}| + |W| \\ &< (|F| + |\overline{H}|)K + K \\ &\le (|\mathrm{hit}(H)| + (n - |H|) + 1)K && \text{(by (2) and } H = U \setminus \overline{H}) \\ &\le (m + n - t)K && \text{(by Lemma 27).} \end{aligned}$$

□

6 Edge Shifts

Thus far we have restricted our attention to vertex shifts, rather than edge shifts, though the latter are perhaps more commonly used in the literature. For various reasons, the problems we consider are in general more appropriate for vertex shifts, as we discuss in the following section. (Vertex shifts are also motivated by applications (§ 7).) Nonetheless, we now give some results for edge shifts, for the first two problems: verifying k-block conjugacies, and testing pairs of shifts for conjugacy. (The third problem remains open.)

In the following, we will leverage our results for vertex shifts, using the standard conversion from edge shifts to vertex shifts: edges become vertices, and pairs of adjacent edges become edges [6, Proposition 2.3.9]. More formally, we recall that given edge shift X^e_G, its vertex shift representation is the shift $X_{G'}$ where $V_{G'} = E_{G'}$ and $E_{G'} = \{(e_i, e_j) : e_i e_j \text{ is a word in } X^e_G\}$. Thus, for any edge shifts X^e_G, X^e_H, there exists a k-block conjugacy $\Phi_\infty : X^e_G \to X^e_H$ if and only if there exists a k-block conjugacy $\Phi'_\infty : X_{G'} \to X_{H'}$ between the vertex shift representations of X_G and X_H.

First, we observe that our verification algorithm for vertex shifts immediately applies to edge shifts.

Theorem 29. *Given directed multigraphs G, H and a proposed k-block code $\Phi_\infty : X^e_G \to X^e_H$, deciding if Φ_∞ is a conjugacy can be determined in $O(|E_G|^{4k})$.*

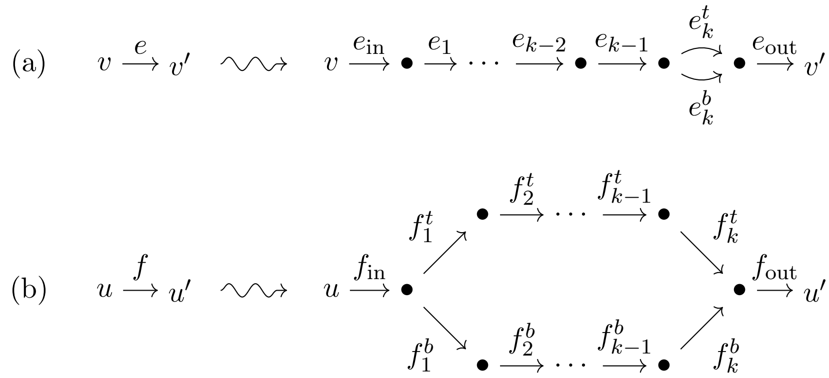

Figure 8: (a) The edge gadget for each pre-image graph. (b) The edge gadget for each image graph.

Proof. Given edge shifts X^e_G, X^e_H, we first construct their vertex shift representations $X_{G'}, X_{H'}$ as above. Letting Φ'_∞ be the corresponding block code between the vertex shifts, by Corollary 13, we can determine if Φ'_∞ is a conjugacy in $O(|V_{G'}|^{4k}) = O(|E_G|^{4k})$ time. $\square$

We now turn to the k-block conjugacy problem for edge shifts, where we again show GI-hardness.

Definition 30. Given directed mult-graphs G, H, the *k-Block Conjugacy Problem*, denoted k-BC^e, is to decide is there is a k-block conjugacy $\Phi_\infty : X^e_G \to X^e_H$ between the edge shifts X^e_G, X^e_H.

Theorem 31. *k-BC^e is GI-hard.*

Proof. We first show that 1-BC^e is GI-hard. Given directed graphs G, H with $|E_G| = |E_H|$, as in the vertex shift case, we will argue that there exists a 1-block conjugacy between the edge shifts if and only if the graphs are isomorphic. Suppose first that G, H are isomorphic. Let G', H' be the directed graphs for vertex shifts, as described above, so that $X_{G'} = X^e_G$ and $X_{H'} = X^e_H$. Since G, H are isomorphic and G', H' are created by the same (deterministic) procedure, G', H' are isomorphic. By Theorem 15, there exists a 1-block conjugacy $\Phi_\infty : X_{G'} \to X_{H'}$, so $X_{G'} = X^e_G$ is conjugate to $X_{H'} = X^e_H$ via a 1-block code.

Now suppose $\Phi_\infty : X^e_G \to X^e_H$ is a 1-block conjugacy. Noting that $\Phi : E_G \to E_H$ is a map on edges, we show that Φ can be realized as a map on V_G. To do this, it suffices to show (i) for any two edges $(v_1, v_2), (v_1, v_3)$ starting at the same vertex, $\Phi((v_1, v_2)), \Phi((v_1, v_3))$ also start at the same vertex and (ii) for any two edges $(u_1, u_2), (u_3, u_2)$ ending at the same vertex, $\Phi((u_1, u_2)), \Phi((u_3, u_2))$ also end at the same vertex. To see condition (i), consider any $(v_4, v_1) \in E_G$. As $\Phi((v_4, v_1)(v_1, v_2)), \Phi((v_4, v_1)(v_1, v_3))$ must both be words in X^e_H, we must have that $\Phi((v_1, v_2)), \Phi((v_1, v_3))$ both start at the same vertex. Similarly, for condition (ii), consider any $(u_2, u_4) \in E_G$, and note that $\Phi((u_1, u_2)(u_2, u_4)), \Phi((u_3, u_2)(u_2, u_4))$ are both words in X^e_H, so $\Phi(u_1, u_2), \Phi(u_3, u_2)$ must end at the same vertex. Thus Φ can be realized as a map $\Psi : V_G \to V_H$ on vertices which is surjective and preserves the edge/non-edge relation. To show Ψ is actually a graph isomorphism, consider the inverse Φ^{-1}_∞. Since Φ_∞ is 1-block conjugacy and $|E_G| = |E_H|$, Φ^{-1}_∞ is a 1-block code. Again, Φ^{-1}_∞ can be realized as a surjective vertex map Ψ' which preserves the edge/non-edge relation. Since both $\Psi : V_G \to V_H$ and

	Block size	Verification (G, H, Φ)	Conjugacy (G, H)	Reduction (G, ℓ)
Vertex	$k = 1$	1-BV: P	1-BC: GI-hard, NP	1-BR: NP-complete
	$k > 1$	k-BV: P	k-BC: GI-hard, NP	k-BR: NP-complete??
Edge	$k = 1$	1-BV^e: P	1-BC^e: GI-hard, NP*	1-BR^e: NP-complete??
	$k > 1$	k-BV^e: P	k-BC^e: GI-hard, NP*	k-BR^e: NP-complete??

Table 1: Summary of results and open questions, for vertex and edge shifts. Question marks denote conjectures, and BV refers to the verification problem (§ 3). The asterisk (*) denotes a subtlety in edge shift representations: the k-block conjugacy problem is in NP when the the representation size is considered to be the number of edges (i.e., a unary representation), but membership in NP is not clear when the shift is given as an adjacency matrix (i.e., a binary representation).

$\Psi' : V_H \to V_G$ are surjective maps between finite sets, we actually have Ψ, Ψ' are bijections. Thus Ψ is a graph isomorphism from G to H.

We now reduce k-BC^e to 1-BC^e, as we did with vertex shifts. Given edge shifts X^e_G, X^e_H, construct $\hat{G}, \hat{H}$ as follows. To form $\hat{G}$, substitute each edge in G with a path of length k followed by two parallel edges and a final edge (Figure 8a). Construct $\hat{H}$ by substituting each edge in H with a single edge followed by two parallel paths of length k followed by a single edge (Figure 8b). Then construct the vertex shift representations $G', H', \hat{G}', \hat{H}'$ of $G, H, \hat{G}, \hat{H}$. By construction of the edge gadget, $\hat{G}', \hat{H}'$ can be formed from G', H' by using the vertex gadget in Figure 3. Thus by Theorem 18, there exists a 1-block conjugacy $\Phi'_\infty : X_{G'} \to X_{H'}$ if and only if there exists a k-block conjugacy $\hat{\Phi}'_\infty : X_{\hat{G}'} \to X_{\hat{H}'}$. Since there exists a k-block conjugacy $\Phi_\infty : X^e_G \to X^e_H$ between edge shifts if and only if there exists a k-block conjugacy $\Phi'_\infty : X_{G'} \to X_{H'}$ between the vertex representations, there exists a 1-block conjugacy $\Phi_\infty : X_G \to X_H$ if and only if there exists a k-block conjugacy $\hat{\Phi} : X_{\hat{G}} \to X_{\hat{H}}$. Thus k-BC^e is 1-BC^e-hard and, in particular, GI-hard. □

7 Discussion

We have addressed several variants of the conjugacy problem restricted to k-block codes, with new algorithms to verify a proposed conjugacy, and hardness results for k-block conjugacy and representation reduction via 1-block codes (Table 1). Below we discuss subteties of input representation, followed by applications and open problems.

Representations of SFTs. When considering how to describe a subshift of finite type (SFT), three representations come to mind: a vertex shift, an edge shift, and a list of forbidden words $\mathcal{F}$. As our results pertain to vertex and edge shifts, we now discuss some nuances in these two representations, leaving lists of forbidden words to future work.

Perhaps the central advantage of edge shifts over vertex shifts is their compact representation size: a shift on n symbols can be represented in size as small as $O(\log n)$ by writing the multi-graph as a integer adjacency matrix, as opposed to $\Omega(n)$ for vertex shifts. This compact representation size can have important implications on the computational complexity. In the verification problem, for example, writing down a k-block code Φ naïvely takes $\Omega(n) = \Omega(|E_G|)$ space, which can be exponential in the size of the graphs G, H. (One can improve this by encoding Φ as a integer $|V_G| \times |V_G| \times |E_H|$ tensor, specifying how many $(u, v) \in E_G$ edges map to a given $e \in E_H$, but this can still be exponential.) Thus, while our algorithm remains polynomial-time, it would not be for cases allowing a compact representation of Φ.

Similarly, for the conjugacy problem, we only know k-BC^e to be in NP if we consider the graphs G, H to be represented in adjacency list form, which takes $\Omega(|E_G|)$ space, rather than the typically more compact integer adjacency matrix form taking $O(|V_G| \log |E_G|)$ space, as the natural certificate is the block map Φ witnessing the conjugacy. For the matrix representation of edge shifts, membership in NP would require a certificate exponentially smaller than the naïve representation of the block map Φ.

Finally, what "size reduction" means for edge shifts depends on the choice of adjacency list or matrix above. For the adjacency list, we have that the problem of reducing the number of vertices in the graph is in NP, but it is less motivated, as the size is dominated by $|E_G|$, not $|V_G|$. On the other hand, while the adjacency matrix representation size is dominated by $|V_G|$, it is not clear whether the problem of reducing the number of vertices is in NP, for the same reason as above.

Motivation from Markov partitions. As noted in [1], variants of the conjugacy problem for vertex shifts have applications in simplifying Markov partitions, a tool to study discrete-time dynamical systems via symbolic dynamics. Briefly, a Markov partition is a collection C of regions of the phase space, satisfying certain properties, which induces a conjugacy to a vertex shift X_G where $G = (C, E)$, i.e., the vertices are labeled with the regions of the phase space. In applications, one can encounter Markov partitions with thousands of regions, thus motivating the problem of simplifying the partition. Without additional information about the dynamical system, essentially the only way to do this while preserving the relevant geometric information is to *coarsen* the partition, by replacing sets of regions with a single region which is their union. This operation is exactly a 1-block code. Our results therefore give an efficient algorithm to test whether a proposed coarsening (1-block code) is valid (yields a conjugacy). Our results also imply that the problem of minimizing the partition size is NP-complete. (Previous work [1] only showed the latter for the case where the 1-block code was a sequence of amalgamations.)

Open problems. Our work leaves several open problems, such as those implied by Table 1: resolving the complexity of the k-block conjugacy problem, and showing NP-hardness of the size reduction problem. The complexity of deciding k-block conjugacy between edge shifts represented as integer matrices is especially interesting, as membership in NP is perhaps unlikely (see above). Regarding the k-block conjugacy problem and resolving where it lies on the spectrum between GI-complete and NP-complete, we conjecture that, similar to the induced subgraph isomorphism problem [7], it is GI-complete when $|V_G| = |V_H|$ and NP-complete when $|V_G| - |V_H|$ is large enough. Beyond these questions, it would be interesting to address the complexity of k-block conjugacy between SFTs given as lists of forbidden words, and the natural variants of the problem for that input (for example, reducing the representation size of the list).

Acknowledgments

We thank Luke Meszar for valuable contributions to the early stages of this work. We also thank Mike Boyle, Josh Grochow, Doug Lind, and Brian Marcus, for several helpful conversations, references, and insights.

References

[1] Rafael M. Frongillo. Optimal state amalgamation is NP-hard. *Ergodic Theory and Dynamical Systems*, 39(7):1857–1869, 2019.

[2] Shui-Hung Hou. Classroom note: A simple proof of the Leverrier–Fadeev characteristic polynomial algorithm. *SIAM Review*, 40(3):706–709, 1998.

[3] Richard M Karp. *Reducibility among combinatorial problems.* Springer, 1972.

[4] Johannes Kobler, Uwe Schöning, and Jacobo Torán. *The graph isomorphism problem: its structural complexity.* Springer Science & Business Media, 2012.

[5] Urbain Le Verrier. Sur les variations séculaires des éléments des orbites pour les sept planétes principales. 5:220–254, 1840.

[6] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding.* Cambridge University Press, 1999.

[7] Maciej M. Syso. The subgraph isomorphism problem for outerplanar graphs. *Theoretical Computer Science*, 17(1):91 – 97, 1982.

[8] Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2):146–160, 1972.

[9] R. F. Williams. Classification of subshifts of finite type. *Annals of Mathematics*, 98(1):120–153, 1973.

A Algorithms

Function `IsInjective`(G, H, Φ):

  Input: irreducible graphs G, H and a 1-block code Φ

  Output: true, if $\Phi_{\mathsf{c}} : \bigcup_n C_n(G) \to \bigcup_n C_n(H)$ is injective; false, otherwise

  `/* Construct the meta-graph` M `*/`

  $V_M \leftarrow V_G \times V_G$;

  $E_M \leftarrow \{((v_1, v_2), (u_1, u_2)) : \Phi(v_1) = \Phi(u_1), \Phi(v_2) = \Phi(u_2), \text{ and } (v_1, u_1), (v_2, u_2) \in E_G\}$;

  `/* Decide if` M `has a cycle passing through` (v_1, v_2) `with` $v_1 \neq v_2$ `*/`

  $\mathcal{S} \leftarrow$ `GetStronglyConnectedComponents`(M) ; `/* Tarjan's */`

  foreach *subgraph s in $\mathcal{S}$* **do**

    if *s is a singleton* **then**

      continue;

    end

    foreach *vertex (v_1, v_2) in s* **do**

      if $v_1 \neq v_2$ **then**

        return *true*;

      end

    end

  end

  return *false*;

Algorithm 1: Determine if Φ_{c} is injective

```
Function IsConjugacyIrreducible(G, H, Φ):
    Input: irreducible graphs G, H and a 1-block code Φ
    Output: true, if Φ_∞ is a conjugacy; false, otherwise
    if not IsInjective(G, H, Φ) then
        return false;
    end
    for i ∈ {1, ..., |V_G|} do
        if tr(A(G)^i) ≠ tr(A(H)^i) then
            return false;
        end
    end
    return true;
```

Algorithm 2: Determine if Φ_∞ between irreducible graphs is a conjugacy

```
Function AddSinkComponents(G, H, Φ):
    Input: reducible graphs G, H and a 1-block code Φ
    Result: (1) alters G, H so each sink component T in H has the property
            |V_{Φ^{-1}(T)}| = 1, and (2) extends Φ to the new graphs so Φ_∞ : X_G → X_H is
            a conjugacy if and only if the original 1-block code was a conjugacy
    𝒯 ← GetSinkComponents(H);
    foreach subgraph T in 𝒯 do
        T′ ← Φ^{-1}(T);
        if |V_{T′}| = 1 then
            continue;
        end
        /* Find the subgraphs C and C′                                   */
        v ← GetRandomVertex(T);
        c ← GetShortestCycleStartingAt(v);
        V_C ← {u ∈ V_T : u ∈ c};
        E_C ← {(u, u′) : uu′ is a word of length 2 contained in c^∞};
        V_{C′} ← {u ∈ V_{T′} : Φ(u) ∈ V_C};
        E_{C′} ← {(u, u′) ∈ E_{T′} : (Φ(u), Φ(u′)) ∈ E_C};
        /* Attach the new vertices t and t′                              */
        V_G.Add(t′);
        N^+(t′) ← {t′};
        N^-(t′) ← {t′};
        foreach vertex u in V_{C′} do
            if Φ(u) = v ∧ there is a path in C′ from u to a cycle then
                N^-(t′).Add(u);
            end
        end
        V_H.Add(t);
        N^+(t) ← {t};
        N^-(t) ← {t, v};
        Φ(t′) ← t;
    end
```

Algorithm 3: Turn every sink component into a single vertex

Function `AddSourceComponents(`G, H, Φ`):`

Input: reducible graphs G, H and a 1-block code Φ

Result: (1) alters G, H so each source component S in H has the property $|V_{\Phi^{-1}(S)}| = 1$, and (2) extends Φ to the new graphs so $\Phi_\infty : X_G \to X_H$ is a conjugacy if and only if the original 1-block code was a conjugacy

G.`ReverseEdges();`
H.`ReverseEdges();`
`AddSinkComponents(`G, H, Φ`);`
G.`ReverseEdges();`
H.`ReverseEdges();`

Algorithm 4: Turn every source component into a single vertex

Function `IsConjugacyReducible(`G, H, Φ`):`

Input: reducible graphs G, H and a 1-block code Φ

Output: true, if Φ_∞ is a conjugacy; false, otherwise

`AddSinkComponents(`G, H, Φ`);`
`AddSourceComponents(`G, H, Φ`);`
V_G.`Add(`v_G`);`
$N^-(v_G) \leftarrow$ `GetSinkVertices(`G`);`
$N^+(v_G) \leftarrow$ `GetSourceVertices(`G`);`
V_H.`Add(`v_H`);`
$N^-(v_H) \leftarrow$ `GetSinkVertices(`H`);`
$N^+(v_H) \leftarrow$ `GetSourceVertices(`H`);`
$\Phi(v_G) \leftarrow v_H$;
return `IsConjugacyIrreducible(`G, H, Φ`);`

Algorithm 5: Determine if Φ_∞ between reducible graphs is a conjugacy