

Embedding in Switching Classes with Skew Gains

Andrzej Ehrenfeucht², Jurriaan Hage¹, Tero Harju³, and Grzegorz Rozenberg^{2,4}

¹ Inst. of Information and Computing Sci., Univ. Utrecht, P.O.Box 80.089, 3508 TB Utrecht, Netherlands, jur@cs.uu.nl

² Dept. of Computer Science, University of Colorado at Boulder
Boulder, CO 80309, U.S.A.

³ Dept. of Mathematics, University of Turku, FIN-20014 Turku, Finland

⁴ Leiden Institute of Advanced Computer Science
P.O. Box 9512, 2300 RA Leiden, The Netherlands

Abstract. In the context of graph transformation we look at the operation of switching, which can be viewed as an elegant method for realizing global transformations of (group-labelled) graphs through local transformations of the vertices.

Various relatively efficient algorithms exist for deciding whether a graph can be switched so that it contains some other graph, the query graph, as an induced subgraph in case vertices are given an identity. However, when considering graphs up to isomorphism, we immediately run into the graph isomorphism problem for which no efficient solution is known. Surprisingly enough however, in some cases the decision process can be simplified by transforming the query graph into a 'smaller' graph without changing the answer. The main lesson learned is that the size of the query graph is not the dominating factor, but its cycle rank.

Although a number of our results hold specifically for undirected, unlabelled graphs, we propose a more general framework and give some preliminary results for more general cases, where the graphs are labelled with elements of a group.

1 Introduction

The material in this paper is motivated by a quest for techniques which enable the analysis of certain networks of processors. Our starting point is that the vertices of a directed graph can be interpreted as processors in a network and the edges can be interpreted as the channels/connections between them, labelled with values from some (structured) set, call it Δ , to capture the current state. The dynamics of such a network lies in the ability to change the labellings of the graph which is done by operations performed by the processors. A major aspect of the model here presented is that if a processor performs an input action, it influences the labellings

of all incoming edges in the same way; the same holds the output actions which govern the outgoing edges. In other words, we have no separate control over each edge, only over each processor. On the other hand, actions done by different processors should not interfere with each other, making this model an asynchronous one.

Ehrenfeucht and Rozenberg set forth in [3] a number of axioms they thought should hold for such a network of processors.

- A1 Any two input (output) actions can be combined into one single input (output) action.
- A2 For any pair of elements $a, b \in \Delta$, there is an input action that changes a into b ; the same holds for output actions.
- A3 For any channel (i, j) , the order of applying an input action to i and an output action to j is irrelevant.

Although each processor i was to have a set of *output actions* Ω_i and a set of *input actions* Σ_i , in [3] (see also [2]) it was derived that under these axioms the input (output) actions of every vertex are the same and form a group. Also, the sets of input and output actions coincide, but an action will act differently on incoming and outgoing edges, as evidenced by the asymmetry in (2) in Section 2. The difference is made explicit by an anti-involution δ , which is an anti-automorphism of order at most two on the group of actions. The notion of anti-involution generalizes that of group inversion. The result of this will be that if a channel between processors i and j is labelled with a , then the channel from j to i will be labelled with $\delta(a)$. The model generalizes the gain graphs of [9] and the voltage graphs of [4].

As we shall see later the graphs labelled with elements from a fixed group Δ (and under some fixed anti-involution of that group), called skew gain graphs in the following, are partitioned into equivalence classes. These equivalence classes capture the possible outcomes of performing actions in the vertices, i.e., the states of the system arrivable from a certain 'initial' state. The transformation from one skew gain graph to another, is governed by selecting in each vertex an operation, which corresponds to an element of the group. Although the equivalence classes themselves are usually considered static objects, it is not hard to see that there is also a notion of change or dynamics: applying a selector to a skew gain graph yields a new skew gain graph on the same underlying network of processors, but possibly with different labels. For this reason the equivalence classes were called dynamic labelled 2-structures in [3].

Consider now the problem where we have a (target) skew gain graph h which represents our network, and a skew gain graph g , the query

graph, which represents a fragment of a network which to us has a special meaning, for instance, it describes a deadlock situation. A question to ask is then: is there a way to transform h by applying a selector, such that in the result we can detect a subgraph similar to g ? In terms of the example: is there a possible state in the system, derivable from h which contains a deadlocked subgraph. If the embedding from g into h is known, then this can be (in many cases) efficiently solved by applying the results of Hage [5]. However, the large number of possible embeddings of g into h remains a problem. In fact, we quickly run into the Graph Isomorphism problem which does not have a known efficient solution. In this paper, we seek to alleviate this problem by seeing how we might reduce the skew gain graph g to a different, simpler graph without changing the outcome, i.e. if the reduced graph can be embedded, then so can g .

After introducing our notation for groups, skew gain graphs and switching classes thereof, we continue by formulating a general framework for reasoning about reductions between skew gain graphs, and give a some illustrative and even surprising examples of such reductions. In some cases they work irrespective of the group and involution. In the case of bridging on the other hand, where we can 'shorten' the lengths of cycles in our query graph, they generally work only for certain groups. We give examples of these for the group \mathbf{Z}_2 and for the group \mathbf{Z}_3 . The correctness of these reductions follow from rather surprising combinatorial results. We then show how these results can be used to derive an algorithm for the embedding problem, showing that the complexity of the embedding problem depends on the cycle rank of the query graph and not on the number of vertices. Finally, we prove some impossibility results for bridging.

2 Preliminaries

In this paper, we use both elementary group theory and graph theory. In this section we establish notation, and introduce the concept of switching classes of gain graphs with skew gains. For more details on group theory we refer the reader to Rotman [7].

For a group Γ we denote its identity element by 1_Γ . Let Γ be a group. A function $\delta : \Gamma \rightarrow \Gamma$ is an *anti-involution*, if it is an anti-automorphism of order at most two, that is, δ is a bijection and for all $x, y \in \Gamma$, $\delta(xy) = \delta(y)\delta(x)$ and $\delta^2(x) = x$. We write (Γ, δ) for a group Γ with a given anti-involution δ .

Define $E_2(V) = \{(u, v) \mid u, v \in V\}$, the set of nonreflexive, directed edges over V . We usually write uv for the edge (u, v) , but note $uv \neq vu$. For an edge $e = uv$, the *reverse* of e is $e^{-1} = vu$

We consider graphs $G = (V, E)$ where the set of edges $E \subseteq E_2(V)$ satisfies the following *symmetry condition*:

$$\text{if } e \in E \text{ then also } e^{-1} \in E.$$

Such graphs can be considered as undirected graphs where the edges have been given a two-way orientation. We use $E(G)$ to denote the edges E of G and similarly $V(G)$ to denote its vertices V .

Two vertices $v, v' \in V(G)$ are *adjacent* in G if $(v, v') \in E(G)$. The *degree* of a vertex is the number of vertices in the graph it is adjacent to. A vertex of degree zero is called *isolated*, a *leaf* has degree one, a *chain vertex* degree two, and all other vertices are called *dense vertices*.

A sequence of vertices $p = (v_1, \dots, v_k)$, $k > 0$, is a *path* in G if v_i is adjacent to v_{i+1} for $i = 1, \dots, k-1$ and all vertices are distinct. By $E(p)$ we denote the set of edges $\{(v_1, v_2), \dots, (v_{k-1}, v_k)\}$. Additionally, p is called a *chain* if all vertices v_2, \dots, v_{k-1} are chain vertices. The chain p is maximal in G if the endpoints v_1 and v_k are not chain vertices. A *cut edge* in a graph is an edge which is not on any cycle.

Let $G = (V, E)$ be a graph and (Γ, δ) a group with anti-involution. A pair (G, g) where g is a mapping $g : E \rightarrow (\Gamma, \delta)$ into the group Γ is called a (Γ, δ) -*gain graph* (on G) (or a *graph with skew gains* or a *skew gain graph*), if g satisfies the following *reversibility condition*

$$g(e^{-1}) = \delta(g(e)) \quad \text{for all } e \in E. \quad (1)$$

In the future we will refer to a skew gain graph (G, g) simply by g unless confusion arises. We adopt in a natural way some of the terminology of graph theory for graphs with skew gains. For instance, every path in G is also a path in g , and we can use $E(g)$ to denote the set of edges of the underlying graph G .

The class of (Γ, δ) -gain graphs on G will be denoted by $\mathbf{L}_G(\Gamma, \delta)$ or simply by \mathbf{L}_G . More importantly, $\mathbf{L}(\Gamma, \delta) = \bigcup \{\mathbf{L}_G(\Gamma, \delta) \mid G \text{ a graph}\}$. A *gain graph* is a $(\Gamma, -1)$ -gain graph; these are also called *inversive skew gain graphs*.

With a path $p = (v_1, \dots, v_k)$ in $g \in \mathbf{L}_G(\Gamma, \delta)$ we can associate the sequence of labels $\lambda(p) = (g(v_1v_2), \dots, g(v_{k-1}v_k))$. Now, p is an a -path if every value in $\lambda(p)$ is equal to a . Secondly, p is a b -summing path for some $b \in \Gamma$ if $g(v_1v_2) \cdot g(v_2v_3) \cdot \dots \cdot g(v_{k-1}v_k)$ equals b . (We often denote this fact by writing $g(p) = b$.) In other words, evaluating the product of

values found along p using the group operation \cdot of Γ evaluates to the group element b .

Let $g \in \mathbf{L}_G(\Gamma, \delta)$. A set $X \subset V(G)$ is an a -clique if for all $x, y \in X$: $x \neq y$ implies $g(x, y) = a$. Also, for $X, Y \subseteq V(G)$, X is said to be a -connected to Y , if $X \cap Y = \emptyset$ and $g(x, y) = a$ for all $x \in X, y \in Y$.

A function $\sigma : V \rightarrow \Gamma$ is called a *selector*. For each selector σ we associate with g a (Γ, δ) -gain graph g^σ on $G = (V, E)$ by letting, for each $uv \in E$,

$$g^\sigma(uv) = \sigma(u)g(uv)\delta(\sigma(v)) . \quad (2)$$

The class $[g] \subseteq \mathbf{L}_G(\Gamma, \delta)$ defined by

$$[g] = \{g^\sigma \mid \sigma : V \rightarrow \Gamma\}$$

is called the *switching class* generated by g .

It is not difficult to prove that a switching class is an equivalence class of skew gain graphs. The underlying equivalence relation on $\mathbf{L}_G(\Gamma, \delta)$ is that $g \equiv g'$ for $g, g' \in \mathbf{L}_G(\Gamma, \delta)$ if and only if $\exists \sigma : V(G) \rightarrow \Gamma$ such that $g' = g^\sigma$. Obviously $g \equiv g$ and if $g_1 \equiv g_2$ then also $g_2 \equiv g_1$, because $g_1^\sigma = g_2$ if and only if $g_1 = g_2^{\sigma^{-1}}$, where the σ^{-1} is such that $\sigma^{-1}(v) = \sigma(v)^{-1}$ for all $v \in V$.

Closure under composition of selectors is something that we would expect in our model: it is a consequence of Axiom A1 of the introduction. If we define the composition of two selectors σ and τ to be $\sigma\tau(v) = \sigma(v)\tau(v)$, then we can prove that for each $g \in \mathbf{L}_G(\Gamma, \delta)$ and selectors σ, τ , $g^{\sigma\tau} = (g^\tau)^\sigma$.

If the group Γ is the cyclic group of order 2, \mathbf{Z}_2 , then by necessity the involution is the identity function and the skew gain graphs are exactly the undirected simple graphs of, e.g., [1, 6]. Directed graphs are obtained by choosing $\Gamma = \mathbf{Z}_4$ and we take the involution δ to be the group inversion.

3 The general framework

In the following let Γ be a fixed, but arbitrary group and δ a fixed, but arbitrary involution of Γ .

Let $g \in \mathbf{L}_G(\Gamma, \delta)$ and $h \in \mathbf{L}_H(\Gamma, \delta)$ be skew gain graphs. An injection $\psi : V(G) \rightarrow V(H)$ embeds g into h , denoted by $g \xrightarrow{\psi} h$, if

$$g(uv) = h(\psi(u)\psi(v)) \text{ for all } uv \in E(G).$$

If we do not care what ψ is, we write $g \hookrightarrow h$ instead. Note that in some definitions of embedding there is also an injection on the labels, but since our application attaches meaning to the labels, we do not allow that here.

The embedding ψ is an *isomorphism* from g to h if $g \xrightarrow{\psi} h$ and $h \xrightarrow{\psi^{-1}} g$. We denote this fact by $g \xrightarrow{\psi} h$, or, equivalently, $h \xrightarrow{\psi^{-1}} g$.

The definition of embedding can be extended to switching classes in a natural way:

$$g \hookrightarrow [h] \text{ if and only if there exists } h' \in [h] \text{ such that } g \hookrightarrow h'.$$

In this and the following sections, the central problem is to decide whether the *query* skew gain graph $g \in \mathbf{L}_G(\Gamma, \delta)$ can be embedded in a switch of the *target* skew gain graph $h \in \mathbf{L}_H(\Gamma, \delta)$.

We assume for the remainder of the paper that the target skew gain graph is *total*, meaning that $H = (V, E_2(V))$ for some set of vertices V .

We now come to the definitions central to this paper. We are interested in establishing for a certain query graph g into which other skew gain graph g' it may be transformed so that the ability of embedding g into h is preserved into g' . More formally, we define $\mathcal{R}_{(\Gamma, \delta)}$ as the set of *embedding equivalent* pairs $(g, g') \in \mathbf{L}(\Gamma, \delta) \times \mathbf{L}(\Gamma, \delta)$ such that

$$\forall h : g \hookrightarrow [h] \iff g' \hookrightarrow [h].$$

Note that in our definition we have left the embedding itself unspecified, meaning that in general we do not care whether g and g' are embedded “in the same place”. It also implies that g and g' may have different underlying graphs.

Although we have just defined the largest possible (equivalence) relation relating skew gain graphs from $\mathbf{L}(\Gamma, \delta)$ to each other, it does not give us any concrete information which pairs are actually in the relation for a given group and involution. In the remainder of this paper we shall establish a number of results which either show that some pairs are definitely in this relation, or that some pairs can never be.

Let R be any equivalence relation on $\mathbf{L}(\Gamma, \delta)$. R is an *embedding invariant relation* (*emir*) if $(g, g') \in R$ implies $(g, g') \in \mathcal{R}_{(\Gamma, \delta)}$.

We now give some examples of emirs that occur in the literature. The following easy lemma shows that for embedding the identities of the vertices of the query graph are unimportant.

Lemma 1.

For two isomorphic (Γ, δ) -gain graphs g and g' (with isomorphism ϕ from g to g'): if $g \xrightarrow{\psi} h$, then $g' \xrightarrow{\psi \cdot \phi^{-1}} h$.

The second example is that embedding a query graph g is the same as embedding one of its switches:

Lemma 2.

If $g \xrightarrow{\psi} [h]$, then also $g^\sigma \xrightarrow{\psi} [h]$ for any selector $\sigma : V(g) \rightarrow \Gamma$.

Note that Lemma 1 implies the existence of an emir $R_{\text{IR}}: (g, g') \in R_{\text{IR}}$ if and only if $g \cong g'$. Another example comes from Lemma 2 where it is proved that in fact \equiv is an emir.

We shall now give a slightly more complicated example.

Define R_{DCR} such that $(g, g') \in R_{\text{DCR}}$ if g' can be obtained from g by removing any number of cut edges of g . The symmetric closure of this relation, R_{CR} , is an equivalence relation on (Γ, δ) -gain graphs. So any two g and g' are related if and only if they have exactly the same cycles and the same domain. A basic result from the theory of switching classes proves that this relation is in fact an emir (see for instance [8]):

Theorem 1.

Let H be a graph and let T be a subgraph of H that is a forest. For every (Γ, δ) -gain graph g on T and every h on $H: g \xrightarrow{\text{id}} [h]$.

The proposition states that any acyclic structure can be embedded in a switching class. Note that by removing edges we do not change the size of the domain of the (Γ, δ) -gain graph; this is necessary for establishing embedding invariance.

To combine two emirs into one we can use the join operation: for two emirs R and R' on (Γ, δ) , the join of R and R' , denoted by $R \vee R'$, is the smallest equivalence relation including both R and R' .

Lemma 3.

If R and R' are emirs, then the join of R and R' is an emir.

The join can be used to combine various emirs into a larger one. For instance, joining an emir such as R_{DCR} with R_{IR} yields an emir that “incorporates” removing of cut edges and isomorphisms. In such a way we can define various emirs and compose these to come as close as possible to the largest of emirs, $\mathcal{R}_{(\Gamma, \delta)}$.

4 Bridging

In this and the coming sections we assume that the group Γ is abelian and that the involution δ is the group inversion $^{-1}$; we will denote the identity of the group simply by 0.

The reason for the restriction is the Cyclic Sum Invariance, which holds for switching classes with abelian groups and involution equal to

the group inversion. It means that if one takes any cycle and computes the sum along that cycle for the labels on that path, then this value does not change when the skew gain graph is switched [2].

Let $g, g' \in \mathbf{L}(\Gamma, \delta)$ be such that g contains a 0-chain $p = (x_0, \dots, x_k)$. Then, for integers k and ℓ with $k \geq \ell$, g' is an (k, ℓ) -bridging of g , denoted $gB_\ell^k g'$, if g' is a (Γ, δ) -gain graph on $V(g)$ with


$$E(g') = (E(g) - E(p)) \cup E(p') \text{ for } p' = (x_0, \dots, x_{\ell-1}, x_k)$$

and

$$g'(e) = \begin{cases} g(e), & \text{if } e \in E(g) - E(p) \\ 0, & \text{otherwise} \end{cases}$$

We additionally define B_ℓ^k to be equal to $(B_k^\ell)^{-1}$ for $k \leq \ell$.

For the following two (Γ, δ) -gain graphs g and g' it holds that $gB_3^5 g'$. In this case $p = (x_0, \dots, x_5)$:



Note that we can assume that the chain is labelled with zeroes, since if it does not we can always switch it so that it does. Also, the fact that we have chosen certain vertices (x_ℓ, \dots, x_{k-1}) to become isolated might seem arbitrary, but joining with $R_{\mathbb{R}}$ removes this restriction.

In what follows we are interested in determining for which groups we can always (i.e., for any (Γ, δ) -gain graph $g \in \mathbf{L}(\Gamma, \delta)$) change bridges of length k to bridges of length ℓ . For this we introduce the following relation $R_\Gamma \subseteq \mathbf{N} \times \mathbf{N}$, where $(k, \ell) \in R_\Gamma$ if and only if B_ℓ^k is an emir on $\mathbf{L}(\Gamma, \delta)$. Obviously, for any group Γ it holds that $(k, k) \in R_\Gamma$ where $k > 0$.

The following lemma couples the concept of bridging to something we can more easily verify. Implicitly we allow the embedding only to be changed on the chain vertices that occur on the bridge.

A (Γ, δ) -gain graph on $\{0, \dots, n\}$ for some n is an (n, k) -bridge structure if it has a 0-path $(0, \dots, k)$.

The following lemma shows that to decide whether we can bridge paths of length k into ℓ , we can look at total skew gain graphs which

have a 0-labelled path $(0, \dots, k)$ and show that whatever labels are on the other edges, we can always find a 0-summing path from 0 to k of length ℓ .

Lemma 4.

Let k and ℓ be natural numbers, and let $n = \max(k, \ell)$. It holds that $(k, \ell) \in R_\Gamma$ if and only if for every (n, k) -bridge structure g there is a 0-summing path p in g of length ℓ from 0 to k .

Proof. For the if-part we need only note that we can replace a path of length k by one of length ℓ which has the same starting and end-point without changing the sum on any of the cycles of which these vertices are part. Because a bridge is part of a chain, every vertex on the bridge is part of exactly the same cycles.

The only if part follows from the fact if we cannot replace the path of length n by a path of the same sum of length k , then we change the cyclic sum along at least one cycle, which contradicts the Cyclic Sum Invariance.

Theorem 2.

For natural numbers $k_1 > 1$ and $k_2 > 2$: $(k_1, 1) \notin R_\Gamma$ and $(k_2, 2) \notin R_\Gamma$, if Γ is not the trivial group, $\{0\}$.

Proof. Let $a \in \Gamma$ with $a \neq 0$. Let g_2 be a (Γ, δ) -gain graph on $\{0, \dots, k_2\}$ such that for $1 \leq i \leq k_2 - 1$ $g_2(0, i) = 0$, $g_2(i, k_2) = a$ and all other edges are labelled arbitrarily. Hence for all i , $g_2(0, i, k_2) = a \neq 0$. The same kind of reasoning can be applied to the other case.

Example 1.

If we know that $(5, 3) \in R_\Gamma$, then it is easy to see that $(k, k - (5 - 3)) = (k, k - 2) \in R_\Gamma$ as long as $k - 2 \geq 3$: if g contains a chain of length greater than 5, then we can take any part of this chain of length 5 and reduce it to 3 and thereby reduce the length of the entire chain from k to $k - 2$. We can repeat this process until the length of chain path is not sufficiently long anymore. We conclude that if we prove that $(5, 3) \in R_\Gamma$ then $(k, k - 2) \in R_\Gamma$ for $k \geq 5$ and even $(k, k - 2\ell)$ for $k - 2\ell \geq 3$. Using similar reasoning we conclude that $(3, 5) \in R_\Gamma$ implies that $(k, k + 2\ell)$ for $k \geq 3$. \diamond

In general we have

Lemma 5.

If $(k_1, \ell_1) \in R_\Gamma$ then $(k_2, \ell_2) \in R_\Gamma$ where $\ell_2 = k_2 - (k_1 - \ell_1)m$, $m \geq 1$ and $\ell_2 \geq \ell_1$.

If $\Gamma = \Gamma_1 \times \Gamma_2$ then $(k, \ell) \in R_\Gamma$ implies $(k, \ell) \in R_{\Gamma_i} (i = 1, 2)$, but not vice versa, not even if $\Gamma_1 = \Gamma_2$ (see Theorem 4). The positive result is easy, because the identity of Γ maps to the identities of the factors. Hence the 0-summing paths stay 0-summing in the projection. The following result says that if a bridging is not possible for a given group, it automatically precludes bridging in groups of which it is a factor.

Lemma 6.

If Γ is a group such that $(k, \ell) \notin R_\Gamma$, then this also holds for all groups of which Γ is a factor.

4.1 The case for \mathbf{Z}_2

In view of Theorem 2 it may be surprising that bridgings do exist.

Lemma 7.

$(5, 3) \in R_{\mathbf{Z}_2}$

Proof. Let b be a (5,5)-bridge structure (recall that the path $(0, \dots, 5)$ is a 0-path, and the other edges are arbitrarily labelled by elements of \mathbf{Z}_2). Now, if $b(0, 3) = 0$, then $b(0, 3, 4, 5) = 0$. The same reasoning applies to $(2, 5)$. In the other cases, $b(0, 3) = 1 = b(2, 5)$ and $b(0, 3, 2, 5) = 0$.

Knowing that bridging is possible under \mathbf{Z}_2 we can now illustrate that the necessity of the target skew gain graph being total: take a cycle on six vertices – call it c . We can bridge c into c' which consists of two isolated vertices and a cycle on four vertices. Obviously, $c \leftrightarrow [c]$, but $c' \not\leftrightarrow [c]$. The reason is that the target graph, in this case c , does not have all its edges present.

Lemma 8.

$(k, \ell) \notin R_{\mathbf{Z}_2}$ if k and ℓ are of opposite parity.

Proof. Let k and ℓ be of opposite parity. We may assume $\ell, k \geq 3$, because of Theorem 2.

Let $n = \max(k, \ell)$, b be a (n, k) -bridge structure and $V = V(b)$. By Lemma 4, we only need to exhibit one such structure which has no path of length ℓ from 0 to k which sums to 0. For that, choose b such that the sets $K \subseteq V$ and $V - K$ are 0-connected 1-cliques. Here, $K = \{x \mid 0 \leq x \leq k, x \text{ even}\}$. Note that there is a 0-path $(0, 1, \dots, k)$.

We are interested in paths of length ℓ which go from 0 to k and sum to 0. If k is even, then ℓ is odd, and the path is one that starts in K

and ends in K . Since we must switch from K to $V - K$ an even number of times, we traverse an an odd number of edges within either K and $V - K$. Since these edges each contribute 1 to the sum, and they are the only edges which contribute, the sum along the path equals 1. If k is odd and hence the path starts in K and ends in $V - K$ similar reasoning leads to a sum of 1.

Theorem 2 and Lemmas 5, 7 and 8 lead to the following.

Corollary 1.

If $k \geq \ell > 2$ then $(k, \ell) \in R_{\mathbf{Z}_2}$ if and only if k and ℓ have the same parity. Also, $(k, \ell) \in R_{\mathbf{Z}_2}$ for $1 \leq \ell \leq 2$ if and only if $k = \ell$.

4.2 The case for \mathbf{Z}_3


The following result was quite a surprise.

Lemma 9.

$(6, 4) \in R_{\mathbf{Z}_3}$ and $(6, 5) \in R_{\mathbf{Z}_3}$, but $(5, 4) \notin R_{\mathbf{Z}_3}$.

Proof. The positive results have been obtained by a computer check of all paths of length 4, resp. 5, from 0 to 6 in a (6, 6)-bridge structure.

The counterexample for $(5, 4)$ is given in the following figure, here the solid edges are labelled with 0 and the dashed edges (in the direction of the arrow) with 1. The reader may verify that indeed no path from 0 to 6 of length 5 sums to 0.



5 An algorithm for checking configuration containment

In this section we use the fact that $(5, 3) \in R_{\mathbf{Z}_2}$ (Lemma 7) to derive an algorithm for checking that $g \hookrightarrow [h]$ for $g, h \in \mathbf{L}(\mathbf{Z}_2, -1)$ where h is total. The result is mainly based on the following graph theoretical argument which shows that if we consider graphs that do not have any isolated vertex or leaves, and every chain contains a bounded number of chain number, then the number of vertices in the graph can be bounded by a constant multiple of the cycle rank of the graph. The *cycle rank* of a

graph G is defined as the size of its cycle base, and equals $e - n + k$, where $n = |V(G)|$, $e = |E(G)|$ and k is the number of connected components of G .

Lemma 10. *Let $G = (V, E)$ be a connected graph which has only vertices of degree at least two and at least one dense vertex. If every maximal chain in G has at most $c > 0$ chain vertices, then $|V(G)| \leq 2c\xi$, where ξ is the cycle rank of G .*

Proof. We first make an estimation for graphs which only contain dense vertices. Let $d_G(v)$ denote the degree of the vertex v of G . Then, by the handshaking lemma of graph theory, $2e = \sum_{v \in V} d_G(v) \geq 3n$, since $d_G(v) \geq 3$ for all v . Hence $\xi = e - n + 1 \geq 3n/2 - n + 1 = n/2 + 1$, so that $n \leq 2\xi$ as required. Now, any edge between two dense vertices can be replaced by a chain of a most c chain vertices, which adds to n and e in equal amounts, so that $n \leq 2c\xi$.

Lemma 11.

Let $g \in \mathbf{L}(\mathbf{Z}_2, -1)$ and let ξ be the cycle rank of g . Then, there exists a g' embedding equivalent with g such that $\text{ni}(g') \leq 6\xi$, where $\text{ni}(g')$ is the number of non-isolated vertices of g' .

Proof. Remove cut-edges, isolated vertices and use the (5,3) bridging to change g into g' which has the property that it consists only of a number of isolated vertices, dense and chain vertices. Neither of these operations change the cycle rank of g . Now apply Lemma 10 to each of the components of the graph (the cycle rank of a disconnected graph equals the sum of the cycle rank of its components) to obtain the given bound for the number of chain and dense vertices, where we use Lemma 7 to limit the number of chain vertices in any chain to 3. We omit in this reasoning components which are simple cycles, connected graphs which have only chain vertices. These, however, can all be reduced to cycles of length at most six, again using Lemma 7.

Finally, we can formulate a bound on the time complexity of the embedding problem for \mathbf{Z}_2 as follows:

Theorem 3.

Let $g, h \in \mathbf{L}(\mathbf{Z}_2, -1)$ where h total, $n = |V(h)|$ and ξ is the cycle rank of g . It can be decided in $\mathcal{O}(n^{6\xi+2})$ time whether $g \leftrightarrow [h]$.

Proof. After checking that $|V(g)| \leq h$, we can find an embedding equivalent g' such that $\text{ni}(g') \leq 6\xi$ through Lemma 11. Now, we actually remove

the isolated vertices from g' (we have already checked that g does not have more vertices than h does). The number of possible injections from g' into h is bounded by $n^{6\xi}$, for each of which we have to do at most $\mathcal{O}(n^2)$ work to see if under the injection, we can switch h so that it contains g' (using the results of [5]). The preprocessing of g , which consists of removing leaves, isolated vertices and shortening chains, can easily be done in time $\mathcal{O}(n^2)$.

6 Some impossibility results

In this section we are interested in determining, given a natural number ℓ , for which finitely generated abelian group Γ it holds for every $k > \ell$, that $(k, \ell) \notin R_\Gamma$. In Lemma 2 we found two such examples, $\ell = 1$ and $\ell = 2$, in which case impossibility was obtained for all groups. Since we have already treated the cases for $\ell \leq 2$, we assume $\ell \geq 3$, and hence $k > 3$. From Lemma 6 and the fundamental result on finitely generated abelian groups, it follows that we can restrict ourselves to solving this question for the cyclic groups (of order a prime power) and \mathbf{Z} .

Since we are interested in proving the impossibility of bridging, we have to show that we can always find (k, k) -bridging structure in which there is no 0-summing path of length ℓ from 0 to k .

First we investigate which edges in the bridge structures *must* be labelled with a non-identity element. These are exactly the edges that are on a path of length ℓ from 0 to k which traverse only edges on the path $(0, \dots, k)$, except for one edge which has an undetermined label. We observe that these edges are those of the form

$$(i, i + (k - \ell + 1)), i = 0, \dots, \ell - 1. \quad (3)$$


We shall next prove that the only bridging $(k, 3)$ for $k > 3$ occurs if the group is trivial or the group is \mathbf{Z}_2 . The main technique used here is to generate a family of skew gain graphs, depending on k , which contains a large 0-clique, and only relatively few edges besides. Parts of the paths in the 0-clique contribute nothing to the sum along a path, so only the values on the other edges really matter. To simplify the proof any vertex outside X is connected in a uniform way to all vertices in X and (by reversibility) the other way around. In 2-structures jargon, such a set X is called a *clan* (see [2]). The next theorem is a typical example of this kind and can be viewed as an illustration of the proof technique.

Theorem 4.

If for $k > 3, (k, 3) \in R_\Gamma$ for a finitely generated abelian group, then Γ is \mathbf{Z}_2 or the trivial group.

Proof. Like in Lemma 2 the idea is to find a (k, k) -bridge structure which does not exhibit an 0-summing path of length $\ell = 3$ from 0 to k . Because of Lemma 6 and the fundamental theorem on finitely generated abelian groups, we start by considering the cyclic groups of order larger than two and the group \mathbf{Z} of integer addition.

Consider the following graph in which all edges whose value is as yet unknown is labelled with a variable label a_i for some i , and the vertex X represents a 0-clique on $k - \ell$ vertices.



By (3), a_0, a_1 and a_2 should be labelled by values different from 0. It is easily seen that also $a_4 \neq 0$ (for paths through X). We also find that $a_0 \neq a_4^{-1}$, because of the path $(0, k-2, x \in X, k)$. In fact if we set $a_3 = 0$, a_0, a_2 and a_4 to the generator of the group, 1, and a_1 to 1^{-1} there is no path of length $\ell = 3$ which sums to 0. It is important to note that since the group has order at least three, $1 \neq 1^{-1}$.

Since a $(k, 3)$ bridging existed for \mathbf{Z}_2 , we should also show that such a bridging is not possible for $\mathbf{Z}_2 \times \mathbf{Z}_2$. Taking the same graph as our starting point, we choose a_3 the identity $(0, 0)$ and set $a_0 = a_1 = a_2 = (0, 1)$ and $a_4 = (1, 0)$. Again, the reader can verify (there are only a finite number of cases), that no path of length 3 from 0 to k sums to $(0, 0)$.

7 Conclusions and future work

Taking the model of Ehrenfeucht and Rozenberg as our starting point, we have considered the embedding problem in detail. We have set up a framework to establish results about reducing query skew gain graphs to smaller ones and proved some general results in this matter. Then we concentrated on bridging, which, for \mathbf{Z}_2 at least, results in an algorithm for the embedding problem which is dominated not by the size of the query graph, but by its cycle rank, corresponding to the general intuition in switching classes that cycles make life harder.

We have not completed a full investigation of all possible bridgings for all possible finitely generated abelian groups, although we have the full picture for \mathbf{Z}_2 and \mathbf{Z}_3 . We do conjecture that for every such group Γ there is for every ℓ a k such that $(k, \ell) \notin R_\Gamma$. Note by the way, that bridging is just one possible reduction strategy and others might exist. In that sense, the research in this area is still very much open, especially for non-abelian groups where bridging is not an option.

Acknowledgements

We thank Hans Bodlaender for some helpful discussions.

References

1. A. Ehrenfeucht, J. Hage, T. Harju, and G. Rozenberg. Complexity issues in switching of graphs. In H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors, *Theory And Application Of Graph Transformations - TAGT '98*, volume 1764 of *Lecture Notes in Computer Science*, pages 59–70, Berlin, 2000. Springer Verlag.
2. A. Ehrenfeucht, T. Harju, and G. Rozenberg. *The Theory of 2-Structures*. World Scientific, Singapore, 1999.
3. A. Ehrenfeucht and G. Rozenberg. Dynamic labeled 2-structures. *Mathematical Structures in Computer Science*, 4:433–455, 1994.
4. J. L. Gross and T. W. Tucker. *Topological Graph Theory*. Wiley, New York, 1987.
5. J. Hage. The membership problem for switching classes with skew gains. *Fundamenta Informaticae*, 39(4):375–387, 1999.
6. J. Hage, T. Harju, and E. Welzl. Euler graphs, triangle-free graphs and bipartite graphs in switching classes. *Fundamenta Informaticae*, 58(1):23–37, November 2003.
7. J. J. Rotman. *The Theory of Groups*. Allyn and Bacon, Boston, 2nd edition, 1973.
8. T. Zaslavsky. Signed graphs. *Discrete Applied Math.*, 4:47–74, 1982. Erratum on p. 248 of volume 5.
9. T. Zaslavsky. Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory, Ser. B*, 47:32–52, 1989.