

THE EMBEDDING PROBLEM FOR SWITCHING CLASSES

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

Jurriaan Hage
Institute of Information and Computing Science
Universiteit Utrecht
P.O.Box 80.089, 3508 TB Utrecht, Netherlands
jur@cs.uu.nl

Tero Harju
Department of Mathematics
University of Turku
FIN-20014 Turku, Finland
harju@utu.fi

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

February 1, 2005

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.

In case vertices are given an identity, 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. 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 many positive and negative results for more general cases, where the graphs are labelled with elements of a (finitely generated abelian) 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 for 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 [6] 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.

A fourth axiom stated that every network/graph should have at least three processors. This was to make sure that no exceptions arose when deriving the most general model that upholds the axioms above; in the paper of Ehrenfeucht and Rozenberg it turned out that the model of a network of two processors is more flexible than that of more than two processors. We do not state the axiom, but shall use the model for more than two processors also for networks of two processors. As an aside, the book by Ehrenfeucht, Harju and Rozenberg [5] proposes another axiom: every processor can choose to remain inactive. There is however no need for this axiom: it is implied by the others.

Although each processor i was to have a set of *output actions* Ω_i and a set of *input actions* Σ_i , in [6] (see also [5]) 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 (3) in Section 4. 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 [12] and the voltage graphs of [7].




Figure 1: The query and target skew gain graphs respectively

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 reachable 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: transforming 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 [6].

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. An example can be found in Figure 1 which features g on the left and h on the right. Here the intended meaning of g signifies the existence of a deadlock.

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 the subgraph g ? In terms of the example: is there a possible state in the system, derivable from h , which contains a deadlocked subgraph somewhere. If the embedding from g into h is known, then this can be (in many cases) efficiently solved by applying the results of Hage [8]. 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 (and vice versa). As it turns out, reduction is possible if the cycle rank (or cyclomatic index) of a graph is low. In a nutshell, our result says that the embedding problem is exponential not in the number of vertices of the graph, but in a different graph measure, the cycle rank. This may be compared to such measures as treewidth [1], in

which case there exist efficient for NP-complete problems working on graphs of bounded treewidth. The main advantage of our work is that the cycle rank, and the corresponding decomposition, can be computed very efficiently. This is not the case for treewidth, for example. On the other hand, the cycle rank as a measure is also quite a bit more restricted. As far as we know it only applies to the embedding problem.

An earlier version of this paper was presented at the 2nd International Conference on Graph Transformations in Rome [4]. The main changes with respect to that paper are that in this paper we start by considering the special case of undirected graphs, which conveys many of the essential ideas, but tries to avoid depending on any knowledge of group theory. Even for this simplest of cases, the embedding problem is NP-complete [3]. Secondly, Section 8 contains additional impossibility results that were not in the first version, in particular, a result that shows with rather tight bounds where looking for the given optimizations is certainly not profitable.

As a result, the paper is now structured as follows: after some general preliminaries where we also explain how to switch (partial) undirected graphs, which we call 0,1-graphs, we show in Section 3 how to derive an algorithm for verifying whether such a graph, the query graph, can be embedded in an undirected graph, the target graph. The resulting complexity depends not on the number of vertices of the graphs in question, but rather on the number of cycles in the query graph. The crucial ingredient we use here is the concept of bridging, which is an operation that shortens cycles in the query graph.

Subsequently, we introduce the full mathematical model of switching classes of graphs with skew gains, and reconsider the embedding problem. We formalize the idea of embedding invariance, and give (im)possibility results for groups other than \mathbf{Z}_2 , which show in which cases the technique of Section 3 may be applied as well. Most of the results in this part of our work are negative: they show that the bridging operation does not readily extend to other groups. These negative results are important in that they show where not to look for savings. In view of the fact that we have nothing significantly better than exhaustive search to look for optimization, this is certainly good to know.

2 Preliminaries

In this section we introduce some general notation on functions, sets and graphs, and a special form of undirected graph, the 0,1-graph, in which we label the edges of an undirected graph with either 0 or 1. We conclude the section with the definition of the switching of 0,1-graphs and some basic results from the literature.

For a (finite) set V , let $|V|$ be the *cardinality* of V . We shall often identify a subset $A \subseteq V$ with its characteristic function $A : V \rightarrow \mathbf{Z}_2$, where $\mathbf{Z}_2 = \{0, 1\}$ is the cyclic group of order two. We use the convention that for $x \in V$, $A(x) = 1$ if and only if $x \in A$. The restriction of a function $f : V \rightarrow W$ to a subset $A \subseteq V$ is denoted by $f|_A$. We denote set difference by $A - B$. It contains the elements

in A which are not in B . If B is a singleton $\{b\}$, then we may write $A - b$ for brevity.

The set $E(V) = \{\{x, y\} \mid x, y \in V, x \neq y\}$ denotes the set of all unordered pairs of distinct elements of V . We write xy or yx for the unordered pair $\{x, y\}$. The graphs of this paper will be finite, undirected and simple, i.e., they contain no loops or multiple edges. We use $E(G)$ and $V(G)$ to denote the set of edges E and the set of vertices V , respectively, and $|V|$ and $|E|$ are called the *order*, respectively, *size* of G . Analogously to sets, a graph $G = (V, E)$ will be identified with the characteristic function $G : E(V) \rightarrow \mathbf{Z}_2$ of its set of edges so that $G(xy) = 1$ for $xy \in E$, and $G(xy) = 0$ for $xy \notin E$. Later we shall use both notations, $G = (V, E)$ and $G : E(V) \rightarrow \mathbf{Z}_2$, for graphs.

Before we go on to introduce the 0,1-graphs which are in fact our focus of investigation, we first introduce some rather standard notation for graphs.

Let $G = (V, E)$ be a graph. A vertex $x \in V$ is *adjacent* to $y \in V$ if $xy \in E$. The *degree* of x in G , denoted $d_G(x)$, is the number of vertices it is adjacent to. The *neighbours* of u in G , denoted $N_G(u)$, or $N(u)$ if G is clear from the context, is the set of vertices adjacent to u in G . A vertex which is not adjacent to any other vertex in a graph is called *isolated*, a *leaf* has degree one, a *chain vertex* degree two, and all other vertices are called *dense vertices*.

For a graph $G = (V, E)$ and $X \subseteq V$, let $G|_X$ denote the *subgraph* of G induced by X . Hence, $G|_X : E(X) \rightarrow \mathbf{Z}_2$.

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 cycle (v_1, \dots, v_k) is different from a path in that $v_1 = v_k$. We naturally extend all notation for paths to cycles.

A *cut edge* in a graph is an edge which is not on any cycle.

Now we may continue and introduce the 0,1-graphs which are in fact slight generalizations of the graphs given above. These graphs will be used in Section 3. Afterwards, we shall generalize the graphs to our full model for this paper. From a graph $G = (V, E)$, we can obtain a 0,1-graph g by labelling its edges with either 0 or 1: $e \in E(G)$ if and only if $g(e) = 0 \vee g(e) = 1$. Such a G is called the *underlying graph* of g . For a graph G , we use \mathbf{L}_G to denote the set of 0,1-graphs with underlying graph G .

A tricky aspect of this definition is the following: every graph G has a natural counterpart g which is a 0,1-graph: taking the complete graph $K_{|V(G)|}$ as the underlying graph we add labels as follows:

$$g(e) = 1 \text{ if } e \in E(G) \text{ and } 0 \text{ otherwise}$$

Such a 0,1-graph will be referred to as a *total* 0,1-graph.

Hence, every 0,1-graph on a complete underlying graph has two graphs associated with it: its underlying graph which is a complete graph, and the graph with which it is associated through its chosen labelling. In order not to confuse the reader in the following, simple undirected graphs shall only arise in the role of underlying graph.

Now, let $g \in \mathbf{L}_G$ and $h \in \mathbf{L}_H$ for some graphs G and H . An injection $\psi : V(G) \rightarrow V(H)$ embeds g , the *query* graph, into h , the *target* graph, (denoted $g \hookrightarrow h$) if

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

A useful intuition behind 0,1-graphs is that the absence of an edge in the underlying graph of G means that we do not care what the corresponding edge of h is labelled with.

The question whether a given simple undirected graph can be embedded in another, entails mapping both graphs to 0,1-graphs with the underlying graph a complete graph, and applying the definition above. As we shall see later, this situation cannot be improved using our results. It is the case when the query graph is relatively sparse (in a sense to be made precise later on), that something can be gained.

With a path $p = (v_1, \dots, v_k)$ in $g \in \mathbf{L}_G$ 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 if $g(v_1v_2) + g(v_2v_3) + \dots + g(v_{k-1}v_k)$ equals b modulo 2. We often denote this fact by writing $g(p) = b$.

We now move on to the definition of switching. Let $g \in \mathbf{L}_G$. A function $\sigma : V(G) \rightarrow \{0, 1\}$ is called a *selector*. For each selector σ we associate with g a 0,1-graph g^σ on $G = (V, E)$ by letting, for each $uv \in E$,

$$g^\sigma(uv) = \sigma(u) + g(uv) + \sigma(v) . \tag{1}$$

where $+$ is addition modulo 2. The *switching class generated by g* is then $[g] = \{g^\sigma \mid \sigma \text{ a selector}\}$. In Section where we introduce switching more generally, we shall prove that switching classes are equivalence classes of graphs.

Example 2.1 Below we have depicted a typical example of a switch. In the first graph, the vertices are labelled with their name, in the second graph we have labelled them with the value selected for the given vertex by the selector, as in $\sigma(u) = 0$ and $\sigma(v) = 1$. Note that the path (u, v, w, x, u) is 0-summing both in g and g^σ .



The fact that the cycle (u, v, w, x, u) gave the same sum in g and g^σ turns out not to be a coincidence. A crucial property used in this paper is the Cyclic Sum Invariance (cf. [11], [9]), the proof of which is part of folklore of the theory of switching classes:

Theorem 2.2 (Cyclic Sum Invariance)

For every cycle c in a 0,1-graph g , and selector σ on g , $g(c) = g^\sigma(c)$.

A stronger version of this theorem can be formulated, since it turns out that the cyclic sums of triangles which involve any given vertex v uniquely determines the switching class.

Related to this result is the Forest Forcing Lemma (cf. Hage [9]):

Lemma 2.3 (Forest Forcing Lemma)

Let $g \in \mathbf{L}_G$ for some graph $G = (V, E)$, and let T be any acyclic subgraph of G . For every $t \in \mathbf{L}_T$, there exists a $h \in [g]$, such that for all $e \in E(T) : h(e) = t(e)$. If the acyclic subgraph T is a spanning subgraph of G , then there is exactly one such h .

We shall now extend the embedding problem for graphs in a natural way to switching classes:

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

Obviously, $g \hookrightarrow [h]$ if and only if $g^\sigma \hookrightarrow [h]$.

The Forest Forcing Lemma has the following consequence for the embedding of acyclic 0,1-graphs:

Corollary 2.4

For every acyclic graph T , and every $t \in \mathbf{L}_T$, t can be embedded in any g which has a (non-induced) subgraph isomorphic to T .

In other words, if the query graph is acyclic, the labels do not matter. The fact that it seems easier to embed structures with few cycles, indeed holds true as we shall show in the next section.

3 An efficient embedding algorithm for 0,1-graphs

In this section we show how to derive a rather efficient algorithm on 0,1-graphs which decides $g \hookrightarrow [h]$ for the case that h is a total 0,1-graph. In addition to the Cyclic Sum Invariance and the Forest Forcing Lemma, the result is based on two more ingredients: the first of these is the following graph theoretical argument which shows that if we consider graphs that do not have any isolated vertex or leaves, and every chain has bounded length, 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 (see Harary [10] for more details).

Lemma 3.1

Let $G = (V, E)$ be a connected graph without leaves 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. Note that in this case we can choose $c = 1$. 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 at most c chain vertices, which adds to n and e in equal amounts, so that $n \leq 2c\xi$. Additionally, we may introduce chains between two such dense vertices, increasing n with at most c , and ξ by one, which keeps the invariant intact.

Before making use of the previous lemma, we have to prove that if we consider the underlying graph of a certain 0,1-graph and find chains of at least a certain length, that we can replace these by shorter ones. In the case of 0,1-graph this 'certain' length turns out to be 5.

Lemma 3.2

Let g be a 0,1-graph on the domain $\{0, 1, \dots, 5\}$ such that the path $(0, 1, \dots, 5)$ is zero-labelled. Then whatever the labels on the other edges in g , there is always a 0-summing path of length 3 from 0 to 5.

Proof: Consider the following graph b , where a solid line indicates a label 0, and all other edges are labelled by something thus far unknown.



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 $b(0, 3, 2, 5) = 0$.

The above result also holds in the other direction:

Lemma 3.3

Let g be a 0,1-graph on the domain $\{0, 1, \dots, 5\}$ such that the path $(0, 1, 2, 3)$ is zero-labelled. Then whatever the labels on the other edges in g , there is always a 0-summing path of length 5 from 0 to 5.

Proof: We depend here on a computer program to try all cases.

Based on this we can show that bridging is a sound operation, in that it does not change the ability to embed:

Lemma 3.4

Let $g, g' \in \mathbf{L}_G$ be 0,1-graphs which only differ in the following way: g has a 0-labelled chain $p = (u_0, u_1, u_2, u_3, u_4, u_5)$ which is part of a cycle in G , where g' has a 0-labelled chain (u_0, u_1, u_2, u_5) , and u_3 and u_4 are isolated vertices. Then g embeds in $[h]$ if and only if g' embeds in $[h]$, where h is a total 0,1-graph.

Proof: Let $h_1 \in [h]$ be such that it has a subgraph isomorphic to g . We show that there is a switch of h_1 which has a subgraph isomorphic to g' . Consider the vertices $V = \{v_0, \dots, v_5\}$ in h_1 corresponding to $U = \{u_0, \dots, u_5\}$. Lemma 3.2 says that whatever h_1 is like, the subgraph $h_1|_V$ has at least one path $p' = (v_0, x, y, v_5)$ which sums to 0. Then g' is isomorphic to a subgraph of h_1^σ by switching an appropriate $\sigma \subseteq \{x, y\}$ to turn the non-zero values from p' in h_1 into actual zeroes. Such a switch however does not endanger the embedding of g' into h_1^σ , because x and y are the images of chain vertices in g' , and because of the Cyclic Sum Invariance it is guaranteed not to harm the sums along any of the cycles, including the one of which the path p' is a part. The vertices u_3 and u_4 can simply be mapped to the two still unused vertices in V .

The same reasoning can be applied in the reverse direction, this time using Lemma 3.3.

It is worth noting that we have two degrees of freedom here: we can take another switch to embed in, and we can change our embedding. We did both in this case.

Lemma 3.5

Let $g \in \mathbf{L}_G$ 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 first from g all cut edges and then use repeatedly switching and Lemma 3.4 as many times as possible to change g to g' . A switching is performed to force paths of 6 vertices to be 0-labelled, and then Lemma 3.4 is applied to the result. Then g' has no vertices of degree 1, and every edge belongs to a cycle. Moreover, in g' no chain has more than 3 chain vertices.

Now we can apply Lemma 3.1 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. 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 3.4, after turning all, except maybe one, label into zero by an appropriate switch.

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

Theorem 3.6

Let g, h be 0,1-graphs with h a total one, $n = |V(h)|$ and ξ is the cycle rank of the underlying graph of g . It can be decided in $O(n^{6\xi+2})$ time whether $g \hookrightarrow [h]$.

Proof: After checking that $|V(g)| \leq n$, we can find an embedding equivalent g' such that $\text{ni}(g') \leq 6\xi$ through Lemma 3.5. Now, we actually remove the isolated vertices from g' . 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 $O(n^2)$ work to see if under the injection, we can switch h so that it contains g' (using the results of [8]).

The preprocessing of g , which consists of removing leaves, isolated vertices and shortening chains, can easily be done in time $O(n^2)$.

The above result has obvious links with the basic idea behind treewidth [1], which is a well-known measure of graphs. Essentially, many graph theoretic NP-complete problems are feasible for graphs of bounded treewidth. The same is shown here for the embedding problem of switching classes with query graph of bounded cycle rank.

A difference with our situation is that, because we deal with switching classes, which gives us an added flexibility, the measure itself, which is the cycle rank, is a much easier notion than that of treewidth. Moreover, the cycle rank of a graph is easy to compute efficiently, something which is not the case for treewidth. On the down side of course, the use of the cycle rank measure is restricted to the embedding problem for switching classes (at least, as far as we can tell).

4 Switching classes of graphs with skew gains

In this section, we introduce the general model for switching classes based on the axioms listed in the introduction. It generalizes the model from the first part of this paper in that the labels are now taken from an arbitrary group, and the label on the edge from vertex u to vertex v , is related to the label on the edge from v to u by an anti-involution, i.e. anti-automorphism, of order at most two. This, instead of simply being identical.

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 .

Since our underlying graphs now become directed graphs to allow different labels between two vertices depending on the direction of the edge, we *redefine* $E_2(V) = \{(u, v) \mid u, v \in V, u \neq v\}$, the set of nonreflexive, directed edges over V . We continue to write uv for the edge (u, v) , but in this and later sections $uv \neq vu$. 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 viewed as undirected graphs where the edges have been given a two-way orientation.

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*

$$g(e^{-1}) = \delta(g(e)) \quad \text{for all } e \in E. \quad (2)$$




Figure 2: Two elements of $\mathbf{L}_G(\mathbf{Z}_4, id)$

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, $\mathbf{L}(\Gamma, \delta) = \bigcup \{ \mathbf{L}_G(\Gamma, \delta) \mid G \text{ is a graph} \}$. A *gain graph* is a $(\Gamma, ^{-1})$ -gain graph; these are also called *inversive skew gain graphs*.

The notion of b -summing path extends naturally to arbitrary groups: a path $p = (v_1, \dots, v_k)$ is a b -summing path for some $b \in \Gamma$ if $g(v_1 v_2) \cdot g(v_2 v_3) \cdots g(v_{k-1} v_k)$ 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 .

Furthermore, 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 $uv \in E$,

$$g^\sigma(uv) = \sigma(u)g(uv)\delta(\sigma(v)) . \quad (3)$$

Example 4.1

To illustrate switching, consider g_1 and g_2 , the (\mathbf{Z}_4, id) -gain graphs of Figure 2(a) and (b) respectively (the group \mathbf{Z}_4 is the group of addition modulo 4; the anti-involution is the identity function giving rise to a symmetric graph). The second of these, g_2 , can be obtained from g_1 by applying the selector σ that maps both 1 and 3 to 3, and both 2 and 4 to 1. For example, the label of the edge $(1, 3)$ is computed as follows: $g_2(1, 3) = g_1^\sigma(1, 3) = \sigma(1)g_1(1, 3)\delta(\sigma(3)) = 3 + 1 + \delta(3) = 3 + 1 + 3 = 3$, where $+$ is addition modulo 4. The path $(1, 2, 3, 4)$ is a 0-path in both g_1 and g_2 . The cycle $c = (1, 3, 4, 1)$ is 3-summing in g_1 (here $\lambda(c)$ equals $(1, 0, 2)$) and 1-summing in g_2 (here $\lambda(c)$ equals $(3, 0, 2)$).

Note that this is an example where the Cyclic Sum Invariance fails. The reason is that the anti-involution is not the group inversion in this case. Therefore we shall later consider only abelian groups with anti-involution equal to the group inversion.

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 for $g, g' \in \mathbf{L}_G(\Gamma, \delta)$

$$g \equiv g' \text{ if and only if } \exists \sigma : V(G) \rightarrow \Gamma \text{ such that } g' = g^\sigma. \quad (4)$$

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 , $g^{\sigma\tau} = (g^\tau)^\sigma$.

Indeed, let $uv \in E(G)$. Then

$$\begin{aligned} (g^\tau)^\sigma(uv) &= \sigma(u)\tau(u)g(uv)\delta(\tau(v))\delta(\sigma(v)) \\ &= \sigma(u)\tau(u)g(uv)\delta(\sigma(v)\tau(v)) \\ &= (\sigma\tau)(u)g(uv)\delta((\sigma\tau)(v)) = g^{\sigma\tau}(uv). \end{aligned}$$

If the group Γ is the cyclic group of order 2, \mathbf{Z}_2 , then by necessity the anti-involution is the identity function and the skew gain graphs are exactly the 0,1-graphs of the first part of the paper. Directed graphs are obtained by choosing $\Gamma = \mathbf{Z}_4$ and we take the anti-involution δ to be the group inversion.

5 The general approach

As we hinted in the first part of the paper, we now introduce a general framework in which operations like bridging (in the sense of Lemma 3.4) and switching can be modelled as operations which preserve the ability to embed.

In the following let Γ be a fixed, but arbitrary abelian group and δ a fixed, but arbitrary anti-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 and reflected 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 anti-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 5.1

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, one that we already glanced at for the case of 0,1-graphs, is that embedding a query graph g is the same as embedding one of its switches:

Lemma 5.2

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

Note that Lemma 5.1 implies the existence of an emir R_{IR} : $(g, g') \in R_{\text{IR}}$ if and only if $g \cong g'$. Lemma 5.2 shows that \equiv as defined in (4) is also 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. The Forest Forcing Lemma extended to arbitrary switching classes proves that this relation is in fact an emir (see for instance [9]). 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 5.3

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 cut edges and taking isomorphisms. In such a way we can define various emirs and compose these to come as close as possible to the largest of emirs, $\mathcal{R}_{(\Gamma, \delta)}$.

6 The general approach to bridging

In Section 3 we considered the possibilities for bridging in \mathbf{Z}_2 and its consequences. In this section define bridging more generally, followed by some general results. We follow up in later sections by considering the case of the group \mathbf{Z}_3 , and give some limits of possibility results, which show in which cases bridging are not possible at all.

In the case of \mathbf{Z}_3 we can use the same approach as in Section 3 to obtain an efficient algorithm for embedding (although the constant in the exponent will be a somewhat larger).

In this and the coming sections we assume that the group Γ is abelian and that the anti-involution δ is the group inversion $^{-1}$; we will denote the identity of the group simply by 0. This is necessary since the bridging operation depends on the Cyclic Sum Invariance and the Forest Forcing Lemma. We refer the reader to the formulation in Section 2, which hold as formulated there also for arbitrary abelian groups with the anti-involution set to the group inversion.

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 a (k, ℓ) -bridging of g , denoted $gB_\ell^k g'$, if g' is a (Γ, δ) -gain graph on $V(g)$ with


$$E(g') = (E(g) - E(p)) \cup E(p')$$

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 (left) and g' (right) it holds that gB_3^5g' . 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.

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 chains of length k to chains 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 reduce chains of length k to 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 . It is essentially a more general version of Lemma 3.4.

Lemma 6.1

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 b there is a 0-summing path p in b of length ℓ from 0 to k .

Proof: By definition $(k, \ell) \in R_\Gamma$ if and only if for all $gB_\ell^k g'$, for all h , $g \hookrightarrow [h]$ if and only if $g' \hookrightarrow [h]$.

For the if-part the proof is a repeat of that of Lemma 3.4 which shows the validity of replacing a path of length k in g by one of length ℓ without changing the ability to embed.

The only-if-part follows from the fact that 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. It is possible that changing the embedding completely compensates this fact in some cases, but this cannot work uniformly.

Theorem 6.2

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 $1 \leq i \leq k_2 - 1$, $g_2(0, i) = 0$, $g_2(i, k_2) = a$. 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 6.3

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 chain 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 $(k, k - 2\ell) \in R_\Gamma$ for $k - 2\ell \geq 3$. Using similar reasoning we conclude that $(3, 5) \in R_\Gamma$ implies that $(k, k + 2\ell) \in R_\Gamma$ for $k \geq 3$. \diamond

In general we have

Lemma 6.4

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 8.1). The positive result is easy, because the identity of Γ maps to the identities of the direct 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 direct factor.

Lemma 6.5

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

In terms of the notation just introduced we have the following result for \mathbf{Z}_2 , which is in fact a restatement of Lemma 3.3.

Lemma 6.6

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

We can go even further and show that this is the best bridging possible for \mathbf{Z}_2 .

Lemma 6.7

$(k, \ell) \notin R_{\mathbf{Z}_2}$ 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 6.2.

Let $n = \max(k, \ell)$, b be a (n, k) -bridge structure and $V = V(b)$. By Lemma 6.1, 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 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 6.2 and Lemmas 6.4, 6.6 and 6.7 lead to the following.

Corollary 6.8

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 $k = \ell$.

7 Bridging in \mathbf{Z}_3

For \mathbf{Z}_3 , it turns out that there is a result similar to the one for \mathbf{Z}_2 , which as a result allows us to find an optimized embedding algorithm for what might be called 0, 1, 2-graphs (of a kind). Again, the embedding problem is restricted to embedding in total 0, 1, 2-graphs.


The following result was quite a surprise.

Lemma 7.1

$(6, 4) \in R_{\mathbf{Z}_3}$ 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 and 5, respectively, from 0 to 6 in a $(6, 6)$ -bridge structure.

The counterexample for $(5, 4)$ is given in the following figure, where 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 5 of length 4 sums to 0. Therefore Lemma 6.1 gives the claim.



In other words, in general the best we can do is to shorten chains to a length of 5 which yields the following result

Theorem 7.2 Let $g, h \in \mathbf{L}(\mathbf{Z}_3, -^1)$ with h total, $n = |V(h)|$ and ξ is the cycle rank of the underlying graph of g . It can be decided in $O(n^{10\xi+2})$ time whether $g \hookrightarrow [h]$.

Proof: Analogous to the approach of Theorem 3.6, this time however the bound is somewhat higher, because at worst, we have $c = 5$ in Lemma 3.1.

8 Impossibility results

In the first part of this section we are interested in determining, given a natural number ℓ , for which finitely generated abelian group Γ^1 it holds for every $k > \ell$, that $(k, \ell) \notin R_\Gamma$. In Theorem 6.2 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 6.5 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 structures 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

$$(i, i + (k - \ell + 1)), \quad i = 0, \dots, \ell - 1. \quad (5)$$

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 X , and only relatively few other edges. Parts of the paths in X 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 [5]). The next theorem is a typical example of this kind and can be viewed as an illustration of the proof technique.

Theorem 8.1


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

Proof: Like in Lemma 6.2 the idea is to find a (k, k) -bridge structure which does not exhibit a 0-summing path of length $\ell = 3$ from 0 to k . Because

¹We continue to insist that the anti-involution is the group inversion; bridging can only be applied if that is the case, anyway.

of Lemma 6.5 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 are labelled with a variable label a_i for some i , and the vertex X represents a 0-clique on $k - \ell$ vertices.



By (5), 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 $a_0 \neq -a_4$, because of the path $(0, k - 2, x, k)$ where $x \in X$. 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 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, $1 \neq -1$.

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 $(0, 0)$.

In the second part of this section, we turn the question around. Bridgings, say from n to k , are discovered by taking a group, generating all (n, n) -bridge structures and verifying for each that there is a 0-summing path from 0 to n of length k . For large groups there are many such structures, so it would be interesting to know which groups we can omit from our search. The following results give a large part of the answer, capturing the intuition that the larger the group, the more likely are we to encounter an (n, n) -bridge structure which does not have the property we are looking for.

Lemma 8.2

For any given n , $(n, k) \notin R_\Gamma$ for all $k < n$ for all groups $\Gamma = \mathbf{Z}_{p_1} \times \dots \times \mathbf{Z}_{p_{n-1}}$, where each $p_i \geq 2$.

Proof: Consider any combination n, k with $k < n$, and Γ as defined in the theorem. Let g be (n, n) -bridge structure with the following additional properties: for every vertex u , we define the labels on the edges to all vertices $v > u + 1$ uniformly: $g(uv) = (0, \dots, 0, 1, 0, \dots, 0)$, where the only non-zero element is in the u position, and in fact equal to the generator 1 of the group \mathbf{Z}_{p_u} . Because we only look at the case that the anti-involution is the group inversion, we have $g(vu) = -g(uv)$.

Consider now any path p of length $k < n$, leading from 0 to n . Since $k < n$, we must traverse at least one edge uv of which the endpoints are not neighbours

on the path $(0, 1, \dots, n)$. In fact, we assume that uv is the earliest such pair, meaning that p is of the form $p = (0, 1, \dots, u, v)$ for some $u \geq 0$ and $v > u + 1$. We know that $u < v$, so that to the sum along p , the value $(0, \dots, 0, 1, 0 \dots, 0)$ is added (the 1 in the u th position). However, since we have now both arrived in u and left it again, we may never return, since otherwise p is not a valid path. But this means that the u th component of the sum along p will continue to be non-zero, since the only edge label which may cancel it is the label on the edges going into u from vertices $v > u + 1$.

Lemma 8.3 For all n , $(n, k) \notin R_\Gamma$ where $\Gamma = \mathbf{Z}$, the group of integers under addition.

Proof: Similar to the previous proof, now taking the labels from u to $v > u + 1$ equal to 2^u (and -2^u for the reverse edge). The independence between the various values in the sum comes from the fact that we use different powers of two, which in fact simulates the use of tuples of the previous lemma.

Lemma 6.5, Lemma 8.2, Lemma 8.3, and the Fundamental Theorem On Finitely Generated Abelian Groups give the following:

Corollary 8.4

Let n be a fixed natural number. For every group Γ that has at least $n - 1$ non-trivial direct factors, it holds that $(n, k) \notin R_\Gamma$ for all $k < n$.

The usefulness of these results is that they limit the number of finitely generated abelian groups we have to look at, when we take a certain fixed value n , and are interested in finding a bridging (n, k) for some k . Our current set-up is to simply generate all (n, n) -bridge structures and show that each of them has a 0-summing path of length k . Hence, this saves us from looking for a long time in a number of certifiably wrong places.

A little bit more thought allows us to derive the following improved result:

Theorem 8.5

For every group Γ that has at least k non-trivial direct factors, it holds that $(n, k) \notin R_\Gamma$ for all $n > k$.

Proof: The proof is very similar to that of Lemma 8.2, except that edges outgoing from i to $j > i + 1$ are labelled with 0 if $i \geq k$. In other words, only the vertices 0 to $k - 1$ actually have non-zero elements on the edges to their successors. The reason why the result continues to hold, is that we can be sure that the first non-zero edge uv (from the proof of Lemma 8.2) leaves a vertex in $\{0, \dots, k - 1\}$. And that edge is all we need to obtain a non-zero sum.

By Corollary 8.4 and Theorem 8.5 we obtain the following.

Theorem 8.6

Let Γ be a group with p (non-trivial) direct factors in its decomposition (according to the Fundamental Theorem Of Finitely Generated Abelian Groups). Then $(n, k) \notin R_\Gamma$ for all $n > k$, $n \geq p$ and $k > p - 1$.

As a result, for every group we only have to check a finitely number of combinations for n and k .

9 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 and \mathbf{Z}_3 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. The measure we need for switching classes is rather straightforward, compared to measures like cliquewidth and treewidth which are in use for graphs. On the other hand, we have only investigated the embedding problem, and not other infeasible problems for switching classes. This is certainly an area worthy of investigation.

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 and many cases in which bridging is certainly not possible. 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 or abelian groups with arbitrary anti-involutions where bridging is not even an option.

A different way of approaching the problem, is to investigate embedding preservation and not embedding invariance. In that case we look for pairs (g, g') such that $g' \hookrightarrow [h]$ implies $g \hookrightarrow [h]$, but not necessarily vice versa. Another aspect is that in all the cases described above, we can also reconstruct the embedding for the larger structure from the embedding for the smaller structure. Maybe if all we want to know is whether embedding is possible, but not where exactly, we may get more possibilities for optimization.

Acknowledgements

We thank Hans Bodlaender for some helpful discussions on an earlier version of this paper, and Nikè van Vugt for reading through the manuscript.

References

- [1] H. L. Bodlaender. A partial k -arboretum of graphs with bounded treewidth. *Theoretical Computer Science*, 209:1 – 45, 1998.
- [2] D. G. Corneil and R. A. Mathon, editors. *Geometry and Combinatorics: Selected Works of J. J. Seidel*. Academic Press, Boston, 1991.
- [3] 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.
- [4] A. Ehrenfeucht, J. Hage, T. Harju, and G. Rozenberg. Embedding in switching classes with skew gains. In H. Ehrig, G. Engels, F. Parisi-Presicce, and G. Rozenberg, editors, *Graph Transformations, Second International Conference, ICGT 2004*, volume LNCS 3256, pages 257 – 270. Springer Verlag, September/October 2004.
- [5] A. Ehrenfeucht, T. Harju, and G. Rozenberg. *The Theory of 2-Structures*. World Scientific, Singapore, 1999.
- [6] A. Ehrenfeucht and G. Rozenberg. Dynamic labeled 2-structures. *Mathematical Structures in Computer Science*, 4:433–455, 1994.
- [7] J. L. Gross and T. W. Tucker. *Topological Graph Theory*. Wiley, New York, 1987.
- [8] J. Hage. The membership problem for switching classes with skew gains. *Fundamenta Informaticae*, 39(4):375–387, 1999.
- [9] J. Hage. *Structural Aspects Of Switching Classes*. PhD thesis, Leiden Institute of Advanced Computer Science, 2001. <http://www.cs.uu.nl/people/jur/2s.html>.
- [10] F. Harary. *Graph Theory*. Addison Wesley, 1972.
- [11] J. J. Seidel. A survey of two-graphs. In *Intern. Coll. Teorie Combinatorie (Roma, 1973)*, volume I, pages 481–511, Rome, 1976. Accad. Naz. Lincei. Reprinted in [2].
- [12] T. Zaslavsky. Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory, Ser. B*, 47:32–52, 1989.