The Membership Problem for Switching Classes with Skew Gains

Jurriaan Hage
Dept. of Computer Science, Leiden University
P.O. Box 9512, 2300 RA Leiden, The Netherlands,
e-mail: jhage@wi.leidenuniv.nl
home-page: http://www.wi.leidenuniv.nl/~jhage

January 12, 1998

Abstract

Gain graphs are graphs in which each edge has a gain (a label from a group so that reversing the direction of an edge inverses the gain). In this paper we take a generalized view of gain graphs in which the label of an edge is related to the label of the reverse edge by an involution, i.e., an anti-automorphism of order at most two. Switching classes are equivalence classes of gain graphs under the switching operation.

The membership problem is important to the theory of switching classes. The problem is to decide for two given gain graphs, on the same underlying graph and having gains from the same group, whether or not they belong to the same switching class. We give efficient and often optimal algorithms.

If a certain gain graph in the switching class has only abelian labels, then we can reduce the membership problem to an elegant problem on involutions.

Finally we show that the word problem can be reduced to the general membership problem, thereby establishing undecidability of the latter for some groups.

1 Introduction

Gain graphs are graphs in which each edge has a gain (a label from a group so that reversing the direction of an edge inverses the gain), see Zaslavsky in [12]. In this paper we take a generalized view of gain graphs in which the label of an edge is related to the label of the reverse edge by an involution, i.e. an anti-automorphism of order at most two as in Hage and Harju [5].

In the special case when the group equals $\mathbf{Z}_2$, gain graphs are known as signed graphs the study of which was originated by Harary [6]. For results on signed graphs, we refer to papers by Zaslavsky, [11] and [10].

The present work is mainly motivated by the study of dynamic labeled 2-structures of Ehrenfeucht and Rozenberg [3].

We assume that the reader is familiar with the basic notions of graph theory and group theory. We give now a number of basic notions in order to establish the notation and unambiguous terminology for this paper.

Let Γ be a group. A function $\delta : \Gamma \to \Gamma$ is called an *involution*, if it is an anti-automorphism of order at most two, that is, 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 involution δ.

Let V be a finite set, and denote by $E_2(V) = \{(x,y) \mid x, y \in V, x \neq y\}$ the set of all nonreflexive ordered pairs of V. For a pair $e = (x,y) \in E_2(V)$, we denote by e^{-1} its reverse pair (y, x).

We consider graphs $G = (V, E)$ where the set of edges $E \subseteq E_2(V)$ satisfies the condition: if $e \in E$ 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 $E(G)$ and $V(G)$ to denote E and V respectively.

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

$$g(e^{-1}) = \delta(g(e)) \tag{1}$$

for all $e \in E$. The class of (Γ, δ)-gain graphs on G will be denoted by $\mathbf{L}_G(\Gamma, \delta)$ or simply by $\mathbf{L}_G$. The set of gains of g is

$$A(g) = \{g(e) \mid e \in E(G)\} \subseteq \Gamma \ .$$

We adopt in a natural way some of the terminology of graph theory for graphs with skew gains.

For each function $\sigma : V \to \Gamma$, called a *selector*, we associate with g a new (Γ, δ)-gain graph g^σ on $G = (V, E)$ by letting for each $(x, y) \in E$,

$$g^\sigma(x, y) = \sigma(x) \cdot g(x, y) \cdot \delta(\sigma(y)) \ .$$

We use $\mathrm{S}(V,\Gamma)$, or simply S, to denote Γ^V, the set of selectors from V into Γ.

The relation $g \sim h$, which holds if there exists a selector σ such that $h = g^\sigma$, is an equivalence relation on the (Γ, δ)-gain graphs, and we denote by

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

the equivalence class of g determined by this relation; it is called the *switching class* generated by g (see Seidel [9]).

Further definitions and basic results on (Γ, δ)-gain graphs are given in the preliminaries. We shall now give a short account of the results of later sections.

This paper considers the membership problem for switching classes of gain graphs, i.e., given two gain graphs $g, h \in \mathbf{L}_G(\Gamma, \delta)$, decide whether $h \in [g]$. This problem is very well motivated from the point of view of dynamic labeled 2-structures (see [3]). The original motivation behind them was the formalization of a specific kind of networks. From the point of view of the networks the question whether or not a network can end up in a specific configuration is certainly a central question.

We start with a general treatment of the problem that reduces its complexity. Based on this we construct an efficient algorithm of time complexity $\mathcal{O}(k|V(G)|^2)$ for finite groups of order k.

Then, we look at various optimizations that can be made for particular kinds of groups, involutions and underlying graphs. In this way we get more efficient algorithms for abelian groups when the involution is the group inverse or the underlying graph is bipartite. Also, if the switching class contains a certain gain graph with only abelian labels, then we reduce the membership problem to an elegant problem on involutions. This reduction takes $\mathcal{O}(|E(G)|)$ time.

Finally, we prove that in general there are groups that yield an undecidable membership problem already for very simple underlying graphs. In particular, these groups have an undecidable word problem.

2 Preliminaries

Let $\mathbf{Z}$, $\mathbf{R}$ and $\mathbf{R}^+$ be the sets of *integers*, *reals* and *positive reals*, respectively. For a finite set X, $|X|$ denotes its *cardinality*. The identity function on X is denoted id_X or simply id.

For a group Γ we denote its identity element by 1_Γ. The *centralizer* of $A \subseteq \Gamma$ $C(A) = \{x \in \Gamma \mid ax = xa \text{ for all } a \in A\}$ contains all elements of Γ that commute with each element of A. Note that $C(A)$ is a subgroup of Γ. The *center* of Γ, denoted $Z(\Gamma)$, equals $C(\Gamma)$.

In this paper Γ always denotes a group, δ an involution of Γ and $G = (V, E)$ a graph.

We note first that g^σ satisfies the reversibility condition (1), by the following lemma.

Lemma 2.1
For each $g \in \mathbf{L}_G(\Gamma, \delta)$ and selector $\sigma : V \to \Gamma$, also $g^\sigma \in \mathbf{L}_G(\Gamma, \delta)$.
Proof:
Indeed,

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

which shows the claim. □

The set S(V,Γ) of selectors can be made into a group in a natural way by defining for all selectors σ and τ,

$$(\sigma\tau)(u) = \sigma(u)\tau(u)$$

for all $u \in V$, but note that $g^{\sigma\tau} = (g^\tau)^\sigma$. Hence, S($V$,$\Gamma$) is a group that *acts on* the (Γ, δ)-gain graphs, that is, S(V,Γ) can be thought of as a permutation group on $\mathbf{L}_G(\Gamma, \delta)$. It follows then, that each switching class $[g]$ is generated by each of its elements:

Lemma 2.2
For all σ and g, $[g^\sigma] = [g]$. □

In the group S(V,Γ) the *trivial selector* σ_1, for which $\sigma_1(u) = 1_\Gamma$ for all $u \in V$, is the group identity of S(V,Γ); and the inverse of a selector σ, denoted σ^{-1}, is found by inverting the selected values in the vertices, that is, $\sigma^{-1}(u) = \sigma(u)^{-1}$.

For $g, g' \in \mathbf{L}_G$, $g^{-1} \in \mathbf{L}_G$ is such that $g^{-1}(u,v) = g(u,v)^{-1}$ for all $(u,v) \in E(G)$, and gg' is defined edgewise by

$$(gg')(u,v) = g(u,v)g'(u,v) \text{ for all } (u,v) \in E(G) \ .$$

Note that gg' does not necessarily satisfy (1).

A *rooted tree* is a tree T with an indicated vertex $u = \mathrm{root}(T)$, so that for each vertex v, there exists a unique path from u to v in T. For simplicity we will refer to rooted trees as trees. We say that a vertex $v \in V(T)$ is *odd (even)* if the distance (the number of edges on the shortest path) between $\mathrm{root}(T)$ and v is odd (even) in T. If two vertices are both even or both odd, they have the same *parity.* We use $\mathrm{odd}(T)$ and $\mathrm{even}(T)$ to refer to the sets of odd and even vertices of T respectively.

For graphs $G = (V, E)$ and $G' = (V, E')$ we denote with $G - G'$ the graph $(V, E - E')$ that contains all edges of G that are not in G'. This definition generalizes in the obvious way to gain graphs.

Let $g \in \mathbf{L}_G$ and let $t \in \mathbf{L}_T$ where T is a spanning tree of G. Define the selector $\sigma_{g,t}$ recursively,

$$\sigma_{g,t}(u) = \begin{cases} 1_\Gamma & \text{if } u = \mathrm{root}(T) \\ \delta(t(v,u)\sigma_{g,t}(v)^{-1}g(v,u)^{-1}) & \text{otherwise, where } v \text{ is the father of } u \text{ in } T \end{cases}$$

Additionally, define $g_t = g^{\sigma_{g,t}}$.

By an easy induction on the distance of a vertex from the root of the tree, the following can be proved (see, for instance, Zaslavsky [11]).

Lemma 2.3
Let $g \in \mathbf{L}_G$ and let $t \in \mathbf{L}_T$ where T is a spanning tree of G. For all $e \in E(T)$, $g_t(e) = t(e)$. $\square$

If t is only labeled by 1_Γ we write g_T instead of g_t and call g_T the *T-canonical (Γ, δ)-gain graph* of g. We sometimes say that T is *1_Γ-labeled* in g_T.

3 General theory

Let $g, h \in \mathbf{L}_G(\Gamma, \delta)$. If G consists of connected components G_i, for $1 \leq i \leq c$ (inducing in this way components g_i and h_i of g and h respectively), we can reduce the problem $g \in [h]$ to the connected case: $g \in [h]$ if and only if $g_i \in [h_i]$ for $1 \leq i \leq c$. Therefore we may concentrate on connected graphs G.

Now, let $G = (V, E)$ be connected and let T be a spanning tree of G. By Lemma 2.3 there exists a T-canonical (Γ, δ)-gain graph $g_T \in [g]$.

For a tree T, a selector σ is *alternating* in T, if for every edge (u, v) in the tree, $\sigma(u) = \delta(\sigma(v))^{-1}$. We denote with $\sigma_{T,a}$ a selector that is alternating in T and selects a in $\mathrm{root}(T)$. The selector $\sigma_{T,a}$ is clearly well defined and unique. The central property of alternating selectors is that applying one to a T-canonical (Γ, δ)-gain graph yields a T-canonical (Γ, δ)-gain graph.

For $u \in V$, a selector σ with $\sigma(u) = 1_\Gamma$ is a *u-selector.* For a given $u \in V$, the u-selectors form a subgroup of the group of selectors and, hence, they partition the switching classes into subclasses. We use $\langle g \rangle_u = \{g^\sigma \mid \sigma \text{ is a } u\text{-selector}\}$ to denote the *u-subclass* generated by g. Note that $\sigma_{g,t}$ is a $\mathrm{root}(T)$-selector.

The u-subclass generated by g contains exactly one (Γ, δ)-gain graph with T 1_Γ-labeled, which is g_T. This follows from the fact that the only selector that is both alternating and a u-selector is the trivial selector.

A property of a u-subclass is that it contains as many elements as there are u-selectors, i.e., each u-selector maps to a different (Γ, δ)-gain graph.

Lemma 3.1
For each $g \in \mathbf{L}_G$ and vertex u of G, $|\langle g \rangle_u| = |\{\sigma \mid \sigma \text{ is a } u\text{-selector}\}|$.
Proof:
Let σ and τ be different u-selectors. Because they differ on at least one vertex, say z, and they correspond on the value selected in u, there must be an edge (v, w) on the path between u and z such that they select the same value in v and different value on w. Consequently $g^\sigma(v, w) \neq g^\tau(v, w)$. □

Corollary 3.2
If Γ is of order k and G has n vertices, then $|\langle g \rangle_u| = k^{n-1}$ for $u \in V(G)$. □

A simple induction on the distance of a vertex to $\text{root}(T)$ proves the following lemma.

Lemma 3.3
If $g_T^\sigma = g_T$ for a selector σ, then σ is alternating in T. □

Summarizing, we have found that every switching class on a connected graph consists of a number of u-subclasses that all have the same size. We also know that each such subclass is generated by a T-canonical (Γ, δ)-gain graph. The only remaining question is to decide which u-subclasses constitute a switching class or, equivalently, which T-canonical (Γ, δ)-gain graphs belong to the same switching class.

In the following we will try to formulate an answer to the following question: given two different T-canonical (Γ, δ)-gain graphs, is there a selector that maps the one into the other? This question is simpler than the original, because we need not consider u-selectors. In fact, we need only consider alternating selectors.

We introduce some definitions based on definitions from Hage and Harju [5]. We denote by

$$\text{EO}_T(g) = \{a \in \Gamma \mid g(u, v) = a, (u, v) \in E(G), u \in \text{even}(T), v \in \text{odd}(T), (u, v) \notin E(T)\}$$

the labels of the edges of g that are not in T and that start in an even and end in an odd vertex with respect to $\text{root}(T)$. In a similar way we define

$$\text{EE}_T(g) = \{a \in \Gamma \mid g(u, v) = a, (u, v) \in E(G), u, v \in \text{even}(T)\} \text{ and}$$

$$\text{OO}_T(g) = \{a \in \Gamma \mid g(u, v) = a, (u, v) \in E(G), u, v \in \text{odd}(T)\} \ .$$

Further, let

$$C_T^\delta(g) = C(\text{EO}_T(g_T)) \cap \text{OO}_T^\delta(g_T) \cap \text{EE}_T^\delta(g_T)$$

where

$$\text{OO}_T^\delta(g) = \{x \in \Gamma \mid \delta(x)a = ax^{-1} \text{ for all } a \in \text{OO}_T(g)\}, \text{ and}$$

$$\text{EE}_T^\delta(g) = \{x \in \Gamma \mid a\delta(x) = x^{-1}a \text{ for all } a \in \text{EE}_T(g)\} \ .$$

Note that $\text{OO}_T^\delta(g)$ and $\text{EE}_T^\delta(g)$ are subgroups of Γ. Hence, $C_T^\delta(g)$ is also a subgroup of Γ. Note also that if the involution is the group inverse, then the definitions of all three sets that determine $C^{\delta}_{T}(g)$ coincide and $C^{\delta}_{T}(g)$ can be defined as $C(A(g_T - T))$. This is also true if G is bipartite, because then $\mathrm{EE}_T(g)$ and $\mathrm{OO}_T(g)$ are empty, hence $\mathrm{OO}^{\delta}_{T}(g) = \mathrm{EE}^{\delta}_{T}(g) = \Gamma$. We return to this in Section 5.

The following lemma generalizes Lemma 3.2 from Hage and Harju [5].

Lemma 3.4
Let $g \in \mathbf{L}_G$ and let T be a spanning tree of G. Then $g_T = g^{\sigma}_{T}$ if and only if σ is alternating in T and $\sigma(\mathrm{root}(T)) \in C^{\delta}_{T}(g)$.
Proof:
Assume that $g_T = g^{\sigma}_{T}$. By Lemma 3.3, σ is alternating.

Let (v, w) be an edge of $g_T - T$ and let $a = g_T(v, w)$. We must consider three cases: both v and w are odd, both are even, or one is even and the other is odd (all with respect to $r = \mathrm{root}(T)$). We will consider here only the case that both vertices are even. The other cases are treated similarly.

Because $g_T(v, w) = g^{\sigma}_{T}(v, w)$, $a = \sigma(v)a\delta(\sigma(w))$. This together with the fact that σ is alternating and the positions of v and w relative to the root, yields $a = \sigma(r)a\delta(\sigma(r))$. Hence, $\sigma(r)^{-1}a = a\delta(\sigma(r))$ or equivalently, $\sigma(r) \in \mathrm{EE}^{\delta}_{T}(g_T)$. The claim now follows.

The converse is proved similarly. □

The following result characterizes the subclasses that constitute a switching class, straightforwardly generalizing Theorem 3.6 from Hage and Harju [5]. Recall that a transversal of a subgroup is such that it contains one element of each of its cosets.

Theorem 3.5
Let T be a spanning tree of G and $g \in \mathbf{L}_G$ such that T is 1_Γ-labeled in g. Also, let $\mathcal{T}$ be a transversal of $C^{\delta}_{T}(g)$. Then $[g] = \bigcup_{a \in \mathcal{T}} \langle g^{\sigma_{T,a}} \rangle_{\mathrm{root}(T)}$. Moreover, all $g^{\sigma_{T,a}}$ are different.
Proof:
We apply Lemma 3.4 to conclude that $g^{\sigma_{T,a}} = g^{\sigma_{T,b}}$ if and only if $b^{-1}a \in C^{\delta}_{T}(g)$, using the fact that $C^{\delta}_{T}(g)$ is a subgroup of Γ. Then $g^{\sigma_{T,a}}$ $(a \in \mathcal{T})$ are different, because $\mathcal{T}$ is a transversal of $C^{\delta}_{T}(g)$.

For $h = g^{\tau}$, let $b = \tau(u)$. Now, $\tau = (\tau{\sigma_{T,b}}^{-1})\sigma_{T,b}$, where $\tau{\sigma_{T,b}}^{-1}(u) = 1_\Gamma$. Hence, $h \in \langle g^{\sigma_{T,b}} \rangle_{\mathrm{root}(T)}$. The claim follows, because for b there exists an $a \in \mathcal{T}$ such that $g^{\sigma_{T,a}} = g^{\sigma_{T,b}}$ by the first part of the proof. □

For the membership problem, the following corollary can be useful.

Corollary 3.6
Let $g, h \in \mathbf{L}_G$ and let T be a spanning tree of G. Also, let $\mathcal{T}$ be a transversal of $C^{\delta}_{T}(g)$. Then $g \in [h]$ if and only if $h_T = g_T^{\sigma_{T,a}}$ for some $a \in \mathcal{T}$. □

Sumarizing, we get the situation as depicted below

$$
\begin{tikzcd}
 & G \text{ with spanning tree } T \arrow[dl] \arrow[dr] & & \\
g \arrow[d] & \in & {[h]?} \arrow[d] & k^n \text{ selectors} \\
g_T \arrow[d] & \in & {[h_T]?} \arrow[d] & k \text{ alternating selectors} \\
g_T^{\sigma_{T,a}} & = & h_T? & k/C_T^{\delta}(g_T) \text{ alternating selectors}
\end{tikzcd}
$$

4 An algorithm for finite groups

In this section we develop an algorithm for finite groups (and, if necessary, it can be modified to work for infinite groups as long as the transversal $\mathcal{T}$ (see Theorem 3.5) is finite). We only give the algorithm for connected graphs; for unconnected graphs it should be applied to each component.

Algorithm 4.1

`SameSwitchingClass?` (g, h)
(* Here, g and h are elements of $\mathbf{L}_G(\Gamma, \delta)$. *)
`begin`
    $T =$ `a spanning tree of` G;
    `Compute` g_T `and` h_T;
    `for all` $a \in \Gamma$ `do`
        `if` $h_T = g_T^{\sigma_{T,a}}$ `then return true;`
    `od;`
    `return false;`
`end;`

Please note that we do not use the theory to the fullest, in the sense that the transversal $\mathcal{T}$ of Theorem 3.5 is not used. The reason is that constructing this transversal in a straightforward way takes as much time as the entire loop, so it is better to just apply selectors $\sigma_{T,a}$ for all values $a \in \Gamma$ and not just $a \in \mathcal{T}$.

Recall that the *cyclomatic number* of a graph is $e - n + c$, where e, n and c are the number of edges, vertices and components of the graph respectively.

Theorem 4.2
Let Γ be a finite group of order k and G a graph on n vertices. Then the membership problem for (Γ, δ)-gain graphs on G is in $\mathcal{O}(k \max(\xi, n))$, where ξ is the cyclomatic number of G.

Proof:
We determine the complexity of the algorithm by counting the number of edge comparisons. We apply at most k selectors and must change and compare at most ξ edges every single time. If $\xi > n$ we get a complexity of $\mathcal{O}(k\xi)$, but if $\xi < n$, then constructing the selectors $\sigma_{T,a}$ dominates and we get a time complexity of $\mathcal{O}(kn)$. So, depending on their respective sizes we express the complexity in either the number vertices, or the number of edges outside the chosen spanning tree. □

In the introduction we mentioned the problem of determining whether a particular configuration can occur in a network. We now address this problem, which is in fact equivalent to determining whether a gain graph can be embedded into some gain graph in a switching class.

In the embedding problem we are given two (Γ, δ)-gain graphs g and h on, possibly different, graphs G and H. The question is whether there exists a $j \in [g]$ such that h can be embedded in j, that is, whether there exists an injective function $\phi : V(H) \to V(G)$ such that

$$h(u, v) = j(\phi(u), \phi(v))$$

for all edges $(u, v) \in E(h)$.

Once we have fixed an injection ϕ from h into g, we can restrict g to g' by removing edges that are not in h', where h' is the image of h under ϕ. (Note that if h' contains an edge that is not present in g', then we know that ϕ does not embed h in g.) With the algorithm described in this section we can answer the question whether $h' \in [g']$. If the answer is affirmative, then we have our embedding ϕ; if the answer is negative, then we should try the next embedding. Note that although the membership check can be efficient, there may be many possible injections. In fact, the embedding problem is a hard one, even if restricted to $\mathbf{Z}_2$ (see Ehrenfeucht, Hage, Harju and Rozenberg [2]).

5 Improvements in the abelian case

In this section we first look at (Γ, δ)-gain graphs with labels from the center of Γ. The involution is first still arbitrary, but later on we also give a further optimization when the involution is the group inverse. The theory in this section differs from the treatment in the previous section, because here we shall *construct* the selector that maps g_T into h_T if it exists, instead of *trying* a number of different selectors. We could have chosen to just improve the algorithm of the previous section for abelian groups, but the treatment here results in a more widely applicable algorithm, which we shall demonstrate by means of an example.

5.1 Improvements when g_T has abelian labels

We call g *abelian* if $A(g) \subseteq Z(\Gamma)$. It is called *inversive*, if the involution is the group inverse. Note that the latter property does not depend on $A(g)$, but is a statement about the context of g.

Let G be a graph, let T be a spanning tree of G and let $g, h \in \mathbf{L}_G$.

To improve on our algorithm for checking that $h \in [g]$, we need to know that at least one of g_T and h_T is abelian. We can assume, without loss of generality, that this is the case for g_T.

Let σ be a selector. Define $G_\sigma \in \mathbf{L}_G$ such that

$$G_\sigma(u, v) = \sigma(u)\delta(\sigma(v)), \text{ for all } (u, v) \in E(G) \ .$$

It is easy to prove that if $j \in \mathbf{L}_G$ is abelian, $jG_\sigma = j^\sigma$. This means that applying a selector is equivalent to applying the group operator edgewise.

Now, assume that $h_T \in [g_T]$ and let σ be the corresponding alternating selector.

Let $(v, w) \in E(G)$. If v and w are of different parity, then $G_\sigma(v, w) = \sigma(v)\delta(\delta(\sigma(v)^{-1})) = 1_\Gamma$. If $v, w \in \text{even}(T)$ then $h_T(v, w) = g_T^\sigma(v, w) = \sigma(v)\delta(\sigma(v))g_T(v, w)$, or equivalently, $G_\sigma(v, w) = a$, where a can be written as $b\delta(b)$ for some $b \in \Gamma$. It can then be deduced that the label between vertices from $\text{odd}(T)$ should be a^{-1}.

Summarizing, for labels of the edges of $G_\sigma = {g_T}^{-1}h_T$ we have the following situation for some $a \in \Gamma$.

parity	odd	even
odd	a^{-1}	1_Γ
even	1_Γ	a

The *decomposable set* of (Γ, δ) is the set $\text{dec}_{(\Gamma,\delta)} = \{a \in \Gamma \mid a = b\delta(b) \text{ for some } b \in \Gamma\}$. For example, $\text{dec}_{(\mathbf{Z},id)}$ is the set of even numbers, $\text{dec}_{(\mathbf{R},id)}$ where the operation is addition equals $\mathbf{R}$, and $\text{dec}_{(\mathbf{R}^+,id)}$ where the operation is multiplication equals $\mathbf{R}^+$.

Since G_σ can contain $a, a^{-1}, \delta(a)$ and $\delta(a^{-1})$ for some $a \in \Gamma$ we show next that if one of these is in $\text{dec}_{(\Gamma,\delta)}$, then they all are. Consequently, it does not matter which of these is chosen to be a.

Lemma 5.1
For a group Γ, $\text{dec}_{(\Gamma,\delta)}$ is closed under taking inverses. Moreover, $\delta(a) = a$ for all $a \in \text{dec}_{(\Gamma,\delta)}$ and if Γ is abelian, then $\text{dec}_{(\Gamma,\delta)}$ is a subgroup of Γ. $\square$

The above reasoning can also be applied in the reverse direction, which yields the following result.

Theorem 5.2
For g such that g_T is abelian, $h \in [g]$ if and only if the edges between vertices of different parity in ${g_T}^{-1}h_T$ are labeled by 1_Γ and there exists a $b \in \Gamma$ such that the edges between two even vertices (with respect to T) are labeled by $b\delta(b)$ and the edges between two odd vertices (again, with respect to T) are labeled by $(b\delta(b))^{-1}$. $\square$

Theorem 5.3
For $g, h \in \mathbf{L}_G$ with g_T abelian, $h \in [g]$ reduces in time $\mathcal{O}(|E(G)|)$ to the characteristic function of $\text{dec}_{(\Gamma,\delta)}$. $\square$

Example 5.4
If $\Gamma = \mathbf{Z}$, the involution δ is the identity, and the underlying graph G is not bipartite, then $C_T^\delta(g)$ equals $\{1_\Gamma\}$, where g is any element of $\mathbf{L}_G$. Consequently, the transversal $\mathcal{T}$ as defined in Section 3 equals Γ and thus is infinite; hence, the switching class consists of infinitely many subclasses. On the other hand, the switching class is not equal to $\mathbf{L}_G(\mathbf{Z}, \delta)$. This means that an algorithm like that in Section 4, even if it uses the information about the transversal, is not able to solve the membership problem in finite time. On the other hand, the theory developed in this section can be applied, since we know that $\text{dec}_{(\mathbf{Z},id)}$ contains exactly the even numbers.

5.2 Switching classes consisting of one subclass

In the following we will investigate in which cases a switching class consists of only one u-subclass. The advantage is that in these cases we need not apply any selectors, because $h \in [g]$ if and only if $g_T = h_T$.

By the fact that $\mathrm{dec}_{(\Gamma,^{-1})} = \{1_\Gamma\}$, the following holds.

Lemma 5.5
Let $g \in \mathbf{L}_G$. If g_T is abelian and g is inversive then $\langle g \rangle_{\mathrm{root}(T)} = [g]$. □

The same conclusion can be drawn if the underlying graph is bipartite.

Lemma 5.6
Let $g \in \mathbf{L}_G$. If g_T is abelian and G is bipartite then $\langle g \rangle_{\mathrm{root}(T)} = [g]$.
Proof:
Theorem 5.2 of Hage and Harju [5] states that $|[g]| = k^n/|C(A(g_T))|$, where k is the order of Γ and $n = |E(G)|$. The result now follows from Corollary 3.2. □

Theorem 5.7
Let $g \in \mathbf{L}_G$. If g_T is abelian, and G is bipartite or g is inversive, then $\langle g \rangle_{\mathrm{root}(T)} = [g]$.□

Lemma 5.8
Let $g \in \mathbf{L}_G$. If g_T is abelian and $\langle g \rangle_{\mathrm{root}(T)} = [g]$ then G is bipartite or g is inversive.
Proof:
Assume that G is not bipartite; we prove that g is inversive.

Let (v, w) be an edge in g_T between two vertices of the same parity. Because g_T is the only (Γ, δ)-gain graph in $[g]$ in which T is 1_Γ-labeled, it holds for every alternating selector σ that $g_T(v, w) = g_T^\sigma(v, w) = \sigma(v) g_T(v, w) \delta(\sigma(w)) = \sigma(v)\delta(\sigma(v)) g_T(v, w)$ and this holds if and only if $\sigma(v)^{-1} = \delta(\sigma(v))$. □

Lemma 5.9
Let $g \in \mathbf{L}_G$. If $\langle g \rangle_{\mathrm{root}(T)} = [g]$ then h is abelian where h equals g_T, but with all edges between vertices of the same parity deleted.
Proof:
Let $(v, w) \in E(h)$, with h defined as above. Because for all alternating selectors σ, $h(v, w) = h^\sigma(v, w) = \sigma(v) h(v, w) \delta(\delta(\sigma(v)^{-1})) = \sigma(v) h(v, w) \sigma(v)^{-1}$, it follows that $h(v, w)$ commutes with each element of Γ. Hence, h is abelian. □

Lemma 5.10
Let $g \in \mathbf{L}_G$. If $\langle g \rangle_{\mathrm{root}(T)} = [g]$, and g is inversive or G is bipartite then g_T is abelian.
Proof:
If G is bipartite then we apply Lemma 5.9 to find that g_T is abelian.

Let g be inversive and not bipartite. First of all, if g is not abelian, this is because of an edge between vertices of the same parity, by Lemma 5.9. So let (v, w) be an edge between vertices of the same parity. Then, because $\langle g \rangle_{\mathrm{root}(T)} = [g]$, for an alternating selector σ, $g_T(v, w) = g_T^\sigma(v, w) = \sigma(v) g_T(v, w) \delta(\sigma(v)) = \sigma(v) g_T(v, w) \sigma(v)^{-1}$. For this to hold, $g_T(v, w)$ must commute with each $\sigma(v)$. Hence g_T is abelian. □

Corollary 5.11
Let $g \in \mathbf{L}_G$. If $\langle g \rangle_{\mathrm{root}(T)} = [g]$, and g is inversive or G is bipartite, then for all $h \in [g]$ and spanning trees T' of G, $h_{T'}$ is abelian.
Proof:
Let $h \in [g]$. First of all, the bipartiteness of G and the inversiveness of g is independent of the labels of g. Because $\langle g \rangle_{\mathrm{root}(T)} = [g] = [h]$, it holds for all spanning trees T' of G that $\langle h \rangle_{\mathrm{root}(T')} = [h]$. The result now follows from Lemma 5.10. □

In the above we have considered three predicates: g is inversive or bipartite (P1), $\langle g \rangle_{\mathrm{root}(T)} = [g]$ (P2) and g_T is abelian (P3). Through manipulation of the previous lemma's we get the following result.

Corollary 5.12
Each pair of predicates P_i, P_j is equivalent under the condition that the remaining predicate, P_ℓ, holds, where $\{i, j, \ell\} = \{1, 2, 3\}$. □

Example 5.13
If $\Gamma = \mathbf{Z}_3$, $\delta = id$ and G is complete on three vertices, then for any $g \in \mathbf{L}_G(\Gamma, \delta)$, g_T is abelian, but it is possible that $\langle g \rangle_{\mathrm{root}(T)} \neq [g]$.

For the same G, but with group S_3 and involution the group inverse, it is also possible that $\langle g \rangle_{\mathrm{root}(T)} \neq [g]$.

6 Undecidability for arbitrary groups

The following theory about groups, presentations and the word problem comes from Rotman [8].

In the previous sections we omitted the task of specifying the group and the involution. We assumed it was given and that we could compute with it. Usually a group is specified by means of a *presentation*

$$\Gamma = \langle x_1, x_2, \ldots \mid w_1, w_2, \ldots \rangle ,$$

where the x_i are the *generators* and the w_j, being words over $\gamma = \{x_1, {x_1}^{-1}, x_2, {x_2}^{-1}, \ldots\}$, are the *relations*. The idea is that by the relations we define a number of sequences of generators that equal the identity of the group (but in general not all of them). We shall use α to denote the, in general non-injective, mapping of the strings over γ into (an element of) the group.

Every presentation determines a group, but a group can have a number of presentations. Because a presentation of a group should in fact be a parameter to the membership problem, we assume the number of generators and relations to be finite.

The *word problem* for presentations of groups is the following: given a word w over γ, does it define the identity of the group? Novikov and Boone have independently proven that there are finitely presented groups that have no presentation for which the word problem is decidable. In fact, if the word problem is undecidable for a presentation of a group, it is undecidable for all presentations of that group.

With this background we shall continue now by proving that the word problem can be reduced to the membership problem, showing that in general the membership problem for groups specified by means of a presentation is undecidable.

Let w be the word for which we would like to know whether it defines the group identity 1_Γ. Let g_w be the following (Γ, δ)-gain graph on $V = \{u, v, z\}$

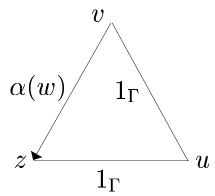

and let T consist of the edges (u, v) and (u, z).

If we can decide whether $g_w \in [d]$, where d is the complete, 1_Γ-labeled (Γ, δ)-gain graph on V, we can solve the word problem for γ, because

$$\alpha(s)\alpha(w)\alpha(s)^{-1} = 1_\Gamma$$

reduces to $\alpha(w) = 1_\Gamma$. Note that by selecting the same value $\alpha(s)$ (specified by means of a word s over γ) in all vertices we guarantee that the selector is alternating (remember that the involution is the group inverse).

Theorem 6.1
There exist pairs of groups and involutions for which the membership problem is undecidable. □

7 Conclusions and future work

A part of our future work will be to determine the general nature of $\mathrm{dec}_{(\Gamma,\delta)}$ and will address questions relating to its existence and the possibilities of its effective construction.

Although we know now that the general problem is undecidable, it is clear that by certain restrictions we might find classes of groups and involutions that always have a decidable membership problem. To give but an example, if we restrict ourselves to the involution being the identity function (and thus restricting ourselves to abelian groups), we find that the problem seems to be easier. We might therefore look at other types of involutions for which good results can be obtained.

Another problem to investigate is the influence the choice of tree has on the results obtained here. More specifically, the question is whether we can easily find T such that g_T is abelian (if such a T exists).

Acknowledgements

The author thanks Tero Harju and Nikè van Vugt for their suggestions and comments.

References

[1] D.G. Corneil and R.A. Mathon. *Geometry and Combinatorics: Selected Works of J.J. Seidel.* Academic Press, Boston, 1991.

[2] A. Ehrenfeucht, J. Hage, T. Harju, and G. Rozenberg. Complexity problems in switching classes of graphs. Technical Report 15, Leiden University, Department of Computer Science, 1997.

[3] A. Ehrenfeucht and G. Rozenberg. An introduction to dynamic labeled 2-structures. In A.M. Borzyszkowski and S. Sokolowski, editors, *Mathematical Foundations of Computer Science 1993*, volume 711 of *Lecture Notes in Computer Science*, pages 156–173, 1993.

[4] J. Hage and T. Harju. The size of 2-classes in group labeled 2-structures. Technical Report 17, Leiden University, Department of Computer Science, 1996. Generalized in [5].

[5] J. Hage and T. Harju. The size of switching classes with skew gains. Technical Report 2, Leiden University, Department of Computer Science, 1997. Generalization of [4].

[6] F. Harary. On the notion of balance of a signed graph. *Michigan Math. J.*, 2:143–146, 1953–1954. Addendum in same journal preceding page 1.

[7] J.S. Rose. *A Course on Group Theory.* Cambridge University Press, Cambridge, 1978.

[8] J.J. Rotman. *The Theory of Groups.* Allyn and Bacon, Boston, 2nd edition, 1973.

[9] J.J. Seidel. A survey of two-graphs. In *Proc. Internat. Colloq. Theorie Combinatorie*, Rome, 1973. Acc. Naz. Lincei. See also [1].

[10] T. Zaslavsky. Characterization of signed graphs. *J. of Graph Theory*, 5:401–406, 1981.

[11] T. Zaslavsky. Signed graphs. *Discrete Applied Math.*, 4:47–74, 1982. Erratum on p. 248 of volume 5.

[12] T. Zaslavsky. Biased graphs. I. Bias, balance, and gains. *J. of Combin. Theory, Ser. A*, 47:32–52, 1989.