Universiteit Leiden
The Netherlands

Decompositions in algebra

Gent, D.M.H. van

Citation

Gent, D. M. H. van. (2024, March 5). *Decompositions in algebra*.
Retrieved from https://hdl.handle.net/1887/3720065

CHAPTER 1

Nonabelian flows in graphs

1.1   Introduction

This chapter is based on [20]. In this chapter, *graph* will mean a simple graph with a finite number of vertices. We consider groups which are not required to be abelian and therefore write our group operations multiplicatively.

With Γ a group and $G = (V, E)$ a graph, we call a map $f\colon V^2 \to \Gamma$ a Γ-*flow in G* if for all $u, v \in V$ we have $f(u, v) = f(v, u)^{-1}$, and $f(u, v) = 1$ if $\{u, v\} \notin E$. This definition agrees with the classical definition of a graph flow when $\Gamma = \mathbb{R}$.

Non-abelian graph flows, i.e. flows where Γ need not be abelian, were first considered by M.J. DeVos in his PhD thesis [9] and later by A.J. Goodall et al. [22] and B. Litjens [36]. They consider graphs embedded on surfaces and ask whether flows exist which are nowhere trivial, i.e. $f(u, v) \neq 1$ if and only if $\{u, v\} \in E$. Although likewise our main result involves planar embeddings of graphs, we instead ask to which extent Kirchhoff's law of conservation holds.

Let $G = (V, E)$ be a graph, Γ a group and f a Γ-flow in G. We call f *tractable* if for each $v \in V$ the subgroup $\langle f(u, v) \,|\, u \in V\rangle$ of Γ is abelian. For tractable f we define the *excess* $e_f\colon V \to \Gamma$ to be the map given by $v \mapsto \prod_{u \in V} f(u, v)$ and we say f is *conserving* in v if $e_f(v) = 1$. In the classical case, we have the following lemma.

Lemma 1.1.1. *Let Γ be an* abelian *group, let f be a Γ-flow in a graph $G = (V, E)$ and let $w \in V$. If f is conserving in all vertices of $V \setminus \{w\}$, then f is conserving in w.*

Proof. We have

$$e_f(w) = \prod_{v \in V} e_f(v) = \prod_{(u,v) \in V^2} f(u, v) = \prod_{\{u,v\} \in E} f(u, v)f(v, u) = 1,$$

so f is conserving in w.   □

We will show that Lemma 1.1.1 can fail for non-abelian Γ. We say a flow f *leaks* if it is tractable and conserving in all but precisely one vertex and we call a graph G *leak-proof* if there exist no flows in G that leak for any group Γ. Our main result, proven in Section 1.4, is as follows.

Theorem 1.4.3. *A graph is leak-proof if and only if it is planar.*

We say a flow f of G has a *binary leak* at distinct vertices $u, v \in V$ if it is tractable and conserving in all vertices of $V \setminus \{u, v\}$ while $e(u)e(v) \neq 1$. Here u and v can be thought of as a source and sink of the flow. We call

G *binary leak-proof* if no binary leaking flows exist for G. Analogously to Lemma 1.1.1 one can show that a flow cannot have a binary leak when the group is abelian. We also prove the following analogue to Theorem 1.4.3 in Section 1.5.

Definition 1.1.2. We call a graph $G = (V, E)$ *extra-planar* if for all pairs of distinct $u, v \in V$ the graph $(V, E \cup \{u, v\})$ is planar.

Theorem 1.5.2. *A graph is binary leak-proof if and only if it is extra-planar.*

Instead of studying leak-proof graphs, one could also study leak-proof groups, where we call a group Γ *leak-proof* if for all graphs $G = (V, E)$ no tractable flows $f \colon V^2 \to \Gamma$ of G leak. Theorem 1.4.3 shows that the decision problem 'Is this graph leak-proof?' can be decided in time $O(|V|)$, as Hopcroft and Tarjan gave an algorithm to test graph planarity in [24] of this complexity. For leak-proof groups, we prove the following in Section 1.6.

Theorem 1.6.4. *The decision problem 'Is this finite group leak-proof?' is decidable.*

The present work, in particular Theorem 1.5.2, was inspired by a problem the author encountered in his Master's thesis [17] on graded rings. Here a flow with a binary leak gives rise to an example (Example 2.17 of [17]) of an efficient ring grading with a non-abelian group that cannot be replaced by an abelian group.

1.2 Definitions and basic properties

We briefly go through some basic definitions. Let $G = (V, E)$ be a graph. With $H = (W, F)$ a graph we call a map $f \colon V \to W$ a *morphism* from G to H if $f[E] \subseteq F$. We call this f an *embedding* if it is injective and an *isomorphism* if f and its induced map $E \to F$ are bijections. A *path from $u \in V$ to $v \in V$* in G is a finite sequence of vertices $(x_0, \ldots, x_n)$ for some $n \in \mathbb{Z}_{\geq 0}$ such that $x_0 = u$, $x_n = v$ and $\{x_i, x_{i+1}\} \in E$ for all $0 \leq i < n$. We call this path *non-trivial* if $n > 0$ and *closed* if $x_0 = x_n$. We say $u \in V$ is *connected* to $v \in V$ in G if there exists a path from u to v in G. The 'is connected to' relation is an equivalence relation on V and we call its equivalence classes the *connected components* of G. For $u \in V$ we call $v \in V$ a *neighbor* of u if $\{u, v\} \in E$ and we write $N_G(u) \subseteq V$ for the set of neighbors of u. An edge $\{u, v\} \in E$ is called a *bridge* if all paths in G from

u to v contain the edge $\{u,v\}$. A *forest* is a graph in which every edge is a bridge.

We now give some facts about (non-)planar graphs.

Definition 1.2.1. For $A, B \in \mathbb{R}^2$ write $\ell(A, B)$ for the line $\{tA + (1 - t)B \mid t \in (0,1)\}$. Let $G = (V, E)$ be a graph. A *planar embedding* of G is an injective map $\varepsilon \colon V \to \mathbb{R}^2$ such that for all distinct $\{a, b\}, \{c, d\} \in E$ we have $\ell(\varepsilon(a), \varepsilon(b)) \cap \ell(\varepsilon(c), \varepsilon(d)) = \varnothing$, and for all $\{a, b\} \in E$ we have $\ell(\varepsilon(a), \varepsilon(b)) \cap \varepsilon[V] = \varnothing$. We call G *planar* if it has a planar embedding.

The above definition of a planar embedding has been simplified for our purposes, which is justified by Fáry's Theorem [13].

Definition 1.2.2. Let $G = (V, E)$ be a graph with a planar embedding ε. The *orientation* of (G, ε) at $v \in V$ is the clockwise permutation $\rho_\varepsilon(v)$ of $N_G(v)$. A *boundary walk* of (G, ε) is a non-trivial closed path $(x_0, x_2, \ldots, x_n)$ in G such that for all $i, j \in \mathbb{Z}/n\mathbb{Z}$ we have $x_{i+2} = \rho_\varepsilon(x_{i+1})(x_i)$ and if $(x_i, x_{i+1}) = (x_j, x_{j+1})$, then $i = j$.

Lemma 1.2.3. *Let ε be a planar embedding of a graph $G = (V, E)$ and let $p = (u_1, u_2, \ldots, u_n)$ be a boundary walk. If $(u_i, u_{i+1}) = (u_{j+1}, u_j)$ for some $i, j \in \mathbb{Z}/n\mathbb{Z}$, then $\{u_i, u_j\}$ is a bridge.*

Proof. To show that $e = \{u_i, u_j\}$ is a bridge, it suffices to show that u_i and u_j are disconnected in the graph $G' = (V, E')$ with $E' = E \setminus \{e\}$. Note that $a, b \in V$ are connected in G' if and only if $\varepsilon(a)$ and $\varepsilon(b)$ are connected in the topological space $X = \varepsilon[V] \cup \bigcup_{\{x,y\} \in E'} \ell(\varepsilon(x), \varepsilon(y))$. Hence it suffices by the Jordan curve theorem to show that there exists a loop C in $\mathbb{R}^2 \setminus X$ separating u_i and u_j, as any path from u_i to u_j must intersect this loop.


Figure 1.1: Boundary walk

We informally construct this loop as follows (see Figure 1.1). Place yourself at the midway point between u_i and u_j. Walk along the path p in G in

the direction of u_j and while doing so draw a continuous curve C on your left hand side, being careful not to let C intersect itself or the graph. That this is possible follows from the definition of a boundary walk. Stop once you have reached your starting point for the first time again, and note that this time you are facing u_i by the assumption that $(u_i, u_{i+1}) = (u_{j+1}, u_j)$. Thus on your right hand side is the start of your curve C, and connect the endpoints, crossing $\ell(\varepsilon(u_i), \varepsilon(u_j))$ once. Then C satisfies the requirements, so e is a bridge. □

Definition 1.2.4. Let $G = (V, E)$ be a graph. We call a subgraph $H = (W, F)$ of G a *spanning forest* if it is a forest and $W = V$. For a spanning forest $H = (W, F)$ of G we define $G_H = (C, D)$ to be the *contraction of H in G*, where C is the set of connected components of H and $D = \{\{X, Y\} \in \binom{C}{2} \mid (\exists\, u \in X,\, v \in Y)\, \{u, v\} \in E\}$. A graph M is a *minor* of G if it can be embedded in some contraction of G.

Note that the 'is a minor of' relation is a partial order (up to graph isomorphism). In particular, if I is a minor of H and H is a minor of G, then I is a minor of G. Write K_5 for the complete graph on 5 vertices and $K_{3,3}$ for the complete bipartite graph on 3 and 3 vertices.

Proposition 1.2.5 (Kuratowski, Theorem 4.4.6 in [10]). *A graph G is planar if and only if G does not have K_5 or $K_{3,3}$ as a minor.*

1.3 Non-planar graphs

First we show that all non-planar graphs can leak.

Lemma 1.3.1. *A graph is leak-proof if and only if all its subgraphs are leak-proof.*

Proof. Since each graph is its own subgraph, the implication ($\Leftarrow$) is trivial. Let $G = (V, E)$ be a graph with a subgraph $H = (W, F)$ and assume that there exists some group Γ with a leaking Γ-flow $g\colon W^2 \to \Gamma$ of H. Then we consider $f\colon V^2 \to \Gamma$ by taking $f(u, v) = g(u, v)$ when $\{u, v\} \in F$ and $f(u, v) = 1$ otherwise. Then f is a leaking flow for G, proving ($\Rightarrow$). □

Proposition 1.3.2. *A graph is leak-proof if and only if all its minors are leak-proof.*

Proof. Let $G = (V, E)$ be a graph. By Lemma 1.3.1 it suffices to show that if a contraction of a spanning tree H in G admits a leaking flow, then so does G. By induction we may even assume H has only a single edge

$e = \{a, b\}$. Then $(W, F) \cong G_H$ with $W = (V \setminus e) \cup \{e\}$ under the natural isomorphism $e \mapsto e$ and $w \mapsto \{w\}$ for $w \in V \setminus e$. Assume (W, F) admits a flow $f \colon W^2 \to \Gamma$ leaking at $w \in W$ for some group Γ. Let $X = N_G(a) \setminus e$ and $Y = N_G(b) \setminus (e \cup X)$. We define a flow $g \colon V^2 \to \Gamma$ such that for $u, v \in W$ it is given by

$$
\begin{aligned}
g(u, v) &= f(u, v) & u, v \notin e, \\
g(a, u)^{-1} = g(u, a) &= f(u, e) & u \in X, \\
g(v, b)^{-1} = g(b, v) &= f(e, v) & v \in Y, \\
g(b, a)^{-1} = g(a, b) &= \prod_{u \in X \setminus \{b\}} f(u, a),
\end{aligned}
$$

and $g(u, v) = 1$ otherwise. Note that g agrees with f outside of e and that the flows going to e have been divided among a and b. Thus g is tractable and $e_g(u) = e_f(u)$ for $u \notin e$. By definition of $g(a, b)$ we have that $e_g(a) = 1$ and $e_g(b) = e_f(e)$. Hence g is a leaking flow for G. $\square$

To show that non-planar graphs are not leak-proof, it now suffices by Proposition 1.2.5 to show that K_5 and $K_{3,3}$ admit a leaking flow.

Definition 1.3.3. Let C_2 be the cyclic group with two elements. Let $n \in \mathbb{Z}_{>0}$ and consider the groups $N = C_2^{n+1} = \langle z, x_1, \ldots, x_n \rangle$ and $G = C_2^n = \langle x_{n+1}, \ldots, x_{2n} \rangle$. Consider the action $\varphi \colon G \to \mathrm{Aut}(N)$ defined on the generators as

$$
x_{n+i} \mapsto \left(x_j \mapsto x_j z^{\delta_{ij}}, \quad z \mapsto z \right) \quad \text{for all } 1 \leq i, j \leq n,
$$

where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise. Then define the group $\mathrm{ES}_n = N \rtimes_\varphi G$.

Although we will not use the fact, the ES_n are all *extraspecial 2-groups*.

Example 1.3.4. Consider the utility graph $K_{3,3} = (V, E)$ with $V = \{1, 2, 3, 4, 5, 6\}$ and $E = \{\{u, v\} \mid u \in \{1, 2, 3\}, v \in \{4, 5, 6\}\}$. We define a flow $f \colon V^2 \to \mathrm{ES}_2$ which we specify by an ES_2-valued (symmetric) matrix where the omitted entries are trivial:

$$
f = \left(
\begin{array}{ccc|ccc}
 & & & x_1 & x_2 & x_1 x_2 \\
 & & & x_4 & x_3 & x_4 x_3 \\
 & & & x_1 x_4 & x_2 x_3 & x_1 x_4 x_2 x_3 \\
\hline
x_1 & x_4 & x_1 x_4 & & & \\
x_2 & x_3 & x_2 x_3 & & & \\
x_1 x_2 & x_4 x_3 & x_1 x_4 x_2 x_3 & & &
\end{array}
\right).
$$

For the first 5 columns it is easy to see that multiplying the first two non-trivial entries yields the third. Thus for the first five vertices v we have $\langle f(u,v) \, | \, u \in V \rangle \cong C_2^2$, which is abelian, and $e_f(v) = 1$. For $v = 6$ we observe that $(x_1 x_2)(x_4 x_3)(x_1 x_4 x_2 x_3) = z$ and thus $\langle f(u,6) \, | \, u \in V \rangle = \langle x_1 x_2, x_4 x_3, z \rangle \cong C_2^3$ is abelian, and $e_f(6) = z \neq 1$. Hence f is a tractable flow that leaks at 6 and $K_{3,3}$ is not leak-proof.

Example 1.3.5. Consider the complete graph $K_5 = (V, E)$ with $V = \{1, 2, 3, 4, 5\}$. Now we consider $f \colon V^2 \to \mathrm{ES}_3$ given by

$$
f = \begin{pmatrix}
 & x_1 & x_2 & x_3 & x_1 x_2 x_3 \\
x_1 & & x_6 & x_5 & x_1 x_6 x_5 \\
x_2 & x_6 & & x_4 & x_2 x_6 x_4 \\
x_3 & x_5 & x_4 & & x_3 x_5 x_4 \\
x_1 x_2 x_3 & x_1 x_6 x_5 & x_2 x_6 x_4 & x_3 x_5 x_4 &
\end{pmatrix}.
$$

For each of the first four columns one notes that its first three non-trivial elements commute pairwise, while multiplying them yields the fourth. Thus for the first four vertices v the group $\langle f(u,v) \, | \, u \in V \rangle \cong C_2^3$ is abelian and $e_f(v) = 1$. For the last column, note that each pair (a, b) of entries is of the form $a = x_i x_j x_k$ and $b = x_i x_{j+3} x_{k+3}$ with $i, j, k, j+3, k+3 \in \mathbb{Z}/6\mathbb{Z}$ distinct. Hence

$$
ab = x_i^2 (x_j x_k)(x_{j+3} x_{k+3}) = x_i^2 (x_{j+3} x_{k+3})(x_j x_k) = ba,
$$

so each pair commutes. Finally, one computes

$$
e_f(5) = (x_1 x_2 x_3)(x_1 x_6 x_5)(x_2 x_6 x_4)(x_3 x_5 x_4) = z \neq 1.
$$

Thus f is a tractable leaking flow and thus K_5 is not leak-proof.

Both examples were found by starting with the free group F with symbols V^2 and dividing out the relations $N \trianglelefteq F$ required to make the obvious map $f \colon V^2 \to F/N$ a tractable flow that is conserving in $\#V - 1$ vertices. Adding the restriction that the generators have order 2 gives us the groups ES_2 and ES_3.

It now follows that all non-planar graphs leak, so we are half-way done proving Theorem 1.4.3.

1.4 Planar graphs

Now we will prove that all planar graphs are leak-proof by induction. For this we require a definition of the excess for non-tractable flows.

Definition 1.4.1. Let $G = (V, E)$ be a graph with planar embedding ε and let $f \colon V^2 \to \Gamma$ be a flow of G. Write $C(\Gamma)$ for the set of conjugacy classes of Γ and $\equiv$ for equality up to conjugation. Then we define for (G, ε, f) the *round flow* $r \colon V \to C(\Gamma)$ as $r(v) \equiv 1$ if $N_G(V) = \varnothing$, and otherwise

$$r(v) \equiv f(\rho_\varepsilon(v)^0(u), v) \cdot f(\rho_\varepsilon(v)^1(u), v) \cdots f(\rho_\varepsilon(v)^{n-1}(u), v),$$

where $u \in N_G(v)$, $n = \#N_G(v)$ and ρ_ε is as in Definition 1.2.2.

Note that choosing a different $u \in N_G(v)$ in the above definition results in a cyclic permutation of the factors, hence the products are conjugate in Γ. Thus the round flow is well-defined. Since $1 \in \Gamma$ is only conjugate to itself, we have that $r(v) \equiv 1$ if and only if $e(v) = 1$ when the latter is defined.

Proposition 1.4.2. *Let $G = (V, E)$ be a graph with planar embedding ε and let $f \colon V^2 \to \Gamma$ be a flow of G. Let $u \in V$ and assume $r(v) \equiv 1$ for all $v \in V \setminus \{u\}$. Then $r(u) \equiv 1$.*

Proof. Firstly, if G is the singleton graph, then $r(u) \equiv 1$ is the empty product, so we are done. We now apply induction and thus assume that the statement holds for all strict subgraphs (W, F) of G with planar embedding $\varepsilon|_W$. We may now assume $\#V > 1$.

Secondly, we consider the case where G is not connected. Here we may apply the induction hypothesis to the induced subgraph of G with as vertex set the connected component of u to conclude that $r(u) \equiv 1$. We may now assume G is connected.

Thirdly, we consider the case where G is a forest. Then G has at least two vertices of degree 1, of which one, say v, is not u. Let $e = \{v, w\} \in E$ be the unique edge incident to v, and note that $1 \equiv r_f(v) \equiv f(w, v)$ implies $f(w, v) = 1$. Hence f is a flow of the subgraph H of G obtained by removing e. Note that ε is a planar embedding of H with the same round flow in each vertex, hence by the induction hypothesis we have $r_f(u) \equiv 1$.

Lastly we consider the case where G not a forest. Then G has an edge $\{v, w\} \in E$ that is not a bridge. Then by Lemma 1.2.3 the boundary walk $p = (x_0, \ldots, x_n)$ of (G, ε) with $x_0 = v$ and $x_1 = w$ satisfies $(w, v) \neq (x_i, x_{i+1})$ for all $i \in \mathbb{Z}/n\mathbb{Z}$. Let $b \colon V^2 \to \{0, 1\}$ be the map such that for all $s, t \in V$ we have $b(s, t) = 1$ if and only if there exists some $i \in \mathbb{Z}/n\mathbb{Z}$ such that $(s, t) = (x_i, x_{i+1})$. Now consider $\gamma = f(v, w)$ and $g \colon V^2 \to \Gamma$ given by

$$(s, t) \mapsto \gamma^{b(t,s)} \cdot f(s, t) \cdot \gamma^{-b(s,t)}.$$

Firstly note that g is a flow of G: For all $s, t \in V$ we have

$$g(s, t)^{-1} = \gamma^{b(s,t)} \cdot f(s, t)^{-1} \cdot \gamma^{-b(t,s)} = g(t, s)$$

since f is a flow, and if $\{s, t\} \notin E$ we have $g(s, t) = f(s, t) = 1$ as $b(s, t) = b(t, s) = 0$. Secondly, we have that $g(v, w) = \gamma^0 \cdot \gamma \cdot \gamma^{-1} = 1$ by choice of $\{v, w\}$, so g is even a flow of the subgraph H of G obtained by removing $\{v, w\}$. We now show that the round flows r_f and r_g of f respectively g in (G, ε) are conjugates in Γ at each vertex. Then by the induction hypothesis applied to H it follows that $r(u) \equiv 1$. Note that for all $s, t \in V$ we have by definition of b that $b(t, s) = b(s, \rho_\varepsilon(s)(t))$. Using this, we now simply verify for $\{s, t\} \in E$, $n = \#N_G(s)$ and $\rho = \rho_\varepsilon(s)$ that

$$
r_g(s) \equiv \prod_{k=0}^{n-1} g(\rho^k(t), s) \equiv \prod_{k=0}^{n-1} \gamma^{b(s, \rho^k(t))} \cdot f(\rho^k(t), s) \cdot \gamma^{-b(\rho^k(t), s)}
$$

$$
\equiv \gamma^{b(s, t)} \left(\prod_{k=0}^{n-1} f(\rho^k(t), s) \gamma^{-b(\rho^k(t), s)} \gamma^{b(s, \rho^{k+1}(t))} \right) \gamma^{-b(s, \rho^n(t))}
$$

$$
\equiv \gamma^{b(s, t)} \left(\prod_{k=0}^{n-1} f(\rho^k(t), s) \right) \gamma^{-b(s, t)} \equiv \prod_{k=0}^{n-1} f(\rho^k(t), s) \equiv r_f(s),
$$

as was to be shown. We conclude that the statement holds for all planar graphs by induction. □

An earlier proof of Proposition 1.4.2 was due to H.W. Lenstra. In his version he does not remove edges in the inductive step but contracts them in the sense of Definition 1.2.4. This proof turned out to be more difficult to formalize.

Theorem 1.4.3. *A graph is leak-proof if and only if it is planar.*

Proof. A non-planar graph has either K_5 or $K_{3,3}$ as minor by Proposition 1.2.5. Both K_5 and $K_{3,3}$ are not leak-proof by Example 1.3.5 respectively Example 1.3.4, so by Proposition 1.3.2 neither are the non-planar graphs. Let $G = (V, E)$ be a planar graph with $u \in V$ and let f be a tractable flow of G such that $e(v) = 1$ for all $v \in V \setminus \{u\}$. After choosing a planar embedding for G we have $r(u) \equiv 1$ by Proposition 1.4.2 and thus $e(u) = 1$. Hence f does not leak and G is leak-proof. □

1.5 Extra-planar graphs

In this section we will prove Theorem 1.5.2, classifying the binary leak-proof graphs. To do this we first prove a 'Kuratowski's Theorem' for extra-planar graphs. Write K_5^- and $K_{3,3}^-$ for the graphs obtained from K_5 respectively $K_{3,3}$ by removing a single edge, which by symmetry we do not have to specify.

Proposition 1.5.1. *A graph G is extra-planar if and only if G does not have K_5^- or $K_{3,3}^-$ as a minor.*

Proof. ($\Rightarrow$) This follows directly from Kuratowski's Theorem: If K_5^- or $K_{3,3}^-$ is a minor of G, then we may add a single edge to G such that K_5 respectively $K_{3,3}$ becomes a minor of this new graph, which is then non-planar.

($\Leftarrow$) We proceed by contraposition, so assume that G is not extra-planar. Let $u, v \in V$ be such that $G^+ = (V, E \cup \{\{u, v\}\})$ is non-planar and let $H^+ = (V, F)$ be a spanning forest of G^+ such that K_5 or $K_{3,3}$ embeds into G_H^+. Consider the spanning forest $H = (V, F \setminus \{\{u, v\}\})$ of G. Then H has the same connected components as H^+ with the exception that if H^+ has a connected component containing both u and v, it might have been split into two. Let T_u and T_v be the connected components of u respectively v in H.


Figure 1.2: Case $K_{3,3}$

Figure 1.3: Case K_5

Case $K_{3,3}$: First consider the case where $K_{3,3}$ embeds into $G_{H^+}^+$, meaning there is a subset $C = \{L_1, L_2, L_3, R_1, R_2, R_3\}$ of size 6 of the set of connected components of H^+ such that $S^+ = (C, \{\{L_i, R_j\} \mid i, j \in \{1, 2, 3\}\})$ is a subgraph of $G_{H^+}^+$. If all elements of C are also connected components of H, then G_H has the graph S^+ minus possibly a single edge induced by $\{u, v\}$ as subgraph, hence G has $K_{3,3}^-$ as a minor. Otherwise, for some $X \in C$ we have $X = T_u \sqcup T_v$ and without loss of generality $X = L_1$. Then the subgraph S of G_H induced by $\{T_u, T_v, L_2, L_3, R_1, R_2, R_3\}$ is as in Figure 2, where the dashed lines indicate edges which are possibly present. Merging T_u and T_v in S yields $S^+ \cong K_{3,3}$, hence for each $i \in \{1, 2, 3\}$ the edge $\{T_u, R_i\}$ or $\{T_v, R_i\}$ is present. Thus T_u or T_v has degree at least 2, which without loss of generality is T_v. It follows that $K_{3,3}^-$ embeds into the subgraph of G_H induced by $\{T_v, L_2, L_3, R_1, R_2, R_3\}$, so $K_{3,3}^-$ is a minor of G.

Case K_5: Now consider the case K_5 embeds into $G_{H^+}^+$, meaning there is a subset $C = \{P_1, \ldots, P_5\}$ of the set of connected components of H^+ such that the subgraph of $G_{H^+}^+$ induced by C is isomorphic to K_5. As before, the only interesting case is where $P_5 = T_u \sqcup T_v$. Then the subgraph S of G_H induced by $\{T_u, T_v, P_1, P_2, P_3, P_4\}$ is as in Figure 3. Since merging T_u and T_v in S yields K_5, for each $i \in \{1, \ldots, 4\}$ the edge $\{T_u, P_i\}$ or $\{T_v, P_i\}$ is present. If both T_u and T_v have degree 2, then without loss of generality S contains the edges $\{T_u, P_3\}$, $\{T_u, P_4\}$, $\{T_v, P_1\}$ and $\{T_v, P_2\}$. Now note that S contains a $K_{3,3}^-$ which partitions its vertices as $\{\{T_u, P_1, P_2\}, \{T_v, P_3, P_4\}\}$. Hence G contains $K_{3,3}^-$ as a minor. Otherwise, without loss of generality T_u has degree at least 3 in S and the subgraph of G_H induced by $\{T_u, P_1, \ldots, P_4\}$ is either K_5 or K_5^-. Hence G has K_5^- as a minor.

As G has $K_{3,3}^-$ or K_5^- as a minor, the claim follows.   □

We are now able to prove Theorem 1.5.2.

Theorem 1.5.2. *A graph is binary leak-proof if and only if it is extra-planar.*

Proof. ($\Leftarrow$) Let $G = (V, E)$ be an extra-planar graph and let $f : V^2 \to \Gamma$ be a tractable flow of G such that there are distinct $u, v \in V$ with $e_f(w) = 1$ for all $w \in V \setminus \{u, v\}$. Consider the graph $H = (V, E \cup \{\{u, v\}\})$ and let ε be a planar embedding of H. Now let $g : V^2 \to \Gamma$ be the map such that $g(s, t) = f(s, t)$ if $\{s, t\} \neq \{u, v\}$ and $g(u, v) = g(v, u)^{-1} = f(u, v) r_f(v)^{-1}$, where $r_f(v)$ is computed by starting from the vertex right after u in the ordering of $N_H(v)$. Then g is a (not necessarily tractable) flow in H such that $r_g(w) = 1$ for $w \in V \setminus \{u\}$. From $g(v, u) = r_f(v) f(v, u)$ it follows that $r_g(u)$ differs from $r_f(u)$ by a factor $r_f(v)$ when starting the multiplication at v. By Proposition 1.4.2 we have $1 \equiv r_g(u) \equiv r_f(u) r_f(v)$ and thus $e_f(u) e_f(v) = 1$. Hence G is binary leak-proof.

($\Rightarrow$) If $G = (V, E)$ is not extra-planar, then it has K_5^- or $K_{3,3}^-$ as minor by Proposition 1.5.1. It is straightforward to generalize Proposition 1.3.2 to show that a graph is binary leak-proof if and only if all its minors are too. It therefore suffices to show that K_5^- and $K_{3,3}^-$ have a binary leaking flow. Simply take the flow f as defined in Example 1.3.4 which leaks at vertex 6 of $K_{3,3}$ and consider $K_{3,3}^-$ as the $K_{3,3}$ with the edge $\{3, 6\}$ removed. Then the flow f^- of $K_{3,3}^-$ which equals f except for $f^-(3, 6) = f^-(6, 3) = 1$ has a binary leak at 3 and 6. Using Example 1.3.5 for K_5^- can be done analogously.   □

1.6 Leak-proof groups

In this section we prove Theorem 1.6.4 and give some computational results.

Definition 1.6.1. Let Γ be a (not necessarily finite) group. Write $V(\Gamma)$ for the set of maximal abelian subgroups of Γ. We define the group

$$F(\Gamma) = \left\{ (f_{u,v})_{(u,v)} \in \bigoplus_{(u,v)\in V(\Gamma)^2} (u \cap v) \,\middle|\, (\forall\, u,v)\; f_{u,v} = f_{v,u}^{-1},\; (\forall v)\; f_{v,v} = 1 \right\}$$

and the homomorphism

$$e_\Gamma \colon F(\Gamma) \to \bigoplus_{v\in V(\Gamma)} v, \quad (f_{u,v})_{(u,v)} \mapsto \left(\prod_{u\in V(\Gamma)} f_{u,v} \right)_{v\in V(\Gamma)}.$$

One can think of $V(\Gamma)$ as the vertex set of a complete graph, $F(\Gamma)$ the set of tractable flows in this graph, and $e_\Gamma(f)$ to be the excess for such flow $f \in F(\Gamma)$. However, $V(\Gamma)$ need not be finite. For example $\Gamma = \mathrm{GL}_2(\mathbb{R})$ has a maximal abelian subgroup $\{ \left(\begin{smallmatrix} a & b \\ 0 & a \end{smallmatrix} \right) \mid a, b \in \mathbb{R},\, a \neq 0 \}$ with infinitely many conjugates.

Lemma 1.6.2. *Let Γ be a group. For $u \in V(\Gamma)$ and $\gamma \in u$ let $[\gamma]_u \in \bigoplus_{v\in V(\Gamma)} v$ be the vector consisting of all-ones except for a γ at coordinate u. We write $\Gamma^\bullet = (\bigoplus_{v\in V(\Gamma)} v)/\mathrm{im}(e)$. Then the map $d \colon \Gamma \to \Gamma^\bullet$ given by $\gamma \mapsto [\gamma]_v$ for any choice of v containing γ, does not depend on the choice of v.*

Proof. Let $\gamma \in \Gamma$ and suppose $u, v \in V(\Gamma)$ are such that $\gamma \in u$ and $\gamma \in v$. Then $\gamma \in u \cap v$, and $f = (f_{s,t})_{(s,t)\in V(\Gamma)^2}$, with $f_{u,v} = f_{v,u}^{-1} = \gamma$ and $f_{s,t} = 1$ for $\{s,t\} \neq \{u,v\}$, is an element of $F(\Gamma)$. We have $e(f) = [\gamma]_v \cdot [\gamma]_u^{-1}$, so $[\gamma]_u$ is equivalent to $[\gamma]_v$ in the quotient $\Gamma^\bullet$. $\qquad\square$

An example one can consider is where Γ is abelian. Then $V(\Gamma) = \{\Gamma\}$ and $\Gamma^\bullet = \Gamma$ and d is the identity. Note that d is (in general) not a group homomorphism.

Proposition 1.6.3. *A group Γ is leak-proof if and only if $d(\gamma) = 1$ implies $\gamma = 1$.*

Proof. Suppose Γ is leak-proof and $d(\gamma) = 1$ for some $\gamma \in \Gamma$. Then there is some $u \in V(\Gamma)$ and $f \in F(\Gamma)$ such that $[\gamma]_u = e(f)$. Note that $E = \{\{u,v\} \in V(\Gamma) \mid f_{u,v} \neq 1\}$ and $V = \{u \mid \{u,v\} \in E\}$ are finite. Now f is a Γ-flow in (V, E) which is preserving in all vertices except possibly u. Since Γ is leak-proof, f is also preserving in u and $1 = e(f) = [\gamma]_u$, so $\gamma = 1$.

Conversely, suppose f is a tractable Γ-flow in some graph (V, E). Pick some map $c \colon V \to V(\Gamma)$ such that for all $v \in V$ we have $\langle f(u,v) \mid u \in V \rangle \subseteq c(v)$. Then f induces a tractable Γ-flow f' in the complete graph with vertex set $\{c(v) \mid v \in V\}$ where

$$f'(s,t) = \prod_{u \colon c(u) = s} \prod_{v \colon c(v) = t} f(u,v) \in s \cap t.$$

Hence $f' \in F(\Gamma)$. Moreover, if f leaks, then so does f'. Assume f' is preserving in all vertices except potentially $v \in V$. Then $e(f') = [\gamma]_v$ for some $\gamma \in v$ and $d(\gamma) = 1$. If $d(\gamma) = 1$ implies $\gamma = 1$, we obtain that f' and hence f does not leak, so Γ is leak-proof. $\qquad\square$

Similarly, one can consider binary leak-proof groups. With a proof analogous to that of Proposition 1.6.3 one obtains that Γ is binary leak proof if and only if d is injective.

Theorem 1.6.4. *The decision problem 'Is this finite group leak-proof?' is decidable.*

Proof. Simply note that for finite Γ the corresponding group $\Gamma^\bullet$ is finite abelian and can thus be computed explicitly. In particular, we can decide for each $\gamma \in \Gamma$ whether $d(\gamma) = 1$. The theorem thus follows from Proposition 1.6.3. $\qquad\square$

From Lemma 1.1.1 it follows that abelian groups are leak-proof, but they are hardly the only ones. By computer search we found the two extraspecial groups of order 32 to be the only smallest leaking groups, one of which we encountered in Example 1.3.4. The smallest leaking groups of order greater than 32 occur at order 64. That there are groups of order 64 that leak was to be expected, because a group leaks when it has a leaking subgroup. The smallest leaking symmetric group is the S_6 and the smallest leaking alternating group is the A_7. That for sufficiently large n the group S_n leaks is to be expected by Cayley's theorem, but interestingly no strict subgroup of S_6 leaks. It would be interesting to have a classification of leak-proof groups or to know whether there is some equivalent, better understood property of groups which is equivalent to being leak-proof like planarity is for graphs.

Our code is available online [21] and is written in GAP [15].