Contents

Chapter 1

Two good counting algorithms

Counting problems that can be solved exactly in polynomial time are few and far between. Here are two classical examples whose solution makes elegant use of linear algebra. Both algorithms predate the now commonplace distinction between polynomial and exponential time, which is often credited (with justification) to Edmonds in the mid 1960s; indeed our first example dates back over 150 years!

1.1 Spanning trees

Basic graph-theoretic terminology will be assumed. Let $G = (V, E)$ be a finite undirected graph with vertex set V and edge set E. For convenience we identify the vertex set V with the first n natural numbers $[n] = \{0, 1, \dots, n-1\}$. The *adjacency matrix A of G* is the $n \times n$ symmetric matrix whose ij'th entry is 1 if $\{i, j\} \in E$, and 0 otherwise. Assume G is connected. A *spanning tree* in G is a maximum (edge) cardinality cycle-free subgraph (equivalently, a minimum cardinality connected subgraph that includes all vertices). Any spanning tree has $n-1$ edges.

Theorem 1.1 (Kirchhoff)**.** *Let $G = (V, E)$ be a connected, loop-free, undirected graph on n vertices, A its adjacency matrix and $D = \mathrm{diag}(d_0, \dots, d_{n-1})$ the diagonal matrix with the degrees of the vertices of G in its main diagonal. Then, for any i, $0 \le i \le n-1$,*

$$\#\text{ spanning trees of } G = \det(D - A)_{ii},$$

where $(D - A)_{ii}$ is the $(n-1) \times (n-1)$ principal submatrix of $D - A$ resulting from deleting the i'th row and i'th column.

Since the determinant of a matrix may be be computed in time $O(n^3)$ by Gaussian elimination, Theorem 1.1 immediately implies a polynomial-time algorithm for counting spanning trees in an undirected graph.

Example 1.2. Figure 1.1 shows a graph G with its associated "Laplacian" $D - A$ and principal minor $(D - A)_{11}$. Note that $\det(D - A)_{11} = 3$ in agreement with Theorem 1.1.

Remark 1.3. The theorem holds for unconnected graphs G, as well, because then the matrix $D - A$ associated with G is singular. To see this, observe that the rows and columns of a connected graph add up to 0 and, similarly, those of any submatrix

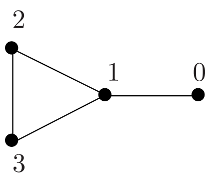

Graph G

$$\begin{pmatrix} 1 & -1 & 0 & 0 \\ -1 & 3 & -1 & -1 \\ 0 & -1 & 2 & -1 \\ 0 & -1 & -1 & 2 \end{pmatrix}$$

$D - A$

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & -1 \\ 0 & -1 & 2 \end{pmatrix}$$

$(D - A)_{11}$

Figure 1.1: Example illustrating Theorem 1.1.

corresponding to a connected component add up to 0. Now choose vertex i and a connected component C such that $i \notin C$. Then, the columns of $(D-A)_{ii}$ that correspond to C are linearly dependent, and $(D - A)_{ii}$ is singular.

Our proof of Theorem 1.1 follows closely the treatment of van Lint and Wilson [79], and relies on the following expansion for the determinant, the proof of which is deferred.

Lemma 1.4 (Binet-Cauchy). *Let A be an $(r \times m)$- and B an $(m \times r)$-matrix. Then*

$$\det AB = \sum_{\substack{S \subseteq [m], \\ |S| = r}} \det A_{*S} \det B_{S*},$$

where A_{*S} is the square submatrix of A resulting from deleting all columns of A whose index is not in S, while, similarly, B_{S*} is the square submatrix of B resulting from B by deleting those rows not in S.

Remark 1.5. Typically, r is smaller than m. However, the lemma is also true for $r > m$. Then the sum on the right is empty and thus 0. But also AB is singular, since $\operatorname{rank} AB \leq \operatorname{rank} A \leq m < r$.

Let H be a directed graph on n vertices with m edges. Then the *incidence matrix of H* is the $(n \times m)$-matrix $N = (\nu_{ve})$ where

$$\nu_{ve} = \begin{cases} +1, & \text{if vertex } v \text{ is the head of edge } e; \\ -1, & \text{if } v \text{ is the tail of } e; \\ 0, & \text{otherwise.} \end{cases}$$

The *weakly connected components of H* are the connected components of the underlying undirected graph, i.e., the graph obtained from H by ignoring the orientations of edges.

Fact 1.6.

$$\operatorname{rank} N = |V(H)| - |\mathcal{C}(H)| = n - |\mathcal{C}(H)|,$$

where $V(H)$ is the vertex set of H and $\mathcal{C}(H) \subseteq 2^{V(H)}$ is the set of (weakly) connected components of H.

Proof. Consider the linear map represented by $N^\top$, the transpose of N. It is easy to see that, if h is a vector of length n, then

$$N^\top h = 0 \Leftrightarrow h \text{ is constant on connected components},$$

i.e., $i, j \in C \Rightarrow h_i = h_j$, for all $C \in \mathcal{C}(H)$. This implies that $\dim \ker N^\top = |\mathcal{C}(H)|$, proving the claim, since $\operatorname{rank} N = \operatorname{rank} N^\top = n - \dim \ker N^\top$. □

Fact 1.7. *Let B be a square matrix with entries in $\{-1, 0, +1\}$ such that in each column there is at most one $+1$ and at most one -1. Then, $\det B \in \{-1, 0, +1\}$.*

Proof. We use induction on the size n of B. For $n = 1$, the claim in trivial. Let $n > 1$. If B has a column which equals 0, or if each column has exactly one $+1$ and one -1, then B is singular. Otherwise there is a column j with either one $+1$ or one -1, say in its i'th entry b_{ij}, and the rest 0's. Developing $\det B$ by this entry yields $\det B = \pm b_{ij} \det B_{ij}$, where B_{ij} is the minor of B obtained by deleting row i and column j. By the induction hypothesis, the latter expression equals -1, 0 or $+1$. □

The ingredients for the proof of the Kirchhoff's result are now in place.

Proof of Theorem 1.1. Let $\vec{G}$ be an arbitrary orientation of G, N its incidence matrix, and $S \subseteq E$ be a set of edges of $\vec{G}$ with $|S| = n - 1$. Then, by Fact 1.6,

$$\text{rank}(N_{*S}) = n - 1 \;\Leftrightarrow\; S \text{ is the edge set of a tree.} \tag{1.1}$$

(The condition that S is the edge set of a tree again ignores the orientation of edges in S.) If N' results from N by deleting one row, then

$$\text{rank}(N'_{*S}) = \text{rank}(N_{*S}). \tag{1.2}$$

This is because the deleted row is a linear combination of the others, since the rows of N add up to 0. Combining (1.1) and (1.2) with Fact 1.7 gives us

$$\det N'_{*S} = \begin{cases} \pm 1, & \text{if } S \text{ is a spanning tree;} \\ 0, & \text{otherwise.} \end{cases} \tag{1.3}$$

Now observe that $D - A = NN^\top$, since

$$(NN^\top)_{ij} = \sum_{e \in E} \nu_{ie}\nu_{je} = \begin{cases} -1, & \text{if } \{i, j\} \in E; \\ d_i, & \text{if } i = j; \\ 0, & \text{otherwise.} \end{cases}$$

Clearly, $(D - A)_{ii} = N'(N')^\top$ where N' results from N by deleting any row i. Thus,

$$\begin{aligned} \det(D - A)_{ii} &= \det(N'(N')^\top) \\ &= \sum_{|S|=n-1} \det N'_{*S} \; \det((N')^\top)_{S*} && \text{by Lemma 1.4} \\ &= \sum_{|S|=n-1} \det N'_{*S} \; \det(N'_{*S})^\top \\ &= \#\text{ spanning trees of } G && \text{by (1.3).} \end{aligned}$$

□

It only remains to prove the key lemma on expanding determinants.

Proof of Lemma 1.4. We prove a more general claim, namely

$$\det A\Delta B = \sum_{\substack{S\subseteq[m],\\|S|=r}} \det A_{*S} \det B_{S*} \prod_{i\in S} e_i,$$

where $\Delta = \mathrm{diag}(e_0, \ldots, e_{m-1})$. The lemma follows by setting all e_i to 1. Observe that entries of $A\Delta B$ are linear forms in $e_0, \ldots, e_{m-1}$. Thus, $\det A\Delta B$ is a homogeneous polynomial of degree r in $e_0, \ldots, e_{m-1}$, i.e., all monomials have degree r. Comparing coefficients will yield the desired result. First we observe that every monomial in $\det A\Delta B$ must have r distinct variables. For if not, consider a monomial with the fewest number of distinct variables, and suppose this number is less than r. Setting all other variables to 0 will result in $\det A\Delta B = 0$, since $\mathrm{rank}\, A\Delta B \le \mathrm{rank}\, \Delta < r$ and $A\Delta B$ is singular. But $\det A\Delta B = 0$ implies that the coefficient of the monomial is 0. Now look at a monomial with exactly r distinct variables, say $\prod_{i\in S} e_i$. Set these variables to 1 and all others to 0. Then, $A\Delta B$ evaluates to $A_{*S}B_{S*}$, and hence the coefficient of $\prod_{i\in S} e_i$ is $\det A_{*S}B_{S*} = \det A_{*S} \det B_{S*}$. □

It is possible to generalise Theorem 1.1 to directed graphs $G = (V, E)$, where a directed spanning tree (or *arborescence*) is understood to be a subgraph $(V, T\subseteq E)$ where (i) (V, T) with the orientation of edges ignored forms a spanning tree of the unoriented version of G, and (ii) the orientations of edges in T are consistently directed towards some distinguished vertex or *root* r. Equivalently, it is an acyclic subgraph in which every vertex other than the distinguished root r has outdegree 1, and the root itself has outdegree 0. (There does not seem to be agreement on whether edges should be directed towards or away from the root; towards seems more natural — corresponding as it does to functions on $[n]$ with a unique fixed point — and in any case better suits our immediate purpose.)

An *Eulerian circuit* in a directed graph G is a closed path (i.e., one that returns to its starting point) that traverses every edge of G exactly once, respecting the orientation of edges. (The path with not in general be simple, that is to say it will visit vertices more than once.) The number of Eulerian circuits in a directed graph is related in a simple way to the number of arborescences, so these structures also can be counted in polynomial time. For details see Tutte [74, §VI.3, §VI.4].

Open Problem. To the best of my knowledge, it is not known whether there exists a polynomial-time algorithm for counting Eulerian circuits in an undirected graph. Note that the usual strategy of viewing an undirected graph as a directed graph with paired anti-parallel edges does not work here.

Exercise 1.8. Exhibit an explicit (constant) many-one relation between the Eulerian circuits in a directed graph G and the arborescences in G. Hint: use the arborescence to define an “escape route” or “edge of final exit” from each vertex.

1.2 Perfect matchings in a planar graph

Let $G = (V, E)$ be an undirected graph on n vertices ($V = [n]$, for convenience). A *matching* in G is a subset $M \subseteq E$ of pairwise vertex-disjoint edges. A matching M is called *perfect* if it covers V, i.e., $\bigcup M = V$. Note that n must be even for a perfect matching to exist.

Around 1960, Kasteleyn discovered a beautiful method for counting perfect matchings in a certain class of "Pfaffian orientable" graphs, which includes all planar graphs as a strict subclass. Linear algebra is again the key.

Fact 1.9. *If M, M' are two perfect matchings in G, then $M \cup M'$ is a collection of single edges and even (i.e., even length) cycles.*

Let $G = (V, E)$ be an undirected graph, C an even cycle in G, and $\vec{G}$ an orientation of G. We say that C is *oddly oriented by* $\vec{G}$ if, when traversing C in either direction, the number of co-oriented edges (i.e., edges whose orientation in $\vec{G}$ and in the traversal is the same) is odd. (Observe that the direction in which we choose to traverse C is not significant, since the parity in the other direction is the same.) An orientation $\vec{G}$ of G is *Pfaffian* (also called *admissible*) if the the following condition holds: for any two perfect matchings M, M' in G, every cycle in $M \cup M'$ is oddly oriented by $\vec{G}$. Note that all cycles in $M \cup M'$ are even.

Remark 1.10. The definition of Pfaffian orientation given above is not equivalent to requiring that all even cycles in G be oddly oriented by $\vec{G}$, since there may be even cycles that cannot be obtained as the union of two perfect matchings.

Let $\vec{G}$ be any orientation of G. Define the *skew adjacency matrix* $A_s(\vec{G}) = (a_{ij} : 0 \le i, j \le n-1)$ *of* G by

$$a_{ij} = \begin{cases} +1, & \text{if } (i,j) \in E(\vec{G}); \\ -1, & \text{if } (j,i) \in E(\vec{G}); \\ 0, & \text{otherwise.} \end{cases}$$

Theorem 1.11 (Kasteleyn)**.** *For any Pfaffian orientation $\vec{G}$ of G,*

$$\#\text{ perfect matchings in } G \;=\; \sqrt{\det A_s(\vec{G})}\,.$$

Our proof of Theorem 1.11 borrows from Kasteleyn [52] and Lovász and Plummer [56]. Denote by $\overleftrightarrow{G}$ the directed graph obtained from G by replacing each undirected edge $\{i, j\}$ by the anti-parallel pair of directed edges $(i, j), (j, i)$. An *even cycle cover* of $\overleftrightarrow{G}$ is a collection $\mathcal{C}$ of even directed cycles $C \subseteq E(\overleftrightarrow{G})$ such that every vertex of G is contained in exactly one cycle in $\mathcal{C}$.

Lemma 1.12. *There is a bijection between (ordered) pairs of perfect matchings in G and even cycle covers in $\overleftrightarrow{G}$.*

Proof. Let (M, M') be a pair of perfect matchings in G. For each edge in $M \cap M'$ (i.e, each edge in $M \cup M'$ that does not lie in an even cycle) take both directed edges in $\overleftrightarrow{G}$. Now orient each cycle C in $M \cup M'$ (with length ≥ 4) according to some convention fixed in advance. For example, take the vertex with lowest number in C and orient the incident M-edge away from it. The resulting collection $\mathcal{C}$ of directed cycles is an even cycle cover of $\overleftrightarrow{G}$.

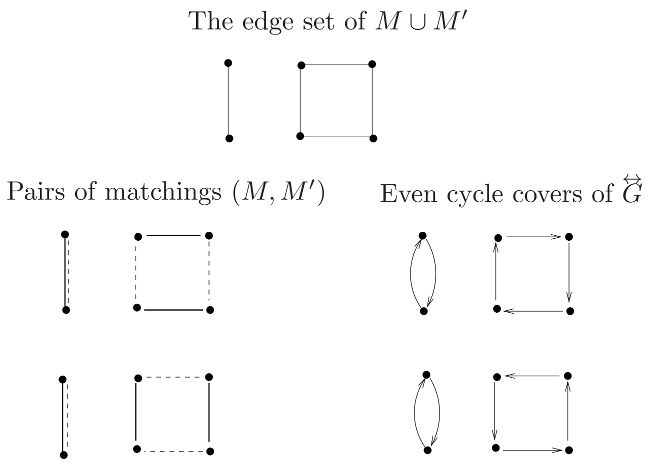

Figure 1.2: Bijection between pairs of matchings in G and even cycle covers of $\overleftrightarrow{G}$.

The procedure may be reversed. First, each oriented 2-cycle in $\mathcal{C}$ must correspond to an edge that is in both M and M'. Then, each even cycle $C \in \mathcal{C}$ of length at least four may be decomposed into alternating M-edges and M'-edges; the convention used to determine the orientation of C will indicate which of the two possible decompositions is the correct one. □

Proof of Theorem 1.11. In view of the previous lemma, we just need to show that $\det A_s(\overrightarrow{G})$ counts even cycle covers in $\overleftrightarrow{G}$. Now,

$$\det A_s(\overrightarrow{G}) := \sum_{\pi \in S_n} \operatorname{sgn} \pi \prod_{i=0}^{n-1} a_{i,\pi(i)}, \tag{1.4}$$

where S_n is the set of all permutations of $[n]$, and $\operatorname{sgn} \pi$ is the sign of permutation π.[1] Consider a permutation π and its (unique) decomposition into disjoint cycles $\pi = \gamma_1 \cdots \gamma_k$. Each γ_j acts on a certain subset $V_j \subseteq V$. The corresponding product $\prod_{i \in V_j} a_{i,\pi(i)}$ is non-zero if and only if the edges $\{(i, \pi(i)) : i \in V_j\}$ form a directed cycle in G, since otherwise one of the $a_{i,\pi(i)}$ would be 0. Thus, there is a one-to-one correspondence between permutations π with non-zero (i.e., ± 1) contributions to (1.4) and cycle covers in $\overleftrightarrow{G}$.

We now claim that sum (1.4) is unchanged if we restrict it to permutations with only even length cycles. To see this, consider a permutation π and an odd length cycle γ_j in π, say the first in some natural ordering on cycles. Let $\pi' = \gamma_1 \cdots (\gamma_j)^{-1} \cdots \gamma_k$ be identical to π except that γ_j is reversed. Then, $\prod_{i=0}^{n-1} a_{i,\pi(i)} = -\prod_{i=0}^{n-1} a_{i,\pi'(i)}$. Moreover, since both π and π' are products of cycles of the same lengths, $\operatorname{sgn} \pi = \operatorname{sgn} \pi'$. Thus, the contributions of π and π' cancel out in (1.4). (Note that for this part of the argument, we do not need that $\overrightarrow{G}$ is Pfaffian.) Thus we may pair up permutations with odd cycles so that they cancel each other.

[1]The sign of π is $+1$ if the cycle decomposition of π has an even number of even length cycles, and -1 otherwise.

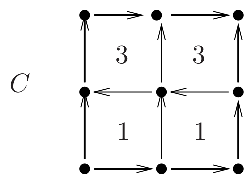

Figure 1.3: Example graph illustrating various quantities in the proof.

Now consider a permutation π which consists only of even length cycles and does not vanish in (1.4). As remarked above, π corresponds to an even cycle cover of $\overleftrightarrow{G}$, which, by Lemma 1.12, corresponds to a pair of perfect matchings in G. Because $\vec{G}$ is Pfaffian, each cycle C_j corresponding to a cycle γ_j of π is oddly oriented by $\vec{G}$. Thus, each γ_j contributes a factor -1 to $\prod_{i=0}^{n-1} a_{i,\pi(i)}$ while it also contributes a factor -1 to $\operatorname{sgn}\pi$, being an even cycle. Therefore, overall, π contributes 1 to the sum (1.4). □

Theorem 1.11 provides a polynomial-time algorithm for counting perfect matchings in a graph G, provided G comes equipped with a Pfaffian orientation. But which graphs admit a Pfaffian orientation?

Lemma 1.13. *Let $\vec{G}$ be a connected planar digraph, embedded in the plane. Suppose every face, except the (outer) infinite face, has an odd number of edges that are oriented clockwise. Then, in any simple cycle C, the number of edges oriented clockwise is of opposite parity to the number of vertices of $\vec{G}$ inside C. In particular, $\vec{G}$ is Pfaffian.*

Proof. First, let's see why the condition on simple cycles implies $\vec{G}$ is Pfaffian. Consider a cycle C created by the union of a pair of perfect matchings in G. Then C has an even number of vertices inside it, since otherwise there would be a vertex inside C which is matched with a vertex outside C, contradicting planarity. Thus, the number of edges in C oriented clockwise is odd, implying that $\vec{G}$ is Pfaffian.

We now prove the main part of the lemma. Take a cycle C. We need the following definitions:

$$\begin{aligned} v &= \#\text{ vertices inside } C, \\ k &= \#\text{ edges on } C = \#\text{ vertices on } C, \\ c &= \#\text{ edges on } C \text{ oriented clockwise}, \\ f &= \#\text{ faces inside } C, \\ e &= \#\text{ edges inside } C, \\ c_i &= \#\text{ clockwise edges on the boundary of face } i \text{ for } i = 0, \dots, f-1. \end{aligned}$$

In the example graph illustrated in Figure 1.3, the cycle C is denoted in bold face. Here, $v = 1$, $k = 8$, $c = f = e = 4$, and the various c_i are included in the figure.

According to Euler's formula,

$$\underbrace{(v+k)}_{\#\text{ vertices}} + \underbrace{(f+1)}_{\#\text{ faces}} - \underbrace{(e+k)}_{\#\text{ edges}} = 2,$$

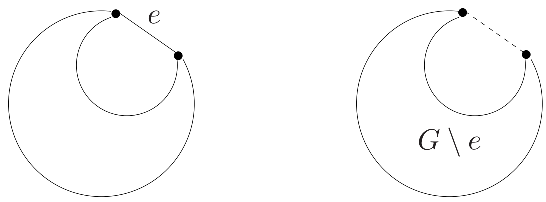

Figure 1.4: Orient e according to the condition of Lemma 1.13.

which implies

$$e = v + f - 1. \tag{1.5}$$

Now, for all i, by assumption, $c_i \equiv 1 \pmod 2$, and thus $f \equiv \sum_{i=0}^{f-1} c_i \pmod 2$. On the other hand, $\sum_{i=0}^{f-1} c_i = c + e$, since each interior edge borders two faces, and in exactly one of these it is oriented clockwise. So,

$$\begin{aligned} f &\equiv c + e \\ &\equiv c + v + f - 1 \pmod 2 \qquad \text{by (1.5)}, \end{aligned}$$

and hence $c + v$ is odd. □

Theorem 1.14. *Every planar graph has a Pfaffian orientation.*

Proof. Without loss of generality, we may assume G is connected, since we may otherwise treat each connected component separately. We prove the theorem by induction on m, the number of edges. As the base of our induction we take the case when G is a tree, and any orientation is Pfaffian. Now, look at a planar graph G with $m \geq n$ edges, and fix an edge e on the exterior (i.e., e borders the infinite face of G). By the induction hypothesis, $G \setminus e$ has a Pfaffian orientation. Adding e creates just one more face; orient e in such a way that this face has an odd number of edges oriented clockwise. (Figure 1.4 illustrates the situation.) Then, by Lemma 1.13, the orientation is Pfaffian. □

Open Problem. The computational complexity of deciding, for an arbitrary input graph G, whether G has a Pfaffian orientation is open. It is neither known to be in P nor to be NP-complete. The restriction of this decision problem to *bipartite* graphs was recently shown to be decidable by Robertson, Seymour and Thomas [68], and independently by McCuaig.

Note however, that the proof of Theorem 1.14 gives us a polynomial algorithm for finding a Pfaffian orientation of a planar graph G, and hence for counting the number of perfect matchings in G.

Exercise 1.15. In the physics community, perfect matchings are sometimes known as "dimer covers." It is of some interest to know the number of dimer covers of a graph G when G has a regular structure that models, for example, a crystal lattice. Let Λ to be the $L \times L$ square lattice, with vertex set $V(\Lambda) = \{(i,j) : 0 \leq i,j < L\}$ and edge set $E(\Lambda) = \big\{\{(i,j),(i',j')\} : |i-i'| + |j-j'| = 1\big\}$. Exhibit a (nicely structured!) Pfaffian orientation of Λ.

Exercise 1.16. Exhibit a non-planar graph that admits a Pfaffian orientation.

Exercise 1.17. Exhibit a (necessarily non-planar) graph that does not admit a Pfaffian orientation.

Exercise 1.18. The dimer model is one model from statistical physics; another is the Ising model. Computing the "partition function" of an Ising system with underlying graph G in the absence of an external field is essentially equivalent to counting "closed subgraphs" of G: subgraphs $(V, A \subseteq E)$ such that the degree of every vertex $i \in V$ in (V, A) is even (possibly zero). Show that the problem of counting closed subgraphs in a planar graph is efficiently reducible to counting perfect matchings (or dimer covers) in a derived planar graph. The bottom line is that the Ising model for planar systems with no applied field is computationally feasible.

Valiant observes that in the few instances where a counting problem is known to be tractable, it is generally on account of the problem being reducible to the determinant. All the examples presented in this chapter are of this form. This empirical observation remains largely a mystery, though a couple of results in computational complexity give special status to the determinant. For example, around 1991, various authors (Damm, Toda, Valiant, and Vinay) independently discovered that the determinant of an integer matrix is complete for the complexity class GapL under log-space reduction [60, §6].[2] Although this is certainly an interesting result, it does beg the question: why do natural tractable counting problems tend to cluster together in the class GapL? For a further universality property of the determinant, see Valiant [75, §2].

In the other direction, Colbourn, Provan and Vertigan [18] have discovered an interesting, purely combinatorial approach to at least some of the tractable counting problems on planar graphs. In a sense, their result questions the centrality of the determinant.

[2] A function $f : \Sigma^* \to \mathbb{N}$ is in the class #L if there is a log-space non-deterministic Turing machine M such that the number of accepting computations of M on input x is exactly $f(x)$, for all $x \in \Sigma^*$. A function $g : \Sigma^* \to \mathbb{N}$ is in GapL if it can be expressed as $g = f_1 - f_2$ with $f_1, f_2 \in \#\mathrm{L}$.